ASSET RISK MANAGEMENT: ISSUES IN THE DESIGN AND USE OF THE RISK MATRIX

Ype Wijnia [a,b]

[a] *Delft University of Technology, Faculty of Technology, Policy and Management, Jaffalaan 5, 2628BX Delft, The Netherlands*

[b] *D-Cision bv, PO box 44, 8000 AA Zwolle, The Netherlands*

Within the profession of asset management many practitioners use a risk matrix in order to prioritize attention. As a risk matrix is an expression of the value system, each organization has own version: there is no universal truth. However, in developing such a custom made matrix, many things may go wrong. Wrong in this sense is that the decisions made with the matrix are not generally perceived as being good decisions. This is often due to errors in the design of the matrix, and misconceptions about its use. In this paper, guidelines for the correct design and use of a risk matrix will be presented and report on some common errors found in the asset risk management practice will be given.

Key Words: Risk Matrix, Decision making, Asset Management

1 INTRODUCTION

The management of risk within the profession of asset management has gained in importance over the years as it has been advocated by numerous relevant parties and sources. The Infrastructure management Manual [1], mentions risk management for example in its definition of advanced asset management (suggesting at least that it is the way forward), PAS55[2] includes the risks of assets in its definition of optimal management. In parallel to this growing significance within asset management, general risk management (or enterprise risk management) has gained in importance as well, no doubt influenced by the Sarbanes Oxley act in the USA. Within risk management, an essential stage is assessing the risk level. As risk is a two dimensional entity (often referred to as risk equals impact times probability), assessing the risk level means assessing both impact and probability and judging their combination against some risk criteria. A tool often used in this exercise is the risk matrix. This is a table which has several categories for probability on one axis and several categories for impact on the other axis, and a level indication per cell, often both in description (low, medium or high) as in colour (green or red risks). As such a risk matrix is an easy to use tool, many practitioners have adopted it and there are few asset managers around that do not have a risk matrix.

However, despite its widespread use, there is surprisingly little standardisation of risk matrices across industries, nor is there a solid framework for constructing a risk matrix. For example, the ISO 31000 standard [3] on risk management does not even mention risk matrices. The accompanying document on risk assessment techniques [4] mentions risk matrices in its appendices but does not really go beyond mentioning that the format depends on the context and that it is important that an appropriate design is used. A few design options are considered, and a list of limitations with using risk matrices is given, but again little guidance on how to do it correctly. This lack of guidance seems to be typical for any publication on risk matrices. As a result, practitioners are left to themselves in developing the risk matrix they are using.

From one perspective, this is a good thing. A risk matrix is an expression of the value system (including risk attitude) of the organisation that uses it, and thus is it at least specific to the organisation. However, it is a bad thing in a different perspective. First of all, even though organisations have different value systems, are they really that different? Organisations operate in the same society, and their attitude with regard to financial losses for example may be more aligned than many think as they may have (in the end) the same shareholders. Additionally, there is a market for insurance, which to an extend is a standardized product and thus presumes comparable risk attitudes. Furthermore, as all kinds of regulations limit the risks organisations may impose on others (like regulations on external safety, operational health and safety , emissions) one might expect (or hope?) that their attitude with regard to those values is not that different from what society thinks is acceptable.

Secondly, the lack of well defined procedures for establishing the risk matrix may mean that the practitioners get it wrong. That is, the resulting risk matrix does not capture the value system of the organisation, and decisions based on the matrix are not perceived as good decisions.

With regard to the first point one might wonder why it is a problem. Some may even argue that going through the process of defining the risk matrix is vital for creating ownership of the matrix and thus building a risk management culture, even though the end result might be achieved by copying a risk matrix from another organisation. Besides, the effort required for creating a risk matrix is not that big that it has a catastrophic (or even noticeable) impact on the organisation. However, it is serious because of the second point. If there are no standard procedures, and no template to start with, it may be a stroke of good fortune to get it right. As a result, one might expect many decisions based on those risk matrices to be wrong.

This paper deals with the issues in designing and using a risk matrix in asset risk management. First, the basic concepts of risk management will be defined, including the definition of risk, the risk matrix and the risk management process. This is followed by a section on the design decisions to be made when developing a risk matrix. Several options will be discussed, illustrated with some practical examples and real life experiences. The next section is on problems in and criticisms on the use of risk matrices, including a discussion on their validity. This results in a review of the risk management process. The paper ends with conclusions on the value of a risk matrix in asset management decision making.

2 BASIC CONCEPTS

Risk has many meanings in everyday use. For example, Beer and Ziolkowski [5] mention 13 different definitions. Most significant difference amongst these definitions is that between risk as the entity (the threat, danger or event with a potential for undesired effects) and risk as a measure of bad fortune, often defined as risk = impact *probability. To prevent ambiguity, within the context of this paper **risk** refers to the entity "event" and **exposure** to the expected bad fortune. As mentioned, the exposure is related to the probability. However, in a strict mathematical approach probabilities cannot be larger than 1, whereas within asset risk management risks may materialize more than once. In such an operational risk management environment, it is better to make use of the term frequency, like in "this asset fails twice a year". Frequencies, on the other hand, lose meaning for unique or very rare events and it is better to use the term probability, even though numerically frequency and probability are virtually equal for small numbers. To avoid referring to probability/frequency the more general term of likelihood will be used in this paper, referring to both concepts. The formula then becomes exposure = impact*likelihood.

Likelihood

Impact

	Very low	Low	Moderate	High	Very High
Very High	M	H	H	H	H
High	M	M	M	H	H
Moderate	L	L	M	M	H
Low	L	L	L	M	M
Very Low	L	L	L	L	M

Figure 1: Example risk matrix

The risk matrix is a table used to characterize risks (according to the definitions above, is to characterize the exposure of risks). A risk matrix is a table with impacts on one axis and the likelihood on the other. The diagram on the left is a typical example. The colour coding and characters represents the risk level, with red for high risks (H), yellow for medium risks (M) and green for low risks (L). This classification (including the names to specify categories) is the expression of the value system of the decision maker and thus subjective. In the example, risk levels more or less comply with the exposure, but this does not need to be the case. Risk levels may depend solely on the impact (resulting in horizontal bands) or solely on the likelihood (resulting in vertical bands).

Within the field of risk management, different views exist with regard to the objectivity of risk assessment[6]. Some claim the exposure can be quantified objectively, as if risks are out there to be measured. On the other hand, many risks to be assessed have not yet materialized (like what is the exposure of the risk Nuclear War?) and quantification is based on models that are inherently assumption laden and thus subjective. Furthermore, there is ample of evidence that risks are not always judged on the exposure[7]. This is for example visible in the example risk matrix of figure 1, where very high impact low likelihood risks are judged as red and low impact very high likelihood risks yellow. Within this paper the objective approach is followed.

With regard to the risk management process, there is no standard terminology. For example, the Risk management standard [8] follows the ISO terminology and uses risk assessment as the bundling of Risk identification, Risk analysis and Risk evaluation, whereas the COSO framework [9] separates event identification from risk assessment (a bundling of risk analysis and risk evaluation). The steps are thus very similar, even though the terminology differs. Within this paper, the ISO terminology is used, as it recognizes risk evaluation (where the risk matrix is used) as a separate step. The diagram below shows the risk management process (based on [8]).

Figure 2: Risk management process

Within this process, the risk matrix is defined in step 1, and applied in step 4. It should be stressed that the risk management process does not require the use of a risk matrix, it only requires that some criteria are defined against which the risk tolerance can be judged.

The final part of this section is on the purpose of risk management and the risk matrix. In short, the purpose of risk management is to limit the exposure caused by risks that could have been efficiently mitigated. As there are many potential threats, the risk matrix is used to prioritize attention [10]. Within this perspective, an incorrect risk matrix is a matrix that does not prioritize attention correctly

3 DEVELOPING THE RISK MATRIX

The risk matrix, as shown in figure 1, is a conceptually simple tool, just a colour coded table with likelihood on one axis and impact on the other. However, there are surprisingly many decisions that have to be made to make a matrix. In this section these will be treated.

3.1 Orientation of the matrix

There are three decisions to be made with regard to the orientation of the matrix.

1. What axis is used for likelihood and what axis is used for impact?
2. What is the direction of the horizontal axis?
3. What is the direction of the vertical axis?

Each of these issues has two options, resulting in a total of 8 different orientations. In many sources (e.g. [6]) impacts are on the horizontal axis and likelihood on the vertical, with the most important risks in the top right corner, but it is no standard. The example risk matrix of figure 1 has the axis reversed, and Cox[10] uses both orientations in his paper. Furthermore, if risks are mapped together with opportunities, it makes sense to use likelihood on the vertical axis, have the risks on the left hand side (impacts increase from right to left), and opportunities on the right hand side of the vertical axis (potential benefits increase from left to right). Most important risks would then be in the top left corner. If the likelihood is on the horizontal, one could make a case for the impact scale increasing top to bottom having the most severe (negative) impacts at the lower end of the matrix, though this does not seem to be common practice.

With regard to prioritization, the orientation is not relevant. High risks are more important than low risks, no matter where they are plotted in the matrix. In terms of communication it could be important, though. If risks are usually plotted in the example matrix (high risks in the top right corner), then reversing the directions of the axis would result in the most important risks being in the bottom left corner and decision makers could get a wrong impression of the importance of a risk. Key with the orientation of risk matrices is therefore consistency.

However, this is not all than can be said about the orientation of risk matrices. First of all, plotting risks and opportunities in the same diagram is a bad idea. People judge opportunities completely different from risks, to the extent that a risk formulated as an opportunity could result in radically different decisions (see the Asian disease problem [11]). A risk matrix should thus really be a risk only matrix, and the most important risks could be plotted top right (apparently the most natural position). The second issue regards the choice of the axis. Although it is very common to plot the impacts on the horizontal axis, there are practical arguments against it. In general, there are multiple values that could be affected, whereas there is only one likelihood scale. Thus, if impacts are on the vertical and likelihood on the horizontal, one could plot the effect category description per value besides the matrix on a single landscape sheet of paper. This is easier to read than a portrait orientation, with impact descriptions above the matrix, and thus more likely to be used. Multiple page risk matrix documents (for example, the matrix, impact descriptions and likelihood descriptions all in separate tables on different pages) are a nuisance in practice and not likely to be used.

3.2 Dimension and resolution of the matrix

The dimension of the matrix is the number of impact bands, likelihood bands and risk levels to be used. The number of risk levels is often limited to 3: low risk (L), medium risk(M) and high risk(H). Low risks are then acceptable without further notice, high risks require action and medium risks should be considered for action. These three levels could be combined with any numbers of impact and likelihood bands, though 3 by 3, 4 by 4 and 5 by 5 seem the most common. There is no necessity to have equal numbers in rows and columns. However, thinking about the resolution, how much sense does it make to distinguish 9 different likelihood bands (extreme example) if it results in only 3 different risk levels? The same holds for impact bands. One could expect the risk level to increase if the impact moves along the categorizations for the same likelihood level, or vice versa. The key question therefore is what the resolution is with regard to the potential impacts. If there is a meaningful distinction in the severity of impacts, it should be reflected in the risk levels. This will be illustrated with the impact categorisation of safety.

Suppose an aircraft operator has several types of aircrafts, ranging from 747s to cargo carriers. Probably the worst event that could happen would be the crash of a plane. Would it, with regard to the categorization of the event, matter how many people would die in the crash? Most likely it would. Two 747's crashing into each other (near 1000 fatalities) would be worse than the crash of a regional plane (100 fatalities), which would be worse than the crash of a cargo plane (10 fatalities). But would it matter if a plane was only used at half the capacity? That would most likely be more difficult to judge. The scale continues downward. The next level would for example be a mechanic slipping during maintenance and breaking his neck (single fatality), followed by serious injury (lost time incidents with extended absence), lost time incidents, near misses or first aid, and dangerous situations. With regard to the top event, anything without fatalities could be considered insignificant, but in safety management the key control is the number of unsafe situations (iceberg theory of Heinrich [12]). Thus, proper risk management would require the impact bands to span all levels, which would mean 8 levels. A similar exercise could be held with regard to the likelihood axis. Unsafe situations might occur 1000s of times per year, but aircraft crashes are generally below once a year but probably above once every 100 years (few airline operators will be without any crash). Two planes crashing into each other would be extremely rare, perhaps below once every 1000 or even 10000 years (the latter figure means that less than 1% of the operators would experience such a crash in 100 years). This would mean 8 or 9 orders of magnitude in the likelihood range, which is quite similar to the number of impact bands (presuming the orders of magnitude would be used as likelihood bands). However, using a 8 by 8 matrix may be not really necessary. Risk characterisation is essentially a decision driven activity [13], and it is not likely decisions regarding small impact high likelihood (bottom right) would end up on the same table as the high impact low probability risks. The risk matrix could be split into two aligned sub matrices, one for the (internal) maintenance service provider and one for the asset strategy department. This is shown in the figures below. However, care should be taken that the sub matrix encloses all relevant risks. Otherwise mischaracterisations could occur. For example, if the service provider matrix was used to classify an impact(9)/ likelihood(v) risk it would result in a Medium(M) risk, whereas the full matrix would characterise it as Unacceptable (U).

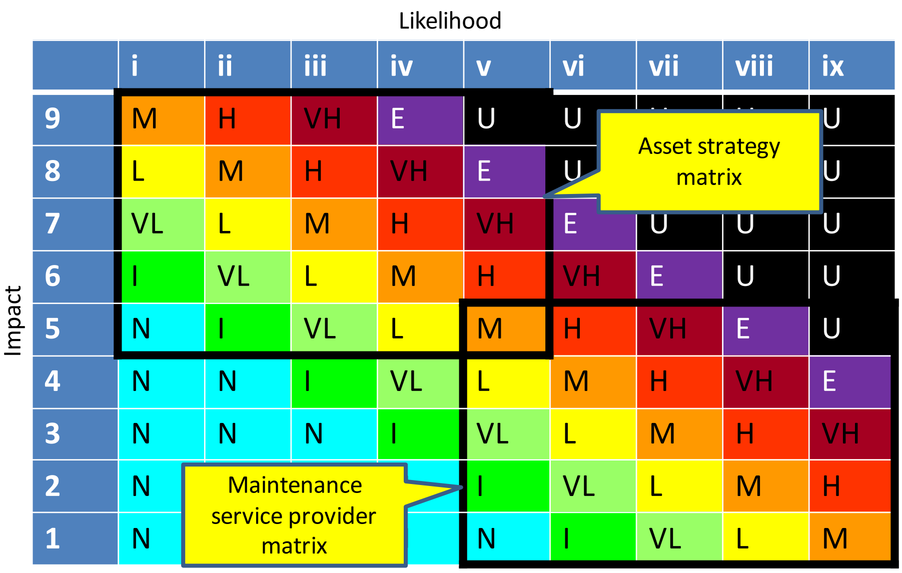

Figure 3: Overall risk matrix split into 2 sub matrices for different departments. Coding: Unacceptable, Extreme, Very High, High, Medium, Low, Very Low, Insignificant and Negligible

A key lesson in of the example above is the extreme spread (many orders of magnitude) in impact and likelihood that can occur within risk assessment. This represents a "challenge" in visual representation, as it requires a logarithmic scale to be meaningful. Using linear scales would mean that only a fraction of the risks could be distinguished from each other. For example, if the impacts of the aircraft safety would be plotted on a linear scale, the most extreme event (1000 fatalities) would be plotted on axis' end (either up or right), and the regional crash (100 fatalities) would already be at 10% of the origin. Anything below 10 fatalities would be compressed into 1% of the scale and thus would by no means be distinguishable. The same holds for the likelihood. Using logarithmic scales is not exactly the discovery of a lifetime (reason why challenge is put between quotation marks), yet some matrices neglect this knowledge. Cox [10] for example bases his arguments why risk matrices are wrong on a 5 by 5 matrix with linear scales for impact and likelihood. Others are more or less logarithmic but not constant in their scaling factor between categories.

Additionally, linear scales have a problem in the margins of uncertainty within the cells. Suppose a 5 by 5 matrix with linear scales. The top cell ranges from .8 to 1. A risk put in the middle of this category would have an accuracy of about 10%: 0.9 ±0.1. But a risk put in the lowest category would have an uncertainty of 100%: 0.1 ±0.1. As uncertainty is a central element of risk it would be surprising if an assessment could be made within a 10% margin of error. Logarithmic scales are constant in their margin of error. If a risk is in the middle of any band, it could be threefold higher or lower before it would be categorized differently.

Summarizing, the resolution of a risk matrix should be on the order of magnitude level (=a logarithmic scale), and the dimension should fit the orders of magnitude between the most and least serious impact. If the scaling factors are constant, the top left bottom right diagonals are levels with equal exposure, and should be used for the risk classification, as has been done in figure 3.

3.3 Scaling multiple values

Risk matrices often express the risk tolerance on multiple values. Risk matrices require the values to be translated into Key Risk Indicators (KRIs). For some values, like financials, the KRI is quantitative. For other values, like reputation, , the KRI may be qualitative or descriptive. Some KRIs mix qualitative and quantitative elements, like safety in the airline example. For fatalities the KRI could be made quantitative (number of fatalities, though the events are still described qualitatively), whereas the non-fatal incidents are qualitative (though there is a trade-off between the number of people and the severity of the accident).

Developing an impact scale for a single quantitative KRI is not that difficult. The orders of magnitude approach should be used, so that the exposure in each cell with the same risk level is equal. Thus, if the likelihood categories are once every 100 years, once every 10 years, once per year, corresponding financial impact categories would be 1 million damage, 100k damage, 10 k damage (or any other value scaled similarly). For practical purposes, quantities in a risk matrix are often not represented by single numbers but by bands, like 1 to 10 times per year, damage between 10k and 100k.

However, for qualitative scales things are more complicated. The problem is not in defining categories, but to assure they are an order of magnitude apart. For the safety example, one could distinguish the loss of a left hand, right hand, arm, leg, eye, ear, burns on the body, hands, face, loss of eyesight, hearing, and so on. Insurance companies may do this in order to determine the compensation fee. However, are they distinctive in terms of a risk matrix? Is the loss of a right leg comparable in severity as ten incidents where people lose their left hand? This probably is not the case, dismemberments (or amputations) are comparable in severity, though it is a difficult subject. Fortunately, the severity and likelihood of incidents are negatively correlated. For every 10 first aid incidents there is one lost time incident [14]. Therefore, categorizing the impact according to this iceberg theory would result in equal risk levels at every impact level of the iceberg.

After developing impact schemes per value, the next step is alignment. Theoretically it is possible to develop a risk matrix for every value separately (possibly with different orientations, dimensions, resolutions and classifications) , but that would be very impracticable and error prone, as risk matrices could get mixed up. Key question with regard to alignment is the comparability of the impact levels: is the most serious impact on finance really as serious as the comparable impact on safety? In other words, if the organisation has two choose between two incidents on the same impact level, would it really be indifferent? Another approach is that of substitution, though it generally only works when comparing to the financial value. The question then is whether the organisation is willing to spend the equivalent financial amount to prevent the impact. As precise valuation is difficult, this can be replaced by two questions: is the organisation willing to spend (without a doubt) the financial amount one impact category lower? And would the organisation accept the impact if the costs would be one impact category higher? However, in this comparison one has to be aware that valuation is highly imprecise, and that the willingness to pay is different from the willingness to accept (the amount that has to be spend to prevent the incident versus the amount that has to be gained in order to accept the incident). Nevertheless, it provides reasonable guidance.

4 PROBLEMS AND CRITICISMS

There are several criticisms on and problems in the use of risk matrices. They can be organized along three lines:

1: the risk matrix does not prioritize at all

2: the risk matrix does not prioritize risks correctly

3: the risk matrix results in incomprehensible decisions

4.1 The risk matrix fails to prioritize the risks

Typically, the distribution over the risk levels is not uniform: it is very rare to find unacceptable risk, high risks are more frequent, medium risks tend to be the most numerous and the numbers start declining below medium again. The reason unacceptable risks are rare to find is that companies that do not manage them are more likely to go out of business. The same holds for high risks: a company does not necessarily need to have formal risk management in place to understand it has a problem that is worth mitigating. For risks below medium another mechanism is active. People may feel the risk is not that important, and simply do not bother to mention it. This is especially true in risk identification workshop settings, where the aim is to capture the most relevant risks in a limited timeframe.

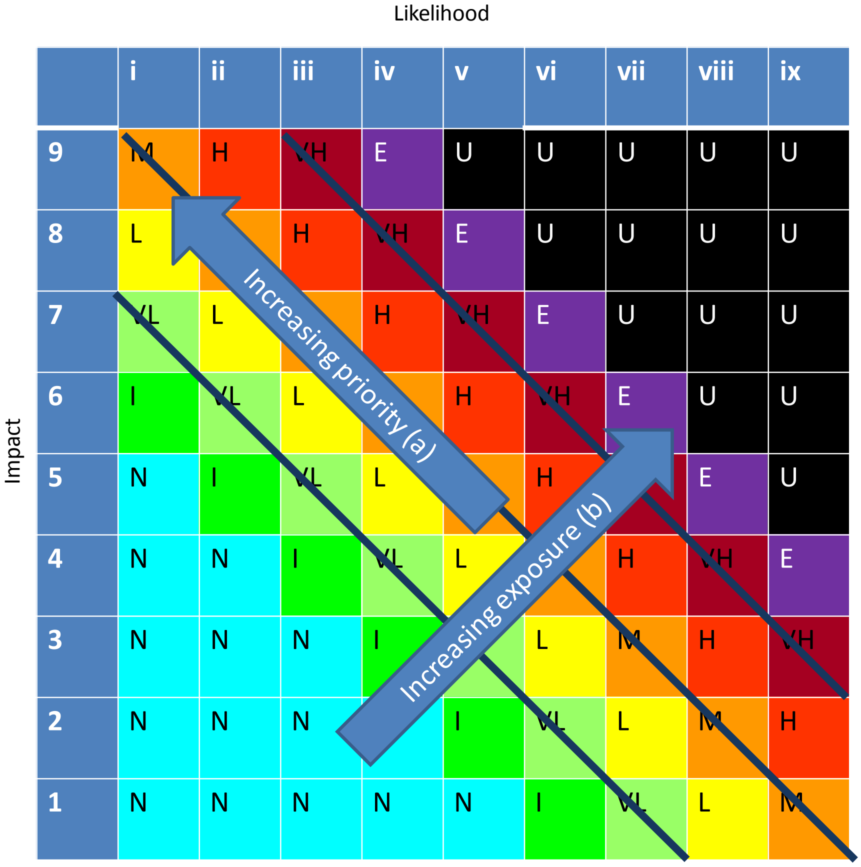

Figure 4: diagram showing equal-exposure lines (blue lines), direction of increasing exposure and priority direction on equal exposure line. Coding: Unacceptable, Extreme, Very High, High, Medium, Low, Very Low, Insignificant and Negligible

As a result, many risks may end up on the same level (medium) or in the same cell. After all, even in the 9 by 9 matrix of figure 3 and 4 there are only 81 cells and risk registers easily contain 1000s of risks [15-17], therefore it is inevitable some cells may be a bit crowded. This could be regarded as a problem for prioritization, but that is only the case if the risk matrix is applied very rigidly. First of all, if the exposure is equal, the risks with the higher impact deserve more attention (shown in figure 4 by arrow a). This is simply because unforseen high impacts cause more disturbance. It is like comparing losing a euro a day (e.g. an unused subscription) to a 1000 euro damage (e.g. broken washing machine) once every 3 years. Furthermore, because of the logarithmic scales, risks with equal exposure are situated on lines parallel to the top left- bottom right diagonal. That means that moving in parallel with the bottom left –top right diagonal means increasing risk (shown by arrow b). This logic in clear in the table, but could be applied within cells as well (though it requires much more accurate information), or within the negligible area on the bottom left.

Another cause for failing to prioritize clearly distinguishable risks is that some risks may impact only one value, whereas others impact more values. Compare the risk af a large customer bankruptcy (generally only financial impact) with the risk of on site explosions, resulting in injuries, repair costs, environmental impacts and reputational damage. Assuming all impacts and likelihoods are comparable, the risks would end up in the same cell. In human judgment though the explosion risk would be considered more serious because the total impact is larger, but the matrix is not able to differentiate single value impact risks from multi value impacts, and there are no rules to add risk levels. The only way around it is to express all impacts in a single unit, multiply them by their likelihood and add the expected values to arrive at the total annually expected misery. This is a continuous scale that thus has a much better resolution than the risk matrix, which is discrete. The unit in which impacts are expressed could be dimensionless, but as the impacts scales have been aligned with the financial scale, expressing everything in monetary equivalents is a practical solution. Besides, if the alignment is based on the willingness to pay, the monetary equivalent is also an indication of the maximum amount that could be spent on mitigating the risk.

The final cause for the risk matrix failing to prioritize risks is that it is used for the wrong type of risk. An example is the system performance risk, like the number of outages that occurs per year, the number of delays, the total financial loss and so on. But these are variations around some mean. Even though this variation around a mean corresponds to risk in a different context (the capital asset pricing method [18]), that is not the context where risk matrices are used. Risk matrices exist to characterize event risks, which can be expressed and measured in terms of impact and likelihood. System level performances cannot be expressed in likelihood, as they have a fixed timeframe (in many cases a year). The risk matrix is thus compressed into a single column. Furthermore, system level performances typically have variances measured on a percentage scale, not on an order of magnitude scale. Thus, if a system performance is plotted in the matrix, it will always end up in the same cell, year after year. Besides, one should not be surprised to find an high level system performance in the intolerable area.

A typical example is the total costs which could be easily 80% of turnover. If an event would occur that would cost that much it would no doubt be in the highest impact category, and if it would occur every year it would be intolerable (in fact, the company would go bankrupt). But that is an additional cost, not the normal cost, as the normal cost has been included in the business plan. Finally, unweighted performance indicators (like the number of delays) lose information on the severity of impacts in the indicator. A 30 minute delay is completely different from a multiple day delay, yet they could be counted as equal in the indicator.

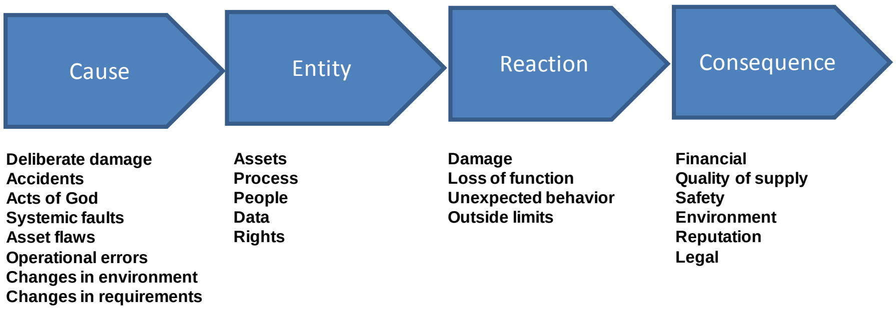

Figure 5: the risk process, based on [16]

A good way to prevent risks being defined wrong is to think of risks as a chain from cause to consequence, as shown in figure 5. A cause happens, impacting an entity, which shows a reaction which has value consequences. For assets the reaction is typically the failure mode [19].

4.2 The risk matrix does not prioritize correctly

Even if the risks are defined properly so that the risk matrix can be applied, occasionally the risk matrix may produce incorrect prioritizations. In this context, incorrect prioritization is that a risk with a higher exposure is assigned a lower risk level [10]. This is in a sense a more serious form of failing to prioritize, as it can result in important risks (in term of the exposure) not getting the attention they deserve. Suppose one risk has a likelihood of 0.9 and an impact of 9.9, whereas another risk has a likelihood of 1.1 and an impact of also 1.1. If the borders are 1, 10 etcetera the first impact would be a lower level risk than the second (same impact band, lower likelihood band), yet the exposure is 9 versus 1. This is a clear mismatch between exposure and risk level. Unfortunately, this potential for incorrect prioritization is inevitable. Both likelihood and impact are an order of magnitude wide, resulting in the risk level being two orders of magnitude wide (1*1 versus 10*10). Yet, risk levels are only one order of magnitude apart, as moving up one category in either direction should result in a higher risk level. It does not matter whether scales are linear or logarithmic, nor if the risk levels are two orders of magnitude apart. At the edges the risk levels overlap. This is not a mere theoretical fact, it really happens in practice, as shown in figure 6.

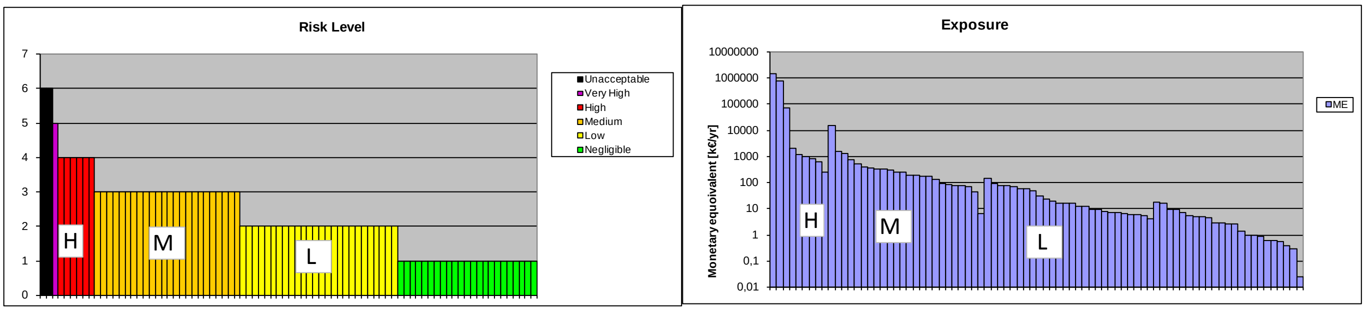

Figure 6: Risk level and exposure for a real collection of risks (both graphs have same order of risks)

Even though there is a mild correlation between exposure and risk level, it is by no means perfect. At the edges of the risk levels, the exposure suddenly goes up, sometimes more than an order of magnitude. On a closer inspection, it becomes apparent that about half the Medium risks are higher in exposure than the lowest High risk. Only Low risks are certainly lower in exposure than High risks. But does this mean that risk matrices are fatally flawed? No it does not. As can be seen in figure 4, the Low risk band is precisely between very Low and Medium, but there is a significant overlap with these categories. However, there is no overlap with the High risk band, it is just contacted in the edges. Secondly, if risks are placed into an impact or likelihood band, they are probably in the middle, and then the distinction seems reasonable. Only if with reasonable accuracy the risk is near the edge of a band, the risks are prioritized incorrect. But why would one use a rough method as a risk matrix is more detailed information is available?

Another problem in incorrect prioritization can occur in the lowest bands of the matrix, as they generally are limited on the lower side by 0. Yet, in orders of magnitude that is an infinite number of bands. Thus, in the example of figure 3 with the border between impact 1 and two being 1000 euros, , a 1 eurocent loss occurring more than 1001 times a year would be a medium risk, whereas a 999 euro risk occurring 99 out of 100 years would be negligible. The exposure is 100 times higher, the risk level 3 categories lower. But this is simply an incorrect design of the matrix. The bands are orders of magnitude, so even the lowest band has a lower limit. Anything below that limit cannot be accurately measured by the risk matrix. If the risk is relevant, the number of bands should be expanded until the risk fits into the matrix.

4.3 The risk matrix results in incorrect decisions

Risk matrices are often advocated as tool for decision making in terms of risk mitigation. The risk level determines whether the risk should be mitigated, and the position in the matrix determines what type of mitigation is preferable. For high impact low probability risks **transfer** (insurance) is recommended, for more likely high impact risks the impact should be **reduced** (like emergency response plans), high likelihood risks should be **prevented** and low impact low likelihood risks should be **accepted**. This is shown in figure 7.

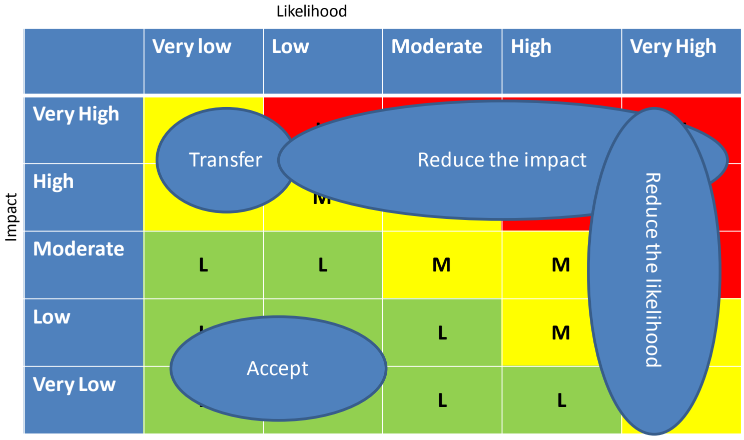

Figure 7: Risk mitigation options in the matrix

A well known criticism on the mitigation scheme is that is does not appreciate the costs nor the availability of a measure. It is not all high risks that have to be eliminated. If mitigations are very costly, one could be better off by accepting the risk. Only intolerable risks have to be mitigated (hence the name intolerable), though even this is relative. Some risk may be branded intolerable but prove impossible to mitigate. They have to be accepted and thus are de facto not intolerable (an illustration that care should be exercised in determining what is intolerable). Risk acceptability depends as much on options to mitigate the risk as on the risk level[13]. Yet, the rule of thumb of figure 7 has perseverance, presumable because the rule works reasonably well in situations where the costs of mitigations are comparable. This condition is not realistic in asset risk management, some mitigations may be relatively cheap (e.g. extra inspections) and others very expensive (e.g. requiring complete redesign). But the condition seems to apply in deciding what mitigations to develop (but not yet execute). The analysis costs of a risk are not very dependent on the risk level, as are the design costs. And if no other information is available than estimates what order of magnitude impact and likelihood are, it may be the best there is. However, the most important task of risk analysis is resolving the uncertainties around the impact and likelihood, or at least make them explicit. Once the information is available, there seems to be no reason to base decisions on the risk matrix, which is (given the previous to criticisms) a rough tool. There are much better approaches available, based on rationales like cost benefit theory, utility theory or social choice theory [20]. Expressing the exposure in terms of monetary equivalents is already a first step towards cost benefit analysis, which has (from a pragmatic viewpoint) the advantage that it results in a clear yes or no for the mitigation based on the equivalent net present value. If there are budget constraints, the best set of mitigations could be determined with a portfolio decision[21]. To do the standard [3] full honour, it should be noted that it prescribes balancing the costs and efforts against the benefits.

Additional to this point with regard to the accuracy of the matrix as a decision tool, there is the problem of insufficient mitigations. Insufficient in this sense is that the risk level before and after the mitigation are the same. This could happen for example if a mitigation “only” halves the exposure. In terms of the risk matrix it would be an insignificant mitigation (no effect) and thus be rejected, whereas it could be a very good mitigation in terms of the cost benefit ratio. This problem only applies to risk matrices with logarithmic scales, but it has been argued why risk matrices need to be constructed this way.

Summarizing, the criticism on the risk matrix as a decision tool for implementing mitigations is completely correct. Risk matrices are not useful for decision making if more than an order of magnitude accuracy is required. This is generally the case in decisions on actual mitigations. Nevertheless, the risk matrix is a useful tool in determining what risks to analyse and what mitigations to develop. As the costs of these actions are comparable for all risks, the categorization into risk levels is often good enough. This makes one wonder whether the risk management process introduced in section 2 should be revised. The stage of risk evaluation (judging exposure against risk tolerability criteria) suggests the use of a risk matrix, whereas it became clear that the characterization of a risk in that stage (as preparation for risk treatment) is perhaps only useful in a communicative perspective. Because of the analysis, much better information on the exposure is available than order of magnitude classification, with the additional benefit of being directly useful in decision making over mitigations. The only place in which order of magnitude information is good enough is in judging which risks need to be analysed. That, is between risk identification and risk analysis. This would require the stage of risk analysis and risk evaluation to be reversed, as shown in figure 8. Such a reversal shows large similarities with the process described for example by Korn [17], where it was introduced to filter the enormous amount of identified risk, which seems to be typical for the asset risk management context.

Figure 8: Risk management process

5 CONCLUSION

The risk matrix is commonly used in (asset) risk management. Despite its widespread use, problems do occur and the concept of the matrix itself as a useful tool has been criticised. The problems and criticisms concentrate on 3 points: (i) the risk matrix does not prioritize distinguishable risks, (ii) the risk matrix prioritizes risks incorrectly, and (iii) the risk matrix results in incorrect decisions on mitigations. In an analysis of the principles in risk matrix development and of the criticisms on the risk matrix, there prove to be two major causes for these problems, which are design errors and incorrect use of the matrix

Design errors

First of all, the matrix format can be incorrect. The risks a company faces may spread orders of magnitude in both impact and likelihood, combined with (at least initially) a significant uncertainty in what the actual values are. The scales of the risk matrix have to appreciate this, which requires the use of logarithmic scales. Linear scales fail in distinguishing risks in the lower bands, that still may be orders of magnitude apart. Furthermore, linear scales are constant in absolute uncertainty, whereas logarithmic scales are constant in relative uncertainty, which is more applicable.

Secondly, risks matrices are only useful if the different values are properly aligned. This is when the decision maker is indifferent between impacts materializing on different values within the same impact band (i.e. a serious impact on safety is comparable to a serious impact on the financial scale). Ideally, this indifference should be a willingness to pay (the decision maker is willing to spend an amount to prevent an equally classified impact on another value), as that helps in developing risk mitigations. If this is not the case, a high risk (even in the same cell) may not have the same meaning for all values.

Use errors

The risk matrix is an instrument to evaluate risks: that is events that have an impact and a likelihood. It is not designed to evaluate (system level) performances, as those are no events. Using it for evaluating the system may result in “intolerable” risks which are just part of the normal operation, or indiscrimination between (from a management perspective) very different levels of performance. Reason is that performances are measured in percentages, risks in orders of magnitude. Squeezing a performance criterion into the impact scale would result in a serious deformation of the matrix and misalignment of the values.

Furthermore, the risk matrix gives an order of magnitude result, which can be one order off. Using it beyond this resolution may result in incorrect outcomes. This almost certainly occurs in decision making on the mitigation (with the exception of mitigations being orders of magnitude cheaper than the risk) and also when prioritizing risks within the same risk level. If a high resolution is needed the actual exposure should be used, though this only makes sense if the information has better than order of magnitude accuracy. But there is no reason to use order of magnitude tools when better information is available. The order of magnitude information is good enough in determining which risks to analyse.

Summarizing, given that the guidelines for developing a risk matrix are followed, a risk matrix is a useful tool in prioritizing attention with regard to analysis of the risk. However, it fails in decisions requiring more than order of magnitude accuracy like decisions on mitigations. Yet, in many risk management process descriptions risk evaluation only takes place after analysis in preparation of risk treatment. It might be better to reverse the steps risk evaluation and risk analysis in the process descriptions, and make perfectly clear at what level of uncertainty an order of magnitude is useful. Once the analysis has been performed, much better information is available and it would be a waste of effort not to use that.

6 REFERENCES

1 IAM, International Infrastructure Managament Manual. V 2.0, UK 2002 ed. 2002: UK Institute of Asset Management.

2 Institute of Asset Management, Asset Management Part 1: Specification for the optimized management of physical assets, in *PAS 55-1*, BSI, Editor. 2008.

3 ISO 31000, Risk Management- Principles and guidelines. 2009.

4 ISO/IEC 31010, Risk management: risk assessment techniques. 2009.

5 Beer, T. and F. Ziolkowski, Environmental risk assessment: an Australian perspective. 1995, Supervising scientist: Barton.

6 Klinke, A. and O. Renn, A new approach to risk evaluation and management: Risk Based, precaution based and discourse based strategies. *Risk analysis*, 2002. **22**: p. 1071-1094.

7 Slovic, P., Perception of risk. *Science*, 1987. **236**: p. 280-285.

8 The Institute of Risk Management (IRM), The Association of Insurance and Risk Managers (AIRMIC), and The National Forum for Risk Management in the Public Sector (ALARM), A Risk Management Standard. 2002.

9 COSO, Enterprise risk management- Integrated Approach: Executive summary. 2004.

10 Cox Jr. , L.A., What’s Wrong with Risk Matrices? *Risk analysis*, 2008. **28**(2): p. 16.

11 Tversky, A. and D. Kahneman, the framing of decisions and the psychology of choice. *Science*, 1981. **211**(4481): p. 453-458.

12 Heinrich, H.W., Industrial accident prevention: a scientific approach 1931, New York: McGraw-Hill, .

13 National Research Council (Committee on Risk Categorization), Understanding Risk: Informing Decisions in a Democratic Society, ed. P. C.Stern and H. V.Fineberg. 1996, Washington D.C.: National Academy Press.

14 Wijnia, Y.C. and R.J.M. Hermkens. Measuring safety in gas distribution systems. in *First World Congress on Engineering Asset Management (1st WCEAM)*. 2006. Brisbane, Australia.

15 Morgan, M.G., et al., Categorizing risk for risk ranking. *Risk analysis*, 2000. **20**(1): p. 49-58.

16 Wijnia, Y.C. and P.M. Herder. Modeling Interdependencies in electricity infrastructure risk. in *1st Annual CZAEE International Conference “Critical Infrastructure in the energy sector: Vulnerabilities and protection”*. 2004. Prague.

17 Korn, M.S. and E. Veldman, The benefits of continuous risk management, in *International conference on infrastructure systems: Building networks for a brighter future*. 2008: Rotterdam.

18 Brealey, R.A. and S.C. Myers, Principles of corporate Finance. 2000, New York: McGraw-Hill.

19 IEC 812, Analysis techniques for system reliability- procedure for failure mode and effect analysis. 1985.

20 Merkhofer, M.W., Decision sciences and Social Risk Management: A comparative approach of cost-benefit analysis, decision analysis and other formal decision-aiding approaches. 1987, Dordrecht: D. Reidel Publishing company.

21 Wijnia, Y.C. and J.P. Warners. Prioritizing investment. The value of portfolio decisions in electricity infrastructure management. in *29th IAEE Annual International Energy Conference 2006: 'Securing Energy in Insecure Times'*. 2006. Potsdam.

Acknowledgments

This research has been sponsored by the Next Generation Infrastructure Foundation.