Back to Basics: Risk Matrices and ALARP

Rhys David

Safety Assurance Services Ltd. Farnham, UK

rhys@safetyassuranceservices.co.uk

Glen Wilkinson

GKW Consulting Ltd, Godalming UK

glen@gkw-consulting.com

Abstract Risk matrices are tools for comparing risks relative to one another (e.g. for a single system or single activity) and hence being able to 'rank' them for the purposes of risk mitigation and the allocation of safety resources. Risk matrices are not tools for determining the tolerability, or otherwise, of 'single risks'[1].

However, in some sectors risk matrices are being used to judge whether 'single risks' are tolerable, and ALARP arguments are being made for each separate 'single risk' rather than an individual's aggregate risk from all causes. This stretches the Risk Matrix concept beyond its breaking point, and is leading to potentially misinformed decisions by senior managers regarding the overall level of risk present, and hence whether risk reduction options are either needed or are reasonably practicable.

1 Risk Management Basics

1.1 Background

The Management of Health and Safety at Work Regulations 1999 (MHSWR, 1999) [1] requires all employers to assess the risks to employees and any others who may be affected by their undertaking, to enable them to identify measures necessary to comply with their duties under Health and Safety Law.

There are various models of the activities involved in Risk Management and the terms "Analysis", "Assessment" and "Control" are used in a variety of ways. Regardless of the model and the terminology, Risk Management is an iterative process, where the results of activities are considered in the repetition and refinement of previous activities.

One generic model of Risk Management for an organisation is presented in Figure 1 below. This shows how the organisation's Risk Policy is shaped by external factors such as Regulations and public expectations, as well as experience from reviewing its own Risk Management processes.

---

[1] We have chosen the term "Single Risks" in preference to "Individual Risks" to avoid confusion with the level of risk of death an individual is exposed to as the result of an activity or operation. "Single Hazards" is used by some people, but is not considered appropriate, because several Hazards may be involved in the Accident Sequence leading to the outcome of interest.

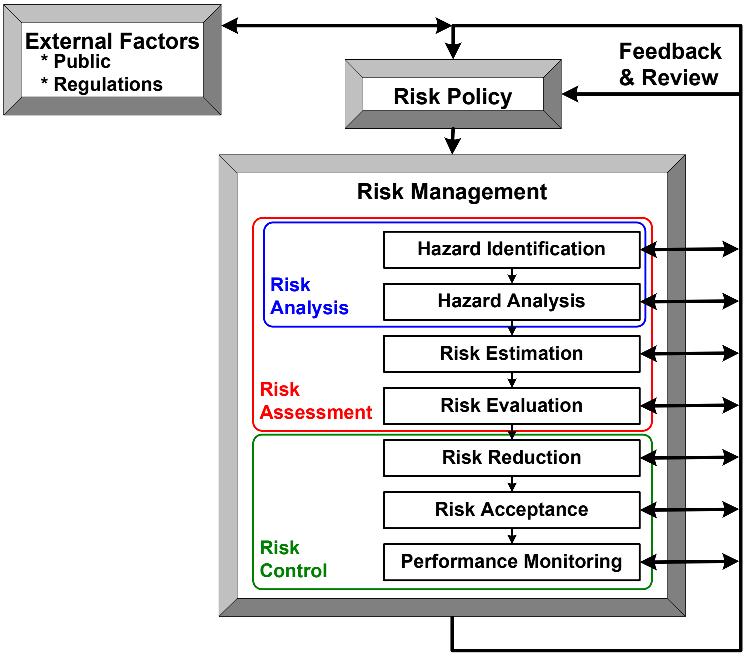

Fig. 1. Example Model of Risk Management Activities for an Organisation

1.2 Risk Assessment

The aim of the risk assessment is to identify the significant risks to health and safety to any person arising out of, or in connection with any work activity. It should identify how the risks arise, and how they impact on those affected. The information is needed so that decisions can be made about how to manage the risks in an informed, rational and structured manner and so that the action taken is proportionate.

Risk assessments are required by MHSWR (1999) [1] to be 'suitable and sufficient'. This means they should:

- identify the significant risks arising out of the work activity;
- consider all those who may be affected;
- be appropriate to the nature of the work; and
- be such that they remain valid for a reasonable period of time.

The level of detail in a risk assessment should be proportionate to the type and significance of the intrinsic hazards. Once the risks have been assessed and taken into account, insignificant risks can usually be ignored as can risks arising from routine activities associated with everyday life, unless the activity compounds or significantly alters those risks.

Employers are expected to take reasonable steps [2] to help themselves identify risks, e.g. by looking at appropriate sources of information, such as relevant legislation, appropriate guidance, supplier manuals and manufacturers' instructions and reading trade press. They should also look at and use relevant examples of good practice from within their industry. The risk assessment should include only what an employer could reasonably be expected to know; they would not be expected to anticipate risks that were not foreseeable.

Risk assessment can be a qualitative, semi-quantitative or quantitative process. Any assessment should begin with a simple qualitative assessment, including consideration of whether any relevant good practice is applicable. In some cases it will be appropriate to supplement the qualitative assessment by a more rigorous semi-quantitative or quantitative assessment, depending upon the level of risk identified.

Regardless of whether the approach to risk assessment is qualitative, semi-quantitative or quantitative, a logical and systematic process must be adopted. The depth of the analysis should be proportionate to the nature and magnitude of the hazards involved and the complexity of the systems being considered.

Descriptions and some examples of these different approaches to risk assessment are given below:

- **Qualitative risk assessment** is based on a representative selection of specific examples for comparison with standards and relevant good practice. Techniques can include:
  - i. Simple methods requiring a basic level of risk-based judgement and suitable for relatively minor hazards, such as described in "Five Steps to Risk Assessment" [2];
  - ii. Hazard Identification techniques which have a qualitative evaluation of significance of the hazards, such as FMECA;
  - iii. Risk Matrices where the categories of consequence and likelihood are defined only descriptively (e.g. "Significant" consequence and "Reasonably probable" likelihood).
- **Semi-Quantitative**: This can include techniques for supplementing qualitative techniques (see above) with, for example, measurements to identify the presence of hazards from chemicals or machinery, or the use of simple modelling techniques. These modelling techniques may be used to derive order of magnitude estimates of the severity and likelihood of the identified possible accidents. These estimates can be combined to obtain estimates of the order of magnitude of the accident risk. Ways of carrying out semi-quantitative risk assessment include:
  - i. Risk Matrices where the descriptive definitions of consequence and likelihood are given numerical interpretations;
  - ii. Layers of Protection assessment;
  - iii. Lines of Defence assessment.
- **Quantitative**: Quantified Risk Assessment (QRA) involves obtaining a numerical estimate of the risk through quantitative consideration of event probabilities and consequence. The results of QRA will be numerical estimates of risk, which can be compared against numerical risk criteria at the risk evaluation stage. QRA can draw on a range of techniques whose purpose is to provide numerical estimates of the probability and/or the severity of possible accidents. Such techniques can include:
  - i. Analysis of historical accident frequencies or system Reliability;
  - ii. Fault Tree Analysis;
  - iii. Event Tree Analysis;
  - iv. Simulation modelling;
  - v. Human Reliability Analysis;
  - vi. Bow-tie Analysis;
  - vii. Cause-consequence diagrams

2 Risk Matrices

2.1 Background

Risk Matrices are tools for use in the Risk Evaluation phase (Figure 2) of the Risk Management process. They are used once Hazard Identification, analysis and Risk Estimation has been undertaken and their primary purpose is to:

- Determine how significant each Risk is (e.g. High, Medium, Low)
- Prioritise/Rank Risks relative to one another for the purpose of Risk Control
- Highlight areas for further more detailed risk assessment (e.g. fully quantitative rather than qualitative for higher level risks)

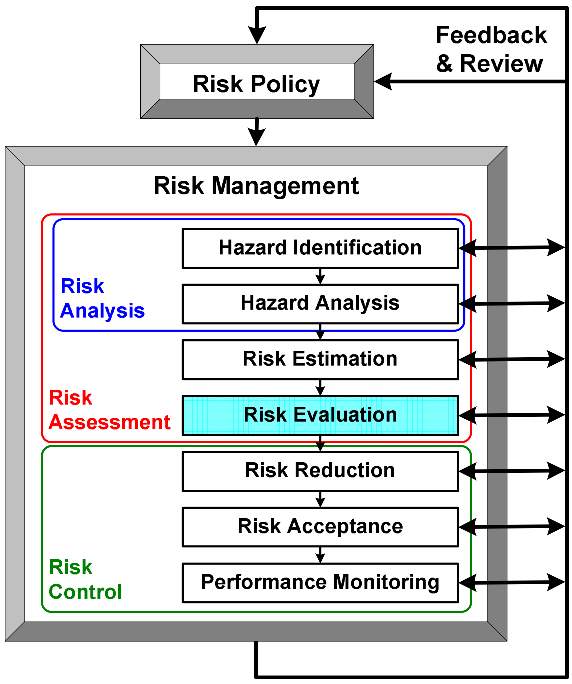

Fig. 2 Risk Evaluation within a Risk Management Process

2.2 Risk Continuum

Risk is usually considered to have two components, namely event likelihood and its severity of consequence (i.e. the scale of the loss). Risk can be visualised on a two dimensional space constructed with axes of likelihood (how often) and severity of consequence (how bad) as shown in Figure 3.

In a Cartesian plot of the Risk space continuum, severity is conventionally plotted increasing left to right on the x-axis and likelihood increasing vertically on the y-axis. Risk is therefore lowest at the bottom left of the plot and increases towards the top right.[2] Where the likelihood and severity scales are quantified, it is very common practice for logarithmic scales to be used.

A Risk Matrix may be overlaid on this continuum, with cells defined in terms of severity columns and likelihood rows. Useful information about Risk significance is available by knowing the cell in which a particular Risk lies, even without knowledge about its precise location.

---

[2] Confusingly, several Standards use Risk Matrices with the likelihood on the x-axis and, worse, with one or both axes reversed.

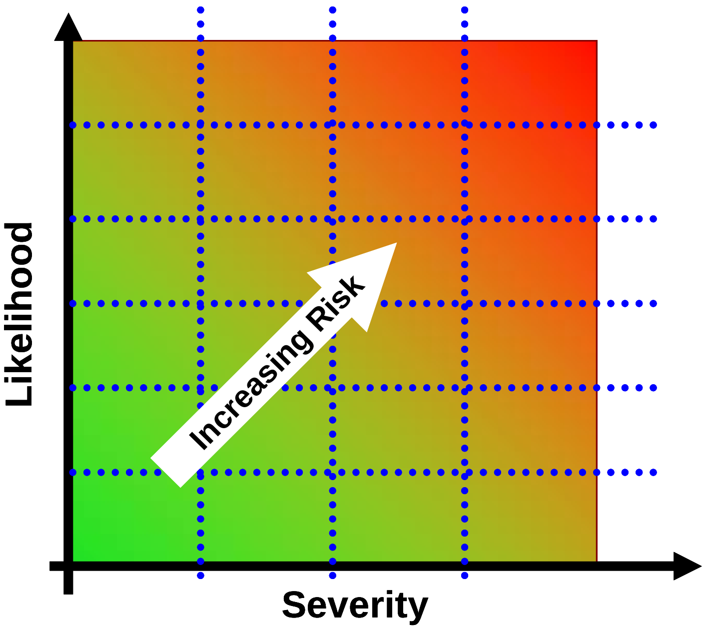

Fig. 3. Division of Risk Continuum

2.3 Risk Evaluation

Risk Evaluation follows after Risk Estimation in which the event likelihood and severity of consequence are each estimated. This may be through analysis of historical data, through use of generic data, by applying models or based on the knowledge and experience of Subject Matter Experts (SME), often working in a committee.

With estimates of likelihood and severity for each 'single risk', its location may be plotted on the Risk Matrix (Figure 4). A Risk Ranking, Risk Score/Index or Risk Class is associated with each cell of the Risk Matrix and this provides a relative indication of Risk significance.

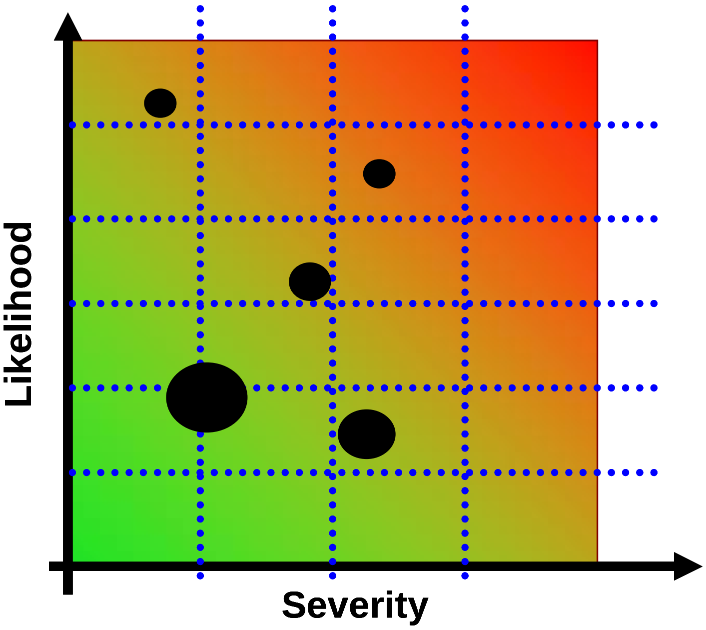

Fig. 4 Plotting Single Risks on the Risk Matrix.

The estimates of likelihood and severity are each subject to uncertainty and this may be indicated by plotting an area rather than a specific point or through use of error bars.

There are many different Risk Matrices in existence and countless standards devoted to the application of these matrices. However, it is noted that whilst the guidance on their use is prevalent, the guidance on how to develop them is actually quite sparse.

Risk Matrices can range from simple 2 x 2 matrices (Figure 5) only used to rank 'single risks' into High, Medium or Low to hugely complex matrices of 14 x 14 (Figure 6), which are attempting to plot a range of different Severities (impacts/ losses) on the same Risk Continuum.

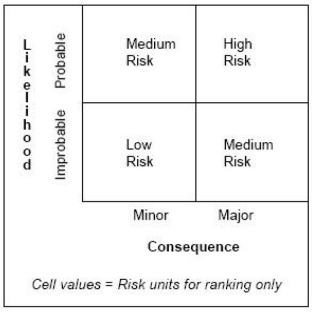

Fig 5. Simple 2 x 2 Matrix

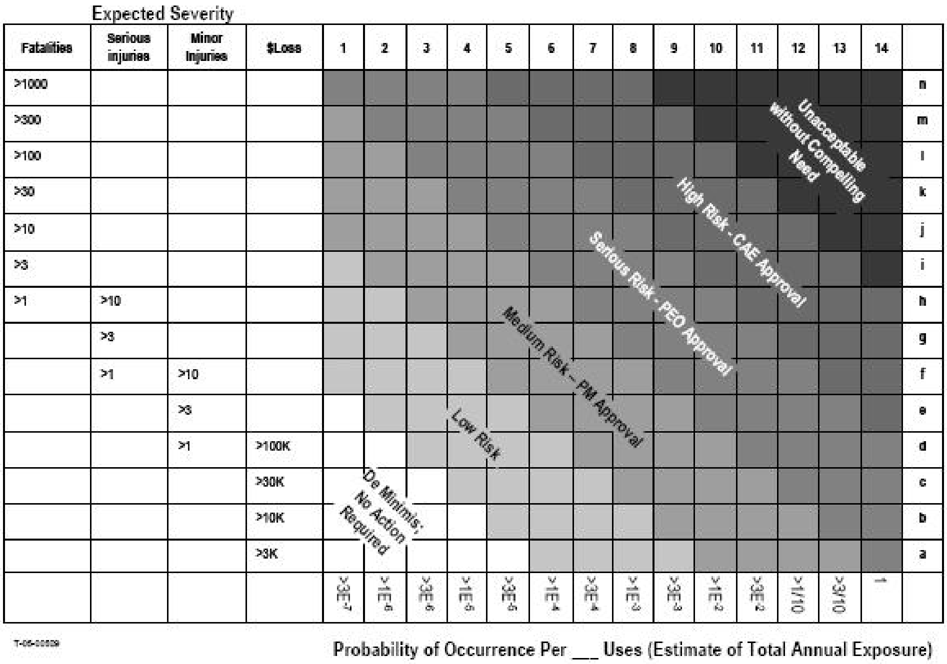

Fig. 6 Complex 14 x 14 Matrix

2.3 Single Risks

There is no 'correct' definition of what constitutes a 'single risk'. Different analysts may each define different 'safety issues' as 'single risks' (e.g. Aircraft Loss, Controlled Flight Into Terrain (CFIT) and CFIT due to Human Error). The authors have seen Risk Assessments of similar systems of medium complexity, one of which had some twenty 'single risks' and the other well over 1,000. In each case the most significant safety issues had been identified and were being managed, but the Risk Matrix gave quite different impressions about the "safety" and Tolerability of the two systems.

At the level of the single system or activity, this is acceptable, providing that safety issues are being recognised and managed. However, for a senior manager with responsibility for safety, this lack of consistency makes it impossible to have a consistent comparative view of risks across multiple projects or systems within an organisation.

A Risk Ranking, Risk Score/Index or Risk Class should be considered to provide a relative rather than absolute measure of Risk significance. Using the same Risk Framework (e.g. a common matrix with quantitatively defined bands) may permit direct comparison of Risks on different systems of the same generic type, but not across dissimilar types.

Where senior managers need to compare exposure to possible loss across multiple systems/facilities/operations, they require metrics which can be directly compared. This would give managers improved appreciation of the context or implications of 'single risks' and might be presented in terms such as:

- Exposure to Loss (calculated in terms of predicted equivalent fatalities per person-year exposed);
- Exposure to Loss (calculated in terms of number of predicted events in each severity category, per person-year exposed);
- Exposure to Loss (calculated in terms of predicted equivalent fatalities per system year or per fleet/inventory year);
- Exposure to Loss (calculated in terms of number of predicted events in each severity category, per system year or per fleet/inventory year).

2.4 Stretching the Matrix

2.4.1 All Things to All Men

As we have already seen the Risk Matrix is a relatively easy tool to use and its visual reference allows experts and non-experts alike to plot and consider impacts Vs likelihood with some degree of understanding. However, it is this apparent ease of use and power that has led to the Risk Matrix being seen as the holy grail of Risk Management. It is sometimes considered to be "all things to all men" and it is now very common to see a single Risk Matrix being used to measure a whole range of unrelated impacts on a single continuum (Figure 7) such as:

1. Fatalities (single and multiple) of both directly involved personnel and 3rd parties
2. Injuries (major and minor) – directly involved and 3rd parties
3. Environmental Impact
4. Financial Loss
5. Reputation

CONSEQUENCE					INCREASING PROBABILITY				
Severity Rating	People	Assets	Environ-ment	Reputation	A Rarely occurred in E&P industry	B Happened several times per year in industry	C Has occurred in operating company	D Happened several times per year in operating company	E Happened several times per year in location
0	Zero injury	Zero damage	Zero effect	Zero impact	Manage for continued improvement				
1	Slight injury	Slight damage	Slight effect	Slight impact					
2	Minor injury	Minor damage	Minor effect	Limited impact					
3	Major injury	Local damage	Local effect	Considerable impact					
4	Single fatality	Major damage	Major effect	Major national impact	Incorporate risk reducing measures				
5	Multiple fatalities	Extensive damage	Massive effect	Major international impact				Intolerable	

Fig. 7 Example of Matrix with Impacts of Different Types

Risk Matrices which cover impacts of different types have an implicit (sometimes explicit) assumption that losses of different types can be equivalent in their significance and hence value. A single metric, often financial, is applicable to different types of loss so that they can be considered on a single Risk Matrix.

2.4.2 Risk Matrices to Assess Acceptability

It is becoming common to see Risk Matrices used not only for the purposes of determining significance, risk ranking and areas for more detailed assessment but also for the determination of Tolerability or Acceptability.

One of the first standards to make this connection between the Risk Matrix and Acceptability was Mil-Std-882 [2] which defined the use of 'Hazard Risk Indexes' (HRI), as demonstrated in Figure 8, and gave examples of how these HRI could be used to determine Acceptability[3].

It is clear from the wording and categories used in the standard that the matrix is being used in a system-centric manner (i.e. risks from the system) rather than a person-centric manner (risks to the individual.)

HAZARD CATEGORY FREQUENCY	CATASTROPHIC	CRITICAL	MARGINAL	NEGLIGIBLE
FREQUENT	1	3	7	13
PROBABLE	2	5	9	16
OCCASIONAL	4	6	11	18
REMOTE	8	10	14	19
IMPROBABLE	12	15	17	20

Fig.8 Mil-Std-882 Example Matrix 2

Hazard Risk Index Suggested Criteria

- 1 - 5 Unacceptable
- 6 - 9 Undesirable (Management Authority decision required)
- 10 - 17 Acceptable with review by Management Authority
- 18 - 20 Acceptable without review

In this example the HRI gives the ranking of 'single risks' which are also sorted into one of four "Acceptability" categories..

The major issue that this approach highlights is Acceptability is now being judged for each of the 'single risks' in turn rather than for the aggregated risk from the system or to an individual. Because the 'single risks' are a subjective human view of part of the overall risk, there will never be consistency in how they defined.

Whilst ranking 'single risks' for the purposes of risk reduction is worthwhile, there is no clear way of knowing whether the aggregate risk from the system under assessment is Acceptable.

2.4.3 Absolute Risk

In the UK the Health & Safety Executive (HSE) states figures for the tolerability of individual risk of fatality from all risks (Individual Risk) per annum [4]. The HSE also states criteria for Societal Risk (e.g. an event leading to 50 or more fatalities)., These criteria are not directly comparable to many Risk Matrices, specifically those which are system-centric and use probabilities of various units, appropriate to the system but not directly relevant to person-centric assessment and criteria.

There have been many attempts to combine the two approaches of system-centric assessment and person-centric criteria. The preferred method has been to take the implied levels of Acceptability defined on many Risk Matrices and overlay the HSE Individual Risk boundaries onto the matrix. Hence, it is seen that many standards imply that the HSE TOR boundary limits equate directly to the 'single risk' Risk Class boundaries of the Risk Matrix.

Attempts are made to map person-centric risk requirements on to the system, so that Tolerability of Risk (TOR) threshold values may be used to derive appropriate 'surrogate' system requirements. This has to take account of many factors which are difficult to understand or estimate and which may be beyond the control of a project. These include:

- Risk exposure time for an individual (e.g. proportion of the working year);
- Number of people exposed to the system risk (e.g. including visitors, transients, general public);

[3] Acceptability here is determined by a Senior Manager/ Officer or panel of safety experts, rather than the use of some industry standard criterion, and does not mean the same thing as Tolerability where a residual level of risk is implied.

- Simultaneous sources of risk beyond the system of interest (e.g. other equipment being used at the same time, other co-located systems, other non-equipment based activities).

The system may affect different people or groups of people in different ways and with different exposure times. Therefore it is not possible to calculate a "system risk budget" by summing the "risk budgets" of all those exposed to its risks.

Taking account of the factors above is a process of "Risk Apportionment" from an Individual Risk (IR) budget but often with no consistency or overall control. The result is frequently inappropriate apportionment and even allocating of a whole year's "Risk budget" to a single system that is used only infrequently. It is noted that there is often no single body responsible for "Risk Apportionment" between different Activities or Projects and no single body responsible for the management of an IR budget.

These IR criteria relate to the total exposure to Risks from all work-related sources and not to 'single risks' one at a time. In order to use these criteria it is necessary to establish the "most-at-risk hypothetical person". A "Person-centric" risk assessment would then be conducted, considering all the sources of risk for this "most-at-risk hypothetical person", including the activities they do, the places they work and/or the "Major Accident Hazards" to which they are exposed.

It is noted that where Safety Requirements are stated quantitatively, a Quantitative Risk Estimation, including Quantitative Aggregation where relevant, will be required. For example if a Basic Safety Level (BSL) Individual Risk of no more than 1 fatality in 1,000 man-years is a stated requirement, quantitative presentation of the total Risk exposure of the "most-at-risk hypothetical person" will be necessary. An exception to this might be where semi-quantitative estimates indicate that the expected IR is significantly less than the BSL.

Furthermore, there is a misconception in many industries that the probability of a Single Fatality (regardless of the number of people exposed) is the equivalent of an Individual's Risk of fatality.

The implication of this is that risk is measured differently from project to project, not just across different sectors or domains but even within the same organisation. Therefore, it is impossible to accurately prioritise risks across a number of projects as they are all measured differently, although this may not be obvious to the casual observer, see Figure 9.

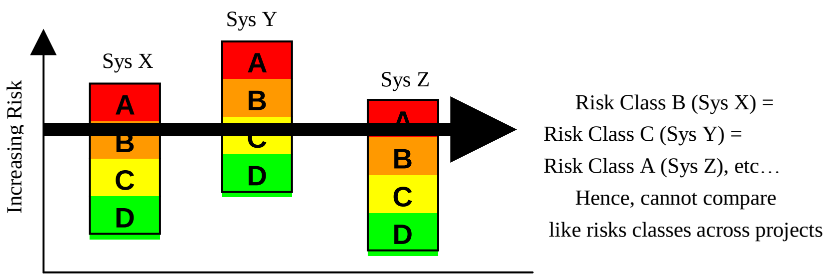

Fig. 9. Comparing Similar Risks

Figure 9 above illustrates the problem: just because the different example systems use the same ranking system (i.e. Risk Class A, B, C or D) they do not actually mean the same thing.

In the above example, it may be considered that all Risk Class As are by definition a higher safety priority than Bs or Cs. However, it can be seen that due to areas of overlap there could be Risk Class Cs in Sys Y that are actually a higher priority (in terms of risk reduction and hence time, effort and cost) than some Risk Class As in Sys Z.

As a tool purely for the purpose of prioritising risks relative to one another the Risk Matrix approach is effective. However, this is not how the matrices are generally employed by organisations, and hence the decisions being made by duty holders are based upon poor data.

Using Risk Matrices alone it is almost impossible to actually prioritise the risks across all the projects or systems managed by a single organisation. This cross-project or system appreciation of risk would require additional measures of risk exposure to be calculated, as suggested in Section 2.3 above.

As there is poor or no traceability to the HSE TOR levels or any equivalent, then there is no justification for any ALARP-style arguments that are applied to single risks, (i.e. there is no proof that the aggregate risk actually sits in the tolerable region).

The implication (or certainly the way it is interpreted in some Standards) is that these Risk Class categories align with the HSE TOR boundaries and hence Risk Classes B and C (Medium) equates to the ALARP region. This is actually not true: ALARP can only be applied where a risk falls into the tolerable region for an Individual's aggregate risk from all causes.

This is shown in the abstract from Def Stan 00-56 [5] in Figure 10 below:

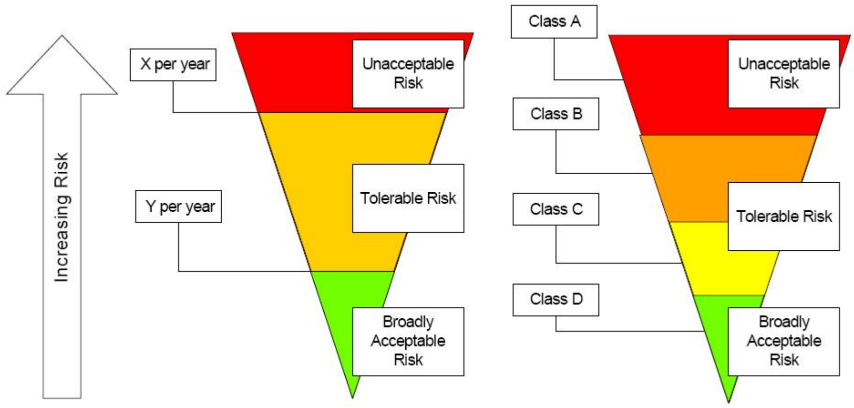

Fig. 10 Risk Classes versus ALARP regions (from Def Stan 00-56)

There are several problems with this approach; the Tolerability boundaries, shown on the ALARP triangle on left, use target levels of risk from the HSE that have the units of measurement as "the risk of fatality to an individual from all risks per annum", whereas the Risk Class boundaries on the right of the figure use whatever units the project has deemed are most suitable and only apply on a 'single risk' basis.

Clearly the two triangles are not measuring the same thing and hence cannot be compared as implied by the figure. Effectively the left-hand triangle sets the boundaries for the 'hypothetical person's' tolerable levels of risk from all risks, which are from all systems and activities that they are exposed to in a typical year. The right-hand triangle is actually only indicating the level of significance of a 'single risk', and from a single system or activity, which could potentially affect numerous different individuals.

If every possible 'single risk' for a system or an activity falls below the TOR threshold, then the aggregate risk for a system is often judged to be tolerable, this is clearly not true in all cases.

Sometimes multiple thresholds may be defined (e.g. intolerable, tolerable, and broadly acceptable) with the location of each single risk relative to the thresholds determining the management level that is authorised to give approval.

The total risk presented by the system of interest is a parameter that should be understood by risk managers and risk acceptance authorities, but it is seldom calculated and presented explicitly. Instead, there may be an implicit assumption that if all of the single risks are tolerable, then the total aggregate risk must be tolerable.

3 Arguing ALARP

3.1 Background

The Health and Safety at Work etc Act (1974) [6] includes the general duty on all employers to reduce risks So Far As Is Reasonably Practicable (SFAIRP). HSE considers that this will be achieved if risks are reduced 'As Low As is Reasonably Practicable' (ALARP).

The concept of "reasonable practicability" recognises that absolute safety cannot be guaranteed in some circumstances, and permits the "duty holder" a defence in law for choosing not to adopt certain risk reduction measures. For example, the risk reduction measures were not technically possible (i.e. not "practicable") or were inappropriate and out of proportion to the safety benefits (i.e. not "reasonably practicable").

The key case from the courts regarding reducing risks as low as reasonably practicable is **Edwards v. The National Coal Board** (1949) [7] in which the Court of Appeal held that:

> 'Reasonably practicable' is a narrower term than 'physically possible' and seems to me to imply that a computation must be made by the owner in which the quantum of risk is placed on one scale and the sacrifice involved in the measures necessary for averting the risk (whether in money, time or trouble) is placed in the other, and that, if it be shown that there is a gross disproportion between them – the risk being insignificant in relation to the sacrifice – the defendants discharge the onus on them'.

The Court further stated:

> '…in every case, it is the risk that has to be weighed against the measures necessary to eliminate the risk. The greater the risk, no doubt, the less will be the weight given to the factor of cost'.

Thus, determining that risks have been reduced ALARP involves an assessment of the risk to be avoided, of the sacrifice (in money, time and trouble) involved in taking measures to avoid that risk, and a comparison of the two. This process can involve varying degrees of rigour, which will depend on the nature of the hazard, the extent of the risk and the control measures to be adopted. The greater the initial level of risk under consideration, the greater the degree of rigour required of the arguments purporting to show that those risks have been reduced ALARP.

3.2 Risk Tolerability

HSE describe in R2P2 [4] the TOR concept (see Figure 11), which sets out the framework for reaching decisions on whether risks from an activity or process are unacceptable, tolerable or broadly acceptable. In this context, 'tolerable' does not mean 'acceptable'. 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, i.e. that everyone would agree without reservation to take the risk or have it imposed on them.

Most decisions on whether risks are ALARP are made by exercising professional judgement on whether the risks are reasonable when set subjectively against the cost of further risk reduction.

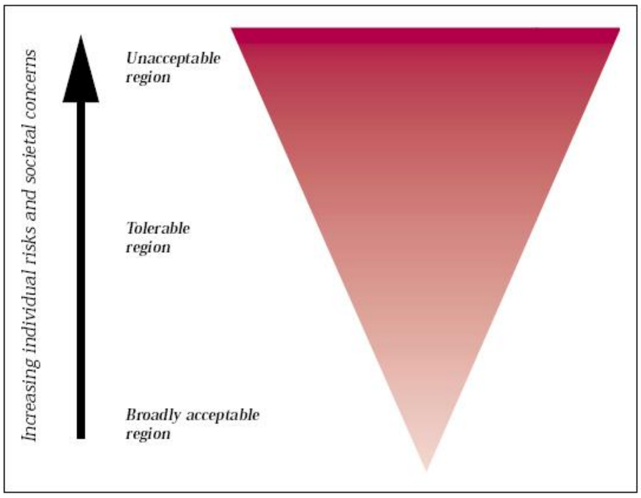

Fig.11. HSE Framework for the Tolerability of Risks (R2P2)

3.3 Approaches to Arguing ALARP

There are three strategies by which duty holders can argue that risks have been reduced to ALARP, shown below in order of precedence:

- **Good Practice Arguments:** In most situations ALARP should be argued by a comparison between the control measures a duty holder has in place and the measures HSE would normally expect to see in such circumstances (i.e. relevant good practice).

- **Qualitative First Principles Arguments:** The second approach to arguing ALARP is from "first principles", and HSE stress that often "first principles" comparisons can be done qualitatively (i.e. by applying common sense and/or exercising professional judgement, or experience).
- **Quantitative First Principles Arguments:** HSE state that there are some instances (often in high hazard industries or where there is a new technology with possibly serious consequences) where the situation is less clear-cut. In these instances, a more formal Cost Benefit Analysis (CBA) may provide additional insight to help come to a judgement. In any case, the outcome of a CBA is only one of several considerations that go towards the judgement that a risk has been reduced ALARP.

The level of detail in a risk assessment should be proportionate to the level of the intrinsic hazards. In general, the greater the magnitude of the hazards under consideration, and the greater the complexity of the systems being considered, the greater the degree of rigour and robustness (and hence the greater the level of detail) HSE requires in arguments to show that risks have been reduced ALARP. The level of risk arising from the undertaking should therefore determine the degree of sophistication needed in the risk assessment.

HSE RR151 [8] identifies a common pitfall of carrying out a detailed Quantified Risk Assessment without first considering whether any relevant good practice was applicable, or when relevant good practice exists. In such cases, resources are expended needlessly on quantified analysis, sometimes at the expense of identifying and implementing risk reduction measures which are known good practice.

The concept of tolerability implies that existing control measures should be periodically reviewed to ensure they are properly applied and still appropriate. Whether they are still appropriate will depend on matters such as the availability of new options for reducing or eliminating risks due to technological progress, changes in society's perception of the particular risks, changes in our understanding of the risk analysis, the uncertainty attached to the risk estimates, and new lessons from accidents and incidents etc. Such reviews should figure prominently in updates of the risk assessments (e.g. Safety Case Reports).

3.4 Gross Disproportion

One of HSE's principles stated in R2P2 [4] is that:

> "there should be a transparent bias on the side of health and safety. For duty holders, the test of 'gross disproportion' implies that, at least, there is a need to err on the side of safety in the computation of health and safety costs and benefits."

Where Cost-Benefit Analysis is used to justify that Risks are ALARP, there is a need to apply a "Disproportion Factor" (DF) which reflects this "bias on the side of health and safety". In consideration of what DF should be considered appropriate, HSE have said:

> "Although there is no authoritative case law which considers the question, we believe it is right that the greater the risk: the higher the proportion may be before being considered 'gross'. But the disproportion must always be gross."

> "HSE has not formulated an algorithm which can be used to determine the proportion factor for a given level of risk. The extent of the bias must be argued in the light of all the circumstances. It may be possible to come to a view in particular circumstances by examining what factor has been applied in comparable circumstances elsewhere to that kind of hazard or in that particular industry."

> "Taking greater account of the benefits as the risk increases also compensates to some extent for imprecision in the comparison of costs and the benefits. It again errs on the side of safety, since the consequences of the imprecision have greater impact, in terms of the degree of unanticipated death and injury, as the level of risk rises."

Widespread practice is for the value of the Disproportion Factor (DF) to increase for Risks further away from the Broadly Acceptable region. Generally DF values between 1 and 10 are used, as illustrated in Figure 12 below.

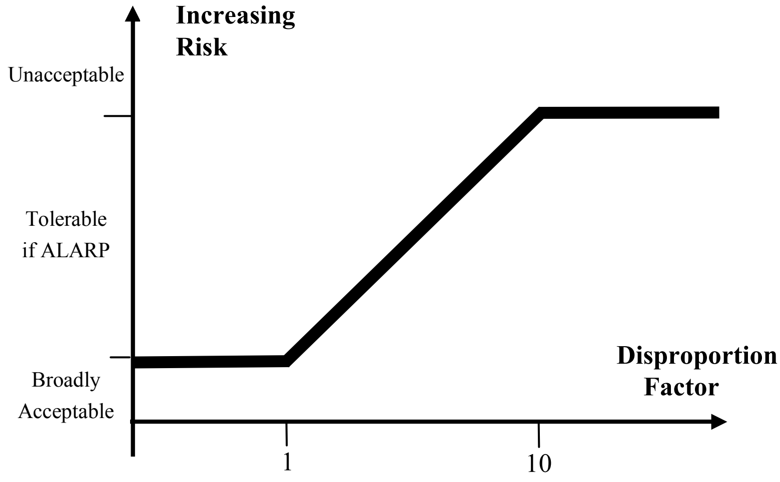

Fig. 12. Example of how DF for CBA ALARP Arguments Increases with Increasing Risk

It should be noted that Gross Disproportion only applies legally to the human aspects of the possible loss (the fatalities and injuries). It does not apply to other elements such as financial loss, asset damage or reputation degradation.

3.5 Risks One-by-one

ALARP arguments made on the basis of "Good Practice" are often produced at the level of "single risks". This can be because "Good Practice" is defined at a comparable level and shows duty holders how to control that type of "single risk". For example, different Codes of Practice or Standards must be consulted to understand how to control hazards relating to electricity or working at height or explosive atmospheres.

Any Quantified Risk Assessment should consider the total Risk exposure of individuals and groups, resulting from all possible Accidents. Without this, there is the possibility of "Accident splitting" or "salami slicing" whereby intolerable overall Risks are made to appear less significant by considering them separately and comparing them separately against overall Risk Criteria.

HSE have identified this flawed approach in HSE RR151 [8] which includes the following Case Study:

Case Study 20: 'Inappropriate use of risk criteria'

A company's COMAH safety report tried to demonstrate that the site's risk was ALARP by comparing the risk from the identified hazardous scenarios to the tolerability criteria for individual and societal risk found in R2P2). The risk from each scenario was found to be below the tolerable level and therefore judged to be broadly acceptable.

Pitfall - The criteria in R2P2 are intended to be used for a site's overall individual risk or societal risk. They are therefore not suitable when trying to determine whether or not the risk from individual hazardous scenarios is ALARP. The risk from all hazardous scenarios should be summed before comparing to the criteria.

Other problems with RA - The societal risk criteria point in R2P2 is for the unacceptable level. Risk below this point is in the ALARP region and therefore cannot necessarily be considered to be broadly acceptable.

If 'single risks' are compared with risk tolerability criteria defined for overall risk (e.g. total individual risk per working year), then they will seem to be much more acceptable than they should be. If there are several ‘single risks’, then each may separately seem to be ‘broadly acceptable’ whereas the individual is exposed to an overall risk that should be judged only ‘tolerable’, or even ‘unacceptable’.

Furthermore, if ALARP arguments based on Cost Benefit Analysis (CBA) are made for ‘single risks’ without appreciating the aggregated risk, too low a Disproportion Factor (DF) will be used and incorrect decisions may be reached to reject risk reduction measures as being ‘grossly disproportionate’. Where ALARP arguments based on CBA are made, they should be based on the aggregated/cumulative risk, compared against the appropriate criteria for overall risk.

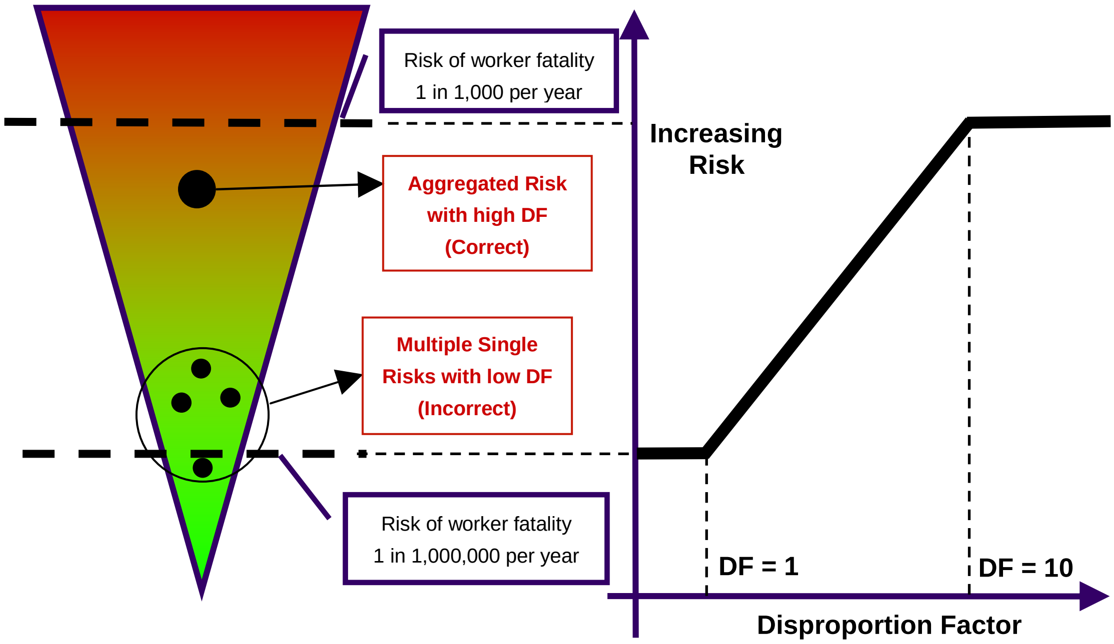

Fig. 13. Comparing Single Risks with Overall Criteria Gives Misleading Tolerability and Incorrect Disproportion Factor

It is noted that comparing the aggregated risk (if known) against overall risk criteria will provide a DF that should be used for CBA on any safety improvements that are being considered. It is the absolute position of the overall risk that determines the DF, rather than that of a ‘single risk’. It is the incremental improvement in the aggregated risk that is of interest, rather than the change in the ‘single risk’ issue. These incremental improvements may be the same, but they could be different if one safety improvement affects more than one ‘single risk’.

4. Conclusions

Risk Matrices are a flexible and powerful tool for evaluating the significance of “single risks”. They provide a readily-understood visual representation of the Risk, as well as its components (namely likelihood and consequence).

Risk Matrices may be Qualitative or Semi-quantitative, but provide only a **relative** indication of the significance of “single risks”, This indication of significance is typically used for prioritising issues for risk reduction and identifying areas for more detailed assessment.

Because Risk Matrices are so useful, they have become very widespread and there are many variants. There have also been many attempts to extend their function. In this paper we have examined the dangers of two adaptations which have led to problems.

Firstly, Risk Matrices are sometimes used to attempt to evaluate Risk in an absolute way against quantitative Risk requirements. This is flawed because:

- The “single risks” evaluated by Risk Matrix are not uniquely defined and are a human view of part of the overall risk exposure. As such, their selection is a matter of style, varying from analyst to analyst: they have no fixed meaning.

- Criteria for overall risk are usually defined in a person-centric way and it is not possible to translate these values consistently into boundaries for a Risk Matrix that is system-centric or activity-centric.
- Risk Criteria stated as quantitative values will require Quantitative Risk Assessment to judge whether they have been met. This goes beyond the limits of Risk Matrices which are semi-quantitative at most.

Secondly, Risk Matrices are sometimes used as a way of determining whether 'single risks' are tolerable and ALARP on a one-by-one basis. Where this is done there is a very real risk that poor decisions will be taken and safety resources will be improperly allocated.

ALARP should be applied at the level of the aggregated/accumulative total risk from a system/activity, but must be done within the context of the UK HSE Tolerability of Risk framework for risk to individuals (real or hypothetical).

Risk assessment centred on the person is the only way to truly understand the total level of risk to which an individual is exposed throughout a typical working year. This requires a risk assessment of all the individual systems/activities to which a 'most-at-risk hypothetical person' is exposed, but the next step of assessing the level of accumulated risk from all those systems/activities is the key step in understanding the level of risk that senior managers are really approving.

Without taking this next step and assessing the aggregated individual risk, and then taking ALARP decisions at this level, the duty holder will potentially make misinformed decisions on risk reduction and the allocation of safety resources to reduce risk to a level that is both tolerable and ALARP.

References

1. Her Majesty's Stationery Office, The Management of Health and Safety at Work Regulations 1999 (MHSWR, 1999)
2. HSE Five steps to risk assessment. HSE Books, INDG 163(rev2)
3. Department of Defence. System Safety Program Requirements. Military Standard Mil-Std-882C, 19 January 1993.
4. UK Health & Safety Executive (HSE). Reducing Risks Protecting People. HSE's decision making process. ISBN 0 7176 2151 0. HSE 2001.. HSE Books.
5. Defence Materiel Standardisation Committee, Defence Standard 00-56: Safety Management Requirements for Defence Systems, Issue 4, June 2007.
6. Her Majesty's Stationery Office, UK Health and Safety at Work Act, 1974.
7. Asquith LJ, Edwards v National Coal Board (1949) 1 KB 704 at 712, CA.
8. HSL report for HSE 2003 Good Practice and Pitfalls in Risk Assessment, HSE Research Report 151 ISBN 0 7176 2732 2.