UNIVERSITY OF Southampton

The Risk Radar and the View of Risk

A Centre for Risk Research

Discussion Booklet

FOREWORD

It is a pleasure, as Director of the Centre for Risk Research (CRR), to welcome this latest CRR discussion document on the *Risk Intelligence* theme. What we present here is perhaps best described as a book or booklet because it is more lengthy and substantial than the previous two CRR discussion documents from 2017 and 2020 that dealt with risk intelligence. This new discussion document seeks to balance both theory and detailed practical guidance.

The work is inspired by a project which the authors, Dr Alasdair Marshall and Dr Felipe Costa Sperb, undertook together in 2018 for *Munich Re*. On that project, Felipe was heavily involved in the detail of designing a risk category solution capable of housing the diverse risk contents of SEC 10-K Corporate Risk Reports. In the present document, he shares his experience in the form of very specific design principles. Complementing this, Alasdair calls upon two management metaphors which, taken together, are intended to capture much of the challenge presented by organisational risk management. He theorises interdependency between the *risk radar* and the *view of risk*, further arguing that both require ongoing critical reflection. More fully, he argues that the *view of risk* requires careful design and maintenance, which can be achieved by following the design principles outlined by Felipe. Such improvements, he argues, entail that the *risk radar* can then improve its scanning of the risk environment, requiring further improvements to the *view of risk*, and so on in what can amount to a continuous benevolent cycle.

When explaining his theorising to me, Alasdair made the point that a professional approach to organisational risk management is partly a question of *intellectual mood*; that is, certain *feelings* can help to guide us through the challenges at issue. Specifically, a dedication to assiduousness in scanning the risk environment should, ideally, combine with a recognition that the *view of risk* always remains corrigible and warrants critical scrutiny. This same combination of resolve and humility is, notably, found in critical realist thought, and in realism more generally. It is also discernible in work by Dylan Evans on what it means to be *risk intelligent*, and it is similarly fundamental to Tetlock and Gardner's psychological profile of the *superforecaster*.

It then becomes interesting to consider that when risk management is approached with a critical concern focused directly towards the *risk radar* and the *view of risk*, it is being approached with exactly the professional attitude which risk management needs.

I recommend this text to all risk management practitioners, students and academics, and hope that it will provide much food for thought.

Dr Ian Dawson
Director of the Centre for Risk Research

Risk management is a young profession and still exploring the body of theoretical knowledge and evidence to support its practices and approaches. One of the first steps that organisations will take is to look around at their risk landscape, which they soon find to be constantly and rapidly changing, baffling in its complexity and full of blind spots. And yet an effort still has to be made. In the honourable tradition of borrowing from other disciplines, the concept of deploying a risk radar to help detect threats is an attractive one. We welcome this research paper as a robust unpacking of what this actually means, and whether it has sufficient intellectual coherence to be a foundation for good practice.

Carolyn Williams CMIRM, ACII
Director of Corporate Relations, Institute of Risk Management

AUTHORS
Dr Luis Felipe Costa Sperb
Dr Alasdair Marshall

CONTRIBUTING EDITORS
Dr Ian Dawson
Professor Ming-Chien Sung
Dr Mario Brito
Fragkiska Roussou

Co-published in 2020 by the Centre for Risk Research, Southampton Business School, University of Southampton.

INTRODUCTION: SOME THEORETICAL PRELIMINARIES

The Risk Radar as a Scanning Device

Two previous (2017, 2019) Centre for Risk Research (CRR) discussion papers have already discussed the risk radar in complementary respects. Both considered risk radars as producing risk intelligence. It is certainly worth consulting these before reading the present booklet. However, this is not strictly necessary, as we do summarise these two earlier papers' contributions in our section below entitled Scanning and Risk Intelligence. The basic proposition to be developed throughout the present booklet is that the ideal-typical organisational or project risk radar is, metaphorically speaking, a device whose essential purpose is to scan risk environments in order to improve what we will call the view of risk. More fully, the relationship we envisage is one of mutual dependency on an ongoing iterative and cyclical basis. More specifically, we will theorise risk radars as requiring some organisational of project view of risk in the first place, in order to scan the risk environment, so that that same view of risk can then have its mettle tested and be improved. We envisage the resulting benefits as fanning out across all risk management process stages, perhaps most importantly better risk assessment and control, but also of course better scanning by the risk radar.

In other words, we consider risk radars as constantly using, testing and striving to improve their associated views of risk. Essentially, then, what we are offering is a simplifying managerial perspective based on close examination of these two interdependent concepts, to help conceivably anyone involved in the risk profession think through what managing risk should very fundamentally involve, or at least what an essential part of it should involve. We think there is potentially much advantage to be gained when everyone who contributes to risk management in an organisation, or on a project, comes to appreciate that scanning for risk, as well as using, testing and improving the view of risk, are all distinct and yet logically interdependent goals whose balancing and co-achievement they should seek to remain focused on throughout their risk management participation.

To first broach the idea of the risk radar by conceiving of it as a device is also helpful, we would suggest, as that term connotes some application of design ingenuity and imitable technical knowledge to enhance management practice. Correspondingly we will be exploring the key technical design issues that we think require consideration for getting the best possible risk radar up and working to full effect in its scanning, testing and improving roles. Organisational and project circumstances, and associated risk management scope and purpose, can vary enormously, of course. We will therefore work towards proposing five basic design principles for the view of risk whose relevance can be considered on an individual organisation or project basis.

The View of Risk as a Risk Ontology

Our approach in the present booklet is to regard any risk radar as being only as good as the view of risk which it scans with. Accordingly, our suggested design principles will relate directly to the view of risk. Risk radar design is certainly conceivable as a distinct topic which we might have focused on more directly that we proceed to do below. Perhaps the best way to address that separate design principles challenge would be to adapt management research dealing with optimising allocations of scanning responsibilities to managers within organisations, considering also issues of adaptive fit against a simplifying typology of different environmental contexts - Smith & Raspin (2008) provide a very good example of this - so as to recognise some nuances of scanning for risk in particular. Those design issues, important as they may be in their own right, are for another day. We will give them only very limited attention, and in passing. Instead we will be more concerned with the designing the view of risk that is used by all managers who scan the risk environment in any circumstance, quite irrespective of how their scanning responsibilities have been allocated, or indeed what basic form of management practice their scanning takes.

Having cleared up those possible grounds for confusion, we can now further elaborate our general approach as follows. We theorise the risk radar as always primed with a view of risk which directs the search for risk, and which always then tests and improves that view of risk in the light of whatever combination of new theoretical insight and new information it finds. In other words, we regard the risk radar as expressing intelligent and focused risk anticipation, and as amenable to a combination of theory-and-evidence led testing, modification and improvement. The progress we make in this third CRR discussion document is therefore to regard the risk radar as not just gathering and churning out risk intelligence information, but, more fully, as converting that information into integrated risk knowledge, formed as a precise and intelligently structured view of risk, which it then harnesses for more effective and efficient functioning. Not wishing to dogmatise on the finer detail of what the view of risk is best understood to mean for all organisations and projects, we offer now some further thoughts beyond these basics, in order to invite debate.

When we refer to the view of risk, we are primarily concerned with use of language for risk description; that is, with specifications of risk categories and their associated contents whose interrelationships can be displayed visually in numerous ways, especially to show what risks are perceived to matter within various risk exposure contexts across an organisation or project. In other words, we are concerned with the challenge of summarising risk within language prior to its formal assessment and evaluation within risk management processes - and indeed we are more generally concerned with the challenge of rendering risk communicable for conceivably any other management purpose. We therefore conceive of the view of risk as a joined-up risk ontology; that is, it expresses what risk is fundamentally perceived to be, within some organisational or project context. By centralising the risk radar and its associated view of risk, then, what we offer in the present document can reasonably be called a risk ontology focused theory of risk management.

Obviously, in order to form each view of risk there must be at least some prior screening, resulting occasionally in exclusion. In other words, there must be some pre-assessment, in some form, at least some of the time. What gets left out might conceivably say a lot about the minds of managers, and about the organisational politics of what they feel they can and can't talk about. Hence, sometimes a view of risk might be most effectively critiqued with reference to its exclusions and blindspots. We are nonetheless principally concerned with the fundamental risk ontology, or set of risk objects, that has passed any initial screening. Our view of risk concept therefore comprises the mental objects for risk that are perceived to matter, to require estimation or forecasting, and indeed which get culturally primed, perceived, communicated and deliberated on various collective management levels in organisations and on projects, conceivably with reference to a wide variety of risk attributes such as controllability, naturalness, equitability or voluntariness of exposure, etc.

We need to recognise that risk ontologies come in very different forms. For more complex risks in particular, to segregate risks as distinct things is to reduce complexity and simplify for various practical management purposes. Hence risk is commonly expressed using many different forms of words. For example, in common language, risk can mean threats, hazards, vulnerabilities events or consequences. These forms of expression are all limited in scope in the precise sense that they all dip into the causal story of a risk at different capture points. Then there is the classic risk ontology of 'X risk', e.g. reputational risk, people risk, operational risk, etc., to consider. Such abstractions supply basic organising categories which, taken together, build towards exhaustive coverage of risk environments and are arguably indispensable for that purpose. Or, drawing on MacKenzie's (2014) theoretical study on risk index development, it can also be argued that probability distributions provide basic descriptive detail, simultaneously visual and statistical in form, in order to render risks comprehensible. Then there is also yet another approach to risk ontology which we will discuss shortly in our section concerned with risk realism, where causal theory, events and statistics are differentiated. And before that, we will explain in our section below dealing with cognitive diversity that both abstract and concrete forms of risk knowledge provide yet another way to nuance risk ontology. Similarly, there are also more implicit risk ontologies, for example where Boardrooms constantly signal risk and its management using the language of finance, or, increasingly, of sustainability, resilience, high reliability, etc. Clearly, there is much to choose from when seeking more and better ways to think about risk.

We have no wish to contest the possible value offered by all of these highly varied approaches to descriptive expression for risk. In fact, there is perhaps a strong argument for developing a more exhaustive classification system in order to be to support risk managers with a checklist of contributing perspectives which they might consult in order to think more thoroughly and critically about risk. However, the key point we need to make here is that we regard the view of risk as what has become settled and recognisable within organisational or project discourse, as the favoured simplifying mental map to the risk environment that is deemed to matter. Acknowledging the validity of all of the above forms of expression for risk, we regard the view of risk as the organisational or project risk ontology solution which allows individual risks to be paraded together within a coherent, or at least partially coherent, overall expression, and which is suitable for carrying forward, conceivably for a range of management purposes. Such purposes may involve ad hoc consideration of individual risks, or they may involve more routine recourse to some more thorough, structured and conveniently visualised view of risk, for example within decision support and executive information management systems, enterprise resource planning systems, or indeed risk management and internal control systems.

The Risk Radar and the View of Risk: metaphors to be used cautiously

While still introducing our subject, we think it is also worth recognising a problem for risk management professionalism which, arguably, deserves to be aired far more often. That problem relates to risk management in general, and the risk radar in particular, as special cases of the more general management reliance on metaphor. Rigour in thought demands that it is better to say what things are rather than what they are like. However, this is very commonly not always possible. Gareth Morgan is well known for contributing to management theory by emphasising management's heavily reliance on metaphor to make sense of the complex dynamics of organisational life. For example, Morgan (1986) theorises the likeness of various organisational dynamics to biological organisms, mechanisms, political systems, etc. His theory invites a broad generalisation whereby the greater the complexity of what we seek to understand, the more we need inventive use of not just one but layers of metaphor to reduce that complexity down into accessible and actionable management narratives. We would suggest that Morgan's arguments extend, and indeed grow more compelling, when we consider management's need for metaphor to address the even greater complexity reduction challenges associated with organisations at risk, where it is not just the organisation itself, but also a potentially enormous variety of influences upon it, that require capture within simplifying risk and risk management narratives. Recognising the complexity of organisations at risk, we next elaborate the point that, quite unsurprisingly, risk management terminology is replete with metaphors which can subtly influence thought in unintended ways. And yet this brings with it a particular problem. Risk management metaphors are rarely recognised as such. Entrapment within habitually used and professionally sanctioned metaphor, inviting cognitive failure which can potentially lead to management failure, is therefore a real possibility. Yet risk management professionals often fail to recognise critical skill in their effective use as an important part of the professional oeuvre they should be striving to develop, both in themselves and in those they support.

The risk radar is a prime example of one such metaphor. We remain within the metaphor when we represent it as scanning for risk within risk environments. The view of risk is a metaphor too. It suggests use of eyesight to ascertain - or perhaps check upon or confirm the presence of - threats or hazards in the physical world, thereby inviting use of further visual expressions such as risk landscape or risk panorama (perhaps especially for external risks), or risk panopticon (perhaps especially for internal risks). Indeed, it might also be argued that risk identification is another very similarly limiting visual metaphor. So too is foresight. By connoting a passive sensory experience, these visual metaphors all arguably conspire to draw attention from the value which proactive investigative effort can contribute to risk management. Yet how many times in the many guidance documents that refer to risk identification as a risk management process stage, are problems associated with its metaphorical expression addressed? Taking stock, the key introductory point we need to make here is that we are about to embark on a discussion of two powerful metaphors which, taken together, constitute a highly distinctive albeit limited perspective on what risk management fundamentally involves. We therefore need to emphasise that advocacy of simplifying management metaphor and advocacy of critical faculty for their effective use must go hand in hand.

Risk Radar use requires cognitive diversity

A further important preliminary point we should make concerning the risk radar is that a particular axis of cognitive diversity is likely to prove beneficial for its effective operation. Following Gregorc's (1984) mind styles model, a very general distinction can be drawn between abstract and concrete mind styles. We will have cause to refer to this distinction throughout our discussions of the risk radar. Risk management can be theorised as requiring an abstract mind style insofar as there is advantage in risk anticipation becoming more imaginative and thorough, both in terms of the larger and the longer view, owing to a recognition that some risks might matter in the abstract, even where there is a relative lack of concrete evidence for them either having occurred or being in the making. Of course, a concrete mind style will tend to complement this, because there are also sometimes very strong cases to be made for cutting through all the abstract possibility to focus on the detail of evidence and experience. Every view of risk is perhaps best developed as a balanced settlement of the two.

Arguably, all risk management teams are likely to benefit from the dynamic tensions that arise with this cognitive diversity. A concrete mind style may focus foresight towards giving regard to the detail of what has happened in the recent past and is happening presently. An abstract mind style may focus foresight more towards future possibility that is relatively novel and unrelated to past or present experience. Working together, a more complete prudence, considered as a mental acuity for foresight that balances out appropriate consideration of past, present and future (see Marshall et al. 2006, p.13), arguably becomes possible - at least where there is effective dialogue and the complementarity of the contributing mind styles is understood and valued. Notably, then, the argument being made here encompasses benefits of cognitive diversity for historians, for commentators on the present, and for those engaged in foresight. All face similar challenges of balancing abstract theory with concrete evidence – and both historians and commentators on the present are, arguably, far more useful to those engaged in foresight, to the extent that they have succeeded in doing so well.

General implications for risk radar use might be expressed as follows. Primed with abstract risk knowledge emphasising certain risk categories and contents, the risk radar undertakes the search for concrete evidence, in the form of data and information, which can then be used to test the mettle of the abstract view of risk. As a complicating factor, the risk radar can then be further conceived as scanning for concurrent abstract views of risk (e.g. by considering the risk priorities of other organisations, regulators, academics, issuers of guidance etc.) which might enrich its own abstract view, even in the absence of concrete evidence to suggest that the newly acknowledged risk is anything more than an abstract possibility for the organisation or project in question. The risk radar metaphor may be construed as starting to wear thin at the point where we regard risk radars as communicating directly with one another – and yet here the extended metaphor of a risk radar stack (or array) can save the day.

Hence, going beyond metaphor, what we can conclude is that the (sometimes subtle) blending of abstract and concrete mind styles may matter greatly, and their relative contribution may be well worth reflecting upon, both when scanning for risk and when testing and modifying the view of risk. We will look in more detail, later on, at the related notion that rigorous discussion of risk requires careful mental partitioning between abstract risk knowledge of what might happen, and concrete risk knowledge of what has happened or is in the making.

In overview, then, we will discuss risk radar practice (vis à vis its associated view of risk) as a metaphorically simplified expression of recommended best practice. In the light of the foregoing introductory discussion, we can now summarise our advocacy of such practice as pointing the way forward towards clearer, simpler, more effective and efficient, better integrated, more participative, more cognitively diverse, more prudent, more communicative, more critically reflective and more proactive risk management practice. In other words, we envision a critical mass of benefits all co-emerging through a range of broadly

participative risk management practices inspired or improved by a shared management narrative which gives careful regard to the risk radar and its associated view of risk - and which, to reiterate, we earlier called our risk ontology focused approach to risk management.

Promoting Realism in Risk Management

However, one remaining important introductory point is as follows. The management narrative based on the risk radar and the view of risk, whose widespread use we advocate, can also be understood as providing simple terms of reference for leveraging the philosophical, psychological and methodological concerns of realism within risk management practice. We contend that when managers are trained to think in terms of there being a risk radar, dedicated to updating and sharpening the view of risk, then their thoughts will, in effect, come to appreciate the centuries-old realist concern that the world and the mental representations we make of it are two very different things - and that considerable effort, in particular involving critical self-reflection aimed at correcting bias and false assumption, are necessary for slowly building a more accurate view of the world.

Theories of critical realism in particular, tend to combine an ontological realism which is committed to understand the complexities of external objective reality, with an interpretivist or relativist epistemology, which is concerned to look critically across various forms of knowledge so that more of that complex reality becomes revealed. Crucially, these concerns lead critical realism to what is termed a depth (or layered) ontology which differentiates between three different types of thing (or ontological realm) in order to develop a more integrated and causal understanding of the world, namely: (1) the realm of generative forces or mechanisms at play in the world, (2) the realm of events which these produce and (3) the realm of statistics, which by recording regularities for events can help reveal underlying causes. For critical realism, then, rather than think within a flat ontology, we can know more of the complex world by disciplining the mind to cross-reference between these three realms in order to test and advance knowledge.

There are some important - and also very straightforward - implications arising here, for how we can think more causally, and more thoroughly, about risk so as to reveal more of its complexity. In fact, this depth ontology is perhaps best summed up as nudging us to focus more rigorously on complementary aspects of risk which are very obviously important. Take, for example, a new home owner who is surprised when their property on a tidal riverside is flooded (realm 2 event knowledge). This may well lead them to think more critically about the environmental report which very recently reassured them that the river in question is likely to flood only once every 200 years (realm 3 statistical knowledge). To make sense of this clash of realm 2 and realm 3 knowledge, the new home owner may then begin to think more critically within causal-theoretical realm 1 about what might explain the apparent discrepancy. Perhaps their thoughts will turn first to global environmental considerations such as rising sea levels, or to more local considerations such as reduced river dredging or increased pluvial run-off upstream; or indeed perhaps to some combination of global and local factors. Turning to consider practical managerial applications within organisations, arguably it is financial risk management, which naturally focuses towards realm 3, that stands to benefit most from this simple technique for thinking more thoroughly about risk. However, it can also be argued more generally that this simple depth ontology might easily help any manager reflect on any tendencies to focus excessively within any one of the three realms. Perhaps, moreover, it should be in everyone's toolkit for critical self-reflection vis a vis any risk.

Hence we can conclude that a layered risk ontology, which juxtaposes causal-theoretical, event-based and statistical forms of knowledge, for any given risk, seems likely to make a very worthwhile contribution indeed to any risk ontology focused approach to risk management - such as the one we present here. Accordingly, the expression risk realism becomes particularly useful. This philosophical position might be summarised as calling attention to the need for careful and cautious applications of depth ontology when reviewing and improving the view of risk, preferably in conjunction with an energetic determination which believes that more is always likely to be necessary, and an optimism which anticipates that further careful effort will yield not just changes but improvements. We advocate for our conceptual innovations of risk radar and the view of risk, then, partly on the basis that their effects are likely to include the inculcation of this risk realism as we have very briefly summarised it, in organisations and on projects.

A very final point about risk realism which we also consider useful in allowing us to introduce our basic subject matter is as follows. It relates back to that combination of ontological realism with epistemological relativism which distinguishes critical realism, and whose main practical implication for risk management is perhaps best summed up as a call to cautious optimism for pressing on relentlessly with operating the risk radar and improving the view of risk. Academic studies variously treat risks as external powers or potentialities which exist in the real world, and as more subjective knowledge propositions. Common sense dictates that both are extremely important. Risk realism in organisations or on projects can usefully be viewed as pivoting on a managerial resolve to match the two as precisely as can reasonably be achieved, recognising and respecting both the difficulty and the permanence of that challenge. We would suggest that our combined risk radar and view of risk management narrative offers simplifying and practical ideation for addressing precisely that challenge.

Our Approach: five basic design principles for the view of risk

We will introduce our more detailed subject matter gradually, working from the above preliminaries. First, we will outline the ambitions we have for CRR discussion documents in general, as a means to establish some further context for appreciating the contribution offered by our simplifying metaphors. Having then looked in some detail at what risk radars and their associated views of risk are, and at key theoretical and practical issues they raise, especially for what we will describe as risk intelligence work, we will at last propose and critically discuss five basic design principles for the risk radar's view of risk. The final conclusion will then emphasise that participation in risk radar practice should serve to sharpen participants' critical faculties and mental dexterity for discussing and managing risk in many different management contexts. To be clear, however, this being a discussion booklet comprising new and exploratory thinking, aimed at risk practitioner and risk academic discussion in particular, we will work through these various stages at a very leisurely pace so as to address a broad range of issues which we think deserve far more attention from risk professionals and academics. And we are particularly keen to receive feedback on anything we say that proves thought provoking or helpful.

SIMPLIFYING AND INTEGRATING RISK MANAGEMENT

Centre for Risk Research (CRR) discussion documents aim to align novel theoretical insights to practical suggestions for enhancing risk management practice (as broadly conceived) in all sorts of organisational and project contexts. What makes them discussion documents is their critical and exploratory approach whereby they address important questions and invite feedback from practitioners and academics in risk management and related specialisms. One problem they all address, to varying extents, is that of how risk management can contribute to management as a whole. More fully, they all seek to advance our understanding of how best to align/embed/integrate/hybridise risk management, beside or within all other specialised management domains - recognising that precise use of terminology matters when exploring the interface between risk management and management more generally. In the present booklet, our discussion of the risk radar and its associated view of risk are intended to offer innovative terminology which captures only limited aspects of risk management practice itself, but which may nonetheless prove helpful within any broader context of management practice wherever risk matters.

A further consideration, underlying the various suggestions made in CRR discussion documents, is that all managed entities have limited resources and therefore require management solutions sufficiently clear and straightforward as to elicit that broad participation upon which mature risk management depends. The risk management profession has become cluttered, we think, with much pertinent guidance pulling it in multiple directions and heaping ever more conceptual overlays upon the same risk management issues. Recognising the resulting problems, and working towards simplifying solutions, has become, we think, an important challenge in itself. Correspondingly, the simplifying metaphors discussed here as constitutive of our risk ontology based approach to risk management, are intended as a focus for design imagination in the creation of clear risk management solutions to the often highly complex organisational or project challenge of risk management's relatedness to management as a whole.

THE VIEW OF RISK: EMERGENCE, AGGREGATION, POLITICISATION, LEADERSHIP

To reiterate, the basic proposition we develop here is that the risk radar comprises heterogeneous management activities which, in the light of newly sourced theoretical insight and all sorts of risk information, continually reproduce, through testing, modification and improvement, the organisational or project view of risk. And to further reiterate, we are concerned only with the view of risk that has been selected as mattering for risk management, prior to its quantification within formal risk management processes. Of course, managing anything competently, in the absence of at least some view of risk, i.e. some risk ontology solution, would be impossible. Applied views of risk are, of course, inevitably ubiquitous throughout organisations and on projects. Overall organisational or project views of risk might plausibly be regarded as continually emergent through clashes of contributing diverse views, leading to the ascendency of whatever is most sharp or accurate, most actionable, most politically or reputationally astute, most popular, most preferred by leadership or by certain shareholders or by other stakeholders, etc. The points of emphasis and underlying selection pressures for each unique and transient criterion mix may, of course, vary considerably - and furthermore these may, sometimes at least, not even be noticed at all.

Recognising the above as an aggregation problem which might be solved in various ways, we suggest that production, testing and revision of the view of risk might be enhanced greatly through some dedicated leadership and coordination which raises awareness that there is a view of risk in the first place, which can be designed to meet various specifications, and which everyone can help test and improve. All risk management participants can provide such input, we further suggest, simply through critical reflection on their own preferred view of risk, giving regard to what it comprises, what instrumental value it offers in various respects, and indeed how and why it may contrast with alternative views reflecting the needs, perspectives and experiences of others.

Obviously, organisational or project views of risk can be read directly, to some extent at least, from risk registers, probability-impact grids, heat maps, risk-based internal control designs and a range of other variously formed lists, indices, measures, visualisations, etc. These can all provide helpful touchstones for anyone seeking to apply some critical thought to their own preferred views, perhaps also giving regard to further prevailing or clashing views. However, what we mean by the organisational or project view of risk cannot be fully equated to these simplifying representations, at least not where these are created in the absence of that rigorous and critical reflection we which think designing the view of risk should entail. The more rigorous and critical, and indeed principles-based the design effort, we would suggest, the more it becomes appropriate to refer to organisational or project management as (purposefully) taking a view of risk. We would suggest that the categories and contents offered within commonplace risk lists and visualisations are rarely as precise as they might be, were rigorous principles-based design processes followed in order to improve them. Hence our chief technical concern, which with which we conclude the present booklet, is that of proposing core design principles which risk leadership might both use, and promote more general and collaborative use of, for proactively taking a view of risk in any organisation or on any project.

The typical reality in organisations and on projects is arguably that risk lists and visualisations will tend to offer at least some value as judgement and experience-based aggregations. Most risk management participants may be perfectly happy with these and may not think to criticise them at all. Some may even naively assume that the risks can simply be read out from the risk documentation, such that no further critical scrutiny or exercise of risk imagination is required (see Marshall & Ojiako, 2013). Such lazy intellectual foreclosure might be considered a problem to the extent that it might dull and desensitise ongoing risk radar use. An important role for risk leadership, arising in response, may be to explain why effort aimed at further improving the view of risk is always likely to be worthwhile, and to point the way forward towards how this might be achieved.

Consider, for example, that such views might often exclude some valuable socially distributed explicit and tacit knowledge of risk that arises through diverse professional encounters with it, and which is likely to be embedded within very specific risk narratives that are distributed widely around organisations and on projects. The overall view of risk might therefore be considered corrigible by engaging with the complex aggregation challenges this diversity brings, and indeed perhaps through some finely tuned reflection on the possible benefits of some limited pluralism. Such improvement effort might entail exploring the practical-managerial or financial implications of any differences found (e.g. for risk communication or for efficiency in insurance coverage or for other forms of risk control).

Furthermore, each contribution towards, as well as each successive iteration of, the view of risk, can also be considered artefactual in character. Participants in the organisational or project enterprise of improving the view of risk may strive hard to be technically rigorous in pursuit of scientific objectivity, and yet they are arguably likely to reveal at least some socio-technical shortfall, whether intentionally or otherwise. This possibility arguably merits careful contemplation, in the form of reflective humility that is vigilant towards management or leadership hubris, in order to steer the view of risk towards a more accurate mapping of what is actually in the risk environment. Arguably risk leadership can do a lot here by setting an appropriate risk-cultural tone at the top. This might happen, for example, simply through some leadership acknowledgement that a professional attitude towards risk management should entail being seen to take sensitive issues of socio-technical manipulation relating to organisational politics, corporate reputation and image maintenance very seriously - at least to the limited extents that circumstances allow.

THE VIEW OF RISK: STRUCTURE AND LOGIC

We thus regard each organisational or project view of risk as amenable to development towards clearer, as well as more precise and accurate, articulation, partly in order to raise critical awareness of the basic subject matter of risk management: the risk ontology; which is to say, the risks themselves. However, to progress our discussion by further clarifying what we mean by risk ontology, we need to address issues of basic structure, and of logical interrelationships between the parts of which the structure is comprised, especially if we are to advise on articulating the view of risk in any more practical detail.

In 2018, the authors of the present document were involved in a project with Munich Re Insurance Company (see Costa Sperb et al. 2018), which involved mapping out the risks that are reported within SEC 10-K corporate risk reports. The purpose was to capture the diverse risks within a layered category solution with a small number of general risk categories at the top and a greater number of specific risks at the bottom. The selections from that report's risk mapping, which we set out below, may serve to illustrate structure and logic for a view of risk. To be clear, this is mainly for illustrative purposes, yet It does establish some fundamental design considerations which we explore later in more detail when we formalise our five core design principles.

Firstly, here are our first and second layers. Notice that our top layer comprised four broad areas of corporate concern, spanning operational, commercial, strategic and financial performance domains. Consider in particular that each equates to a distinctive and broad area of risk exposure. Our experience, based on a combination of iterative fit testing and reflection on the high level risk categories that are commonly used in corporate risk management, convinced us that focusing within risk ontology at the capture point of general risk exposure would work best for our top layer. More fully, we discerned that this would provide us with the most viable high level category solution under which we could then sub-categorise the thousands of diverse risks that were mentioned, or covered more implicitly, in the SEC 10-K reports.

Notice also how our second layer (comprising the bulleted items under the four risk exposure headings) focuses towards more specific areas of corporate concern, sensitising us to specific points of vulnerability within businesses in particular. Or to put this a slightly different way, in contrast to very general top layer specifications of what is exposed to risk, it is discernible that on the second layer there is a tendency towards specifying more precisely what can fail or be lost.

Operational capability	**Commercial standing**	**Strategic management**	**Financial performace and viability**
– Facilities and processes – Sourcing, storage and distribution – Human capital	– Competitive environment – Intellectual property – Product conformity	– Business strategy – Corporate growth – Governance matters	– Operating expenses – Credit – Finacing and capital management

Figure 1. First and Second Levels

Facilities and equipment	**Sourcing, storage and distribution**	**Human capital**
– Business interruption risk – Manufacturing competence risk – Cyber and information security risk – Business process implementation risk	– Supplier risk – Input risk – Distribution risk – Inventory risk	– Key worker risk – Talent and staffing risks – Stakeholder health and safety risk – Employee bargaining power risk

Figure 2: Second and Third Levels

Next, the following table shows how we were able to then further sub-categorise the three main points of vulnerability for operational capability in order to create a third layer. Interestingly, this is where we discovered the classic risk ontology of 'X Risk' as providing the most appropriate form of words for expressing what we felt needed to be drawn into focus here, to more fully unfold the detail of what the 10-K reports were telling us. Notice how the third layer contents all seem to address what is vulnerable to failure or loss in even more precise detail, particularly through more organisational-functional specificity. They stop short of actually specifying the associable threats or hazards themselves; nonetheless they concisely summarise sufficient organisational-functional context as to direct the risk imagination squarely towards at least some of the more obvious ones.

We also discerned that beneath this third layer there are, generally speaking, too many specific threats or hazards to classify in fine detail. We did create what we called a 3.5 layer based on limited and sporadic scope for application of still more specific categorisation principles which focused still further back along the risk-causal chain towards what might variously be called the hazards, threats, root causes, sources of risk, perils, etc. For example, we found it sensible to subdivide industrial unrest risk into specific forms such as striking and working to rule. However, we also became aware, when trying to populate the 3.5 layer with content, that such content is often likely to be transient. In other words, it became clear that the 3.5 layer was likely to require a lot more frequent revision than the layers above it. This is perhaps best understood as the layer where views of risk destabilise and become too awkward to maintain, at least for some corporate purposes.

In summary, what we learned is that risks in the SEC 10-K reports were best categorised in layers by moving from generalities to specifics, and at the same time from exposures through vulnerabilities to further risk ontology capture points that move further back along the risk-causal chain to where specific associable threats and hazards start to become clear to the risk imagination. Essentially what results is a view of risk that is more fully a view of why the risks matter in terms of their risk-causal traceability through the vulnerability and exposure categories with which they are associated. We can illustrate this simply below.

This permits us to now summarise what we mean by structure and logic for any given view of risk. Drawing on the above Munich Re project illustration, we envision overall organisational or project views of risk as structured by category and as cascading down through various sub-categories so as to reveal not just logical coherence by conformity for listed content within each category but also the underlying logic of the cascades as they thread from small numbers of general risk exposures at the top to more numerous and very specific risks at the bottom. Arguably, these threads can also be regarded as aligning relatively fixed and inflexible top-down C-suite/Board views of risk to more profuse and changeable bottom-up views of risk that can arise with some limited autonomy within the specialised functions and professions. They also have risk-causal significance; that is, by allowing exposures at the top to be aligned towards root causes at or below the bottom, they permit the causal story of every risk to be told with reference to why it matters within its organisational or project context.

To be clear, these are just basic considerations for understanding the issues of structure and logic that perhaps need to be addressed for any view of risk. Conceivably, various risk ontology capture points, or, more vaguely, areas of focus, can be used to differentiate the category layers. We have certainly not exhausted the full range of potentially viable ones. These might further relax the temporal frames for the risk causality to be covered by the view of risk, by sub-categorising aspects of the risk environment from which the risks emerge (at the front end), or by considering short or long term consequences or outcomes (at the back end). Exposures and vulnerabilities are, however, what we would regard as necessary basic categories requiring logical alignment within every view of risk. It is hard to imagine a coherent view of risk that doesn't specify and seek to align these basic inclusions as a preliminary to accommodating larger numbers of more specific risks. Essentially, these permit the materiality of the risk to the organisation or project to be expressed.

Fundamentally, we view precision in defining the various contributing risk categories and contents as deserving far more attention within the risk management profession, especially given that when aligned together they may form a helpful visual map, of sorts, for understanding where broadly participative risk management begins and ends within organisations or on projects.

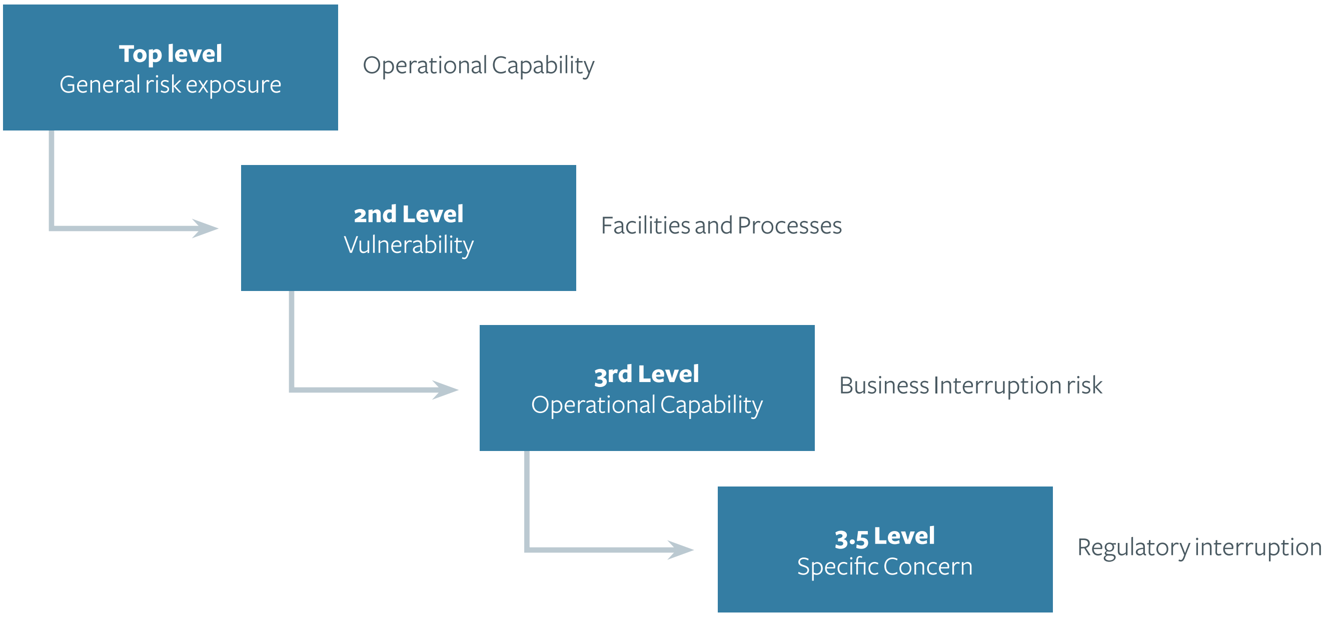

Figure 3: A risk category cascade

THE RISK RADAR METAPHOR: A CLOSER EXAMINATION

Our main technical goal, as explained earlier, is to set out and critically discuss five basic design principles for organisational or project risk radars, focussing attention very specifically towards the views of risk that are both captured by the risk radar, and which point the radar towards the risk environment with specific expectations and dedicated scanning resources. Before that, however, the present section will look more closely at what practices the risk radar metaphor encompasses. After all, it is worth reiterating that the risk radar is first and foremost a very rich organisational metaphor - and our purpose is to explore it both theoretically and in terms of its practical management implications. As with all metaphors, the risk radar metaphor reduces complexity. We conceive of it as a particularly helpful metaphor for designing, scrutinising, aligning, simplifying and otherwise improving often highly complex and diverse organisational practices whose interrelatedness, in the absence of the metaphor to serve as a frame of reference, might be far less apparent. Of course, we do need to look, at least in outline, at what these practices are.

To begin to explore what management commitment to risk radar design and use may actually entail, we would emphasise that anyone contributing to the total stock of organisational or project risk knowledge, in any way, or drawing upon that stock of knowledge, for any reason, is potentially a risk radar participant. Our reasoning is as follows. Risk is ubiquitous in organisational and project management. Anyone who has cause to give regard to risk, within any management context, should construe this role as entailing critical reflection on the abstract views of risk which are available for purposes of contextualising and structuring the discussion of the risk issue. They should also take into separate account any concrete experience of risk which the evidence is pointing towards. Following through, they should also be concerned to highlight any awkwardness of fit between the abstract view and the concrete reality, in which case the risk management issue arising would also become a risk communication and organisational learning issue.

A key point arising here that we think there may often be advantage in managers prioritising their understanding of the risk radar and the associated view of risk in their professional understanding of what risk management fundamentally involves as a critical management practice. Even if they know little else about risk management arrangements within their organisation, we would suggest that if they understand these two metaphors then they should at least be able to discuss risk with a basic level of critical awareness and competence, recognising in particular the clashes and complementarities of abstract and concrete risk knowledge that require critical attention within risk discussions. As with all good metaphors, of course, what matters is that we do not get trapped within them but instead derive inspiration from them and specify the real issues they raise. As we have suggested previously, that is another important critical awareness issue for risk radar use.

In order to further unpack the risk radar metaphor, for purposes of clarifying the required management commitment, we focus next on two very closely related key questions: what it does and how it does it. Looking at the what issue, we will inquire deeper beyond our earlier generalisation whereby risk radars require to maintain and improve a view of risk if they are to detect and protect against risk. The main line of argument here will be that risk radar use must be, to some extent at least, an ongoing exercise in social construction, where perceptual priming and modification in the light of risk experience influence how contextual understandings of risk are captured within simplifying and logically aligned risk categories/contents for various practical management purposes. We will emphasise that, crucially, these social constructions must delineate risk itself against some broader risk environment context. Looking at the how issue, we will emphasise that many different management practices can contribute to the risk radar by counting as scanning for risks within risk environments. Indeed, anyone who enriches the management information base may, perhaps inadvertently, also be contributing to the stock of useful risk information. Nonetheless, we regard information gathering to be just one part how risk radars go about differentiating risks themselves from their associated risk environments. As we have just indicated, there is a further question of how to socially construct risk so as to differentiate it from its risk environment that needs some focus. Hence we inquire a little deeper into this ongoing epistemic project, centring on the production, testing and improvement of risk knowledge.

Accordingly, the following two sections will now emphasise that risk radar use very much comprises thinking about, over and above looking for, risk, which is necessary given the nature of what is being looked for in the first place. Essentially this distinction between the two activities recognises the risk that being scanned for as something that is both objectively real (and hence searchable for using risk identification techniques) and as composed of subjective knowledge propositions whose testing and improvement require human capacities of intellectual rigour and mental dexterity, all rooted in a deep understanding of the risk ontologies that are available.

As we will now see, our more detailed inquiry into these what and how issues for risk radar use will draw attention to to some surprisingly important-and-yet often-neglected issues which should be of interest to any risk manager seeking terms to help them stretch their professional understanding further.

WHAT RISK RADARS DO

Thinking within the terms of the metaphor itself, it makes sense to conceive of organisational or project risk radars as devices that scan for risk within what we might further call risk environments. But following the metaphor's distinction between the thing itself and the place where the thing can be found, what, precisely, are these risk environments? How might we theorise them to differentiate them from the risks that emerge within them - if, indeed, emerge is the best term to use? And what does scanning for risks actually mean, when this question is posed within the problem context of how best to conceptually differentiate risks from risk environments? These are, we would suggest, important questions for anyone seeking rigour in their approach to risk management - yet they are rarely asked. Here we take a closer look.

One way to circumvent this challenge of inquiring into what scanning for risk within risk environments might mean in richer theoretical terms, would to frame scanning very narrowly as an information search, such that the risk environment is simply regarded as comprising sources of risk information. These sources might be people, places, databases, etc. On that view, scanning would simply mean looking for and accessing information from such sources. However, we defer our consideration of such activities until the next section, where they are recognised as an obviously important component part of how risk radars achieve their basic goal of producing a view of risk, differentiated from its associated risk environments.

Accordingly, here we focus on scanning for risk as an activity that seeks to increase the total stock of risk knowledge, such that scanning becomes fundamentally concerned with producing meaningful understanding of risk over and above all contributing information about risk. Taking our approach, scanning for risk within a risk environment is not just searching or looking, it is also very fundamentally about striving to know risk, through various forms of experience with it, and through incremental familiarisation and learning - all of which might be incrementally and progressively resolved within and communicated through a view of risk.

Considered within this broad context of ongoing knowledge production, and simplifying down somewhat, scanning the risk environment for risk can therefore be regarded as comprising activities that strive towards not just more and better information, but also more and better theory, in order to produce meaningful understanding of what the risks are. We would suggest that this, by necessity, entails further constructing or at least making working theoretical assumptions about the broader risk environment within which risks become meaningful and attain distinctiveness as objects of managerial attention.

Consider, for example, how world leaders and corporate leaders may very differently conceive of risk within broader risk environments. Michaelson (2020) observes that at the January 2020 World Leaders' summit in Davos, Switzerland, the findings of the World Economic Forum's 2020 Global Risk Survey were presented, as were findings from the 23rd PWC Annual Global CEO Survey. The World Economic Forum findings emphasised 'climate action failure' as the top global risk. THE CEO survey, by contrast, emphasised 'low CEO confidence', linked to uncertainties spanning regulation, trade and economic growth – which are all, of course, factors that are further bound up with global climate change concerns. In effect, then, what the world leaders at Davos were being asked to view as the risk itself, the surveyed CEOs seemed more likely to regard as secondary risk environmental context for making sense of the top global risk summarised as 'low CEO confidence'. This illustrates well that what counts as risk, as opposed to risk environmental context, can be very much a matter of perspective, whereby risk ontologies might focus anywhere along the causal chains, or within the various reductionist layers of explanation which they deem relevant. Recognising that such differing perspectives might arguably sometimes very subtly influence, and be influenced by, how managers scope their management responsibilities and options, we think it is important that risk managers address these differentiation issues.

As we discuss below, there may be many explanatory frameworks to choose from when delineating risks within risk environments. Working within any given theoretical framework (or hybridised sets of frameworks) for understanding the dynamics of the complex social world, one particularly important question becomes, to what extent can we usefully specify causal context for understanding the narrower causal stories of the risks we are interested in? To be clear this might matter greatly for what the view of risk ends up looking like. Broader views which reach out further into the risk environment may provide richer causal-explanatory contextual detail, while narrower views, limited more to particular management perspectives, may be more focused towards practical management calls to action. Hence there are balances to be struck, between contextual understanding and actionability. These issues are considered in more detail below.

The idea of scanning for risk within the risk environment might set us thinking, in the first instance, in some combination of functionalist and adaptive terms about ongoing reflexivity between organisational/project entities and all sorts of threats or hazards which are only comprehensible with reference to at least some broader environmental context. This might often entail differentiating between the risk environment, theorised as fundamental, and risk itself, theorised as particular.

Taking that approach, we might focus our concerns towards certain highly permanent structural factors deemed to exist as fundamental properties, or indeed trends, comprising the risk environment, and then separately theorise particular threats themselves as triggered or accelerated, within those broader contexts. Or, to put this slightly differently, we might say that in every view of risk there is a balance to be struck between recognising more fundamental risk, supplying valuable and sometimes essential explanatory context, and recognising more particular risk, which is more likely to be actionable (i.e. controllable).

Hovering ambiguously between these two levels, however, there many different concepts that may be considered useful for describing risk itself, by some, or for describing risk environmental context, by others. These include propensities, mechanisms, trends, practices, tendencies, conflicts, volatilities, circumstantial hazards, systems, social equilibria, contingencies, fundamental (or other significantly scaled) events, causes, sources, antecedents, determinants, precipitants or influencers of risk, or perhaps management problems, issues, weaknesses or vulnerabilities which can each associate to various individual risks over time. All of these concepts are potentially very useful for helping to socially construct the causal stories of individual risks, at the explanatory context end, entailing that they all might, at times, be incorporated to at least some extent within views of risk.

Hence there is clearly a problem when it comes to clarifying what scanning for risk within risk environments actually means. Leaving aside simple exceptions, for example where the risk environment is a clearly defined place such as a factory floor, what emerges from the above discussion is that artistry in social construction is likely to matter greatly when discussing risk within the context of any specified environment for it. Risk is a term used so promiscuously in common language as to equate to conceivably any of above synonyms whenever a negative - or indeed positive - influence on organisations or projects becomes discernible. The difference between a risk, and an environmental circumstance from which it becomes distinct within the mind, might arise through nothing more than a decision to socially construct a boundary at a single link in a greater causal chain or nexus, the fullness of which anyone seeking to achieve a causal understanding of risk ought to be interested in.

Hence a perceived risk might emerge as conceptually distinct where a manager, recognising the need for simplifying management narratives, has chosen to simplify complex causation by declaring a root cause (another simplifying management metaphor) where the risk is construed to emerge in its distinctiveness from its environment. Such specifications may be made lazily and arbitrarily, or perhaps with recourse to the but for or unbroken-chain-of-events test associated with the insurance doctrine of proximate cause, or indeed by drawing on that doctrine in other ways such as by taking into account concepts such as dominance and concurrency as a means to better reflect subtleties of causation within the complex social world.

Ultimately, the choice of risk categories and risk names is likely to determine what understandings of risk environment, and risk, shape what is meant by scanning for risk in the first place, thereby rendering risk management meaningful as a practice that goes out into the world, with purpose and direction, in order to simultaneously discover and theorise what is there. Consider, for example, that one way to map the risk environment is by using the widely used PESTLE acronym and its many variants. Clearly, we can think within or across the PESTLE categories about very diverse risks. In other words, the categories supply a sensitising framework for risk identification such that the random creativity of risk imagination is focused and compartmentalised more effectively. However, there are many further ways in which we can usefully sub-categorise risk environments. The simplest and arguably most important of these is the logically complete distinction between the internal and external context of the organisation, recourse to which

can help counteract tendencies to overlook internal risks that may be reputationally embarrassing, or external risks where prospects for gathering useful management information may be relatively poor. The view of the risk environment as comprising a combination of natural hazard and human threat is also commonplace; however the boundary arising with this distinction is arguably more ambiguous. Here we encounter a particular problem whereby risks often do not respect the category boundaries we draw. It would be foolish to partition human risks from technology risks and then neglect human-technology interaction risks.

More fully, it can be shown that some broad risk category architectures, designed with category juxtaposition in mind, are able to provide rich environmental context within which the risk imagination can be sensitised towards specific risks that might easily remain unanticipated otherwise. Marshall et al. (2019) offer a simple four element high level risk matrix design for managing major airport projects so as to avoid disastrous openings at commencement of operations. Drawing on a literature review of relevant readiness and commissioning factors by Al-Mazrouie et al. (2020), they call attention to the need for the abstract risk imagination to traverse four key operational readiness categories as follows:

FACILITY READINESS - the availability and appropriate functioning of physical assets

PEOPLE READINESS - human skill, commitment, and adaptability

TECHNOLOGY READINESS - availability, back-up, infusion within operations

ORGANISATION READINESS - relevant structures, relationships, processes and cultural influences

Facilities, people, technologies and organisations are, of course, different abstract contexts for thinking about risk. The key point made by Marshall

et al. (2019) is that all four categories will inevitably matter together and in their various interrelationships as complex airport projects move towards commencement of operations. Hence recourse to the above four high level risk categories can be viewed as bringing into play the necessary mental dexterity for a rich understanding of what might go wrong. What this illustrates well, more precisely, is that just as the panorama of a physical environment can be scanned with binoculars, so too a complex risk environment can be scanned using various combinations of the above four conceptual lenses. This is fundamentally how scanning gets done for risks in risk environments.

More problematic still are categorisations that simply refer to different types of risk. For example, regulated vs non-regulated and financial vs non-financial risk can each be construed as providing logically complete coverage of the different sorts of risk that can populate a risk environment. However, once more, risks will not always respect these boundaries and hence there are important category design implications arising which we will come to shortly. Similarly, we can also render more of the risk environment visible through a series of further categorisations which employ inventive use of abstraction to specify what is at risk: systems, processes, assets, technology, etc. (perhaps also considering successful interaction for combinations of these). Clearly it is important to design the best set of these, for any given overall view of risk. Arguably, a well chosen set of exposure categories can inspire a more nuanced (and yet still structured) understanding of risk in its environmental context. Admittedly, to map a risk environment using exposure categories is to raise once more the important question of what that environment is perceived to be in the first place. To map something (i.e. what is out there) by offering a partitioning of something else (i.e. what is exposed to what is out there) is arguably not ideal. However, this necessity does reflect what we have already said about the dual nature of risk itself, whereby each risk proposition is a bridging concept; risk always (perhaps tentatively) claims a match between some objective external reality and more subjective knowledge about that reality. Recognising this point, we think that careful consideration of exposure categories is likely to prove fundamental for creating an effective overall view of risk in most cases. This is because it is reasonable to categorise things on the basis of why they matter. Each risk exposure category will express a separate and distinct reason for being concerned about risk (environments) in the first place.

Deepening this point that risks might usefully be categorised by why they matter, it is perhaps worth adding that a structured view of risk might also arise on a per objective or per performance/success factor basis. There are, however, some miscellaneous further theoretical approaches that might be taken which model risk exposures within broader risk environment contexts. For example, conceiving of organisational or project environments in terms of system or equilibrium metaphors can also prove helpful for revealing more of risk. Working within Parsonian structural-functionalism, for example, managed entities become systems which need to adapt to their system environments on a teleological hierarchy of levels; more specifically, they must do what it takes to survive, to pursue goals, to integrate internally, and to develop culturally so as to integrate within broader institutional fields. Or, working with an equilibrium model, for example, an organisation might theorise itself as acting in such a way as to perturb some social equilibrium and thus face risk expressed metaphorically as re-equilibriating force. Clearly, all such rather unconventional approaches to risk categorisation may at times bring something useful to the table. Each may have very different implications for where risk radar effort is focussed, and hence for what risk management is perceived to be about in the first place. Arguably the place of social theory within the risk management knowledge set deserves to be considered more often, in the light of these possibilities.

Our concluding point, then, is that when risk radar participants reflect on the purpose of the risk radar, in the light of the issues we have covered above, addressing the challenge of specifying some partition between risk and risk environment may prove very helpful. In particular, this may focus the mind on the necessity, artistry and frailty of social construction, and on the corresponding need for relentless pursuit of the view of risk that works best.

HOW RISK RADARS SCAN THE RISK ENVIRONMENT

The term scanning extends the risk radar metaphor, inviting us to reflect on the diverse management practices of which it can comprise in the real world. Here we return to the idea mentioned at the beginning of the last section, that scanning for information from information sources is a big part of what is at issue. The contributions offered by various forms of proactive intelligence gathering work might be accentuated here. Everything that competitive intelligence professionals do, for example, springs to mind. Likewise, all the customer insight work undertaken by marketing professionals can be counted. Even representation at professional networking events can be regarded as scanning activity. And then, of course, there are all the arrangements commonly put in place to feed organisational risk management with the risk information it needs, such as fostering a speak-up culture and routine use of various risk identification and assessment techniques, often drawing on routinely available financial and performance data, etc.

Clearly, valuable risk information can be sourced from anywhere, by anyone, following any method or technique, and for conceivably many different reasons within organisations or on projects. Having the risk radar in widespread corporate usage, as a simplifying metaphor, reminding all the diverse organisational participants of the common risk radar enterprise they are engaged in, can, we think, prove highly significant for purposes of building integrated risk management. That common enterprise can be further theorised, of course, with reference to the views of risk that need to be harnessed and aggregated so that people can communicate with shared understandings of precisely what risks are at issue.

There is potentially a lot of diversity to consider here, as we explained earlier with reference to various risk ontologies. Separate views of risk, each with high context specificity and its own blending of risk ontologies might emerge in stark contrast with one another across contexts of governance, reporting, operations, strategy, audit, quality, performance, marketing and the like. In addition, a number of more complexity-reducing and holistic organisation or project wide views of risk might also very usefully be available to house these - as we explained earlier when we gave our argument for equating higher level risk categories with broad areas of risk exposure.

Accordingly, risk radar design and use is very tightly associable within the challenge of compartmentalising some organisational or project vision of risk on a per risk basis - bearing in mind diversity of possible applications. Linked to this is the challenge of producing compartmentalised organisational or project views of risk on varying levels of granularity which integrate together and yet remain fit for serving diverse organisational purposes.

Linking the above what and how discussions, then what emerges is that the risk radar metaphor is useful because it can help many and varied participants in risk management to develop an important and yet typically neglected element of risk management expertise. As soon as risk management participants begin to appreciate that they are also risk radar participants, what may then develop is a more critical understanding of risk categories and why they matter, extending to encompass an appreciation of the need to participate in the ongoing

testing and improvement of a clear and precise organisational or project view of risk by focusing carefully on pertinent issues of social construction. This subtle shift of management focus towards greater recognition of the importance of social construction when managing risk, can arguably do much to enliven how people experience and contribute to risk management in their daily working lives. We can neatly conclude this section by using a particular expression for summing up both the what and the how of risk radar use. Risk radars, we would suggest, can usefully be described as producing risk intelligence.

SCANNING AND RISK INTELLIGENCE

The centrality to the risk radar metaphor of organisational scanning allows us to root our ideas within the history of management academia. In 1967, Francis J. Aguilar wrote a seminal management text called 'Scanning the Business Environment'. This focused on strategic planning applications of environmental scanning and was followed up by much literature concerned with complexities of environmental scanning aligned to applications of resulting knowledge for strategic planning and various other corporate purposes. Generally speaking, our present interest in the risk radar can be situated and perhaps developed much further than we can offer here, with reference to this rich history of ideas which has developed within what is now a voluminous management literature.

Two further very recent influences, however, require special mention here. Two predecessor (2017 and 2019) CRR discussion documents dealing with designing and using risk radars to produce risk intelligence have already explored the risk radar's contribution to environmental scanning in some detail. These documents looked at how environmental scanning can be improved in order to manage risk more effectively, especially by applying the military intelligence process logic of proactive direction and collection of risk intelligence information. The present section summarises this risk intelligence context for risk radar use, especially because we think it is important to enliven our discussion of the risk radar by laying a strong emphasis on how it can be used ambitiously, through proactive engagement with, over and above passive exposure to, the risk environment. Indeed, it might even be considered a weakness of the risk radar metaphor that it suggests passive detection when, we would suggest, a proactive scouring of the risk environment for risk intelligence information, as might be summed up in the expression risk intelligence work, is arguably what is really required if risk radars are to be used to full effect.

This lineage of ideas, which conveys us to our present concern to discuss design principles for risk radars, can be traced very briefly as follows. The (2017) document was called simply Risk Intelligence. It proposed that risk identification practice, positioned close to the front end of traditional risk management processes, warrants criticism as tending to be too passive and sedentary; which is to say, too reliant on existing management information. The proposed solution was risk intelligence practice where risk identification is enhanced through professionalised competitive intelligence work, closely aligned to marketing intelligence work and business intelligence work, such that environmental scanning consolidates through the direction-collection logic and phasing of these processes' common ancestor: the military intelligence process.

Crucially, this re-imagining of risk management does not just entail greater proactivity, but more fully a concern to engage reflexively with the full range of human threats spanning competitors, hackers, criminals, terrorists and various other categories of mal-wisher that might seek to undermine corporations. The result envisioned was risk management that looks more like security management focused on use of intelligence analysis and related techniques such as red teaming and credibility/reliability rating procedures, in order to protect critical assets. Drawing on the popular resilience management metaphor of the risk radar, explained as deriving initially from environmental scanning literature, the (2017) CRR document used the expression boosted risk radar to refer to the above enhancements to risk management.

This (2017) CRR discussion document was further developed as an academic paper in the International Journal of Forecasting by Marshall et al. (2019). The academic paper further elucidated risk intelligence work, framed metaphorically in terms of the boosted risk radar. It contended that just as military intelligence famously deals in unknown-unknowns, known-unknowns, unknown-knowns and known-knowns, so too risk knowledge can always usefully be broken down into its abstract and concrete components. Running with this theory, a known-unknown risk is one that is discussed as an abstract theory of what might happen, yet in the relative absence of concrete experience of the risk in the making. Hence discussions of known-unknowns may well focus on taking stock of risks which have struck elsewhere under at least partially similar circumstances. An unknown-known risk, conversely, is where there is experience and awareness of some risk in the making. There may be weak signals which give cause for concern, one might say, and yet the risk at issue continues to defy sufficient abstract categorisation so as to determine how it is viewed. An example might be set within the very commonplace internal risk scenario where there is escalating and increasingly worrying employee turnover and yet leadership finds it difficult and potentially embarrassing to come to terms with the reasons for it. Is this best categorised as a leadership risk? Perhaps a reputational risk, an HR risk or an operational risk? Faced with risk knowledge in such an unbalanced state, there is arguably a good case for developing appropriate abstract categorisation so as to produce balance. This transports us directly back to the concerns of the present document, relating to the need to reconcile abstract theoretical understanding with more concrete information and experience within any view of risk. The nomenclature of (un)known-(un)knowns might, we think, provide very effective terminology for expressing perceptions of imbalance pertaining to either the overall view of risk or to the current state of risk knowledge for any risks therein.

Indeed, the main risk practitioner guidance point made by the (2019) academic paper was to urge that risk knowledge always deserves to be critically scrutinised by considering it in its abstract and concrete aspects. The paper argued that routine use of (un)known-(un) known nomenclature in conceivably any risk discussion, can act as a sensitiser to the current state of risk knowledge by drawing attention to any imbalance between the two theorised polar extremities of high-abstractness-low-concreteness and high-concreteness-low-abstractness, thereby providing not just impetus for but also, crucially, direction for, ongoing risk radar scanning effort. For example, abstract risk knowledge deficits might be corrected through isomorphic learning (perhaps via risk consultants) from how other organisations have regarded similar problems, or by seeking solutions within academic or practitioner literatures; concrete risk knowledge deficits, on the other hand, might be corrected by involving more managers with experience of the risk within risk management processes, and indeed by fostering healthier risk cultures where managers feel they can come forward to share their experiences and concerns.

That brings us to the ensuing (2019) CRR discussion document on risk intelligence production, called Risk Intelligence Production on Complex Projects: practical discussion points. This looked more closely at how risk management can learn, not just from intelligence management and analysis practice in its various corporate and military forms, but also from the academic discipline of knowledge management has offered, for example pertaining to how data can be raised into information, how information can be raised into different forms of knowledge, and how that knowledge can be further raised into actionable insight delivering competitive advantage. This (2019) discussion document argued that developing risk knowledge within organisations requires selective and targeted engagement with multiple management processes that handle all sorts of information. It presented the cultivation of risk knowledge as far more than a matter of raising unknown-unknowns into known-knowns, which it regarded as desirable but not always possible. For example, it explained that organisational distributions of what is termed tacit or ineffable risk knowledge should be valued. In terms of the present discussion, what this effectively means is that it may make sense for the organisational or project view of risk to make some provision for tacit risk management knowledge that may only be articulable in abstract terms (i.e. as known-unknown risk). Speaking more practically, the view of risk would then become one that, to some extent at least, is sufficiently flexible as to accomodate some risk concerns of experienced managers and technical specialists even though these are not fully explicated in terms of detailed evidence and with reference to the precise causes and effects at issue.

To conclude, when risk radar use is equated with risk knowledge production, and when this challenge is further regarded as one of balancing abstract and concrete risk knowledge, there is a further useful expression that we think can usefully be brought into play to remind managers of the importance of optimising the abstract theory, which is to say, of incorporating the best risk ontology within the view of risk. That expression is what gets theorised gets managed. Before finally outlining our five core design principles, we will look briefly at what this expression can mean.

WHAT GETS THEORISED GETS MANAGED

The expression what gets measured gets managed, widely associated with Peter Drucker and Henry Ford, calls attention to the practical necessity of a management focus on measurable performance that is only possible because the matters selected present some objectively measurable variability. We would argue that this idea finds its risk management corollary in the very different social constructionist idea that what gets theorised gets managed. The point, here, is not to suggest that individual risks are intangible and that risk reduction efforts are immeasurable. That is clearly not so. Rather, it is as follows. In the physical world, risks can be seen, heard, felt or smelled, at least insofar as they are in the making and are therefore amenable to direct experience. Everything else can be categorised as foresight reliant on imagination. In the complex social world that all organisations and projects need to concern themselves with, foresight reliant on imagination becomes far more important, especially where the major risks are concerned. Possibility is usually expressed - and even experienced to some extent - as abstract categories that are theorised rather than perceived, in order to reduce complexity on a per risk basis.

This point can be further elaborated with reference to the need for at least some creative or artistic part within risk identification. We would suggest that use of risk radars to produce risk intelligence will tend, to some extent at least, to involve applications of abstract theory so as to provide what might loosely be called structure, direction, liberation, focus, sensitisation, and indeed basic ideation, for the risk imagination. Widely used business ethics conceptions of moral imagination become important within that context. There is also a chicken-and-egg issue arising here between risk imagination and moral imagination, reflecting exactly the chicken-and-egg relationship between the risk radar and the view of risk. On the one hand, risk imagination can be viewed as energised by moral imagination, which might be experienced as care, moral courage or fortitude. Looking from this standpoint, moral imagination is articulated within the view of risk as specific risk vulnerabilities or exposures which operate as categorisation principles to structure and direct scanning effort. However, reversing this, it is also arguable that this moral imagination itself hinges on some prior risk imagination striving towards the long and large view of things, which perhaps always ought to be a governing principle for risk radar use. Clearly this is an important consideration too, and one that might enrich many discussions about risk identification or scanning. The key conclusion arising here is, thankfully, very straightforward: the basic theoretical categories upon which risk management depends can be designed, tested or improved to incorporate ethical concerns to any extent that is deemed appropriate. In other words, ongoing management focused towards the risk radar and the view of risk can provide a valuable opportunity for leveraging more ethical management. High level risk categories which describe risk exposures in particular, can be viewed as expressing what people fundamentally care about and wish to protect through risk management effort. The risk category architectures used for scanning risk environments are, because of this, always considerable for the extent to which they make ethical statements, revealing where the organisation's risk imagination and moral imagination meet.

FIVE DESIGN PRINCIPLES FOR THE VIEW OF RISK

Finally, we come to our five design principles which arise from the foregoing discussions. Up to this point, we have focussed on conceptual innovation, drawing out the issues involved in some detail from a range of perspectives. The five subsections that now follow, however, are likely to be of particular academic and practitioner interest. Hence they are referenced more intensively, drawing attention to their indebtedness to pertinent academic literatures. To reiterate what we said at the start of this booklet, we advocate that our five design principles are considered for their possible relevance on a case by case basis. They are strongly influenced by our own recent experience of designing a generic view of risk to capture the diverse risk contents of SEC 10-K risk reports across multiple industry sectors. That gives us some confidence to expect that our design principle recommendations for generating effective risk categories will very often prove useful.

1. Risk Scoping (and aggregation)

Given that in principle, risk is inherently a multi-attribute concept (Fischhoff et al., 1984), several value judgments must be made when constructing risk categories. The first and most important one, is related to defining a set of core contextual principles to be used to guide the risk identification process. Judgments about the selection of abstract and concrete knowledge on risk attributes and their relative importance are likely to differ between people and organisations (Slovic, 2010; Marshall et al., 2019). For instance, studies have also shown that even simple factors such as differences in demographics, culture or even environmental conditions, may influence the manner risks are assessed and evaluated (Johnson and Bruner, 2003; Costa Sperb et al., 2019). Thus, in categorisation processes that involve multiple stakeholder groups, it is to be expected that each may bring their own risk definition and frames that may be projected when identifying, scoping and conceptualising risks. Plurality of perspectives and value judgements are generally associated with improving the informational richness to which risks are decoded and characterised (Johnson et al., 2006). However, they do not guarantee the adoption of risk lenses that adequately encompass relevant contextual information and/or to a risk conceptualisation process that effectively weighs the relevant constituent components of risks that really matter to an organisation or task. This can result in a disjointed risk communication framework made up of various siloed perspectives on risk, posing a potential threat to the achievement of a common (risk) understanding (Boholm, 2019). Given that the coherence and effectiveness of subsequent categorisation efforts are highly dependable on how risks are identified and framed, it is necessary for risks to be identified in a common, clear and consistent format if there is to be an overall view of risk.

We propose that the foundations of a clear and consistent risk identification/scanning process can be achieved by identifying the core contextual principles to guide the subsequent direction of value judgements and risk framing efforts. Much like a radar, the identification and use of core contextual principles should provide a solid foundation to incite and direct the formulation and adoption of specific attributes and value judgments that are relevant to the existence and prosperity of the company. Under a different analogy, it provides the glue that holds the process of identifying and framing risks in a way that is meaningful to the organisation.

One common practical approach for helping to establish the core contextual principles is to identify the goals and objectives of the company and the risk management decisions that the organisation faces (Morgan et al., 2000). This approach can provide the necessary support to ensure that all risk scoping and identification/scanning efforts are well-aligned within the realms of the strategic direction and operational needs of the company. That is, it offers the benefit of providing an objective criterion to identify and frame risks in a format that is consistently relevant to the organisation. Furthermore, as objectives and needs change, so does the relevance of risks and their attributes. Thus, this approaches offers the additional benefit of providing an enduring, cohesive and adaptive framework to support the adoption of a congruous risk language, ensuring that a common risk understanding is consistently secured over time and that organisations continuously identify and characterise risks in a manner that is relevant to them.

Finally, we propose that each risk identified must be recorded on a risk review sheet. We recommend for the risk review sheet to be designed to help stakeholders to quickly learn the necessary information on each risk as to guide and support the subsequent stages of the categorisation process (i.e., risk review sheets should be concise, consistent and, more importantly, effectively capture risk information in terms that are compatible to the core contextual principles). Then, risk review sheets should be logged in a risk repository database as to provide a readily available information knowledge base on each risk to those who will be performing the categorisation process. The risk repository provides the practical foundation to support a systematic and consistent risk categorisation process.

2. Clarity

The goal of good categorisation frameworks is to generate categories that are effective at explaining the relevant details and attributes that constitute the risks under the category. Importantly, a categorisation framework has to be successful in producing its intended results from inception (i.e., it must be effective from its first introduction). To achieve this, we propose that effective risk categories must follow from their elementary level, the basic principle of *clarity. Deriving from its conventional definition*[1], the principle of clarity offers the necessary direction to develop categories that are coherent and intelligible. In practical terms, categories may be defined based on a common set of topics or themes that emerge from the assessment of a collection of risk review sheets (Morgan et al., 2000).

Clarity begins by having an encompassing understanding of the risk categorisation context. That is, the context of the categorisation process must be defined as it provides a consistent foundation to guide the conceptual levels of clarity desired at each stage of the development process of the categorisation framework. An encompassing understanding of context should adequately balance the objectives of the categorisation process, the visions and needs of the stakeholders involved and the available resources and time constraints[2]. An adaptation of the seven basic questions proposed by Chapman and Ward (2011) provides a practical framework to balance these factors to ensure that the intended risk categorisation purpose is aligned with its motivations. These questions, labelled as the seven Ws, are shown in

[1] From the English dictionary, 'clarity' is the quality of being easily expressed, seen or understood; clearness or lucidity as to perception or understanding.

Table one. The key words in italics highlight the core elements of each of the seven (W)amigos. By using the seven-Ws as a compass when defining the categorisation context, it is more likely that all relevant (contextual) components are pondered and balanced and that their interdependencies are considered. This framework provides structure to develop categories that are feasible, practical and useful. In other words, they facilitate consideration of the best risk ontology solution for the view of risk.

In the real world, it may be challenging and often impractical to attain full clarity under categorisation processes that are deployed on complex or opaque contexts or that involve less mature risk functions. A useful aim to employ in these cases is that of seeking to achieve a 'minimum clarity' acceptable (Chapman and Ward, 2011). A minimum clarity categorisation process is one that accomplishes its intended purposes without compromising its contextual requirements. Put simply, its purposes should be achieved by providing the minimum level of support and/or basic necessities demanded by each of the contextual elements described in Table one. As a cautionary tale, it is important to note that when applying the minimum clarity concept, the end result should be to make the categorisation process more efficient and simpler, but never simplistic[3]. Hence, the seven Ws are aimed at providing direction to guide the trade-offs between the relevant contextual elements necessary to develop categories that possess at least a minimum clarity acceptable. In reference to the opening sentence of this principle, we propose that at the fundamental level, a minimum clarity approach should lead to the generation of categories that are effective at explaining the relevant details and attributes that constitute the risks under the category in a contextually meaningful manner.

Notably, an adjacent benefit of following a minimum clarity approach is the improvement in efficiency brought to the process of constructing the categorisation framework. Embedding a minimum clarity approach across all stages of the categorisation process may prevent the issue of 'paralysis by analysis', where unnecessary and disproportionate attention and resources are allocated to particular stages and issues, hence preventing the efficiency of the development of the process altogether (Forbes, 2005). Importantly, if further clarity over and above the minimum clarity approach is desired, the further principle of iteration (discussed earlier in relation to risk realism) provides the necessary support to refine and improve clarity where deemed necessary. Specifically, if the concept of minimum clarity and the principle of iteration are used in synergy, the end result is a categorisation process where improvements are possible when necessary, while ensuring the effectiveness of the categorisation process design is achieved from its first introduction.

1. **W**ho	Who are the stakeholders involved?	*Stakeholders*
2. **W**hy	What do stakeholders want to achieve from the categorisation process?	*Purpose*
3. **W**hat	What are the desired categorisation design features that stakeholders are interested in?	*Design*
4. **W**hichway	What are the tasks and procedures and necessary to develop the categorisation framework?	*Activities*
5. **W**herewithal	What are the resources necessary to develop the categorisation framework?	*Resources*
6. **W**hen	What is the timeline to deliver the categorisation framework?	*Schedule*
7. **W**here	What are the functions or settings where the categorisation process will be used?	*Setting*

Table one: Key questions to guide the definition of context – The seven (W)amigos

[2] This is not to be confused by the core contextual principles proposed in the risk scoping principle: the latter is intended to help identify 'risks' that matter to the organisation, while the *context in the clarity principle is intended to provide more encompassing guidance to support the development of 'categories' that matter.*

[3] The meaning of the adjective *simple is related to plain, easy, ordinary, or uncomplicated. From this definition, a simple solution to a problem is still an effective solution. On the other hand, the meaning to the adjective simplistic is a pejorative connotation related to disproportionate and oversimplification (i.e., characterized by extreme and often misleading simplicity, leading to an ineffective solution to a problem).*

3. Logical Completeness: Mutual exclusivity and collectively exhaustive

In general terms, the logic of the mutually exclusive premise specifies that two or more propositions cannot be simultaneously true in the same sense. To say that risk categories are mutually exclusive means that any one risk cannot be allocated at two or more categories at the same time. Motivated by this premise, mutual exclusivity proposes that risk categorisation processes should aim to formulate risk categories that are unambiguously distinguishable amongst the spectrum of categories established (i.e., ensuring that risks can be placed into one category and no other). However, it is important to note that mutual exclusivity does not necessarily result in categories that are collectively exhaustive of the risks to be categorised. Using a generic example to illustrate this, when categorising a standard 52-card deck in odds and evens, we achieve categories that are mutually exclusive (as card numbers can only be even or odd), however not collectively exhaustive (as the Ace, Jack, Queen and King would not fit any of these two categories[4]). The latter premise in a categorisation process is referred to generating a portfolio of categories that cluster the entire realm of possible risks identified. Together, these two premises compose the concept of a facet, which is defined as "a clearly defined, mutually exclusive, and collectively exhaustive aspect, property, or characteristic of a class or specific subject" (Taylor, 2004). Therefore, to ensure that categorisation structures are logically complete and effective at clustering risks, it is necessary for the premises of collective exhaustive and mutual exclusivity to work in tandem, in a process to generate 'facet-based' categories.

In reality, due to the often complex and ambiguous nature of risks, it may be challenging to attain both premises, for any given overall view of risk. Although it may be simpler to generate collectively exhaustive categories (as at the very basic level, it is possible to attain a portfolio of collectively exhaustive categories as it is possible to generate categories that encompass a singular risk), the premise of mutual exclusivity is more difficult to achieve. For instance, it is possible that one or more risks share attributes (i.e., risks may not be independent in their entirety, such as sharing a root cause or posing identical symptoms and impacts). This means that there will be instances when subjective judgments will have to be made in order to support a clear distinction of risk categories. In practical terms, structuring the categories into a hierarchical framework can greatly improve the consistency to which such judgments are made, hence improving the prospects of accomplishing the mutual exclusivity premise across risks that may display a degree of commonality (Morgan et al., 2000). For instance, hierarchies provide a sequentially objective decision support framework, where conflicts are likely to be contained and more objectively identified at a particular 'level' of commonality. This greatly helps the categorisation developer to more precisely identify the areas (e.g., categories and their attributes) demanding improvements to achieve the premise of mutual exclusivity.

In practical terms, categorisation frameworks that are organised hierarchically should be aimed at composing categories that, when read in conjunction, provide sufficient clarity and enable prompt understanding of the nature of the risks disclosed. In the hierarchy, top-level categories could be viewed as directing risk awareness towards a very general understanding of the broad environmental (and especially, exposure) contexts of the risks, before then being refocused by lower-level categories towards more specific and numerous relevant risk issues and attributes linked to these general contextualising understandings.

One initial guidance suggestion for achieving mutually exclusive categories when using hierarchical frameworks, is to deconstruct any identified ambiguity and commonality of (overlapping) risks. The aim here should be to reconceptualise these risks in terms of their unique attributes relevant to the hierarchical level where the conflict exists (Haimes et al., 2002). If this initial approach is unsuccessful at generating mutually exclusive risk categories, the alternative is to redefine and aggregate these risks as there would be no reasonably clear basis to distinguish them, hence resulting in categories with a greater degree of generality. As a last resort solution, if overlapping categories still exist after significant attempts are made to address the issue, a category labelled 'other' should be used. It is preferable to adopt the generic risk category 'other' relative to having overlapping categories, as the former is more likely to ensure that the integrity of the remainder (mutually exclusive) categories is preserved.

In summary, the principle of logical completeness may offer practical guidance to support the creation of complete and accurate categorisation frameworks, hence providing a fundamental step in the direction of the development of the overall view of risk.

[4] This is based on considering that their 'face value' are letters, so in their fundamental sense they do not fit an even/odds categorisation. However, it is customary in some card games to represent these cards as numbers, making them suitable to an even/odds categorisation. Importantly, this is only possible when there is a collective understanding about the assumptions made regarding the changes in their face value. This implies that such categorisation structures would only hold and remain consistent for as long as there is full consensus and understanding by stakeholders on the changes necessary to make the categorisation effective.

“Plurality must never be posited without necessity.”

William of Ockham

4. Parsimony

The extreme boundaries of the optimal number of risk categories must lie between one (i.e., a single category encompassing all risks) and the number of risks to be categorised (i.e., each risk being its own risk category)[5]. Following Aristotelianism, parsimony is related to the principle that the simplest explanation is the preferred one. When applying this principle to categorisation processes, the goal of parsimony is to develop the least number of categories that accurately and coherently encompass all risks. As discussed in the risk scoping principle, conceptualising risks may be an often difficult and complex task. Therefore, parsimonious categorisation frameworks can greatly enhance the intelligibility and communication of risk categories by minimizing the complexity and overcomplication of the categorisation structure.

In the context of risk categorisation frameworks, complexity may be related to a range of factors. These factors include heterogeneity levels of risk attributes (Doerner 1980), the volume of the information processing necessary to conceptualise risks (Timmermans 1993), the ability to conceptualise risks based on the complicated and often overlapping relationships between the nature and attributes of risks (Sung and Johnson 2007; Tversky 1972) and the dynamic nature of risks (e.g., recognising the emergence of new risks; Hogarth 1987). These factors are generally associated with two types of complexity: alternative and attribute-based complexity. The former relates to increases in complexity associated with the volume of risk categories. The latter relates to the ease to which risk categories can be discriminated, based on the attributes that define each category.

With the rises in these complexity levels, individual capacities to effectively decode information significantly decrease (Brehmer and Allard 1991), making it difficult to learn as well as to employ information in a logically correct manner (Berry and Broadbent 1984). Increasing complexity may also render categorisation stakeholders more reliant on heuristics when using the categorisation framework (Payne et al. 1993). This, in turn, may negatively influence the usefulness of the risk categorisation relative to its intended purpose, as increases in heuristics are highly correlated with increases in the number of cognitive errors (Reason 1990) and systematic biases (Harvey et al. 1994; Kahneman and Tversky. 1982). Therefore, it is critical for a categorisation framework to minimise complexity so as to deliver on its clear and intended purpose (i.e., its conceptual purpose), while also remaining practically useful for stakeholders. This need for balance might be regarded as an important motivating factor for the parsimony principle introduced here.

From the discussion above, the functional goal of parsimony is to minimise alternative and attribute-based complexity whenever possible. Although it may be difficult to reduce both forms of complexity simultaneously, the following approach may provide some practical guidance for sequentially reducing them to achieve parsimony. We might also refer to the initial stage of the categorisation process as the construction phase, as it is in this stage that the majority of categories are formed. We recognise, of course, that any organisation or project seeking to actively take a view of risk may well not be doing so entirely from scratch.

During this initial construction phase, it is preferable to increase the alternative-based complexity, in order to obtain an initial level of specificity and clarity when defining the attributes that compose each category formed. However, with increasing maturity of the categorisation process; that is, increasingly over the course of successive iterations, the relative importance of the two concerns should be reversed. The key concern should now be to increase attribute-based complexity by increasing the attributes that constitute each category. More fully, at these more mature stages of the categorisation design process, the greater concern should become that of reducing the number of categories by increasing the attributes that define each category. This might usefully be called the deconstruction phase to differentiate it from, and clearly position it as consequent to, the construction phase.

Hence, for example, a sufficient number of categories that exhaustively encompass all risks identified should be formed in the initial phases of the categorisation process. Specifically, each category should initially be formulated on the criterion that it should viably and effectively explain all risks to be categorised. At each iteration through the risk repository, the aim should be that of increasing the specificity of the risk categories so as to increase the number of risks encompassed in each category. To conclude, then the underlying purpose of parsimony is to balance the complexity and completeness of the categorisation structure, thus rendering it effective for its intended purpose while still being comprehensible and useful for a wide array of stakeholders. In succinct terms, this illustrates the functionality of the principle of parsimony.

[5] This is based on the mutually exclusive principle, where the assumption is that each risk can only pertain to one category, as otherwise it is highly unlikely that risk categories are coherent and intelligible.

5. Iteration

As introduced in the clarity principle, the general purpose of a categorisation process is to assist with the development of an appropriate and sufficiently clear understanding of all relevant risks as to improve the effectiveness of subsequent risk management processes. Establishing the iteration principle as a central mechanism underlying the construction of risk categorisation frameworks is critical to the formulation of coherent and effective categories. Indeed, this follows directly from what we said earlier in our section on risk realism. In summary, the realist epistemology of risk realism applies caution, care and as much mental dexterity for thinking about risk as it can muster, to the task of testing and improving its risk knowledge base. Risk knowledge is regarded as corrigible and therefore a methodological realism founded upon iterativity becomes a plain necessity.

The importance of this principle is likely to be recognised very strongly by participants in the categorisation design process. Going through categorisation phases multiple times generates more in-depth understanding of risks and their individual attributes, partly because it encourages decision makers to maximise the value of information available, and also because it promotes a systematic learning process cycle for continuously improving categories as new evidence and knowledge emerge within the constantly changing risk environment (Grondahl et al., 2011). More fully, however, the practical purpose of the iteration principle is to allow the development, restructuring, refinement, alignment, correction and reconsideration of various aspects and assumptions undertaken to date when formulating risk categories. In other words, methodological refinement is at issue over and above better use of a constantly improving and changing information base.

Categorisation processes generally follow two fundamental approaches. They are either based on a linear procedure, where all categories are formulated on a 'one-pass and right-first-time' basis, or they follow an iterative approach, where planned and unplanned iterations are carried to revise the composition and decomposition of categories. We propose that the former approach is likely to be highly inefficient and unrealistically attainable in most categorisation processes. Due to the complex nature of risks, it is likely that there will be insufficient levels of information on relevant aspects for each risk to be categorised in accordance with the previous principles. This may lead decision makers to waste time and effort on unimportant issues and features of the categorisation process, potentially directing attention away from important issues not anticipated when the categorisation process started, perhaps leading to seriously ineffective categories (Chapman and Ward, 2011). Thus, linear categorisation procedures are inherently prone to take on an 'unplanned' iterative character. This accidental stumbling into good practice may occur because shortcomings and deficiencies of categories are discovered, or due to the need to align and correct assumptions related either to new risks, or indeed to the evolution of previously identified ones.

Conversely, planned iteration involves purposely embedding, within the categorisation process, the practice of revisiting earlier developments within a phase or looping back to earlier steps. Planned iterations can be embedded as a systematic review of the categorisation process (e.g., iteration occurring at pre-specified milestones or intervals), or they may be included in an unsystematic fashion (e.g. 'on demand', whenever new risks and related information/knowledge emerge to demand attention).

Recognising the complementarity of both approaches, we further propose that flexible and phased iteration is critical for maximising the benefits provided by the iteration principle, while minimising its potential costs. A flexible approach refers to the process of carrying planned iterations throughout the categorisation process, while allowing for reactive iterations should unplanned shortcomings or deficiencies be discovered. A phased approach concerns the employment of different iteration targets and points of emphasis that are contingent on the maturity and lifecycle stage of the categorisation process. In earlier stages and for less mature categorisation processes, iterations should be aimed at establishing the foundations of the categorisation process. In other words, they should be based on refining the scoping of risks and the understanding of context in order to compose a general categorisation framework (Morgan et al., 2000). This may, generally speaking, demand several complete loops through all phases of the categorisation process, where the practical emphasis should be to develop a solid and intelligent categorisation foundation for the view of risk. For example, a natural initial iteration may be based on looping through risk summaries in the risk register database in order to identify general risk topics (i.e., 'themes'). As for the second iteration, risk themes could be aggregated and contextualised collectively. Following this example, each subsequent iteration should aim to improve the resolution of each risk in order to allow the initial construction of general categories. Recognising the opportunity cost associated with earlier iterations, the aim should be to adequately balance the resources and attention employed in relation to the desired level of clarity attainable at early stage iterations. Late iterations and more mature processes should progressively allocate more attention and resources towards refining and enhancing specific features, and addressing critical issues within the categorisation process. Characteristically, later iterations warrant a profound and thorough attention and understanding of very specific components and issues derived from the constructs developed from earlier iterations or from the emergence of new information or risks. In practical terms, this should be achieved via partial loops within specific stages or issues within the categorisation process, where resources should be strategically placed to refine constructs and consolidate the categorisation frameworks developed from earlier iterations. The end goal of later iterations should be to ensure the formulation of a parsimonious framework that contain categories that that are clarity efficient, logically consistent and mutually exclusive.

In sum, earlier iterations should focus towards providing a general structure and foundation for the categorisation process before allocating larger resources and attention to refine in depth, while later iterations should be aimed at gradually becoming more focused on key issues of the categorisation process. In other words, earlier iterations should be aimed at constructing and composing a general categorisation framework, while later iterations should be aimed at improving specific issues and processes by decomposing and reconstructing initial constructs.

CONCLUSION

Many approach risk management with an insurance focus, an internal control focus, or indeed a focus on narrow risk management processes and their associated risk matrices and risk registers, which arguably remain the most highly visible features of risk management practice within organisations and on projects. We conclude this booklet by observing that no matter what focus is adopted, there will always remain a need to engage with the complexities of risk ontology. Every approach to risk management is, by default, also risk ontology focussed. There is simply no taking the risk out of risk management. Arguably, it is surprising that there is so little practitioner guidance today which addresses the challenges that relate to simply thinking about risk, conceivably within any management context where risk is deemed to matter.

Our booklet has striven to address some of the complexities involved in thinking about risk, and in creating and maintaining a view of risk as the central repository for what is known about risk. We hope that our simplifying risk management narrative of mutual support between the risk radar and the view of risk, further aligned to some very practical design principles for creating and improving the view of risk, will provide some food for thought. Likewise, we also hope that our discussions have helped to elucidate what risk intelligence production can usefully be understood to involve.

Arguably, within what we might call professional attitude towards risk management (which might be understood to comprise many and varied aptitudes), a central place should be granted to what is best called a professional wit for dextrous thinking about risk itself. Our various recommendations have all been concerned, in effect, with the cultivation of this professional wit. We envisage such cultivation as very much an ongoing and collaborative enterprise within organisations and on projects. We have looked at how this professional wit can be cultivated through giving more serious regard to the problematic nature of scanning for risk within risk environments, and indeed by taking the view of risk as both its end product and its touchstone for ongoing improvement. We welcome all comment and feedback.

REFERENCES

Al-Mazrouie, J.R., Ojiako, U., Williams, T., Chipulu, M. & Marshall, A. (2020). An operations readiness typology for mitigating against transitional 'disastrous openings' of airport infrastructure projects. Production, Planning and Control. In press.

Berry, D. and Broadbent, A. (1984). On the relationship between task performance and associated verbalised knowledge. The Quarterly Journal of Experimental Psychology A, 36, 209-231.

Boholm, M. (2019). Risk and quantification: A linguistic approach. Risk Analysis, 39(6), 1243-1261.

Brehmer, B. and Allard, R. (1991). Dynamic Decision Making: The Effects of Task Complexity and Feedback Delay. In J. Ramussen, B. Brehmer, and J. Leplat (eds.), Distributed Decision Making: Cognitive Models Of cooperative Work. Chichester: Wiley

Chapman, C. and Ward, S. (2011). How to manage project opportunity and risk. Jonh Wiley and Sons Ltd.

Costa Sperb, F., Roussou, F., Marshall, A. & Mues, C. (2018). Risk Radar: a methodology to categorise risks disclosed on SEC 1--K reports. University of Southampton.

Costa Sperb, L.F., Sung, M.C., Johnson, J.E. and Ma, T. (2019). Keeping a weather eye on prediction markets: The influence of environmental conditions on forecasting accuracy. International Journal of Forecasting, 35(1), 321-335.

Doerner, D. (1980). On the difficulties people have when dealing with complexity. Simulation and Games, 11(1), 87-106.

Fischhoff, B., Watson, S. and Hope, C. (1984). Defining Risk. Policy Sciences, 17 ,123–139.

Forbes, D. (2005). Managerial determinants of decision speed in new ventures. Strategic Management Journal, 26(4), 355-366.

Gregorc, A.F. (1984). Gregorc Style Delineator: development, technical and administration manual. Gregorc Associates, Inc.

Grondahl, I.H., Lund, M.S. and Stolen, K. (2011). Reducing the effort to comprehend risk models: Text labels are often preferred over graphical means. Risk Analysis, 31(11), 1813-1831.

Haimes, Y.Y., Kaplan, S. and Lambert, J.H. (2002). Risk filtering, ranking, and management framework using hierarchical holographic modeling. Risk Analysis, 22(2), 383-397.

Harvey, N., Bolger, F., and McClelland, A. (1994). On the nature of expectations. British Journal of Psychology, 85(2), 203-229.

Hogarth, R. (1987). Judgment and choice: The psychology of decision. Oxford: John Wiley and Sons.

Johnson, J.E.V., Jones, O. and Tang, L. (2006). Exploring decision makers' use of price information in a speculative market. Management Science, 52(6), 897-908.

Kahneman, A. and Tversky, A. (1982). Judgment under uncertainty: Heuristics and biases. In Utility, Probability, and Human Decision Making, 141-162. Amsterdam: Springer.

MacKenzie, C.A. (2014). Summarizing Risk Using Risk Measures and Risk Indices. Risk Analysis, 34(12), 2143-2162.

Marshall, A. & Ceylan, S. (2020). Risk Intelligence Production on Complex Projects: practical discussion points. A Centre for Risk Research Discussion Document. Southampton, GB: University of Southampton.

Marshall, A., Johnson, J. (Ed.), Dawson, I. (Ed.), Lin, F. (Ed.) & MacCrae, C. (Ed.). (2017). Risk Intelligence: a centre for risk research discussion document. Southampton, GB: University of Southampton.

Marshall, A., Johnson, J. (Ed.), Sung, V. (Ed.), Ashleigh, M. (Ed.), Baden, D. (Ed.), Brito, M. (Ed.), & Dawson, I. (Ed.) (2016). Why risk cultures need prudence. (A Centre for Risk Research Discussion Document). Southampton, GB: University of Southampton.

Marshall, A., Ojiako, U., Wang, V., Lin, F. and Chipulu, M. (2019). Forecasting unknown-unknowns by boosting the risk radar within the risk intelligent organisation. International Journal of Forecasting, 35(2), 644-658.

Marshall, A. & Ojiako, U. (2013). Managing Risk Through the Veil of Ignorance. Journal of Risk Research, 16(10), 1225-1239.

Michaelson, C. (2020). How CEOs, experts and philosophers see the world's biggest risks differently. Theconversation.com. Online article posted on January 27th, 2020.

Morgan, G. (1986). Images of Organization. Sage Publications.

Morgan, M., Florig, H., DeKay, M. and Fischbeck, P. (2000). Categorizing risks for risk ranking. Risk Analysis, 20(1), 49-58.

Payne, J., Bettman, J. and Johnson, E. (1993). Adaptive decision making. Cambridge: Cambridge University Press.

Reason, J. (1990) Human Error. Cambridge: Cambridge University Press.

Slovic, P. (2010). The Feeling of Risk. London and Washington: Earth-scan.

Smith, B. & Raspin, P. (2008). Creating Market Insight: how firms create value from market understanding. John Wiley & Sons.

Sung, M. and Johnson, J. (2007). The influence of market ecology on market efficiency: evidence from a speculative financial market. Journal of Gambling Business and Economics, 1(3), 185-198.

Taylor, A. (2004). Wynar's introduction to cataloguing and classification. Libraries Unlimited.

Timmermans, D. (1993) The impact of task complexity on information use in multi-attribute decision making. Journal of Behavioral Decision Making, 6(2), 95-111.

Tversky, A. (1972). Elimination by aspects: A theory of choice. Psychological Review, 79(4), 281-299.