Confusion over Risk Criteria

Carolyn Nicholls & Jordan Smith, Risk and Hazard Management Consultants, RAS Ltd, 73 New Crane Street, Chester, CH1 4JE

Risk assessments come in all shapes and sizes. They range from the estimation of risk using a team's experience, through to semi- and full quantification. Once the assessment is complete and there is a robust understanding of the risk, the next stage in the decision-making process is to decide if the risk is tolerable. If the risk is tolerable, we move on to satisfy the challenge of demonstrating the risk is ALARP. Those key decisions require operators to be using appropriate risk criteria to support the decisions. There are multiple measures of risk and none are one size fits all. Individual Risk (IR) of fatality is not always the appropriate risk criteria, yet it is the go-to for many carrying out risk analysis. It may be that IR is the most easily found within guidance published by the UK regulator, or that on the surface it seems most simple to apply.

There is a general lack of appreciation for the limitations in using this measure. Examples include:

- Risk Calculated in LOPA. In calculating risk within a LOPA study, perhaps for SIL determination, analysis is made of a single event, and so is calculating just part of an individual's risk. It is therefore not appropriate to compare the risk results of LOPA to IR criteria.
- Representative Scenario Risk. Looking at risk from a single scenario again cannot be compared directly to IR criteria. It may be more appropriate to carry out a 'deep dive' analysis of a scenario or produce analysis for a representative set of scenarios. If the risk is left as a scenario risk, it can't be compared to IR criteria, unless the scenario is modified to account for the other events it represents, and an estimate is made of the cumulative effect from other scenarios.
- Societal Risk. For high hazard events societal risk is more appropriate, and often the more stringent tolerability measure.
- Matrices. Often matrices are not adjusted for scenario risk and are instead only appropriate for whole site risk.
- IR for ALARP decisions. If IR is insisted on as the criteria, the whole risk to the individual needs to be calculated or estimated, which may lead to disproportionate analysis effort. In demonstrating ALARP is it even necessary to calculate IR? If cost benefit analysis is carried out it is the aggregated scenario risk that is used, not IR.
- IR contours. Contour plots can be valuable tools when comparing locations and layouts but adding the risk element doesn't necessarily add value.

Existing guidance is very dispersed. Terminology can often be used interchangeably by those that are not clear or aware of the importance of the subtle differences, so exacerbating the confusion.

1. Introduction

Before identifying which risk criteria will be used, it is important to reflect on why we are doing our risk analysis, and what questions we are looking to answer. It sounds obvious, but if we don't set off with clear intentions, the decisions we make along the way could lead to unsupported conclusions, and ultimately inappropriate risk management decisions. There are many questions to be answered when reflecting on why we are doing the analysis, such as what can go wrong, where do we need to invest further risk reduction effort, and are we comfortable to continue operating? So it is unlikely that a single set of criteria can be used to inform them all. Rather than looking to dumb down assessments to achieve this one size fits all ideal, we should be looking to increase awareness and knowledge of the differences and the essential reasons for them.

The ubiquitous use and acceptance of Individual Risk (IR) as the single criteria for acceptance of risk will be discussed and challenged within this paper, with an aim of providing clarity, and suggestions of more appropriate alternatives.

1.1 Key Definitions

As discussed by Stanley et al (2018) there is a general misuse or merging of terminology throughout the industry. To assist in the understanding of this paper, a list of terminology definitions is provided below. This is not an exhaustive list, intended instead to define the main terminologies of this paper. Other terms and anacronyms are defined within the main body of text. These definitions have been simplified and are based on accepted industry guides such as R2P2, they are also grouped to show the relationships between each term.

Risk Tolerability Criteria – Defined criteria giving (in the UK) broadly acceptable, tolerable if ALARP and intolerable levels of risk. These are set values for individual risk. For societal risk, they are usually defined on a risk matrix or FN curve, with acceptable frequency (F) reducing as the potential number of fatalities (N) increases. The published criteria are set maximums from all hazards and must therefore be adapted to produce target risks for single events or scenarios.

Individual Risk (IR) or Individual Risk of Fatality (IRF) - Total risk to a specified individual from all hazards to which that individual is exposed.

Societal Risk – Total risk to all populations from all hazards to which they are exposed, based on the frequency of each hazard and total number of people affected. Common synonyms for Societal Risk include:

Potential Loss of Life (PLL) - Calculated as the probability of a fatality (per year). It may be presented as the probability of fatality per year from a single scenario, or from all scenarios at a facility.

Aggregated Risk (AR) – Total risk from a hazard based on the potential number of people affected and the frequency of that event per year.

Group Risk - Risk to a group or population based on the number of people affected and frequency of each hazard which can affect said group.

Risk Target – The defined maximum allowed risk of a single event, which can be compared to the actual risk of an event to assess whether it is at an acceptable level or not. This can be used to define the Frequency Target for an event once the severity (number of people affected) is taken into account. These are set by the individual companies and should not be confused with regulator driven risk tolerability.

Frequency Target – The defined maximum allowed frequency for a single event, which can be compared to the actual frequency of an event to assess whether the frequency is acceptable or not - usually extracted from corporate risk targets (often from a risk matrix) and changed depending on event severity. These are set by the individual companies and should not be confused with regulator driven risk tolerability.

2. Background

There are two main underlying issues requiring attention, the first is using the right risk assessment tool, and the second is presenting the risk against the right criteria.

2.1 Risk Assessment Tools

There is a huge number of risk assessment tools available, ranging from qualitative methods through to fully quantified methods. Process safety hazards are complex and full of uncertainty, no two incidents are the same. Further uncertainty and simplifications are added within the assessment process itself, whether that be due to judgements and experience being used in the hazard identification step, or limitations on validation of models used to understand consequences. The objective of using risk assessment tools is to add some clarity to support decision making (Thomas and Bratvold, 2014), but we must be aware of limitations and only use a tool as it was originally intended (Aven, 2016). We must not rely on tools having powers beyond being an aid or support to provide consistency and structure within risk assessments. Any risk assessment is only as good as the people involved (Pritchard et al, 2010), applying tools correctly with appropriate awareness of their limitations. Anyone involved in risk management should be aware of the intentions of these tools to provide guidance, but be clear that the question of 'how safe is safe enough?' is not one that should be treated as a true science (Hansson and Aven, 2014).

One clear example of where a tool has been stretched over time to deliver more than it was intended is HAZOP, and similarly other hazard identification techniques. This is another example of where terminology can lead to confusion, as people may refer to HAZOP as a risk assessment tool. In our experience the main reason that HAZOP is becoming so entwined and linked to risk assessment is that there is a drive to slimline the process and save resources by combining methods into a single function. The ever growing desire to develop unified methods which carry us through the whole risk assessment process in one easy step, is leading to blurring of the lines that distinguish between a hazard identification tool and a risk assessment technique. In a drive for efficiency, we are forgetting the purpose of each process, and looking to tick the box that they are done, often providing meaningless numbers that do not help with risk management decisions. The mention of tolerability and risk criteria are sneaking into the assessment ever earlier in the process.

With such a wide spectrum of risk assessment tools available, it should be no surprise that a single set of risk acceptance criteria such as Individual Risk of Fatality (IRF) per year is not appropriate for use in all of them.

2.2 Risk Criteria

Individual risk of fatality is widely used as a single set of criteria for tolerability of risk. As a result, we are driven to tolerability criteria based on numerical risk values, which can lead to over complication and over-quantification of risk with little benefit to industry and those charged with making decisions. Numerical quantification, with ever more layers of assumption and uncertainty, can fool us into a false sense of accuracy and security when it isn't necessarily needed or justified.

Depending on the risk assessment technique employed, it may be more appropriate to use alternative criteria. For example, simple criteria associated with the number of independent protection layers required, depending on the level of harm that is foreseen. The UK regulator for the chemical industry is clear that assessments should be proportionate, so for low and medium proportionality sites, this would be a supported approach.

Confusion of what tools are appropriate, lack of clarity in the objectives of particular pieces of analysis, and further confusion over the most appropriate criteria with which to compare results all lead to a common misuse of Individual Risk as criteria. To address the issues of using the right risk assessment tool, and presenting the risk against the right criteria, the concepts behind the tools and criteria need to be understood and a clear idea of why the risk assessment is being carried out is needed.

3. Application of Common Tools

The following table provides a summary of commonly used tools with a reminder of their intended use, and some indication of appropriate risk presentation forms:

Table 1 Summary of tools and processes

Tool / Process	Objectives	Appropriate risk presentation
HAZID/HAZOP	Identify what can go wrong. Qualitative judgement of safeguards. Identify improvement actions.	Severity Only.
LOPA	Assess layers of protection. Identify risk gaps by comparing result to a target event likelihood. Inform SIL requirement.	Frequency of event per year.
ALARP Review	Demonstrate taken all measures necessary. Identify further risk controls. Cost benefit for risk reduction options.	Cost vs Benefit in Monetary terms calculated from PLL.
QRA	Deeper understanding of consequences through modelling. Establish a total site risk figure. Establish if risk is tolerable. Understand key contributing events.	Individual risk of fatality. FN Curve. Individual risk contours. PLL.
Scenario Assessment	Detailed analysis of a single key event(s). Reduce or refine judgements using quantification or semi-quantification.	Scenario risk contribution. Risk matrix. PLL.
Occupied Buildings Study	Locate buildings safely. Understand risk to building occupants.	Consequence contours. Individual risk to building occupants.
Layout Study	Optimise layout. Improve inherent safety. Avoid escalation.	Consequence contours.

4. Tolerability of Risk

The concept of tolerability of risk is one that sits outside of the regulations themselves. The regulations do not use language such as acceptability or tolerance, in terms of managing risk to an appropriate level. The following are two key quotations from UK legislation:

- Health and Safety at Work (HSW) etc. Act 1974 "*It shall be the duty of every employer to ensure, **so far as is reasonably practicable**, the health, safety and welfare at work of all his employees."* **SFAIRP**
- The Control of Major Accident Hazards Regulations 2015, SI 2015/483 *"Every operator must take **all measures necessary** to prevent major accidents and to limit their consequences for human health and the environment."* **ALARP**

The focus of the legislation is on reducing risk to **As Low As Reasonably Practicable (ALARP)**; the concept of tolerability and acceptance criteria has since been added within guidance as an aid to both the regulator and industry. Our experience has shown that this intention has been lost, with an obsession for numbers and meeting target numbers meaning that we lose focus on demonstrating ALARP, looking to meet criteria targets rather than seeking the best solutions for managing risks on our facilities (Aven, 2016).

4.1 UK Risk Tolerability Criteria

If we are to use any criteria, we must understand where they came from, and indeed their intended purpose. In 2001 the HSE published 'Reducing Risks, Protecting People- HSEs decision-making process' (R2P2) which sets out the HSE's approach to making decisions about SFAIRP and ALARP.

Individual Risk Criteria

Individual risk of fatality (IRF) is calculated by summing all events on a facility that have the potential to lead to a fatality in the case of a specified or theoretical individual, taking into account their occupancy and chance of exposure to relevant hazards.

The R2P2 document clearly defines tolerance criteria for individual risk of fatality (IRF) as follows:

- Individual risk over 1E-04 per year is intolerable for the general public.
- Individual risk over 1E-03 per year is intolerable for workers.
- Individual risk below 1E-06 per year can be considered as broadly acceptable without the need for further expenditure on risk reduction measures, as it is considered comparable to everyday risks that people consider insignificant.
- Depending on the above, individual risk between 1E-03 and 1E-06 per year is tolerable provided all practical measures to reduce risk have been considered and implemented if cost effective. This is referred to as the ALARP (As Low As Reasonably Practical) region.

These tolerability criteria are commonly shown on a risk triangle, see Figure 1 below.

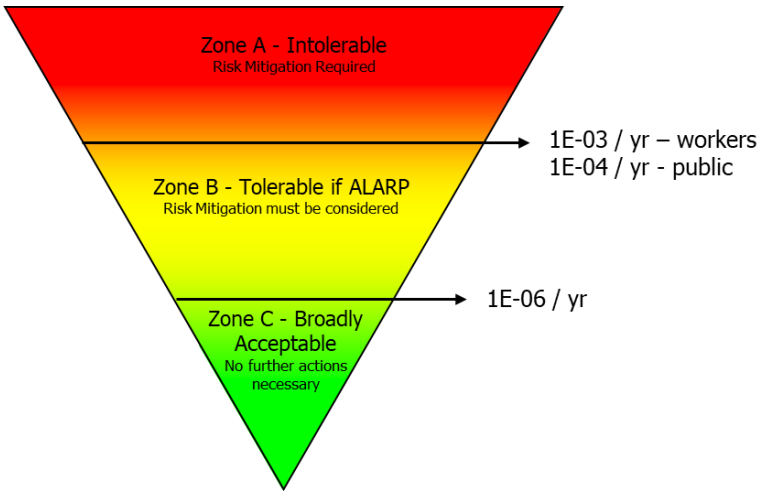

Figure 1 Individual Risk Triangle

Societal Risk Criteria

Societal risk is calculated based on the predicted frequency of an event (per annum) and the potential number of people affected by it. Societal risk is expressed as a single figure, calculated by multiplying the potential number of fatalities of an event by the frequency per year. This can be expressed as a numerical value and can be referred to as Aggregated Risk (AR), group risk, or Potential loss of life (PLL). These can be calculated per event, per scenario, or for the facility as a whole. To illustrate societal risk for an establishment in a graphical form, an FN curve can be used. This is a plot of the cumulative frequency (F) of scenarios against the number of people that could potentially be affected (N).

Reducing Risk, Protecting People (R2P2) is less clear in explaining the criteria for societal risk. One data point is provided for guidance on the level of societal risk that would be considered intolerable. It states that the frequency of an accident occurring which results in 50 or more fatalities will be intolerable if it occurs more than once in five thousand years (i.e. 2E-04 per year). This point is expanded on within HSE's 'Guidance on ALARP Decisions in COMAH' (2012) which takes the single data point from R2P2 and the work of Ball and Floyd (1998) to define tolerability criteria for societal risk. This work expresses tolerability criteria through definition of the bounds on an FN curve. It describes that those bounds should be drawn with a slope of -1 (as confirmed by the HSE's Risk Assessment Policy Unit (RAPU)) and sets the boundary between the broadly acceptable region and the tolerable if ALARP region at two orders of magnitude below the uncomfortably high (sometimes referred to as the intolerable) boundary. The resulting FN graph with acceptable and uncomfortably high frequency limits for a range of fatalities (1 - 1000) is provided in Figure 2 below. Plotting a facility's calculated risk as an FN curve overlaid on these tolerability bounds allows the tolerability of societal risk to be determined.

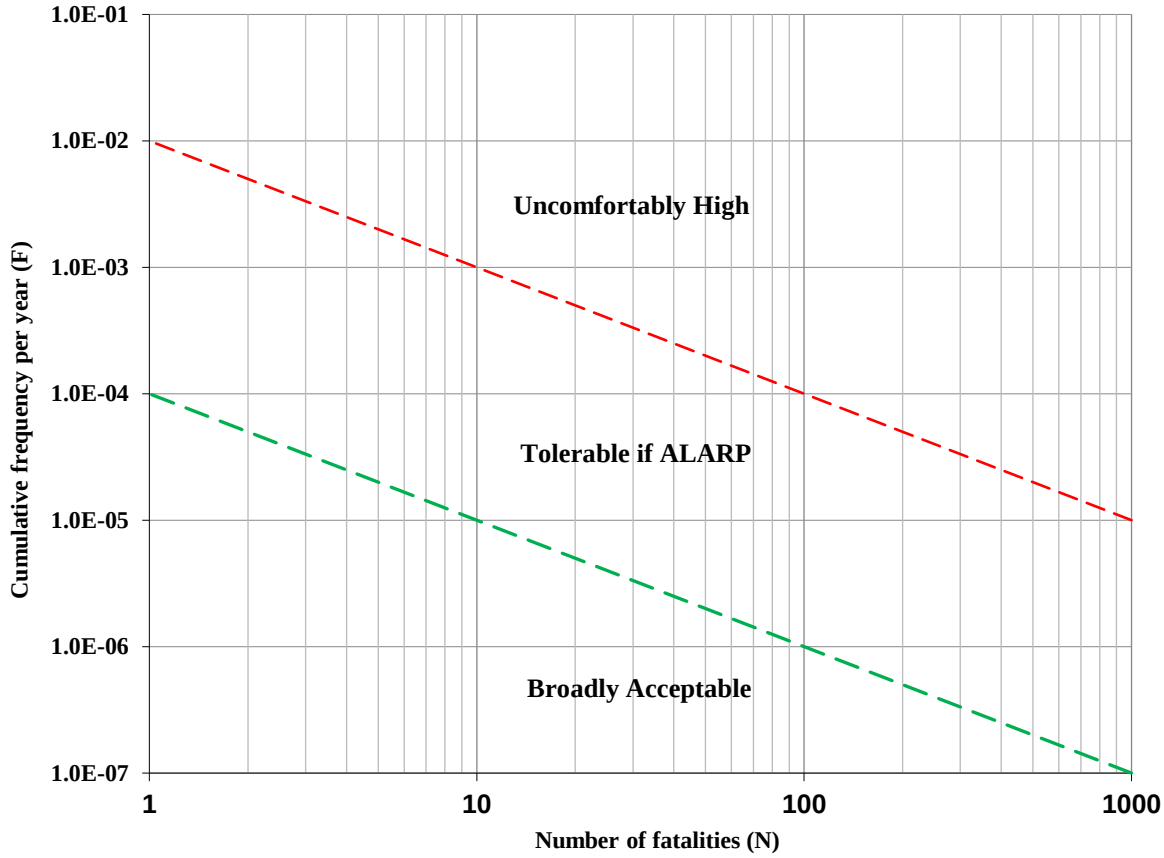

Figure 2 Societal risk tolerability FN graph (log-log scale)

Both of the risk criteria for IRF and societal risk provided by the UK regulator are intended for judgement of facility risk. This is a very important point that should not be lost. In the majority of risk assessment work only a portion of total facility risk is analysed as only a particular event, or small number of key events are considered, and so using these risk criteria as they are without adjustment is meaningless (Baybutt 2017).

5. Application of risk criteria

Individual risk of fatality (IRF) is not the appropriate criteria to be used in the majority of typical risk assessments; as such, the following sections look to describe particular applications, these sections give further discussion of why IRF should not typically be used in those situations, and offer an acceptable alternative.

5.1 Risk criteria in Hazard Identification Studies

Risk criteria should not be used in Hazard Identification Studies. In our experience, a huge amount of time is wasted in debating the frequency of events that are beyond our experience and judgement capabilities.

The purpose of hazard identification tools is to identify what can go wrong, not assess risk. However, there is a common temptation to add risk judgements at this early stage of the process. Clearly there may be a need to screen those events identified and select a few to interrogate further, but that screening could be done on a broad ranking of potential severity. Instead we see overly complex risk judgements being made. For complex and high hazard events, we should not be using our experience alone, and while we recognise that and pass events forward to further quantification in LOPA for example, we still see many HAZOPS where teams of people have wasted huge amounts of time trying to categorise frequency with a misleading amount of resolution. Many of the events discussed are well outside the experience of the team, so cannot be categorised with any accuracy. The use of risk criteria within a HAZOP is therefore ill-conceived and should be avoided. Screening of events should be based on broad judgements of severity only, with the use of a high, medium or low to distinguish between events with potential for multiple fatality put forward for quantification, single fatality events put forward for semi-quantification and low severity events screened out of any further quantification, respectively.

5.2 Risk criteria for LOPA and other single event assessments

As individual risk criteria are well known, it is often misapplied within LOPA or single event assessments. Baybutt (2017) discusses the challenges of risk practitioners mistakenly using overall facility risk criteria. Stanley et al (2018) looked to establish appropriate risk targets for use in LOPA recognising this common pitfall.

In calculating risk within a LOPA or scenario study it is important to recognise that only a single event is being studied, and just part of an individual's or group's risk is being calculated. It is therefore inappropriate to compare the risk results of LOPA to IRF. There are several alternative options to determine tolerability a) qualitative judgement, b) comparison to specially derived scenario criteria, c) enhancement and aggregation to total facility numbers to compare to facility criteria.

A) Qualitative Acceptance Criteria

If quantification is not required, a more proportionate approach may be to demonstrate that there are sufficient numbers of independent protection layers, without making attempt to quantify their reliability beyond the standard assumptions of orders of magnitude reduction.

ONR's Safety Assessment Principles (SAPs), together with supporting Technical Assessment Guides (TAGs), discuss criteria which are focused on the target reliability for safety systems. Using this approach, a severity rather than risk based system could be developed. Such a method could be used to demonstrate that for a low severity event a single IPL is required, for a medium severity event two IPLs are required and so on. The production of qualitative criteria focusing on the number of IPLs warrants further exploration in order to provide an alternative less complicated solution.

B) Scenario Criteria

If quantification is needed, perhaps to establish a risk gap and identify a specific PFD for SIL determination, then risk criteria for a single event (with acknowledgement of risk contribution from other events) is needed.

Stanley et al (2018) proposed the following criteria for use in LOPA and other single event assessments, based on the criteria provided for FN curves, and then adjusting for scenario use with an assumption of a total of ten scenarios contributing to the risk exposure of those impacted by the event being studied:

Table 2 LOPA Target Frequencies (Societal Risk)

Severity / Harm	Target Frequency (per year)	
	Existing Establishments Target Middle of TifALARP	New Establishments Target Broadly Acceptable
Injury (0.1 Fatality)	1E-03	1E-04
1 Fatality	1E-04	1E-05
2 Fatalities	5E-05	5E-06
10 Fatalities	1E-05	1E-06

They recognise that for new establishments more effort should be invested in reaching the broadly acceptable point and set the target there. For existing establishments, an allowance is given for demonstration of ALARP by use of measures that may not be fully independent or easily quantified, and therefore set the target to be the middle point between intolerable and broadly acceptable frequencies for the given number of fatalities of each LOPA (including the factor reduction to account for other scenarios).

C) Enhancement and Aggregation

For high hazard scenarios, it may be appropriate to carry out a 'deep dive' quantified analysis of that scenario. A detailed quantified analysis may be produced for a representative set of scenarios. When detailed analysis is carried out, there is a need to compare the calculated numerical value of risk to numerical values for tolerance of that risk such as IRF. Here it must be recognised that if the risk is left as a scenario risk, it cannot be compared directly to IRF criteria. If this is to be done it must be modified to account for the other events it represents, and an estimate must be made of the cumulative effect from other scenarios. Only by aggregating in this way will a total facility risk, that can be compared to IRF and societal risk criteria, be achieved. Furthermore, to establish the risk to a specified individual there is a need to understand site occupancy and time spent in various areas, through time in motion assessments or similar, in order to properly proportion up a worker's time spent within the hazard ranges of each scenario. This is difficult, time consuming and unlikely to be correct due to huge variabilities.

It should be noted that generating the risk figures in this way is subject to multiple layers of adjustment, so it may be more appropriate to use scenario based risk criteria if this needs to be avoided.

5.3 Societal Risk

Stanley et al (2018) quote several sources all agreeing that in assessing risk for high hazard events, societal risk is more appropriate, and often the more stringent tolerability measure, though as discussed earlier it is a measure that is less well described, and less well used.

In the majority of risk assessment, societal risk will be calculated, even if it is not expressed as such. This is true for the LOPA and single scenario cases described above. The suggested targets for LOPA are derived from the FN Curve, and adjusted for use in a scenario assessment by an order of magnitude to account for risk contribution within the facility from a total of ten similar scale events.

If a single risk value is preferred rather than a frequency paired with a severity band as shown above, then PLL is calculated by multiplying the number of fatalities by the frequency of event. PLL criteria can be derived and adjusted for scenario use in the same way as the frequency targets described in the earlier section, giving the following values:

Table 2 PLL Criteria

Criteria	Whole Facility	Single Scenario
Intolerable / Uncomfortably high	1E-02 / yr	1E-03 / yr
Broadly Acceptable	1E-04 / yr	1E-05 / yr

These values are based on the same broad assumption of 10 similarly sized scenarios contributing to overall site risk, and therefore may not be appropriate to all sites.

5.4 Risk Matrices

The most appropriate use of a risk matrix is to give a visual indication of the relative positioning of various events, to allow prioritisation or risk management effort. It is important that matrices are built with sufficient resolution, otherwise they fail in this purpose (Nicholls and Carroll 2017).

Matrices can be used as a visual presentation of societal risk, and used to judge tolerability, however they must be calibrated for the type of risk that has been calculated. Baybutt (2015) discusses the issues around the calibration of risk matrices. He puts forward pitfalls relating to the lack of a link between risk tolerance criteria and the events to which they are applied, making it clear that there is a common issue in using facility criteria for scenarios.

Risk matrices are effectively a low resolution FN curve. In the same way as described in earlier sections, if single scenarios are to be plotted on the matrix it will need to be adjusted and calibrated for the scale of event in question, while bearing in mind the number of other events that might be present at the facility as a whole.

Alternatively, all scenarios could be treated as FN pairs, and if the frequency is presented in a cumulative way, it can be shown on a matrix derived directly from the FN Curve (though in this case it may be as effective and more efficient to show directly on the FN curve).

5.5 ALARP Decisions

Although risk can be quantified to determine tolerability, it does not demonstrate that risk is ALARP.

In most cases, quantification of risk is not required to make this demonstration. Demonstration of ALARP can be done qualitatively showing relevant good practice has been adopted, and there has been effort to identify what more could be done. Where further measures are identified, they should be easily accepted or rejected on clear merits or excessive cost without the need to quantify. This is supported by Hansson and Aven (2014) who stated that the question of ‘how safe is safe enough’ is not something that should be answered purely by science.

Only in a small number of cases will cost benefit analysis be needed to help aid decisions. For cost benefit, the total risk benefit of the measure under consideration needs to be quantified, and so a total risk, usually expressed as PLL is required. In our experience the desire to present IRF leads to confusion, whereby IR is used in the cost benefit analysis calculations in error, or PLL is derived from IR in unusual ways.

This again suggests that individual risk is an inappropriate risk measure when considering regulations whose focus is on demonstrating ALARP.

5.6 IR Contours

Contour plots can be a valuable tool when comparing locations and layout options for new facilities.

However, in plotting individual risk contours there are a number of simplifying assumptions made around occupancy which need to be recognised. In many cases it would be more appropriate to consider layout with consequence contours instead, the additional quantification to establish risk is not helpful or meaningful in this context.

IR contours can be misleading, for example showing intolerable IR frequencies, but without taking into account occupancy this does not equate to intolerable risk in reality. Conversely, IR contours can show acceptable risk across the whole site, but when PLL is calculated it may be at an intolerable level if there are large groups of people exposed.

This reemphasises the need to ensure that the correct risk criteria are used, and that the methods for analysis and criteria chosen are meaningful and useful to those making decisions based on them.

5.7 Communicating Risk

There are many stakeholders and audiences of the risk assessments produced, and so we should keep in mind that the risk criteria and presentation measures adopted must be appropriate to those audiences and the information they need.

Those involved in process safety need to be aware of their responsibilities, so they can be sure they are asking the right questions, at the right time, of the right people. As risk analysts we need to be sure we are providing information that answers those questions.

We should not be tempted to dumb down risk presentation into what looks to be a simple or more visual form if it leads to meaningless results. The familiar look of a risk matrix, and its relatively simple appearance may lead to those making risk decisions to ask for them to be used, without proper understanding of the implications. Instead we should look to educate our audiences in the correct terminology and tools. Process hazards are complex, and by dumbing down the presentation of them we can inadvertently give false impressions and suggest overly simple solutions are adequate to manage the complex risks.

6. Conclusion

There is clearly a large amount of confusion, not helped by an industry that is constantly looking for ways to increase efficiency and speed, leading to a loss of focus. Process safety and understanding of complex hazards needs to be given proper due care and attention. Proper attention does not necessarily mean quantification, efforts to understand risk should be proportionate.

Confusion of what tools are appropriate, lack of clarity in the objectives of particular pieces of analysis, and further confusion over the most appropriate criteria with which to compare results all lead to a common misuse of Individual Risk as criteria. In common application, societal risk is the more appropriate risk measure to use, yet the guidance available on how to apply it is dispersed. This paper aimed to raise awareness of this issue and provide some clarity by proposing appropriate solutions for the most common risk assessment applications.

To address the issues of using the right risk assessment tool, and presenting the risk against the right criteria, the concepts behind the tools and criteria need to be understood and a clear idea of why any risk assessment is being carried out is needed. This paper urges all process safety professionals to not just accept the request to use a set tool, or a set of criteria and use them blindly, instead confirm they are appropriate for your specific applications.

Further investigation is needed to develop more qualitative tolerability criteria, perhaps focusing on the number of IPLs required for various severity levels of scenarios, to allow unnecessary quantification to be avoided.

There is a need for the issues raised in this paper to be recognised more widely and for joint industry and regulator guidance to be provided to make the solutions more readily accessible for all.

7. References

Aven (2016), Risk assessment and risk management: Review of recent advances on their foundation.

Ball, D. J. and Floyd P. J. (1998) 'Societal risk,' Crown Copyright, available from Health & Safety Executive, Risk Assessment Policy Unit, 2 Southwark Bridge, London SE1 9HS.

Baybutt, P (2015), Calibration of Risk Matrices for Process Safety, *Journal of Loss Prevention in the Process Industries*, 38: 163-168.

Baybutt (2017), Overcoming challenges in using layers of protection analysis (LOPA) to determine safety integrity levels (SILs), *Journal of Loss Prevention in the Process Industries*, 48: 32-40.

Hansson and Aven (2014), Is Risk Analysis Scientific? *Risk Analysis*, 34(7): 1173-1183.

Health and Safety at Work (HSW) etc. Act 1974

HSE (2001), Reducing Risks, Protecting People- HSEs decision-making process, Norwich: Her Majesties Stationery Office.

HSE (2012), Guidance on ALARP Decisions in COMAH - SPC/Permissioning/37. [Online] Available from: http://www.hse.gov.uk/foi/internalops/hid_circs/permissioning/spc_perm_37/ [Accessed 16 January 2020].

Nicholls and Carroll (RAS Ltd, UK), 2017. Is there value in a 'one size fits all' approach to risk matrices? 27th Institution of Chemical Engineers Symposium on Hazards - symposium series 162.

ONR (2020), Safety Assessment Principles for Nuclear Facilities. [Online] Available from: http://www.onr.org.uk/saps/saps2014.pdf [Accessed 16 January 2020].

Pritchard, York, and Hannegan (2010), Drilling hazard management: The value of risk assessment. *World Oil*, October 2010: 43–52. [Online] Available from: http://www.successful-energy.com/wp-content/uploads/ 2011/02/WO1010_Series_2_Final.pdf [Accessed 16/01/2020].

Stanley, Nicholls, Burnup and Smith (RAS Ltd, UK), 2018, Risk tolerability targets; misconceived, misunderstood and misapplied. 28th Institution of Chemical Engineers Symposium on Hazards - symposium series 163.

The Control of Major Accident Hazards Regulations 2015, SI 2015/483

Thomas and Bratvold (2014), The Risk of Using Risk Matrices. *Society of Petroleum Engineers Economics and Management*, 6(2): 55–66.