Is there value in a 'one size fits all' approach to risk matrices?

Carolyn Nicholls and Jennifer Carroll, Risk and Hazard Management Consultants, RAS Ltd, 73 New Crane Street, Chester, CH1 4JE

Risk matrices are one of the most favoured and established tools within the fields of risk assessment, analysis and evaluation. They are applied widely within industry, evidenced by several papers documenting them as the principal risk management tool. This is because risk, being a two dimensional concept, lends itself well to the matrix structure and a matrix itself appears simple and transparent in design. However, despite a general acceptance, their use is still disputed amongst professionals.

This paper looks to explore the risk matrix concept and provides several discussion points relating to the use of risk matrices with the ultimate aim of addressing the following question "is there value in a one size fits all approach to matrices?" A 'one size fits all' approach is perhaps an ideal solution as it would, at first glance, enable consistent risk-based comparisons and decisions to be made. In answering this question, efforts are focused on discussing the potential pitfalls of a 'one size fits all' approach, as well as alluding to the general consensus on the positives and negatives of matrices in general.

The main outcome of the review is that the intended use of the matrix is a pivotal factor in defining its content and structure. One of the key pitfalls in trying to achieve a common approach lies in the granularity of the matrix, which in turn depends on how broad or specific the matrix needs to be and ultimately its use.

The complexity of the issues surrounding risk matrices and the ensuing controversy is evident. Regardless of whether they are used for communicating levels of risk or assisting with risk ranking and prioritisation, careful consideration is needed as to whether a universal matrix is a practical solution or whether, in reality, it is better for risk to be assessed for each individual case and for the nature and complexity of the accompanying presentation tool to be proportionate to its intended use.

Keywords: risk, assessment, matrix, matrices, presentation

Introduction

Risk matrices are one of the most favoured and established tools within the fields of risk assessment, analysis and evaluation. They are applied widely, evidenced by several papers documenting them as a principal risk management tool (Center for Chemical Process Safety/American Institute of Chemical Engineers, 2008) (Ellis, 2003). Risk is essentially a two-dimensional concept, characterised by specific events and the level of harm and likelihood of occurrence these events hold. It can often be a difficult concept to visualise but by being two-dimensional lends itself well to being displayed in the form of a matrix.

Whilst simple in appearance, the judgement and ideas behind risk matrices can be complex and their application can present several challenges. As a risk management consultancy, we have recently had numerous discussions with clients surrounding the use of risk matrices, and from exploring what has been written on the subject, realise that we are not alone in debating their use. Cox (2008) amongst others, demonstrates limitations of risk matrices with implications for risk management. This paper looks to explore the risk matrix concept and provides several discussion points relating to the use of risk matrices, with the ultimate aim being to address the following question, "is there value in a one size fits all approach to risk matrices?" On the surface, this idea does not seem unwise and people may be tempted to use a multi-purpose matrix for ease above other reasons; after all, risk is risk, right?

In answering this question, a particular thread will be followed, beginning with a brief background to risk assessments as the foundations on which matrices are based and concluding whether there is feasibility in a one size fits all approach. Along the thread, the thought processes behind the use of matrices will be looked at, followed by a consideration of risk matrices along the entire spectrum of risk assessments ranging from qualitative to quantitative. Several other interesting points relating to the use of matrices will be explored, in particular the structure of risk matrices and the calibration of them in relation to tolerability criteria.

Background to Hazard, Risk and Risk Assessment

Risk assessments are now a routine aspect within most jobs, ranging from dynamic on the spot assessments to semi-quantified and quantified risk assessments such as those needed for Control of Major Accident Hazards (COMAH) risk assessments. But do we really appreciate how useful they can be or are they just a tick-box exercise?

Hazard and risk are used interchangeably in everyday vocabulary, but they are two different notions. Distinguishing between the two is important, especially for the purpose of a risk assessment. A hazard is something that has potential to cause harm, whilst risk is the chance or likelihood that someone or something that holds value is adversely affected by a hazard (HSL, 2003). Risk, as a two-dimensional concept, is not fully understood or appreciated until both of these aspects have been considered.

Ultimately, a risk assessment is about understanding where potential hazards lie and deciding whether the risk is acceptable (Pitblado and Turney, 1996) or whether more needs to be done to manage the risk to a level that is 'as low as reasonably practicable' (ALARP). Within the COMAH Regulations (2015), Regulation 4 requires operators to "take all measures necessary to prevent major accidents", which is equivalent to the ALARP principle (HSE, 2010) and the So Far as is Reasonably Practicable (SFAIRP) principle is central to several other regulations, including the Management of Health and

Safety at Work Act Regulations (1999). Each hazard or event being assigned a level of risk also enables prioritisations to be made over where resources should be directed and focus be held.

Events with a high consequence and high likelihood (and sometimes low likelihood) are of major concern whilst those of low severity that are less likely to occur are of little concern and not particularly subject to risk reduction efforts. It is the intermediate group which creates the most interest and needs further consideration, especially in relation to the concept of ALARP. For risks that fall into this category, it has to be demonstrated that the risk has been reduced to a level deemed tolerable and that there are no additional control measures which would be grossly disproportional in cost versus the risk reduction benefits they could provide (HSE, 2001). Incorporating ALARP also prevents risk assessments from being a yes or no decision making exercise. The topic of ALARP was something we explored in 2012.

Thought processes behind the use of risk matrices

Risk matrices are so widely used and appealing that in some circumstances they may be selected as a tool within a risk assessment without being given much thought. It is extremely important to stop and think about what the purpose and the objectives of the risk assessment is before designing and implementing a matrix, adopting one from elsewhere, or even using one at all. Without a thorough prior understanding of the purpose and objectives, a matrix can be incorrectly designed, leading to incorrect implementation and subsequently unavailing results being taken forward for evaluation. It is worth spending time front-end loading when opting to use a matrix in order to produce the most beneficial results. The Health and Safety Laboratory (HSL) (2003) advises answering the following questions before conducting a risk assessment; risk 'of what', 'to what (whom)' and 'from what'. These questions can also assist decisions relating to risk matrices.

Risk matrices have two predominant applications and overriding objectives (Dujim, 2015):

- Risk ranking and the prioritisation of where resources should be directed
- Decision making about the tolerability of risk

These objectives should govern the thought processes behind the final product of the matrix and unless a matrix answers the questions you are asking from the risk assessment and is informative, it is effectively of no use.

It is important to note here that matrices themselves are not a tool for decision making, they just provide visual assistance, whether by highlighting priority areas or displaying levels of risk and where within the tolerability regions that the risk is concentrated. They are ultimately communication tools, thus the process of developing and using a matrix is the most important aspect.

Given the level of thought that is needed to produce an effective risk matrix, is there any value in a one size fits all approach? Being able to communicate several aspects of risk on one matrix, or even using one matrix for several purposes within an organisation would be beneficial and perhaps save a lot of work. Organisations may deliberate about using just one type of matrix on the basis that it could be simpler to compare risks that are displayed uniformly. However, a 'one size fits all' approach to risk matrices is risky in itself because the use of a matrix is tailored to and governed by the particular parameters of a risk assessment. The feasibility of this idea is open to question.

The spectrum of risk assessments

Risk matrices are commonly found in both qualitative and semi-quantitative risk assessments but with differing styles. They offer benefits across the entire spectrum of risk assessments ranging from qualitative risk assessments through to quantitative risk assessments (QRA) despite a general belief that they are not quantifiable structures.

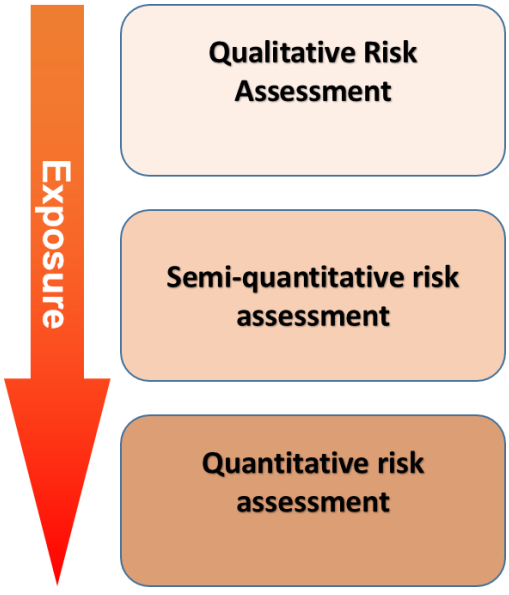

Figure 1 The Spectrum of Risk

At the lower end of the spectrum, a qualitative risk assessment uses a matrix where the likelihood and consequence categories are descriptive. Risk matrices here can be useful to derive a structure and consistency within the risk assessment. People are likely to find it easier to assess the risk related to something if they have categories to choose from as opposed to selecting from an open pool. With regards to consistency, people are more likely to select similar severity and frequency categories and thus levels of risk if they have a framework to use. There are however several papers debating the use of matrices, focusing on discrete categories providing little benefit. Dujim (2015) argues that qualitative risk assessments using discrete categories inject too much subjectivity whilst Prince (2005) proposes that labels such as 'frequent' and 'severe' carry a level of presumed meaning because of their typical use in everyday language. Borysiewicz (2006) states that when using discrete categories, they should be explicitly defined in the hope of causing as little deliberation over which category to select as possible.

Within the process safety industry, qualitative risk matrices are a common attribute of Hazard and Operability (HAZOP) studies (HSE, 2010) and are especially useful for risk ranking (Crawley and Tyler, 2015). HAZOP studies predominantly serve to assess the risks associated with consequences resulting from particular individual parameters through the use of particular guidewords but they can also attempt to identify potential operating problems (Crawley and Tyler, 2015). The severity categories can relate to the potential for different numbers of fatalities whilst the frequency categories can relate to whether the event has been seen within the site itself, or within industry as a whole.

Selecting the appropriate category for severity is often a fairly simple task, despite being based on subjective but educated estimations. However, when it comes to frequency, some can find it more difficult to select a category, potentially due to the frequency of an event being out of our comfortable realm when it comes to decision making, resulting in lengthy HAZOP sessions. If it is agreed that hazard identification (HAZID) is the key step of the risk assessment process, (unless all risks are identified, the biggest hazard could have been missed), is it worth spending so much time on a HAZOP deliberating over something which ultimately is no more than an educated guess? How many people have been sat in a hazard study arguing over categories for extended periods of time; is this good use of time?Overall, the level of assessment using a qualitative risk matrix is coarse, the matrix has little risk ranking capability and calculating cumulative risk is difficult (Worksafe Victoria, 2011). However, for a small, simple facility dealing with less severe hazards and with low exposure of the public or workforce to these hazards, this matrix style could be ideal and would serve its purpose well.

Moving higher up the spectrum of risk, an organisation may need better estimates of risk from hazards at the facility and judgement alone may not be proportionate. In this case, a semi-quantitative risk assessment may be more appropriate, using a matrix displaying numerical values. Semi-quantitative risk matrices are a common feature within COMAH reports (HSE, 2010) and also a popular tool for making risk decisions within Layer of Protection Analysis (LOPA). LOPA studies are often the next step after a HAZOP, and semi-quantitatively assess the frequency aspect in more detail, bringing up the discussion point as to whether both severity and frequency are needed at an early stage. It is also often the case in a HAZOP study that determining which events to take forward for rigorous frequency analysis via LOPA is carried out based on the severity categories selected (Summers and Hearn, 2012). The severity of an event occupies an entire row of the matrix, regardless of the frequency selected. Prince (2005) however states that a one-dimensional matrix based on severity only provides little risk ranking opportunity. This is because with regards to risk reduction measures, it is "obvious" that priority lies with high likelihood events.

In risk matrices used for LOPA, each cell is associated with a particular degree of risk reduction for a scenario which falls into that cell (Center for Chemical Process Safety, 2001) based on its severity and frequency categories. In some circumstances, the level of risk reduction inferred from a matrix can be associated with a particular number of Independent Protection Layer (IPL) credits required in order to satisfy the tolerability levels (Center for Chemical Process Safety, 2001). Risk matrices can be well suited to LOPA studies as they are both acknowledged as being not so precise and based on simplistic and often conservative assumptions (Center for Chemical Process Safety, 2001).

Refining events before taking them to LOPA can however create problems when it comes to estimating levels of risk. Being conservative and potentially negative when it comes to assessing events where levels of judgement and uncertainty are high can result in events being carried forward to LOPA unnecessarily. For example, an intolerable risk event carried forward may, after a more accurate assessment, prove to be a broadly acceptable risk event, whereas other events which were previously deemed tolerable if ALARP and not progressed further now appear to be of higher risk than those focused on in LOPA. In cases like this, the focus has been misdirected and the outcomes of the risk ranking process generated within the HAZOP and LOPA studies when recombined are inaccurate.

At the top end of the spectrum, within a fully quantified assessment for example, it is often common for individual or even societal risk to be calculated. These indicators of risk are high level and to be carried out by experts, but if the picture of risk within an organisation is to be communicated to people who may have less expertise in the area, a risk matrix can be a useful communication tool and help to focus on the risks and depict where they lie. Matrices enable a visual representation of what events sit where and their relativity to other events. They are useful to risk managers and provide more meaning than just a number produced within the calculations.

Within the risk assessment process, as discussed earlier in the paper, it all boils down to what the objective of the risk assessment is. Is a HAZOP needed to understand the hazards, the risk, or both? The answer to this question changes depending on the position within the plant lifecycle, whether pre-construction, during operation, or after a particular incident for example. Similarly, the most suitable matrix to be used depends on the position within the risk assessment. Qualitative risk matrices may be suitable at the outset where refinement is low, but as the need for refinement and quantification increases, these matrices will be redundant, highlighting the point that one matrix cannot be used at all stages. The image of a funnel is a good analogy of what is meant here in terms of levels of refinement. The top of a funnel is the widest point,

indicating a broad, unrefined level of assessment. The width decreases towards the bottom of the funnel, indicating a more refined, detailed level of assessment. The challenge lies in bringing each stage back together at the end in order to present a simple but whole site risk picture.

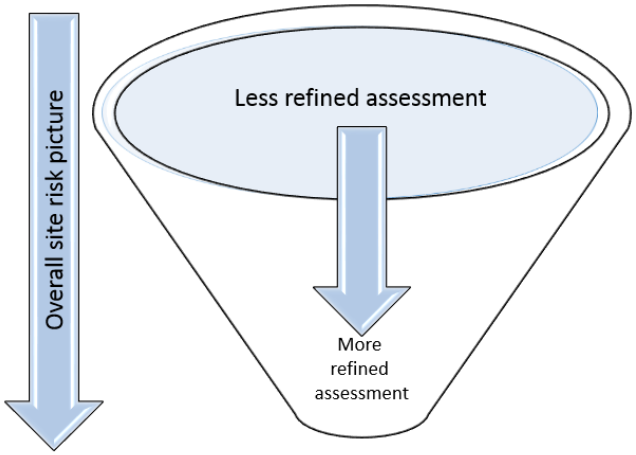

Figure 2 The funnel analogy

The structure of risk matrices

A common structure for matrices is a 5 x 5 square, but depending on the use of the matrix, this may be too broad or too specific; the level of detail within a matrix is known as granularity (Talbot,2011). Using a matrix with an increased number of categories has the potential to give an impression of accuracy, which within this context is often false (HSE, 2010). Using a matrix which is too broad loses the detail and can often result in range compression (Cox, 2009) or risk ties (Ni et al., 2010), where many of the events being assessed fall into the same risk region, or cell, when in fact, their consequences may differ greatly. If, for example, all hazards are taken out of the context of the matrix and just listed with their risk level, risk ties are not useful when it comes to prioritisation and risk ranking, especially within the Tolerable if ALARP region. Within the Broadly Acceptable region this is less important as there is no need for measures to be implemented, and within the Intolerable region it is less important as measures **have** to be implemented to lower the risk.

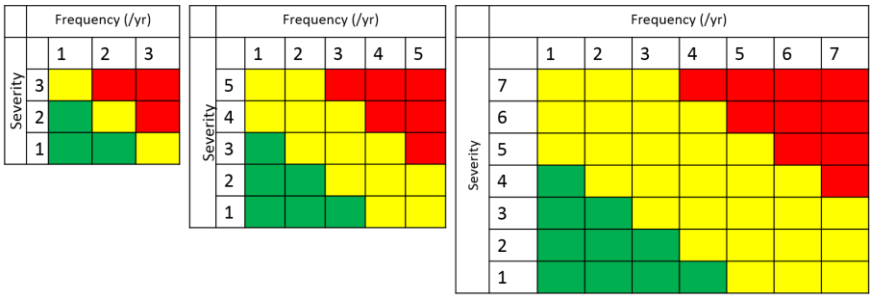

Figure 3 The granularity of risk matrices

If the sole purpose of the matrix is to aid prioritisation, there are options which can remove the likelihood of risk ties. The first is to select a larger number of severity or frequency categories in order to obtain a sufficient resolution of detail; a larger number of categories is also often preferable with matrices in general (HSE, 2010), but ultimately depends on the level of risk assessment. In a subjective or qualitative risk assessment, too many options may not always be practical.

Further options to prevent risk ties that can result from discrete categories is to use risk indexing, whereby risk is calculated based on scores assigned to several factors within each severity and frequency category. Risk is often presented as severity multiplied by frequency, based on its two-dimensional structure, but sometimes it can be more than that. Rather than simply placing an event on the matrix within a general risk region, indexing offers a more precise point. Both severity and frequency categories are broken down into several components, each rated and summed to give a score. Indexing however, can be seen by some as too complicated, because of the multiple components involved. Ultimately, the calculations are carried out in the background and at face value, the matrix has to be no more complicated. If anything, these options are preferable if the matrix is to be used for risk ranking and prioritisation, due to removing the issue of risk ties. Both enable each event to have a precise point within the risk region, enabling easier prioritisation decisions to prevail if the matrix is to be used in that way.

With regards to risk indexing, it can also be useful in that behind the scenes, if necessary, the data can be analysed further to see what aspect of each is the driver behind the location on the matrix where priorities have been highlighted.

Risk tolerability and the calibration of matrices

The Health and Safety Executive (HSE) tolerability framework, presented in their document Reducing Risks, Protecting People (R2P2) stipulates three regions of risk that can be applied to all hazards: broadly acceptable, tolerable and unacceptable. Levels of tolerable risk for individuals are presented for both the general public and the workforce of a facility. For individual risk of fatalities, a tolerable level lies anywhere between 1 in 1 million fatalities per year and 1 in 10,000 fatalities per year for members of the public, with an increased value for workers to 1 in 1000 per year (HSE, 2001). Above this, risk is unacceptable. The HSE framework is conceptual, and the factors, processes and thoughts which rule what is deemed unacceptable or tolerable are dynamic and depend largely on the hazard.

With regards to environmental harm or harm to business, these topics are ones that people are largely detached from and perhaps less sensitive to, often resulting in a greater tolerance. The Chemical and Downstream Oil Industries Forum (CDOIF) guidelines 'Environmental Risk Tolerability for COMAH Establishments' present tolerability criteria for harm to environmental receptors, to be used within environmental risk assessments. The concept of tolerability is a difficult one and in reality, nobody wants to tolerate any level of harm, or decide who or what should have different levels of tolerability, but for risk assessments to take place, a baseline is necessary. The HSE model in practice ensures that the risk from any hazard is controlled to a degree whereby the residual risk either falls into the broadly acceptable region, or is as close to the bottom of the tolerable region as practicable (HSE, 2001).

For matrices to be used to depict the overall outlook of risk tolerability, they have to be calibrated appropriately to align with the agreed tolerability thresholds. Calibration of matrices is a tricky art and the tolerability of an event often depends on its size. Different types of hazards potentially have different thresholds thus it is not advisable to try to plot multiple types on one matrix, nor should the risk to different groups of people be presented on one matrix for the same reason. From a process safety background it is also questionable whether it is feasible to compare scenarios of differing nature on one matrix. The risk associated with a release from a reactor could be different to a release from a standard process vessel, and different tolerability levels could be defined.

It is also important to ensure risk is being plotted onto a matrix suitably calibrated for the type of risk being considered. Equipment risk, scenario risk and site risk are all different, and the tolerability levels for each will differ, thus each should have its own individually calibrated matrix; a matrix calibrated for scenario risk cannot be used to present overall site risk. It is also important to ensure that individual risk criteria from the HSE is not translated onto a risk matrix without thought being given to what the matrix is displaying; often matrices which show individual events have to be made cumulative in order to match the HSE criteria (HSE, 2010). The HSE provide tolerability for a single point but for a company to manage risk properly, they need to understand the wider risk picture. When trying to understand this picture through the use of a matrix, the matrix is effectively adopting the position of an FN curve presenting cumulative risk. The calibration of risk matrices is tricky largely due to them incorporating so many corners as opposed to taking the form of a definitive line, where each side has a different level of risk. It can also be said that an FN curve is an extremely refined risk matrix whereby so many pixels have been introduced that the corners have been smoothed out and a curve is formed.

Baybutt (2015) offers an extensive discussion on the calibration of risk matrices and the complications that calibration presents. He agrees that pitfalls in calibration relate to the nature of risk tolerance criteria and the events to which they are applied; he concurs that one matrix cannot be calibrated for multiple purposes. Baybutt (2013) also presents a paper on the allocation of risk tolerance criteria where he delves further into methods calculating scenario risk versus overall site risk.

Risk aversion

Risk aversion is the concept of the extent to which an organisation has appetite for or wishes to disfavour risk (ISO, 2010). With regards to a risk matrix, aversion generally results in a larger portion of the matrix making up the Intolerable region, whereby high severity but low frequency events are given greater consideration than would otherwise be given to an event believed to have an extremely remote chance of occurring. Of course, this largely depends on the nature of the consequence and in some circumstances how political the subject matter may be (Ale et al., 2012).

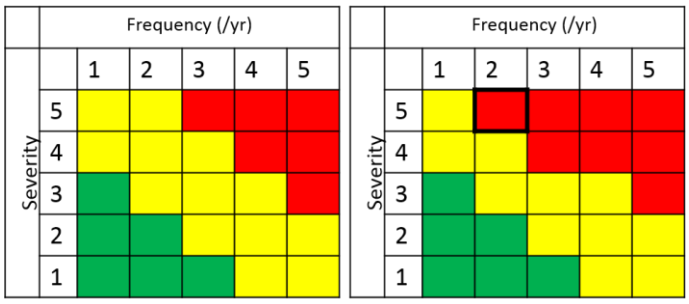

Figure 4 Building risk aversion into a matrix

There is currently little guidance on risk aversion. With regards to the process safety industry, a key question is whether building aversion into a matrix will result in over-engineering of a facility. Risks lying in the intolerable region of the matrix are subject to risk reduction efforts to lower the risk, regardless of any cost etc. Being extremely conservative by building in

risk aversion is likely to result in the risk of more hazards being rated as intolerable and more measures having to be implemented, for the sake of the company's risk attitude and aversion to high severity events. There is no question that organisations should strive for continuous risk reduction, but resources are not infinite.

An organisation may have different levels of aversion determined by the topics being assessed. Consequently, a matrix which indicates a risk averse attitude may not concur with a topic where levels of aversion are low, meaning that each topic should have its own risk matrix. This may be a complication with little value.

Conclusion

The purpose of this paper was to discuss the use of risk matrices with the aim of deciding whether there is any feasibility in a one size fits all approach. The complexity of the issues surrounding risk matrices and the ensuing controversy is evident. The arguments presented here lead to the conclusion that there are several factors which mean that a universal approach is an unlikely outcome, and it will prove difficult to apply a common approach across a range of circumstances (ISO, 2010). The underlying factor in all instances is the use of a risk matrix and its objectives. The use of a risk matrix ultimately governs its structure, the way categories are defined, how it is calibrated and whether aversion is built into the matrix or not. The use of risk matrices has also sparked interesting discussion points relating to HAZOP and LOPA studies, reinforcing that the utilisation of matrices is no simple task despite their simple appearance. All of these factors need careful consideration in order to produce the correct outcomes from the risk assessment within which the risk matrix is found. Whilst Cox (2008) has criticised the use of matrices, if used correctly they can provide benefits across the entire spectrum of risk. The key is to spend time getting the logistics correct; getting them wrong could mean applying unnecessary control measures or failing to apply important ones. Whilst a one size fits all matrix could potentially enable more consistent risk based decisions to be made within an organisation, it is more favourable that each risk be assessed individually and the nature and complexity of the accompanying risk matrix should be proportionate and fully suited to its intended use. It is essential that risk matrices are given more thought and not selected for use within risk assessments without care, just because of their wide appeal.

References

Ale, B., Burnap, P and Slater, D., 2012, *Risk Matrix Basics*, UK [Online]. Available at: https://www.researchgate.net/publication/305889949_Risk_Matrix_Basics (Working Paper)

Baybutt, P., 2013, Allocation of Risk Tolerance Criteria, *Process Safety Progress*, 33: 227-230

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

Borysiewicz, M.J., 2006, *Models and techniques for health and environmental hazard assessment and management*, Poland: Institute of Atomic Energy, 378-412

CDOIF, 2015, *Environmental Risk Tolerability for COMAH Establishments*, UK: HSE

Center for Chemical Process Safety, 2001, *Layer of Protection Analysis, Simplified Process Risk Assessment*, New York, 1-269

Center for Chemical Process Safety, 2008, *Guidelines for Hazard Evaluation Procedures*, 3rd Edition, New York, 1-576

COMAH, 2015, *Control of Major Accident Hazards (COMAH) Regulations 2015*, UK: HSE

Cox, L.A., 2008, 'What's wrong with risk matrices?'*Risk Analysis*, 28: 497-512

Cox, L. A., 2009, *Risk Analysis of Complex and Uncertain Systems*, New York: Springer

Crawley, F., and Tyler, B., 2015, *HAZOP Guide to Best Practice*, 3rd Edition, USA: Elsevier

Dujim, N.J., 2015, Recommendations on the use and design of risk matrices, *Safety Science*, 76: 21-31

Ellis, G., 2003, The Ongoing Challenge of Demonstrating ALARP in COMAH Safety Reports, *Symposium Series*, 149: 103-115

HSE, 2001, *Reducing Risks, Protecting People*, Norwich: Crown Copyright, 1-78

HSE, 2010, *COMAH Safety Reports: Technical policy lines to take for predictive assessors*, SPC/Permissioning/11, 1-31

HSE, 2010, *Guidance on 'as low as reasonably practicable' (ALARP) decisions in control of major accident hazards (COMAH)*, SPC/Permissioning/12: 1-17

HSL, 2003, *Good Practice and Pitfalls in Risk Assessment*, Norwich: Crown Copyright, 1-54

ISO, 2010, *IEC/ISO 31010:2009: Risk management - Risk assessment techniques*, Brussels: CENELEC

Ni, H., Chen, A., Chen, N., 2010, Some extensions on risk matrix approach, *Safety Science*, 48: 1269-1278

Pitblado, R.M. and Turney, R.D. eds., 1996, *Risk Assessment in the Process Industries*, 2nd Edition, Warwickshire: Institute of Chemical Engineers

Prince, M., 2005, *Developing and Using Risk Matrices*, UK: Atkins Consultants Ltd, 1-17

Summers, A.E. and Hearn, W.H., 2012, Risk Criteria, Protection Layers and Conditional Modifiers, *Process Safety Progress*, 31: 139-144

Talbot, J., 2011, What's right with risk matrices?, Australia: Jakeman Business Solutions

The Management of Health and Safety at Work Act, 1999, *The Management of Health and Safety at Work Regulations 1999*, UK

Worksafe Victoria, 2011, *Safety Assessment for a major hazard facility*, Australia: Worksafe Victoria, 1-32