national research centre for
OHS regulation

RegNet

Working Paper 72

Risk Management and Rule Compliance Decision Making in Hazardous Industries

Professor Andrew Hopkins
(Australian National University)
(contact: Andrew.Hopkins@anu.edu.au)

February 2010

THE AUSTRALIAN NATIONAL UNIVERSITY

About the Centre

The National Research Centre for Occupational Health and Safety Regulation (NRCOHSR) is funded by WorkCover New South Wales, WorkSafe Victoria and Workplace Health and Safety Queensland to work to achieve excellence in OHS research and regulation development. The NRCOHSR is a research centre within the Regulatory Institutions Network (RegNet) at The Australian National University (Canberra), and operates in association with Griffith University (Brisbane).

The NRCOHSR conducts and facilitates high quality empirical and policy-focused research into OHS regulation, and facilitates the integration of research into OHS regulation with research findings in other areas of regulation. We encourage and support collaborating researchers to conduct empirical and policy-focused research into OHS regulation. The NRCOHSR also monitors, documents and analyses Australian and international developments in OHS regulation and research, as well as related areas of regulation, and produces a web-based series of working papers reporting on research into OHS regulation.

Address for correspondence

National Research Centre for OHS Regulation
Regulatory Institutions Network
Coombs Extension
Cnr Fellows and Garran Road
The Australian National University
Canberra, ACT, 0200
Email: nrcohsr@anu.edu.au.

Disclaimer

The views expressed in this paper are the authors' alone and do not reflect any formal opinion of the National Research Centre for OHS Regulation, the Regulatory Institutions Network or the Australian National University. They are provided for the purposes of general discussion. Before relying on the material in this paper, readers should carefully make their own assessment and check with other sources as to its accuracy, currency, completeness and relevance for their purposes.

Abstract

Risk-management and rule-compliance are inter-related strategies for promoting safety in hazardous industries. They are co-existing and complementary, not contradictory. However risk-management offers very little guidance to end point decision-makers; they need rules to guide their decisions. Accordingly, it is important, even within a risk-management framework that risk-management be translated into rule-compliance for end point decision-makers, where possible. The paper demonstrates that this is what in fact happens for a wide range of operational decision-making.

For non-operational decisions, such as investment and design decisions, the need to convert risk-management into rule-compliance is equally important, although more controversial. Nevertheless the authorities have shown that they are willing to impose prescriptive technical rules on duty holders in relation to non-operational decisions, in the interests of safety.

These points are illustrated using a variety of empirical examples and materials, most particularly, the BP Texas City accident, the Buncefield accident, and the Australian pipeline standard.

1. Introduction

Two broadly contrasting methods of assuring safety in hazardous industries can be identified: risk-management and rule-compliance. These are not mutually exclusive approaches; they are complimentary. The issue therefore is not to decide between the two; it is to get the balance right. Debate about the relative merits of these approaches has been on-going for years, but there appears to be a resurgence of interest at the present time.[1] This paper is an intervention in the debate and seeks to re-emphasise the importance of rule-compliance.

The debate can be traced back at least as far as the early 1970s, when Lord Robens in the UK proposed a dramatic shift away from so-called prescriptive regulation, that specified in great detail the rules to be followed, to an all-encompassing requirement that employers ensure the safety of workers "so far as is reasonably practicable". This requirement was subsequently enacted in many jurisdictions around the world. The Robens requirement was not specifically couched in terms of "risk", but it is precisely equivalent to the requirement that risks be "as low as reasonably practicable". Nowadays, major hazard facilities in many countries are regulated under safety cases regimes that require facility operators to demonstrate that the risks are as low as reasonably practicable.

This background makes it clear that there is an historical dimension to the debate: rule-compliance is sometimes described as the traditional approach and risk-management as the new or modern approach. For those who see history as progress, the implication is that the risk-management approach is to be preferred, and any attempt to move the balance in the other direction is somehow retrograde. But for those who see change more as a swinging pendulum, it is conceivable that the pendulum has gone too far in the direction of risk-management and that the time has come for a swing back in the direction of rule-compliance.

Metaphors aside, the thesis of this paper is as follows. Rule-compliance is a vital component of any safety strategy, and in the movement towards the risk-management of safety, we have tended to lose sight of this fact. The reality is that risk-management needs to be converted to rule-compliance wherever possible. Moreover, risk-management *has indeed* been converted to rule-compliance in many cases. This is not an argument for abandoning risk-management, but rather for recognizing the importance of rule-compliance within an overall risk-management framework.

Nor is it simply an argument for more or better legislatively created rules. Other writers have made the point that we need a wide array of government-made rules, ranging from detailed prescription in some cases through to general outcome requirements that leave it to the duty holder to decide how to achieve the required outcomes.[2]

---

[1] See for instance material produced by the European Process Safety Centre, www.epsc.org.

[2] Bluff and Gunningham identify four different kinds of rules or standards that governments can impose: specification standards, general duties, performance standards and process standards. The requirement to

The point I want to make here is that rules are not only made by governments. There are many other relevant rule makers, including industry associations, companies, and even individuals, who sometimes create rules for themselves as a way of dealing with the uncertainties of risk-management. When examined more closely, many of the arguments against detailed rules are really arguments that *governments* should not be involved in detailed rule making, not arguments against rules as such. This paper is concerned with the need for rules, regardless of who makes them.

There is an in-principle reason why rules are necessary. Consider the following decision–making dilemmas.

- Should I wear a hard hat on a production site, to reduce the risk of being injured by falling objects, or not?
- Is it too hot to work in the normal way, or not?[3]
- Am I too fatigued to fly this aircraft, or not?
- Should I stop a process now because of the risks involved, or not?
- Should I authorize this expenditure now as a means of reducing risk, or not?
- We already have several safeguards in place. Do we need one more, or not?

In all these cases, the risk level falls somewhere on a continuum from extreme to insignificant. Yet in each case the risk has to be judged as falling on one or other side of a line: on one side, a certain action is required; on the other, it is not. In other words, the risk continuum must be converted into a dichotomy for the purposes of decision-making. How is this to be done?

The risk-management approach does not in general provide much guidance to those faced with these decisions, that is, it does not offer a way of deciding whether the risk is acceptable or not. There is one obvious exception to this proposition. Where it is possible to carry out a thorough quantitative risk-assessment, determine the numerical risk and then compare this with some predetermined acceptable risk limit, then risk-management does in principle provide an unambiguous way of making the decision. But apart from the inherent limitations of this approach,[4] this is not a practical possibility in most of the decision-making situations described above. Generally speaking decision-makers need rules, not numerical risk acceptance criteria, to guide their decisions.

That does not mean that risk-management is irrelevant. On the contrary, determining the appropriate rules will most likely depend on an assessment of the risks. But this risk-

---

engage in a risk-management process is an example of what they call a process standard. See Chapter 1 in L Bluff, N Gunningham and R Johnstone, *OHS Regulation for a Changing World of Work*. The Federation Press, Sydney, 2004.

[3] A particular issue in Australia.

[4] I have written extensively about these limitations in my *Safety, Culture and Risk*, (CCH, Sydney, 2005), chap 12, and will not rehearse them here. See also a useful discussion of these issues in Inger L Johansen, *Foundations and Fallacies of Risk Acceptance Criteria*, Masters thesis, Norwegian University of Technology, Trondheim, 2008.

assessment is one step removed from the end point decision-maker. Furthermore, behind this risk-assessment may lie another rule, this time a legislative requirement that risk-assessments be conducted and controls put in place. From this point of view, risk-management and rule-compliance are intertwined and complementary strategies.

Given that decision rules serve to dichotomise the risk continuum, they are inherently arbitrary to some degree. What this means is that for cases that fall immediately on one side or the other of the cutting point, the rule may seem unnecessarily strict or alternatively unreasonably weak. For instance, why do I need to wear my hard hat everywhere on the construction site, even though there are many places on site where there may be no equipment or people immediately above me? Surely a more realistic rule could be devised? However, once these kinds or arguments are entertained, we are back to case-by-case risk-assessment and the rule loses its simplicity and enforceability. This point applies as much to decisions about whether to invest in a particular piece of safety equipment, as it does to the decisions about whether to wear a hard hat.

One more introductory point. The preceding discussion has talked about rule-*compliance* as a strategy, not just rule-*making*. The reason should be obvious. Accident investigations routinely identify non-compliance with rules as a contributory factor. Clearly, rule-making by itself is of little value in achieving safe operation. What is required is compliance with those rules. This in turn depends on adequate enforcement or accountability mechanisms. In other words, the rule-compliance approach requires both rule-making and rule-enforcement.

2. Front line decision-making

Let us begin by asking whether it is reasonable that safety-relevant decision-making by front line workers or operators be based primarily on their own risk-assessments? For a number of reasons, the answer is: no.

In the first place, workers may not fully understand the hazards and the controls that have been put in place to deal with those hazards. This is especially true of process hazards which can generate major accident events. Engineers may have done complex calculations to identify the operating risks and to determine a safe operating envelope, that is, a set of temperature, pressure, flow, and other limits that need to be observed if a process is to be carried out safely. Frontline operators cannot be expected to appreciate the full significance of these limits and must simply regard them as rules governing the operation of the plant. In short, safety depends on operators complying with a set of operating rules which themselves are derived from a complex technical risk-assessment.

This idea is nowhere more clearly set out than in the UK Nuclear Installations Act. Condition 23 requires that

> the licensee shall, in respect of any operation that may affect safety, produce an adequate safety case to demonstrate the safety of that operation and to identify the conditions and limits necessary in the interests of safety. Such conditions and limits shall hereafter be

> referred to as operating rules… (and) the licensee shall ensure that operations are at all times controlled and carried out in compliance with such operating rules.[5]

But even when the hazards are understood, there are good reasons for seeking to replace risk-management by rule-compliance for front line workers. The problem is that there are various factors conducive to risk-taking behaviour by workers, such as: a desire to make life easier for oneself, a perceived pressure to get the job done, and a preference for working skillfully, which may mean "closer to the edge".[6] All these things are likely to result in workers accepting risks that are considerably higher than public policy or companies themselves are willing to accept.

Let us reflect for a moment on this discrepancy. The individual worker may decide that the probability that a certain kind of risky behaviour will result in a fatal accident is acceptably low. However, if the company has a thousand workers behaving in this way, the probability that the *company* will incur a fatality is 1000 times greater. The largest companies carrying out hazardous operations must realistically expect several fatalities every year. For a large company, then, a fatality is a high probability event. It would also be a relatively low consequence event for the company, were it not for a various consequence-amplifying factors that come into play when a fatality occurs. The fatality is likely to traumatize workmates and managers who are directly involved; the company may be prosecuted for failure to ensure the safety of the worker; and bad publicity may have a variety of intangible but profit-threatening consequences. As a result, companies cannot afford to leave it to individuals to assess risks for themselves and act accordingly. They must ensure that the risks are as low as companies can reasonably achieve. For all these reasons, companies require workers to comply with a variety of rules, ranging from very simple prohibitions and requirements, through to complex operating procedures.[7]

3. Protection against unscrupulous employers

One factor conducive to risk-taking behaviour by workers is pressure from employers. This is touched on above but needs further emphasis. It should not be forgotten that prescriptive safety legislation was originally introduced to protect workers from unscrupulous employers who regarded their workforce as an expendable resource.[8]

Here is a recent example of how, in the absence of clear-cut rules, employer pressure can induce risk-taking behaviour among workers. The example concerns changes in fatigue

---

[5] Quoted in Jan Hayes, "Operational decision making in high hazard organizations", PhD thesis, Australian National University, 2009, p170.

[6] Research on motorcyclists shows that many like to drive as fast as their skill level will allow. Lower speeds are regarded as boring, precisely because they do not make use of the rider's skills. K Natalier, (2001), "Motorcyclists' interpretations of risk and hazard". *Journal of Sociology*, 37(1):65-80. As car drivers we can all relate to this motivation. See also P Hudson et al *Bending the Rules II: Why People Break Rules of Fail to Follow Procedures*, University of Leiden, no date. p22.

[7] Angiullo notes that in nuclear power stations he has visited, the more safety-critical the task, the greater the amount of detail included in the procedure. R Angiullo, "Operational discipline", chapter 7 in A Hopkins (ed) *Learning from High Reliability Organisations*, (CCH, Sydney, 2009).

[8] N Gunningham, *Safeguarding the Worker: Job Hazards and the Role of Law*, Law Book Company, Sydney, 1984.

management regulation in the aviation industry in Australia.[9] Until about 2000, fatigue was managed by government regulations specifying the maximum numbers of hours that pilots were allowed to fly. These so called prescriptive rules proved to be inappropriate in many circumstances, particularly in the general aviation sector (as opposed to the regular public transport sector). For instance, rules governing commercial airline pilots may not be appropriate for emergency service pilots who spend much of their on-duty time awaiting call out, or for balloon pilots who must begin work very early in the morning. The government therefore began moving away from its prescriptive approach, allowing aviation companies to develop their own fatigue risk-management systems. It was hoped that companies would develop their own hours-of-duty rules that would effectively manage fatigue, while taking account of the particular circumstances of the operation. However the new regime was widely perceived as freeing operating companies from almost any requirement to limit the number of hours flown by its pilots.[10] The study found that companies responded in one of two ways. Larger companies, tended to continue with existing hours-of-work limits, especially where these had been agreed with a unionized workforce, without attempting to go through a real risk-management process themselves. On the other hand, some smaller operators saw it as an opportunity to have their pilots fly longer hours. Furthermore, they devolved responsibility for managing fatigue to those very same pilots, expecting them "to put up their hands" and decline to fly when they felt too fatigued to fly safely. Given the precarious nature of employment in the general aviation sector, this was an entirely ineffective way to manage fatigue, since pilots understood that they risked losing their jobs if they refused to fly. In summary, the actions of the larger companies were an implicit recognition of the need for some kind of rule set in managing pilot fatigue, while the actions of smaller companies demonstrated how the complete absence of any rules can result in irresistible pressure on front line operators to behave unsafely.

4. Rule management

By their very nature, rules are general in their application. It is almost inevitable therefore that there will be situations where workers will judge the rules to be unnecessary or inadequate in some way. These situations will create routine non-compliance unless carefully managed. At BP Texas City, operators were expected to follow sets of procedures and tick off that they had complied with them. They were also able to tick N/A (not applicable) if they regarded a procedural step as inappropriate. As a consequence, procedures were seen as, at best, guidelines, and certainly not requirements to be followed. This casual attitude to compliance was one of the factors that contributed to the explosion at the refinery in 2005.[11]

---

[9] Fiona Keer, "Organisational cultures of safety and regulatory effectiveness: The Civil Aviation Safety Authority's fatigue risk management systems in Australian general aviation". PhD Thesis, Australian National University, 2009.

[10] The regulator did promote the use of a particular software package, FAID, that calculated fatigue scores and it set a score of 80 or 85 as the maximum allowable. But by itself this limit did not provide adequate protection against fatigue.

[11] A Hopkins, *Failure to Learn: the Texas City Refinery Disaster*. CCH, Sydney, 2008.

What is required is that when workers come across procedures that they regard as unworkable or inapplicable, rather than devising their own solutions, they should notify management of the situation and request a review of the rule. Management should then respond rapidly and flexibly, in a way that recognizes the concerns of the workers, while not losing site of the purpose the rule was intended to serve, that is, the risk to be controlled. In short, a rule-compliance strategy requires that management recognize that the regime of rules is always a work in progress and that it needs to be actively managed.[12] In this way, respect for rules can be maintained. A study of a nuclear power station in the US has shown that when rules are actively managed in this way, non-compliance can be completely eliminated.[13]

5. Rule-compliance and risk awareness

It is sometimes suggested that a strategy of relying on rules to direct workforce behaviour creates a compliance culture, in a negative sense, because it encourages workers to think that safety is merely about compliance with rules and that there is no longer any need for them to take responsibility for their own safety or to maintain an awareness of the risks of what they are doing. In short the suggestion is that a strategy of rule-compliance undermines risk awareness on the part of front line workers. Historically, this has been a problem for railway companies, for example.[14]

However, it should be obvious from the comments in the preceding section that this is not a necessary outcome. An organisation that has a regime of rule management in place relies on workers to call attention to situations where they believe that rule-compliance is inappropriate, so that the rules can be examined and changed if necessary. In particular it relies on workers to call attention to situations where they believe rule-compliance would lead to unsafe outcomes. This requires workers to remain risk aware and not simply to follow rules blindly. In short, provided there is a possibility of rule modification built into the management system, there is no incompatibility between rule-compliance and risk awareness.

There is one other way in which risk-assessment remains relevant for front line workers, even within what is, from their point of view, a rule-compliance framework. Many companies require workers to carry out risk-assessments before new tasks are begun. But these are not risk-assessments from first principles. They are better viewed as exercises to raise the awareness of risk and perhaps to identify the relevant safety rules that need to be complied with.

---

[12] Hale, A., T. Heijer, F. Koornneef (2003), "Management of safety rules: The case of railways". *Safety Science Monitor*, 7 (1):1-11.
[13] M Bourrier (1998) "Elements for designing a self-correcting organisation: examples from nuclear plants": In A Hale and M Barram (eds) *Safety Management: The Challenge of Change*. Pergamon: Oxford, pp13-146.
[14] A Hopkins, *Safety Culture and Risk*, CCH, Sydney, pp38-9.

6. The on-going quest for rules to guide decision-making

Not all decision-making can be proceduralised. There will always be situations not covered by the rules, or perhaps where quick decisions are needed, which require individuals to draw on their own expertise to assess risks and act appropriately. The theory of High Reliability Organisations holds that this is not only inevitable but also desirable.[15] In particular, decisions to interrupt an ongoing process, for example to abort a space shuttle launch or to close down a nuclear power plant, need to be taken by those with the greatest expertise, whomever they may be. Importantly, such people may be in quite lowly organisational positions.

Nevertheless, given the difficulty of making formally unstructured decisions, it is not surprising that there is an on-going quest for additional guidance. This guidance amounts to further sets of rules about how decisions are to be made in complex situations. Sometimes the procedures that are developed are highly innovative. Here are two interesting examples.

6.1 *The rule of three*

The first of these is the "rule of three", developed by Shell. Suppose there are several risk-enhancing factors present. No one factor poses a significant threat by itself, and therefore no one factor requires action to reduce the risk, but it is clear that the greater the number of such factors, the greater the risk. What is the decision-maker to do? This is a common dilemma, and it is a dilemma precisely because there are no rules that govern the situation. The rule of three states simply that if three or more such factors are present, this is to be taken as a trigger to stop the activity or take some risk reducing action. The rule is conveniently stated in traffic light terms: three orange lights are the equivalent of a red. For instance, a helicopter pilot is almost out of allowable hours; the weather has closed in but is still just within company rules; the flight plan has been changed at the last minute. Each of these is a risk factor and none by itself is sufficient to stop flight operations, but together, according to the rule of three, they pose an unacceptable risk. This is a real case, in which the rule of three was not in use and the flight ended in tragedy with 13 lives lost.[16]As this example makes clear, just what the orange lights are will depend on the operation, and operators will need to establish beforehand, perhaps in the calm of the office, just what the risk enhancing factors are that will be taken into account in applying the rule. This is an excellent small group exercise for promoting risk awareness in hazardous situations.

Three oranges do not necessarily mean that the whole operation must be terminated. It may be that there are opportunities for "managing" one of the oranges to green, that is, eliminating one of the risk factors. For instance, in this case, there might be a relief pilot available who might be asked to take over the flight.

---

[15] Weick K, K Sutcliffe & D Obstfeld (1999), "Organising for high reliability: processes of collective mindfulness", *Research in Organisational Behaviour*, vol 21, pp81-123.

[16] Hudson, P et al "The rule of three: situational awareness in hazardous situation". Society of Petroleum Engineers, SPE 46765.

The rule of three provides a mechanism that converts the risk continuum into a dichotomy for the purposes of decision-making. It does not, however, ignore the expertise of operators; it draws on that expertise in identifying relevant risk factors. Put another way, the rule of three structures expert decision-making; it does not replace it.[17]

6.2 *TARPS (Trigger Action Response Plans)*

The mining industry in Australia provides a second example of an attempt to devise rules to assist decision-making in situations that might otherwise be thought too complex to proceduralise, and in which it might be thought that decision-making must be left to professional judgment. The companies and the regulators in this industry have developed the system of "trigger action response plans" (TARPs) to cope with the uncertain hazards that confront miners.[18] These plans indentify warning signs, or risk factors (triggers) of increasing concern and a corresponding set of actions that must be taken when these factors are present. Take the hazard of a collapsing roof. As mining moves into new areas, the nature of the roof can change; in particular, it can become less secure and need more intensive support in order to protect miners working beneath it. Miners must therefore be alert to the changing nature of the roof and must take appropriate action to support the roof as it changes. Indicators of increasing danger include falling flakes of rock, increasing quantities of water dropping from the roof, and the appearance of certain geological formations. This is the kind of situation in which it might well be thought best to rely on the expert judgement of those concerned, but it is the very type of situation that TARPs are designed to deal with. One mine I visited had identified four states of increasing concern, labelled: green, yellow, orange and red. Triggers were specified for each state, along with the actions required by miners and their managers. Here again a complex risk-management problem has been converted into a question of rule-compliance.

7 Self-imposed rules

Sometimes decision makers deal with the uncertainties of risk-management by developing and applying their own rules. This interesting pattern of behaviour was discovered by Jan Hayes in her recent study of on-site operations managers.[19] In multi-shift operations these are in effect the shift managers.

The respondents in the study were drawn from industries that were regulated to varying degrees under safety case regimes, which required that risks be identified and managed down to acceptable levels. Such regimes explicitly accept that there is no such thing as absolute safety; that safety is a matter of degree, and that safety-relevant decision-making

[17] The figure three is not entirely arbitrary. There is evidence that once three risk-enhancing factors are present the risk level rises appreciably. K. v. d. Merwe, "Testing the rule of three", Master's thesis, University of Leiden, 2004; H. Jonker, "Cockpit decision making", Master's thesis, University of Leiden, 2000.

[18] A Hopkins, *Safety Culture and Risk*, CCH, Sydney, Chapter 8.

[19] Jan Hayes, "Operational decision making in high hazard organizations", PhD thesis, Australian National University, 2009.

involves balancing safety against cost or production. Despite this framework, the operations managers in the study generally did not think in these terms. They did not see themselves as accepting a certain level of risk so as to facilitate production. Their thinking was more dichotomous. The system was either safe or unsafe and if it was unsafe, they would close operations down until safety could again be guaranteed. What was particularly remarkable about this finding was that many of these operations managers had been trained to think in risk-management terms, yet when it came to the point the training seemed to be irrelevant.

The key to this remarkable situation lies in the mental model that these managers use. For many of them, safety is not a question of risk-management; it is about ensuring that the system safeguards or barriers against failure are all in place.[20] If they are in place, the system is safe, absolutely safe, so safe that these managers would be happy to bring their children on site and show them around, they said. But if not all the barriers are in place, the system is not safe and something must be done. According to Hayes,[21] when not all barriers are in place, managers adopt one of two options:

- Stop/limit/curtail production to within the limits of the remaining barriers, or
- Provide a temporary replacement barrier, which might simply be increased monitoring by the operational team.

Hayes observes that

> This barriers approach is not inconsistent with broader organisational risk-management (which should have been the way the design of the various barriers was developed in the first place), but it explains why operational managers do not use risk directly in reaching safety decisions. Consideration of barriers is less subjective than trying to assess risk in a dynamic operational environment and explains why operational managers do not see a technical conflict between safety and production.[22]

From the present point of view, the significance of the barrier approach is that it enables managers to develop rules about what to do when barriers are compromised, rules which they impose on themselves. They are rules that emerge from the mental models and the experience of these people and they provide them with decision-making criteria they need to do their job.

One of Hayes' research sites was a nuclear power station and here she found that in some circumstances shift managers invoked a different kind of rule to deal with the increased risk associated with barrier failure. It can be illustrated in the following way.

---

[20] The UK HSE comes to arguably similar conclusions. "Optimising hazard management by workforce engagement and supervision". Norwich, UK Health and Safety Executive, 2008, RP637. See Hayes op cit. pp 267-8.
[21] Ibid., 267.
[22] Ibid., 284.

Suppose an emergency backup pump had been found to be non-operational. It might be one of several redundant backup systems, but to continue production in these circumstances means operating with a smaller safety margin. Should the process be closed down, or is it reasonable to accept the slightly higher risk and to continue operating while the issue is resolved? The problem is that it is seldom possible to know how long this will take. Managers know that, at any one point in time, a few more hours operating with the smaller margin of safety will not increase the risk appreciably. Accordingly, if the question is whether to stop now, or to continue operation a little longer in the hope that the problem will soon be fixed, it is reasonable to continue.

However, they also know that the longer they continue operation in a degraded state, the greater the tendency to "normalize" the situation, that is, to accept the greater level of risk as normal. The normalisation of risk has been a significant factor in many major accidents. For example, prior to both the Challenger and Columbia space shuttle accidents, a certain level of equipment malfunction came to be accepted as normal because it had not in the past led to disaster. People became desensitized to the risks of operating in this way. Ultimately these malfunctions proved fatal.[23]

The way the nuclear power station managers dealt with their dilemma was to draw what they called "a line in the sand": if the matter was not resolved within say 24 hours, they would stop production. In this way they created a rule for themselves: if 24 hours passed without a resolution, there was no longer a risk-assessment to be carried out, there was a rule to be complied with, and the decision was clear-cut. Similar line-in-the-sand thinking was evident at all research sites in the study, although not as clearly articulated as in the case of the nuclear power station.

In summary, Hayes' research shows that operations managers are not using the risk-management framework to balance safety and production in the way that the overarching legal framework envisages. Instead, they see their job as ensuring that safety is not compromised and they develop decision rules to ensure that it is not. In short, they have converted risk-management into rule-compliance.[24]

8 Non operational decisions

Discussion to this point has been largely about decisions in an operational environment, that is, by people who are monitoring or controlling on-going operations. There is, in addition, a whole set of what might be described as non-operational decisions that impact on safety, for example, planning, design, and investment decisions. I shall argue that it can be dangerous to allow those who make these decisions to carry out their own risk-assessments, because these assessments are likely to be biased in the direction of

[23] Columbia Accident Investigation Board, *Report, Volume 1*, NASA, Washington, August 2003
D Vaughan, *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*. Chicago: Univ of Chicago Press, 1996.

[24] Studies of safety management in the airline industry come to similar conclusions about the way risk assessment is translated into barrier analysis. C Macrae, "From risk to resilience", and A Hopkins, "Identifying and responding to warnings", both in A Hopkins (ed) *Learning from High Reliability Organisations* (Sydney, CCH, 2009).

allowing decision-makers to do what they are already predisposed to do. In particular, there are constant pressures on these decision-makers to minimize cost and hence to under-state the risks of the lowest cost option. As one risk analyst acknowledges: "risk-assessment is like torturing a spy. If you do it for long enough you get the answer you want!"[25] The argument here will be that these kinds of decisions often need to be constrained by rules. Of course such a rule may be determined by a higher level risk-assessment, but as long as this risk-assessment is independent of any particular application of the rule, it stands a much better chance of being unbiased.

The claim that it is necessary to translate risk-management into rule-compliance appears to be more controversial in this non-operational context than for operational decision-making. One can speculate about the reasons. First, investment and planning decisions are likely to have a greater impact on company profit than many operational decisions. Second, the decision makers are likely to be more powerful, and hence more resistant to limitations on their decision-making freedom and more able to point to the inevitable inconsistencies and inefficiencies in all blanket restrictions. Third, the rules concerned are more likely to be externally imposed rules and thus in conflict with prevailing presumption in favour of self-regulation.

Two recent major accidents - BP Texas City in the US and Buncefield in the UK - have highlighted the issue of non-operational decision-making and the need for these decisions to be constrained by technical rules. This will be developed at some length in what follows.

9. BP Texas City - the absence of prescriptive rules

An explosion occurred at the BP Texas City refinery in 2005, killing 15 people. I shall not describe the accident in any detail here but concentrate on two features that are relevant to this discussion. First, as a result of a process upset, the equivalent of a tanker load of petrol escaped through an open stack, a tower, over a period of nearly two minutes. It cascaded down to the bottom of the stack where it formed a massive vapour cloud that subsequently exploded. Best practice in the industry is not to allow flammable material to escape in this way but to ignite it, as it is released, by means of a flare. Had such a flare system been operating at Texas City, there would have been no explosion and no deaths. Texas City acknowledged that flares were best practice and its policy was that new process units would be equipped with flares, not open vents. But it had chosen not to replace existing vents with flares, for reasons of cost. On one occasion the vent in question had been completely rebuilt from the ground up, but Texas City had not seen this as an opportunity to implement best practice, and had rationalized its decision by saying that this was not a new vent but merely the replacement of an old![26] The US Occupational Health and Safety Administration regarded the vent as unsafe and had tried

[25] P Webb, "Prescription – a step on the road to dependence or a cure for process safety ills?" Paper delivered at the Hazards XXI conference, Manchester, 10/11/09.

[26] It has been argued that in so doing Texas City was violating its own rule (Webb Ibid.), but it is not at all clear that that the company rule applied in this situation. Texas City was taking advantage of an ambiguity in the rule so as to avoid the cost of converting to flare.

to get Texas City to replace it, without success. The problem was that there was no regulation or otherwise enforceable rule that specifically required the vent be replaced, and it was therefore a matter of assessing the risks. Texas City argued that the risks were adequately controlled and ultimately OSHA was not able to over-ride this judgment. Interestingly, Texas City management recognized that they would eventually be required by the Environmental Protection Agency to replace the vent with a flare, for environmental reasons, but their position was that they would not make the change until required to by law.

There is an implication that can be drawn from this story. Where it is clear what good practice is, as it was in this case, it needs to be enshrined as a rule, for example in company standards or industry standards or perhaps in government regulations, in such a way that it can be enforced by government inspectors, if necessary. The problem at Texas City was that there was room for argument, which site management was able to exploit in such a way as to avoid the expenditure required to bring the site up to standard.

There is a second aspect of the Texas City disaster that is relevant to this discussion. The people killed were located in flimsy, temporary accommodation units – trailers – that were located much too close to process equipment. Company engineers had done risk calculations and determined that the risk was acceptably low if the trailers were at 350 feet. However this was not imposed as a company rule. Trailers could be sited closer if a local, site-specific risk-assessment indicated that this could be done safely. In the Texas City case, a site specific risk-assessment was done, as a result of which trailers were located much closer, in one case, within 120 feet of the vent. This came about because the people who conducted the local risk-assessment already knew where they wanted to locate the trailers and, from their point of view, the risk-assessment was designed to justify the proposed location. This introduced a powerful confirmation bias into the decision-making process, turning it into little more than a legitimation ritual. This is not to point a finger of blame at these people. They were not equipped to carry out the required risk-assessment and they should never have been put in this position. The company would have been much better off treating the 350 foot figure as a rule to be followed rather than as a trigger to a site-specific risk-assessment. Such an approach would combine risk-assessment and rule-compliance in an optimal way. Site-specific decisions about trailer location would be governed by a rule, while the rule itself would have been determined by a risk-assessment done by company risk engineers well away from site specific pressures.[27] This is an important example because the rule referred to is a company rule. The argument here is about the need for rules, not necessarily government imposed rules.

The Texas City saga continues, and subsequent events provide further evidence of the importance of rules. In late 2009 the US Occupational Safety and Health Administration proposed that BP be fined another $87 million, over and above the initial $21 million

[27] In this case, the risk engineers made various assumptions that turned out to be incorrect. As a result a 350 rule would not have been sufficiently conservative. However it would have been a great deal better than no rule at all. Had a 350 rule been in place, and complied with, it is possible that no one would have been killed in the Texas City explosion.

fine. The new fines were for BP's failure to implement certain risk reduction strategies at Texas City. BP's lawyers contested the new fines on various grounds. One of these is interesting from the present point of view, because it is precisely analogous to the reasoning that led to the failure to replace the vent with a flare.

The issue is whether existing relief valves should meet a certain performance standard. The American Petroleum Institute has formulated the performance standard as "a recommended practice" (RP 520). BP states that "as a recommended practice API RP 520 is not a mandatory standard in the refining industry". It agrees to comply with the standard for *new* installations but not for *existing* ones. OSHA has insisted that *existing* relief valves at Texas City should comply with the standard, on the grounds that it is "recognized and generally accepted good engineering practice" (RAGAGEP).[28] BP has countered that existing relief valves at most refineries across the US do not in fact comply and therefore this cannot constitute RAGAGEP. In short BP is resisting OSHA's attempts to force Texas City to adopt the performance standard in question on the basis that there is not a rule that unequivocally requires it to do so.[29] This is a graphic illustration of the importance of rules for regulatory effectiveness.

10. Some objections

This argument about the need for more prescriptive rules was advanced in my book *Failure to Learn: The BP Texas City Refinery Disaster*, and it has proved to be controversial. Various writers in Australia and the UK have understood me as arguing for a return to the pre-Robens era of prescriptive regulation. This is a misreading of the argument. I said in the book that "this is not a recommendation for the abandonment of existing legislative frameworks, but it is a suggestion that, in some cases, regulatory objectives may be better achieved by converting risk-management requirements into requirements for rule-compliance".[30]

A criticism by Peter Webb[31] needs to be dealt with in more detail here, because it necessitates a clarification of the argument. Webb focuses on the trailer siting issue at Texas City. The trailer siting risk-assessment was specified in a complex set of procedures, known as a process hazard analysis (PHA). The team that performed the PHA did not fully understand the procedures and did not comply with them adequately. The whole process, Webb argues, was over-proceduralised, leading to a box-checking mentality rather than any real consideration of risks. For this reason, the failure is better viewed as a failure of the rule-compliance approach rather than of the risk-management approach. He goes on:

---

[28] OSHA 29 CFR 1910.110 (D)(3)(ii) states that "The employer shall document that equipment complies with recognized and generally accepted good engineering practice".

[29] Letter from Thomas Wilson of Vinson and Elkins to Mark Briggs, OSHA, dated October 5 2009. Some insight into the enforceability of standards in the US can be gained from "Secretary of Labor v. Luna Tech, 2002 OSHRC No 3, dated September 11, 2002. It appears that RAGAGEP may include "an employer's own appropriate internal standards, as well as industry consensus standards". 57FR at 6390.

[30] A book review in *OHS Professional*, October 2009, p7, asked for a more detailed and systematic analysis of the topic. The present paper is that analysis.

[31] P Webb, "Prescription".

> Their goal was to comply with the …procedure rather than manage the risks. …if they had regarded (their goal) as managing risks they could apparently have thrown away the procedure and done a better job.[32]

In support of this latter claim he notes that there was one individual at Texas City, a technician, who had expressed concern about the trailer siting. This man had a standard of comparison, having worked at other refineries where trailers were located a lot further away, as a precautionary measure.

However this man was exceptional. The evidence is overwhelming that the level of risk awareness at Texas City was low (for instance, the explosion was triggered by a vehicle near a process unit that had been left with its engine idling), and it is quite implausible to suggest that relying on unstructured and informal risk-assessments would have yielded a better outcome.

But let us return to Webb's fundamental claim that what happened at Texas City was a failure of the rule-compliance approach. He is able to make this claim because he conceptualises the risk-assessment practiced at Texas City as a question of rule-compliance. In short he makes no distinction between rules that require that risk-assessments be carried out, and rules that require that specific technical risk controls be adopted. Yet this is what the debate is really about. If the failure of the trailer siting procedures is treated as a failure of rule-compliance then the debate must be recast as a debate about the relative merits of different kinds of rules.

11. Types of rules

The distinction in question is well described in the system of classification for safety rules proposed by Hale and Swuste.[33] They identify three categories:

1. rules defining goals to be achieved, eg duty of care requirements;[34]
2. rules that define the way decisions about a course of action must be arrived at;
3. rules defining concrete actions or required states of the system.

The transition from type 1 to type 3 amounts to a progressive limitation on the freedom of choice for the rule follower.

---

[32] Ibid.

[33] A Hale and P Swuste, "Safety rules: procedural freedom or action constraint? *Safety Science* 29 (1998): 163-177.

[34] Many writers include in this category rules that specify maximum acceptable concentrations of toxic chemicals. However such rules specify required system states and are better seen as type three rules. Consider the following rule: If the concentration of flammable contaminant in the atmosphere of a confined space is found to be 10% or more of its lower explosive limit, no person may enter or remain in the space" This would seem to be a classic example of the type of rule in question. Yet Bluff and Gunningham classify it as a type 3 rule, op cit, p18.

It is clear that the two types or rules discussed above correspond to types 2 and 3 in this scheme. Rules requiring risk-assessments are examples of rules specifying ways that decisions are to be made, while rules about which technical risk controls are to be used are specifying "concrete actions or required states of the system". Using this scheme, the trailer siting represents a failure of type 2 rules, but not type 3, and the Texas City experience points to the need for a more extensive use of type 3 rules.

12. Buncefield – official recognition of the need for a prescriptive rule

The need for more prescriptive technical rules has emerged in the wake of another very high profile accident, at Buncefield, in the UK. The Buncefield case is in many ways analogous to the Texas City accident. One important difference is that it occurred in a Robens-inspired jurisdiction, that is, where the overarching legal framework requires that risks be as low as reasonably practicable. A second relevant difference is that in this case governmental authorities explicitly concluded that there was the need to replace case-by-case risk-assessments with a rule. This is what makes Buncefield particularly interesting from the present point of view.

Buncefield is a very large petroleum storage depot, a tank farm, not far from London. It was destroyed by fire in 2005, the same year as the Texas City disaster, playing havoc with England's fuel supplies, particularly to Heathrow airport, yet by sheer good luck, no one was killed or injured. One of the tanks was being filled by fuel coming by pipe line from a distant refinery. For reasons that were not made clear in the official report,[35] the tank overfilled, resulting in a massive vapour cloud that drifted away, found an ignition source and exploded.

Risk analysts had not previously focused on the possibility that a spillover could generate a vapour cloud and hence a vapour cloud explosion, with catastrophic consequences.[36] The realization that this was indeed possible led the Buncefield Major Incident Investigation Board to recommend that all such tanks be fitted with level detection equipment that would *automatically* cut off supply in the event of an overfill event. They explicitly rejected the idea that operators should be relied on to interrupt flow *manually.* The following passage explains why.

> While the application of BS EN 61511 provides a risk-based target for the integrity of an overfill protection system, and hence for the reliability of the constituent components, it does not require such systems to be automatic in operation. However, if the overfill protection system relies on human operation, the possibility of human failure remains, resulting in common cause failure of both the tank gauging system and its overfill protection system. This possibility is very difficult to quantify but is likely to be a critical factor in determining the likelihood of overfilling. For this reason it is felt necessary to

[35] This was a deliberate omission so that the report would not jeopardise the subsequent criminal prosecution. However the result is that the report fails to provide a complete account of the accident.
[36] Buncefield Major Incident Investigation Board (MIIB), The Buncefield Incident, 11 December, 2005, Final Report, Volume 2, "Initial Report" , p19.

> make the additional recommendation that overfill protection systems should be automated.[37]

This recommendation explicitly rejects the risk management approach. More to the point, it will be relatively expensive to implement. Not surprisingly, UK industry was at first reluctant to accept it, preferring case-by-case risk-assessments that would allow tank farm operators to argue that the other risk control strategies they might have in place reduced the risk to acceptable levels without the need for automatic cutouts. Eventually, however, the UK Petroleum Industry Association and the Tank Storage Association announced that they would adopt the recommendations.[38] Finally, the British government announced that it would require all sites to move to fully automatic shutdown systems on tanks storing gasoline.[39] The requirement for automatic cutout was now to be a rule.

However, the details of the rule were still to be worked out. Interestingly it was not the regulator that carried out this work, but a group representative of industry, unions and the regulator, known as the Process Safety Leadership Group. It produced a document entitled "Safety and environmental standards for fuel storage sites".[40] Among other things this detailed the precise nature of the tanks to which the requirements for automatic cutout applied. It also allowed for the possibility that there might be technical reasons as to why automatic cutout systems were not appropriate, but it stated in these circumstances, "duty holders will need to prepare a robust demonstration that alternative measures are capable of achieving the same …outcome as an overfill protection system that is automatic…".[41]

The document was published by the regulator, the Health and Safety Executive (HSE), and it notes that, while it is not an authoritative statement of the law, compliance with its requirements would normally ensure that the duty-holder was in compliance with the law. This raises the question of its enforceability. Perhaps the most important point here is that, as an industry agreed document, it is an industry agreed statement of good practice. And to the extent that the regulator is capable of enforcing good industry practice, the requirements specified in the document are indeed prescriptive rules that must be complied with. As one UK inspector said to me, "this is a big step towards prescription".[42]

It seems at first sight remarkable that industry groups such as the Tank Storage Association were willing to commit themselves to a document that would require some of

[37] Buncefield MIIB, Final Report, Volume 2, Recommendations on the Design and Operation of Fuel Storage Sites, p33. See also pp26-7.
[38] Buncefield MIIB, Final Report, Volume 1, p29.
[39] The Buncefield Investigation: The Government and Competent Authority Response. November 2008, p10.
[40] Process Safety Leadership Group, "Safety and environmental standards for fuel storage sites" London, HSE, 2009.
[41] Ibid., 29.
[42] The chair of the Process Safety Leadership Group says in the introduction that the standard had managed to avoid prescription. It is not clear what he could mean by this.

their members to make costly modifications to their tanks.[43] The explanation appears to be that, given the public pressure for action following the Buncefield incident, and given the government's commitment to create an effective rule, the regulator was able to persuade recalcitrants that it was ultimately in their interest to accept the higher standard.

13. Enforcing good industry practice – a step towards rule-compliance

The preceding section left an important question unanswered. To what extent are regulators in a position to enforce good industry practice? Interestingly, in the UK the enforcement of good practice is now an important and explicit element of regulatory strategy. The overarching legislative requirement is that risks be reduced as low as reasonably practicable (ALARP). In the past, for major hazard industries, this has been interpreted as a requirement to carry out a quantitative risk-assessment (QRA) to demonstrate that the risk of fatality is below some target figure, eg 1 in 100,000 per annum. This approach has proved problematic and in recent years the HSE has de-emphasised it. Its guidance now state that "where the law requires risks to have been reduced ALARP, the HSE may accept the application of relevant good practice… as a sufficient demonstration… and does not normally accept a lower standard of protection than would be provided by the application of current good practice".[44] Here is a more extended passage from the HSE document, "ALARP at a Glance".[45]

> In most situations, deciding whether the risks are ALARP involves a comparison between the control measures a duty-holder has in place or is proposing and the measures we would normally expect to see in such circumstances i.e. relevant good practice. "Good practice" is defined … as "those standards for controlling risk that HSE has judged and recognised as satisfying the law, when applied to a particular relevant case, in an appropriate manner." We decide by consensus what is good practice through a process of discussion with stakeholders, such as employers, trade associations, other Government departments, trade unions, health and safety professionals and suppliers.
>
> Once what is good practice has been determined, much of the discussion with duty-holders about whether a risk is or will be ALARP is likely to be concerned with the relevance of the good practice, and how appropriately it has been (or will be) implemented. Where there is relevant, recognised good practice, we expect duty-holders to follow it. If they want to do something different, they must be able to demonstrate to our satisfaction that the measures they propose to use are at least as effective in controlling the risk.

It is clear from this description that good practice is close to a set of prescriptive rules as far as the HSE is concerned. The standards for fuel storage sites are a case in point. Such rules can be enforced, if necessary with improvement and prohibition notices. Of course duty holders may choose to contest these notices in court, but a plaintiff would have a

[43] For an estimate of these costs see Buncefield MIIB, Final Report, Vol 1, p83.
[44] HSE "Assessing compliance with the laws in individual cases and the use of good practice", May 2003, p3.
[45] Ibid., 3-4.

hard time convincing a court that it had done all that was reasonably practicable if it was not complying with an agreed statement of good industry practice or some demonstrable equivalent.[46]

Many regulators use good industry practice as a bench mark against which to judge whether risks are as low as reasonably practicable, and in various ways HSE inspectors have for years been drawing on their knowledge of good industry practice to make such judgements.[47] What is different is that the HSE has now articulated this as its formal policy in a quite dramatic fashion. This is a significant philosophical shift.

14. Rules within rules

There is of course nothing new about technical prescription in standards. In terms of the Hale and Swuste typology there are numerous type 3 rules in a host of standards applicable to hazardous industries. There are also numerous type 2 rules, about how companies should make decisions. These rules require that companies develop various detailed type 3 rules for themselves. It is important that these company-made technical rules be seen as the prescriptive rules that they are, and that regulators be willing to enforce them.

The preceding points are rather abstract. In order to give substance to them, I want to illustrate them using the Australian standard for gas and liquid petroleum pipelines (AS2885), hereafter the pipeline standard. I do not claim that this is typical of all industry standards, but it is by no means exceptional.

14.1 *The status of the pipeline standard*

The legal status of industry standards varies, but there are many situations where they are directly enforceable. For instance, in many safety case regimes, where ever a safety case makes reference to a standard, that standard is enforceable. Moreover, in some situations standards are called up in legislation in a way that gives them the full force of law. The Australian pipeline standard is directly enforceable in this way.[48]

---

[46] It is interesting to contrast this situation with the difficulty that OSHA has in enforcing RAGAGEP (recognized and generally accepted good engineering practice). The overarching legislative requirement in the US is to provide a workplace that is "free from recognised hazards that are causing or are likely to cause death or serious physical harm." Sec 5(a)(1) of the OSH Act. It would be difficult to demonstrate that failure to live up to some good industry practice, such as the installation of automatic cutouts, was "*likely* to cause death or serious harm" (emphasis added), since with or without such cutouts, death or serious harm is unlikely. In short, the overarching statute in the US does not support regulatory attempts to enforce good practice in the way that the over-arching statute in Robens-inspired jurisdictions does. In fact enforcement is under Sec 5(a)(2), which requires that employers "shall comply with occupational safety and health standards promulgated under this Act". See further, Secretary of Labor v. Luna Tech, 2002 OSHRC No 3, under the heading "citation no. 1".

[47] Hutter B, *Compliance: Regulation and Environment* (Oxford: Clarendon, 1997), pp95-6.

[48] For instance Victorian legislation specifically requires that pipelines be constructed and operated according to AS2885. See Pipelines Regulation 2007, section 21. Even where standards are not directly enforceable they generally have evidentiary value, in that non-compliance is prima facie evidence that a duty holder is not doing all that is reasonably practicable to ensure safety.

This is an interesting situation. It means that in many cases governments have in effect delegated their legislative powers to non-governmental bodies. Rather than governments themselves making the detailed rules, they are requiring compliance with rules created by others. It is important therefore to consider, if only briefly, the processes by which these rules come into existence. Standards are generally written in such a way as to reflect whatever consensus can be achieved by those involved in writing the standard. Where standards are written by industry associations, without input from regulators and other groups, there is a risk that the so-called "lowest common denominator approach",[49] will prevail, that is, that the standard will contain only those things that all parties are happy to be bound by.[50] On the other hand where standards are written under the auspices of standards associations, with a requirement that all stakeholders are involved, in particular regulators, they have a better chance of representing good industry practice. The Australian pipeline standard is such a standard.[51]

14.2 *Prescription in the pipeline standard*

One of the greatest threats to underground high pressure gas pipelines is accidental damage by third parties who may be digging nearby. In 2004, a pipeline in Belgium was damaged in this way, causing an explosion that killed 24 people and injured 132. The Australian standard is therefore highly prescriptive in this matter. It specifies the depth at which pipelines must normally be buried (1200 mm in residential areas). It allows for reduced cover in areas of rock (900 mm in residential areas). Moreover, it specifies in remarkable detail exactly what constitutes rock and how extensive the rock needs to be before pipe can be laid at the shallower depth.[52] This degree of prescription was made necessary because some pipeline laying companies were taking advantage of the shallower option at the first sign of any rock.[53]

Again, the standard requires that in certain areas, pipelines must be designed so that, if a pipe is accidentally damaged, perhaps by unauthorized digging, it will "leak not rupture". This means that the pipeline must be designed so that a single hole will not weaken the pipeline to such as extent that it splits open (either around the circumference of the pipe or along its length).[54] This requirement adds to construction costs, but it limits the potential for a catastrophic explosion. The requirement means that no matter how low the probability of damage by third parties may be, pipeline operators cannot argue that this is low enough to justify using lower quality pipes. This is a significant and costly move away from a pure risk-management approach.

---

[49] To anyone who remembers their high school mathematics, this is an irritating expression; it should be the "highest common factor approach".
[50] According to Bluff and Gunningham, allowing industry to write its own standards is "not regarded as the preferred approach" op cit p 39.
[51] In this particular case, the quality of the standard is also attributable to the quality and commitment of industry participants.
[52] AS 2885.1 –2007, table 5.5.2 and figure 5.5.3. See also section 4.3.4.
[53] Interview with regulator.
[54] AS 2885.1 1997, sec 4.7.

14.3 *The further delegation of responsibility for technical rule-making*

The pipeline standard also contains type 2 rules, that is, rules that specify how technical decisions are to be made, without specifying what those decisions must be. So, for example, when it comes to the operation and maintenance of pipelines, the principle requirement of the standard is that operators carry out a thorough risk-assessment and that they devise a "safety and operating plan" to manage identified risks. The standard draws attention to a variety of risk controls that need to be considered, but stops short of mandating particular actions. For instance, it specifies that the plan should include pipeline maintenance and inspection schedules, but it does not itself specify maintenance or inspection intervals.[55]

Let us focus here on the issue of inspection frequency. This is a topical issue. BP has been accused repeatedly of not inspecting its Alaskan pipelines frequently enough, with the result that they were able to corrode through and leak oil onto the snow, creating significant environment damage. Moreover, a recent gas pipeline rupture at Varanus Island, off the coast of West Australia, which severely damaged the economy of that state, has been blamed on insufficient frequency of inspection, among other things.[56] Here is what the Australian standard has to say about frequency of inspection:

> The frequency of inspection and assessment should be documented and approved and based on the past reliability of the pipeline, historical records, current knowledge of its condition, the rate of deterioration (both internal and external corrosion, coating degradation and the like), and statutory requirements.[57]

This is a classic type 2 rule and requires companies to carry out their own risk assessments. On the basis of these assessments they must make their own type 3 rules, about frequency of inspection, and then of course, comply with those rules. This amounts to an additional step in government delegation of responsibility for technical rule-making.

It is in principle more difficult for regulators to ensure that companies are complying with a type 2 decision-making rule, such as the one above, than to ensure compliance with a type 3 technical rule specified in the standard. Checking on compliance involves two steps. First, it would be desirable to be able to evaluate the adequacy of company-made technical rules. However, the decision-making process involved in determining the frequency of inspection may be technically complex and in practice the regulator may be restricted to evaluating the competency of those involved in making the rules, in order to gain some assurance that the rule-making process is one of integrity. Second, when it comes to evaluating compliance with the inspection frequency schedule, the regulator may find that there are no readily accessible documents and that the details of what is supposed to occur and what is in fact occurring may be buried somewhere in a computer database. The regulator may therefore have to request specially prepared documentation to ensure that companies are complying with their own rules.

---

[55] AS 2885.3 –2001, sec3.2, 3.4.1, 4.2.2, especially (d).
[56] Report by NOPSA into the Varanus Island fire, final revision, 7/10/2008, p4.
[57] sec 5.3.2. There are no relevant requirements in Victorian legislation, the one case I have investigated.

All this requires a highly "engaged regulator".[58] In fact not all regulators get down to this level of detail. Furthermore, some regulators would argue that it is not their responsibility to ensure compliance at this level, that compliance is ultimately the responsibility of the company, and that all they can be expected to do is to provide "some level of assurance"[59] that companies are managing their affairs appropriately.

Of course the best companies need no such regulatory oversight; they will enforce their own rules. However, unless regulators are willing and able to enforce company-created rules, such as inspection frequency schedules, there is a real danger that these rules will be regarded as dispensable. I recall speaking to a plant manager who told me that he was under pressure to cut costs and that one way he could see to do this was to lengthen the period between inspections of certain items of equipment. There was no corresponding pressure to comply with the previously accepted inspection program. As a result he found himself beginning to wonder whether he was being too pedantic in sticking to the schedule and whether he could afford to prolong the periods between inspections without significantly increasing the risk. This man was back to case-by-case risk-assessments, made under financial pressure. This is the very situation rules are designed to avoid.

14.4 *Summary and implication*

This discussion of the pipeline standard demonstrates that there may be a complex hierarchy of rules and rule-makers, starting with legislation at the top, but unless the end point decision-maker is confronted with a prescriptive technical rule and unless there is some mechanism to ensure compliance with that rule, the relentless pressure to minimize costs is likely over time to erode commitments to safety.

Where regulators are, for whatever reason, unable to ensure that companies have devised appropriate rules and are complying with them, and where this is shown to have contributed to a major accident, the public may demand that more uniform and restrictive technical rules be included in the standards themselves or even in regulation. That is certainly one of the lessons of Texas City.[60] Evidently, those who advocate that companies should have the freedom to make their own technical rules should also be encouraging regulators to diligently monitor and enforce those very same rules.

15. Conclusion

This discussion has ranged widely over the relevance of prescriptive technical rules and has stressed the desirability of using such rules rather than risk-assessment to guide end point decision-making, whenever possible. Pure risk-assessment provides little if any

[58] A phrase used by one regulator at interview.
[59] A rather ambiguous phrase used by another regulator at interview.
[60] The "Eva Bill", introduced into the Texas legislature, proposed that all temporary buildings should be located at least 1000 feet from process units. See *Failure to Learn,* op cit, p151.

guidance to such decision-makers. This is not just my view. It is clear that in many quarters there is a quest for such rules. Importantly, this is not just a quest for rules that can be imposed on others; in some cases, people develop rules that they apply to themselves.

The argument about the importance of prescriptive technical rules seems more controversial in the case of non-operational decision-making, where significant sums of money may be involved. In this situation companies have an incentive to resist the more expensive risk controls and to argue for the less expensive, on the grounds that risk-assessment in their particular circumstances reveals that the cheaper course of action poses an acceptable risk. It is here that we see governmental authorities advocating stronger technical rules and endorsing standards that contain a wide array of specific technical rules, in order to manage major hazards more effectively.

This discussion enables us to draw some tentative conclusions about the circumstances in which it is appropriate that prescriptive technical rules be imposed on companies, in standards or possibly regulations, even within an overall risk-management framework. First, where industry good practice is agreed, it makes sense to formulate it as a clear rule so that laggards can be forced into line. The HSE's decision to interpret ALARP in this way demonstrates the point. We saw, too, that the gas pipeline standard contained some highly detailed prescriptive requirements in order to deal with laggards. Second, where regulators and others are seeking to nudge standards higher, as they were following the Buncefield explosion, the best way to do this may be to formulate the higher standard as a rule to be complied with. Finally, where the consequences of failure are potentially catastrophic, the public may not be willing to accept any risk analysis that treats a certain level of risk as acceptable. In these circumstances it may be necessary to devise rules to give expression to this concern. This was an explicit reason given by Buncefield Board for insisting on its automatic cutout rule.[61]

On the other hand, some of the big design and siting decisions made at the outset of a new project may be one-of-a-kind decisions. It may not be possible to convert risk assessments into technical decision rules in these circumstances.[62]

In Robens-inspired jurisdictions, where the overarching legal framework is one of risk-management, there has been a long term tendency to de-emphasise rule-compliance. Yet safety depends on both risk-management *and* prescriptive rule-compliance. The aim of this paper has been to demonstrate the inter-relatedness of these two concepts and to re-emphasise the importance of technical prescriptive rules.

[61] Buncefield MIIB, Final Report Vol 2, p 26.
[62] Decision making here may revert to numerical risk acceptance criteria.