national research centre for
OHS regulation

RegNet

Working Paper 89

How Much Should be Spent to Prevent Disaster?

A Critique of Consequence Times Probability

Professor Andrew Hopkins
Emeritus Professor of Sociology, Australian National University, Canberra
Andrew.Hopkins@anu.edu.au

November, 2014

Australian National University

About the Centre

The National Research Centre for OHS Regulation (NRCOHSR) is a research centre within the Regulatory Institutions Network (RegNet) at The Australian National University (Canberra), and operates in association with the Queensland University. The Centre monitors, documents and analyses Australian and international developments in work health and safety regulation and research, and related areas of regulation, and publishes a web-based series of working papers relating to work health and safety regulation. The Centre also conducts and facilitates high quality empirical and policy-focused research into work health and safety regulation, facilitates the integration of research into regulation with research findings in other areas of regulation, and encourages collaborating researchers to conduct empirical and policy-focused research into work health and safety regulation.

Address for NRCOHSR correspondence

National Research Centre for OHS Regulation
Regulatory Institutions Network
Coombs Extension, Building 8
Cnr Fellows and Garran Road
The Australian National University
Canberra, ACT, 0200
Email: nrcohsr@anu.edu.au.

Disclaimer

Abstract

Major hazard facilities, such as nuclear power stations or chemical manufacturing sites, frequently engage in cost benefit analysis (CBA) of measures designed to prevent rare but catastrophic events. These analyses are beset with difficulties, among them the difficulty of costing disastrous events and the difficulty of estimating their probability (likelihood or frequency in some models). But CBAs in these circumstances have a far more fundamental problem: they rely on a calculation of consequence times probability. I argue that this figure is quite irrelevant to the owner of a hazardous facility and of absolutely no use in deciding how much to spend on disaster prevention. I further argue that owners should stop thinking of catastrophes as chance events and think of them instead as *caused* by failures of control systems, and therefore preventable by ensuring that controls are working as intended.

Acknowledgements

I am extremely grateful to Jan Hayes and Peter Tuft for their advice and input into this article. I could not have written it without them. They do not agree with all my conclusions.

Feedback welcome

Comments, criticisms, corrections are invited. Please email: andrew.hopkins@anu.edu.au

Introduction

Major hazard facilities such as nuclear power stations, chemical manufacturing sites, gas processing plant and high pressure gas pipelines are relatively safe. Yet, no matter how safe they may be, the potential for disaster remains. Moreover, no matter how safe they may be, they can always be made safer, at a price. The question therefore arises: how much should the owners of such assets spend in the pursuit of greater safety? When is it reasonable to conclude that it is not worthwhile to seek to lower the risk further? In this paper I wish to explore how risk analysts go about answering these questions. I want to suggest that current methods are seriously flawed, in fact fatally flawed. In particular my claim is that calculations that rely on probability times consequence are fundamentally misleading for the owners of hazardous facilities. Finally I want to suggest ways out of the problem.

Broadly, the way these decisions are made is by means of a cost benefit analysis (CBA). This involves quantifying the risks, calculating the benefit of some risk reduction measure, usually in terms of the number of lives saved and the value of each such life saved, and then comparing the benefit with the cost of the risk reduction measure. If the cost outweighs the benefit, (or in some interpretations, is grossly disproportionate to the risk[1]) then the risk reduction measure is judged to be not reasonably practicable.

This kind of CBA poses methodological and ethical problems. But the end point of my argument is that, regardless of these problems, CBA in relation to high consequence low probability events is of absolutely no use to those faced with the decision of how much to spend on safety. This is clearly a controversial argument, so I will need to make it very carefully.

Consequence and Probability Matrices

Risk analysis usually involves hypothesising a scenario, such as an explosion, and making an assessment of both possible consequences and the probability of an event with such consequences. Risk is then determined by considering consequence and probability. This leads naturally to the idea of a consequence by probability matrix. It will be useful to begin by outlining some of the features and problems of these matrices, before introducing the main argument of the paper.

Figure 1 below is a risk matrix in use in the gas pipeline industry.[2] Some of its features relate specifically to pipelines but similar matrices are used in most if not all major hazard industries. The probability dimension in this case consists of five ordered categories. It is important to understand that whereas frequencies are often per annum, in the case of risk matrices for major hazard facilities the time frame is the life time of the facility. So, for example, "frequent" here means "expected to occur several times" in the life time of the facility, which might be 50 years. In addition to this qualitative description, each of these categories is associated with a specific numerical probability. This latter feature of the matrix is advantageous for present purposes. The probability dimension of risk matrices is often

---

1 Edwards v. National Coal Board.(1949) All ER 743 (CA).

2 It is taken from Australian Standard 2885, with the exception of the numerical frequencies. The latter are derived from an Australian Standards handbook and were added to the matrix by a practising risk analyst to facilitate quantitative cost benefit analysis.

treated as an ordinal scale only, without specific probabilities attached. But numerical probabilities are required for CBA calculations and, as Figure 1 includes numerical values, this matrix can be used for the type of CBA calculations to be discussed later in the paper.

		CATASTR-OPHIC	MAJOR	SEVERE	MINOR	NEGLIGIBLE
PEOPLE:		Multiple fatalities	Few fatalities, or several people with life-threatening injuries	Injury or illness requiring hospital treatment	Injuries requiring first aid treatment	Minimal impact on health
SUPPLY:		Long term interruption	Prolonged interruption or long-term restriction	Short term interruption or prolonged restriction	Short term interruption or restriction but shortfall met from other sources	No interruption or restriction
ENVIRONMENT:		Effects widespread, viability of ecosystems or species affected, permanent major changes	Major off-site impact or long-term severe effects or rectification difficult	Localised (<1 ha) & short-term (<2 yr) effects, easily rectified	Effect very localised (<0.1 ha) and very short term (weeks), minimal rectification	No effect, or minor on-site effects rectified immediately with negligible residual effect
FREQUENT	Expected to occur several times (≥ 10 events)	Extreme	Extreme	High	Interme-diate	Low
OCCASION-AL	May occur occasionally (0.1 - 10 events)	Extreme	High	Interme-diate	Low	Low
UNLIKELY	Unlikely to occur but possible (0.1% - 10% probability)	High	High	Interme-diate	Low	Negligible
REMOTE	Not anticipated for this pipeline at this location (0.001 - 0.1% probability)	High	Interme-diate	Low	Negligible	Negligible
HYPOTHE-TICAL	Theoretically possible but has never occurred on a similar pipeline (<0.001% probability)	Interme-diate	Low	Negligible	Negligible	Negligible

Figure 1: Risk Matrix

Although the probability and consequence dimensions are in principle quantitative variables, in practice they are usually treated as ordered categories only, as in figure 1. The outcome of considering consequences and probability cannot therefore be a number but is simply another ordinal variable, in this case: negligible, low, intermediate, high, extreme. In figure 1 the risk contours - lines of equal risk - run diagonally from bottom left to top right, so that low consequence high probability scenarios have roughly the same risk ranking as high consequence but low probability scenarios (intermediate or low). I return to this in a moment. The way these particular rankings are determined is not generally explained and they often appear to be based on the judgement of the matrix designer.[3]

An important feature of the risk matrices is that they reduce two independent variables to one: probability and consequence are reduced to the single variable – risk. This is seen to be a useful simplification because it prioritises threats clearly and indicates to decision makers unambiguously (although not always appropriately) where their attention needs to be focussed. But collapsing two dimensions into one involves an unavoidable loss of information. It treats low consequence high probability events (eg first aid events) as equivalent to high consequence low probability events (such as explosions). This disguises the distinctive nature of high consequence, low probability events. Worse than this, high consequence events that are judged to have very low probability are treated as intermediate level risks rather than extreme. This has important practical consequences. Some companies have a policy that the greater the risk, the higher is the level in the organisation at which a decision about its acceptability must be made. For example, extreme risks may need to be considered and accepted by the CEO or the Board of Directors, while intermediate risks may need to be considered only at site manager level. In other words the fact that scenarios with catastrophic consequences can be classified as intermediate risks means that they do not come to the attention of the most senior people in the organisation, who are the very people most able to do something effective about them. A case in point was the BP Gulf of Mexico blowout which killed 11 men in 2010. For various reasons the risk that a blowout might kill numbers of people had been judged to be "low". According to BP's policy this meant that the decision about its acceptability was made by the most senior BP person on the drilling rig a – a foreman! (Hopkins 2012, p 68) It can be argued that risk acceptance for extremely high consequence events, no matter how improbable they may be, is a matter for boards of directors; in other words, that the determinant of where in the organisation decisions are made should be the potential consequence, not risk. I return to this point later in the paper.

One of the confusing things about the concept of risk, as used by risk analysts, is that it is used in two quite different ways: one, risk as consequence and probability, and the other, risk as probability alone, for example the individual risk of death per annum (IRPA). In a sense the second is a special case of the first – it holds consequence constant (individual death) and deals with variations in probability. For example, the widely quoted UK Health and Safety publication R2P2 specifies that a risk of death greater 10^{-3} is intolerable and must be reduced no matter what, whereas a risk in the range of 10^{-3} to 10^{-5} must be reduced only if it is reasonably practicable to do so (HSE 2001). Risk is being used here as a synonym for probability. It is strange that practitioners accept this ambiguity without comment. To reiterate, in this paper I am using risk in the more general sense of consequence and probability.

---

[3] For a mathematical critique of these qualitative risk matrices see Cox (2008), and Pickering and Cowley 2010. Cox (2008, p 500) argues that where probability and consequence are negatively correlated, as they are here, these matrices are "worse than useless".

Before moving on it is worth noting that risk matrices are used not only for risk assessing hypothetical scenarios, they are also widely used for assessing the significance of incidents that have occurred. Take the following example recounted to me by the manager of an outback production facility. A company vehicle, travelling on a road near the facility, had hit a pothole and swerved to the other side of the road before the driver recovered control. The incident was reported to the manager, who in turn had the job of determining a risk score. The facility was in a remote location and the chances of hitting an oncoming car were very slight. But if there had been a collision, it might have resulted in a fatality, he said. If he took this into account the matter would receive a relatively high risk ranking and would need to be reported to corporate headquarters. How was he to assess the risk associated with the incident?

There is a fundamental problem with risk assessing incidents in this way. The matrix is designed for risk assessing hypothetical events. But reported incidents are not hypothetical; they have occurred. Furthermore, in the typical case, no harm has occurred. Strictly speaking it makes little sense to carry out a risk assessment for an incident that has already occurred and which itself caused no harm. One can however treat the incident as having identified a hazard – the pothole. One can then imagine a scenario, namely a car hitting the pothole and swerving into the path of an oncoming car and killing someone. This imaginary event can then be sensibly risk-assessed. On the other hand, such a risk assessment seems like overkill. A hazard has been identified, it is easily eliminated by filling the hole and it would make sense to do so without the need for further analysis. This idea of bypassing the risk assessment process and focusing on the implementation of obvious controls is something I shall return to later. The main point here, however, is that any attempt to risk assess an incident using a risk matrix is misconceived and leads almost inevitably to confusion.[4]

Cost-Benefit Analysis

The discussion so far has been preliminary to the main topic of the paper – cost-benefit analysis of measures designed to prevent rare but catastrophic events. Generally speaking what makes a scenario catastrophic is the large scale loss of life. Of course in some situations, such as the Gulf of Mexico oil spill, the loss of life is overshadowed in the public mind by the scale and cost of the environmental damage. But for present purposes I shall take loss of life as the most significant consequence of major accident events. To carry out a cost-benefit analysis therefore requires that a value be placed on human life. Placing a monetary value on human life raises numerous ethical and social issues. These are so serious as to call into question the whole enterprise (Heinzerling and Ackerman 2002).

. However I shall not canvas these issues here because it would be too much of a diversion from the main argument to be made in this paper.[5] For the sake of that argument I shall take the value of life to be $5 million.[6]

---

4 For a more extended discussion see Hopkins (2009, pp42-46).

5 One remarkable example of where this thinking can lead is provided by Harvard Professor Kip Viscusi, who undertook a cost-benefit analysis of smoking. He concluded that states *saved* money as a result of smoking by their citizens! This was because smokers died early, thus reducing the expense of providing nursing homes and other services for the aging. On the basis of this analysis he suggested that cigarette smoking should be subsidised rather than taxed! (Viscusi 2005)

The other number which is important in any CBA is probability. As noted above, probability in figure 1 refers to the probability of a scenario occurring during the lifetime of the facility. This scale is basically an ordered set of categories including descriptors such as "unlikely to occur but possible", to which a set of numerical equivalents has been be attached. The descriptors are quite vague. In particular the category "hypothetical" talks about "similar pipelines". The word "similar" is critical here. There may have been several disasters previously that are similar in some respects but not others to the scenario under consideration. However, if by similar one means similar in all respects it may be easy to conclude that the particular scenario to be assessed has never occurred before. This problem is generic to the probability dimensions of risk matrices. The matrix in use by BP at the time of the Gulf of Mexico blowout had 4 categories defined in part as follows:

Very low - no comparable occurrence is known;

Low - comparable events are known to have occurred in the past

Moderate - comparable events are within the team's experience

High - comparable events are frequent (Hopkins 2012, p 183).

Clearly a great deal hinges of the meaning of "comparable" which, being an undefined term, leaves plenty of scope for arbitrary judgement.

To return to figure 1, in addition to these ambiguities in the qualitative descriptors, when it comes to the numerical equivalents there is considerable arbitrariness about the ranges chosen, particularly at the low probability end of the scale. This arbitrariness will necessarily find its way into any CBA. But as with the value of life, I shall not dwell here on the problems of the probability scale, but rather take these categories at face value. I do this so as to be able to focus on what I see as the fundamental problem of CBA in the context of the decision-making about the prevention of rare but catastrophic events.

To develop the argument let us now consider a specific and very simplified CBA. Suppose the scenario of concern is one that could cause 40 deaths. At $ 5million per death, the total cost of this incident is $200 million (ignoring costs other than death).

Suppose the probability of such an event is one in a million. (In the "hypothetical" category in figure 1, <0.001% equates to $<10^{-5}$, so 10^{-6} lies squarely in this category). The CBA process requires that the cost of the scenario be multiplied by the probability. Multiplying $200 million by 10^{-6} gives us a figure of $200. The big question is: what meaning can be given to the figure obtained in this way and to what use can it legitimately be put?

---

[6] Those who defend the method sometimes argue that they are not valuing life, but rather, statistical life, since what is at stake is not the saving of identifiable lives but rather reducing the risk of death slightly for each of a large number of people (Wolff 2011). However the overall result of a slight risk reduction for each of a large enough number of people is indeed a small reduction in the number of actual deaths. In these circumstances the distinction between real life and statistical life seems illusory. Of course in these circumstances there is no way of knowing which particular lives were saved, but the claim must be that some lives were saved, otherwise there is no point in spending money at all. However, the distinction makes more sense if we are talking about a slightly reduced risk of death for each of a small number of people, since then the reduction in risk may amount to only a fraction of a life saved, which of course is an impossibility.

To answer this question we need to be clearer about what it means to say that the probability is one in a million: one in a million "whats"? The initial statement, made more carefully, is that the probability is one in a million that a facility will "blow up" during its lifetime. (I will use "blow up" as shorthand for experience a catastrophic event such as described above). So the "whats" above are facilities, and the probability statement becomes: "It is probable, ie we expect on average, that one in every million such facilities will blow up in its lifetime." An average of 1 is difficult to work with, so let's multiply by 1000. The probability statement now becomes: "We expect on average that of every *billion* facilities, a thousand will blow up during their life time." In reality there will never be anything like a billion facilities or even a million so we can never test these claimed probabilities against any real data, but it is the logic of an argument I am following here.

Now here is one way to give meaning to the $200 figure. Suppose I am a super insurance company, responsible for insuring all one billion facilities against blow-up. What will I charge them as a premium?

I know that the cost per blow-up is $200 million (assuming the above example is typical). And of my one billion clients, approximately 1000 of them will blow up and require a payout. So the total payout for my clients is $200 million x 1000 = $200x$10^9$. What do I need to charge each of my clients to cover this cost? Answer: $200x$10^9$ divided by 1 billion = $200

Of course I will need to charge a bit more to cover the possibility that more than 1000 facilities will blow up during their lifetime. Perhaps 1200 or 1300 will blow up. But if the probability statement can be taken at face value, the expected number is 1000. I will also need to charge more to allow for profit. But the figure of $200 is the break-even figure.

The conclusion here is that if I am the super insurer in the position described above, it makes sense to multiply consequence by probability to derive a figure that will be used as the basis of an insurance premium. In this highly artificial context the $200 provides a guide to action for the super-insurer.

We can frame this super-insurer argument a little differently if we substitute *society* for the super-insurer. This involves personifying society in a way that is highly questionable, but I leave that aside for present purposes. Society has an interest in preventing the single blow-up costing $200 million, not because it is a disaster for those concerned, but simply because it is a loss to society. Assuming one million facilities, if society can eliminate the risk by requiring each facility to spend an additional $200, that will be value for money, just. Why do I say this? The assumption here is that the cost is passed on to the rest of us, as consumers of the facility's products. So we, society, will be spending $200 million to prevent a loss of $200 million. Arguably that is just worth our while. But it will not be cost effective to spend more. If the risk can only be eliminated by spending say $2000 per facility, society will be better off accepting the loss of one facility, and the associated deaths, and putting the money to some other use, for example medical research, or even risk reduction in some other area of human activity. As members of society we have an interest in the million facility owners being compelled to spend up to $200 to eliminate the risk, *but it is not in our interest that they spend more.* $200 is the "maximum justifiable spend" *from the point of view of other members of society*. ("Maximum justifiable spend" (MJS) is a standard term in use in CBA).

Now imagine a single facility owner. Can the $200 figure be given a meaning and/or be used as a guide to action?

We can answer this by reverting to statistics for a moment. What is the *statistical* meaning of the $200? The cost of an actual blow-up is $200 million. Statistics talks about the *expected* cost given a certain probability distribution. In simple terms the expected cost is the weighted average of the cost across all the facilities. So if one facility blows up and experiences a cost of $200 million, while all the others avoid blow-up and hence have no cost, the "expected cost", that is, the average cost across the million facilities is $200. How useful is this knowledge to the single facility owner? The answer is: not at all. The average across all million facilities give no guidance about what the single facility owner should spend to prevent a blow-up.

To see why, let us suppose that all agree that the risk cannot be eliminated for $200 and that it is as low as it can be without spending a great deal more money. Is it sensible for the facility owner to simply accept this risk on the basis that it is not in society's interest to spend more than $200? Such an owner is in a very different position from society. If the facility blows up it will be a disaster for those directly involved and may destroy the company. For society on the other hand, $200 million is a completely insignificant figure when set against the total value of all economic activity. The stakes in other words, are quite different. It is therefore not sensible for the facility owner to simply accept this risk on the basis of what is in society's interest. The individual hazardous facility owner might reasonably take the view that it is worth spending a good deal more than the expected cost to eliminate the risk of a blow-up killing 40 people.

In support of this last statement, consider an analogy suggested to me by a colleague. Suppose the chance of having a write-off accident in my car is one in a thousand, that is, 10^{-3}, and my car is worth $20,000. The break-even insurance premium to protect me against this loss (the maximum justifiable spend) is $\$20{,}000 \times 10^{-3} = \$20$.

From my colleague's perspective, losing $20,000 is pretty significant; in fact it would be a disaster for him, in a relevant sense. Accordingly, he told me, he would probably be prepared to pay $200 per year for insurance, despite knowing this is 'over the odds'. But if the insurance company tried to charge $2000, his view was that he would carry the risk himself because the cost of the risk mitigation measure (the insurance) was too high when compared to the benefit, he said.

Let us consider this analogy. My colleague's position is that if it takes $200 to insure his car and so avoid financial disaster, then he will pay. This is because he simply cannot afford to run the risk of losing $20,000. So he will pay whatever it takes, although there is a limit of $2000. Why does he stipulate this limit? He did not say, but we can speculate. A premium of $2000 per annum would mean that in 10 years he had paid out the total value of the car in premiums (I am assuming for the sake of the discussion that the value of the car does not change over this period.) In this way what is intended to be an insurance against loss of $20,000 has become a certainty that he will lose $20,000 over a 10 year period. That makes no sense, so of course there is a limit to what he is willing to pay in premiums. But what is interesting about this situation is that the amount he is willing to pay to avoid disaster bears no relationship to the break-even insurance value; it is largely determined in this case by the total cost of the disaster.

The analogy with the single major hazard facility may not be complete, but it is illuminating. It suggests that it is in the interests of the owner of such a facility to pay whatever it takes to avoid disaster, subject to the limitation that this payment remains a small fraction of the total cost of disaster, should it occur. There is no obvious way of deciding just what this limit

should be, but one thing is clear: the single facility owner's decision about how much to spend will be quite independent of the "expected value" or "maximum justifiable spend".

As a matter of fact, risk analysts generally concede that the single facility owner should not limit spending to the expected value. The law in many countries requires them to spend whatever it takes, provided that the expenditure is not "grossly disproportionate" to the risk. To take account of this, risk analysts sometimes add a "proportionality factor" to the calculation,[7] a number ranging from 1 to 10. The result is that the "maximum justifiable spend" increases by a factor of up to 10. However this in no way alters my argument. If we start with an irrelevant figure and multiply it by 10, the resultant figure is still irrelevant.

An Alternative Way of Doing the CBA - Implied Cost of Averting a Fatality

THE MJS method requires us to multiply probability by consequence to get an expected cost, which I have argued is irrelevant for present purposes. There is another way of doing the calculations that appears to avoid this requirement and even avoids putting a value of life. This is to calculate the Implied Cost of Averting a Fatality (ICAF). In the next paragraphs I want to consider whether this method really does solve the problem.

ICAF starts with the cost of a proposed risk reduction measure and asks, ultimately: is it worthwhile? To avoid the problem of an expected frequency of less than one, let us suppose we have a million facilities. And suppose there is a measure costing $2000 per facility which, if implemented, will eliminate the risk described above. The expected number of blow-ups prevented is then 1 and the expected number of lives saved is 40. The total cost to all facilities is $2000 million. So the cost per life saved is $2000 million divided by 40 = $50 million. If we treat this cost to all million facilities as the cost to society, because ultimately it is passed on to customers, we can say the cost to society is $50 million per life saved. This is the cost of averting a fatality implied by the particular risk reduction measure.

So far so good. We have not multiplied consequence by probability. Nor have we put a value on life. But where do we go from here? From society's point view we can only progress the argument by putting a value on life. Given that a typical value of life is $5 million and that the implied cost of this measure is $50 million per life saved, we would conclude that the proposed risk reduction is not value for money from a societal point of view, that is, we'd all be better off accepting a blow-up with 40 fatalities, rather than spending the money to avoid this outcome. So from society's point of view, the end point of the ICAF approach is the same as for the MJS approach.

Suppose now we look at this from the point of view of the single facility owner. Again the stakes are quite different. While for society the scenario under discussion is of infinitesimal significance, for the facility owner it is a disaster. But we are back to the problem of very low probabilities. The question for the facility owner is: do I want to reduce the risk from a very small number to zero? If I do nothing I will almost certainly be fine; but if my luck runs out,

---

7. For example, for offshore installations in the UK the regulator suggests a proportionality factor of 6 (see HSE 2005)

the outcome is almost impossible to contemplate. There is no way the ICAF can give me guidance on this decision. Again, the end point is just as it was for the MJS method.[8]

There is one way in with the ICAF can be an aid to decision-making, without putting a value on life, and that is, identifying spending priorities. Let us change the question from: "How much is it worth to eliminate the risk?" to: "Given that I have a certain amount of money to spend on risk reduction, how can this be most usefully spent?" The lower the ICAF of a proposed risk reduction measure the larger will be the number of lives saved for any given level of expenditure. The ICAF therefore provides an indication of how money earmarked for safety can be most effectively spent. Moreover it does so without at any point placing a value on life. Using ICAF in this way therefore avoids the moral problems I have alluded to but not discussed in this paper.

Consequence-Based Decision Making

Let us revert to the main focus of this paper, the individual hazardous facility owner faced with the question of how much to spend to prevent a particular disaster scenario. If we accept that consequence times probability provides no basis for decision-making, what is to be done? One possibility is to make decisions on the basis of consequence only. This approach repudiates any idea of acceptable risk. The philosophy is that, if the consequences are severe, then people must be protected from them, no matter how unlikely they may be. BP moved to consequence-based decision making in relation to the siting of occupied portable buildings (eg trailers) following its 2005 Texas City explosion. That explosion killed 15 people who were housed in trailers located close to the blast source. BP decided that henceforth any building located in a potential blast zone must be built to withstand the blast pressure. To implement this principle, it divided each refinery into zones, based in part on calculations of the possible blast pressures: a red zone, in which no buildings may be located; an orange zone in which buildings may be located if they are designed to withstand the blast pressures that are possible in that zone; and a yellow zone beyond that(BP Report). An immediate consequence of this was that trailers were restricted to the yellow zone.

The principle of consequence-based decision-making is not one that can be universally applied. In many, if not most, situations it is not possible to protect people from all risk. However in the case of occupied portable buildings the risk to occupants can easily be eliminated by remote siting. What makes it relatively easy to adopt the principle in this case is that the financial cost of moving trailers out of range is not significant; all that is really at stake is questions of convenience.

---

[8] I reserve to a footnote the question of whether the ICAF method avoids having to multiply consequence by probability. According to the method, a risk reduction measure (RRM) will be justified if the following inequality holds:

ICAF< VoL (Value of Life)

i.e cost of RRM/(probability X no of deaths) < VoL

i.e cost of RRM < probability X no of deaths X VoL

In our example, no of deaths x VoL is the cost of the blow-up, that is, the consequence.

So, RRM is justified if cost of RRM < probability X consequences.

In short the ICAF method requires us to calculate probability X consequence.

But even where the cost is high, it is still possible to adopt a consequence-based approach. The issue of occupied buildings (not just temporary or portable buildings) has been exercising the minds of board members of a very large company I have been working with. They have instructed company officers to reduce exposure to explosion hazards by relocating occupied buildings, even though this will be extremely expensive. Furthermore, some of the company's constituent businesses are more affected than others and would bear a heavy cost burden if required to fund this relocation themselves. Accordingly the company has decided to make the funds available centrally.

A final example of consequence-based decision making relates to explosives. In manufacture and storage of explosives, the military use the concept of 'safeguarding distance'. This determines layout of sites including issues of separation of occupied buildings from explosives storage and plant, as well as separation between parts of the plant to prevent escalation. The safeguarding distance is determined by possible consequences of an explosion and is independent of the probability of explosion.

A Strategy for Risk Analysts

CBA is intended to be a tool to aid top decision makers. But the guidelines and standards routinely say that CBA should not be the only input into their decision-making. Decision makers are exhorted to take account also of less quantifiable matters such as reputational risk and perhaps public outrage. In practice they tend not to, for the following reason. The result of the CBA process for a high consequence low probability event is often that the risk is already as low as reasonably practicable, meaning that it is not worth spending a significant amount money to reduce it further. From the point of view of the busy decision-maker, the risk analyst has provided technical advice which the decision-maker is happy to accept. In this way the CEO and board may never get to consider the consequences of the hypothesised failure. Risk analysts may feel frustrated that they are in effect being used to make decisions which rightfully should be made much higher in the organisation, but they feel powerless to do much about it.

Here is a suggested solution. First, analysts should avoid summarising their analysis in a single indicator: consequence times probability. They should provide information about consequences and probability separately, with information about probability being in qualitative rather than quantitative terms. This has two advantages. It means that top decision makers are not presented with a decision "ready-made" by the risk analyst; they must therefore apply their minds to making decisions themselves. It also avoids the erroneous (I have argued) implication that consequence times probability is of relevance to people making decision about single facilities.

Information about consequences should include a description of the hypothesised event, and the likely damage to the business, to public property and to the environment. If the environmental damage has implications for people's livelihoods this should also be noted. Finally information should be provided on likely numbers killed and injured, together with relevant details, such as whether children attending a nearby school are expected to be among the victims. There would be no need to monetise the value of life. The documents or standards that lay out this procedure should invite decision makers to pay attention also to the

reputational costs, legal costs and personal costs to top company people subject to legal scrutiny after a major accident.[9]

In addition to outlining consequences in some detail, analysts should identify various measures aimed at eliminating the threat or reducing the severity of the consequences, together with some indication of costs.

It will also be valuable for analysts to identify and describe accidents around the world broadly similar to the one under discussion.

Finally, the documents or standards laying out these processes should make it clear that decisions in relation to very high consequence events should be taken at the highest corporate level. One guiding principle might be that the decision maker should be sufficiently senior to be able to authorise the most costly of the consequence reduction measures identified by analysts.

An Alternative Paradigm

So far I have argued that probability times consequence is not a relevant number for the single major hazard facility owner and that traditional CBA methods cannot therefore be used to determine how much to spend on hazard control. This leaves top decision makers in a difficult position. I want now to suggest a solution to this problem, one that involves abandoning the whole idea that major accidents are simply chance events occurring with a certain probability.

The principle strategy for the prevention of major accident is defence in depth, which means putting in place a series of defences or controls to guard against a major accident event. The corresponding model of accident causation is that accidents occur when all these defences fail simultaneously. This is nicely depicted in the Swiss cheese accident model (Reason 1997): each defence is a slice of cheese with holes in it, representing imperfections, and accidents only occur when the all these holes line up (see Figure 2 below).

---

[9] A vivid example of what can happen can be seen at http://www.texascityexplosion.com/site/interfacerel2?id=1

See John Manzoni – Global Refining Group Vice President (London) Three Little Pigs and The Big Bad Wolf. In this video the deputy CEO of BP at the time of the BP Texas City accident is seen being interrogated about the company's CBA method which appears to equate human life with that of the pigs in the story of the three little pigs. Here is what the lawyer asking the questions had to say on his website:

"BP was even willing to take the callous step of assigning a value to human lives for the purposes of Cost Benefit Analysis. What's more, the analysis was weighted to favor the cheapest method, not necessarily the safest. BP even went so far as to compare their workers to the three little pigs in the famous children's story."

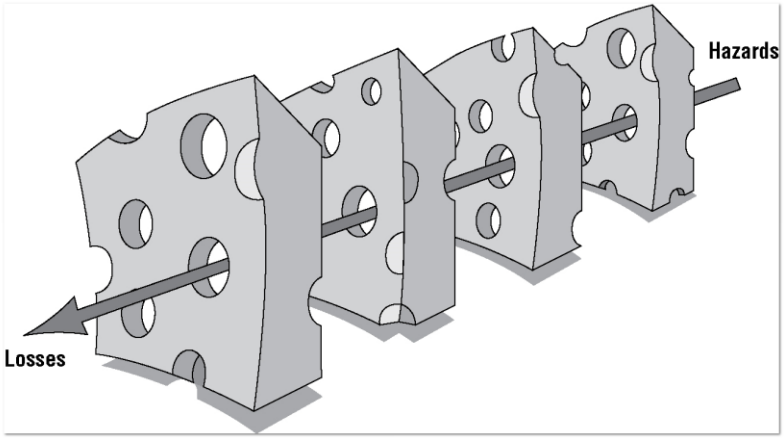

Figure 2 Swiss Cheese model of accident causation.

If we stick with probability for a moment we can say that, if one of the barriers is absent or non-operational for any reason, the probability of a major accident event is correspondingly higher. To take an example, if one of the controls is regular inspection of pipes and tanks for corrosion, yet inspections are not being done, we can infer that the conditional probability of an accident - the probability given that this control is not in place - is much higher. The more defective our system of controls the higher is the probability of accident. On the other hand, if our controls are all being implemented as intended, the probability of an accident is zero. That is the matter of logic. Putting this another way, major accidents can be prevented by ensuring that all controls are working as intended. From this perspective, accidents are no longer chance events, they are caused by control failures which are within the power of the facility owner to prevent. The word "cause" is used here in the sense of *but for* cause. The point is that in relation to each of these control failures we can say that *but for* this control failure that accident would not have occurred. This way of thinking shifts the question for top decision makers from "how much is it worth spending?" to "what are the critical controls that are supposed to be in place and how can I guarantee that they are"?

This approach does not eliminate financial considerations. Some controls are more reliable because once in place they remain in place – for example replacing hazardous equipment with less hazardous - but they may be expensive. Other controls, such as procedural and engineering controls, may be cheaper, but require more on-going effort to ensure that they remain effectively in place. Decision-makers must therefore understand that if they economise on the more reliable controls, they will need to be more vigilant with respect to the controls they in fact adopt. The history of major accidents indicates just how unreliable procedural controls can be (Hopkins 2008). Decision makers need to be armed with this understanding in order to make the best decisions about how to prevent major accidents.

This alternative approach has implications for the issue of ALARP – "as low as reasonably practicable". In many jurisdictions the ultimate requirement imposed on companies is that they ensure the health and safety of work so far as reasonably practicable, or words to that effect. Risk analysts understand this to mean that risks are as low as reasonably practicable. As we have seen, CBA translates this to a question of whether the costs outweigh the benefits. However, using the defence in depth approach, the question is: have we done all that is reasonably practicable to ensure that the required defences are in place? In legal proceedings that follow a major accident, this is always the question that courts focus on, and they almost invariably find that controls that were supposed to be in place were missing.

From this point of view, the Swiss cheese approach is a far more useful guide for decision-making than CBA.

Conclusion

This paper is concerned with cost benefit analysis in a very limited context, namely, cost benefit analysis for measures aimed at reducing the risk of rare but catastrophic events. I have not attempted to generalise beyond this context and there are certainly other areas in which this critique does not apply.

CBA in the present context is controversial because of the way in which it places a financial value on life and also because of the uncertainty of the probability estimates that are made for rare events. But this paper has not been concerned with those questions. I have focussed instead on a very fundamental feature of these CBAs - the fact that they require the analyst to multiply consequence by probability – and I have argued that for the individual facility owner, the result is a meaningless number, in fact that it is worse than useless because it allows top decision makers to avoid making the hard decisions.

It is far better to present top decision makers with information about consequences and probability separately, rather than combining them into to single indicator of risk. Furthermore, rather than thinking about accidents as chance events, it is better to think of them as caused by control failures. They can therefore be prevented by ensuring that controls that are intended to be place really are.

References

BP Report, *The Report of the BP US Refineries Independent Safety Review Panel*, Washington, January 2007, p 53.

Cox L, "What's wrong with risk matrices?" (2008) 28(2) *Risk Analysis*, 497-512.

Heinzerling L and Ackerman F, *Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection*, Georgetown University Law Center, 2002.

Hopkins A, *Failure to Learn: The BP Texas City Refinery Disaster*, CCH, Sydney, 2008

Hopkins A (ed), *Learning from High Reliability Organisations*, CCH, Sydney, 2009

Hopkins A, *Disastrous Decisions: The Human and Organisational Causes of the Gulf of Mexico Blowout*, CCH, Sydney, 2012.

HSE (Health and Safety Executive) *Reducing Risks, Protecting People, HSE's Decision Making Process*, HMSO, Norwich, 2001.

HSE, Offshore Installations (Safety Case) Regulations 2005, Regulation 12. Demonstrating Compliance with the Relevant Statutory Provisions, HSE Offshore Information Sheet No. 2/2006.

Pickering A and Cowley S,"*Risk matrices: implied accuracy and false assumptions"* (2010) 2(1) *Journal of Health and Safety Research and Practice*, 9-16.

Reason J, *The Organisational Accident*, Ashgate, Aldershot, 1997.

Viscusi W (Kip), The Value of Life, Vanderbilt Law and Economics Research Paper No. 08-04. Available at SSRN: http://ssrn.com/abstract=827205.

Wolff J,"Five types of risky situations"(2011)(2) *Law Technology and Innovation*,151-163.