SAFETY SCIENCE Monitor
Special Edition 2006
Article 1
VOL 10

NeTWork WORKSHOP INTRODUCTION

BARRY KIRWAN,
Eurocontrol, Safety R&D

We live in an age of risks, where we must be aware of the various risks in our environment, and control and manage our personal and societal risk levels. Risk management may indeed be considered to be a desirable attribute of modern industrialised society. Yet it is difficult to compare risks from different sources, and even more difficult to make trade-offs between these different risks. It is therefore difficult, at a personal and societal level, to manage risks in a coherent fashion. The result may be incoherent risk strategies, whether at a personal or societal level. We may be killing more people than we 'need', and spending money where there is little comparative safety return.

The question can be raised whether coherency in risk decisions can be objectively described; if it can whether it is desirable to achieve it and if it is why this seems not to be the case at present. In this introduction an attempt is made to set out the issue, to see why we do not have coherence, and how we may move towards such coherence if indeed we want it. This is done in four steps: first outlining some of the issues from a 'popular' perspective, then in the second part trying to capture the 'elements' and roots of the risk coherence problem from a more scientific perspective. The third section tries to model some of these elements to gain further insight and understanding, and then the fourth suggests some ways forward.

1.1 OUTLINE OF THE RISK INCOHERENCE ISSUES

1.1.1 A risky existence...

When you get up in the morning, risk questions very quickly materialise, at breakfast, for example. How many eggs is it safe to eat in a week? Do they increase or decrease cholesterol? Does the breakfast contain any GM soya, and does it matter? How should I travel to work? If I have to go to a conference abroad - which is the safer method of transport? Energy is used all the time. Which is the safest and cleanest form of energy production? How safe are the vaccinations for my children? What are the impacts of living in the middle of a big, polluted city?

The long list of questions is matched by a bewildering array of statistics, from a range of sources. The statistics themselves come in all shapes and sizes, but all of them seem to say that their product or process is safer than any other, and there is generally nothing to worry about. But they can't all be equal risk, can they? And what of exposure rates? These are not usually factored into readily available statistics, but the exposure rate often dramatically affects the real risk level.

There is much information on risks associated with food, medical matters, and other areas, but the amount and variety of information makes it difficult to analyse these risks coherently. Some members of the public no doubt feel bewildered by the various statistics and counter-statistics, the pronouncements of some experts and other experts who disagree with them, and the 'scientific' pronouncements of government ministers which sometimes subsequently turn out to be completely wrong.

Such confusion does not merely add to our 21st century stressful lifetime experience, what may be called 'risk neuroticism', it also affects our behaviour. Recently in the UK for example, a number of parents decided to stop giving their children certain vaccinations (for 'whooping cough'), because of a small number of cases of serious side-effects. According to the scientists, however, the risks of these side-effects are very small compared with the much more likely risks of serious illnesses if not vaccinated. The point is that the potential result of risk confusion is unsafe behaviour, the opposite of the aim of risk assessment.

Within industry, risks tend to be assessed using metrics that have evolved within that industry via the companies involved and their respective regulators. The various metrics used in one industry are not always compatible with metrics from other industries. Furthermore, as concern over environmental issues increases generally in society, there is a need to consider trade-offs between risks and environmental damage. For example, an environmental campaigner might argue for less air travel due to noise and pollution, but less air travel might lead to more travel by road and rail. There could thus be a net increase in risk (fatalities) and even a net increase in environmental harm. But there is rarely a proper interface between these two considerations (risk of injury/death and environmental harm) even though they both ultimately impact on quality and quantity of life, and even though they are both considered during planning phases for large new installations or transport systems.

It is therefore difficult for society, including the public, to judge and balance relative risks in the environment coherently. This lack of coherence logically not only leads to confusion and maladaptive strategies, but also to more fatalities than society need accept, whether in daily life, or at work. Some of the main problems are described below.

1.1.2 Public perception of risk - bias, or poor information?

Previous work (see chapter 4) suggests that people's risk perceptions and preferences are biased, e.g. against nuclear power, for example, leading to higher criteria in such industries than others. But what if this is more due to a lack of coherent information on relative risks of different energy production systems, for example? Can the public be expected to make a reasoned judgement if the reasoning material is not available? Is the public really biased or does it make another judgement. Do we ('scientists'; administrators, etc.) ignore them (the public), attempt to change them, or allocate risk reduction resources with these judgements in mind? The latter options do not appear to have been much considered.

1.1.3 Sources of risk information

There are numerous sources of risk information. Firstly there is information produced by the manufacturer or producer or corporate body associated with the product or process. This obviously will put things in a positive light. For some of these industries, there may be 'opposition' groups, locally formed or more widely organised, who aim to present a different picture of the risks to the public. Then there are government sources, from various ministries, for food, health, or bodies regulating the safety and environmental impacts of various industries. Such governmental bodies aim to provide public information on risks, but some of them may have a vested or conflicting interest in the activity or the product. Then there are the media organisations, who occasionally aim to highlight new risks, or significant changes in established risks, and who also sometimes aim to expose organisations which are perceived by the media to be mis-representing the risks to the public. So, who do we believe?

1.1.4 Trust in Sources

It is difficult to know which of the above sources to believe, as they may not agree. Even the governmental sources sometimes get it wrong, as BSE in the UK dramatically showed. Eggs are another good example, since over the past few decades advice has ranged from eating many, to only 4 per week to avoid cholesterol problems, to 2 per week (well-cooked!) during the salmonella scare, and now more recently to new evidence suggesting that perhaps eggs actually can reduce cholesterol!. How can a source, whose frequent unreliability is only matched by its continual supreme confidence, be trusted?

And as for the company, the 'anti' lobby groups, and the media, they all have clear interests (if only in selling papers!), so some degree of distortion can be expected.

1.1.5 Trust in Data

Distortion is at least partly enabled by the inevitable imprecision of risk predictions, particularly for very low risks. Such risks include uncertainty ranges, and different sources can quote from anywhere in these ranges, which can themselves be quite large.

Additionally, with low probability events and hazards, new data may cause radical updates to the risk estimates. This can happen in the form of a new insight brought about by an accident, which revealed a new risk or cause of risk, which was previously thought to carry no risk. This has been seen in many industries, such as the airframe industry, the chemical and offshore industries, etc. This means that our risk estimates do not exhibit the property of stability over time, as seen by the public. This is a continual paradox in the risk and safety domains, because the risk analysts do not know what they do not know, usually until we all find out via some form of accident, or if lucky, a non-fatal incident.

1.1.6 Industry risks and uneven playing fields

Some industries are regulated to very strict criteria compared to others. For example, in the '70s, there was a rush to gain North Sea oil and gas in the United Kingdom Continental Shelf. The fatality rates were high (for example associated with commercial diving activities) during this period. However, in the early 80's there was then a considerable push to reduce the level of fatalities (from up to 20 per year to less than 2), which was achieved. At around the same time, the fatal accident rates 'allowable' in various industries differed markedly. Nuclear Power stations were being designed with a risk criterion of 1.10^{-6}, whereas in the chemical and offshore industries, the allowable fatality rate was much higher (1.10^{-4} to 1.10^{-5}). Similarly, whilst large construction projects are often designed to have very low risks for their operational lifetimes, they nevertheless lead to a number of fatalities during the construction and commissioning phases. Yet how is it that we reduce the probability of reactor meltdowns, but can't stop people falling off a ladder?

There is clearly not a 'level playing field' across different industries or aspects of working or normal life. It is sometimes assumed that this is in effect due to the public themselves, who for example may trust nuclear power less than coal-generated power. But even if this assumption about public perception of risk were true, what is that perception based on? Is it based on a true reflection of life cycle risks, or on media-reported events? Every day around the world, 50 miners lose their lives. It usually doesn't make the news. Perhaps if the pubic really knew this, perceptions would change. Even if they did not, at least the influence of that perception would be an informed one, attributable to social values rather than media 'hype' and company and other source 'spin'. On the other hand, maybe the public is aware of these differences in risk, but there are other factors that lead them to accept conventional energy production better that nuclear.

1.1.7 Risk Presentation

Risks can be portrayed in myriad ways, but perhaps the most useful way is largely ignored. A risk implies that a decision has to be made, and implies trading off between one risk and another. Therefore, because the essence of risk evaluation concerns trade-offs, risks are best presented comparatively, rather than absolutely. But such a comparison must not be abstract (e.g. the risk is equivalent to being struck by lightning), but should be contextual, i.e., in the context of the likely decision that has to be made. Travel is a good example. Why not show examples that people can relate to. E.g. 10 international business trips per year around Europe, via plane and via train and via car - with the attendant risks calculated. Or, even better, the composite risks of the complete travel trip, including driving to the airport, or walking to the station, etc. This could than be balanced for instance against time consumed during the trip, costs, relative comfort, etc.

1.1.8 Risk compartmentalisation: divide and ignore?

A major stumbling block seems to be compartmentalised thinking, reinforced by compartmentalised government ministries and attendant agencies and regulators. Many industries are regulated to be as low as reasonably practicable (ALARP) within that industry only. They are not compared to other industries. So, for example, car travel is a big killer, but is very under-regulated compared to other modes of public transport (e.g. if we were applying similar safety culture and regulatory controls from nuclear power to road safety, this would probably result in no longer designing cars that can greatly exceed the maximum speed limit, and driving speeds would be automatically limited in the vicinity of built up areas). This is accepted generally, but not always. Many small residential communities lobby to place speed-controlling devices in their roads to ‘kill speed not kids’. This is a clear example of local rejection of risks that are overall accepted, and also a neat understanding of the difference between procedural defences (e.g. fines etc.) and deterministic defences (speed bumps etc.). But at a governmental level, if number of fatalities from transport risks is the key aspect to reduce, then road is obviously the target for radical measures. A common transport ministry, and less compartmentalised regulatory systems, would possibly help in this respect, by highlighting the obvious and leading to logical action. If the governments need encouragement for such radical considerations, they need only to look at the public health bills associated with car accidents etc., in much the same way that airlines have realised that few airlines have recently financially survived a fatal crash. However maybe the costs associated with further risk reduction, such as slower traffic or the need for expensive types of road construction would outweigh the costs of fatalities and injuries that it aims at reducing. This is a societal and moral question, but one that society does not always get the chance to answer or even ask, because the risks are dealt with by compartmentalised bodies managing compartmentalised risks.

1.1.9 Risk target allocation

This leads on to two critical and often under-discussed aspects: who sets the risk targets and how are they derived? This is a question currently facing the aviation industry, for example. When an industry regularly kills people, the target is easily seen: zero fatalities. But once the industry moves into the low risk region, then targets become less intuitive. Frequently, risk targets are used of a fatality less than once in a million years, for example. Such targets however, do not always seem relevant or realistic. Such stringent criteria can in fact lead to risk analysts either saying a new plant (designed to be safer than existing ones) cannot be built, or else ‘massaging’ the figures to achieve the criteria. Risk analysis then becomes a ‘numbers game’ – the criteria are unrealistic, so the numbers can also be unrealistic, because we ‘know’ the plant is safe. Yet still such targets are used, despite several engineers and risk analysts pointing out that such figures are unrealistic and optimistic in the face of incident statistics and common mode failure phenomena. In such industries, there is also concern over the 'cliff edge' scenario, wherein an industry appears to be safe, but in fact is merely heading for an accident – admittedly very unlikely - whose consequences will be on a catastrophic scale. Such failure modes can be over-looked in the fight to reach the 1.10^{-6} or 1.10^{-7} target criteria for a system.

1.1.10 Risk Gradients: rich and poor societies' risks

In some third world countries, risks are seen very differently. There is not concern over whether eggs increase cholesterol, but rather concern over when, or even if, the next meal will occur. Certain countries are regularly plagued by recurring natural disasters, e.g. flooding in Bangladesh, drought and famine in many countries, malaria, etc. The amount of money spent in the industrialised nations on risk assessment alone, must be seen by some as wholly obscene when natural disasters are still able to claim the lives of millions every year. Taking the idea (or ideal) of risk coherence to its logical extreme, with quantity and quality of life as the goal, then such naturally occurring risks should feature in a 'global risk landscape', and might result in a re-thinking of risk priorities. The recent tsunami in Asia has caused a widespread concern over such large scale natural disasters, in particular the need to gain earlier warning of them and to be better prepared should they recur. Risk tolerance of such natural disasters may be expected to decrease both as a function of globalisation and the ability through technological means to make their risks at least partly manageable, and less dismissable as ‘acts of God’.

1.1.11 Extracting the Elements of the Problem

The problems raised above, and some others, may be conceptualised as shown in Table 1 below, in the form of ‘barriers to risk coherence’. This table is divided into five sections reflecting the dominant ‘origin’ of the barrier or bias against risk coherence:

- Individual level – e.g. biases in risk perception
- Organisational level – e.g. vested interests
- Socio-political level – e.g. uneven playing fields
- Cultural level – e.g. risk gradients
- Physics/Nature of the risk domain (e.g. technology complexity; natural risk areas)

These five levels reflect the sources of the problems or ‘barriers’ identified above.

The barriers to risk coherence may also be considered to be ‘**soft**’ or ‘**hard**’ barriers, in terms of their resilience rather than their character – the former are therefore easier to overcome than the latter.

Table 1: Summary of Barriers to Risk Coherence and Their Consequences

Barrier	Consequences
	Individual level
Biases in risk perception	Two classic examples of this aspect are first at an individual level: voluntary risks are tolerated to a higher degree than involuntary ones (i.e. we may accept comparatively high risks in our own sporting pursuits, at a much higher level than we would accept at work for example). Secondly, certain industries (e.g. nuclear power) have been distrusted more than other equivalent industries (e.g. coal power), despite the fact that coal mining for example kills far more people than nuclear power. This may lead to societal risks being handled unevenly, and to more people dying than is required. The solution is firstly to have better information on comparative risks using the same equivalent metrics, and relating these to exposure rates that people can relate to. Secondly, we need a far better understanding of the roots of these biases, such as aversion to certain risks and 'dread' that may be associated with certain accident types, and the stability of these feelings.
Statistic al lives, and real lives	The current value of a statistical life is of the order of 2 million dollars. The concept of a statistical life arose in part because when trying to do cost-benefit analysis in high reliability industries, where no-one is actually likely to be killed, a value nevertheless needs to be placed on potential or hypothetical lives saved, to decide if further risk reduction measures are justified. The problem is then one of accountancy versus personal value. The price of one's own life to oneself is priceless, in most cases. Similarly, there is no value to be placed on the life of one's spouse or child. In a mathematical sense, risk accountancy assumes a 'monotonicity' of risk value of statistical lives, but personally when it is us or those we know, it is not monotonic. The solution is, perhaps simplistically, to deal with what we actually value when it comes to fatalities, which is human lives and not money. If risk is compared across industries and industrial and public sectors, then equivalent lives saved may be a more useful metric. Then, if one's loved one dies in an accident, then at least there may be some comfort from knowing that someone's life was saved elsewhere – rather than knowing that one's loved one's life could have been saved, but instead the money was spent on an expensive nuclear power plant filter that will actually save no-one.
Risk compen-sation	The essence of this theory is that as risks decrease in one area, risks may be (voluntarily) increased in another. This may work at an individual level, but it is not clear that it should work at an organisational or even societal level – for example, seat belts do appear to have reduced risk, despite risk compensation arguments. Similarly, companies and industries that have decreased their risks in one area do not necessarily increase them in another. The consequences of risk compensation are twofold: first if true, then it limits risk reduction, and target levels of safety should reflect these limits; secondly, it can be used as an argument against risk coherence, since risk compensation in certain risk areas may prevent coherent risk levels ever being achieved. This concept is important in the risk coherence argument, but this author contends that it is not a 'hard' barrier to the achievement of risk coherence.

Organisational level	
Biased sources of risk information	There are few 'clean' data sources, often they are produced or interpreted by bodies with vested interests, resulting in 'playing down' or marginalising the risks, or else over-stating the case with rhetoric. The net result is distrust in sources, often leading to a decision to go with one's own intuition. This means risk management behaviour becomes based on ignorance, political and social persuasion and gut feel, rather than science. There is no obvious solution to the cause of this problem – there will always be protagonists and antagonists – but perhaps the symptoms can be alleviated, via truly impartial and independently–funded risk evaluations.
Non-Bayesian biases:& mis-perception of the nature of stochastic events	There is initially shock when accidental events happen, followed rapidly by denial in some form or other. Two frequent forms of denial are that the event that has happened was a freak accident and hence could not happen again and also that it was very localised and so could not occur in similar systems. Even after some time has elapsed, a third reaction (bias) may be present, in two further forms. Now that it has happened, the chances of it recurring are very small, because (firstly) surely we have learned our lesson now, haven't we? Or else (secondly) one can apply the gambler's fallacy in reverse (that's our disaster now for the next 1000 years) due to a mis-appreciation of the nature of random (stochastic) processes. The solution is partly one of education, e.g. along the lines of understanding the nature of accidents and their complex roots (e.g. pathogens in systems), rather than using a simplistic (superstitious) model of being 'unlucky'. In practice accidents often are symptoms of deeper problems, which are often suppressed further by enquiries and blame-allocation, so the risk remains and even augments, and so becomes more difficult to reduce. We still need to find better ways to learn from accidents.
Disprop or-tionate focus	Industries often tend to focus on rare events (including so-called 'cliff edges'), building their defences around such rare event happenings. Sometimes this obsession even slips into common parlance, e.g. the 'China Syndrome' (the mythical result of a meltdown in a US-based nuclear power reactor is that the core melts and heads due to gravity towards China). There can therefore be a disproportionate focus on high consequence very low probability events. This masks our ability to manage risk more effectively, and can lead to risk incoherence. It also makes industrial comparisons more difficult, whereas every industry has much more in common when considering moderate frequency low fatality risks. The solution is not to focus so much on the nature of the death, but on the fact of the death itself. For example, for a new nuclear reprocessing system being built, the risks may be radiological, chemical, physical, etc., and may occur in construction, operation, and de-commissioning. But most effort will be placed on radiological risk assessment and control/reduction. These other various risks should ideally be treated equally (notwithstanding environmental concerns), shouldn't they? This approach would generate a more balanced risk management approach.

Common metrics	Different industries and sectors use different metrics to manage risks. Thus, risk in transport may refer to number of people killed per year, or to fatalities per km, or per unit trip, or per trip hour. However, parts of the same sector (e.g. transport) often use different metrics, e.g. aviation uses risk per flight hour, whereas car and rail may be expressed per km travelled. So, which is safer, travel by plane or train? The answer is not clear. When comparing across sectors (e.g. transport and energy), it becomes even more difficult. Certain metrics seem to be common, though not always commonly used – e.g. fatalities per year. This should be based on the complete life cycle of the industry, e.g. for nuclear power this must include uranium mining (where most fatalities occur) and construction phases. For risks relating to personal life (e.g. travel, food, etc.) the risks have to be described in terms that are meaningful and personally interpretable, allowing the individual to calculate their own exposure rates.
Risk presentation	Presentation of risks is also potentially misleading due to imprecision and non-contextual references. For example, flying is apparently the safest form of transport. But does this mean that if you spent the same amount of time in a train as in a plane, the train will crash first? (the answer is no). But this is a possible interpretation of the first statement. Similarly food risks or other health risks are often represented with non-contextual references: as much risk of X as of being hit by lightning, or of winning the lottery. But this is not a useful reference, as it has nothing to do with the context. In both cases it would be more useful, though perhaps politically difficult, to state the relative risks, or to be very precise about what the risk is and the 'dimensions' under-pinning the risk estimate. Furthermore, the range should be specified, since there is always significant uncertainty in risk estimates.
Socio-political	
Uneven playing fields	Some industries are more safe, and in a sense more 'penalised', than others. Another perspective is that some industries or sectors are less safe, and perhaps less regulated, than they need to be. When viewed societally across industries, money to save lives is not spent efficiently – cost per life saved differs dramatically in different sectors, so the net result is that more people die than is logically required. This uneven-ness across different industries and sectors also affects the costing structure of these industries and their profitability, as well as resulting in more fatalities than necessary, since increasingly large sums have to be spent to have net safety increases in the 'high-risk' but 'super-safe' industries. The solution is to be able to contrast risks across different sectors of industry, so that risks are truly comparable, and decisions made appropriately. The resilience of this barrier to risk coherence relies in turn on some of the other technical barriers under-pinning its resolution (aspects of metrics etc. described above).

Risk target allocation	The allocation of risk targets is a critical area, since it will drive the whole risk management process for an industrial or public sector. Once set, such targets tend to remain fixed for some time. The problem is how to set targets which are leading to an acceptable level of risk for society, whilst maintaining the delivery of the service or product that is also desirable to society. At the same time, there is always a desire to reduce the risk component of this balance where possible, and this drive is managed by the regulators for the sector. If risk target levels are set too lax, then fatalities will arise which could and should have been avoidable. If levels are too stringent, then service delivery may falter, or else risk assessments may divorce themselves from reality resulting in a numbers game which suggests impossible targets are being reached, without any validation. The problem with this latter result is that there is a focus on increasingly incredible accident scenarios, to the detriment of the consideration of, and protection from, more credible and likely accidents. The solution is to consider targets across sectors, and consider risks according to societal functions, rather than specific industries. Transport and energy production are two relatively clear examples. How many people are killed per year in trains, planes, and automobiles? Where is the money best spent? What should the relative targets be to best reduce the overall number of fatalities?
Trust in data – disputed risks	There are numerous risks that are 'accepted' and understood – they have clearly killed people, and the causes are accepted. However, there are many 'new' (and some old) possible risk sources that are in dispute. New ones are understandable, as it takes time to determine whether something is actually causing harm, especially when the mechanisms for causing harm are novel or not well-understood scientifically. However, this also applies to some well-worn risk debates (e.g. linking leukaemia clusters to certain radiological sites, for example). The bases of the difficulties in such disputes are elaborated below, but the basic problem is one of confusion and inaction. The confusion arises because to those outside the arguments, it seems unbelievable that after decades of research into such issues, causal links still cannot be established, and yet alternative arguments are often seen as weak, e.g. relying on 'statistical fluctuations and anomalies. The inaction is a consequence of a failed agreement (see below) about the causes, and therefore not only what action to take, but also the responsibility and hence funding source of any implied action. Since this chapter is arguing from a 'rationalist' perspective, the only solution is to resolve the arguments. This is discussed further in the following boxes.
Arguing from different premises	A main problem arises, however, as elucidated in Adams (1995), in that the different protagonists in such arguments end up arguing from different premises – therefore there can be no direct comparison or contrast of their 'data' and logic. This is compounded by uncertainties in data, discussed below. The solution is common ground, not always easy to find in socio-political situations.

Media short-term hype	Because the media (newspapers etc.) often hype an event over a short period, because it sells papers, this often gives disproportionate influence on the public perception of risk. This is particularly so when horrific images are shown, as these can create a lasting psychological impact. The solution is one of independence and responsible journalism, and also one of defining relative risks and exposure rates. However, a caveat is that sometimes the media are right to 'over-egg' the risks, which the respective companies and often governments, are trying to play down. There are three biases at work here: availability (of information in memory); base rate fallacy, as the opportunity rates are often not disclosed or known, so the public imagine what these might be, often erroneously by orders of magnitude; and conservatism (of the companies and governments).
Bureau-cratisation	Bureaucratisation generally leads to inflexibility, and, *in extremis*, a lack of rationality. When risk management and accounting systems are developed, this represents state of the art, and is bureaucratised to ensure the developed processes are carried out effectively. However, as time changes, and needs change, the bureaucracy will do what it does best – resist change, no matter how rational the arguments are. Bureaucratisation therefore cements processes, and so is a problem when trying to improve or change risk coherence, particularly when dealing with serious matters such as fatalities and accidents. Since bureaucracy is ultimately connected with power, the only way to counteract it is to have sufficiently high-level and influential studies and bodies who can command change.
Responsibility diffusion	Certainly an attribute of bureaucratisation, though it can function outside bureaucracies, is responsibility diffusion. As risk areas are aggregated and compared, it becomes easier to see that someone, or a committee, will have the ability (and implicitly therefore the responsibility) to see how risk levels etc. should be changed or managed better. This can be an unenviable position (e.g. the current transport minister in the UK has resigned due to a catalogue of problems and accidents under his jurisdiction, and his job has been referred to as the 'poisoned chalice' of the government). Keeping risks apart appears to make them more manageable, and also restricts perceived liability for those managing the associated risks. This is most evident when the relatives of fatal victims of an accident which was clearly 'waiting to happen' cannot find anywhere to focus their anguish, because the responsibility is so diffused within the company. The solution is not clear to this problem, though there is perhaps a useful distinction to be made between responsibility and accountability. In any case, as discussed below, there needs to be some consideration of risks in aggregation and compared to each other, in order to improve overall risk levels.
Risk compartment-alisation	Similarly to responsibility diffusion, risk compartmentalisation keeps risks from being assessed more rationally and coherently. As already noted, it is allowable to kill construction workers during the building of something like the Channel Tunnel (where approximately 15 workers lost their lives), but once the Tunnel is operational, the risk targets are much more stringent. Similarly, assessing nuclear power risks without including uranium mining is quite misleading. Risks for systems, including their system boundaries, and potential for risk export (see below), need to be assessed.

Culture	
Cultural biases in risk perception and risk management	There appear to be a number of cultural biases in risk perception and approaches to risk in general. Adams (1995) distinguishes between four main bias categories, or 'myths' that people generally fall into. The first is nature-capricious (not controllable, leading to a fatalistic approach). The second is nature-benign (this leads to optimism, therefore under-managing risks). The third is nature-ephemeral and fragile (hence cautious, perhaps over-cautious). The last is nature-perverse/tolerant or ecological (here nature or systems are controllable and self-rectifiable within certain limits – the risks to worry about are therefore extreme circumstances where the system's natural or designed defences will break down catastrophically). These biases, based on work by Holling (1979; 1986) and then Schwartz and Thompson (1990) (both cited in Adams, 1995) were originally aimed at natural systems rather than high-technology industries, but seem intuitively applicable to many systems, natural or technological. The existence of cultural differences are a potential impediment to risk coherence in an international domain, but our understanding of them is a potential key to managing international risk communication and harmonisation efforts (see 'arguing from different premises' above).
Risk gradients: rich and poor societies' risks	Risks vary dramatically in some different countries and continents – what is tolerated in one place would be intolerable in another. This is very evident when considering the value of a statistical life discussed earlier – American lives appear to have a higher 'world street value' than say, an African life. Also, the amount of money required to save more lives differs dramatically between certain countries and continents. This may be seen as socio-political, though perhaps it has cultural roots too, since there tends to be a slightly fatalistic attitude [culturally learned helplessness?] about certain risks (e.g. especially from natural disasters), that nothing can be done except to provide aid after the event. One aspect of a solution that may help matters is perhaps to draw up a global fault tree or fatality-cause schedule, and prioritise the main killers in the world today, and then consider how to allocate resources on a global scale. Such a study might be seen as more relevant to World Health, but it is a little strange to juxtapose 3 million Euros being spent to save a life in the West, when 300 Euros might save a life in certain other parts of the world.
Physics / Nature of the Risk Domain	
Risk migration	Risk migration is superficially similar to risk compensation, in that it has the same effects, that of preventing systemic or total risk reduction. Risk migration means that when risks are reduced in one area, then new risks occur elsewhere in the system, usually at the boundaries of the system with another system, or in areas of the system not formerly associated with significant risk. As a hypothetical example, risk reduction focusing on the flight part of aviation safety could lead to a migration of risk to the ground control of the aircraft. The consequence of risk migration is a failure to achieve the desired reduction of risk for the system. Just as many people die as before. The solution is to consider the total risk 'landscape', the system and its boundaries with other systems, and all the possible impacts of risk reduction approaches.

Risk Export	This is an extension of risk migration, but is from one complete system to another. For example, if an airport catches fire, then it is shut down and all waiting and future landing aircraft will be diverted to other airports. This exports risk to these other airports, and also to the trains and cars for people who have had their planes cancelled. If public faith is lost, for example, on one area such as the railway industry, this may export and increase risk to other transport media (e.g. road). The potential result is no reduction of risk, and possible increase of short to medium term risk, because the sector that imports the risk is under excess pressure. The solution is to have across-sector risk evaluations, e.g. for transport using a number of media. This in turn, is only likely to happen seriously, if such sectors are ministered as a single large system, or function.
Reflexivity	When efforts are made to change risk levels, or to improve risks in a domain, the very fact of these efforts may alter the balance of risk. Adams (1995) describes this as a risk organisational version of the Hawthorne effect. This means that influencing risks is a tricky (and non-linear) business. Essentially, the organisation becomes 'self-conscious' about the particular risk type, so that it may reduce in the short term, and reporting on the risk may be affected (either increasing or decreasing disproportionately to the change in actual risk-related event occurrence) meaning that feedback on risk reduction impact may be distorted. The solution is continuous risk indicators, so that trends and impacts can be monitored. Also, risk efforts must be maintained in order to lead to sustained risk impact after the area has moved 'out of the spotlight'.
Uncertainty in data	One of the largest technical issues in risk is the uncertainty in the data, and hence in the risks themselves. As systems rarely remain static, data from the past is continually stretched to predict an uncertain future – but the future will always hold surprises. Uncertainty in data means that there is more room to argue over differences in different risk areas, and less easy to compare relative risks. It allows those with polarised vested interests to argue their cases using the same data, but different ends of the uncertainty ranges. The trite solution is less uncertainty, but this is difficult to achieve particularly with high reliability systems and low probability events, or with elements that interact with human physiology in complex and little-understood ways. A secondary solution is via sensitivity analysis that considers the range and statistical distribution of the data and all its attendant uncertainties. This approach aims to remove the polemic and rhetoric from the arguments, to allow relative risks to still be considered appropriately.

Risk versus danger	As systems become more complex, and as risks are suppressed, the risks become more distant, more difficult to see and perceive. Danger when seen causes a natural defence. Risks of less than once in a million years are however a different matter – it is difficult for the public, and even risk managers to deal with such events that are difficult to comprehend accurately. One solution is to use scale to allow comprehension. The risk for a nuclear power plant reactor having a meltdown may be less than once in a million years. However, if there are, say 300 nuclear reactors worldwide (including submarines), then this equates to roughly one reactor having a meltdown approximately every 3000 years. This is still a very remote figure, but is at least within the sphere of human experience, especially as two commercial nuclear power plant meltdowns have already happened. The use of scale encourages assessments of the risks posed by industrial sectors across international boundaries. For example, the maximum risk of loss of an aircraft accident due to air traffic management failure is prescribed as 1.55 x 10^{-8} per aircraft flight hour. This sounds very low, but when the number of flights in Europe is considered, the risk is much higher than expected (approximately 0.6 per year).

1.3 MODELLING RISK PROCESSES

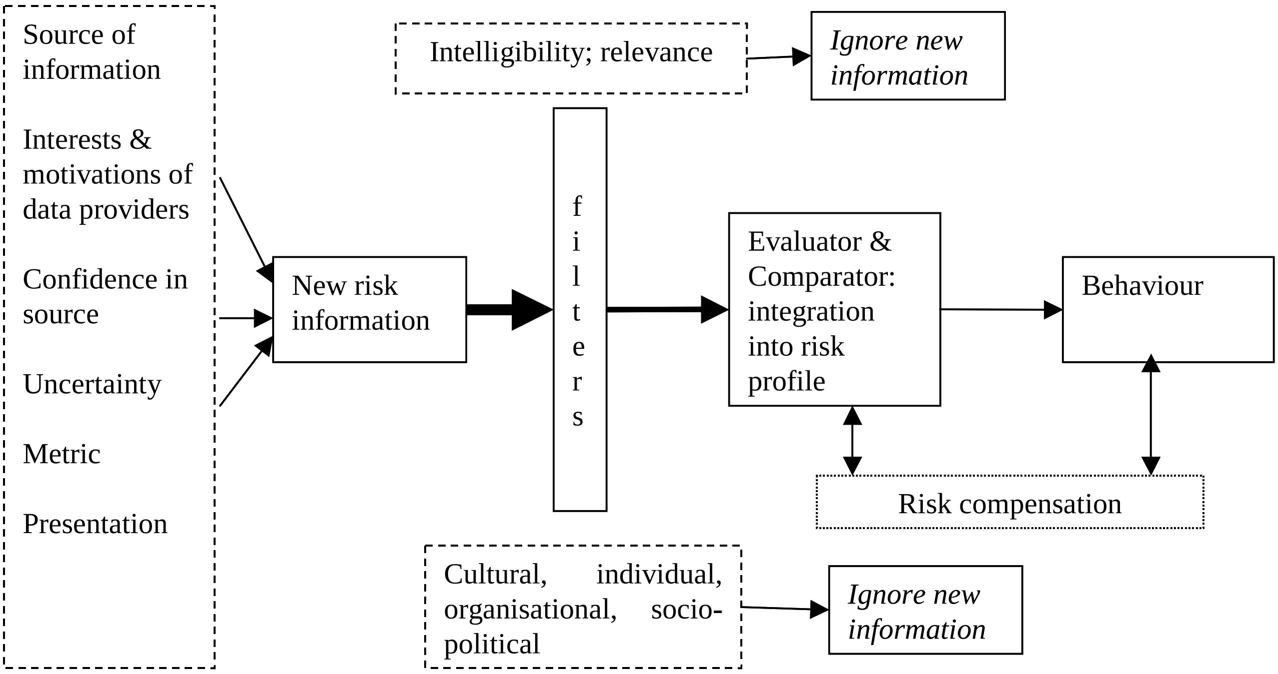

Figure 1. Ad hoc process model of assessing and integrating new risk information

This section attempts to develop some understanding of how we think about risk, in order to gain some insights into how risk coherence can be improved. Two ad hoc models are proposed, simply to add structure to the debate, of risk handling processes, at the individual level, and then at the organisational and socio-political levels for setting risk targets. The first model is shown in Figure 1, and is an ad hoc 'model' of how new risks are assessed by individuals. As new information arises, there are various criteria of interest: the source of information, the interests under-pinning the source, the uncertainty, the presentation, the actual metric, and the confidence in the information.

Various filters may then be applied to screen the information, to give it 'weight', or even reject it. These may be cultural, individual, socio-political, or organisational in nature. Additionally, there is the actual intelligibility of the information, and the apparent relevance to the individual. If the information is not understandable, or appears irrelevant, it will be ignored at this stage. If the information survives the filtering process, the next step is to evaluate the information and to decide what to do about it. This evaluation must have a comparative component, since in today's society we more often trade risks rather than eliminating them (of course this could be my own cultural bias coming through!). Part of this cognitive process is also weighing up the costs of the implications of the risk information, particularly if we must change some aspect of behaviour. The last stage is the actual behavioural change itself (or no change if the risk information is to be ignored, or if the costs of change are perceived to be too high or otherwise unattainable). The phenomenon of risk compensation may affect these last two boxes in the diagram.

This model focuses significantly on the inputs to decision-making, i.e. the sources of information and the filters that can lead to ignoring the information or 'weighting' it prior to decision-making. This explicit focus in the model is due to the author's contention that this is a large part of the problem of risk incoherence. This focus is also due to the lack of knowledge about the later decision-making processes in the model, and how these evaluation, comparison, and trade-off processes actually function when considering one's own 'risk profile'.

The insights from this 'model' are several. First, trust in the data, and clarity of the message, are critical. There is a need for more neutrality in risk data. Secondly, the information must be meaningful and 'integrable' into personal models of risks in one's life. Otherwise, risk comparisons and trade-offs cannot be made. This means that metrics must be 'harmonised' in some way, so that risks can be aggregated and compared, and traded, in one's personal 'risk profile'. Thirdly, the filters are very important, but little understood, especially the cultural ones. Adams (1995) makes it clear that the four cultural stereotypes are of course gross simplifications, and different combinations of their attributes are possible, and they may shift or change with time and events, even in the same person (this means that cultural barriers are in fact soft, not hard barriers). Fourthly, we could probably all do with some help in structuring our 'risk profiles' – we manage our risks heuristically on a personal level, although in our jobs, the risks are generally more structured, even if they are tightly bounded. What is needed therefore is a 'risk profile prototype', which would integrate the risks in modern living, from all sources

The second 'model', again rather ad hoc and simplified, concerns the setting of risk targets. This very simple model suggests that risk targets evolve during the evolution of the sector, and thus are the product of a 'sector-centric' approach. The risk management cycle works well, but remains inward-looking compared to other industries etc.

The alternative is to have a more integrated risk approach considering risk targets from whole sectors (i.e. related or 'bordering' industries) and across different sectors. This would lead to more risk coherence. In practice this needs three things. First, we need wider-scope risk assessments (to show the bordering interactions and impacts). Second, comparative risk assessments are required across functional sectors of life (e.g. transport; energy; etc.). Third, a 'total risk picture' is needed, to consider the risk priorities in modern civilisation.

1.4 WAYS FORWARD

The risk area is not one that is easy to grasp, as it has many facets, and is a complicated field, dealing with complex multi-dimensional problems in an open social and cultural context. Nevertheless, risk coherence is a logical goal, even if it turns out to be highly difficult or to achieve. Some possible ways forward are therefore discussed briefly below.

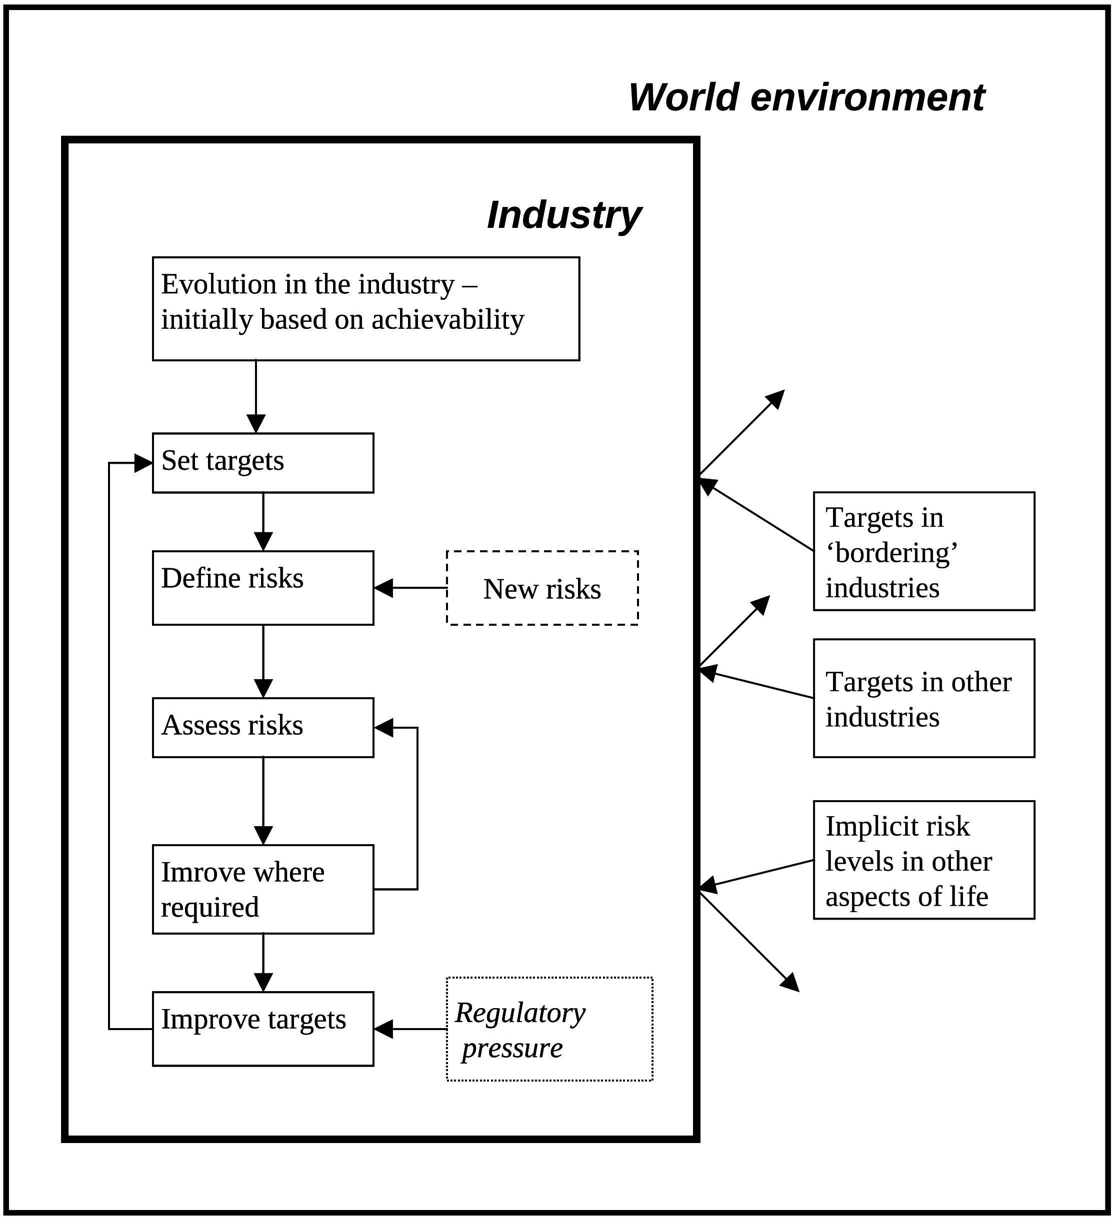

Figure 2. Target setting model

Fatality data – a focus on fatality data across an industry (i.e. so that there is some fatality data even in high reliability companies) is needed. This does not denigrate the utility of incident data, but avoids the complications of risk approaches based on arguable criteria (e.g. when is a near miss in aviation really serious?).

Risk boundary assessments – Risk assessments need to consider their boundaries, and how risks impact bordering industries, and the potential for risk export and import.

Intra-sector analysis - there need to be within sector, but across industry, risk assessments, e.g. contrasting different transport risks, or different energy risks. This in turn necessitates a harmonisation of metrics and determination of exposure rates.

Global fault tree – a project should be carried out to develop a 'total risk picture' at a societal level, and also at an individual level. The aim should be to discern the true priorities at both these levels.

Individual risk profiling – there should be a project to develop a prototype 'individual risk profile'. This in turn can then evaluate the practical utility of all the available risk metrics that purport to inform us about our risks, to see if they can be integrated into the profile.

Ministerial and regulatory responsibilities properly bounded – as ministerial responsibility becomes more 'social-function' oriented, e.g. transport; energy; etc., the need to rationalise the industries within these sectors increases, especially when accidents are occurring. However, there is a need for the regulatory bodies, which are often under a national structure, but working in their own industry compartments, to talk more to each other, and compare notes. The regulatory bodies are perhaps the best place to start the process of achieving risk coherence.

1.5 OVERVIEW OF THE BOOK

The following chapters address various aspects of risk coherence from a number of industrial and socio-political perspectives. First, there is a section on risk management and how it has evolved to where we are today. This is achieved by three chapters on the development of risk frameworks in the Netherlands, the UK, and the US respectively. The Netherlands considers how various events in a number of industries have shaped risk management policies and processes, and public expectations. This chapter highlights the key role that high-profile accidents can take in influencing risk policy. The UK chapter focuses on the nuclear power industry, one of the most mature risk management systems in industry, and warns against too much reliance on quantitative approaches to risk informed decision making. The US chapter considers 'democratic' and legal aspects of risk regulation, and the causes of dissatisfaction with current risk management and regulatory regimes, and calls for harmonisation of risk regulation approaches.

The second main section considers the bridge between experts and stakeholders. This is structured firstly by describing a simple model of interaction between these two groups, albeit noting importantly that experts in one field are inevitably the 'public' in another field where they are non-expert. This chapter in particular considers the psychological and sociological under-pinnings of risk perception, and argues for more study on public risk perception in order that more effective risk communication can take place. Implicit in this chapter and echoed in others, is that the risk communication between the two groups must be a 'two-way street'.

The next chapter considers 'risk consent' in the medical field, highlighting problems of uneven playing fields, risk presentation and risk compartmentalisation. It argues for a shift from doctor-centred to patient-interest-centred risk informed decision-making. The third chapter in this section then considers risks in the transport sector in particular, and succeeds in making some useful and informative comparisons firstly between different transport sectors, and then between other walks of life as well. This chapter argues from a dispassionate view towards risk coherence in transport, but highlights specific problems that can arise in trying to apply similar criteria to different modes of transport (e.g. rail and road). The final chapter in this section considers radiation risk.

The third main section then focuses on the aviation and air traffic management industry, as an example of where risks appear to be relatively well-controlled. The first chapter in this section gives a historical account of how risk management has evolved in Air Traffic management, and how this has in particular led to the development of Target Levels of Safety (TLS) in this industry. This account is instructive also in the pitfalls of deriving such risk criteria in very safe industries where fatal accidents are rare. The chapter also proposes a 'Risk Relatedness Matrix' as a potential solution to considering risks across different industrial and social concerns (e.g. global warming).

The next chapter in this section, also in the context of aviation, considers the issues of risk when using Commercial off-the Shelf (COTS) products, showing the special risk considerations that industries need to be aware of when using such them.

The last section considers special challenges in the area of risk coherence and more generally risk management and regulation. The first chapter returns to the subject of comparability of risks, showing some statistical comparisons of different industries in terms of fatalities, and considering more complex metrics. This leads into the second chapter on occupational diseases, where it is more complex to parameterise risk These two chapters show the necessary complexities of risk metrics when not considering simply 'fatalities', and call for a general 'index of harm' for such risk areas. The third chapter focuses on controlling risk (risk reduction), and in particular how to determine risk-related improvements from the mine of data that industries can produce on incidents and accidents. It proposes a new approach based on so-called 'Swiss Cheese' modelling of accidents due to failures in explicit and implicit barriers. The final chapter then considers risk across system boundaries, and trade-offs that may have to be made at or across such boundaries, as well as potential risk export from one transport mode to another. Two case studies in rail and air traffic management are used to illustrate the problems and potential solutions.

The final section concludes on the desirability and practicability of risk coherence, and some research areas that could usefully be pursued to improve our understanding in this area of societal and individual risk management.