Risk tolerability targets; misconceived, misunderstood and misapplied

Andy Stanley, Director, RAS Ltd, 73 New Crane Street, Chester, CH1 4JE

Carolyn Nicholls, Director, RAS Ltd, 73 New Crane Street, Chester, CH1 4JE

Bob Burnup, Principal Consultant, RAS Ltd, 73 New Crane Street, Chester, CH1 4JE

Jordan Smith, Senior Consultant, RAS Ltd, 73 New Crane Street, Chester, CH1 4JE

Following on from a quantification (or semi-quantification) of risk, the next step in an assessment is to compare the result to a set of risk tolerability criteria. Three regions of risk are usually defined, an unacceptable region and a broadly acceptable region bordering a region of tolerable risk. The tolerability of this middle region is dependent on those risks being As Low As Reasonably Practicable (ALARP). The 2001 HSE publication 'Reducing Risks, Protecting People - HSE's decision-making process' (R2P2) is widely used in industry to set these region boundaries. R2P2 clearly defines the tolerability boundaries for individual risk, and gives guidance regarding societal (or group) risk tolerability.

How these criteria have been applied across the industry varies dramatically, particularly in setting targets for Layer of Protection Analysis (LOPA). LOPA is often used to identify the need for, and define the required integrity of, a Safety Instrumented System (SIS). In this context, it is easy to see how setting the wrong targets for a risk assessment can result in serious consequences. For example, setting a risk target that is too lenient could result in a process that is not adequately protected, and setting a risk target that is too onerous could result in the requirement to needlessly install a high integrity safety system at a significant cost. To set appropriate risk targets, the intricacies between types of risk need to be well understood, and applied correctly for the given situation. Our experience is that this is often misunderstood and misapplied, people take simple rules and apply them in the wrong context.

This paper considers the differences between individual risk and societal (group) risk, and how these are often confused and misapplied. Differences between scenario risk and whole site risk, and how these should be considered in different ways depending on the type of risk and the type of study are then discussed. Risk targets are proposed which may be used in LOPA, including a demonstration of how these were developed and how they meet industry and regulatory standards. How these targets can be used and adapted into existing or new facilities, and how they may interact with existing corporate risk criteria and matrices is also considered.

Keywords: risk, frequency, tolerability, target, LOPA, individual, societal, SIL, ALARP

1. Introduction

Risk tolerability and the concept of ALARP is a subject which has received great attention within the 'major hazard' industries. It is our experience that although the definitions, criteria and guidance on this subject are generally well established, we often find that there are common misconceptions regarding risk, and that the industry is not consistently applying the same rules. All companies are responsible for setting their own risk tolerability criteria, and it should be applicable to the site and hazard in question (BS EN 61511, 2017), however it is our experience that these are often not well justified, sometimes set incorrectly, and often applied in the wrong way.

These issues seem to be highly prominent when used in Layer of Protection Analysis (LOPA). This may be due in part to the added complication of calibrating risk targets for use in single major accident scenarios (or events). We are not the first to notice the frequency of errors in this area, for example Baybutt (2012, 2017) gives multiple examples of the incorrect application of societal and individual risk in LOPA. The UK HSE also note that justification of risk targets used in LOPA is generally poor (HSL & HSE, 2009). A few common real examples of errors which we have experienced are outlined in the box below.

Common mistakes

"if the tolerable line for individual risk is 1E-05 per year, then the tolerable risk to ten individuals must be 1E-06 per year"
- this is incorrect, the tolerability of individual risk does not change based on the number of people potentially affected.

Mixing and matching types of risk within the same assessment
- i.e. using individual risk criteria, and applying occupancy factors for societal risk.

Applying individual risk where societal risk would be more applicable and more stringent, or *vice versa*
- for major accident hazards, societal risks are generally considered more appropriate and are likely to be much more stringent.

Attempting to calibrating a corporate risk matrix to account for individual risk
– as individual risk does not scale based on severity of consequence this is not possible and can result in incorrect application of both types of risk.

"A common error is to use individual risk criteria but calculate group risk measures which can lead to more stringent risk reduction measures than needed" (Baybutt, 2017).

The issues above have highlighted the need to re-affirm accepted industry guidance on the differences between individual and societal risk, and when each should be used. The large variation in company risk targets being used, and methods for setting

them, also suggests that accepted guidance in this area is lacking. It may be the case that the company risk targets have been set incorrectly, that they have been used in error, or that the meaning being them has been lost over time.

In this review paper, LOPA is used as an example throughout, as it is our experience that this is a common place for mistakes to be made, and it is a good demonstration of the scale of the issue. The paper is structured to firstly outline the background within this area of study, followed by detailed descriptions of the different types of risk and how and when each should be used. Following this, a method is outlined for setting risk targets for use in LOPA and some proposed risk targets are provided for existing and new establishments. The aim being that this provides end users with a good understanding of the current best practice, and gives a proposed methodology for setting risk targets which can be applied and adapted as required.

Based on the scope of this paper, operators who are responsible for setting, reviewing or using risk tolerability criteria for their company or establishment would benefit from this review. Likewise, assessors and auditors of risk studies may find this review helpful in spotting and addressing common errors in application.

1.1 Key Definitions

A potential contribution to the confusions outlined above may be to the general misuse or merging of terminology throughout the industry. To assist in the understanding of this review paper, a list of terminology definitions is provided below. This list is not exhaustive but merely aims to define the main terminologies used in this paper, other terms and anacronyms are defined within the main body of text. Frequency and risk figures are quoted throughout the paper in the format 1E-04, which is equal to 1×10^{-4} or 0.0001.

Aggregated Risk (AR) – total risk from a hazard based on the potential number of people affected and the frequency of that event, often used as a synonym for societal risk.

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 changes depending on event severity. These are set by the individual companies and should not be confused with regulator driven risk tolerability.

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, often used as a synonym for societal risk.

Individual Risk (IR) - total risk to a specified individual from all hazards to which that individual is exposed.

Potential Loss of Life (PLL) - calculated as the probability of a fatality (per year) from a hazard, often used as a synonym with societal risk.

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

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.

Societal Risk – total risk to a population based the on the frequency of each hazard and total number of people affected.

2. Background

2.1 Legislation

The UK has in place the Health and Safety at Work (HSW) etc Act (1974) and The Management of Health and Safety at Work Regulations (MHSWR) (1992) that requires workplaces to assess their risks to employees and to other people, and reduce these risks so far as is reasonably practical (SFAIRP). Note that the HSW Act does not explicitly distinguish between types of risks.

For high hazard industries specifically, the Control of Major Accident Hazards (COMAH) Regulations (2015) requires operators of sites with the potential for major accidents to take ‘all measures necessary’ to prevent these occurring. Operators must demonstration that major accident risks have been reduced to As Low As Reasonably Practicable (ALARP).

It is well established (HSL & HSE, 2009) that the application of the SFAIRP and ALARP philosophy are one and the same. To make the demonstration that the duties of SFAIRP and ALARP are met, operators must conduct risk assessments that are proportional to the potential severity of an event, and show that the cost of reducing the risk of the potential event is grossly disproportionate to the cost of introducing further risk reduction measures. Therefore, as concluded by the Societal Risk Technical Advisory Group (HSL & HSE, 2009), for an event capable of harming more than one individual, if societal risk is not considered, then the operator may not have demonstrated compliance with the required health and safety laws. Inversely, if the operator only considers societal risk, for an event which disproportionately affects one individual, the same conclusion can be drawn.

The majority of industry standards, guidance and publications state that both the level of individual and societal risk must be taken into account when deciding whether the risk from an event is unacceptable, tolerable or acceptable (See Section 4). For

COMAH operators, the HSE's COMAH Safety Report Assessment Manual (HSE, 2016) includes guidance that risk assessments for major accident sites should consider both individual and societal risk, and that it should be compared against suitable criteria. This requirement brings in the need to define, justify and then apply risk tolerability criteria in our risk assessments.

2.2 Risk Tolerability Criteria

In 2001 the HSE published 'Reducing Risks, Protecting People- HSEs decision-making process' (R2P2) which sets out the HSEs approach to making decisions about SFAIRP and ALARP.

R2P2 defines the tolerability boundaries for risk to an individual, but is less definite regarding societal risk. One data point is provided for guidance on the level of societal risk that would be deemed intolerable. The HSE's 'Guidance on ALARP Decisions in COMAH' (2012) then uses this data point and the work of Ball and Floyd (1998) to define the criteria that should be used when determining the tolerability of societal risk in risk assessments. Both these concepts are discussed in more detail in the relevant sections of this paper.

As no further criteria has been published for the tolerability of individual and societal risks, it can be inferred that the HSE's criteria is the most appropriate for use in risk assessments by UK COMAH establishments. Industry experience has shown that COMAH sites using company tolerability criteria, are usually required to demonstrate that it is equally or more stringent than that defined by R2P2.

It is worth noting that societal risk has been an area of interest in recent years, particularly regarding the tolerability levels and where it is appropriate to be used. In 2007 the HSE published a consultation document (CD212) that sought views on whether societal risk should be taken into account when assessing safety measures at major accident sites. Based on the results of this consultation, which were that a large majority of responses were in favour of societal risk being a factor in onsite control measures and land use planning decisions, it was agreed that societal risk should be considered going forward (HSL & HSE, 2009). How societal risk is to be practically implemented into government policy is still to be decided. One key issue raised by the UK chemical industry is that risk criteria that are too stringent could put UK businesses at a disadvantage in the worldwide market. At this point in time, there has been no agreement to use local societal risk in LUP, however our experience is that it is used in certain situations e.g. vulnerable populations (schools/hospitals).

3. Scope and Objectives

As can be seen from the background information above, the job of setting risk tolerability criteria for the UK, for both individuals and society is highly complex, and potentially ever changing due to public opinion, fallout following industry accidents, risk aversion, market conditions and industry advances.

The way forward for UK risk tolerability is therefore outside the scope of this paper, the main objective of this paper is to ascribe a method by which appropriate tolerability targets can be selected for any study, based on the type of study, type of event and the people potentially exposed. By providing clear methods and reasoning, risk targets can be set and adapted easily dependent on the latest accepted tolerability criteria of the time.

This paper uses Layer of Protection Analysis (LOPA) as the case study, as it is a relatively simple form of risk assessment that is widely used in industry. It is also our experience that due to the perceived simplicity of LOPA, this is an area where risk targets are often set incorrectly or used inappropriately. When LOPA is then used to set the required Safety Instrumented Level (SIL) for a Safety Instrumented System (SIS), it is clear that errors can be costly to safety and financially.

We hope to add clarity and show simple examples to help reduce the potential for errors, and allow operators to have a better understanding of terminology and methods when working with corporate or third-party methodologies.

4. Risk

Risk in the context of this paper is the risk to people from a particular major accident. Risk to the environment and commercial risks can also be considered in risk assessments for major accident sites, however these are outside of the scope of this review. Safety risks to people can be described in two ways, individual risk and societal risk, both of which take into account the potential consequence of the event and the frequency of the event (per year)

- Individual risk is the risk that a particular individual is affected by an event, and does not take into account the total scale of the event in terms of harm to other people. For example, it is possible to have two events with the same calculated individual risk even if one only results in harm to one person, and the other results in over 100 fatalities.
- Societal risk on the other hand, is calculated taking into account the total number of people affected by an event. For two events of the same likelihood, where one affects one person and the other affects 100 people, the latter would be expected to have societal risk 100 times higher.

As outlined in the Section 2 above, it is clear that both individual risk and societal risks must be taken into account when deciding whether a risk is unacceptable, tolerable if ALARP or broadly acceptable. It is generally accepted that societal risk is much more likely to be relevant to major accident sites in assessing the tolerability of risks, by their very nature high hazard sites have the potential to impact large numbers of people. This is discussed further below in Section 4.3.

These two types of risk to people are discussed in more detail below, and their current industry accepted tolerability levels are described. It should be noted that these criteria are not laid down in law, but have been widely accepted for many years.

4.1 Individual Risk

Individual risk is calculated as the risk of fatality (or serious injury) to a specific person per year. An individual risk calculation is performed for a specified individual. This can be calculated for the perceived most exposed person, but it can also be location specific or for a hypothetical person, as detailed exposure data for a specific individual can be hard to obtain or estimate. There are well established methodologies available for each type of individual risk calculation, for example Baybutt (2012) specifically suggests a method for using hypothetical (representative) individuals in LOPA by summing the frequencies of all LOPA scenarios that potentially impact a specific location.

Individual risk includes the risk from all hazards to which a specific individual is exposed, e.g. J. Bloggs's individual risk is the summation of every hazard to which he is potentially exposed, both at work and at home. It is usually expressed in a single figure, which is the chance of fatality per year for that person. The risk from a single major accident hazard cannot therefore be used to calculate individual risk, as it does not necessarily account for all the sources of harm to which an individual may be exposed from a given establishment. There are methods used in industry to approximate these aspects, which are discussed in Section 5.2 below, and therefore produce an acceptable target risk which can be used in assessments (e.g. LOPA). It is worth noting therefore that a LOPA assessment for a one fatality event will not usually calculate an individual risk value, as the risk of a one fatality event is usually shared between more than one person (i.e. a shift team of operators), and each of those people will be exposed to other risks throughout their day.

If individual risk has been properly calculated, it can be compared with individual risk tolerability criteria. This is clearly defined in R2P2, which is based on individuals' knowledge and acceptance of other day to day risks e.g. car travel. The following individual risks are given for the public, employees and process workers:

- Individual risk over 1E-04 per year is intolerable for the general public
- Individual risk over 1E-03 per year is intolerable for workers (depending on personnel function and location this may be reduced to 1E-04 per year, as per CIA occupied building guidance, 2010)
- 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 are tolerable provided all practical measures to reduce risk have been considered and implemented if cost effective. This is often called the ALARP (As Low As Reasonably Practical) region.

These tolerability limits are commonly shown on a risk triangle, see Figure 1 below. This risk triangle is used rather than a traditional risk matrix as individual risk only has one severity level, the severity of fatality to the individual in question. A traditional corporate risk matrix, or any graph of severity versus frequency, cannot therefore be calibrated against individual risk criteria.

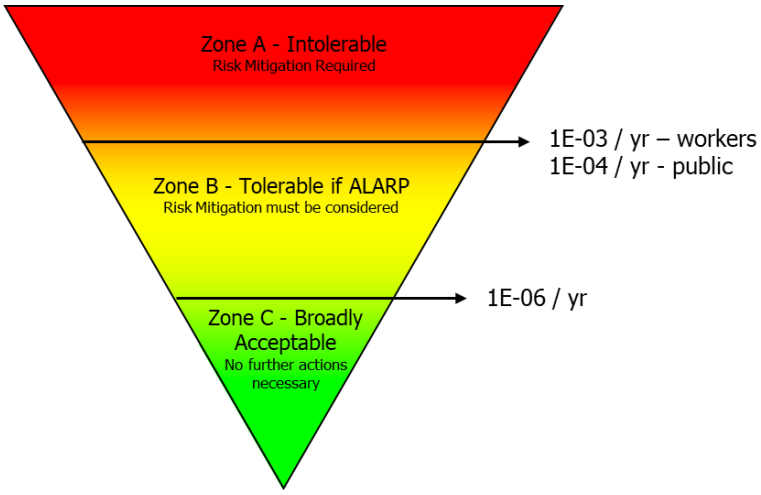

Figure 1 Individual Risk Triangle

Individual risk does not calculate the total risk of events which can potentially result in multiple fatalities, or a single fatality within a group of people. To properly calculate the risk of these scenarios, societal risk (and criteria) must be used.

4.2 Societal Risk

Societal risk is calculated based on the 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 is also referred to as Aggregated Risk (AR), group risk, or Potential loss of life (PLL). To illustrate societal risk for an establishment, an F-N curve is often used. This is a plot of the cumulative frequency (F) of major accident scenarios against the number of people that could potentially be affected (N).

Societal risk can be considered as a subset of societal concern. Societal concern includes societal risk, but also takes into account the public reaction to that event. These aspects are less tangible than the actual calculated societal risk of an event, and can make it difficult to gauge what tolerable levels of societal risk are. This is where disproportion factors and risk aversion can be incorporated, for example differentiating between risk to a school versus a warehouse, or car crash versus train crash.

As discussed in the background to this paper, tolerability criteria for societal risk have been the topic of much discussion and study, and are likely to be investigated much further in the near future. However, the latest accepted guidance on this is still from R2P2, which is therefore used in this review.

R2P2 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). Although only a single point for societal risk tolerability is given in R2P2 which does not allow for criteria lines to be defined, HSE's 'Guidance on ALARP Decisions in COMAH' (SPC/Permissioning/37) defines the criteria that should be used when determining the tolerability of societal risk in risk assessments.

SPC/Permissioning/37 is a semi-permanent circular operational guidance published in 2012 by the HSE, it draws on R2P2 and the work of Ball and Floyd (1998) to set the societal risk criteria. This sets an F-N curve with a slope of -1 (as confirmed by [the HSEs 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 boundary. The resulting F-N graph with acceptable and uncomfortably high frequency limits for a range of fatalities (1 - 1000) is provided in Figure 2 below. By combining this with an establishments' calculated F-N curve, the tolerability of the risk can be determined.

It is worth pointing out the change in terminology between R2P2 and SPC/Permissioning/37, the former uses the term 'Intolerable' whereas the latter uses the term 'Uncomfortably High' to described risks over the Tolerable region. It is not explicitly stated why this change occurs, however it is likely due to HSE policy being to not enforce risk reduction for uncomfortably high societal risk. For the purposes of this review however, the two terms can be considered one and the same, as industry experience seems to show that they are considered so in practice.

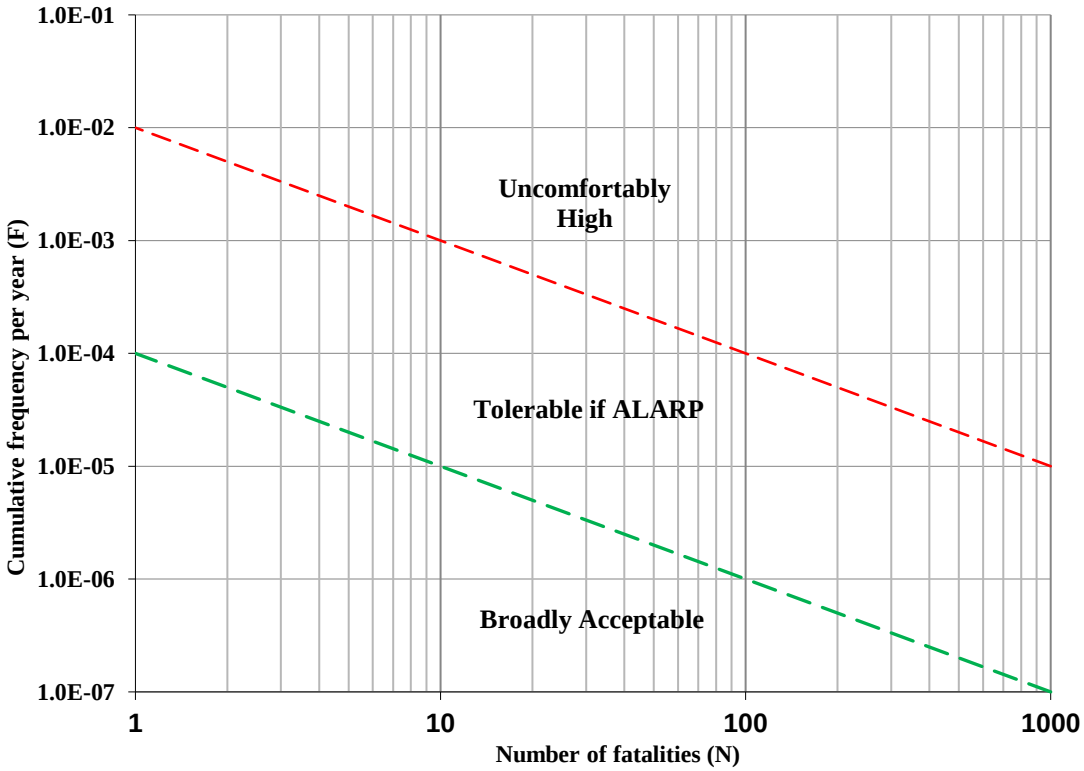

Figure 2 Societal risk tolerability F-N graph (log-log scale)

4.3 Which to use when

Based on the information within this paper, it is considered well established that for major accident hazard sites, societal risk is likely to be the most appropriate and the most stringent criteria to meet. It is also likely to be the easier of the two to properly calculate, and the more useful in making risk reduction decisions i.e. use in cost benefit analysis (CBA) calculations.

However, individual risks must still be considered, as when risks from all aspects of an individual's job are combined the total individual risk may not be tolerable (even if the societal risk is acceptable).

Some extracts from literature are provided in the box below in support of these statements. Despite these sources all agreeing that societal risks are the more appropriate measure for high hazard sites, the only definite set criteria are for individual risk (see Section 2.2).

R2P2

"hazards that give rise to such levels of individual risks also give rise to societal concerns and the latter often play a far greater role in deciding whether a risk is unacceptable or not"

PSLG

"For high-hazard sites, Societal Risks/Concerns are normally much more relevant than Individual Risks, but Individual Risk must still be addressed"

RR716

"…societal risk as well as individual risk should be taken into account, with the more onerous of the two risk targets being applied in any SIL calculations. Consideration of societal risk if found relevant, may often require more stringent safety measures to be applied."

RR703

"For a work activity capable of causing accidents affecting many people, however, the harm would not be limited to an individual. From this it may be argued that the HSW Act itself requires dutyholders to take account of societal risk. Individual risk may aid the design of an industrial plant but it is not necessarily sufficient to demonstrate compliance with Health & Safety Law."

Risk assessments at high hazard sites should therefore include both societal and individual risk calculations, this is most commonly addressed by calculating the societal risk of the establishment and comparing that to societal risk criteria plotted on an F-N curve or a risk matrix. The individual risk of the most exposed person or persons on the establishment is then calculated and compared to IR criteria. Only when both types of risk are demonstrated to be tolerable and ALARP can the duty of the operator be considered met. Once individual risk has been demonstrated to be tolerable, it is not normally used any further and societal risk is the main driver for risk reduction considerations i.e. it is used in CBA.

It is our view that calculating societal risk and including a check of individual is an appropriate technique. In terms of LOPA assessment, the approach adopted by RAS Ltd is to firstly apply societal risk targets to LOPA, and then include a check that the individual risk to the most exposed person does not exceed tolerable levels. This can often be done qualitatively by the LOPA team, by assessing if a specific person or worker group are disproportionately affected by this scenario.

The main aspect to include in any assessment of risk, is to fully explain and justify the type or types of risk that have been assessed, and how the tolerability criteria have been set.

Some examples of where each type of risk is likely to be more appropriate i.e. have low individual risk but high societal risk or vice versa:

Societal risk:

- Chemical reactor overpressure on a highly occupied site
- Gasoline tank overfill scenario on a site within an industrial estate

Individual risk:

- Road tanker overfill scenario where the same operator is exposed for every operation

5. Setting Risk Targets

When conducting risk assessments, we are not always concerned with calculating the total risk from a site, some assessments are required to make informed decisions regarding smaller pieces of the sites' risk picture. Therefore, a risk target must be set which can be used for major accident scenarios, which ensures that no one scenario takes up too big a portion of the total site risk.

A risk target is therefore a defined maximum level of risk that a company is willing to accept from a specific event. This is often used in LOPA to define the maximum frequency that an event can be considered tolerable (dependent on the severity of that event). In LOPA assessments used to establish an integrity level for an instrumented system, this is often referred to as the target event frequency or the target mitigated event likelihood (TMEL). Based on this frequency calculation, it can be identified whether further risk reduction measures are required to protect against this scenario, this may be in the form of a

SIF or other risk reduction measure. The LOPA can be used to define the required level of integrity those further measures need to be, by showing the difference between the target event frequency and the current event frequency.

Companies usually set their own risk tolerability targets as is described within IEC 61511 (2016) / BS EN 61511 (2017), often these are based on the severity of an event. Sometimes this is set using a company's corporate risk matrix. However, these are often designed to be used for all types and levels of risks and risk assessments, which can make it confusing to apply in specific scenarios. End users of corporate risk matrices are also not always aware of which type of risk is being assessed or how these targets have been decided. Calibrating a risk matrix to be used for numerous types of studies, from basic hazard identification all the way to full QRA, can also be difficult. This can also be exacerbated by incorporating environmental, occupational safety, financial and commercial risks into one matrix.

Whilst a one size fits all matrix could potentially enable more consistent risk based decisions to be made within an organisation, it is our view is that risk criteria should be set depending on the type of assessment and type of hazard and that it is more favourable that each risk be assessed individually. Any accompanying risk matrix should be proportionate and fully suited to its intended use (Nicholls and Carroll, 2017).

The two main aspects to consider when setting risk targets are as follows:

1. Where to aim on the tolerability scale
2. How to account for other risks

5.1 Where to aim?

Regardless of the type of risk being assessed, the relative position on the risk scale for which to aim for a given scenario can be considered the same. The exact point on the risk scale to aim often differs company to company, and none of the regulations, literature or guidance appears to give exact criteria for this. Therefore, we have to refer back to R2P2 and make judgements based on the fundamental criteria which defines the three areas of risk; Intolerable, Tolerable if ALARP and Broadly Acceptable.

A review of a range of company's risk target guidance for use in LOPA has shown that some companies prefer to use the Broadly Acceptable line as a target, and some prefer to aim for a point within the Tolerable if ALARP region. There are benefits to both, and which method to use can be dependent on the specific situation.

By aiming for the Broadly Acceptable line, and demonstrating that the risk is beneath this line, it reduces the onus on demonstrating that a risk is ALARP, and also adds a level of conservatism to the risk assessment. This can reduce future costs (both financial and time) in demonstrating that an establishment's risk is ALARP. However, existing establishments with legacy systems may find it difficult to meet these targets, resulting in the need for expensive and unjustified instrumented systems or additional safeguards to mitigate the residual risk.

Aiming for a point within the Tolerable if ALARP region would reduce this onus on existing operations, the remaining risk reduction required to achieve the broadly acceptable frequency can then be provided as part of an ALARP demonstration. This can include both the independent protection layers (IPLs) from the LOPA and 'softer' protection measures which cannot be taken credit for in a LOPA assessment (e.g. control of work and training).

It can therefore be appropriate to use either method, provided that ALARP is demonstrated, it can be common for people to lose sight of the principles of ALARP when conducting LOPA (Nicholls, 2012). A general rule of thumb may be for existing establishments to aim for the middle of the Tolerable if ALARP region, and for new establishments or new projects to aim for the Broadly Acceptable line. The key point to include in all assessments is a transparent description of the method used and why, which will facilitate the demonstration of ALARP.

5.2 Accounting for Other Risk

In assessing the risk from one major accident scenario/event, i.e. that which is done in a LOPA, neither type of risk (individual/societal) can be fully assessed. Individual risk tolerability must include all risks that that person is exposed to, not just one scenario, and societal risk must be aggregated for the whole site. To assess the risk tolerability of one scenario therefore, the targets must be adapted to make an allowance for these other risks.

It is generally accepted across the industry that it is appropriate to reduce the risk target by one order of magnitude to account for these other risks. This is on the basis that for an individual, there may be ten similar major accidents on a site that could contribute to that individual's overall risk, and for societal risk, there may be ten other potential major accidents of similar scale at the establishment.

This rule of thumb is very generic and there are likely to be many instances where these are not appropriate figures. However, based on inaccuracies of risk assessments in general, and the levels of conservatism built into most types of risk assessments, this may be considered a good approximation and is widely applied and accepted.

Note that the concept of occupational risk taking a certain portion of the risk tolerability is also mentioned in some of the literature on this topic, but the majority of references and guidance documents do not discuss it. It is therefore not considered further in this review but could be a topic for further consideration.

5.3 A Proposed Societal Risk Target

Based discussions above, we have developed the following risk targets for LOPA which we suggest as appropriate. As in Section 5.1 above we have suggested different targets for new and existing establishments, the process is therefore shown for an existing establishment and a new establishment.

Existing Establishments

When considering an existing establishment, we are suggesting taking the starting point as the middle of the Tolerable if ALARP region. The 'middle' in this context is taken as the midpoint of the zone when plotted on a log-log scale i.e. one order of magnitude below the Uncomfortably High line, and one order of magnitude above the Acceptable line. This point is shown on the F-N curve (Figure 3) as Point A, for a number of values of N. This equates to a single fatality at 1E-03 per year. As noted above, a single fatality event is not the same as individual risk.

This target is then reduced by an order of magnitude to account for other scenarios on the establishment. This is shown by the points marked B in Figure 3 and for a single fatality equates to a frequency of 1E-04 per year. By drawing a line through all the points marked 'B', a tolerability line is set for which a target frequency (F) for any number of fatalities (N) can be derived for an existing establishment. This tolerability line is shown in grey on the figure below.

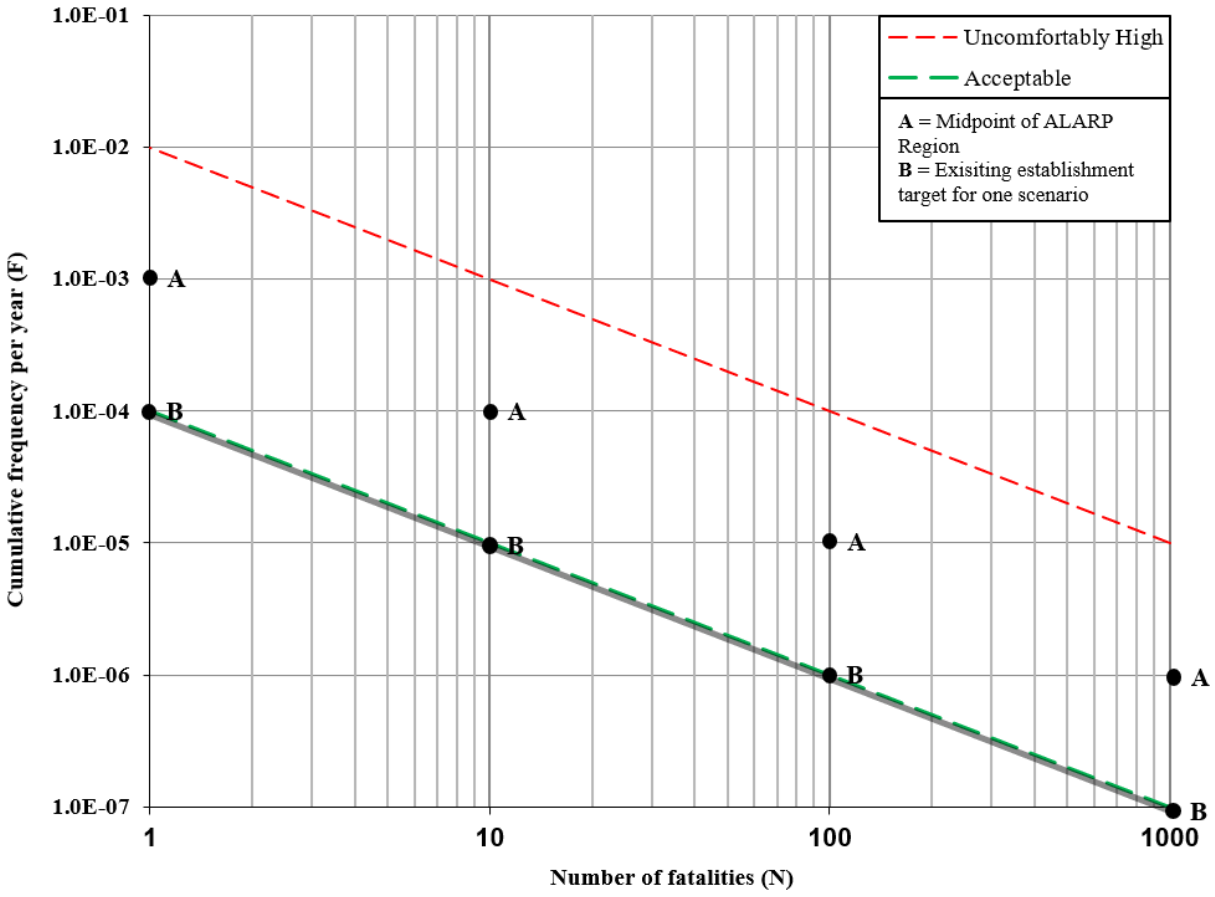

Figure 3 Societal risk tolerability lines – Existing Establishments

New Establishments

When considering a new establishment (or a new project on an existing establishment), we suggest taking the starting point as the Broadly Acceptable line. This point is shown on the F-N curve (Figure 4, below) as Point A, for a number of values of N. This equates to a single fatality at 1E-04 per year. This target is then reduced by an order of magnitude to account for other scenarios on the establishment. This is shown by the points marked B in Figure 4 and for a single fatality equates to a frequency of 1E-04 per year. By drawing a line through all the points marked 'B', a tolerability line is set for which a target frequency (F) for any number of fatalities (N) can be derived for a new establishment. This tolerability line is shown in grey on the figure below.

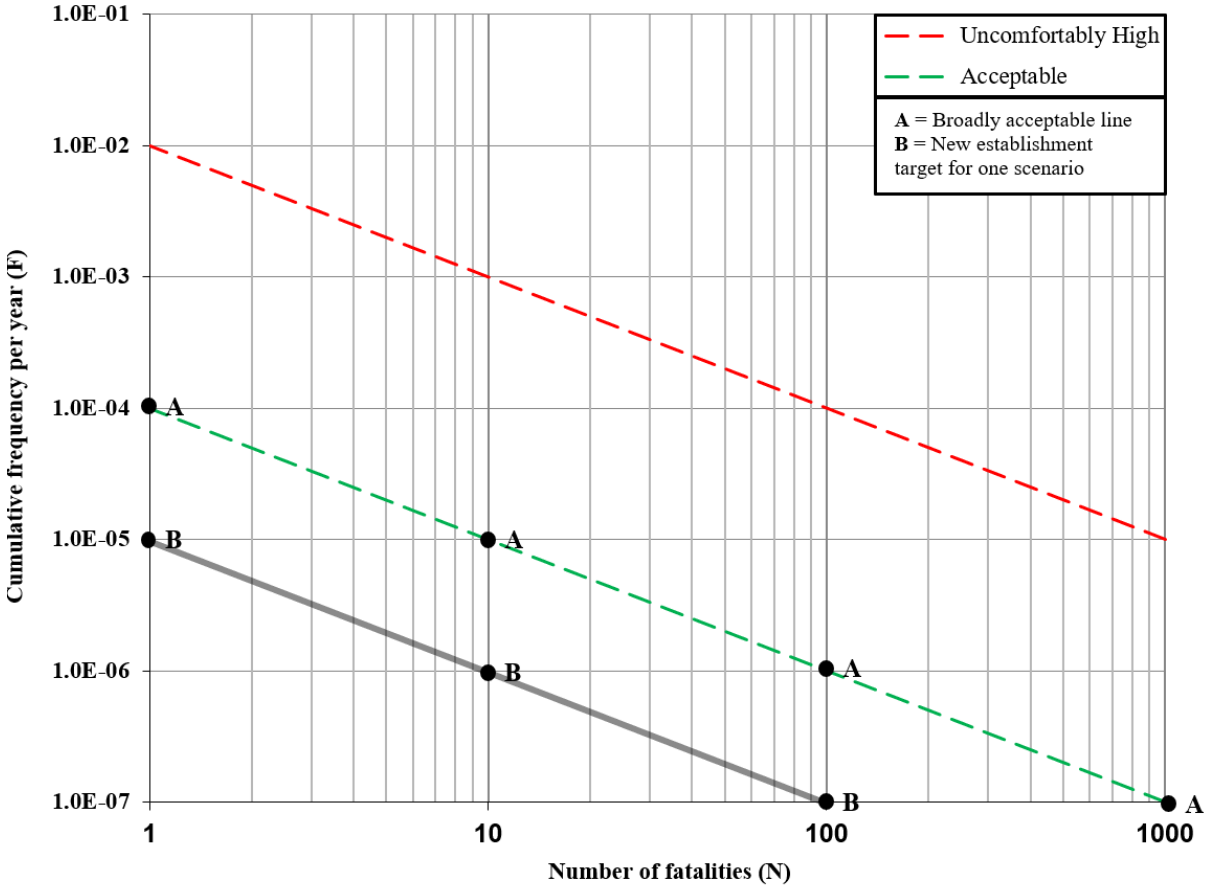

Figure 4 Societal risk tolerability lines – New Establishments

Some example frequencies read of these tolerability lines are given below in Table 1. The target frequencies show that for societal risk, as severity of an event increases, the target frequency is reduced proportionately.

Table 1 LOPA Target Frequencies (Societal Risk)

Severity / Harm	Target Frequency (per year)	
	Existing Establishments	New Establishments
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

As LOPA is generally an 'order of magnitude' tool, and the severity of an event can be hard to accurately estimate in a HAZID, it can be convenient (and conservative) to apply the targets for consequences in terms of orders of magnitude also, i.e. one fatality, >1 but ≤10 fatalities, >10 but ≤100 fatalities. RAS Ltd generally consider that LOPA is not the best tool for assessing events which are predicted to result in more than ten fatalities, these types of event are better assessed by more sophisticated quantitative techniques to demonstrate that risks have been satisfactorily controlled.

5.4 A Proposed Individual Risk Target

We suggest that individual risk targets are set using the same method as above for societal risks. A worked example is given below to show our proposed risk targets for individual risk in LOPA, we have used the tolerability for workers rather than public as this is the more likely scenario.

As an individual's risk should include the risk from all hazards to which that individual is exposed, in LOPA assessments the target frequency should be adjusted accordingly to account for other potential sources of risk to that individual. It is generally considered appropriate to lower the intolerable boundary by an order of magnitude (from 1E-03 per year to 1E-04 per year) to account for the other risks that an individual is exposed to.

As with societal risk, it is considered appropriate to set the target frequency for individual risk in the middle of the Tolerable if ALARP zone. The midpoint between 1E-04 per year (adjusted intolerable line) and 1E-06 per year (broadly acceptable) is 1E-05 per year.

As discussed above, some companies may prefer to aim for the broadly acceptable risk line to reduce the onus on demonstrating ALARP. For LOPA that would equate to a target frequency of 1E-06 per year for individual risk.

These targets for individual risk apply to all scenarios regardless of the potential total number of fatalities. The risk is to an individual therefore there is no risk aversion factor to consider.

The key consideration when assessing individual risk in LOPA is correctly calculating an individual's exposure to the hazard. This can be done manually for each individual (or the expected most exposed individual), using their shift pattern and other occupancy information to calculate a probability that that individual will be exposed to this event. This can then be applied within the LOPA as an occupancy factor for the specified individual.

A broader method is to use a 'sharing' factor which can be applied in the LOPA to account for the sharing of risk with other people (e.g. the other shift operators). If the scenario in question is expected to affect one person, but that one person can be any one of six shift operators, then a broad 1/6 factor can be applied to the LOPA frequency to calculate the average individual risk that each individual operator experiences from this scenario, the calculations would be as follows:

Societal Risk (per year) *= N (Number of fatalities) x F (Frequency of loss of 'N' lives)* ***(Eqn. 1)***

Individual Risk (per year) *= (N x F)/* P_{total} ***(Eqn. 2)***

Where P_{total} is the total persons who are exposed to this hazard and share the risk of the 'N' fatalities specified in the calculation.

This should be considered carefully for each event however, especially where more than one worker group is exposed to the event. For example, in a ten fatality event involving six operators and four maintenance technicians, at first glance the operators may seem like the most exposed worker group. However, there may be 30 operators (five shift teams) sharing this risk, compared to only eight total maintenance technicians. This would result in the maintenance technicians being the most exposed individuals and having a higher individual risk from this event than the operators. It is also worth noting that in this example the operators may have a higher total individual risk than the maintenance technicians, as they are exposed to more risks across the site. It is vital to check how sensitive your assessment is to any of the assumptions made within it, and clearly justify each to facilitate reviews and audits.

We suspect that these subtle complexities of individual risk are party responsible for the level of confusion and misapplication of criteria in LOPAs, and may be the best justification for societal risk being considered the more appropriate criteria to apply in LOPA. Individual risks are better assessed as part of a full quantified risk assessment, taking into account all scenarios that each worker group is exposed to. The only exception to this (as discussed elsewhere) is for events where an individual is quite clearly the main person exposed to a hazard.

6. Discussion

It is clear from legislation, regulatory guidance and industry guidance, that both the levels of risk to an individual and to society must both be shown to be tolerable and ALARP for the duties under health and safety law to be considered met. It is also clear from literature and industry experience, that risk criteria are not used consistently and are often inappropriate or incorrect.

Based on this review, it is obvious that regulators assessing these aspects have a difficult job, as there can be little to no consistency between companies and sites across the UK, and limited justification provided for what is used. This is a fundamental concept as part of a demonstration that risks are firstly acceptable, and secondly ALARP, yet it is not obvious that everyone is comparing apples to apples. How can regulators adequately assess these areas when two similar establishments, with similar hazards, use completely different risk targets (sometimes orders of magnitude apart). This can result in what appear to be the same or very similar processes, with similar risks, having different required levels of protection, with no clear reason why.

Difficulties with assessing these studies can also be exacerbated by general LOPA methodology errors. Although LOPA appears to be a simple methodology on the surface, it requires a trained hand to apply properly. We often find simple errors (e.g. incorrect use of IPLs) are made which can result in a large a difference to the final result. Where these errors in applying LOPA technique are compounded with incorrect risk targets, we have seen LOPAs for low severity events resulting in an incorrect risk gap equating to a requirement for a SIL 3 solution. However, we have also found some lucky escapes where multiple errors have cancelled each other out within the assessment, and after correcting all errors within the LOPA the result of the assessment was relatively similar. Based on these points there may be many LOPA assessments within the industry that are incorrect, resulting in many under-protected processes, or many unnecessary risk reductions implemented and maintained (at significant cost).

We have provided a methodology for assigning targets for use in LOPA, which we deem to be appropriate based on the latest legislation and guidance. The principles behind this method can also be adapted for other assessments as required by the end user, provided the end user has a good understanding of the fundamentals of risk as described within this review. We believe that the idea of aiming for the middle of the Tolerable if ALARP region of risk is a proportionate method for existing establishments, and whilst not a new concept, is not widely applied.

The key recommendation from this review is that all assessments of risk should adequately describe the criteria that are being used, and from what it has been derived. This will aid in assessments, audits and future updates of studies. It is also important to refrain from applying methods that we have seen others use, without fully understanding if the same method applies to the situation in question. When using criteria set by a company, the user should take care to understand where this information

came from and how it has been adapted for use by them. Only once this is fully understood can you be confident that you are applying the correct risk criteria for your given situation.

This review paper has explained the fundamental basis behind the methods we have presented, to aid in both the use of our methodology, and for individuals to have a deeper understanding of methodologies they may already be using. By understanding the background behind risk tolerability criteria, we can make better judgements about the type of risk criteria we should be using, when we should be using it, and how to appropriately apply it.

7. References

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 (2012), Risk tolerance criteria for layers of protection analysis, American Institute of Chemical Engineers Process Saf Prog 31: 118–121

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

BS EN 61511 (2017), Functional safety – Safety instrumented systems for the process industry sector.

Colin Chambers, Jill Wilday and Shane Turner (Health and Safety Laboratory)(2009), RR716 A review of Layers of Protection Analysis (LOPA) analyses of overfill of fuel storage tanks, Available at http://www.hse.gov.uk/research/rrpdf/rr716.pdf [Accessed 12 January 2018]

Chemical Industries Association (CIA)(2010), Guidance for the location and design of occupied buildings on chemical manufacturing sites.

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

Health and Safety Laboratory (HSL) and Health and Safety Executive (HSE) (2009), RR703 Societal Risk: Initial briefing to Societal Risk Technical Advisory Group

HSE (1992), The Tolerability of Risk from Nuclear Power Stations, Office of Public Sector Information.

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 12 January 2018].

HSE (2016), Predictive Aspects of Safety Report Assessment, [online] Available from http://www.hse.gov.uk/comah/sram/docs/s10.pdf [Accessed 12 January 2018].

IEC 61511 (2016), Functional safety – Safety instrumented systems for the process industry sector.

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.

Nicholls (RAS Ltd, UK), 2012, Integration of ALARP into project management processes – how low is low enough? 23rd Institution of Chemical Engineers Symposium on Hazards, Symposium Series 158, Paper 52

Process Safety Leadership Group (PSLG) (2009), Safety and environmental standards for fuel storage sites, Crown Copyright, ISBN 978 0 7176 6386 6

The Management of Health and Safety at Work Regulations 1999 No. 3242

The Control of Major Accident Hazards Regulations, 2015 No. 483