

A Systems Toolbox for Risk Management

Richard Barber

Ph.D. Candidate
University of NSW
Australian Defence Force Academy
barberrb@bigpond.net.au

Abstract

Current risk management practices based upon AS/NZS 4360 are non-systemic and deal poorly with soft or complex issues. This paper illustrates a series of tools designed to enable repeatable systemic analysis of risk. When these tools are combined with conversation based knowledge elicitation and effective social processes, it creates greatly increased capability within the organisation to optimise its risk management effort and to work on its systemic weaknesses. This can be considered in the form of a systems model of risk management, and also as a practical framework for designing risk management systems in organisations.

Keywords Risk management, risk systems, tools and techniques, systems model, risk management framework

Introduction

The management of risk is a major focus for managers of organisations. Because of this, most large organisations have formal risk management processes. In Australia, these are most likely to be based upon the processes described in the risk management standard AS/NZS 4360. However it can be argued that AS/NZS 4360 is based upon non-systemic thinking and that the processes it describes are perhaps “*necessary but not sufficient*” when dealing with risk in a complex, uncertain world (Barber, 2003).

A weakness of current approaches is that the tools and techniques used do not deal well with issues that are difficult to predict, define, or quantify – the so called “soft” issues (Barber, 2003). In response to that problem, this paper describes several new tools designed to improve the identification, analysis and management of systemic risks. It also proposes a new systems model of risk management in organisations, with emphasis on integrating risk analysis processes with appropriate social processes and knowledge elicitation techniques.

An overview of systemic weaknesses in AS/NZS 4360

Using AS/NZS 4360, context is established, risks are identified and each risk is analysed, evaluated and treated. This process is conducted in conjunction with communication, consultation, monitoring and review. This is illustrated in Figure 1.

The AS/NZS 4360 approach is tried and tested. It is simple to apply and it works well for risks that are easily identified, well defined and quantifiable. However it also has weaknesses. One is that it is designed as a self-contained sub-system. That is, it can operate without being integrated into the larger systems that drive organisations. This lack of integration means that risk management outcomes are often determined by other powerful systems or requirements in the organisation.

For example, managers are more likely to actively manage risk if regular risk reports are required. They may also conduct risk assessments and produce risk registers because they are required to for specific milestones - such as when producing a business case. In such cases, it is not the existence or otherwise of a risk management plan that determines the effort put into risk management- but rather whether risk management is an integral ongoing element of management systems. However AS/NZS 4360 processes are capable of being 'stand alone' and this is often what happens - with a commensurate lack of effective application. People are too busy doing their normal work to stop and do something else which appears to them to be an optional extra. Arguably, this is the case when we see risk management plans and processes that are not applied in practice.

The AS/NZS 4360 approach also does not deal well with 'soft' risks - those that are difficult to quantify or to fully define. One reason is the requirement for all risks to be identified and documented - so that they can be analysed, evaluated and treated. This is a practical approach for risks that are not sensitive in nature, but it does not work for any risk that cannot be written down for reasons of human sensitivity - such as when risks arise from the behaviour of a manager. If not documented, they cannot be managed using AS/NZS 4360 processes.

AS/NZS 4360 relies heavily upon simple methods of classification and representation. This is illustrated in the way it documents, classifies and represents individual risks. Every risk is given a risk rating dependant upon only two factors - probability and impact. However, risks can have many other important dimensions - such as urgency (how soon a decision must be made on what to do about that risk), and the likely cost of the risk treatment options available. Despite their importance, such factors do not influence the risk rating. Managers who rely

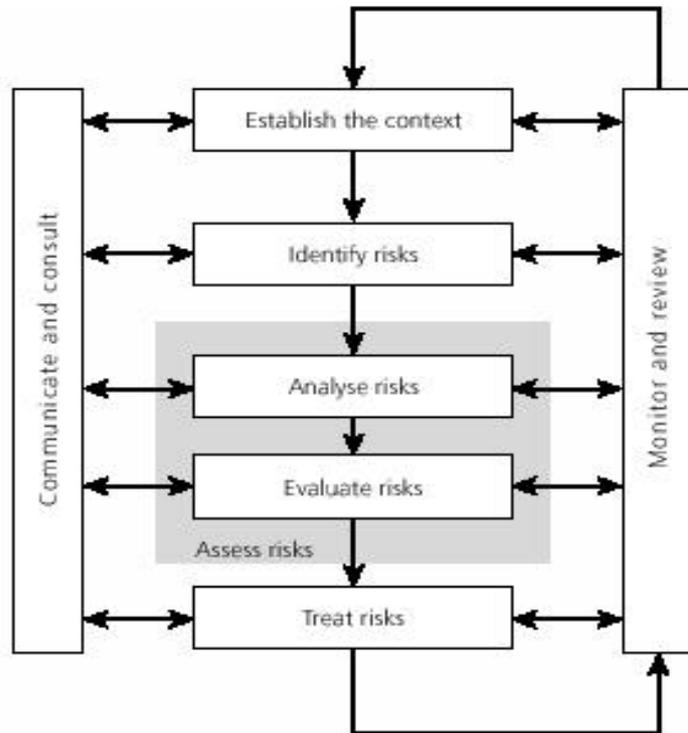


Figure 1: Typical risk management process, taken from AS/NZS 4360

only upon AS/NZS 4360 may therefore make less than optimal decisions on which risks to tackle first, and how to treat them.

Like other attributes of complex organisations, risks are often highly interrelated (Barber et al, 2002). That is, the likelihood or impact of one risk is related to the likelihood or impact of other risks. For example the risk of a cost overrun in a project is usually closely linked to the risk (if one exists) of a schedule blowout. If the schedule goes longer than planned, costs will increase. This problem is not discussed in the Standard. It does not provide guidance on how to identify risk interrelationships, nor on how to use such information when considering risk treatment action. This means that a risk that shares multiple sources or impacts with other risks will still be treated as though an independent risk. Again, the result for the manager is that they are likely to apply less than optimal risk treatments.

No discussion about risk management is complete without looking at how risks are analysed. According to AS/NZS 4360 risk analysis can be either qualitative or quantitative, with quantitative analysis being considered more rigorous and therefore the preferred approach. The method of analysis described is to assess each risk's probability and impact, and then to calculate a risk rating that combines the two to create a "risk rating". The central problem here is that the process described is not analysis of risk – it is risk categorisation.

In fact the credibility of the risk rating given to a risk will depend heavily upon whether any meaningful analysis has first been carried out. For quantifiable technical appropriate analysis tools often exist and are routinely used by competent technical professionals before they provide an estimate of risk probability and impact. However for soft risks involving human behaviour and relationships in organisations, this is less likely to be the case. Practical tools for the analysis of complex qualitative issues are less common and more difficult to apply, so managers very often resort to assigning risk probability and impact based upon intuition and experience. Indeed, the absence of any discussion of analysis within the risk analysis section of AS/NZS 4360 implies that this is the expected approach. This is a problem because whilst intuition and experience serve well in some situations, this is not a viable or reliable approach when dealing with the levels of complexity routinely present in organisations (McLucas, 2003).

On the whole, under AS/NZS 4360 the analysis, classification and representation of risks is simplistic. It may be a reasonable approach for well-defined quantifiable risks, but is unlikely to be effective in the real world of multiple interrelated soft risks. These issues are covered more fully in McLucas (2003), and Barber (2003). The remainder of this paper will illustrate some systemic risk management tools to overcome the weaknesses identified, and will propose a larger systemic framework for risk knowledge elicitation and analysis.

Tool 1: Using modified cause and effect mapping to understand complex risks

If we accept that both organisations and their environment are complex and uncertain, then we should also expect that the risks they face will demonstrate a commensurate degree of complexity and uncertainty. If that is true for a significant number of risks, then we need to have the capability to analyse and therefore understand such risks. We also know from experience that risks may be hard (quantifiable and or easily defined) or soft (qualitative and or difficult to define) in nature, and that in many cases a single risk will involve a mix of both characteristics. Together, these requirements mean that to be useful, risk analysis techniques should be capable of dealing with complex, interrelated, hard and soft issues.

One way to achieve this capability is to use a modified form of concept mapping. The process of mapping tends to force relationship questions to be asked and answered, and is capable of dealing with a combination of hard and soft issues. Mapping does not in itself provide a single risk rating for the risk being mapped, but it does mean that a risk rating (if required) can be based upon greater understanding. However a risk map also represents complex risk source and impact interrelationships, which can then be reviewed and updated as circumstances change. It is important to understand that risk maps are intended as a basis for thinking about and understanding the underlying nature of complex risks. Like other concept or cognitive maps they are fuzzy in nature and do not seek to predict or define exact cause and effect (McLucas, 2003).

Building a risk map commences with an initial risk statement. This may be a risk source (e.g. “the project is short of staff”), a risk impact (e.g. “critical computer data may be lost”), or a risk relationship statement containing both source and impact (e.g. “changes in scope may lead to schedule slippage”). The exact nature of the initial statement is not critical, since it only determines where the mapping process starts, not the final shape or content of the risk map. Other claims, statements, or ideas that seem to be related to the risk may also be placed on the map at this stage, to make them visible and available for inclusion in the risk map as it develops. They may be deleted or modified later if that is appropriate. Each single item of text on the risk map is called an “element” of the risk map.

Once there is at least one initial statement on the map, building the rest of the risk map is a process of adding elements and the links between them. Relationship links (represented by arrows) on the risk map represent influence, rather than a quantifiable specific cause and effect relationship— that is, if A is true then B tends to be true at least some of the time. Each link should remain in place only until it is demonstrated to be flawed or until a better explanation is provided. The risk map as whole is not a flowchart or other process representation, but a set of relationship hypotheses that seek to explain how and why a risk is arising, and how and why it may have impacts. An example of the first stages of a risk map is shown in Figure 2 below. The symbols used in the risk maps that follow are defined in Appendix 1.

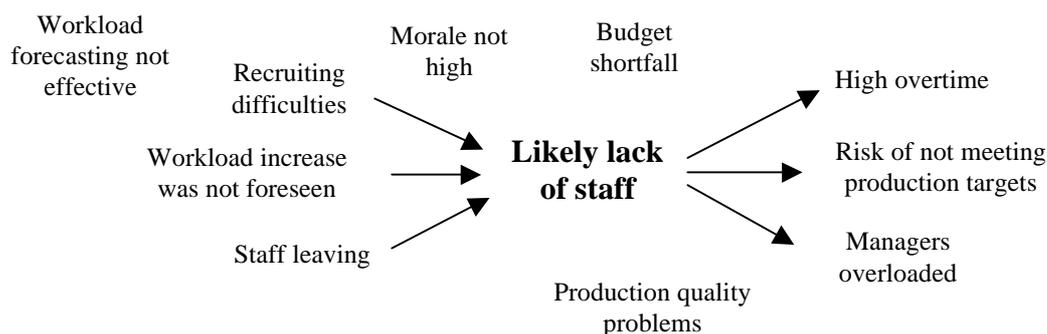


Figure 2: This example of the first stage of a risk map is based upon a “*risk that production targets will not be met because of a lack of staff to deal with increased workload*”. A superficial response could be to hire more staff. This might provide short term respite but the systemic causes mean that the same problem will arise again later. At this early stage of mapping, it appears that there are several contributing factors to the shortage of staff, and several different (but perhaps inter-related) impacts. So far, the diagram raises more systemic questions than it answers – Why was the workload increase not foreseen? Why is morale low? Why is it hard to recruit? Are production quality problems in any way related to this risk? If so, how and why?

Information may already be available to answer the questions raised by the first elements of the map, so that the map can be extended. Alternatively, the map may prompt new conversation and inquiry. The process of building the risk map continues by asking and answering questions until it appears that there is no further value in doing so. This is a matter of judgement and there is no right or wrong time to stop mapping. As the map builds, it is necessary to identify new relationships between each new element being added and the risk elements already in place. It is also often useful to reposition existing elements to make new links easier. Figure 3 shows the risk map of Figure 2 considerably expanded.

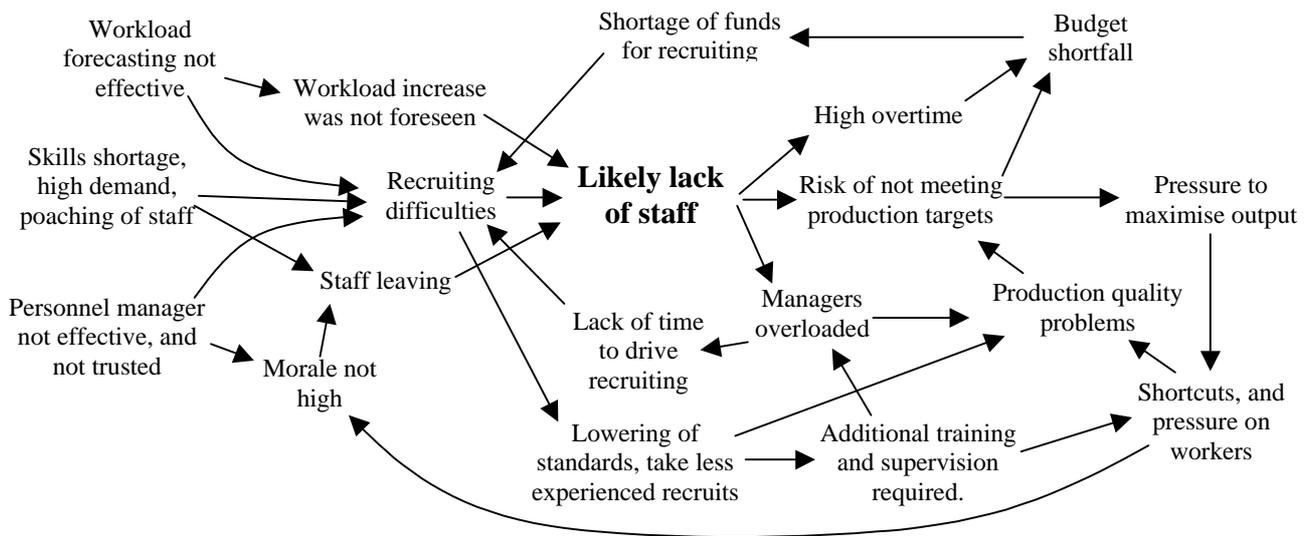


Figure 3: As elements are added and linked, the complexity of the risk map increases. There is no suggestion that any single reader could absorb the whole meaning of the map, or even that there is any such “whole meaning”. The map illustrates imprecise ideas about relationships between elements. These ideas should be seen as propositions or as hypotheses for testing. In this way the risk map promotes conversation, further inquiry, and greater understanding of what is going on.

The ability to identify and include sensitive issues in the risk map depends very much upon the quality of the knowledge elicitation upon which the risk map is based. For example, in this case it could be that the low morale is being contributed to by abrasive management behaviour from the personnel manager. If so, this is not something that can be discussed easily in a public forum, but it may be important information when trying to manage this risk. This means that knowledge elicitation needs to occur in a way that permits such sensitive issues to be identified at least, if not mapped. The problem of how to deal with sensitive issues that are difficult to document will be discussed further later in the paper.

Before considering possible risk treatments, it is important to identify any reinforcing feedback loops that exist. Although such feedback may not cause obvious problems most of the time, if any one element in a loop goes out of control then the whole loop may do so. This can result in a sudden and catastrophic variation in the outputs of a process or system that was thought to be under good control. Figure 4 identifies four feedback loops in this risk map.

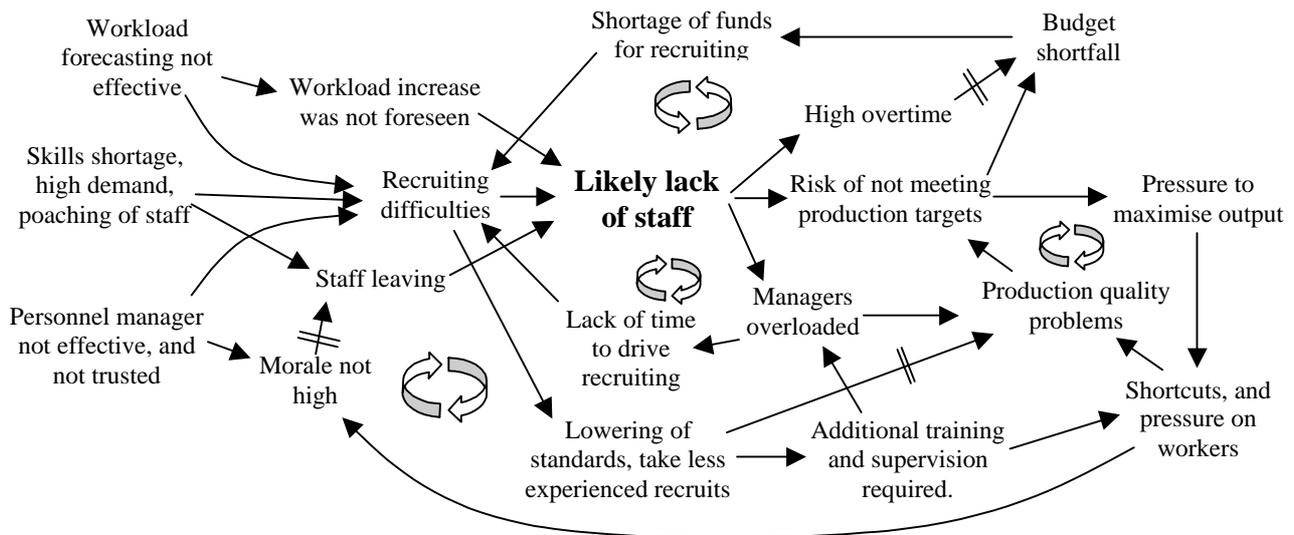


Figure 4: In this case, there are a number of reinforcing loops. As an example, if staff leave this increases the shortage of staff which increases the risk of not meeting production targets and hence also increases the pressure to maximise outputs. This places workers under pressure and reduces morale, making it more likely that staff will want to leave. In many such loops there is an implied time delay – the flow-on impacts are not instantaneous, making it more difficult to comprehend why things are happening. Rather than try to quantify, predict or control outputs, for long term improvement it is more appropriate to modify the system design in order to influence the nature of the dynamics.

Once feedback loops are identified and understood, appropriate points of risk treatment intervention can be considered. In Figure 4, it seems that the impacts on the RHS arise largely almost directly from the shortage of staff although there is some additional feed-forward directly from ‘Recruiting Difficulties’. This means that one necessary outcome of treating the risk is that the staff shortage must be reduced. In the short term it may be appropriate to take direct action to make recruiting more effective – perhaps with additional funding, and management effort. Whilst the shortages remain, we may also have to take steps to ensure that output pressures do not cause product quality standards to drop, and to help managers with the resources they need to conduct additional training and recruiting of new staff as they arrive.

In the medium term, the short term treatments will have limited impact if the abrasive behaviour of the personnel manager continues, because people will continue to leave and recruiting effectiveness will remain weak. In the long term, to avoid repeats of this problem appropriate risk treatments may be to make system design improvements in workload forecasting and in performance management.

All proposed treatment actions should be shown on the risk map, so that they are seen in the same context as the other elements. In this way, the risk map becomes the embodiment of the current understanding of the nature of the risk and how it is to be treated. If circumstances change later, the map can be revised. For the example, the overall map with risk treatments added might be as shown in Figure 5. Through the understanding gained and illustrated in the risk map, the risk treatments proposed are now much more comprehensive and effective than the classic knee jerk reaction of hiring more staff.

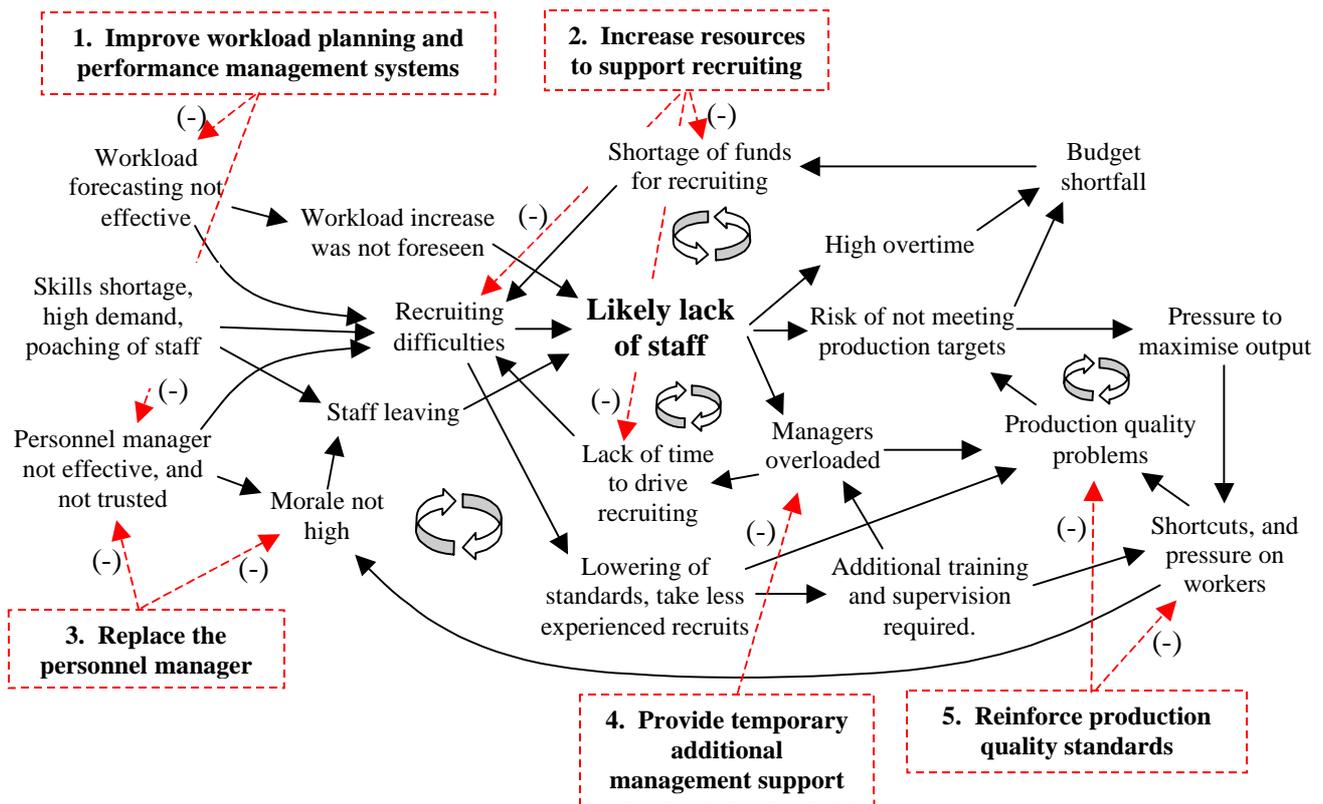


Figure 5: This diagram shows the risk map including proposed treatment options. The mix of risk treatments will often include short, medium and long term interventions. In this case treatments 2, 4 and 5 are short term and Risk Treatment 3 is medium term. The long term intervention is Treatment 1, since it aims to reduce the systemic sources of risk by improving the system design. Treatment 1 is a ‘systemic’ intervention, the others are not.

Risk maps do not exist in isolation. For better overall management, the risk processes described in AS/NZS 4360 are still needed. For example, it is still valuable to list each risk in a risk database or risk register. This is necessary to ensure that risks are managed once identified, but it is through the risk map that a rich and complex understanding is achieved and documented for later review. This is an important distinction, central to understanding why it is not appropriate to rely on the risk standard as the total risk management approach. AS/NZS 4360 proposes risk control processes - it does not provide analytical tools for the examination and critical understanding of risks.

A risk map represents the understanding of its contributors, and should be used to support a conversation and further exploration of the nature of the risks represented. Risk maps provide a documented and reusable form of risk representation, particularly suitable for soft risks that are otherwise poorly managed using AS/NZS 4360 systems alone. Once a risk map has been completed it is more likely that appropriate risk assessments and categorisations can be made, that risk treatments will be systemic, and that unforeseen consequences will not arise.

Tool 2: A risk prioritisation chart in three dimensions

A common characteristic of most risk management systems is that sets of risks are reported to management in some form of probability-impact chart (P/I chart) or an equivalent table. The purpose of such charts or tables is to enable managers to see their risks in relation to one another, so that they can decide which to work on first. Such charts and tables are usually two dimensional, and in their basic form they only provide probability and impact information. Although the exact form varies, a basic chart is shown at Figure 6.

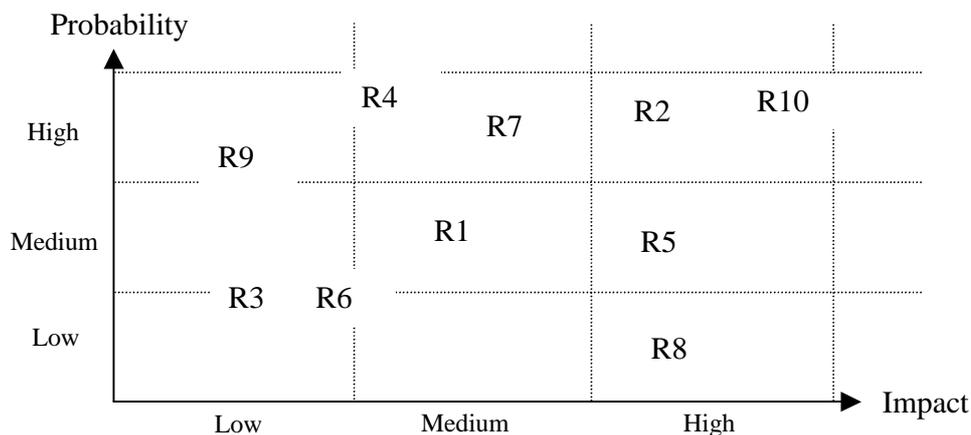


Figure 6. In AS/NZS 4360 systems risks are typically shown in a P/I chart so that the highest overall risks can be seen and worked on first. Other than working directly from the risk register, this is typically how risks are prioritised and reported. The amount of information about the risks and their relative importance is very small. It does not identify whether some of the risks require decisions more urgently than others, nor does it include other important information like treatment costs. In a complex world, a P/I chart is unsatisfactory for making decisions or even for setting priorities.

To add further information, each risk may be plotted with a colour or one or more symbols added to discriminate between risks. This serves to highlight (for example) high, medium and low risks, or it may provide additional data by categorising risks as financial, technical and so on. As colours and symbols are added, the chart or table becomes progressively more cluttered, and more complex to understand and to use.

The key weakness of P/I diagrams is that they inherently display two dimensions only and hence cannot include additional data about each risk without becoming much more difficult to interpret and apply. However including even one additional dimension can greatly increase the value of the risk information present. For example for each risk it is important to consider the optimum time for a risk treatment decision to be made. A relatively low risk may require a decision urgently (if it is to be treated at all) whereas a decision on a much more significant risk might best be left for some time. To take this into account managers need some way to include timing information without losing existing information and without making the chart confusing. Figure 7 shows how this can be done using a quite different type of risk prioritisation chart.

Probability = 1

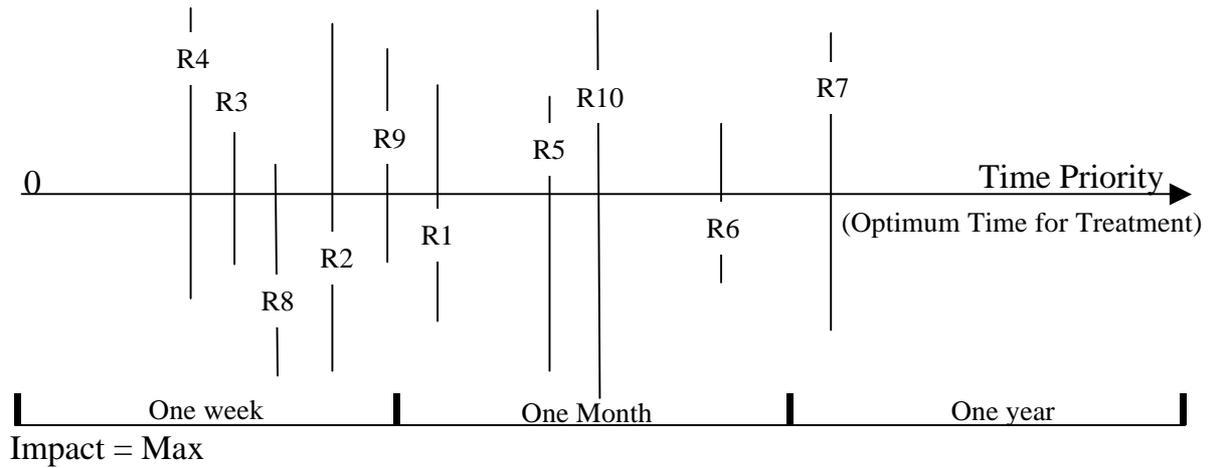


Figure 7. Risk Prioritisation Chart. For comparison, this chart shows the same risks as in Figure 6. The probability and impact dimensions are shown by the length of each risk line above and below the centre line respectively. For example, Risk 5 (R5) is shown as a moderate probability and high impact, as it is in Figure 6. The best time to act is also represented - by how far along the line from left to right the risk is. In this case, R4 is shown as the most urgent risk for a decision to be made, and R7 the least urgent. To prevent congestion on the LHS of the diagram, a non-linear scale has been used, such that the first third of the line represents one week, the second one month and final third one year. This chart shows that several risks require almost immediate treatment, several more need attention in coming weeks and only one can safely be left for more than one month.

Using Figure 7, senior management may well decide to focus their time and attention on quite different risks than if they only had the P/I chart shown in Figure 6. To further help managers to see which risks may be worth their attention now, it is useful to add a priority cone in the diagram. Risks that cut both cone lines are more likely to be worth review now, since less urgent risks must be more probable and have larger impacts in order to extend outside the cone. This concept is shown in Figure 8.

Probability = 1

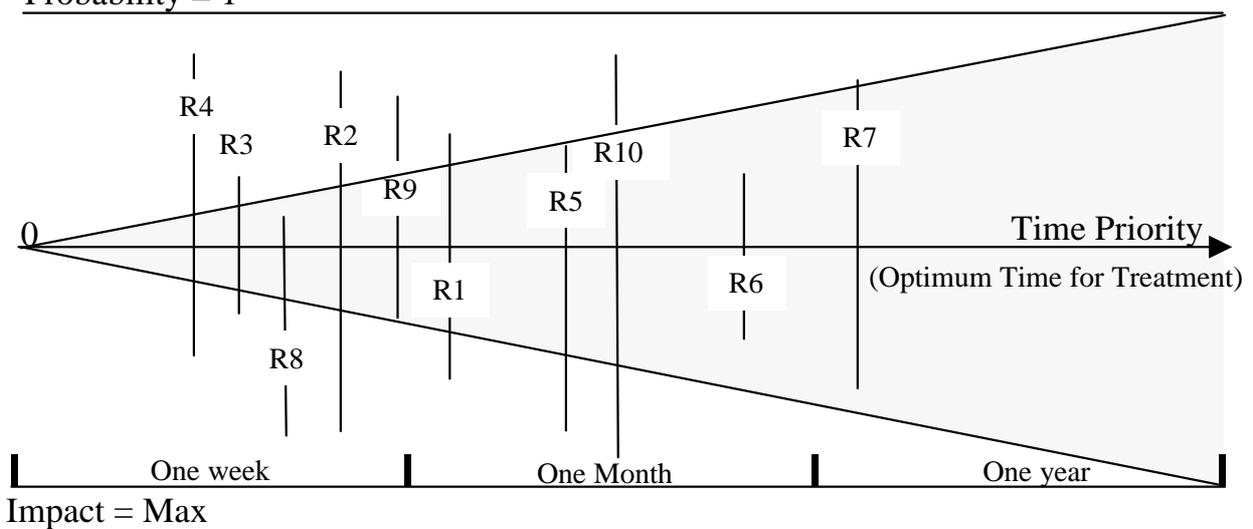


Figure 8. Drawing a cone from left to right in the Risk Prioritisation Chart provides a useful way to assist managers to focus on the most appropriate risks to review. Those risks that extend beyond both sides of the cone (above and below the cone area) are of most interest to senior management.

In this case R4, R3, R2, R1 and R10 stand out as risks that should be a priority. Even though R10 is less urgent, it has such a high P/I rating that it still extends beyond the cone. It deserves consideration in advance. The cone is an arbitrary device, but may be useful in directing attention to both the most important and the most time critical risks. As required, the time priority dimension can be replaced by treatment cost, cost-benefit or any other useful dimension for management consideration. A non-linear scale may or may not be required.

By providing an additional dimension in a diagram that is easily interpreted, the Risk Prioritisation Chart provides a better basis for risk prioritisation and reporting to senior management. It does not replace the need for a rich understanding of each individual risk, and should be used in conjunction with risk mapping.

Tool 3: A relationship matrix can identify the risks that offer the best risk treatment leverage

Individual risks are almost always related to other risks by multiple shared sources and or impacts. For example, if there is a risk that quality failures will lead to a loss of sales, then there are also likely to be related financial risks - such as increased costs of production rework and reduced income because of loss of clients. Working on any one of these risks will impact on the others – they cannot safely be treated in isolation. This being so, it is worth considering which risks to treat in order to obtain the best overall result for the treatment effort applied – i.e. the best leverage on risk as a whole.

The links between risks are often subtle and may become visible unless they are closely scrutinised, such as when safety risks arise indirectly from production difficulties apparently because pressure to increase outputs leads to corners being cut. In fact, the real link may be financial – the risk of decreased profits provides a motive for changes in management behaviour, which may then lead to safety risks. Likewise, the risk of failure of a new technical design in a project is likely to be related to the risk of schedule slippage and of financial problems in the same project.

Regardless of whether risk interrelationships are clear and immediate or subtle and time-delayed, it would be inappropriate to manage one risk without considering the impact (or opportunities) arising from related risks. There may be several different options for influencing all the risks, or it may even be that working on just one of the risks may greatly reduce the others. Unless this is understood and used to advantage, it is likely that both risk treatment effort and risk reduction impact will be sub-optimised.

A new perspective on how and where to work on risk can be achieved if we can assess and represent in a practical form these many-to-many relationships. This would permit identification of the “sources of sources” – that is, those systemic sources of risk that drive multiple risks. A focus on treating such sources can deliver risk reduction benefits disproportionate to the effort or cost involved. To see this, consider the matrix in Figure 9 below.

Risk Relationships Matrix		Risk 1	Risk 2	Risk 3	Risk 4	Risk 5	Risk 6	Risk 7	Risk 8	Risk 9	Risk 10
Risk 1			1		1	3		2	1	2	
Risk 2		2		2		2	3	1			2
Risk 3						2				1	
Risk 4		-1	2			1	1	2		2	3
Risk 5		1		2				1		1	
Risk 6			1		1				1	1	
Risk 7		2		1	1	2					1
Risk 8		1	3	1	1		1	3		2	2
Risk 9		2				2		1			
Risk 10			2		1	1	1		1		

Figure 9. The data in the Risk Relationships Matrix shows the relationship between each individual risk on the left and the same set of risks across the top. For each element of the matrix, the question is “If the risk to the left becomes more likely or increases in impact, might this change the likelihood or impact of the risk above?” 3 indicates a likely strong positive (same sense) correlation, zero or “blank” is no apparent correlation. A negative number indicates an inverse relationship.

Figure 9 indicates that if Risk 4 becomes more certain or increases in its impact, then Risk 10 will become commensurately more likely or its impact will increase. On the other hand, the risk rating of Risk 1 would be expected to reduce somewhat. Risk 6 has no significant relationship with Risk 3 at all (to see this, look at the two elements that combine these risks – i.e. elements 3,6 and 6,3 are blank).

The understanding required to make each assessment (i.e. to fill in each element of the matrix) uses the knowledge gained when mapping the risks. The risk maps described earlier are of single risks and thus do not directly describe interrelationships with other risks, but a review of the elements and links for each risk allows an assessment of whether any two risks are related. The relationship ‘score’ then entered in the appropriate matrix cell is intuitive and subjective, but the existence of the risk map allows it to be reviewed and challenged. However the real value lies in understanding that although each single element of the matrix may not be very meaningful, as a whole the pattern in the matrix can provide valuable insight.

Consider Figure 10, where each row and column has been totalled. On the LHS, Risk 6 has a score of 4 – which indicates that treating Risk 6 will not have much flow-on benefit in reducing other risks. On the other hand Risk 8 has a much larger score on the LHS and also a low score on the top. It is a strong systemic risk driver – with a significant impact upon risk as a whole, even though it is not itself easily influenced by other risks. Risk 8 is a priority candidate for high-value risk treatment action. Its high score on the top line indicates that Risk 5 is influenced by a variety of other risks, but it is not a driver of risk (it scores low on the LHS). The most cost effective strategy to deal with Risk 5 might be to treat it indirectly. That is, by treating other risks that are related to it.

Risk Relationships Matrix		Risk									
		Risk 1	Risk 2	Risk 3	Risk 4	Risk 5	Risk 6	Risk 7	Risk 8	Risk 9	Risk 10
Total		7	10	8	6	13	7	10	3	9	8
Risk 1	10		1		1	3		2	1	2	
Risk 2	12	2		2		2	3	1			2
Risk 3	4		1			2				1	
Risk 4	12	-1	2	2		1	1	2		2	3
Risk 5	6	1		2	1			1		1	
Risk 6	4		1		1				1	1	
Risk 7	8	2		1	1	2	1				1
Risk 8	14	1	3	1	1		1	3		2	2
Risk 9	5	2				2		1			
Risk 10	6		2		1	1	1		1		

Figure 10. The column and row totals are useful indicators of how each risk relates to the others.

The next step for Risk 8 could be to review its risk map, seeking to ensure that its relationships with other risks are understood. If not, further knowledge elicitation and inquiry may be appropriate so that the risk map can be improved. This would be followed by the design of suitable risk treatment interventions, seeking to maximise both the overall risk reduction for Risk 8 and the flow on benefits to related risks.

Using a Risk Relationships Matrix adds a new dimension to understanding risks. It permits the real contribution of individual risks to be assessed within the overall risk profile, rather than dealing with each risk in isolation. This means that managers are able to work on their risk as a whole, seeking to optimize effort and to maximise benefit.

Tool 4: Using pattern analysis of risk treatments, to provide insight into systemic risk sources within the organisation’s own design

Systemic organisational design weaknesses tend to impact broadly and over time, each contributing to a number of the problems faced by the organisation (Sterman, 2000). When this happens and the risks are analysed and treated, we find identical or closely related risk treatments that apply to a number of different risks. Consider an organisation with a weak system of financial management. This could be expected to result in multiple risks, and the appropriate long term risk treatments for each of those risks would be to work on the same system – that is, the system of financial management.

The importance of this is often overlooked by risk managers, who are focused on individual risks rather than the whole suite of risks experienced by the organisation. They fail to understand that if they look for patterns in the treatment of risk, it can lead to the identification of otherwise hidden systemic weaknesses.

In other words, when the treatments for multiple different risks work on the same system, or on the same part of the organisation, it is evidence that a systemic problem exists. Rather than treat each risk individually it may be more effective to work on that system as a whole, in the process not only treating individual risks but also improving the system more broadly. To be able to use this approach, we first need a method of analysis that allows us to identify patterns within the risk treatment data. Table 1 demonstrates one such method.

Risk No	Risk Treatment	Super-ordinate Goal(s)	Strategy	Structure	Systems	Style	Skills	Staff
	Total	4	16	6	39	12	18	10
1	Treatment 1				1	1		
	Treatment 2		1		1			
	Treatment 3	1					1	
2	Treatment 1		1		1			
	Treatment 2				1			1
3	Treatment 1				1		1	1
....
....
29	Treatment 1			1		1		
	Treatment 2			1	1			1
30	Treatment 1		1			1	1	
	Treatment 2		1		1		1	
	Treatment 3				1	1		

Table 1: Risk Treatment Pattern Recognition Table. Table 1 shows risk treatment pattern analysis against the McKinsey 7S design elements. In this case, the analysis shows that the organisation is creating a relatively high amount of risk from its systems in comparison with other elements of design. The organisation appears to have largely effective goals and vision, and its structures are likely to be generally sound. This analysis provides evidence of where leaders should be focusing their design effort, and where effort might be less valuable.

To use this method, all of the identified risks in an organisation are listed down the LHS of a table, with their proposed risk treatments beside them. Exact duplications of treatments may occur, and are allowed. The column headings are then chosen to reflect related areas of interest. In this case, the headings represent the seven elements of the McKinsey 7S organisation design model. Whenever a specific risk treatment works on (i.e. seeks to change) a specific element of design, a "1" is placed in that column. Since not every risk arises from internal design problems, we should expect some rows to be blank. For example, short term treatments and treatments for externally generated risks may not involve any design element from the 7S model. In such cases there would be no entries in that row of the table.

At the end of the process, summing the columns permits us to see where most of the treatments are being applied, and hence draw conclusions on which elements of the organisation's design are relatively stronger or weaker. For example, the pattern of data in Table 1 indicates that the organisation may have relatively effective vision, goals and structures but that its operating systems have weaknesses that are contributing to multiple internally generated risks.

The use of the McKinsey 7S design elements gives a broad indication of where the organisation's design is strongest and where it may be weakest. However having determined

that many risks are arising (for example) from system design, the same technique can then be repeated, to focus more closely upon a specific area of concern to see if further patterns are visible. This can be done by choosing appropriate column headings as shown in Table 2.

Risk No	Risk Treatment	Sales	Finance	HRM	Risk Mgt	Project Mgt	Performance Management
	Total	2	2	11	3	9	3	
1	Treatment 1			1		1		
	Treatment 2		1					
	Treatment 3							
2	Treatment 1			1	1			
	Treatment 2					1		
3	Treatment 1			1			1	
....
....
29	Treatment 1							
	Treatment 2			1	1			
30	Treatment 1							
	Treatment 2	1					1	
	Treatment 3					1		

Table 2: This table shows risk treatment pattern analysis focused on specific systems. The column headings can be chosen beforehand to match the systems of organisation or they can be built during the analysis itself by adding new systems when they are required. In this case, the pattern analysis indicates that the HRM and Project Management systems are creating significantly more risk than other systems. Rather than try to work on specific risks arising from those areas one at a time, it might be more useful to carry out complete system design reviews on those two systems.

As might be expected, the larger the number of risks and risk treatments identified, the greater the potential to reliably discriminate between systems that are performing effectively and those that are not. Applied effectively, risk treatment pattern analysis draws upon the richness of earlier analysis of individual risks to achieve a systems view of the sources of risk. This makes it a powerful tool for leaders of organisations who wish to identify where to focus their system improvement effort.

A systems model of risk management

Earlier parts of this paper identified that the application of appropriate systemic tools can greatly enhance the effectiveness of risk management. However no matter how good the tools used, systemic risk analysis can only be effective if the risk knowledge upon which it is based is complete and meaningful. How we go about gathering data about risks and their context is a critical determinant of whether risk management is effective or otherwise.

For effective risk knowledge elicitation, a key requirement is that people within the organisation are able to speak openly about threats to their success or to the success of the organisation. Especially for sensitive internal issues, this requires the explicit permission and support of leaders, who must also have the courage to permit their own actions, decisions and behaviour to be questioned and scrutinised. Of equal importance is a broad culture of inquiry and acceptance of challenge, supported and rewarded by the systems of organisation (Barber, 2002).

Based upon these ideas, a systems model of risk management is proposed in Figure 11. The first requirement is for appropriate social processes so that the risk conversations can occur. If effective social processes can and do occur, knowledge elicitation processes can be effective in ensuring that the necessary data and understanding is available for managers to be able to manage risk. Finally, the application of systemic risk tools to the data in order to gain insight into the risks faced, requires both risk knowledge elicitation and effective social processes to be ongoing. Risk management control processes are also necessary in order to integrate risk management activities over time and across organisational boundaries.

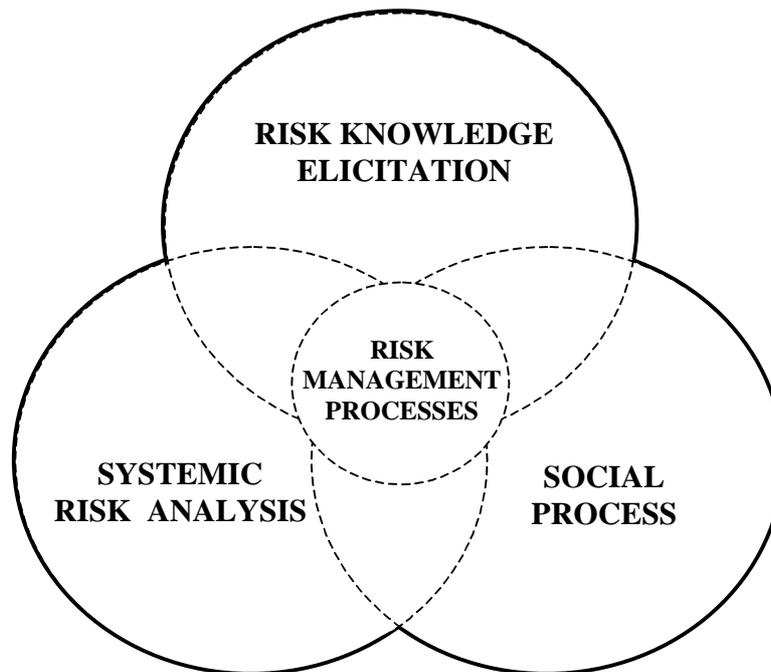


Figure 11. Effective systemic risk management requires that a number of key elements are integrated. Risk knowledge elicitation occurs as a continuing conversation that interacts iteratively with systemic risk analysis. This requires appropriate social processes and a supporting environment, with all three integrated by management using risk management control processes such as risk documentation, risk reporting and higher management review.

In Figure 11, the risk management processes referred to are likely to have some similarities to the AS/NZS 4360 processes shown in Figure 1, but will be more extensive. To work, they need to include elements that drive risk management to occur, and to reward appropriate openness when dealing with internal risks. They should also ensure that risk analysis and knowledge elicitation activities are integrated, and iterative. As part of this, it may still be useful to categorise and classify risks, and to apply risk ratings as described in the Standard.

From a knowledge management perspective, each risk management activity forms part of a framework designed ultimately to deliver risk knowledge to decision makers so that they can make appropriate decisions. As shown in Figure 12 below, the key cycle is one where data is continually gathered, analysed and stored for use. The analysis may be technical or mathematical where this is appropriate, supported by risk mapping and other systemic analysis techniques. As a consequence of this cycle of activities, managers can obtain the benefits described in the list on the RHS of Figure 12.

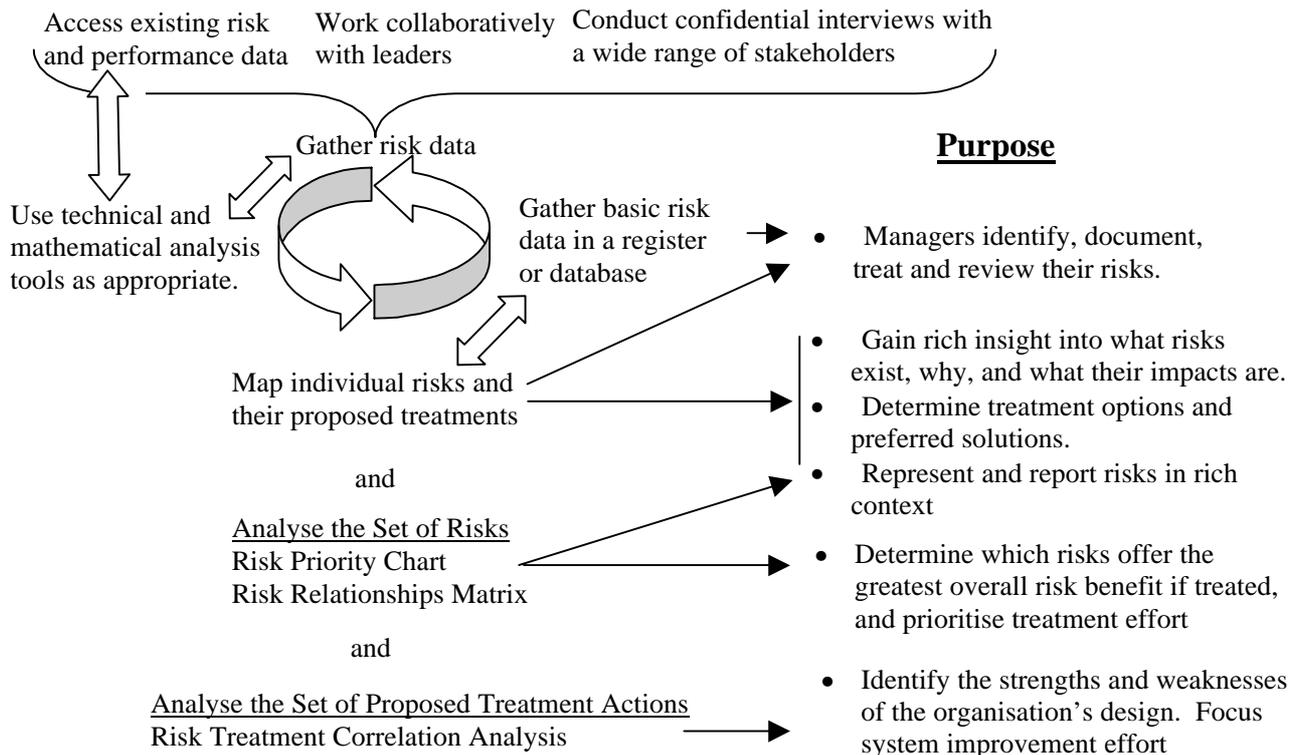


Figure 12: Effective systemic risk analysis is a complex blend of conversation and the use of systemic analysis tools. Technical and mathematical models and tools have a place, but for a large number of risks (especially those that are internally generated) they are not effective. The maximum benefits arise when both risk knowledge elicitation and risk analysis are based upon open conversations over time.

Figure 12 brings together into a single overall system map the key elements of an effective risk management system, where the use of AS/NZS 4360 type processes form but a small (although important) part. It includes systemic risk analysis and risk knowledge elicitation, and it shows how and why risk management activities add value to managers – i.e. their purpose. Figure 12 does not explicitly include the social process part of the systemic risk management model – the need for social processes that permit risk conversations to occur.

Risk accountability transfer is an essential management activity when dealing with sensitive risks

Figures 11 and 12 illustrate that risk management within organisations requires effective social processes as well as effective systemic tools and techniques. However there is one significant issue this does not necessarily resolve – that of the difficulty in documenting sensitive risks. Although effective risk knowledge elicitation within a social process may well allow sensitive risks to be identified, this does not mean that they can be documented and reported easily. Like a standard risk register, a risk map documents risks in writing and this may well be unacceptable to those involved.

However risk maps do have an important advantage over AS/NZS 4360 processes when it comes to dealing with sensitive issues. If an important issue relating to a risk is not documented, it is very likely that the gap will be visible on the risk map. In effect, it would mean that a question implied by the map was not being answered. A link would be missing, and one or more risk sources or risk impacts would be incomplete. Although this does not

reduce the sensitivity of the 'hidden' issue, it does mean that in some cases at least, there will be pressure to identify and deal with the matter.

A more complete response to this problem is possible if we acknowledge that different risks should be the accountability of different managers at different levels. This is true for a number of reasons. Some risks require treatments that are outside the authority of the manager experiencing their impacts, but within the authority of another manager. Others originate within the work area of another manager and are simply best managed by them. Still others are so critical to the organisation as a whole that it is appropriate to elevate their management as a risk treatment strategy in itself. Whatever the reason, the transfer of risk accountabilities to where they are best managed is a necessary capability in any organisation.

If risk accountabilities are able to be transferred routinely, this offers an appropriate and repeatable way to deal with sensitive risks. This is because it is generally possible to identify an appropriate manager even for sensitive issues. For example, consider a case where the (lack of) capability and performance of an individual is creating risk. That person will have a manager to whom they report, and who is accountable for managing their performance. Although it would be socially and politically quite inappropriate to document or to discuss an individual's performance publicly, it is possible for the responsible manager to consider, map and document that risk in a confidential manner. If necessary, the risk can be reported to senior management – provided that the sensitive information is still only visible to those who have line management accountability.

On the assumption that every manager has a manager, and that someone is accountable for every process, system or function in an organisation, it should be possible to assign even the most sensitive risks to an appropriate manager. The transfer of risk accountability should not be confused with the assignment of specific risk treatment actions. The transfer of accountability referred to here is necessarily a complete transfer of the responsibility for understanding, monitoring and treatment of the risk. In organisations where risk management transfer can occur, all managers have an additional option when managing their risks – that is, to negotiate the transfer of risk accountability to a more appropriate manager. In the case of sensitive risks, this mechanism is essential for their effective management.

Summary

This paper discusses systemic approaches to risk analysis and management in organisations, with a focus on practical tools and techniques to deal with soft, complex risks in an uncertain world. The new tools identified in this paper include:

- Risk Maps: A technique for analysing and understanding complex individual risks, based upon an adapted concept mapping approach.
- Risk Prioritisation Chart: A chart designed as a more meaningful way of reporting a set of risks for review and prioritisation by management.
- Risk Relationships Matrix: A matrix that provides insight into which risks to treat to obtain the greatest overall systemic reduction in risk.
- Risk Treatment Pattern Analysis: A technique to help managers identify systemic sources of risk in their organisation.

The paper goes on to propose a systemic model of risk management that integrates knowledge elicitation techniques, systemic analysis, and social processes into a single framework linked by risk management control processes. Finally, the paper addresses the problem of how to

better manage sensitive risks, by suggesting that the key is to use risk accountability transfer to place such risks where they can and should be managed.

Conclusions

Too often, risk management is seen as the identification and management of individual risks, and too often even this is done poorly especially when the risks are soft or complex in nature. For risk management to be effective, we need to understand and deal systemically with the risks we face, not just one at a time but also through analysis of their inter-relationships. Great benefits arise when we use appropriate tools to do this, since we are able to design better interventions both on a case by case basis and across the systems of organisation.

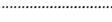
However leaders should use systemic risk analysis for another reason, not directly related to the management of risk per se. Risk analysis provides a unique and very powerful tool for gaining insight into the performance of the systems of organisation, and hence is a tool for continuous systemic improvement of the organisation as a whole.

References

- Barber, R., 2003, Systemic Issues in Risk Management in Projects – Why the AS/NZS 4360 Approach is Not Enough, *Australian Institute of Project Management National Conference Proceedings*, Alice Springs
- Barber, R. and Burns, M., 2002, A Systems Approach to Risk Management, *Proceedings of the 8th Australian and New Zealand Systems Conference 2002*, pp 41-47
- McLucas, A., 2003, *DECISION MAKING: Risk Management, Systems Thinking and Situation Awareness*, Argos Press, Canberra
- Standards Australia, 1999, *Risk Management AS/NZS 4360:1999*, Standards Association of Australia, Strathfield
- Sterman, J., 2000, *Systems Thinking and Modeling for a Complex World*, McGraw-Hill, New York

APPENDIX 1

RISK MAP SYMBOLS

Black block of text	A black block of text indicates an element of a risk map. Elements are linked by one or more lines or arrows to other elements of the risk map, to indicate the nature of the risk relationships.
	A solid arrow indicates that the originating element (cause or source) contributes to the destination element (impact). The reverse does not necessarily apply. If the cause becomes stronger or more evident, the impact is also more likely or becomes stronger.
	A solid line with no arrow indicates an evident strong link between two elements but that the nature and/or direction of the cause and effect is complex or variable.
	A continuous dotted line indicates a link between two elements but that the nature and/or direction of the cause and effect is uncertain, variable or unclear.
(-)	If a (-) is present beside an arrow, the direction or sense of the impact is inverted. That is, if the cause becomes stronger or more evident, the impact tends to become less likely or weaker. The (-) symbol is most often used when a risk treatment is the source.
	A double line across a line or an arrow indicates a delayed effect from the cause to impact.
	A circular arrow indicates a reinforcing feed-forward or feedback loop, with the direction indicated by the arrowheads.
	Directional brackets of the form shown indicate a group of elements that operate together.
<div style="border: 1px dashed red; padding: 2px; display: inline-block;">Risk treatment description</div>	Boxed text (solid lines or dashed) indicates a risk treatment action. This is usually shown in red for emphasis but the colour is not essential.
	A triple lined arrow indicates a direct enabling relationship between elements of risk treatment. The originating element has to be in place first, so that the destination element is possible. This symbol only ever occurs between risk treatment elements. It is usually shown in red for emphasis, but this is not essential.
	A dashed red arrow indicates where a treatment (cause) will impact in the risk map. This symbol is only ever used in conjunction with risk treatments. It is usually shown in red for emphasis, but this is not essential.
	A dashed red line with no arrow indicates a relationship between elements of risk treatment, but where one does not necessarily have to be in place first, for the other to occur.