CRUISSE Pilot – Cabinet Office – Identifying and Addressing Uncertainties in the UK's Cyber Risk Landscape

Summary

Pilot 4. Identifying and Addressing Uncertainties in the UK's Cyber Risk Landscape was developed by Drs. Miles Elsden (*former Chief Scientist and Chief Science Advisor at the UK's Department for Transport*) and Jason Blackstock (*senior academic at UCL*) in collaboration with the National Security Secretariat (NSS) of the UK's Cabinet Office. The aim was to assist the NSS with developing: (i) a deeper understanding of the uncertainties characterising the UK's cyber risk landscape; and (ii) practical strategies for delivering against the NSS's responsibility for *assuring* that all nationally-relevant cyber risks —including complex cross-system, cascading and emergent risks — are being effectively identified and managed by appropriate actors. Based on informal interviews with experts and officials (*as well as literature reviews*), the pilot involved designing and delivering two highly-interactive and interdisciplinary expert workshops for UK government officials from a diverse range of departments and agencies with varying responsibility for delivering aspects of the National Cyber Security Strategy (NCSS). The workshops presented and then applied expertise and frameworks from the CRUISSE network to systematically examine case examples of cyber risks described by participating government officials. These CRUISSE expertise and frameworks included: Tuckett and Nikolic's analysis of decision making under uncertainty; adaptive management frameworks of Petersen, Blackstock and the *International Risk Governance Council*; and cyber security frameworks and as well as the adjacent EPSRC-funded PETRAS programme. These highly lauded expert workshops resulted in targeted briefing memos for the NSS, presenting a range of agreed recommendations, including: (i) a novel matrix for categorising and conveying complex national cyber risks; and (ii) a model for defining the critical functions and approaches necessary for the NSS to *assure* the long-term effectiveness of the UK's national cyber risk management strategy and system. The pilot has demonstrated significant practical value to the NSS of CRUISSE's interdisciplinary approach to examining decision-making under uncertainty and generated new theoretic understanding of the dimensions of assuring national cyber risk landscapes are effectively managed. This work is currently being written up for academic publication and development and further application of the generated recommendations and frameworks are being explored with the NSS and other departments.

CRUISSE Cyber Risk Pilot Project
Partner: National Security Secretariat (NSS) of the UK Cabinet Office

The National Security Secretariat (NSS) of the UK Cyber Office is responsible for assuring that nationally significant cyber risks are effectively managed by appropriate actors within the UK Government. Within this content, the Cyber Risk Pilot Project brought the frameworks, expertise and insights from the CRUISSE network and a wider academic and industry perspective to bear to support the NSS to:

1. develop a deeper understanding of the uncertainties characterising the UK's cyber risk landscape; and
2. explore practical strategies and processes for providing assurance of cyber risks through adaptive risk management processes.

Over a period of four months, this pilot project combined scoping discussions with NSS and other government department/agency officials with desk-based research to craft two interactive, Chatham House Rule workshops (see Appendices for workshop reports).

The first workshop explored the cyber risk landscape from the perspective of the uncertainties underpinning experts' and officials' identification and characterisation of cyber risks. Through discussions leveraging the *three-box* uncertainty framework of CRUISSE, it emerged that, despite a general perception of cyber risk being "highly complex," many of the examples of cyber risk raised in the discussions could in fact be understood and further characterised using existing knowledge and analytic models (i.e. fell into Boxes 1 or 2 of the CRUISSE framework). The utility of an appropriate set of questions to differentiate dimensions of cyber risk that truly manifest radical uncertainty, and therefore require different risk management and oversight approaches, was a key insight from this first workshop.

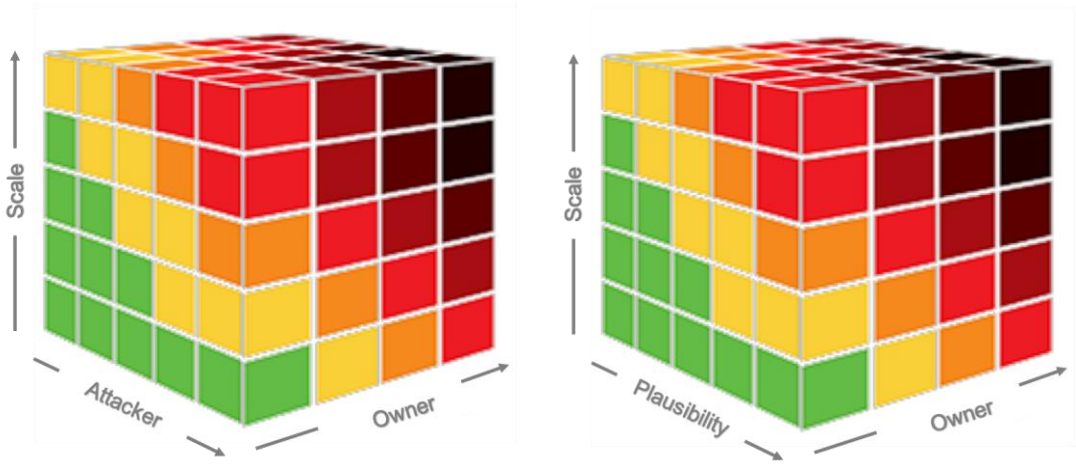

Figure 1: Possible axes for a cyber risk matrix based on three characteristic variables for risk assurance within Cabinet Office. Questions remain about the orthogonality and definition of each of the axes, but initial testing against key currently identified UK cyber risks in the National Risk Register show some promise.

Between the first and second workshops, the project examined different ways of capturing and representing the newly differentiated dimensions of the cyber risk landscape. An initial CRUISSE-based proposal for a national cyber risk matrix was developed and examined in discussion with select NSS colleagues as potential policy tool for capturing and conveying the understanding of both the uncertainties and priorities that emerge from the application of these insights to the cyber risk landscape (see Figure 1). While considerable further development would be required to expand the applicability and utility of this risk matrix as an analytic and policy tool, this approach appears promising, particularly as an outcome from a short pilot project.

Building from the insights of the first workshop (and discussions of the draft cyber risk matrix), the second workshop focused on examining how adaptive management tools could support NSS (and other government colleagues) effectively manage the evolving uncertainties inherent in the cyber risk landscape as an example of risks driven by rapid technological change. A number of contemporary non-cyber examples, along with the Emerging Risk Governance framework of the International Risk Governance Council, were presented by CRUISSE experts and discussed in detail to support government colleagues with developing a deeper understanding of specific adaptive management tools in practice. This led to detailed discussions of how current cyber risk management practices frequently include, often implicitly, adaptive approaches that might be enhanced with conscious attention paid to their potential adaptive (learning) value.

Throughout this second workshop, the value of having a central coordinator (or 'conductor') within the UK Government to integrate and oversee the collective response to cyber risk emerged as a critical theme. In this context, specific tools were discussed that could enhance the ability of the NSS to fulfil at least part of this coordination function via its assurance role. Central to this was ensuring that the NSS had processes and mechanisms for regularly identifying emergent, cascading and systemic risks, and ensuring responsibility for monitoring and managing risks of national importance was delegated to appropriate actors (departments/agencies) within the UK Government landscape.

Through this project, the research team gained deeper understanding of the coordination challenges facing government entities responsible for *assuring* that *all* nationally important cyber risks are being effectively identified and managed. Working with the NSS and other government departments/agencies made clear that, although currently existing (and still evolving) cyber security frameworks are reasonably suited for supporting entities and individuals focused on *management* of cyber risks within well-defined subsystems, additional frameworks are needed to enable higher-level government entities responsible for overseeing the entirety of a nation's cyber risk landscape. In particular, assuring that all risks — including existing systemic and cascading, as well as emergent risks — are being managed by appropriately capable and resourced actors requires an expanded toolkit of information gathering processes and analytic tools, new ways of synthesising and presenting complex, interdisciplinary cyber risk information and new policy response structures within

government. These insights are presently being written up for academic publication and have generated a range of new policy and research questions of direct relevance to the NSS and bodies in other jurisdictions with similar assurance mandates.

At the end of this four-month pilot project, considerable progress has been made in supporting the NSS with the original aims summarised above. Numerous new questions have emerged, however. It was clear from the attendance at the workshops from a wide range of government departments and the resulting discussions that management of rapidly changing, complex, systemic risks (such as cyber) require significant cross-government cooperation. Taking the necessary adaptive approach to managing these risks would require the development and embedding of new capabilities and additional capacity across government. Learning from this pilot project will feed into Cabinet Offices' thinking on development of the next phase of the National Cyber Security Strategy (due from 2021) and may form part of a 2019 Funding Round bid or a National Cyber Security Programme project for FY 2019/2020.

More widely, this pilot has cast new light on how we understand and equip individuals and organisations responsible for *assuring* the and effectively managing large-scale complex risks. While government ultimately acts as 'owners of last resort' for risks that aggregate at the national level, sub-national governments and large corporate entities face similar needs to assure that their risk management strategies and practices are fit for purpose. The outcomes of this pilot have revealed the potential for more generally applicable analytic tools and facilitated processes for supporting those tasked with large-scale risk assurance to effectively audit their organisation(s) risk management strategies and practices to identify opportunities for improvement. More work is clearly necessary to translate the knowledge and expertise embedded in CRUISSE's decision-making under uncertainty and adaptive management frameworks into pragmatic tools and processes; but initial discussion indicate a strong appetite within Whitehall and beyond for partnering on this important next stage of work.

Appendices

Appendix A: Report of Workshop 1 (25 July 2018)

Appendix B: Report of Workshop 2 (10 October 2018)

Appendix A – Cabinet Office Cyber Risk Workshop 1 – Identifying Uncertainty

Intro

The workshop was held at the National Cyber Security Centre (NCSC) office in Nova South, Victoria, London on the 25th July 2018. The meeting was convened as part of CRUISSE pilot project 'Identifying and addressing uncertainties in the UK's cyber risk landscape' together with the National Security Secretariat of the Cabinet Office and the NCSC.

The aims of the meeting were to (i) understand the types and sources of uncertainty in cyber risks for the UK as seen from a national government perspective (ii) consider different approaches to cyber risk management and identify key strengths and weaknesses of current best-practice approaches to assessing and managing risks in the UK National Cyber Security Strategy (NCSS).

The meeting was held under the Chatham House rule.

Summary

- *There was general agreement that an overarching risk management framework for cyber risks would be beneficial that contained a range of approaches, tools and techniques to analyse both component level and system level risks. Defining such a framework, however, was not seen as straightforward and there was no consensus on what such a framework suitable for use in the NCSS could look like.*

- *There was agreement that risk management must be driven by a clear understanding of the intended outcome or outcomes. For example, this would need to define the balance between prosperity and security outcomes.*

- *Cyber is a complex system and there is a need to look at the interconnection and interaction of components to identify key higher-level, systemic risks. Taking such a systems approach to cyber security was seen to be key to developing a more successful approach to cyber risks*

- *The group acknowledged that there can be a tendency to assume all cyber risk is complex, but a majority of national cyber risks are known problems with known solutions. The majority of cybercrime attacks have rational, economic drivers and can generally be predicted, for example.*

- *Similarly, a majority of cyber incidents – including some higher impact / lower probability risks sometimes considered to be 'unexpected', such as WannaCry - are predictable through traditional analytical methods. However there remain some risks that are characterised by radical uncertainty that require new thinking.*

- *There was agreement that in order to define and constrain the problem space, government must focus on those risks where only government can mitigate (or where there is a national or sector-wide market failure).*

- *A clear understanding of what to measure (and what can and cannot be measured) is needed to successfully understand the risk and to develop optimal mitigation options. Poorly defined metrics/incentives were considered to lead to poorly understood risks, badly aligned incentives and undesirable behaviours.*

- *Government will need to develop an understanding of the unintended consequences and impacts of risk mitigation actions taken in a narrow context that, for example, shifts risk across sectors, compounds risk impacts or leads to the creation of new risk factors.*

Discussion Themes

There was a broad ranging conversation during the workshop covering a number of key issues in the understanding and management of cyber risks from a range of perspectives. The main discussion themes are outlined below.

Understanding the Risk

There was agreement that cyber risk management needs to be underpinned by a clear, explicit understanding of intended outcome. The National Cyber Security Strategy has prosperity and security as two of its key drivers. The balance between these drivers will dictate the appropriate risk management approach. The phrase 'follow the pain' was used to describe how different sectors look at cyber risks. In the finance sector, for example, minimising financial loss is generally the overarching driver. The subjective nature of risk was discussed and the desire to protect business and personal reputation were often confounding factors again rational risk response.

A 3-stage model of risk was proposed.

1. **Zone of Routine** – where the risks are economically driven, relatively well understood with the threat changing incrementally and a level of individual risk is being tolerated.
2. **Zone of Surprises** – extrapolating existing threats and using scenarios to extend to things that haven't been seen before – mitigated through capital buffers or insurance (in finance)
3. **Zone of Catastrophe** – looking at feasible risks that have not yet been seen but are not driven by easily understood economic drivers or are systemic in nature. Here organisations look at consequence management identifying impacts and recovery plans. Often regulators and insurers also don't know how manage these types of risks.

All of these types of risks are relevant for the NCSS. Type 1 risks are generally managed at a company or individual level (though NCSC supports a wide range of activities to address these types of risk) and includes the majority of low level cybercrime and risk to government systems. Types 2 and 3 are of more interest at a NCSS level. Type 3 risks are of particular

concern at a national level and correspond to the types of risks generally considered in the Nation Risk Register. The management of type 2 risks, be those to government systems or within the wider economy, present particular challenges to government as many of the standard risk mitigation techniques (such as deferment through insurance) are not available.

Risk evolution was seen to be driven by landmark events such as the widespread use of ransomware after the 'successful' use of WannaCry in 2017. These changes are generally driven by rational drivers (economic or otherwise) that make them possible to foresee. There was agreement that no recent cyber incidents could not have been predicted and that there is often a perception that cyber risks are more complex than they actually are. This means looking at 'simple solutions to simple problems in a simple way' and communicating this approach to a wide range of audiences.

There was a broad discussion on how measurement of risk factors and success metrics needs to be based on an explicit understanding of the underlying problem and the motivation behind the risk strategy. Ill-defined or inappropriate metrics were considered to incentivise behaviours that meet the metrics (compliance) rather than addressing the fundamental problem. A clear understand of what to measure (and what can and cannot be measured) is needed to successfully understand the risk and to developing optimal mitigation options.

The group agreed that the government should focus on what only government can do. Government should encourage the private sector (and individuals) to internalise and manage the cyber risks they can control and focus on addressing the residual risk that is beyond an organisation (or individual) to address. Specific examples of where the government could intervene were mentioned:

- work to ensure there are sufficient cyber skills (at all levels) to meet the national demand
- support a heterogeneous software and hardware environment

Government also has a key role in ensuring there is an appropriate balance between risk minimisation at the individual and the higher, system level. Optimisation of an individual's, company's or sectors cyber risk position can often lead to a sub-optimal solution at a higher level. This requires the government to develop an understanding of the unintended consequences and impacts of a risk mitigation actions taken in a narrow context that, for example, shifts risk across sectors, compound risk impacts or leads to the creation of new risk factors.

Risk Framework

There was general agreement that an overarching risk management framework for cyber risks would be beneficial. It would need to contain a range of approaches, tools and techniques to analyse both component level and system level risks. Defining such a framework, however, was not seen as straightforward. There are a range of different approaches in the literature and at a practitioner level, but there was no consensus on what such a framework suitable for use in the NCSS could look like.

Most current cyber risk approaches are component driven. The aim is to reduce the risk to individual component in the system (servers, switches, software applications, etc) rather than look at the overall sociotechnical system. This is partly due to the background of people involved in the cyber security business who tend to take a technology driven perspective.

The complexity of IT systems, and the multifaceted threat surface they are characterised by, mean there is value in taking a higher abstraction of risk assessment. Looking at the interconnection and interaction of components in such a complex system enables the identification of key higher-level, systemic risks. While component level risk analysis remains important, taking such a systems approach to cyber security was seen to be key to developing a more successful approach to cyber risks. Taking this type of high level system view of risk requires the development of a suitable suite of robust tools and risk analysis techniques (at both a component and systems level) and will entail a significant culture change in the cyber risk community. There are a number of existing high-level frameworks that could be used as a basis for such a risk framework, such as the NIST1 or STAMP2 frameworks. It was also noted that there would be considerable value in looking outside current cyber methods to develop new cyber risk tools as there a likely a range of suitable analysis methods to be found in other disciplines.

Any overarching framework would also need to be supported by well-defined measures, metrics and incentives that support the overall objectives. Poorly defined metrics/incentives were considered to lead to poorly understood risks, badly aligned incentives and undesirable behaviours.

In the long term it was agreed that the aim should be to mainstream cyber risks into the standard risk assessment function. Cyber risks should be considered in the same way as other kinds of business risk, linking key business outcomes to how these risks are managed. This would need to move cyber risk from a technical to a management role and remove the current separation between risk officers and cyber risk professionals. This could be supported by encouraging people from a wider range of backgrounds and perspectives to join the cyber risk profession.

There was agreement that a number of more mature policy areas contain lessons for cyber. Safety risk management was seen as having many parallels with cyber security. Safety risk is significantly more mature in its thinking and moved from a component-based risk assessment to a systems level framework a number of decades ago. There should be clear cut-across of the learning from safety risk assessments to cyber.

Medicine was proposed as another policy area that contains a balance between component level (individual clinical care) and system level (public health) approaches to risk management and would be worthy of comparative study. The functional re-offending risk assessment approach used in offender management was also seen as a model that could provide some learning for the cyber community. Finally, road safety, and in particular, the licensing of both vehicles and drivers was also seen as an area that could be relevant.

---

[1] https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf

[2] http://psas.scripts.mit.edu/home/get_file.php?name=STPA_handbook.pdf

Risk Complexity

The CRUISSE three box model for risk complexity was introduced to the group and seen as a useful way to characterise cyber risks (see below). It was also agreed that although the common perception was that cyber risks generally fall into the third, most complex category, this was not necessarily true. Many risks and risk trajectories are often known and understood before they happen (as per stages 1 and 2 of the cyber risk model outlined above), but that knowledge is not necessarily acted upon.

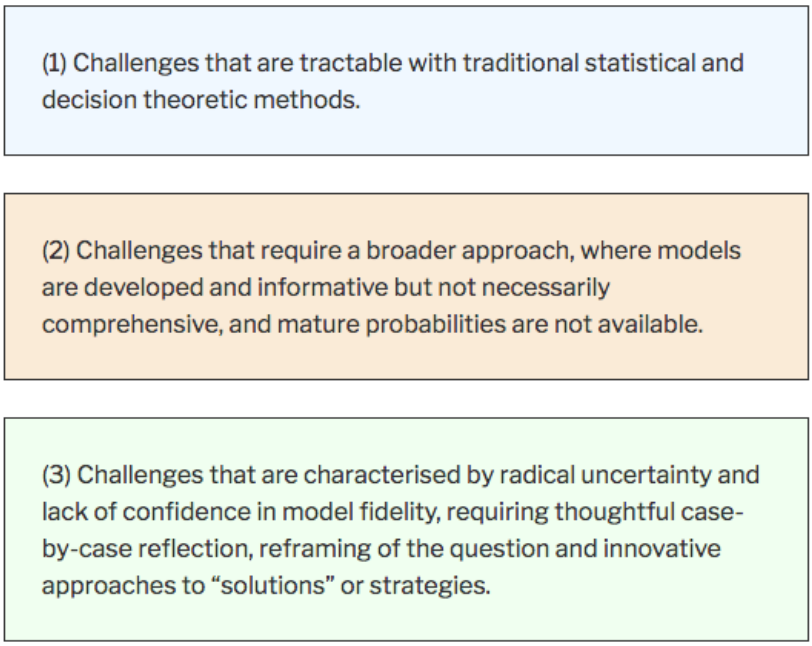

It was noted that the ESRC funded research Centre for Research and Evidence on Security Threats (CREST)3 has already done work in this area and has a range of relevant resources.

Communication

There was much discussion during the workshop around how best to communicate cyber risks to different audiences.

There was general consensus that a Red/Amber/Green (RAG) categorisation of risk was too crude to properly convey the inherent complexity in cyber risk. Other communication approaches were considered more appropriate and that the subjective nature of risk needs to be reflected in any communications approach. As in the discussion on risk frameworks, it was clear that using the right measures and metrics to frame the discussion was very important.

[3] https://crestresearch.ac.uk/

The approach of looking for 'simple solutions to simple problems in a simple way', mentioned above could be supported by taking a narrative approach to risk communication. This would need to be supported by data but could take into account the impact of subjectivity and emotions on risk management decisions.

It was noted that insurance and re-insurance are already addressing many of the same cyber risk questions as government4. However, it was noted that the drivers were potentially very different across the sector – with significant variation between an insurance and re-insurance perspective. The risk mitigation tools, techniques, solutions and incentives in this sector are also significantly different to in government.

Possible Short-Term Actions

The discussions highlighted a number of short-term actions that could start to address some the issues identified in the discussion.

It was agreed that it would be helpful to develop a matrix of national cyber risks that are owned across different departments and their framing. In support of this it was suggested that a common taxonomy be used5 across government departments to classify strategic, national level risks. Both of these activities would: facilitate discussions about common risks; enable better understanding of what is understood as cyber risk within different departments; and assist the understanding of relative risks across government.

It was further proposed that the risks identified above be classified into the 3-box CRUISSE uncertainty model. This would deepen the understanding of those risks are and are not characterised by radical uncertainty, identifying those that may be more amenable to current risk management approaches.

It was suggested that a detailed review be undertaken to understand what is currently measured and what could be measured to better understand cyber risk in the context of the NCSS. This would ensure that measures (and the related metrics) are fit-for-purpose, enable proper characterisation of the NCSS risks, and avoid the risk of driving unwanted behaviours.

Finally, it was suggested that a use-case of a Ministerial (or senior official) briefing on a specific cyber risk issue be analysed to see how information, risk perception, uncertainty and complexity are considered through the reporting process. This could look at how these factors change as information flows from the technical level into the risk assessment process and how this translates into policy recommendations. Understanding this process could help develop new communication techniques, identify points where important caveats may be lost, and ensure risk is well characterised when reported to seniors.

---

[4] For example Lloyds' Risk Reports https://www.lloyds.com/news-and-risk-insight/risk-reports

[5] There are a number of existing taxonomies that may be used as a basis, e.g. http://www.dtic.mil/dtic/tr/fulltext/u2/a537111.pdf from the US

Possible Longer-Term Research Avenues

There were also a number of suggestions for longer term research activities that could be funded from the NCSP.

The expansion of the available set of analysis tools and techniques for cyber risk management was seen as a key area for further study. There are likely a number of suitable analysis methods available in other disciplines that could readily be re-purposed, and this should be investigated, as well as looking to develop new cyber specific analysis processes. Both component driven and systems level cyber risks tools are needed. The development of a suite of analysis tools, techniques, methods and procedures (existing, re-purposed or new) that supports a single risk framework was seen as necessary to encourage the take up of a more sophisticated understanding and management of cyber risks.

There is a clear need to build more detailed background knowledge of the complexity of cyber risks and to explore the wider complexity landscape. This was particularly true of compound risks, and unintended consequences of risk mitigation actions and how there manifest at a systems level.

The development of new communication tools and methods was seen as a key focus area. Better communication approaches and strategies for decision makers, the general public and industry were, together with the development of new analysis tools, seen as essential to develop a superior understanding of the risks and impacts of cyber and to encourage better risk management.

Finally, understanding the behavioural and emotional aspects of cyber and in what way it impacts on how risks are defined, measured and mitigated was also seen as central. This includes understanding new threat drivers, looking at how these factors might drive new threat trajectories, and how to incentivise new, improved behaviours. Game theory may provide a range of methods to understand these issues.

List of Attendees

Non-Government

- Prof Miles Elsden (Chair) STEaPP UCL and CRUISSE
- Prof Lenny Smith, Director Centre for the Analysis of Time Series, LSE and CRUISSE
- Dr David Good, Department of Psychology, Universit of Cambridge and CRUISSE
- Dr Geraint Price, Information Security Group, Royal Holloway, University of London and Chair of RISCS Practitioners Panel
- Prof Emma Barrett, Professor of Psychology, Security & Trust, University of Manchester
- Peter Jaco, Chairman CyberOwl Limited
- David Ferbrache, Chief Technology Officer, Cyber, KPMG UK

Government

- Dr Ian Brown CSA team DCMS
- Will Harvey, Deputy Director Information Security, HMRC
- Emma Green, Head of Regulation & Incentives, DCMS
- Sam B, Sociotechnical Security Group, National Cyber Security Centre
- John Y, Sociotechnical Security Group, National Cyber Security Centre
- Dave M, Assessments team, National Cyber Security Centre
- David Britain, Nation Cyber Security Strategy & Programme, Cabinet Office
- Nikita Shah, National Crime Agency / Cabinet Office

Appendix B- Cabinet Office Cyber Risk Workshop 2 – Adaptive Management

Introduction

This was the second workshop of two planned as part of an EPSRC funded CRUISSE[6] pilot project '*Identifying and addressing uncertainties in the UK's cyber risk landscape'* in collaboration with the National Security Secretariat (NSS) of the Cabinet Office and the National Cyber Security Centre.

The first workshop looked at the types and sources of uncertainty in cyber risks, from the UK national government perspective, and considered different approaches to cyber risk management and how they related to the UK National Cyber Security Strategy (NCSS).

The purpose of this second workshop was to support the NSS with developing practical strategies and processes for providing assurance of cyber risks through adaptive risk management processes. At the heart of 'adaptive management' frameworks are two foundational concepts: (i) identifying and continuously monitoring uncertainties within the (complex) system being managed; and (ii) ensuring each stage of decision making (*organisation, strategic and tactical*) leaves sufficient flexibility for future decisions to change (*adapt*) in response to new evidence (*particularly from the monitored uncertainties*).

The workshop looked to leverage experience from various non-cyber contexts to consider how the heuristics of adaptive risk management could be effectively applied to cyber risks. This will include considering how the NSS could examine the implementation of the NCSS to ensure that sufficient adaptability is built into those processes.

The meeting was held under Chatham House rule.

Summary

- *The majority of cyber risks are addressable with traditional risk management methods but future or system wide risks need different approaches.*
- *Cyber is a complex landscape with a range of emerging, cascading and systemic risk issues. Current best practice recommends taking a systems view of these types of risks and using adaptive techniques in their governance and management.*
- *The UK has a robust process for managing traditional NRA risks, but this process can break down when considering the rapid evolution timescale involved in technology driven risks such as cyber. These risks can be difficult to compare them to traditional risks due to the complexity of cause and affect.*
- *Expanded scenario planning and red teaming could allow better risk identification and prioritisation. This should look at actor motivation and consider social and economic drivers as well as technology.*

[6] http://cruisse.ac.uk/

- *A 'Risk Conductor' (a somewhat expanded one from current Cabinet Office role) could be beneficial to improve the coordination and assure cyber risks in government. This role would coordinate all activities of identifying, managing and communicating around both the risks and the risk management process.*
- *There is an opportunity to apply lessons learnt from previous incidents more widely and existing mechanisms for reporting of near misses could be enhanced.*
- *Significant data is being collected on cyber incidents. This could be a valuable source of 'weak signals' of future risks if its analysis is well coordinated. This could form part of a broader, targeted horizon scanning/technology watch programme.*
- *Some flexibility in research funding would be valuable to enable rapid investigation of new and emerging risks or system behaviours as they are identified.*

Cyber Risk System

The first workshop discussed that many cyber risks, most generally at the component and sub-system level, are often tractable with traditional statistical and decision theoretic methods and standard risk approaches (CRUISSE box 1). At higher systems levels these techniques are often of limited applicability and parts of the problem remains poorly defined but may be amenable to qualitative approaches such as scenario modelling (CRUISSE box 2). At the highest, systems-of-system level, the key risks are more complex still and there is a complete lack of data and limited, or no, understanding of the system in which the decision is being made. We call these situations of decision-making under *deep* or *radical uncertainty* - often characterized as *Wicked Problems* [7] (CRUISSE box 3).

In the first workshop the lack of a clear approach to providing an overarching risk methodology, that can be consistently applied across all levels of cyber risks, was seen as a key issue. Cyber risks generally occur in complex system-of-systems with a large hierarchy of risks ranging from those to individual components, through sub-systems to broader systemic factors including emergent and cascading risks. Each of these risks may have a different owner and a range of mitigation strategies, some of which may compromise responses at other levels in the hierarchy.

The interlinkage of risk, costs and benefits adds more complexity to the management and governance of such systems. Current approaches struggle to fully encapsulate the dynamic interdependencies of cyber risk, particularly when looking at these risks from a national perspective. This implies it may be necessary to consider the overall sociotechnical system rather than the more common approach of addressing risks to individual system components (servers, switches, software applications, etc).

Adaptive Governance

The key principles of adaptive management and governance were discussed in the context of three examples. (i) Air pollution in the US; (ii) Floods in the Netherlands; (iii) Transport

[7] *Churchman, C. West (December 1967). "Wicked Problems". Management Science.* ***14*** *(4)*

safety in the US. The work of the International Risk Governance Council (IRGC)[8] in particular was discussed. The IRGC has done significant work over the last few years looking at the management of systemic[9], emerging[10] and catastrophic[11] risks and their governance[12]. Much of this work is in the context of complex systems, including the effects of rapid technological change, and provides a number of recommendations based on an adaptive approach to risk management, that are highly relevant for the governance of cyber risks.

Adaptive approaches are about learning and trying to anticipate the changing risk landscape. It's about ensuring the right questions are being asked by the right people, at the right time and confirming the right processes are in place and that they are being followed. There is a recognised need to continually search for new evidence and when new data comes to light there needs to be a rapid process to implement changes in how the consequent risks are managed.

The ability to detect and manage early signs of potential issues is a fundamental component of adaptive approaches, be that through identification and analysis of weak signals in data or through widespread, open incident reporting lines (including the reporting of near misses). This is a key component of how the US National Transportation Safety Board manages its risks. There is a clear, simple process for incident reporting, with a no-blame culture, a strong response mechanism to address concerns that are raised and a robust culture of learning lessons.

The general approach to adaptive management, here in the context of emerging risks, is summarised in Figure 1, taken from the IRGC.

The IRGC framework introduces a central role of what the IRGC guidance call a risk 'Conductor' in these complex risk scenarios. The IRGC define the role of the conductor to:

• **Facilitate** interactions between the participants, who may have different objectives and expertise.
• **Validate** technical frameworks and approaches adopted during the process.
• **Monitor** performances to demonstrate their relevance for the organisation and, if required, identify and correct weaknesses. Monitoring the performance of outcomes is challenging, so focusing on the performance and robustness of the process may be more appropriate.
• **Promote** behaviours and attitudes in line with the cultural challenge raised by emerging risk governance. Ensuring transparency, encouraging divergent views, avoiding simplistic visions and eliciting the consideration of extreme and highly unlikely events are examples of what is required to overcome cognitive bias (possible deviation from a standard of rationality) and organisational opposition.

[8] https://irgc.org/
[9] https://irgc.org/risk-governance/systemic-risks/
[10] https://irgc.org/risk-governance/emerging-risk/
[11] https://irgc.org/risk-governance/preparing-for-future-catastrophes/
[12] https://irgc.org/risk-governance/irgc-risk-governance-framework/

• **Communicate** internally and externally to increase awareness of and concern for emerging risks, create a favourable climate for risk governance, explain and clarify decisions, and answer questions.
• **Report** on the potential impact of emerging risks on various sectors and their management.
• **Periodically review** whether the adopted risk management framework and strategy options still conform with the organisation's internal and external context.

This is very similar to the role the Cabinet Office fulfils at present. The IRGS guidelines give more details on the exact role and draws parallels with the role of the UK's Chief Scientific Advisor role in mediating scientific input into departmental policy.

The IRGC guidelines for ERG (as shown in Figure 3 below) comprise five consecutive and interlinked steps. The following sections give a detailed description of each step, its objectives, key participants and expected outcomes.

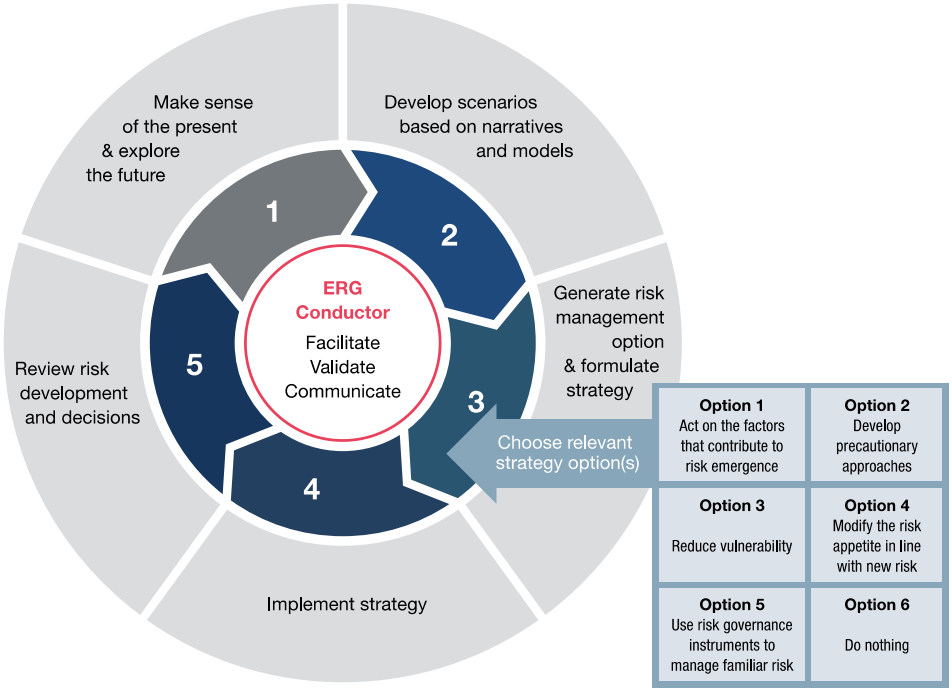

Figure 3: IRGC Emerging Risk Governance Guidelines

Figure 2: From IRGC Guidelines for Emerging Risk Governance

There was a recognition that adaptive approaches are often the most appropriate ones for complex systems, that contain emerging and systemic risks. However, it is not a panacea. It is more a way of thinking about these types of risks, together with a set of tools and approaches that support a systems approach in these complex risk environments than an off the shelf solution.

Discussion

There was a wide-ranging discussion following the introduction of adaptive approaches. Below are some of the key themes that were raised.

The use of scenarios in current planning approaches in the NRA was recognised by the group, however issues were raised around the difficulty in developing credible cyber scenarios and that this sometimes felt strained. In particular there is difficulty in taking the right level of abstraction – is cyber the root cause of the issue or just another attack vector? For example, from a national perspective is the risk of the national electricity grid failing due to a cyber-attack a different risk from a failure caused by a physical attack or from overloading due to plugging in too many electric vehicles?

It was suggested that looking at the actors and motivations were useful lenses to use when considering these risks. Technological feasibility is one risk component, but there are clearly risks where it is technically possible, but without a financial or other driver they are unlikely to materialise. Red teaming was proposed as a useful approach to better understand and prioritise risks, to look at the risk scenarios themselves and also at the risk management processes and governance. This would help identify gaps and weaknesses that may be exploited through, for example, unclear ownership of new or cross-cutting risks.

It was noted that because cyber risks don't generally lead directly to loss of life it can be difficult to compare them with some traditional NRA type risks. There were also questions about the appropriate threshold levels for consideration at a national level. The scenario of hacked pacemakers was used as an example, which has the potential to lead to loss-of-life but not necessarily at scale. Economic, social, political and legal factors can become more important in deciding the relative priority of risks in these contexts.

It was also noted that the current risk review timescales may not be appropriate for rapidly changing technology driven cyber risks, since new risks can emerge much more quickly than can be managed by current annual or biennial review process. The potential for lack of clear ownership of cross-cutting and systemic risk in the current system was raised also as an issue, as was the reactive nature of much risk identification at present. It was clear that there was the desire to move to a more pro-active approach, but this was often constrained by the lack of resources. There was a discussion about lessons learnt and how these could be more widely applied. The learning from the WannaCry attacks lead to recommendations about how to prevent similar future attacks, but there are potentially wider lessons that could be drawn for related risks and in different systems. A more strategic approach to looking at how lessons learned from these types of incidents could be more broadly applied was seen a valuable.

There was a general recognition that while there is a large amount of data available in this space it can be somewhat fragmented – different reporting mechanisms to comply with different legal/regulatory requirements – and is not yet systematically analysed. It was recognised that there would be a clear benefit to enhancing these activities and developing an overarching strategy on how best to use this data. It was recognised that NCSC have been

working hard to encourage better cyber incident reporting, but there are still benefits to be gained from enhancing this capability.

Communication was seen as a key issue (as in the first workshop). Systems approaches are relatively new to policy makers and require a somewhat different way of looking at problems. However, there are already examples in Whitehall, such as in Defra on flood plain risk, where systems approaches are being taken. Cyber risk was seen as an area that would clearly benefit from this approach. The complexity of the systems involved often mean that the risk impact can be widely distributed and there is a disconnect between the mitigation costs and the benefits. This can lead to poor incentives, for example, the cost to an individual of cyber bank fraud is small, as costs are often covered by their bank, this reduces the incentives for positive behaviour change or for the user to enact simple mitigation measures.

There was discussion of the applicability of standard Cost Benefit Analysis to investment in this area. This is a common problem with resilience issues, where the benefit can be difficult to quantify. This applies within organisations, not just at a government level, where cyber security costs fall in one department and the benefits generally fall elsewhere – an example was given of a home bank card reader, a more secure reader would have incurred a small cost increase for significantly better security, but there was no link between the costs in one department and the reduction in fraud costs elsewhere. Alternate, narrative driven decision-making was suggested as an alternative approach, as is sometimes used in innovation investment, however it was recognised that there remains a need for suitable success measures to support spend.

It general it was recognised that there was already a lot of activities addressing many of the issues raised. However, there could be significant benefits gained by stepping back and taking a fresh, strategic view to identify issues and barriers to making the current system perform better and crucially to identify actions that can be taken to reduce these barriers.

Outcomes

There were a number of key points that came out of the discussion, focussing on how the Cabinet Office, in particular, might consider using adaptive approaches in the future management of Cyber risks. These were seen as building on the strong current risk management methodology at a national level and would support a robust assurance role for cyber risks going forward.

It was clear that the role of a ‘conductor’ in this space is a powerful one that builds on much of the current activities of the Cabinet Office, and that could be used to support the assurance role at the centre of government. The value of an individual (or more likely a small team) whose role is to specifically address the challenges of managing complex risks areas that contain emergent, cascading and systemic risk would strongly support a more adaptive approach.

There is a need to develop more systemic activities aimed at looking across the whole cyber risk landscape. This would draw on the work done in government, academia and industry and develop and support activities looking to be proactive in identifying new and unaddressed risks. There are a range of ongoing activities that could be built on within the current NRA approach, such as scenario planning and red teaming, modified to reflect the different characteristics of technology driven risk, such as cyber. This could include, for example:

- commissioning appropriate horizon scanning and technology watch capability from external parties;
- flexible funding for small pieces of research to investigate new threats as information emerges;
- developing a coherent approach to data analysis and reporting, to look for weak signals of emerging threats; and
- developing a range of appropriate measures of success

It was also noted that some current constraints and approaches, such as defining a lead departmental risk owner and the reliance on Cost Benefit Analysis to prioritise funding, present challenges in such complex environment subject to high levels of uncertainty. The pace of change inherent in these risks also presents challenges to existing methods and consideration is needed on how they may need to change to accommodate more rapid decision-making and to introduce the necessary flexibility to respond to changing needs.

Pre-Questions

Attendees were asked to consider a small number of pre-questions prior to the meeting to help focus the discussion. These were:

- What approach(es) to representing cyber risk (*e.g. scenarios; maps of component risk, sub-system risk, systemic risk; analytic risk matrices*) could be most effective for the NSS in fulfilling its assurance function?
- Given the unique characteristics of cyber risks, how can they be best analysed and represented relative to other risks that that NSS manages (*e.g. in the context of the NRA*)?
- What questions does NSS need to ask, and of whom, to ensure that relevant cyber risks, both current and emergent/future, are well managed by the current implementation of the NCSS?
- What aspects/characteristics of 'adaptive management' of cyber risks should the NSS be looking for when evaluating current implementations of the NCSS?
- How could a more adaptive approach to cyber risk management be implemented into the existing and future versions of the NCSS?

Attendees

Non-Government

- Prof Miles Elsden (Chair), STEaPP UCL and CRUISSE
- Prof Arthur Petersen, Professor of Science, Technology and Public Policy, UCL
- Dr David Good, Department of Psychology, University of Cambridge and CRUISSE
- Dr Jason Blackstock, STEaPP UCL and CRUISSE
- David Ferbrache, Chief Technology Officer, Cyber, KPMG UK – not present

Government

- Alexander Brooks, FCO
- Kate R, NCSC
- Maria Egan, ISRC, CCS, Cabinet Office
- Helen Evans, Home Office
- Will Harvey, NSS, Cabinet Office
- Naomi Gilbert, DCMS
- John Friend, CCS, Cabinet Office
- David Britain, NSS, Cabinet Office