Describing Risk – Variations on a Theme

J.R. Inge, MEng PGDip(SCSE) CEng MIET MAPM; Ministry of Defence; Bristol, UK

Keywords: risk assessment, risk management, risk matrix, notional estimated loss.

Abstract

Risk assessment is a fundamental part of safety management but one that generates a huge amount of debate. There are many equally valid, but different, ways of assessing and presenting information about risk. This paper looks at the basics of what we are trying to achieve and highlights some of the variables in the process. It also examines the potential for conflict between various drivers for optimisation when considering risk management from different viewpoints. It also introduces the concept of notional expected loss as a potential mechanism for comparing risks between projects and providing a valid description of an organisation's safety risk level.

Introduction

This paper is about the role of risk analysis in safety management. It examines this context, deriving a number of viewpoints for looking at risk information. These viewpoints are then used to focus a discussion on the variations in how risk can be examined and presented to support the safety management effort, and the pitfalls and problems that can arise. Finally, it presents notional expected loss as a tool to help management of risk at an organisational level.

The point of safety management is basically to make stuff safe, but people differ in their interpretation of just what 'safety' means. It can be defined in absolute terms, such as "*freedom from accidents or losses*" [1], but absolute safety is impossible to achieve. For practical management we need a scale on which to measure the degree of freedom from harm, so that we can define what is safe enough. Standards such as IEC 61508 therefore use the concept of risk, with safety defined as the "*absence of unreasonable risk*" [2]. A more usable working definition of 'safe' is given in Defence Standard (Def Stan) 00-56: "*Risk has been demonstrated to have been reduced to a level that is ALARP[1] and broadly acceptable or tolerable, and relevant prescriptive safety requirements have been met, for a system in a given application in a given operating environment*" [3].

The Def Stan 00-56 definition breaks into four components:

- Making the level of risk as low as is reasonably practicable. i.e. ensuring that reasonable measures have been taken to reduce the risk, and that there are no ways of reducing it further without grossly disproportionate sacrifice.
- Making the level of risk low enough that is either broadly acceptable, or can be tolerated in order to gain some benefit.
- Meeting prescriptive safety requirements. Legislation is a source of such requirements, but other prescriptive requirements may reflect deterministic functions that are necessary to ensure safety, e.g. to allow proper interfaces with other systems.
- Defining the scope of the risk, in terms of the system, application and environment involved.

From this definition, it follows that being able to analyse the risk posed by a system is an important part of being able to manage its safety, and this paper focuses on how this might usefully be done. However, risk is not the only relevant consideration. There are other elements of safety that are prescriptive or deterministic rather than risk-based, or that depend on a subjective judgement of tolerability or reasonableness. Another key point is that just as the level of risk may change with the scope being considered, so may these subjective judgements.

---

[1] As Low As Reasonably Practicable.

What are we trying to achieve?

Safety risk management is about figuring out what could go wrong (risk assessment) and doing something about it (risk control). It is a process of systematically identifying hazards and potential accidents, working out the detail of the accident sequences that link them, estimating the risk that they pose, assessing the acceptability of that risk and mitigating the risk to the point that the residual risk can be accepted. This is shown diagrammatically in Figure 1.

This paper concentrates on the interface between the risk estimation and risk evaluation parts of the process, in particular on how we present risk estimates that let us figure out how bad the problem is. The definition of safety gave us two reasons to examine risk: to contribute to the assessment of whether or not risks are ALARP, and to determine whether the risk is one that we can tolerate.

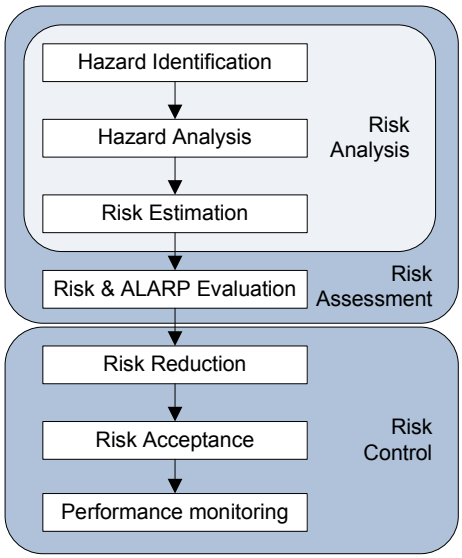

Figure 1. The Risk Management process

The requirement to reduce risks to a level that is ALARP is derived from the UK Health and Safety at Work etc. Act 1974 (HSWA). The requirements of the Act include "*Provision and maintenance of plant and systems of work that are, so far as is reasonably practicable, safe and without risks to health*" and other similar clauses that include the wording "*so far as is reasonably practicable*" [4]. Case law interprets 'reasonably practicable' to be a narrower term than 'physically possible', implying that there must be a balance between the magnitude of a risk and the sacrifice necessary to avert the risk. This balance must be weighted towards safety: the risk may only be allowed to remain if the cost of removing it is grossly disproportionate (in terms of money, time or trouble) [5].

Estimating a measure of the level of a risk does not tell us whether or not that risk is ALARP. That does not depend on the absolute level of risk, but on whether or not it is reasonably practicable to reduce it further – and to understand that we need to use engineering judgement to identify if any options are available. Measuring risk does however give us a metric of the potential benefit of a safety improvement. We can use this metric in cost-benefit analysis when we try to judge what is reasonably practicable, but it is the change in risk that is more important than the absolute level.

Although there is an obligation under the HSWA to reduce all risks so far as is reasonably practicable, there becomes a point where some risks are insignificant or so low that they may be considered broadly acceptable. This concept is called the 'de minimis threshold' in GEI-STD-0010 [6]. While acknowledging that the legal obligation persists, the UK Health and Safety Executive (HSE) advises that they would not usually require further reduction of such risks [7]. At the other end of the spectrum, some risks are unacceptable, whatever the benefit. Between these levels, some risks may be tolerated, if doing so gives sufficient benefit. The HSE explains that tolerance is not the same as acceptance: "*It refers instead to a willingness by society as a whole to live with a risk so as to secure certain benefits in the confidence that the risk is one that is worth taking and that it is being properly controlled. However, it does not imply that the risk will be acceptable to everyone*"[7].

Assuming a risk is ALARP, we need to know about its magnitude in order to decide whether we should tolerate it. The decision is not simple and as the level of risk increases, we are likely to demand increasing levels of benefit to justify the necessity and increasing assurance that the risk is properly controlled. The need for assurance leads to another viewpoint of risk: the potential risk. Normally the risks posed by a system ought to be reduced to a level that is both ALARP and tolerably low before it is put to use. What is needed for assurance about the system does not therefore depend on the level of these reduced risks, but on the risk that would be posed if the controls and mitigations were not present.

In practice, safety management is constrained by available resources. Affordability is not an excuse for failing to carry out risk reduction (an organisation should not carry out an activity that it does not have the budget to make sufficiently safe), but there is a business imperative to make the best use of the resources we have. As well as examining the cost-effectiveness of different methods of improving safety, we should prioritise the way we address risk. Analysis should focus first on the higher-level risks. There is no point in spending lots of money investigating mitigation for risks that are already low, only to find that the enterprise must be abandoned due to high risks that cannot be reduced to a tolerable level. It is therefore useful to be able to compare risks against each other, to determine which is most important.

This discussion highlights four viewpoints for information about risk, which will be examined in detail in the following sections of the paper:

- Absolute levels of risk, to judge tolerability against a standard.
- Relative levels of risk, to allow ranking and prioritisation.
- Changes in level of risk, to assess the impact of different options and hence their cost-effectiveness or reasonableness.
- Potential levels of risk, to determine the appropriate level of assurance.

Absolute risk

So far, we have considered 'risk', or rather its absence, as a kind of proxy measure for safety, but have not actually defined it. 'Risk' is typically thought of as a combination of the chances of something bad happening and the magnitude of just how bad that thing is. In the terms of Def Stan 00-56, it is a "c*ombination of the likelihood of harm and the severity of that harm*", where 'harm' is "*death, physical injury or damage to the health of people, or damage to property or the environment*" [3]. Some definitions include positive risk – beneficial events – but most assume that the event is harmful in some way, and this is generally the case in safety risk management

Beyond this simple understanding of risk there are many nuances of definition, but the key point is that risk is a vector quantity, made up of two parts: severity and likelihood[2]. The severity part can be measured on many different scales, depending on the type of harm you are interested in. Some definitions constrain it purely to physical harm to humans, but many common definitions expand it to other types of loss. The likelihood part can either be a probability (expressed as a number between 0 and 1) or a frequency (which can exceed 1). In either case, the likelihood needs to be bounded by a scope: the probability is the probability of an event happening in a given duration, for a given population. The combination of measures for severity, likelihood and scope come together to set the units of risk, e.g. deaths per year per 10,000 people, or critical accidents per operating hour per system. I shall refer to this combination as the basis for the risk.

For managing the risk posed by a system, the basis ought to be set according to the application, considering the underlying physical processes that could contribute to potential accidents. For systems that operate fairly consistently for long periods of time, it may be appropriate to consider a rate of accidents per calendar year. For others, the rate of accidents may be proportional to the usage, so a rate per mile travelled or per operating hour may be more appropriate. Other systems operate on an on-demand mode, where the risk of accident per operation might be might be more important.

Another consideration is whether you are trying to manage the risk from individual in-service systems one at a time, or whether you are trying to take management actions that affect a whole fleet, for instance through making design changes during procurement. Where the system or hazard is a one-off or the expected rate of accidents is relatively high, the risk might be considered per unit, but when there will be a large number of deployed systems, it can be easier to consider the expected number of accidents across a fleet or batch.

The considerations above are focussed on the risk posed *by* a system. An alternative view is to focus on the risk posed *to* the users of the system, or to other people and things that could be harmed. This approach is more valid when managing the people who use a system, rather than the system itself. It leads to measuring risk in terms such as an individual's probability of death per annum. Note that there is a difference between the probability of someone getting killed, and the probability of a particular person being killed. When taking this approach, one must carefully consider the population at risk, who is actually exposed, how often and for how long. Forecasting risk in this way can lead to some potentially illogical behaviours. It is not acceptable to try to make a system appear safer (per person) by exposing more people to it for shorter periods of time, or to claim a low average risk across a large population, when some individuals are exposed to much higher levels of risk. To get around these pitfalls, it can be useful to define 'hypothetical people' for the purpose of analysis, who have some fixed relationship to the system in question (e.g. pilot, passenger, maintainer, etc.), as described in [7]. Estimation of risk is carried out using the hypothetical person; real people then compare their own circumstances to those of the hypothetical person and decide its relevance.

Accidents are hopefully rare events. Whichever basis is used, the probability component of risk often involves very small numbers. This can make it convenient to use large sampling windows when setting the scope of risk measurements, e.g. per system life rather than per hour, per fleet rather than per unit or per population rather than per person. These approaches make estimation easier, but can invite confusion when system assumptions change.

---

[2] Derivation of estimates for severity and likelihood is beyond the scope of this paper.

A final factor to feed into the decision of how to set the basis for risk measurement is the standard of tolerability that will be used. The state of the art varies between different industries and technologies, affecting the level of risk reduction that can be achieved. The public perception of what can be tolerated also varies, and organisations have varying appetites for risk. This all means that what can be tolerated varies from case to case. In some industries there are particular targets or good practice limits of allowable risk. For instance, targets and limits for risk of radiation exposure are set in the HSE's *Safety Assessment Principles for Nuclear Facilities* in terms of maximum allowable dose per person per calendar year [8], and the European Aviation Safety Agency's AMC 25.1309 sets a target of better than one serious accident per 10 million flight hours for new large aircraft designs [9]. In the UK, when specific standards are not available, guidance from the HSE publication *R2P2* is often used to set limits of tolerability, generally in terms of the risk of death to an individual per year [7]. If the tolerability of risk will be measured against such a standard, it is useful to measure it in terms of the same units.

Decisions on how to measure risk therefore depend on a combination of factors: the physical characteristics of the system in question, the emphasis of the management effort, and the standards being used. Having decided on the units and scope, the absolute level of risk can be calculated simply by multiplying together estimates of likelihood and severity to give a single number. For clarity of interpretation, this number should always be quoted with its basis or units. It is also helpful to identify the level of uncertainty in the risk estimate, although discussion of uncertainty is beyond the scope of this paper.

Relative Risk

In practice, precise numerical representations of risk are not often used. This is partly because of the difficulty in obtaining numerical estimates of probability or severity, and partly because of the need to make comparisons between different risks. It is easy to determine the relative importance of two risks if they are both measured on the same basis. Often this is not the case, and in order to be able to rank or prioritise risks, we need to first convert them to a single basis. Difficulties arise when the risks posed by a system cause different types of harm.

Some types of harm can be measured on continuous scales, such as financial loss or radiation dosage, and it is possible to envisage some kind of mathematical equation to show equivalence between them. Others, in particular harm to humans, cannot easily be measured on a linear scale. A solution that is often used is to categorise the degrees of severity. Multiple definitions can be given for each category of severity, such that different types of harm with the same relative importance can be grouped together. Table 1 shows an example of such categories from GEIA-STD-0010, showing harm to people or equipment systems grouped with equivalent financial loss [6].

Description	Category	Criteria
Catastrophic	I	Could result in death, permanent loss of system function, permanent total disability, or loss exceeding $1M.
Critical	II	Could result in major system damage, permanent partial disability, injuries or occupational illness that may result in hospitalization of at least three personnel, or loss exceeding $200K but less than $1M.
Marginal	III	Could result in minor system damage, injury or occupational illness resulting in one or more lost work days, or loss exceeding $20K but less than $200K.
Negligible	IV	Could result in injury or illness not resulting in a lost work day, or loss exceeding $2K but less than $20K.

Table 1. Severity categories from GEIA-STD-0010 Table A-II .

In principle, making conversions between different bases of likelihood should be easier than converting severity, because likelihood can generally be measured on a continuous scale. This means that likelihood 'per fleet lifetime' can be converted to 'per equipment per hour' or 'per person per year' using simple scaling factors. However, these scaling factors rely on system-specific assumptions about things like fleet sizes and lifetimes, the way the equipment is used (duty cycle, operating tempo), and the number of people that interact with it. They can also vary between different risks in a single system, depending on who or what is involved. This means that when the assumptions change, it is not necessarily a simple matter to determine the effect on the risk.

Although likelihood can be measured on a continuous scale, it is often hard to make accurate estimates. As a work-around, it is common to measure probability on a semi-quantitative scale, making estimates to the nearest order of magnitude. When estimates of likelihood are even hazier, qualitative terms may be used, such as 'Frequent', 'Probable', 'Occasional', 'Remote', 'Improbable', or 'Incredible'. It is also common for these categories to be used as a simple way of converting between different bases of likelihood, similarly to the use of severity categories for different types of harm. The example in Table 2 shows such a conversion between risks based on a single unit and those based on a whole fleet.

Description	Level	Specific Individual Item	Fleet or Inventory
Very Likely	A	Likely to occur often in the life of an item, with a probability of occurrence greater than 10^{-1} in that life.	Continuously experienced.
Likely	B	Will occur several times in the life of an item, with a probability of occurrence less than 10^{-1} and greater than 10^{-2} in that life.	Will occur frequently.
Probable	C	Likely to occur some time in the life of an item, with a probability of occurrence less than 10^{-2} and greater than 10^{-3} in that life.	Will occur several times.
Unlikely	D	Unlikely but possible to occur in the life of an item, with a probability of occurrence less than 10^{-3} and greater than 10^{-6} in that life.	Unlikely, but can reasonably be expected to occur.
Improbable	E	So unlikely, it can be assumed occurrence may not be experienced, with a probability of occurrence less than 10^{-6} in that life.	Unlikely to occur, but possible.
Impossible	F	Incapable of occurrence. This category is used when potential hazards are identified and later eliminated.	

Table 2. Example likelihood categories from GEIA-STD-0010 table A-III.

When different types of harm are associated together by grouping them into categories, the definition of the categories records the organisation's value judgements about the relative importance of those effects. Similarly, the definition of qualitative or semi-quantitative categories of likelihood records assumptions about the way a system will be used.

If these categories of likelihood and severity are represented numerically (e.g. by a mid-point value), they can still be combined simply as before, by multiplication, to obtain a numerical measure of risk, which gives a common basis for ranking and prioritisation. However, if the categories are defined qualitatively, then some form of look-up mechanism is required. Tools such as a risk matrix are often used for this[3], as in the example in Table 3, which builds on the categories defined in Table 1 and Table 2.

		Likelihood Category					
		A	B	C	D	E	F
Severity Category	I	1	2	4	8	12	N/A
	II	3	5	6	10	15	
	III	7	9	11	14	17	
	IV	13	16	18	19	20	

Table 3. Risk matrix showing risk index values, after GEIA-STD-0010 table A-IV.

When these look-up techniques are used, a further variation on risk calculation can be brought into play. The simple multiplication method of calculating risk assumes a "rational" approach to risk: a reciprocal relationship between likelihood and severity, such that one death ten times a year would carry the same weight as ten deaths once a year. When using a look-up table, this relationship can be deliberately skewed. This could be used to weight unlikely high-impact events over frequent low-impact events, or visa versa, to reflect either society's dread of high-impact events or an organisation's appetite for risk.

In the example in Table 3, categories of likelihood and severity are used to determine risk on an arbitrary scale from 1 to 20, known as a risk index. In this index scale, '1' denotes the highest risk, and '20' the lowest. This is the opposite sense to numerically calculated risk, where higher numbers mean greater risk, but ties in with the concept of prioritisation, where the No. 1 priority is the most important. This matrix also exhibits a skewing bias towards severity. A likely loss of $500k every 50 years would could as category II-B or risk index 5, while $50k every 5 years would be III-A or risk index 7, despite the loss rate being the same per year.

The risk index allows risks from different bases to be compared, ranked and prioritised on a common basis, but does not give an absolute indication of the actual magnitude of the risk. Without the link to absolute risk measurements, it is hard to compare a risk against an external standard, or to determine whether it is tolerable. This can be worked around by converting standard limits or targets according to the common basis of the risk matrix, and using the results to divide up the cells. By doing this, a matrix like that in Table 4 could be derived from the example index in Table 3, using an appropriate standard.

[3] Other mechanisms include nomograms and decision trees.

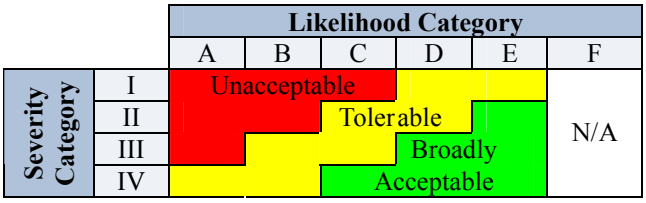

Table 4. Risk Classification Matrix

The new matrix divides up risk indices into classes or categories of risk. At the most basic, a risk class matrix can have just two classes, showing whether or not risks meet a particular standard of tolerability, but more complicated classes can be devised to show multiple thresholds. A common application is to have different classes to represent different levels of authority required for sign-off of risk acceptance. Different sets of classes may also be derived from the same risk index, to show different standards. The downside to this approach is that the use of the matrix obscures the detail of the original standard. When the matrix is used to classify and compare risks from different bases, it is not necessarily apparent whether a given standard should apply to a particular risk. E.g. a purely financial loss may end up being compared against a standard dealing with risk of harm to humans. A common failing of this type is to compare the individual risks due to particular hazards in a system against a standard for the overall risk from the whole system.

The categorization method gives a way of converting different types of risks onto a single basis and allowing their relative importance to be judged. However, the measure of risk that emerges is a rather abstract concept, and often cannot be linked back to a precise set of units. This can mean that the link between such categorised risks and standards of tolerability can become rather tenuous.

Changes in risk

As another viewpoint, we may want to look at changes in risk, to inform analysis of different options for safety improvements. If risk is being measured quantitatively or semi-quantitatively on an absolute scale, then the task is fairly trivial. One can compare a baseline estimate of risk against the estimated risk after implementation of the improvement, to gain an estimate of the benefit in risk reduction.

However, if one of the risk classification systems described in the previous section is being used to measure risk, it can be more complicated. Depending on the matrix being used, some risk reduction measures may not show as an improvement at all. For example, moving from cell A-I to B-II in Table 4 reduces the risk by two orders of magnitude, but leaves the risk class unchanged. Using a risk class matrix with more classes, or using a risk index (where each cell is effectively a class in itself), can improve the situation, but anomalies remain. In the example risk index in Table 3, moving a risk with a fixed severity from column D to E will always reduce the absolute level of risk by an order of magnitude. However, change in risk index is inconsistent. The index changes by 4, 5, 3 or 1 units, depending on whether the severity is I, II, III or IV. This means that a risk matrix alone is often inadequate for assessing potential changes in risk, and one should instead look at the source data for likelihood and severity.

Another potential problem with measuring changes in risk arises from the model that is used to calculate the risk in the first place, which also can mask the benefit of some safety improvements. We have defined risk as a combination of the likelihood of harm and the severity of that harm. A 'hazard' is "*a potential source of harm*" [2]; something that could lead to an accident. But what precisely is it that we are measuring the likelihood and severity of? Take the example of fire as a hazard. What harm does it cause? Depending on the way events play out, a fire might cause negligible harm, or multiple deaths and millions of pounds worth of damage, or something in between. Should we choose the worst possible scenario? Or the worst scenario with a credible likelihood? Or the most likely scenario?

The precautionary principle could lead us to combine the likelihood of a hazard with the severity of its worst credible outcome. This is Leveson's approach – she defines 'hazard level' as a combination of the likelihood of the hazard existing and the severity of its worst possible outcome; and risk as a combination of the hazard level, the exposure to the hazard, and the likelihood of environmental conditions allowing the hazard to develop into an accident [1]. This is slightly unsatisfactory, as it leaves doubt over which likelihood to use: that of the hazard developing into any accident, or of the worst-case scenario. By using the highest estimates of both probability and severity, the first case may significantly over-estimate the risk, but the second ignores the effect of higher-probability but less severe outcomes, giving an under-estimate. This method can still be used effectively for ranking risk despite its potential inaccuracy, but causes problems when analysing changes in risk. By only considering the worst case, it is possible for some safety improvements to fail to impact the risk assessment, if they do not affect the severity of the worst-case scenario, despite an obvious reduction in overall risk.

To get a more accurate measure of risk, the likelihood being assessed must match the outcome. Authors such as Ericson

suggest that hazards and accidents (mishaps) are intrinsically linked: that a hazard is a set of conditions and events that will lead inevitably to a particular outcome [10]. This means that there is a 1:1 mapping between hazards and accidents. This approach is perhaps more pure, but means that each potential outcome arising from a source of harm should be considered a separate hazard, quickly leading to a multiplication of the hazard management effort. Faced with this growth of effort and limited management resource, it is likely again that only the worst-case is assessed. Ekholm claims that if this is the case, typically as much as 50% of the risk may be omitted from the calculation [11].

The helpful concept in Leveson's model is that it breaks down likelihood into a part that is specific to the system and can be independently assessed (hazard level), and a part that depends on personnel exposure and the environment to determine the actual outcome, and can be varied according to circumstances to determine the current risk. Similarly, Ekholm separately calculates the likelihood of the hazard and the exposure of affected people. He then extends the model by making the severity a 'damage distribution' to give an estimate of the expected loss taking into account the many possible outcomes of a hazard.

The damage distribution is calculated by modelling how often the hazard would lead to fatal, serious, less serious or negligible outcomes. The severities of these outcomes are summed, weighted by their relative likelihoods. The weighting factors in the impact of environmental effects that may alter which outcome arises from a hazard. In order to perform the summation, different forms of injury have to be measured on a common basis, for which Ekholm uses equivalent fatalities, holding one fatality equal to ten serious injuries or one hundred less serious injuries. This method allows a composite overall risk to be generated for each hazard; these can be summed to give a total system risk as a single value [11].

This approach seems to give an accurate picture of the overall risk, and allows credit to be taken for risk reductions that do not affect the worst-case scenario. A similar effect could be achieved by using event tree analysis to calculate the range of potential outcomes from a hazard, with their associated likelihoods.

Potential risk

Most risk estimation is based on estimates about the current or predicted state of affairs; either the level of risk currently being experienced, or that which will be experienced after a system has been deployed or modified. Such assessments take into account controls designed to mitigate the risk. A different viewpoint is to look at the potential risk that would be posed if these controls and mitigations were not effective. This viewpoint is useful to allow determination of the significance of different risks, to determine the appropriate assurance regime.

Assurance should be proportionate to the risk involved, with a light touch applied for low-level risks, increasingly stringent measures applied to gain confidence in safety as the potential risk rises, and active regulation and permissioning regimes applied to the highest risks. As deployed systems are supposed to have had all their hazards controlled such that the risk is ALARP and can be tolerated, this could be incorrectly taken to imply that a uniformly low level of assurance could be applied. Instead, assurance should be proportionate to the potential risk, to gain adequate confidence that the controls have been effectively designed and are maintained through the life of the system, and that the post-mitigation assessment of the risk is accurate.

The potential risk is not quite the same as the baseline level of risk that is drawn up in preliminary hazard assessment. Credit can be taken for some mitigation, as hazards that have been totally removed from a system do not need to be included as part of the potential risk. However, the potential risk must reflect hazards that are still present in the system, but are somehow controlled to a tolerable level.

Organisational management of safety risk

This paper has shown some of the variety of how risk estimates can be presented to satisfy different points of view, depending on whether one is trying to compare a source of risk against a standard of tolerability, rank or prioritise the sources of risk within a project[4], assess the benefits of a safety improvement, or determine the level of assurance appropriate for a project. The methods presented have generally been tailored to be project-specific, i.e. to suit particular management arrangements for a given system. This tailoring takes the form of choosing standards, units and conversions appropriate to the underlying mechanisms that cause the risk in question.

It is also sometimes necessary to manage risk at a higher organisational level, for instance to make decisions about which systems to deploy (or not deploy), where to invest resources so that they will have the greatest impact, and which

[4] In the context of this paper, a project is a defined undertaking that has its own management arrangements.

issues deserve most management oversight and assurance. At this organisational level, the project-specific methods of presenting risk estimates are less helpful.

Absolute measures of risk are not helpful at the organisational level. They are not comparable with each other unless they are all measured on the same basis, and if they are all measured on the same basis, it becomes difficult to determine whether they meet the relevant standards of tolerability, which can vary from area to area. Similarly, risk classifications can be helpful to give a rough idea of the magnitude of the risk, but are only comparable if made using the same basis. Unfortunately, the bases for risk ranking and classification are often implicitly project-specific, so "Class A" for one project can be different to "Class A" for another, even when nominally using the same risk matrix. These risk classes should therefore be considered as project-specific measures of significance and priority rather than absolute measures that can be compared together.

As an example, consider the risk posed by a utility truck and a one-off specialist truck to their drivers. When specialist standards do not exist, it is common to consider the risk of a fatality per exposed person per year, to allow assessment of tolerability. Measured on this basis the magnitude of the risk might be identical, and for arguments sake might be allocated Risk Class C. But looking from an organisational viewpoint, if there are 1000 utility trucks and only one specialist truck, the risk arising from the design of the utility truck is clearly higher. There is then a question of how to aggregate the risk. Should it be quoted as "two C-class risks" (which is what, in design terms, is being managed)? Or should the aggregate reflect 1001 instances of a C-class risk?

Differences in bases used for risk estimation appear inevitable in an organisation that manages diverse types of system. The key problem appears to arise when estimates of risk are supplied without detailed information about the basis that they are founded upon. This makes it impossible to reach a sensible conclusion about how the risks should be compared or aggregated. Although such information could be supplied, this would place a huge burden on project teams to report the necessary details, and on management teams to make sense of them.

To try to improve the situation, it is proposed that the following risk information is reported by project teams to their corporate organisations:

- Notional expected loss – defined in terms of expected loss per year across all instances of the risk, converted to a monetary value. This would show the overall significance of each risk in terms of its impact on the organisation, and could be summed by system, business unit or other grouping. It is "notional" because it represents a combination of many types of loss, not a real financial loss.

- Compliance issues – a flag or traffic light value that would show where there is a particular concern that a risk does not meet the appropriate standard of tolerability for the industry or technology in question. It could be set either using a people-focussed measure such as a tolerable fatal accident rate per individual, or a system-focussed measure appropriate to the context.

For projects producing systems or modifications that are not yet deployed, it may be appropriate to also report whether or not the design has matured to a stage that the predicted risk is expected to be ALARP upon deployment. Deployed systems would generally be expected to have their risks controlled to ALARP at all times, but if this was not the case, the deficiency could also be reported via the compliance flag.

At the level of management statistics it may be sufficient for the compliance flag purely to highlight that there is a problem, since the detail of the specific issue can be requested if necessary. A refinement might be to discriminate between compliance with legal requirements and with best practice or self-imposed limits.

Notional expected loss is defined as the equivalent financial loss that could be expected from all instances of a risk, per year. i.e.

N.E.L = exposure x likelihood x severity

Exposure is the number of things (people, artefacts, etc.) exposed to the risk, multiplied by the proportion of the year that they are exposed. It is likely to be proportional to the number of deployed instances of the system that poses the risk. Likelihood is the probability that the risk matures (causes an accident) in one of those systems in a one-year period. Severity is the estimated loss (in financial terms) per risk event.

Where a single hazard has multiple outcomes, either the notional expected loss can be calculated separately for each and then summed, or an average expected severity could be calculated using a damage distribution. Alternatively, a set of likelihood/severity pairs derived from an event tree could be summed together, then multiplied by the exposure value. Damage to property is normally easy to quantify financially. Harm to humans can be converted to financial terms actuarially, or using a Value for Preventing a Statistical Fatality (VPF). Environmental impacts might be quantified using clean-up costs, or arbitrary valuations such as have been proposed for control of carbon emissions. Many other

techniques could also be conceived, and would be compatible so long as they resulted in an expected financial loss per risk per year.

Notional expected loss is a management-focussed measure: it does not relate to a risk that a particular individual might be exposed to, or the risk from a particular system. Instead, it measures the risk that arises from each hazard being managed. Financial value is used, as money is the natural mechanism for assigning relative values to heterogeneous quantities. A year is used as the time-base, for ease of comparison with management and planning cycles. Notional expected loss could be calculated for a single year, or plotted over time, to show how an organisation's risk profile varied as different systems were brought in and out of service.

Project teams would be free to use whatever methods they liked to generate their estimates of risk and conduct their risk management, so long as they reported notional expected loss using the common basis. They would be forced to derive their own project-specific conversion factors to arrive at the correct measures of exposure, likelihood and severity (although guidance from the parent organisation would help maintain consistency). These conversions would be explicit (rather than implicit in the set-up of management tools), and so can easily be changed as assumptions about the usage or deployment of a system change. The conversion factors are also managed and used in the same part of the organisation where they are generated, so there would be less likelihood of their misuse.

Conclusions

The differences between methods for presenting estimates of risk are often subtle, and when considering a single project may not be significant. There are numerous ways of assessing risk and although some may be more useful than others under particular circumstances, most are equally valid. Problems occur when we need to take a high-level view and aggregate or draw comparisons between risks that have been estimated on different bases. Although this can happen when prioritising particular sources of risk in a single project, it is more problematic at a higher level of management when making decisions that span multiple projects, since the information about those original bases may not be available.

Notional expected loss has been proposed as an appropriate management-focussed metric for showing the relative significance of different risks at an organisational level.

References

[1] Nancy G. Leveson. *Safeware: system safety and computers*. Addison-Wesley, 1995.

[2] International Electrotechnical Commission. Functional safety of electrical/electronic/programmable electronic safety-related systems - part 4: Definitions and abbreviations. British Standard BS EN 61508-4:2002, British Standards Institution, March 2002.

[3] Directorate of Standardization. Safety management requirements for defence systems - part 1: Requirements. Defence Standard 00-56, Ministry of Defence, Glasgow, UK, June 2007. Issue 4.

[4] Health and safety at work etc. act. Elizabeth II 1974 chapter 37, 1974.

[5] Tucker L.J., Asquith L.J., and Singleton L.J. Edwards v. National Coal Board, 1949. All ER 743 (CA).

[6] ITAA G48 Committee. GEIA-STD-0010: Standard best practices for system safety program development and execution. TechAmerica standard, American National Standards Institute, February 2009.

[7] *Reducing Risks, Protecting People (R2P2)*. HSE Books, 2001.

[8] HSE Nuclear Directorate. Safety assessment principles for nuclear facilities. Technical report, UK Health and Safety Executive, Bootle, January 2008. 2006 revision 1.

[9] Certification specifications for large aeroplanes. Certification Specification CS-25, European Aviation Safety Agency, Cologne, December 2009. Amendment 8.

[10] Clifton A. Ericson. Two hazards, two mishaps? *Journal of System Safety*, 46(1):31–32, January-February 2010.

[11] Ragnar Ekholm and Arne Börtemark. System safety in the Swedish defence. In *Proceedings of the 25th International System Safety Conference*, pages 181–190, August 2007.

Biography

J. R. Inge, Head of the Ship Safety Management Office, Ministry of Defence, Sea Systems Group, Elm 1c #4134, MOD Abbey Wood, Bristol BS34 8JH, UK, telephone – +44 (0) 30 679 35143, facsimile – +44 (0) 117 913 5943, email – DESSESea-SSMO@mod.uk.

Mr Inge was lead editor for Issue 4 of Defence Standard 00-56 – Safety Management Requirements for Defence Systems. His experience is as a project manager and safety policy specialist in the UK Ministry of Defence. He is a Chartered Engineer and holds a MEng from Durham University and a postgraduate diploma in Safety Critical Systems Engineering from the University of York. He contributes to the IET/BCS Independent Safety Assurance Working Group and the IET's Technical Advisory Panel on Functional Safety.