UNIVERSITY OF Southampton

Risk intelligence:

a Centre for Risk Research discussion document

FOREWORD

This latest addition to the Centre for Risk Research series of discussion and guidance documents is intended to enliven debate about how organisational risk management can be improved to tackle risk associated with both cyber and non-cyber related security threat, and with business competition. It contends that risk managers should focus more on actively gathering 'risk intelligence', and likewise that the 'risk radar' concept deserves closer scrutiny and development in order to recognise the need for proactive exploration of risk environments. These suggestions reflect the importance of developing shared terminology within risk management as a basis for improving practise. They also speak to the need for risk management to integrate with other organisational functions. The document's various illustrations of how this can entail learning from military intelligence techniques and engaging the skills of competitive intelligence professionals will hopefully stimulate students of risk management and practitioners alike to view the risk profession in a new light. We are grateful to all contributors and welcome feedback from everyone.

Professor Johnnie Johnson

Director of the Centre for Risk Research

Professional Association Foreword: "The Institute of Risk Management is keen to stimulate debate about new techniques to manage risk in a world of increasing complexity. We welcome new, interdisciplinary thinking that moves us beyond process and linear analysis and supports effective Enterprise Risk Management. We would encourage our members and the wider risk community to engage with this paper and consider how its concepts might play out in a real world context."

Carolyn Williams CMIRM, ACII

Director of Corporate Relations, Institute of Risk Management

Author

Dr Alasdair Marshall (Centre for Risk Research)

Contributing Editors

Professor Johnnie Johnson (Centre for Risk Research)
Dr Ian Dawson (Centre for Risk Research)
Dr Fenfang Lin (Centre for Digital, Interactive and Data Driven Marketing)
Dr Claire MacRae (Glasgow Caledonian University)

Published in December 2017 by the Centre for Risk Research, University of Southampton

INTRODUCTION

This discussion document asks whether risk management should seek to mature as a more proactive organisational function offering capability for audacious non-routine engagement with non-routine threat involving competitive or other hostile ill will by persons or groups towards organisations. This concern leads it to contemplate how risk management might develop forms of vigilance and investigation which actively explore the social world well beyond the radar scope currently associated with the risk profession. More fully, this central question can be expressed in terms of how risk management might focus more towards providing effective early warning against - and *ex ante* engagement with – threats rooted in ill will towards organisations. The first half of this discussion document will explore this idea using the metaphor of an enhanced *'risk radar'* for organisations. The second half will continue the discussion on a more practical level by looking at how the risk radar can gather and process *'risk intelligence'* to create the *'risk intelligent'* organisation.

Thinking along these lines entails asking the risk profession to look more closely at how 'risk identification' is best understood and undertaken within organisations. Perhaps some healthy introspection is called for on the question of whether there is too much desk-based risk identification practice. In other words, perhaps risk identification is all too often ensconced within organisational 'safe spaces' where it is limited by desk-bound imagination and readily available data. Perhaps what really matters, much more than is commonly acknowledged, is risk identification activity which goes out into the world seeking information which is elusive and yet highly valuable for organisations.

Professional Association Commentary: *"Many risk professionals agree that a healthy approach to risk identification is to be perpetually dissatisfied with existing "legacy" risk information, and to appreciate the necessity of striving to access more and better information from whatever more predictive sources become available. In a world where 90% of data was created in the last two years, the more proactive and creative such effort, the more likely it is that an organisation will be able to test, broaden and update its view of risk to keep pace with the fast-moving, complex and increasingly connected risk environment".*

Julia Graham,

Depute CEO and Technical Director,
Association of Insurance and Risk Managers in Industry and Commerce

Accessing Primary Sources of Risk Information

We contend that purposeful, targeted - and sometimes episodically repeating, learning and transforming - human threat towards organisations deserves more attention as a serious and perhaps even paramount challenge for risk management. We explore various issues arising with that contention, which we think have not yet received sufficient critical and creative discussion within the risk profession.

Perhaps the risk identification issue introduced above can be more precisely framed in terms of whether the modern risk profession might be excessively cautious in its views and assumptions concerning ways in which risk information can and should be sourced. Concern with purposeful and targeted human threat naturally leads us to consider that reliance on secondary sources which bear testimony to such threat may often prove insufficient. This is partly because the information available from such sources will often be out of date by the time it becomes available. For example, whenever magazine articles are written about new forms of cyber threat, whenever security consultants begin advising companies on how to handle these threats, or whenever new and emerging cyber threat information is pooled and circulated across entire industries, those who pose the threats will almost certainly know about any publicity they receive. They might even have anticipated such publicity and subsequent defensive precautions by organisations, in order to plan new and better forms of attack well in advance.

Speaking more generally, no matter what particular form of corporate mal-wisher we have in mind, we can usefully envision sequences of attack/intrusion as repeating episodically and transforming through each side's learning experience – including through insights or even guesses made by each side about their adversary's learning experience. Needless to say, such ongoing reflexive interaction might often be highly complex and characterised by high levels of ambiguity over intentions and capabilities which could become a sensible focus for investigative risk management effort. The fact remains, however, that desk-based interrogation of secondary sources will often prove insufficient as a basis for such effort. Each time publicly available information on a new threat or attack/intrusion episode spurs organisations to shore up their vulnerabilities, they may well remain several steps behind the attackers/intruders who are constantly anticipating and seeking new ways around the strengthening defences.

EARLY WARNING RISK RADARS

Given the above problem of organisational 'lag time' in relation to emerging threat, it makes good sense for organisations to seek to improve their defences by taking whatever reasonable investigative measures they can to reduce lag time for constructing defences against emerging threats. Such improvements can be theorised metaphorically in terms of organisations extending the reach of their 'risk radars' to get as close as they dare to the primary sources who pose the threats. Such improvements can be further theorised as aspiring to enable organisations, wherever possible, to prepare defences against threats before attackers/intruders have time to develop them toward implementation. In that sense, organisational risk radars are perhaps best conceived as things that should be developed towards providing as much enhanced 'early warning' as possible.

Such early warning capacity would of course need to comprise systematic and constantly learning defences. The coordinated participation of various organisational functions would therefore be necessary, based on a common appreciation of the need to act from within each specialism to support or engage in far-reaching investigative risk identification effort. Levels of participation would inevitably vary by function, yet a baseline level of participation for all functions would seem appropriate to ensure a focus on providing resilience against the unexpected. Without this, the result may be risk radar coverage skewed to reflect pre-existing organisational biases.

Risk Management Ownership of the Risk Radar

A strong case can be made that these risk radar defences are best championed and owned by the risk management function in particular, despite protestations that might be levelled by marketing, cybersecurity, competitive intelligence and others. Numerous issues need to be addressed within any new thinking we might do to sketch out what is required here. Firstly, we need to conceive of a suitable risk management process incorporating and moving forward from risk identification, which is integrated clearly, effectively and efficiently – and to achieve clarity efficiency in particular - with various other management and governance processes. At the very least, the situation of such a process within - and across - the organisation's various 'lines of defence' would need to be clear. However the broader challenge would essentially be one of information management, re-envisioned to accommodate the risk radar metaphor as the basic sensory apparatus through which the organisation becomes aware of and responds to what is happening in the world.

Secondly, we would however need to consider that some conceptual awkwardness inevitably arises with our conceptual prism of risk radar designed to receive and relay risk information while also performing a broader information gathering and processing role. In the real world, what would matter is that investigative effort yields information useful for various organisational purposes, and that all gathered information can circulate throughout organisational lattices effectively and with the sometimes highly confidential treatment it requires. However, why should such information always be couched in the language of risk and flow through what are called 'risk management' processes? Surely it would it be foolish to even attempt this? Clarity would be required concerning the scope of each risk radar, with broader scope tending to bring with it more serious fitness-for-purpose issues for the processes concerned. It would be particularly important to guard against any use of unnecessarily restrictive terminology whose effect might be to funnel information through inappropriate semantic and organisational channels. Would the risk radar focus on providing early warning of all sorts of threats, including those arising with human ill will towards organisations? And what about opportunities? Would it also seek to initiate what is often termed competitor intelligence, business intelligence, marketing intelligence, or even consumer insight?

Problems with Looking at Information Gathering through a Risk Lens

It should really be common sense to anyone employing the risk radar concept that it should seek sensitive attunement to all sorts of information, demanding various forms of expression and useful for diverse organisational purposes. It would be ludicrous to build a risk radar which sees only 'risk'; which is to say one that partially blinds itself by insisting upon seeing only what the lens of 'risk' renders visible. Alternatively, it can credibly be argued that risk radars need to gather all sorts of information so that, speaking generally and abstractly, the organisation can manage its risks more effectively.

Perhaps the only reason why 'risk radar' is a helpful metaphor at all is that it can credibly refer to the organisation that is carefully attentive to the world around it so that it can better contemplate and influence its future. Everything which an organisation does in this respect can be considered to involve matching risks to controls, which can be labelled abstractly as 'risk management' activity. It follows that organisational radars are always 'risk radars' because they are always *for* risk management in this broadest conceivable sense. Good ones are likely to be attentive to both threat and opportunity, in recognition of the fact that an astute grasp of both the past and of future possibility requires agile use of both threat and opportunity frames to attune towards complexity within any causal understanding of what is going on in the world, and for all ensuing contemplation of possibilities for action.

Perhaps the biggest problem with the risk radar concept, then, is that people may use it with insufficient understanding of why simplifying metaphors are problematic. The 'risk radar' metaphor is plainly just one of many interrelated metaphors which are fundamental to risk management practice, whose combined effect might sometimes be to normalise their uncritical use. Fortunately there is a fascinating management literature dedicated to debating the pros and cons of management metaphor, which reminds us how we can be critical. One prominent author within this literature, Gareth Morgan, has famously argued that despite problems such as the ease with which the mind can become trapped within metaphor, we nonetheless need it for our metal dexterity, particularly in order to engage with complexity.

This idea arguably has important implications for risk management. The profession's terminology should, arguably, align to a fundamental management concern to achieve a rich causal understanding of what is going on in the world so that actions can be taken to improve organisational prospects. We perhaps begin to disengage from this concern when we view a 'risk radar' as a device that spots 'the risks' – just as a military radar might spot incoming enemy aircraft – so that 'the risks' can then be pushed through risk management processes and dumped into risk registers and internal control systems. The solution is perhaps, very simply, to regard the 'risk radar' as information gathering which, when used well, advances the organisation's concern to have a rich understanding of its environment which can be useful for many purposes.

Generating Consumer Insights

Thinking flexibly about purpose, then, the notion that risk radars can be used to generate consumer insight is particularly intriguing because it leads us to conceive of sustained investigative effort which regards insights as nuggets of gold for organisations, discovered only rarely, yet sometimes offering tremendous competitive advantage. An excellent (2008) text by Brian Smith and Paul Raspin called *'Creating Marketing Insight'* discusses consumer insight along these lines. A very relevant source for further opening out the risk radar concept skimmed over within the present document, it emphasises need for elaborately designed scanning systems characterised by planned variation in allocations of monitoring tasks to managers with diverse skillsets, matched to the changing complexity and volatility levels of competitive organisational environments. Smith and Raspin consider all the elaborately designed scanning effort they propose as worthwhile if it can hit the jackpot just very occasionally. More specifically, they describe the rare moments of insight that scanning should aspire to create as deserving the label 'exquisite' where they meet these four **'VRIO'** criteria:

- Insights should offer **V**alue-added for organisations;
- Insights should be **R**are in the sense that competitors are unlikely to find them;
- Insights should not be **I**mitable; competitors should lack capabilities to either find them or act upon them;
- Insights should be aligned to **O**rganisational capabilities for follow-through actions (i.e. be actionable).

Risk professionals may well raise an eyebrow here and contemplate that perhaps the above VRIO criteria should routinely be considered for their relevance to risk identification activities. To elaborate, we could even define the term 'risk insight' by these same VRIO criteria, and contend that all risk information pushed through risk assessment processes would benefit from routine screening for whether it also constitutes an important 'risk insight'. It could also be argued that risk registers might usefully acknowledge – and perhaps even quantitatively 'score' – the insight value of identified risks. This could help to transform the value of risk registers as decision-making tools, because good strategic decision-makers are likely to value any simple visual representations of knowledge which might give their organisations some advantage over their competitors. The common aspiration – some might say the delusion – of the 'transparent' public-facing risk register would of course need to be reconsidered with a healthy infusion of realism.

However the immediate point at issue here is to note that Smith & Raspin's theory of consumer insight illustrates very ably indeed that not all useful, actionable information derived from organisational scanning effort is best wrapped up in the language of risk as a risk type within an assigned risk category. Risk radars need to feed 'multilingual' communication and dialogue within organisations, where the financial language of the boardroom, the risk language of the risk management department, the marketing language of the marketing department and the technical language of cybersecurity are all accommodated and translated where necessary.

> *Risk professionals may well raise an eyebrow here and contemplate that perhaps the above VRIO criteria should routinely be considered for their relevance to risk identification activities.*

'Boosting' the Risk Radar

Our proposed 'risk radars' will also need some dedicated management infrastructure which can be quickly scrambled to go out and engage with any threat the radar finds – and whose active exploration of the organisation's environment can also be theorised as shaping and boosting the reach of the risk radar. If we wished we could call this 'booster infrastructure'. We could even begin to theorise this as the key missing link within risk management thinking today. Once more, however, the conceptual lens of risk management is problematic. If this is to be considered risk management infrastructure at all, then we will need to rethink the risk management skillset to include, for example, the skillsets of competitive intelligence professionals and other specialists whose bread and butter is to know the dos and don'ts of dealing with primary information sources. This suggests, at the very least, a need for much closer cooperation between these specialisms.

Professional Association Commentary: *"Today's businesses, large or small, operate in an increasingly VUCA world; that is, in a world characterised by **v**olatility, **u**ncertainty, **c**omplexity and **a**mbiguity. Many would describe this as the new normal. In the VUCA world, flows of information offer more business value than ever before. Such information can be assessed by its 4Vs: its **v**olume, its **v**ariety, its **v**elocity and its **v**eracity. This is best achieved by a dedicated organisational function. Competitive intelligence, the ethical and legal collection and analysis of all source public domain information to create the knowledge and the foreknowledge of the market around us, as a prelude to decision-making action, is the ideal antidote to the VUCA world. Analytical frameworks available to the trained and experienced competitive intelligence professional (and/or business professional) help businesses and other public organisations to control the uncontrollable. They can be pointed to larger businesses' risk management processes, and can be adapted by smaller businesses as appropriate, given the resources and infrastructure available. Analytical frameworks include STEEPLE Analysis, Nine Forces' Model and Scenario Analysis to name just three. The risk management process aligns particularly strongly to 'early warning' competitive intelligence applications – which can do much to help businesses navigate through the VUCA world".*

Andrew Beurschgens

Volunteer UK Chapter Chair
Strategic and Competitive Intelligence Professionals (SCIP).

Clearly, the potential advantage to be derived from accessing primary sources of information about human threats to organisations might be enormous. And if new infrastructure developed for this purpose is also useful for enhancing organisational awareness of threat and opportunity more generally, and indeed for gathering all sorts of valuable insights for organisations, then so much the better. This seems to constitute a compelling case for - at the very least – some lively discussion concerning whether the risk profession should today be manoeuvring to progress the agenda for a new and better risk management with a 'boosted' risk radar, benefiting and requiring some participation of all organisational functions. To reiterate in quite strong terms, such an improved risk management function would need to relax any vice-like semantic grip it might hitherto have either deliberately or unwittingly imposed on the flow of information across organisations, in order to elicit enthusiastic participation from across the organisation. It would certainly need to resist any temptation to process information through the conceptual straitjacket of the risk register. However, it might employ a more flexible alternative to the risk register by maintaining an 'action point register' geared towards providing triage for incoming information; in particular, ensuring information flows to where it will offer value, with the urgency and sensitivity appropriate to it.

A Dangerous High Stakes Activity?

Some may argue, however, that the risk profession should wash its hands entirely of this opportunity to raise its profile. They might claim that active investigation of primary sources of risk information is best left to the area of fuzzy overlap between the competitive intelligence, business intelligence and marketing intelligence domains (to be discussed in more detail later on). Nonetheless there seems to be an empire-building opportunity for risk management here. Risk managers might build cases for overseeing the development of their organisation's risk radar as a means to access more resources and strengthen their case for more participation in top level management decisions. This could be of particular interest wherever there is organisational politics characterised by various organisational functions competing for influence and resources.

ERM CONTEXT FOR BOOSTED RISK RADARS

Speaking further to the above issue of overlapping specialisms, it can also be argued that the risk profession is *already* strongly committed to conceiving of itself as aspiring to subsume various other - sometimes competing - organisational specialisms to serve its master concept of a singular, overarching, early warning risk radar for organisations. Thinking from this perspective, the risk profession is fundamentally concerned to promote such a system because without it there cannot be organisation-wide enterprise risk management (ERM) practice, through which organisations become resilient within their ever-changing and hard-to-anticipate risk environments.

Resilience Context for Boosted Risk Radars

Indeed, the term 'risk radar' itself is one that many already consider a basic principle of organisational resilience. The well-known (2014) *'Roads to Resilience'* report produced by Cranfield Management School on behalf of the AIRMIC risk management professional association, recognises five such principles (the 5 **R**s) as follows:

- The Risk radar which anticipates problems before they escalate;
- Resources and assets which are diversified;
- Relationships and networks which allow risk information to flow freely;
- Rapid response before crises or disasters happen;
- Review and adapt, to learn from experience and make changes.

Notice that the first, third, fourth and fifth **R** in particular describe an ordered sequence of activities. Hence we can conceive of the risk radar as driving the resilience process, or indeed the resilience cycle, depending on which metaphor is preferred. Either way, we can conclude that organisational resilience can only be as good as the risk radar which makes the risk environment visible.

The present document's added suggestion that the risk radar for organisational resilience should lay stronger emphasis on the need for people with highly astute people skills to go out and active explore the social world (thereby 'boosting' the risk radar) should be of interest not just with respect to problems of ill will towards organisations, but also in view of the risk profession's growing interest today in what is variously termed third party risk and partnership risk. We live in what is often termed an 'age of access' characterised by highly complex supply chains and fluidity in co-working and partnering between organisations. Many of the risks which matter most today relate to ambiguity in the relationships which organisations rely on. Surely this places a high premium on robust people skills, and a willingness to go out into the social world, in order to build the risk radar which risk management needs.

Developing the case for the Boosted Risk Radar using ERM's Biological Adaptation Metaphor

The argument that organisations need risk radars which actively explore the social (including the 'human threat') environment of the organisation can usefully be linked to the biological adaptation metaphor for enterprise risk management (ERM). This depicts the organisational entity as constantly exploring the organisational environment and feeding the resulting sensory information back through its 'organisational nervous system' to its 'organisational brain' in order to navigate through that environment, just as every human being relies on their individual senses every moment of every day as they move through their physical, interpersonal, and more broadly social environments.

An opportunity to release further value from the metaphor arises when it is considered that some brains are better at handling social complexity than others. As evolutionary psychologists such as Robin Dunbar have famously emphasised, humans have extra neo-cortical brain development over apes in order to move through, and gain advantage within, much larger and more complex social environments than apes can handle. It could be argued that the equivalent organisational development can occur where ambitious risk management departments aspire to develop and improve ERM towards greater attentiveness to the social-interactive aspects of the risk environment in particular. This of course provides a very loose criterion by which ERM systems might be compared and evaluated. Some are likely to be better at handling higher levels of social complexity than others. Arguably, if risk management departments focus on developing the 'boosted risk radars' which we propose, then in terms of the biological adaptation metaphor they will become less 'simian' and more fully 'human' in their social intelligence. Some further - perhaps even tongue-in-cheek - discussion of what might differentiate 'simian' from 'human' ERM might prove useful – but we leave that to the reader.

This subtly enhanced biological adaptation metaphor can easily accommodate many further individual and social psychology-based metaphors for opening our minds to what it can mean for organisations to negotiate the social aspects of their risk environments. Humans are social animals in ways too rich and diverse to begin to list here with any thoroughness. They explore their social environments by getting to know people around them, by talking to them, by finding out who they like, who they can cooperate with and trust, who can grant them access to resources, and, on the flip side, who dislikes them, who they need to protect themselves against, who most wishes them harm or would obstruct their access to resources, etc. When they find hostility within their social environments, it is wise for them to 'know their enemy' as best they can. They need to keep their friends close and their enemies closer still. Who would deny that organisations might find value in contemplating how they might behave similarly within their organisational environments?

To further enhance this metaphorical insight into how organisations are best advised to move through their risk environments, we might even view the organisation which rely heavily on secondary sources of risk information, and which pursue risk identification largely as a sedentary practice, as akin perhaps to an excessively shy or fearful person who tries to ignore anyone who seems threatening in the naive hope that they will just go away. We could also compare organisations whose risk management relies excessively upon quantitative risk data sets to introverted mathematicians who prefer to steer clear of the rough-

and-tumble of interpersonal confrontation because they find social engagement psychologically awkward or painful, or indeed perhaps in some cases because they fear falling foul of the law if any resulting conflict escalates. A more humorous alternative would be to consider such organisations as rather like 'smartphone zombies' who amble down busy pavements engrossed in the data they call up in front of them, all the while remaining dangerously unaware of their surroundings. These suggestions might each in their own ways stimulate some further fresh thinking about challenges and opportunities for the risk radars organisations need.

Some Key Problems with the Boosted Risk Radar

Nonetheless, there remain some very reasonable additional grounds for resisting this discussion document's central proposal. Perhaps the most obvious problem is that primary sources and others close to them will be highly resistant to disclosure of any usable information concerning threats which they themselves either pose or have a stake in concealing. They may take strong measures to prevent or dissimulate any such disclosures. They may well try to take legal action if their activities are within the law and they believe their privacy has been invaded, or that they have been subjected to some uncompetitive behaviour. If they believe organisations are investigating them using any legally or ethically dubious practices, they may go straight to the media to inflict what reputational damage they can. Boosted risk radars, then, deserve to be regarded as potentially hazardous for their users, just as powerful microwave radar systems can be. An apt metaphor, it would therefore seem.

Taking these factors into account, it is perhaps a fair generalisation to propose that the closer an early warning risk radar gets to the human threats it is designed to protect against, the more practical and ethical problems it is likely to run into. This means there is a problem of diminishing returns and increasing legal and reputational risks from greater resource expenditures on boosting risk radars. A key question, certainly for the risk profession as a whole, and perhaps also for the resource management of risk management activity within each individual organisation, might therefore be phrased in terms of the need to find a 'sweet spot' where risk radars are operating as close to primary sources of threat information as they dare. Risk sociologists and psychologists might explore this as an issue of 'organisational edgework'. They might ask some fascinating questions about how the 'edges' at issue might be experienced by managers as intoxicating exhilaration, anxiety, relief, reward etc. An important consideration here is that although these edges will inevitably be circumscribed by law and other social norms, the edgework actually undertaken may often include at least some transgressive boundary work. Hence questions will frequently need to be asked about whether potential reputational damage and/or legal cost caused by a risk radar might threaten to outweigh the potential benefits from the information gained. The more a risk radar is 'boosted' as we propose, the more serious these issues of risk-adjusted return are likely to become.

Strengthening the Case

Taking stock, it is at least clear that if risk management professionals are to be convinced to take on this challenge, then the first hurdle is to convince them of a very compelling need for it. A key argument must surely be that significant non-routine threats to organisations are, as we suggest above, surprisingly often characterised by competitive or malicious hostile intent. Such threats do seem to constitute a blindspot, or at least as an area of particularly weak coverage within risk management practice, auguring against thorough and even-handed risk identification within organisations. A tantalising evidence base which should spur the risk profession to take this idea very seriously is highlighted by Bruce Wimmer, the Senior Director of G4S Corporate Risk Services. In recent articles he has highlighted G4S estimates suggesting that the costs to businesses of non-cyber related business spying may be more than double those of cyber-related espionage.

This should raise numerous questions within the risk profession. Has it focussed too narrowly on cyber espionage in recent years, when what it really needs to do is to get better at tackling espionage or, more fully, hostile intent towards organisations, in all its forms? Should it shift focus from helping organisations to 'succeed' in their performance ambitions, towards helping them 'win' in their struggles against competitors and other parties which harbour ill will towards them? What synergies between risk management and other corporate functions need to be exploited to allow this to happen? Are cyber and non-cyber related threats best handled uniformly within a singular management process adapted towards active engagement with human threat? Do cyber, non-cyber and other threats involving malice towards organisations each require highly distinct management approaches and their own separately engaged management teams? Any comprehensive risk radar for organisations would need to address these integration challenges by making sensible trade-offs between the advantages associated with highly specialised management teams, and the advantages of management teams with broader competencies. Given that cyber and non-cyber security threats will often be closely interrelated (e.g. passwords may be hacked, but they can also be accessed through blackmail or visual surveillance) it may often be more useful to have management teams in place which provide broad coverage of both and are sensitive to their interrelatedness.

It should at least be relatively uncontroversial to contend that threats involving purposeful human threat to organisations may easily evade detection by conventional risk management systems. We have already touched upon key reasons for this. For one thing, the opaque and fast-changing social world provides limitless hiding places from which hostility may be orchestrated against organisations. Secondly, let's reiterate that corporate security may be viewed as forever playing catch-up with threats associated with complex new technology and vulnerabilities arising with human use of technology. And of course thirdly, hostile intent towards organisations may often be accompanied by proactive threat concealment. This might take the form of security measures, deliberate dissimulation or misinformation, corporate feints, and the like. Taken together these problems pose a serious challenge for what we are calling our 'boosted risk radar' – and yet at the same time they seem to establish a plain need for it.

The Ethics of using Boosted Risk Radars

Seeking solutions to the above problems, the remainder of this document focuses towards exploring what proactive and non-routine engagement with purposeful and targeted human threat can mean in simple practical terms. Yet we are also attuned to the need for the risk profession to develop in ways that not only protect but also grow its ethical integrity. The risk profession is, after all, more fully the risk and insurance profession.

Ethical integrity is and has always been a bread and butter requirement for insurance practitioners, given the importance of both seller and buyer trust in insurance contractual relations. Concerned to promote the risk and insurance profession's ethical credentials, a recent (2016) Centre for Risk Research discussion document argued for 'cultures of prudence' which foster ethical-cultural and ethical-psychological context wherein organisational risk management can flourish. Set within the virtue ethics paradigm, the document argued for a vision of what it means to have 'ethical' risk management which emphasised the need for proactive and sometimes very risky truth-seeking and truth-expressive behaviours, such as speaking truth-to-power in acts of whistleblowing and admitting failure or weakness in order to improve the flow of risk information. This vision of ethical risk management based on virtuosity also seems to fit very neatly indeed with the present document's advocacy of active and investigative forms of risk identification. An important consideration here is that virtue ethics arguably comes into its own as an approach to ethics under circumstances where very difficult individualised decisions need to be taken, where boldness is often a prerequisite for success, and where frameworks of rules and regulation offer less practical guidance value than individual conscience and resolve. Arguably ethical culture and ethical psychology become especially important practical considerations if the profession is to grow its ethics credentials while also getting smarter at sourcing vital risk information from sometimes hostile parties resistant to disclosure. Much more consideration of the best ethical frameworks to allow this to happen is certainly called for.

A Respectable Front for Corporate Espionage?

The challenges facing any effort to build ethical risk management under these circumstances are, however, formidable. Behavioural imperatives can crush ethical conscience. Dishonesty might often seem like a very tempting option if it allows high value risk information to be accessed. Another less obvious but nonetheless important challenge arises with the fact that, while our concept of a 'boosted risk radar' refers in part to a capacity for counter-espionage, we should also appreciate that any risk management capacity which subsumes capacity for counter-espionage must also provide some capacity for espionage itself.

KEY QUESTION: Does the risk management profession really need to broaden its scope by engaging with the many dangers associated with going out into the world to gather sensitive risk information? Might this prove an extremely dangerous direction of travel insofar as it entails developing capability that can be used for engaging in, as for protecting against, corporate espionage? This discussion document is intended to get risk professionals thinking more – and perhaps in some cases, for the first time - about these important questions.

Hence there arises the possibility that risk management could conceivably be used as a respectable front for organisational resource geared towards espionage. After all, irrespective of whether espionage or counter-espionage is at issue, the same sensitivity to legal and reputational checks on covert activity, and the same awareness of human and technological vulnerability, supply the essential knowledge and skills base. There may exist similar temptations to subvert ethical norms, or even to break the law, when exploiting human or technological vulnerabilities, no matter whether the intention is to cause harm in acts of espionage or to access vital risk information in acts of counter-espionage. Taking stock, some of the key ethical challenges surrounding this document's advocacy of a more proactive future for the risk profession will be explored later in this discussion paper.

The practical suggestions for improving risk management to be made in the remainder of this discussion document are relevant for dealing with all sorts of hostility towards organisations, stemming for example from NGO activism and orchestrated public campaigns, from disgruntled (former) employees or shareholders, or conceivably even from regulators. However it is written with cybersecurity and corporate espionage threat foremost in mind. Opportunities for risk management to become bolder and more proactive against such threat are explored in the light of how risk managers might learn from the competitive intelligence profession to ensure that protective investigative actions remain within the law and adhere to ethical standards so as not to endanger corporate reputation. Arguably this borderland between the risk profession and the competitive intelligence profession is too often neglected. This discussion document is partly intended to foster greater interest in how the two can work together and find efficiencies within organisations.

‘RISK INTELLIGENCE’: THREE MEANINGS

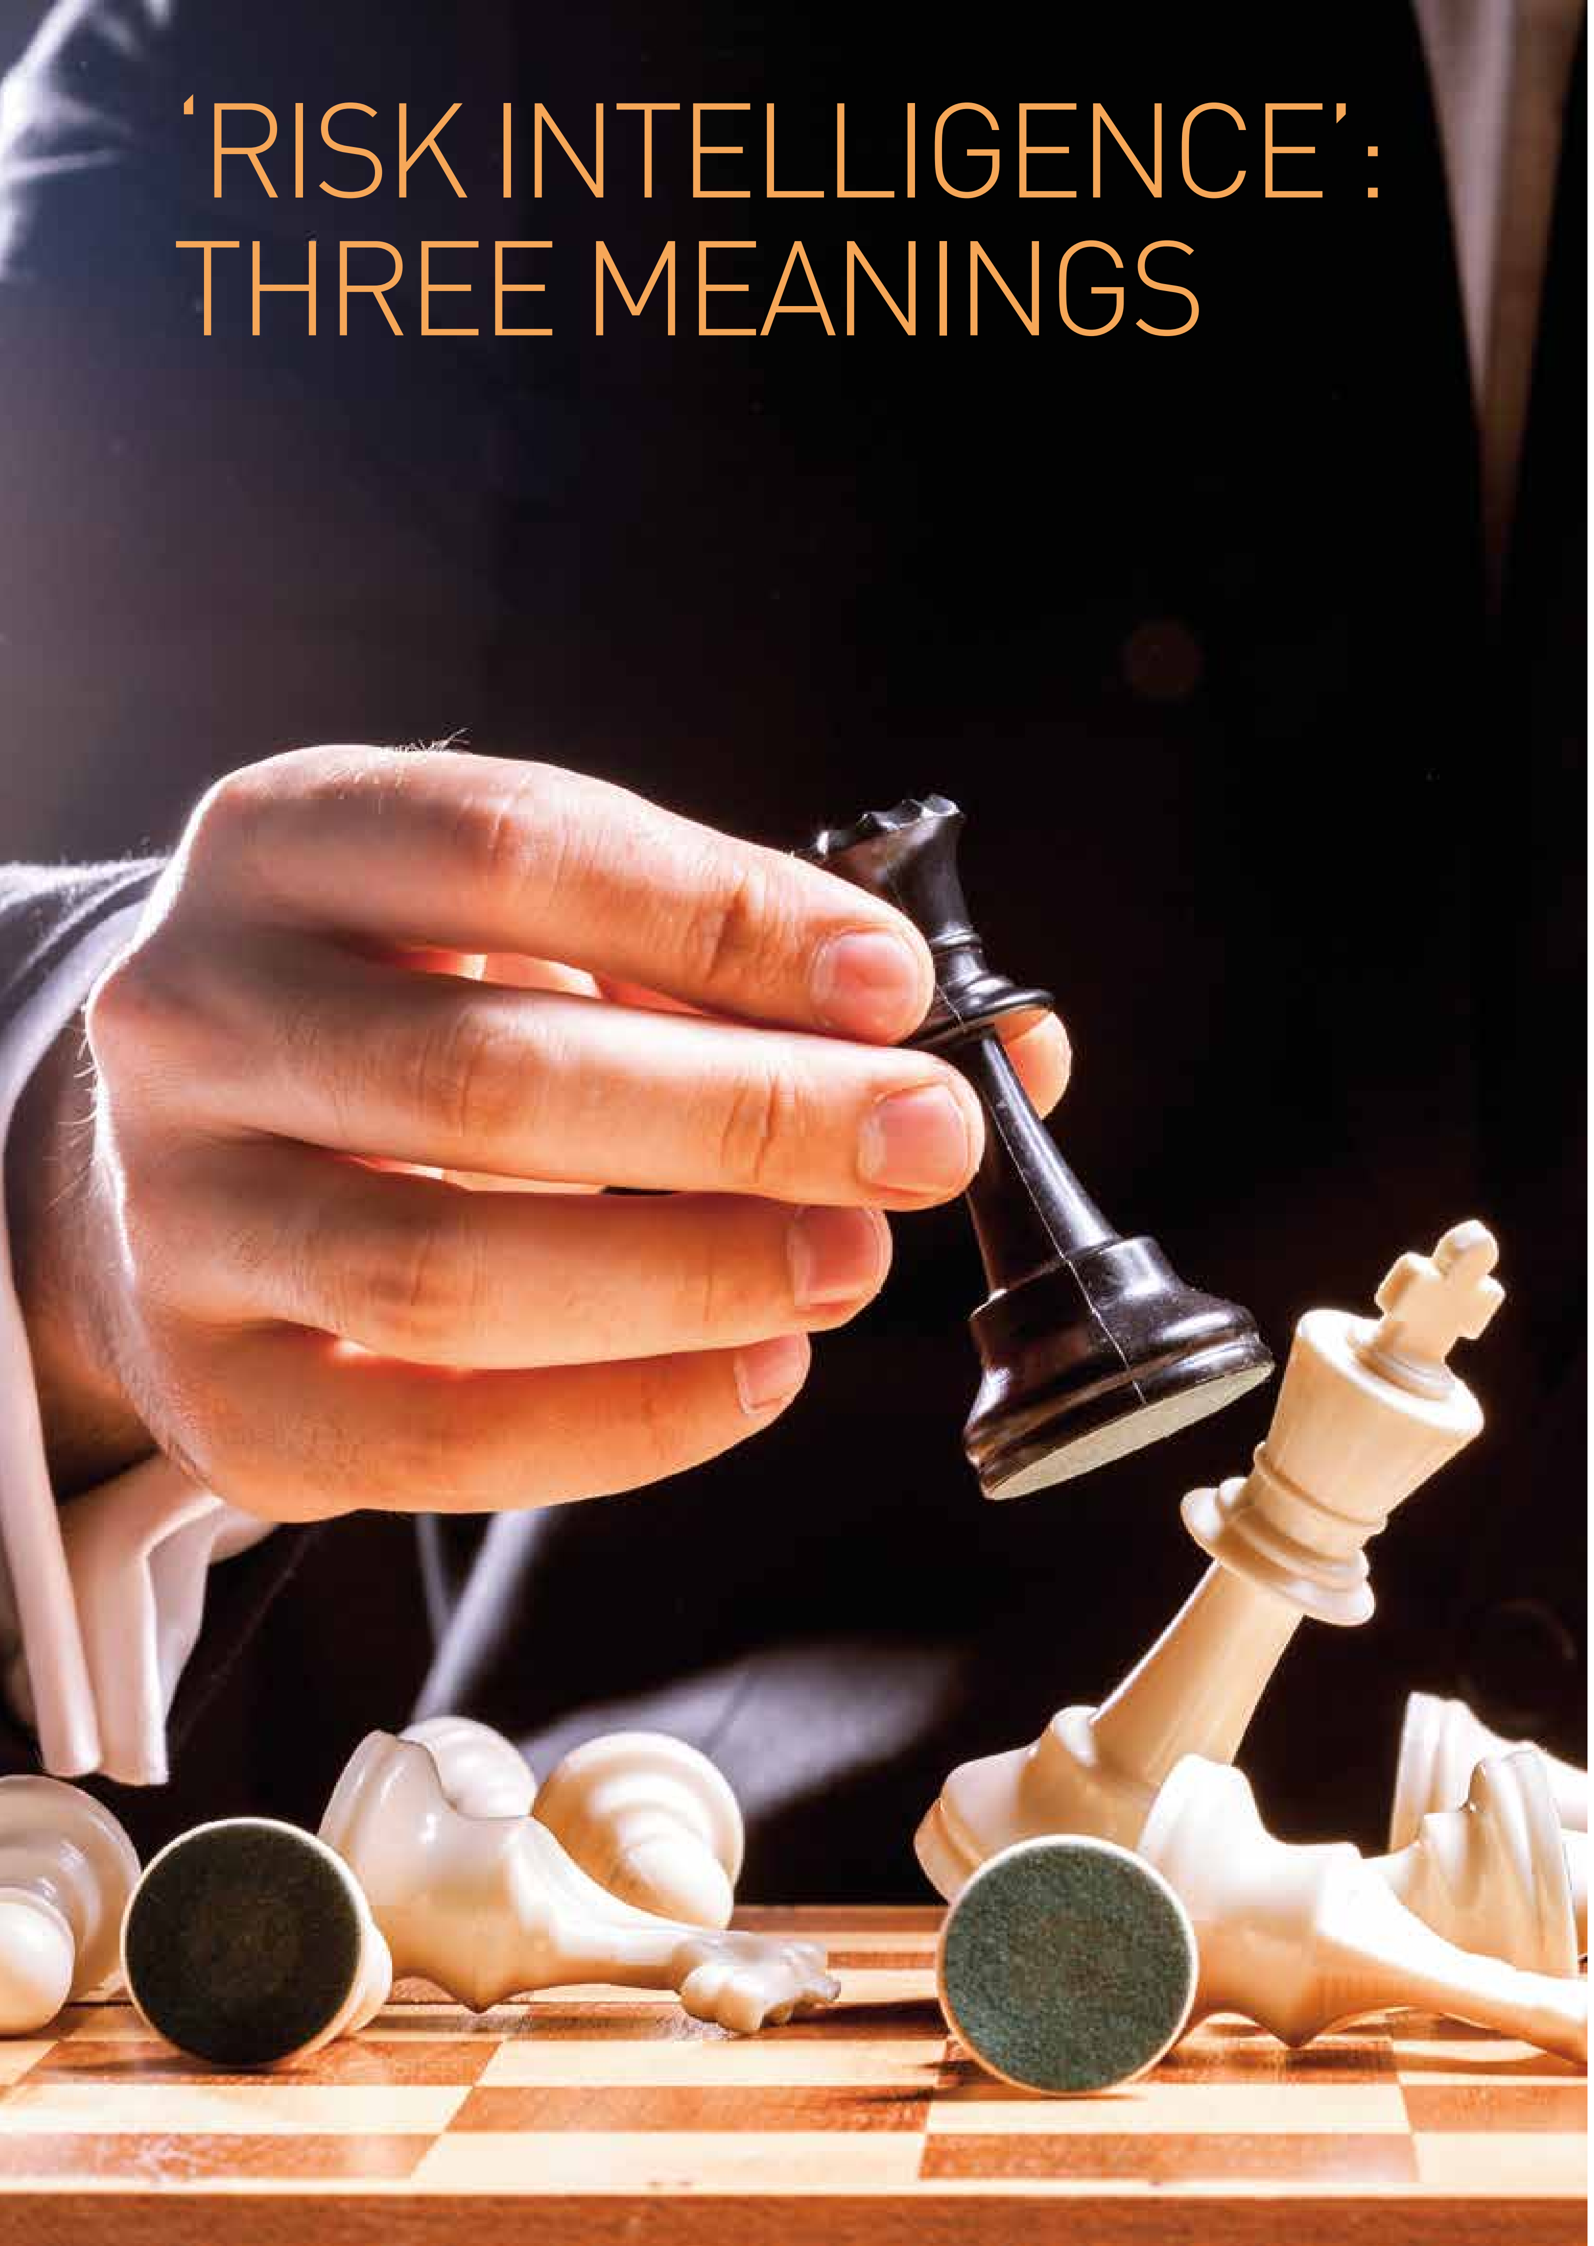

Blending the concerns of risk management and competitive intelligence, we get the term ‘risk intelligence’ which has already received rather different treatments from various authors. Our proposal for the risk profession to develop what our introductory section called a ‘boosted risk radar’ is phrased henceforth in more practically implementable terms of developing ‘risk intelligence’; that is, we explore what risk intelligence can mean and what the practical development options are for organisations. Terminology matters a great deal in risk management.

TERMINOLOGY MATTERS: An important challenge facing risk managers is how to be efficient and effective as a very thinly spread organisational resource. A key consideration here – perhaps under-recognised in most risk management guidance – is that for risks linked to objectives to be managed effectively, it is unwise for the management of each risk to become remote from the management of its associated objective. Where there is careful coordination, we can conceive of risk management as being integral to managing an objective, such that risk management becomes an obviously necessary part of simply ’managing’. By contrast, where there is poor coordination we can instead begin to conceive of risk management as a ‘bolt-on’ activity, vulnerable to being perceived as unnecessary added bureaucracy.

Mindful of the need for such coordination, we can further conceive of risk management activity as intrinsically social in character, involving a coming together of different forms of essential professional experience and knowledge from different sources. Hence shared risk management terminology matters greatly. Accordingly, the present discussion document draws upon the view that important conceptual innovations in risk management should pass basic tests of clarity and simplicity in how they are understood and applied. It further holds that where new developments are driven through injections of new thinking about the nature and scope of the discipline, they should ideally gain acceptance as simple common sense which everyone can agree on. Once incorporated within risk management thought and practice, people should wonder why they were not thought of earlier. It is hoped that ‘risk intelligence’ fits this bill well, as indeed does our development of the concept of a ‘boosted’ risk radar to extend the already widely-used risk radar concept. Might it even be argued that such new terminology offers risk management some glamour that will make people across organisations more likely to want to participate in it?

Whereas risk identification is something which can all to easily be desk-bound, risk intelligence offers an alternative form of words where we naturally contemplate active and energetic intelligence gathering activities.

Thus, we look to ‘risk intelligence’ to supply the pivotal terminological innovation for driving best practice forward. A definition is therefore appropriate. Seeking a working definition of risk intelligence which is clear, and yet whose multiple interrelated meanings are sufficiently rich as to enable risk professionals to rethink best practice, we do not choose between but instead synthesise thought-provoking definitions which have arisen within previous writings on risk intelligence.

We argue that risk intelligence deserves a central place in the risk management lexicon when understood as:

1. *A professional aptitude;* thinking from this perspective, the ‘risk intelligent’ risk manager appreciates the gravity of the challenge in developing a view of risk that is both thorough and detailed, objective and kept up-to-date; crucially, this ‘risk intelligent’ professional aptitude can be expected to manifest, in part at least, as a strong commitment to proactivity tempered with ethical integrity in both the pursuit and critical interrogation of risk information; fundamentally at issue here is aptitude for intelligent – and in particular, socially astute as well as legally and reputationally sensitive - engagement and interaction with potentially hostile primary or near-primary sources of risk information (i.e. that practice which we referred to earlier as corresponding to hazardous and high stakes ‘organisational edgework’);

2. *An organisational role and its output;* thinking from this perspective, the *risk intelligence role* centres upon organisational processes focussed flexibly towards gathering and churning out *risk intelligence information* for use within various management contexts; the flexibility at issue here can be understood as relating not just to flexibility in the use of various possible methods of gathering information, but also flexibility in how that information is triaged onwards so that best use is made of it while privacy and confidentiality issues are respected;

3. *A hallmark of mature organisational risk management;* thinking from this perspective, a vision of mature risk management practice arises where highly developed professional aptitudes (1, above) align with flexible organisational process capability (2, above) to create the *risk intelligent organisation* geared towards responding urgently and effectively to all sorts of non-routine risk issues involving purposeful and targeted human threat, over and above all of those risk issues conventionally handled by risk management. More fully, here we may also conceive of the risk intelligent organisation in more metaphorical terms as relying upon its ‘boosted risk radar’ to navigate all the threats and opportunities which loom up before it within its risk environment.

To summarise, then, the remainder of this discussion document argues that a new professional focus on *risk intelligence* - based on the above definitions which are intended to accommodate and integrate the full range of senses in which the expression may be used within everyday language in organisations - might well prove the best way to grow the risk management role within many organisations today. Its ideas should be of interest to academics, students and practitioners of risk management alike.

The challenge at hand, then, is to explore the above interrelated meanings of ‘risk intelligence’ in order to outline their practical implications for what risk management can mean, and for how risk managers might consider going about their workaday lives differently. We will work towards concluding that the risk profession should give much more attention to the challenge of developing risk intelligence in the third and most ambitious sense of the term. To reiterate, it will be argued that such effort might usefully focus on exploiting organisational efficiencies by combining a number of management processes to form a singular integrated ‘risk radar’ which can be ‘boosted’ to engage with purposeful human threat to organisations while also better fulfilling the organisation’s broader information gathering needs.

Terminology matters a great deal in risk management.

MEANING ONE: RISK INTELLIGENCE IS MANAGING RISK INTELLIGENTLY

Studies treating risk intelligence as a measurable aptitude use the abbreviation 'RQ'. They define 'high RQ' individuals by their ability to estimate probabilities well without letting false certainty cloud their judgments..

PROFESSIONAL ATTITUDE AND KNOWLEDGE

Dylan Evans' RQ scale has been reproduced widely on the internet. It seeks to measure a person's ability to estimate probabilities accurately. Its constituent items achieve this by asking respondents to rate their level of certainty in the truthfulness of some proposition. Hence RQ diminishes with false confidence in wrong answers. For example, one question is *"The "Sorrows of Young Werther" was written by Dante"*. Someone not knowing that Goethe is the true author, who gives a 100% certainty rating to support the statement's claim of Dante's authorship, would achieve a reduced RQ score as a consequence.

It is certainly tantalising to conceive of the RQ scale as calling attention to a soft psychological skill, based on awareness of the self and its limitations, which is vital for effective risk management. Nonetheless, measurable RQ seems to cover just a part – albeit an important one – of the intelligence required for managing risk. Anyone wishing to scope their organisation's 'risk intelligence' in a more broad sense might take a flexible approach focussed towards the needs of their organisation. For example they might aspire to collate and then showcase risk management professional knowledge developed through individual and/or collective experience within their organisation. They might develop this off the back of an organisational learning training event where *'what does risk intelligence mean to you?'* has provided the focus for discussion. The resulting distillation of professional knowledge might then be communicated as widely as possible within that organisation, both in written risk management guidance documentation and in verbal advice given.

Let us further consider that the scientific measurement approach to RQ suffers from narrowness in what it measures. Opening out some parallels with the RQ construct, we could first of all argue that effective risk management also requires 'IQ'. This naturally sets us thinking of cognitive problem-solving skills like pattern recognition and why they might often prove useful. However we could also expand our scope to consider emotional intelligence (EQ), which could lead us to reflect upon why various affective biases often matter for risk management (for example during affective coping with crisis or when relying on heuristics to take decisions under uncertainty). We could then bring in cultural intelligence (CQ) and contemplate the risk profession's rapidly growing interest in cultural contexts likely to advantage or disadvantage effective risk management. Setting RQ within this much broader multiple intelligence context seems useful by providing a loose yet thought-provoking framework for contemplating the quick thinking and self-awareness which risk managers are likely to need, and to aspire to promote in others, in order to succeed in their roles. This might also supply a useful framework for contemplating the particular 'spikes' (an HR term denoting key aptitudes) within risk intelligence required for 'boosted risk radars' to operate successfully. Clearly risk intelligence emphasising cultural and social skill is likely to be more valuable than RQ defined narrowly as aptitude for estimating probabilities, where engagement with primary sources of risk information is at issue. More dialogue between the risk and competitive intelligence professions seems appropriate in order to clarify the spikes needed.

Ethical Rules of Thumb for Risk Intelligence

One of the likely benefits of such dialogue is a clearer understanding, within the risk profession, of the ethical *dos* and *don'ts* of engaging with various primary or near-primary sources in order to elicit risk information. Speaking very generally, no matter what the legal issues involved, organisations may well pay reputational penalties if their information gathering activities are exposed in the media and publicly perceived to breach ethical conventions. Hence the ethical intelligence (our second form of 'EQ') which risk intelligence requires needs to be flexible in its view of what it means to be ethical. Some good ethical 'rules of thumb' are:

From virtue ethics:
Can someone be open, honest and proud of their actions?

From the golden rule:
Are you treating others as you would wish to be treated yourself?

From Kantian universal law:
What if everyone acted as you are acting?

From Utilitarianism:
Will benefits outweigh harms, for the greater number?

From Rawlsian Justice:
What if there were a 'veil of ignorance' preventing the decision-maker from knowing whether they would be among the people impacted by an action?

From Lockean rights:
Will the action violate basic rights of those affected?

Ethics Codes for Risk Intelligence

Further, more practical, guidance on the ethical intelligence that risk intelligence requires, can be derived from codes of conduct developed within the competitive intelligence profession. Here the practical challenges arising with our 'boosted' risk radar concept start to become very clear.

The approach taken by the 'Strategic and Competitive Intelligence Professionals' global professional organisation on their website www.scip.org is to offer a succinct high level code of conduct. As one might expect, this emphasises the importance of honesty, the need for compliance with all laws, and the need to avoid conflicts of interest. It also incorporates a more specific need to "faithfully adhere to and abide by one's company policies, objectives and guidelines", which should be an important consideration for anyone whose initial view of the profession is that it might sometimes operate in the grey areas between espionage and rectitude. Of further interest for anyone concerned to understand the ethical credentials of the profession is the specific need to "accurately disclose all relevant information, including one's identity and organisation, prior to all interviews". Notably, then, the effect of the code is to accentuate the profession's specialisation in legitimate intelligence gathering free from all deceit and subterfuge.

The SCIP website also offers some interesting FAQs which illustrate how the code's ethical requirements can be applied to various circumstances. Here are some examples:

Should Private Investigators (PIs) look at files discarded by a company? The SCIP guidance urges consideration that such activity is likely to be in breach of most organisations' codes of ethics. They recommend applying a 'red face test' (debatable in meaning but perhaps best understood as a 'shame' test contrastable with our test 'from virtue ethics' which might be considered a 'guilt' test) by asking the company employing the services of the investigator to consider whether they would be comfortable with a newspaper report revealing how the files had been sourced and used.

Is it ok to source information from another company by calling them and pretending to be a customer or a student? The SCIP guidance highlights an ethical imperative to disclose all relevant identifying information, further noting that misrepresentation may sometimes be illegal. However it further mentions that 'mystery shopping' is an ethically acceptable practice provided it does not violate retailer-specific rules.

Is it ok to buy or otherwise obtain a password to access a competitor's website? The SCIP guidance advises that this is likely to be illegal, and to violate most organisations' codes of conduct.

Should PIs be hired to investigate the Principals of Companies? Here the SCIP guidance advises that all such activity should 'begin' with public domain information, however further investigations may sometimes be legally and ethically permissible. Some further commentary in addition to what the SCIP say may be helpful here. This is a fascinating ethical problem domain, because there are many subtle techniques which companies can and do use to source intelligence information from Principals at rival companies. For example, If company A wishes to know more about Managing Director X of company B, they may send their loyal employee, Y, who was a friend of X during their undergraduate days, to attend a corporate event where they know X will be present. The resulting ethical ambiguity experienced by Y when meeting X may then prove extremely challenging. Also consider the above conditions repeated, only this time X and Y are senior sales managers. Even what may seem like a perfectly innocent conversation, where X tells Y about their recent travels to exotic destinations, may constitute valuable competitor intelligence by unintentionally revealing where commercial opportunities lie. Furthermore if X gives Y cellphone or social media identity information to allow them to stay in contact, then such information might later be exploited to monitor X's changing geographical locations.

Should Companies hire former government operatives to direct their intelligence operations? The SCIP say it makes good sense to hire people experienced in intelligence processes. Indeed, moving beyond what the SCIP say, it might indeed by claimed that government operatives with experience in military and national security intelligence in particular, may prove to be important assets to organisations. As we illustrate in the next section, long established military intelligence theory and practice can supply a very rich seam of diverse ideas for improving corporate intelligence activities.

Is it ok to post a fake job advert to enable discussions with applicants from competitor companies, or to apply to competitor company jobs simply to learn more about them in the interviews? The SCIP's strong advice is to avoid such practices as they may well violate organisational codes of conduct and local laws. They also note that the strong weight of opinion within the profession is to find such practices unethical. This may lead us to consider that greater dialogue between risk management and competitive intelligence is called for, not just at organisational level, but also at institutional level through dialogue between professional associations where the weight of opinion within both professions can be considered as a basis for practical dialogue focused towards improving best practice at their professional interface.

What information can be gleaned from new employees who formerly worked for competitor companies? Here the SCIP mention the ethical imperative to remind employees of their duty of confidentiality to their former employer. If any information becomes available in what may constitute a breach of such confidentiality, the matter should be referred to management or legal counsel.

A further FAQ discusses the need for experienced intelligence professionals to work with corporate lawyers and with ethics and compliance staff, all of whom need to be sensitively attuned to ethical and legal problems encountered by competitive intelligence professionals, in order to develop codes of conduct tailored to meet the needs of specific organisations. Of course, we can imagine these coordination challenges to be more complex where a broader range of organisational functions must come together to participate within what the present document is calling a 'boosted risk radar'.

Practical Steps for Developing Risk Intelligence

Perhaps organisations should focus effort towards clarifying and supporting the 'risk intelligence' they need, as a means to establish guidance to support their boosted risk radars. Consider for example how a risk management department might launch a campaign and hold workshops within an organisation urging employees to think about and discuss what particular spikes within risk intelligence ought to matter most for them and their colleagues. Arguably this could be a very useful exercise, even where it elicits differing answers from across an organisation. Our point here is that the 'risk intelligence' concept might plausibly be regarded as having enormous untapped potential, even where its sole use is to encourage self-reflection across the broad gamut of organisational risk management practice, which might in turn produce benefits too varied to be summarised here.

Of course, any organisation acting in this way to encourage employees to incorporate 'risk intelligence' within their own personal understanding of what it means to maintain a professional attitude to their work, may derive particular benefits from asking employees to consider their personal 'risk intelligence' within the broader context of the need for organisations to be more risk intelligent, through more proactivity in gathering risk intelligence in particular. In other words, inviting reflection on what 'risk intelligence' means, and steering such reflection towards stronger appreciation of the term as a multi-layered metaphor, might be a great way to draw attention to the problem whereby risk identification effort is all too often a desk-bound safe space activity. This could be a good bottom-up consultative way to tease out all the ethical and other issues that need to be addressed when creating or improving guidance for the boosted risk radar.

MEANINGS TWO AND THREE: RISK INTELLIGENT PROCESSES AND RISK INTELLIGENT ORGANISATIONS

This section assumes the reader's familiarity with the stage-by-stage logic of the risk management process, as found in the (2002) AIRMIC/ALARM/IRM Risk Management Standard, and more recently with some minor changes in the (2009) ISO 31000 guidelines. Throughout the discussion which follows, we selectively juxtapose some basic essentials of the risk management process, as illustrated by these guidance documents, with various similar organisational processes and related activities which all deal with the gathering and processing of intelligence information. Our purpose will be to highlight some learning examples and opportunities for consolidation. This will enable us to outline what a consolidated risk intelligence process, and hence a risk intelligent organisation, might look like.

We cannot undertake exhaustive comparative studies of specific processes. However we can whet the appetite for such studies by illustrating how ideas and practices relating to a range of closely related organisational processes can all be drawn together within a consolidated risk intelligence process geared towards systematic and sustained reaching out much further into the social world to engage with hard-to-access sources than traditional risk management has hitherto dared. A simplifying starting point is to appreciate that business intelligence, marketing intelligence and competitive intelligence processes overlap considerably and are often considered to amount to much the same thing. This is partly because of their common ancestry in military intelligence processes.

Learning from Competitive and Marketing Intelligence Processes

Competitive intelligence processes can straightforwardly be consolidated with general risk management processes where they are already risk management processes to begin with, such that they observe the risk identification, analysis, evaluation and control staging of the risk management process. Consider that competitive intelligence processes are often conceived through a risk lens as centring on risk management of the competitive environment. Such risk-focused competitive intelligence processes may differentiate themselves from general risk management by requiring consideration of a particular category of risk, *'industry dissonance risk'*, as a basis for exploring what a company can do to achieve competitive advantage.

When competitive intelligence is undertaken with a marketing focus, it may employ a whole host of further very specific risk subcategories under such broad risk category headings as *'market risk'* (based on uncertainty over whether new products will have sufficient market sizes), *'product risk'* (based on uncertainty over whether the product strategies will seize sufficient market share), and *'profit risk'* (based on uncertainty over whether that market share will deliver sufficient profit). For further reading on this, a (2005) text by McDonald, Ward & Smith entitled *"Marketing Due Diligence"* is particularly helpful. It discusses six risk subcategories under each of these three risk category headings and urges systematic consideration of the resulting eighteen point probabilistic risk reports within strategic decision-making. All of these risk categories and subcategories can be used to focus creative imagination during risk identification. Hence competitive intelligence cycles, focused on processing such risks, are very strong candidates for being consolidated within risk intelligence processes. Simply making explicit provision for consideration of the above risks within general risk identification might in itself provide a spur for risk management to remould itself upon the proactivity of the competitive intelligence profession. Furthermore, by enhancing the due diligence value of risk information, this would help improve the standing of risk management in its association with strategic decision-making.

Adopting the language of competitive and military intelligence processes

Yet perhaps an even more important consideration when advocating for a risk management process more integrated with - and enhanced through learning from - other information processes, is that the latter commonly refer to intelligence 'gathering' or 'collecting'. The equivalent stage within risk management processes is 'risk identification'. These terms convey markedly different connotations, which we can consider significant because, as Whorfian linguistics and simple common sense tell us, words influence how we perceive and act. 'Collecting' and 'gathering' are words suggesting movement. 'Identification' connotes the observer who can see clearest when stationary. Consider, therefore, the different effects which the words *'let's talk about identifying risk information'* and *'let's talk about gathering risk information'* might have when spoken at a meeting. Which has more meaning and power as a call to action?

This emphasis on 'gathering' or 'collecting' information, which is key to intelligence practice within organisations, has deep roots in decades of intelligence theory and practice developed by the military. Whereas the ISO 31000 risk management process begins with an *'establishing the context'* stage prior to its *'risk identification'* stage, US Military Joint Publication 2-0 moves through parallel *'planning and direction'* followed by *'collection'* stages. The military process captures much of the logic of ISO 31000's first two stages in its view of *'planning and direction'* as being concerned with specifying what information is necessary for the successful achievement of specified military objectives, so that follow through *'collection'* activities can then take place at the next stage. However, in the military intelligence process the intelligence *'collection'* at issue conjoins logically with the prior stage's requirement for *'direction'*. In the equivalent risk management process such directed elicitation of active information gathering effort may not necessarily take place, unless *'establishing the context'* has explicitly concerned itself with identifying important areas of uncertainty so that 'risk identification' then becomes compelled to gather information to resolve that uncertainty.

However, this may not happen. The *'establishing the context'* stage is widely understood to be concerned with exploring various relations, structures, factors, drivers, trends, etc., which compose the *'internal'* and *'external'* context of the organisation, so that risk identification effort can then be focused towards specific areas of concern. More practically speaking, the emphasis may simply be on deciding *who* to invite to risk identification meetings. Hence the first two stages of the ISO 31000 process are not linked by the same ask-the-questions-then-go-and-find-out-the-answers logic as are the first two stages of the military intelligence cycle. As a consequence, the two processes invite very different management approaches at their respective second stages. This need not have been so. Perhaps the ISO 310000 'establishing the context' stage would have been more effective in spurring active investigative effort at stage two had its wording dealt with establishing contextual *uncertainties* so that stage two obligations to go out into the world and find answers might then favour enhanced risk identification.

Notably, however, neither of these processes reflect a resilience philosophy based on a perceived need for ceaseless scanning activities permitting organisations to respond urgently to unexpected threat. Each has a first stage which could conceivably be theorised in metaphorical terms as a decision to situate a risk radar on some terrain where it can most be effective. However this falls short of our earlier discussed need for risk radars to provide early warning of issues which take organisations by surprise. Risk radars need to be able to detect 'black swans' flying towards them as soon as possible, so to speak.

Perhaps an enhancement which both processes may benefit from, is for stage one to specify some practical balancing of stage two activities between risk anticipation effort which directs the risk radar to pay particular attention to certain areas, and some further resilience effort involving ceaseless monitoring which sensitises the organisation to issues it has not anticipated. Stage two activities might then incorporate some commentary on how that balance has been struck and how well it has worked. Perhaps one big advantage of injecting sensitivity towards 'anticipation vs resilience' issues into information processes in this way, is that it potentially opens up a valuable feedback loop at stage two, running back to stage one, whereby managers are quickly informed of any dissonance or failure of imagination which has weakened stage one. The effect of this would be to ensure that the 'boosted risk radar' which we propose in this document is better able to direct and inform the whole risk management process, as opposed to simply providing some stage two enhancement to that process.

Learning from Military Intelligence Practice

Next we turn to a further, this time much simpler and easier-to-implement, example of how risk intelligence can learn from military intelligence. Within the risk management process, risk analysis typically focuses on estimating probabilities and consequences for risks, so that risk evaluation can then consider each risk's significance with reference to pre-established criteria such as risk appetite or tolerance. The more-or-less parallel military practice evaluates collected items of intelligence by estimating the reliability and credibility of sources for filtering and weighting purposes. The well-known 'Admiralty' or 'NATO' classification system uses a lettering form for reliability of sources and a numbering form for credibility of what the sources say, as follows:

Reliability		**Credibility**
Completely reliable	–	Confirmed by other sources
Usually reliable	–	Probably true
Fairly reliable	–	Possibly true
Not usually reliable	–	Doubtful
Unreliable	–	Improbable
Reliability cannot be judged	–	Truth cannot be judged

If we are to re-envision risk management as a practice invigorated by the active information gathering skills of the competitive intelligence profession, it follows that simple reliability ratings for sources, and simple credibility ratings for the information these sources provide, may offer value for enhanced risk management processes. During the risk identification stage of the risk management process, reliability and credibility ratings might be used to score risk information derived from conceivably any secondary or primary sources. These scores might then be used to filter matters for consideration. Later on, during the risk analysis stage they might offer further value as touchstones for estimating risk likelihoods and impacts. And of course, anyone given ownership of a risk control later on within the risk management process would naturally want to reflect on the reliability of sources and the credibility of what they say, as indeed would anyone wishing to use the risk information yielded by the risk management process for decision-making purposes.

Still in this *'learning from the military'* vein, and this time with a concern to understand in very practical terms how risk intelligence can be gathered, we might consider some active intelligence gathering techniques which businesses have increasingly been using (especially for cyber security) over the last few years, and in whose provision risk management consultancy services sometimes specialise. The techniques in question are usually referred to very generally as *'red teaming'*. As Micah Zenko explains in his (2015) book *"Red Team: how to succeed by thinking like the enemy"*, the practice comes in three basic forms which can easily overlap in practice. To conclude the main part of this discussion document we explore how each might be integrated within a consolidated risk intelligence process.

1. RED TEAMING AS CONFLICT SIMULATION

Most people who have heard of 'red teaming' will associate it with military and business wargaming. Red teaming considered specifically as conflict simulation, entails rehearsing conflict events which might happen in real life in order to achieve greater understanding of some human threat environment. Hence we might value such activities as techniques for sensitising organisational risk radars to particular threats within conflict environments. These conflict simulations could be structured at stage one of the risk management process and then undertaken at stage two. Where they involve exploring how competitors might respond to new products, they could further be considered a means of linking product cycles and risk management cycles. However, flexibility is important when situating these simulations within organisational contexts because the information can be useful for many purposes. All sorts of plans and controls can be tested and developed. Participants can develop their team working skills and their second nature responses to threat. Possible unintended consequences of major strategic decisions can be explored before the decisions are taken.

Expert Commentary: *"My experience in the US military and its employment of Red Teaming provided valuable insight during military operations planning. By utilizing a Red Team familiar with tactics and techniques of a potential advisory, the Blue Team played out its operational plan within a given scenario. After a series of moves and reactions, the results of each move could be further analysed, and the results utilized in future iterations of the physical military plan.*

Business could employ military Red Teaming to provide decision makers with potential outcomes, weaknesses or risks that may have been overlooked by individuals with intimate knowledge or personal ties to a plan. This provides a second view point that can be tested which may reduce risk or be used to identify a new solution or path to take on a project.

From my experience, Red Teaming caused us to consider new ways to deploy a force as compared to what we have experienced in the Cold War or even in our recent history in Counter Insurgency and fluid operations that we face in Syria and Iraq. The results of the Red Teaming exercises are leading to advances in troop employment, tactics and techniques and equipment by forcing institutional change in how planners deploy and project forces within military planning activities".

Ross IV, Richard R., MSgt (Ret) USAF.

Former Red Teamer with HQ US Air Forces Europe, Directorate of Logistics, Engineering and Forces Protection, Engineer Division.

The term *'red teaming'* itself derives from the colour-coded nomenclature used by the US military for wargaming during the cold war, where *'red teams'* played the attackers and *'blue teams'* played the defenders. Various other teams also play important roles. *'White teams'* set the parameters and ensure the red and blue teams interact accordingly. *'Purple teams'* shape the learning narratives by listening carefully to the perspectives of all participants. *'Green teams'* extend the simulations to consider conflict impacts upon third parties. In military wargaming these third parties may be neutral military powers, or refugees and other civilians in conflict zones. In business wargaming they may be an organisation's stakeholders. Hence *'green teams'* can be a useful focus for organisational effort to render risk management more sensitive to stakeholder concerns.

Professional Association Commentary : "One important competitive intelligence framework is called war gaming, or competitor role play. This technique can use 'red teaming'. A war game is neither a war nor a game but a structured forum to assess a market's competitors and other stakeholders' (customers, regulator, lobbyists) likely moves in relation to a specific market event or series of events. From a commercial perspective a wide number of stakeholder groups are often in play, therefore making the more binary focussed 'Red Team' concept less relevant.

One of the benefits of war gaming is to gain a better and shared understanding of the competitive arena, therefore making it ideal to the risk management process. Other benefits include sensitising the management war game participants and wider business to competitive moves, identifying 'blindspots', providing the starting point for monitoring competitive risks, and of course we should not forget the team building benefits. Given that the average CEO tenure in the United Kingdom is 4.8 years, there is a strong argument for using wargaming for new incoming CEOs to stress test the robustness of the strategy he/she is inheriting and/or use it to craft a new one. However, operating in a VUCA world implies that this facilitated analytical practice should be taking place more frequently".

Andrew Beurschgens

Volunteer UK Chapter Chair
Strategic and Competitive Intelligence Professionals (SCIP).

Sometimes also included in these simulations are *'tiger teams'*. Here the strength and agility of the tiger symbolises the discretionary powers which can be granted to participating teams whenever it is deemed that they are best placed to follow up on the learning outputs from a simulation. However conflict simulation outputs are often advisory rather than determinative for real world follow up. The case for tiger teams is stronger where the simulations produce technical knowledge best acted on immediately – for example where urgent action is required to address cybersecurity vulnerabilities.

2. RED TEAMING AS VULNERABILITY ANALYSIS

What sets vulnerability analysis apart from the above exercises is that they are less readily described as simulations because they focus on testing real vulnerabilities. Vulnerability analysis is a form of wargaming where the red teams gain advantage from surprise in choosing how, when and where they attack. There are many famous examples of this. For example the US Department of Homeland Security's Office of the Inspector General publishes statistics on the success rates of their red teamers in breaching security measures multiple times at specific airports. Posing as ordinary travellers, these red teams are often successful in smuggling fake weapons, drugs and explosives through airport security. This has the effect of leveraging security improvements through public embarrassment. No airport security service would wish to neglect a security vulnerability which has been exposed in national media.

It is also well known that military and business organisations employ hackers to test cybersecurity arrangements. For example, Peiter "Mudge" Zatko joined google in 2013. He had formerly been associated with the "Lopht" hacker collective whose members testified before US Congress in 1998 that they could shut down the internet in thirty minutes. Lopht's commercialisation during the 1990s towards providing cybersecurity services reflects the good commercial sense of the poacher who discovers there is more money to be made as a gamekeeper. Lopht's history illustrates well that there is value in hiring 'white hat' hackers with murky pasts in 'black hat' activities and communities. Looked at from the standpoint of the boosted risk radar, we can consider such practices an important technique for gaining deeper insight into the possible identities, intentions and capabilities of real adversaries.

Finally, it is interesting to consider that hiring individuals with highly specialised technical knowledge for both attacker and defender roles within red teaming, does seem to bring with it increased opportunity for using secretive red teams as 'fronts' for corporate espionage, or at least for competitive intelligence activity which pushes legal and ethical boundaries. Concerns are likely to be greatest when the individuals hired have previous histories of engaging in such activities.

3. RED TEAMING AS ALTERNATIVE ANALYSIS

Here red teaming is considered as a decision support activity. A (2010) UK Ministry of Defence (MoD) *'Red Teaming Guide'* defines red teaming as "the independent application of a range of structures, creative and critical thinking techniques to assist the end user make a better informed decision to produce a more robust product". The guide notes that NATO prefer to call this 'alternative analysis', defined as "the deliberate application of independent critical thought and alternative perspectives to improve decision-making.

Israel's government and military are reputed to make systematic use of alternative analysis in their major decisions. The red team they use for this is a secret unit within the Israeli Defence Forces called *'Ipcha Mistabra'* (a term of Hebrew and Aramaic origin meaning *'on the contrary'*). Set up following the 1973 Yom Kippur war, the team operates on the principle that it is the duty of the military officers they send to participate within decision-making, to challenge the assumptions and conclusions others make, no matter what their rank or seniority. The social and psychological implications of this *ipcha mistabra* role are of course fascinating. It may be necessary to rotate contributing officers frequently to ensure nobody gets 'stuck' in their adversarial roles and drifts from their mission commitment and loyalty.

Essentially at issue within this view of red teaming, then, is bringing criticism and challenge to bear in order to enhance decision-making. The (2010) MoD guide offers specific guidance on the *'diagnostic'*, *'creative'* and *'challenge'* phases of red teaming practice. It also suggests specific analytical techniques which might be used. Of particular note is one of their 'guidelines for good red teaming' which notes that red teaming needs 'an open learning culture, accepting of challenge and criticism'. Two comments can be made about this particular guideline. Firstly, it is valid for all three forms of red teaming introduced above. Secondly, it also matches what most risk management professionals would look for in a healthy risk culture. Hence it is interesting to consider that wider use of red teaming might in itself provide a technique and spur for developing the risk cultures of organisations so that risk management activities can be more effective. Moreover, in a risk intelligent organisation, it makes sense that all relevant, reliable and credible intelligence information received by the risk radar is passed forward for strategic decision-makers to consider – including information they would rather not have to deal with. Accordingly, we might even advocate for healthy risk cultures manifesting as *ipcha mistabra* stances, whenever risk intelligence is considered within decision-making.

Risk management can learn much from military intelligence practice.

CONCLUSION

Many are drawn to the risk profession because they are fascinated by ideas. They appreciate that it is one of the few professions where skill and innovation in theorising the things that need to be managed is everything. They discern that organisational risk management, while no longer in its infancy, remains far from exhausting all lines of theoretical inquiry and experimentation with new practices through which it might fulfil its potential.

The present discussion document has undertaken to reflect, inspire and to some extent direct this optimism and enthusiasm for a greater future for risk management. It has striven to provoke discussion through a recommendation for improving risk management as an intelligence-gathering practice. This theoretical possibility has been available as low-hanging fruit for risk management theorists for as long as the modern risk management profession has existed, and yet its various strands – including various writings on risk intelligence and related concepts which have proliferated in the last fifteen years – have not been drawn together into a simple proposal for the future of risk management until now. Perhaps if organisational risk management had become an important profession in the post war world of the mid 20th century, deliberately recruiting from demobbed ex-forces personnel with military intelligence experience, it might today seem obvious that this pattern of recruitment was necessary to set the risk profession on firm foundations. This didn't happen, of course, and so we perceive the profession very differently.

Critics of this document's main contention may ask whether scope really remains for the risk management profession to find and incorporate some new and transformative common sense. In the early 2000s the idea that 'opportunity' could be managed in the same way as 'threat' emerged within best practice guidance as what seemed like transformative common sense to many – and yet it didn't quite fulfil all expectations. Critics may correctly remark that this idea met with mixed success and created an odd disconnect between best practice guidance advocating it and real world practice staunchly opposed to it, even in some very large organisations.

Taking stock, it seems highly unrealistic to advocate for untried and untested ideas. Instead it seems much more appropriate to seek opportunities for importing ideas and practices boasting prior history and pedigree in alternative management domains – particularly when these have clear relevance to the work of risk managers. Accordingly, this document has advocated a reshaping of risk management towards a fundamental concern with 'risk intelligence' under the influence of ideas and methods that are already well established. Comments from readers are welcome on their own experiences, or ideas, regarding the synergies we have outlined and tentatively advocated for.