RISCS

CYBER RISK QUANTIFICATION RESEARCH PROJECT

A Review of Cyber Risk Quantification in the Context of Specific Challenge Areas

PROJECT TEAM:
Alpesh Bhudia, Anna Cartwright, Edward Cartwright, Frank Cremer, Tom Meurs, Phillip Samson, Jacob Seifert, Darren Shannon, and Barry Sheehan.

VERSION 1.0
28/02/24

OXFORD BROOKES UNIVERSITY
DE MONTFORT UNIVERSITY LEICESTER
UNIVERSITY OF LEICESTER
UNIVERSITY OF LIMERICK OLLSCOIL LUIMNIGH

RISCS

Table of Contents

Executive Summary

Cyber risk quantification provides a means to measure, and subsequently communicate and manage, the risk from cyber attacks or breaches. This requires evaluating the likelihood or probability of negative events and the loss that would be incurred as a result of those events. Cyber risk quantification methods typically focus on quantifying risk within a specific organisation. In this report, we consider how cyber risk quantification methods can be effectively extended or revised to evaluate aggregate risks, such as those in a sector. We also explore how cyber risk quantification can be used to evaluate cybersecurity programmes. In the final section of the report, we focus on the nuclear, financial, and health sectors as examples.

Our key findings can be summarised:

1. There is no agreed-upon method for quantifying cyber risk at an aggregate level for sectors or industries. Cyber risk quantification methods are primarily designed for specific organisations and are not easily adapted to quantify aggregate risk.

2. A 'bottom-up' approach can be used to quantify cyber risk at an aggregate level, in which risk is quantified for representative organisations and then scaled up to give an aggregate measure. This approach is complicated by economic and technological spillovers as well as the correlation of risk across organisations.

3. A 'top-down' approach can be used to quantify cyber risk by analysing aggregate level threats and contagion networks. Methods to estimate systemic risk tend to adopt this approach. The difficulty in this approach is obtaining reliable information to inform the quantification exercise. Novel approaches such as prediction markets could be trialled to provide the needed information.

4. There are various ways, in principle, to use cyber risk quantification to analyse the cost-benefit implications of cybersecurity programmes. The challenge is to obtain accurate data. Data analysis needs to take account of relative threat levels, the overall combination of cybersecurity measures employed, and the human factors in cybersecurity implementation.

5. There is a need for improved cyber risk quantification methods in key sectors, such as nuclear, finance, and health. While these sectors have a long history of risk quantification, existing methods are not well adapted to cyber risk quantification. For instance, they tend to focus on internal threats about which there is reliable data. New cyber risk quantification methods are needed that are better suited to the challenges these sectors face.

List of Abbreviations

AHP	Analytical Hierarchical Process
ANSI	American National Standards Institute
API	American Petroleum Institute
AT	Attack Trees
BDMP	Boolean Logic Driven Markov Processes
BN	Bayesian Network
CD	Chain Diagrams
CDS	Credit Default Swaps
ENISA	European Union Agency for Cyber Security
ETA	Event Tree Analysis
ESRB	European Systemic Risk Board
EY and IIF	Ernst & Young and Institute of International Finance
FAIR	Factor of Analysis Information Risk
FAIR-CAM	Factor of Analysis Information Risk - Controls Analytics Model
FMEA	Failure Modes and Effects Analysis
FMVEA	Failure Modes, Vulnerabilities and Effects Analysis
FTA	Fault Tree Analysis
GT	Game Theory
HAZCADS	Hazards and Consequences Analysis for Digital Systems
HAZOP	Hazard and Operability Analysis
HHM	Hierarchical Holographic Modelling
IMECA	Intrusion Models + Criticality Analysis
MFC	Mean Failure Cost
NIST	National Institute of Standards and Technology
OCTAVE	Operationally Critical Threat, Asset and Vulnerability Evaluation
PN	Petri-Net
ROSI	Return on Security Investment

RM	Risk Matrix
RS	Risk Score
SAG	Security Argument Graph
STPA	Systems Theoretic Process Analysis
UML	Unified Modelling Language
VT	Vulnerability Tree

Introduction

Cyber risk quantification is designed to enable organisations to identify security risks, outline risk scenarios, identify the consequences and associated costs of such scenarios, as well as their frequency or likelihood and possible interventions (Shamala et al. 2013). As part of a RISCS Cyber Risk Quantification Research Project, we were commissioned to produce a review of cyber risk quantification in the context of specific challenge areas. This includes the application of cyber risk quantification to aggregated risks, such as an entire sector or range of organisations, and comparing risk across sectors; the use of quantification to develop a business case for cyber security programmes and interventions, and evaluating their effectiveness, particularly in an environment with an ever-changing threat landscape; and the application of quantification to critical national infrastructure and/or safety-critical settings.

The implications of cyber-attacks are rarely confined to the organisation directly impacted (MacColl *et al.* 2024). For instance, the ransomware attack on Hackney Council is estimated to have had a direct impact of £12 million for the Council; however, the downstream impact of the attack in terms of lost services is likely to be far worse (Pattnaik *et al.* 2023). Similarly, the NotPetya attack on Maersk is estimated to have had a direct impact of around $250 million for Maersk; however, downstream losses are likely several billions of dollars (Welburn and Strong 2022). Given the wider societal and economic costs of cyber-attacks, it is important to question how cyber risk quantification can be conducted at an aggregate level. This would allow a better understanding of the wider risks and consequences of cyber-attacks. For instance, it can be used to quantify the social and economic benefits of cyber security programmes and/or identify high-risk problems that should be policy priorities.

As we shall discuss in this report, the vast majority of research on cyber risk quantification focuses on quantification within an organisation. There is, therefore, a lack of methods to measure aggregate cyber risk quantification. Moreover, we shall see a recognition that current cyber risk quantification methods (e.g., in nuclear, finance, and health) are inadequate to measure complex risks and there is, therefore, a pressing need for improved methods. Against this negative backdrop, we shall also highlight a range of diverse and novel approaches that are emerging and offer positive prospects for improved cyber risk quantification (e.g., modelling of systemic risk). The optimal approach to aggregate cyber risk quantification will likely involve an amalgam of different approaches that combine to produce robust results and relevant margins of error.

Quantification and Aggregated Risk

The vast majority of research on cyber risk quantification has focused on assessing cyber risk within organisations. Relatively little work has considered aggregate risk, and there is a lack of evidence on how cyber risk quantification can be performed at an aggregate level (Tagatev *et al*. 2020, Welburn & Strong 2022). Existing evidence tends to focus on systemic risk, particularly in the financial sector (e.g. Bouveret 2018, Orlando 2021, Welburn & Strong 2022). There is also work on how to consider aggregate risk to insurers (e.g., Welburn & Strong 2022, Zeller & Scherer 2022) and work looking to compare cyber risk across sectors (e.g. Shevchenko 2023).

In principle, there is no reason why cyber risk quantification methods designed for use within organisations could not be used or adapted for use to quantify aggregate risk. In practice, however, there are many challenges in doing so. At a basic level, to be effective, cyber risk quantification requires considerable expertise and information and ideally would be an ongoing, continuous process within management. The larger the scope of the cyber risk quantification exercise, the more difficult it will be to implement the methods accurately and effectively. There are, as we will now discuss, additional challenges to quantifying aggregate risk.

Broadly speaking, we can distinguish two approaches to quantifying cyber risk at an aggregate level: (a) a bottom-up approach of quantifying cyber risk at representative organisations and then aggregating up, or (b) a top-down approach of directly quantifying risk at the aggregate level, recognising limits in information. We consider each in turn. In doing so, we remark that most work on aggregate risk has focused on a Value at Risk (VaR) approach. Informally, VaR provides a probability distribution over potential cyber losses. Specifically, for a given confidence level p, the VaR is the smallest number L such that the probability losses exceed L is not greater than 1 – p (Orlando 2021). Factor Analysis of Information Risk (FAIR) provides one method of implementing VaR. Only a limited number of papers have studied cyber VaR with aggregated risks (e.g., Pal *et al*. 2021)

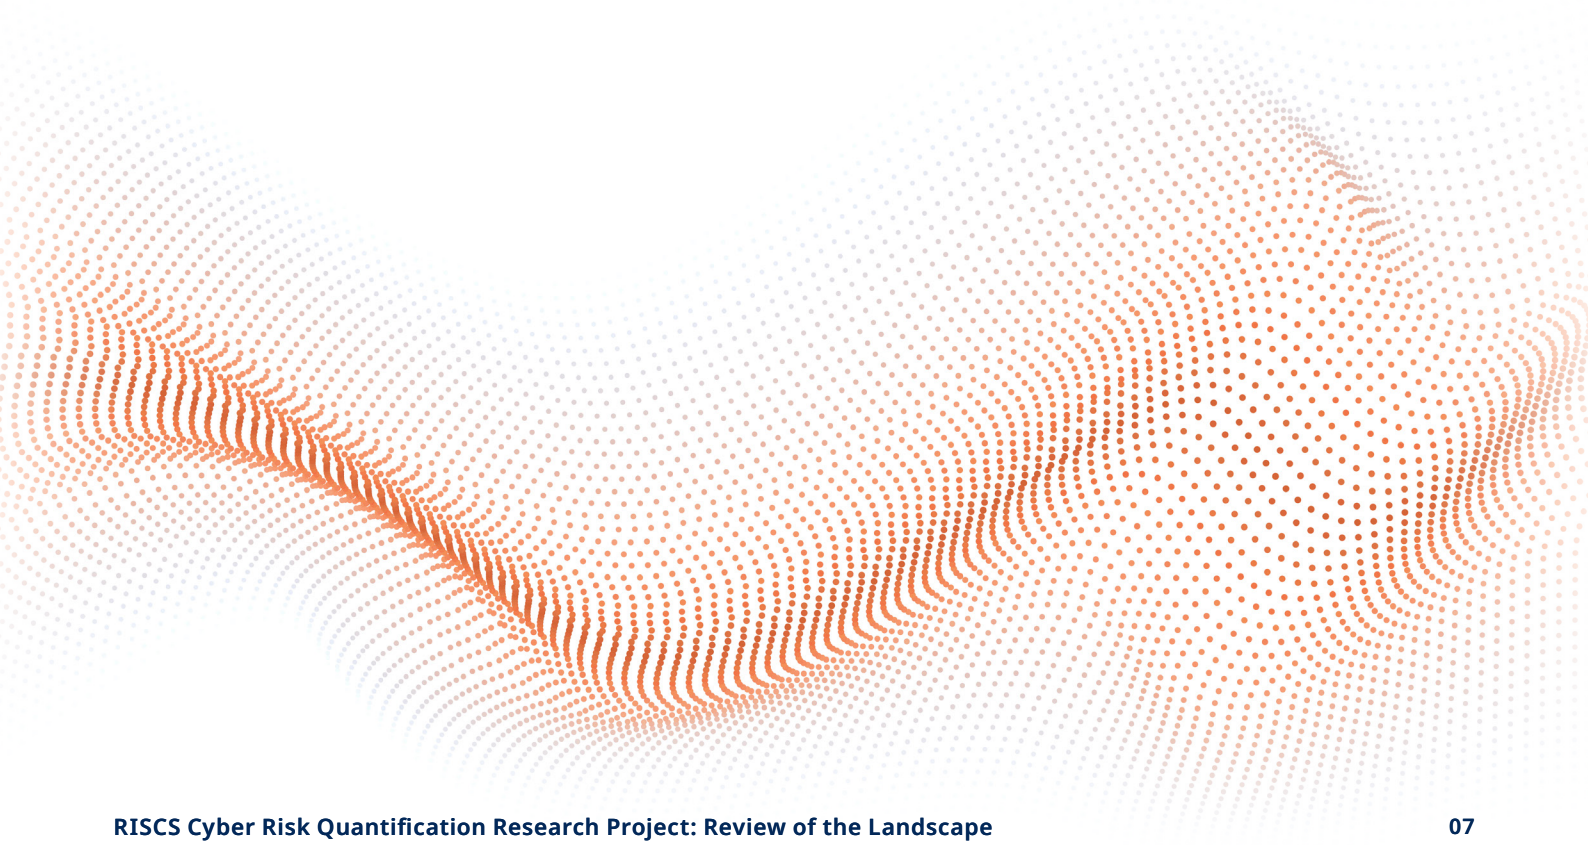

BOTTOM-UP APPROACH TO MEASURING AGGREGATE RISK

One natural approach to assessing aggregate risk, in, say, a sector, is to calculate the cyber risk for a selection of representative organisations within the sector and then use that to formulate the aggregate risk. For instance, in the water sector, a cyber risk quantification for one supplier could approximate for other suppliers and allow the evaluation of an overall sector risk (Zhang *et al*. 2020). In the food sector, one could base the evaluation on a representative distribution of food suppliers (Duncan 2019). Underlying any VaR approach is some variant on the basic equation:

Risk = magnitude loss x probability

As we now discuss, estimating loss and probability involves complications at an aggregated level.

Loss Magnitude and Spillover Effects

If one organisation within a sector is subject to an impactful cyber attack, that will likely have indirect spillover/external effects on other firms within the sector and beyond. These effects can be both negative and positive, as the following examples illustrate.

- **Negative Spillover:** There is a **negative indirect external impact** if an attack on one firm has negative consequences for other firms in the sector. For example, if a producer suffers downtime, then it can negatively impact all organisations in the supply chain. Similarly, if local government is attacked and unable to function as normal, that can negatively impact all businesses that rely on critical functions. In the case of a negative indirect external impact, *the cyber risk quantification of one organisation will **underestimate** the aggregate sector risk*. The left-hand side of Figure 1 provides a stylised example: Firm A suffers a ransomware attack that will cause £1 million worth of downtime loss to the Firm. This would be accounted for in the cyber risk quantification of Firm A. However, Firms B, C and D, in the supply chain of firm A, also suffer downtime losses summing to £1 million. A sector-wide cyber risk quantification exercise should, therefore, account for a £2 million total loss.
- **Positive Spillover:** A threat has a **positive spillover** if an attack on one firm has positive consequences for other firms in the sector. For example, if an organisation suffers downtime and customers go elsewhere, then competitor organisations may benefit from increased trade. The loss in profit to the impacted organisation will be offset by a rise in the profit of competitor organisations. Similarly, an attack on one firm may alert others in the sector to the threat and allow preventative measures that reduce their risk. In the case of a *positive spillover, the cyber risk quantification of one organisation will **overestimate** the aggregate sector risk*. The right-hand side of Figure 1 provides a stylised example: Again, Firm A suffers a ransomware attack that will cause £1 million worth of downtime loss to the Firm. This would be accounted for in the cyber risk quantification of Firm A. However, some customers of Firm A increase purchases from competitor Firms E, F and G. As a consequence, the profits of firms E, F, and G increase by £0.5 million. A sector wide cyber risk quantification exercise should, therefore, account for merely a £0.5 million net loss.

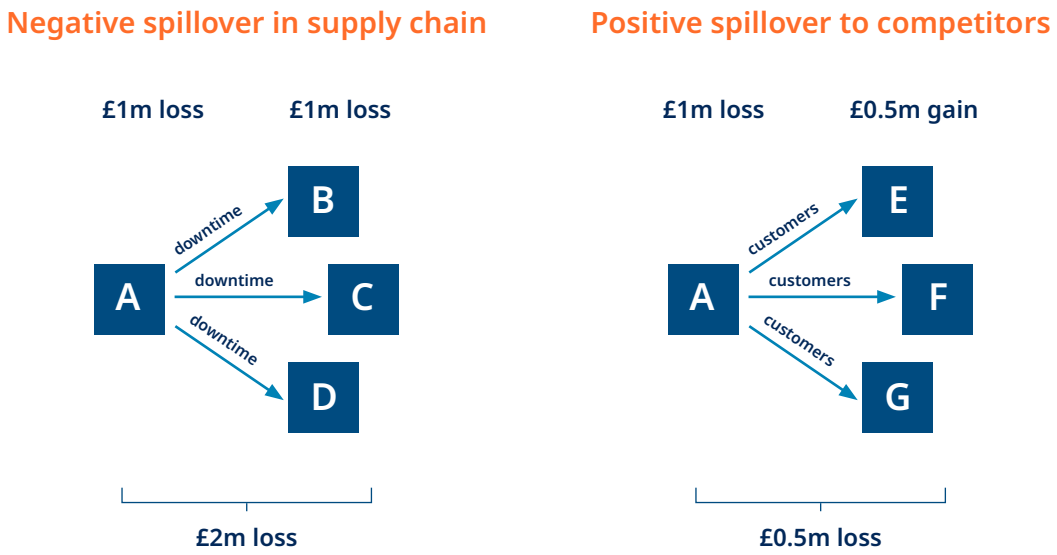

Figure 1: Stylised examples of negative and positive spillover effects.

As may be apparent from our examples, an attack can have both negative and positive spillover effects. For instance, the ransomware attack on Firm A can negatively impact Firms B, C, and D in the supply chain while positively impacting competitors E, F, and G. The sector wide impact of an attack is, therefore, complex to calculate. An additional complexity, as we now consider, is the possibility of an attack that simultaneously impacts multiple Firms.

Event Frequency and Independent or Correlated Risks

In calculating event frequency at an aggregate level, we need to carefully consider a threat that impacts multiple organisations simultaneously.

- **Independent Risk:** Certain threats may pose independent and uncorrelated risks across different organisations. For example, the insider threat from a disgruntled employee can be considered an independent risk because it is unlikely that multiple insiders across multiple organisations will simultaneously launch an attack. Many approaches to VaR and cyber risk quantification assume independent risks. This, however, is a strong assumption in the case of aggregated cyber risk.
- **Correlated Risk:** More common in cyber security are correlated risks in which a threat to one organisation in a sector will likely increase the probability of a threat to other organisations (Amin *et al*. 2013, Pal *et al*. 2020, Welburn & Strong 2022). For example, a zero-day threat can simultaneously impact multiple organisations in a sector. Or, if all organisations in a sector are using the same supplier then an attack on that supplier will impact multiple organisations in the sector. As a final example, if criminals are targeting businesses in a particular sector then other businesses in the sector are also vulnerable.

When looking at an aggregated level one can expect correlated risks need to be taken into account. Given that many models assume independent risks this is a challenge for research on cyber risk quantification (Böhme *et al*. 2019, Orlando 2021, Woods & Böhme 2021).

Spillovers and Correlated Risks Interact

The difficulties of estimating loss magnitude and event frequency are compounded by spillovers and correlated risks often interacting (Welburn & Strong 2022). As an illustration, consider the WannaCry attack on the NHS. If only one hospital or GP surgery in a region was impacted then patients could potentially have been redirected (a positive spillover) to mitigate the risk. However, if multiple hospitals and GP surgeries are impacted within a region then there is less chance to mitigate. Spillovers and correlated risks need, therefore, to be studied in tandem. The following quote from Welburn and Strong (p. 1619-1620, 2022) highlights the importance of this issue:

> ***"The potential direct costs associated with cyber incidents are greatly outweighed by the multiplier effects of up- and downstream connections. While this concern could be true for any firm-level shock, it is particularly relevant given the heavy interconnections of cyberspace."***

The need to take account of spillovers and correlated risks means that the results of a bottom-up approach to cyber risk quantification should be treated with caution. In particular, given the presence of negative spillovers and correlated risk, this approach may underestimate the Value at Risk. We, therefore, suggest the need to consider top-down approaches.

TOP-DOWN APPROACH

Quantifying aggregate cyber risk through a bottom-up approach is complicated by the presence of *economic spillovers* (spillovers in the *effects* of an attack), which may be positive or negative, and *technological spillovers* (spillovers in the *causes* of an attack) that may cause an isolated incident to spread through a sector. The accuracy of aggregate risk estimates based on this approach is also limited to the extent that quantification errors arising at the level of an individual firm are compounded in the process.

An alternative, top-down approach to quantifying cyber risk involves collecting data on cybersecurity breaches at the economic level and then using this data to make inferences about the cyber-risk exposures of individual sectors or even firms.

This approach requires data to be collected on the nature of attacks (scale, type of attack, etc.) and on the characteristics of targeted firms (sector, turnover, employees, regions of operation, etc.). This would allow inferences to be made about:

- The *sectors* that are most frequently subjected to attacks of a certain type. Normalising (dividing) this attack count by the total number of firms in that sector allows an estimate of the probability of attack to be derived and compared across sectors (e.g. Shevchenko 2023).
- The *firm types* that are most frequently attacked. Again, normalising the incident count by the total number of firms within that category allows the probabilities of attack to be compared between firm types (e.g. Jamilov et al. 2021, Malavasi et al. 2022).[1]

Current reporting requirements under NIS Regulations 11(3) and 12(5) are arguably more focused on the technical aspects of a breach than the characteristics of the attack victim (only a description of the types of digital services provided is required). In principle, a regulator can gain an informational advantage over private providers of top-down quantification tools by mandating more in-depth disclosures by firms reporting a cyber-breach.

A further step in this direction would involve an econometric analysis that weighs the relative importance of different firm characteristics in determining their probability of attack, for example using logistic regression.[2] This approach is based on a statistical analysis of the observed variation in attack occurrences, but is complicated by the fact that detailed information needs to be collected, not only about firms that become victims of a cyber-attack but also on firms that are not attacked (Woods & Böhme 2021). Moreover, it is vital to take into account sufficient information to accurately identify the characteristics that determine loss frequency.

As we discuss in the next section (on evaluating cybersecurity programmes), statistical analysis of cyber risk can often produce spurious relations, particularly if the cyber threat level an organisation faces is not accounted for. In the remainder of this section, we first focus on systemic risk, one area where there is more research on cyber quantification, and then cover two more speculative areas where we feel research is needed, namely, prediction markets and market competition.

---

[1] For example, if 150 incidents were recorded in 2023 in relation to manufacturing firms with less than 50 employees based in the West Midlands, and 300 such firms exist in total, the estimated annual attack probability for firms of this type (assuming each firm was attacked only once) is 50%. We currently lack data to make such assessments.

[2] For instance, in the context of the previous example, does the fact that these firms are based in the West Midlands or their nature as manufacturing firms contribute more to their observed probability of attack?

Systemic Risk

Systemic risk is most frequently discussed in the financial context, where it describes the risk that the banking system as a whole will fail as opposed to the risk of failure of an individual bank (European Central Bank 2009). While the failure of an individual bank can be compensated if other banks continue to perform their intermediation role (an example of the economic spillover discussed above), systemic risk matters in the financial context because a general shutdown of the banking sector would entail excessive economic costs.[3] The concepts of financial risk and cyber risk are not mutually exclusive, of course. A cyber-incident may be the cause of a bank failure, but equally, banks may fail for other reasons, and cyber-incidents can also cause firms in other sectors to fail.

The question of what transforms an isolated cyber-incident into a systemic issue, be it in the financial sector or elsewhere, hinges on the notion of amplification (ESRB 2020). Amplification is the process by which a cyber-incident at one firm spreads through the systems and networks that the targeted firm is part of. Amplification is therefore closely related to the economic and technological spillovers discussed above: amplification can occur through the technological channel when the exploited vulnerability is shared by many firms in a sector; it can spread through the economic channel when one firm winding down its operations negatively impacts the business of other firms.[4]

Therefore, systemic risk as a concept should not be seen as distinct from the cyber-risk discussion in the remainder of this report. Rather, it complements the preceding discussion on the economic and technological spillovers of cyber-attacks through the notion of amplification. Moreover, while systemic issues are fundamentally important in the banking sector, the costs of system-wide failures are also high in other sectors. We discuss several categories of critical national infrastructure in this connection below. At this stage, we note that research on systemic cyber risk provides a potentially interesting and alternative (top-down) approach to modelling aggregate cyber risk more generally, particularly in a context of negative economic and technological spillovers. There is wide recognition that systemic cyber resilience scenario stress testing provides one tool to quantify the financial impact of cyber incidents and analyse their potential amplification into a systemic event (Adelmann et al. 2020, ESRB 2022).

As an illustration of the potential for quantifying cyber risk using a top-down systemic approach, consider Welburn & Strong (2022). They propose a specific sector-level input-output analysis to model cascading impacts through a supply chain and quantify aggregate cyber threats. They use the example of the NotPetya attack on Maersk to illustrate how the model can generate estimates of upstream and downstream impact on the supply chain. They estimate upstream losses of between $663 to $773 million and downstream losses of between $16 to $19 billion, compared to a direct loss to Maersk of between $250 and $300 million. These estimates depend on modelling assumptions about the resilience of the network and how 'easily' disruption can flow through the supply chain.

As a second example, consider Kotidis and Schreft (2022). They look at the impact of an attack on a technology service provider that disrupted financial transactions in the US banking sector.

---

[3] A financial crisis can affect the real economy through a contraction of lending as banks seek to maintain a buffer of regulatory capital above their risk-weighted assets, for example (a 'credit crunch'). This affects economic activity and growth because credit-constrained firms tend to reduce technology spending, employment and investment (Campello *et al*. 2010).

[4] This latter channel was important in the 2007-2008 financial crisis, for example, during which liquidity-constrained banks substantially withdrew from the inter-bank lending market, thereby worsening the liquidity problems of other banks.

As a result of the attack, some banks were not able to access the Fedwire settlement systems, resulting in a significant drop in payments being made. Given that some banks were able to access the system and others not, it led to imbalances in cash reserves leading to subsequent spillover effects. Koidis and Schreft (2022) use a statistical model to estimate the aggregate loss in transactions that resulted from the downtime. This could potentially be then used to quantify financial loss.

These two examples show how theoretical and statistical modelling of past attacks can be used to quantify aggregate risk that can inform on the threat from future attacks. These approaches focus on the positive and negative spillovers and correlated risk. As such, they fundamentally differ from the bottom-up approach; in particular, they take as given the direct impact of an attack and then model the aggregate indirect spillover impacts. Such methods can be seen as complementary to more 'standard risk quantification approaches' which focus on measuring the direct impact of an attack on the organisation. It would be beneficial, therefore, to develop these methods further in the future (Amin et al. 2013).

The Role of Markets in Cyber Risk Quantification and Aggregate Risk

The preceding discussion of risk quantification, aggregate risk, and systemic risk has considered the object of the quantification exercise – the risk of a cyber-attack – to be fixed. In reality, quantifying cyber risk is complicated by the fact that a company's vulnerability is not constant but rather interacts in potentially complex ways with its wider business decisions, including the company's decision to invest in cyber risk quantification and mitigation tools. Firms' decisions (incentives) are shaped by the objective to maximise profits while competing with rival firms in a given market setting. As such, understanding the nature of these market incentives is crucial to understanding firms' cyber-risk exposures and the effectiveness of mitigation tools in reducing this risk. Two important concepts from information can be applied here, namely moral hazard and adverse selection. We will discuss moral hazard in the next section (on the quantification of cybersecurity programmes). Here we focus on adverse selection.

Adverse selection arises as a consequence of an informational asymmetry between the buyers of cyber-risk quantification and mitigation tools, who cannot perfectly observe the quality of the product, and the vendors of those products, who are better informed on the quality of their product. Such asymmetrical information can lead to a market dominated by low-quality products that can be sold profitably at a lower price (Anderson and Moore 2006). This, of itself, could result in an aggregate (maybe even systemic) risk that is not accounted for. By way of illustration, suppose multiple organisations within a sector are buying and using the same cyber risk quantification tool. Moreover, suppose that the tool is systematically biased in underestimating the loss from a zero-day ransomware attack. Then there is an unaccounted for correlated risk that, crucially, results from the cyber risk quantification market itself.

Information and transparency about the effectiveness of cyber risk quantification and mitigation tools are therefore crucial. In some cases, adverse selection can be avoided, even in the absence of perfectly informed buyers. This occurs when buyers can infer the quality of the product they are buying from its price. A separation can then occur between sellers of high-quality products, who sell at a higher price, and low-quality products, who sell at a lower price. The specific reasons for which such a separation can occur differ (Wolinsky 1983, Milgrom & Roberts, Bagwell & Riordan 1991, Jones & Hudson 1996, Janssen & Roy 2010), but may be illustrated as follows. If high-quality products are more costly to produce, then the loss of sales that goes along with charging a high price is less damaging for high-quality than low-quality sellers (they save more costs when their sales are reduced). Moreover, if at least some buyers are informed, then falsely signalling a high quality by charging a high price is more damaging to low-quality Firms since this puts off the informed buyers.

Prediction Markets

We have so far highlighted a number of challenges with both the bottom-up and top-down approaches to quantifying cyber risk. Primarily these challenges revolve around estimating the complex impacts of cyber-attacks while only having limited information. Financial markets can provide accurate estimates of risk by aggregating widely dispersed information. Prediction markets extend this market mechanism beyond the financial context in order to predict uncertain events that are not associated with existing financial instruments (Wolfers & Zitzewitz 2004). Participants in a prediction market trade in contracts, the payoffs from which are linked to the outcome of these uncertain events. Decentralised trading between market participants determines the market price of the contract, which in turn allows information about the underlying risk or uncertainty to be inferred. The precise nature of the information that can be inferred from market prices depends on the design of the contracts. Several common examples are summarised in Table 2.

Contract	Example	Details	Reveals market expectation of
Winner-take-all	Event *y*: Firm suffers business disruption because of ransomware	Contract costs £*p*, pays £1 if and only if *y* occurs	Probability that event *y* occurs
Index	Contract pays £1 for every day of business disruption	Contract pays £*y*	Mean value of outcome *y*, *E*(*y*)
Spread	Contract pays if business disruption more than y* days	Contract price fixed at £1. Pays £2 if *y*>*y**, pays £0 otherwise	Median value of *y*

Table 2. Contract Types in Prediction Markets (Wolfers & Zitzewitz 2004)

Prediction markets have been used to forecast various outcomes, including election results, economic data releases, corporate sales, and the box office success of films (Wolfers & Zitzewitz 2004). They have also been considered for information security risks (Pandey & Snekkenes 2014). The appeal of using prediction markets to quantify risk in general is that, if the market is efficient, the market price represents the **best available predictor of that risk**. No combination of available polls, surveys or other information can improve these forecasts: if they could, this would generate arbitrage opportunities for better-informed traders, which would immediately lead to a correction in the market price. (Of course, as we discuss further below, prediction markets are unlikely to be truly efficient in practice.) Prediction markets have been shown to reveal useful information to defenders in aggressor-defender contests (Deck *et al*. 2015), which is particularly relevant in the cybersecurity/hacking context.

Prediction markets may provide a novel way of measuring aggregate risk that may be particularly useful for policy and regulation. As such, we briefly discuss some advantages and disadvantages:

Advantages of using prediction markets to quantify cyber-risk

1. Accuracy of forecasts: Evidence suggests that prediction markets can outperform competing predictors of uncertain events (e.g. Berg *et al*. 2008, Chen and Plott 2002).[5] Another way to assess the accuracy of forecasts in prediction markets is to look at the efficiency of the market in terms of available arbitrage opportunities. Evidence suggests that arbitrage opportunities are virtually absent when comparable contracts are traded on separate exchanges and that the pricing of different contracts on the same exchange tends to be internally consistent (Wolfers & Zitzewitz 2004).

2. Application to Government Departments: The design of the contracts that are traded on prediction markets is very flexible. As such, prediction markets can be applied to the evaluation of cyber-risks facing a Government department or the NHS, for example.

3. Depth of Information Revealed: Standard contracts reveal information about a probability or a mean/median value (see Table 2). Prediction markets can also be used to assess the market uncertainty around these point estimates. For example, a family of winner-take-all contracts can reveal information about the probability distribution of the market's expectations (Wolfers & Zitzewitz 2004).

4. Potential for Risk Hedging: While a prediction market in its initial stages would likely be thin in terms of total trading volumes, a mature market offers the potential for firms and the Government to use contracts to offset (hedge) their cyber risk exposures.

Issues to consider when using prediction markets to quantify cyber-risk

1. The Potential for Market Manipulation: Prediction markets should be resilient against attempted market manipulation. Available evidence from theory, real-life prediction markets and markets set up in economic laboratory experiments suggest that this is typically the case (Wolfers & Leigh 2002, Wolfers & Zitzewitz 2004, Deck & Porter 2013). For example, Deck et al. (2013) show that, while manipulation of prediction markets is possible, it only succeeds in the rather extreme context where traders have deep pockets and only get returns from manipulation.[6]

2. Departures from Market Efficiency: While evidence suggests that prediction markets are capable of delivering accurate forecasts, they are unlikely to be perfectly efficient. For example, there is evidence that people overestimate the probability of highly unlikely events and underestimate the probability of near-certainties (Wolfers & Zitzewitz 2004). These behavioural biases do not rule out the application of prediction markets to the quantification of cyber risk but rather highlight that the risk estimates derived from prediction markets must be interpreted carefully.

---

[5] While many studies find evidence in support of the accuracy of prediction markets, this support is not universal. See Deck and Porter (2013) for a detailed survey.

[6] In this setting, even when manipulators affect the market price, outstanding bids and asks remain informative (Deck *et al*. 2013).

3. Ethical Concerns Around Betting on Harmful Outcomes: Early plans for prediction markets in the US ran into opposition because they were concerned with politically sensitive geopolitical risks (Wolfers and Zitzewitz 2004). There were objections to the idea that perpetrators of harmful actions might use these markets to profit financially from their activity, and a similar concern can be imagined in the cyber risk context. These concerns are likely to be exaggerated for several reasons. Firstly, prediction markets tend to be rather small in scale, so that the overall gains from insider trading (market manipulation) on the part of the perpetrators of cyber attacks are limited. In all likelihood, it would be more profitable for these actors to short the stock market, for example. It may also be argued that this form of market manipulation is actually desirable, as it may reveal otherwise secret information about the intentions of cyber-criminals.

Beyond these general points, there are several specific questions to consider around the design of prediction markets to quantify cyber risk. These include: the specification of contracts, who is allowed to trade in the market, how buyers and sellers are matched, and incentives to trade (e.g. is real or imaginary money used). These issues have been discussed in very general terms in the cyber-risk context. To the best of our knowledge, detailed work around the design and application of prediction markets to the quantification of cyber-risks remains to be completed, however.

SUMMARY

We have argued that there are many challenges in measuring aggregate risk. Moreover, there is limited research on how to accurately measure aggregate risk. Instead, there is a range of different methods, fitting within bottom-up and top-down approaches. These range from a simple aggregation of firm level cyber risk quantification, assuming no spillovers and independent risk, to measuring aggregate loss from past events, measuring spillovers and systemic risk. Given the complexities in measuring aggregate risk, we have argued that novel approaches, such as prediction markets, may be useful. Ideally, a combination of methods, including both bottom-up and top-down approaches, is almost certainly needed for accurate cyber risk quantification.

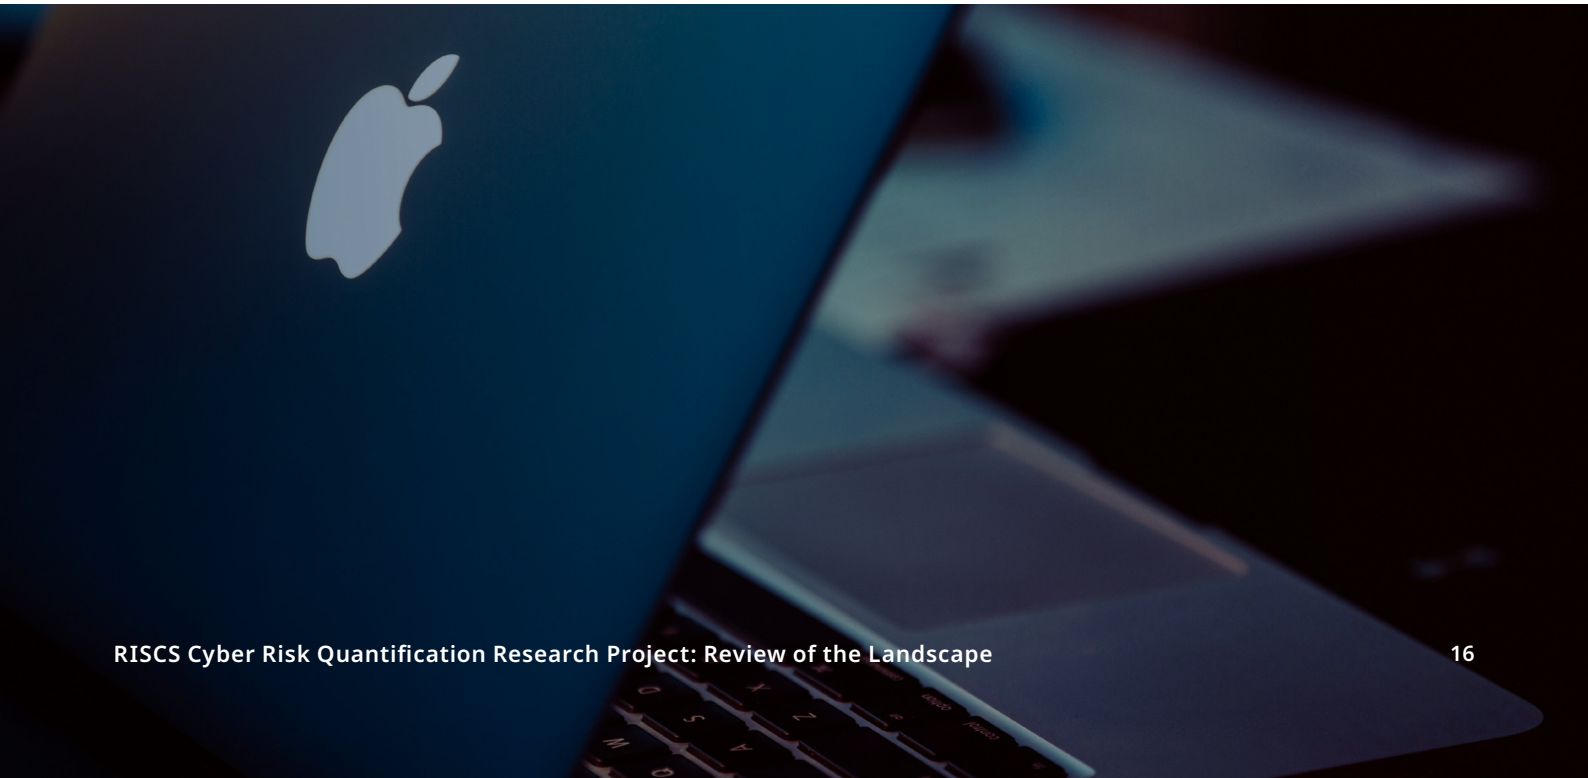

Quantification and Evaluation of Cyber Security Programmes

There are various approaches to using cyber risk quantification to evaluate the potential costs and benefits of a cyber security intervention. For instance, Erola et al. (2022) discuss how VaR can be used to quantify in financial terms the potential gains from introducing cyber security controls. They demonstrate that quantifying and classifying cybersecurity risks using a likelihood-impact analysis matrix provides a way to delineate and communicate in a simple way the financial returns from cybersecurity interventions. Their approach is primarily based on using data from past events to extrapolate future threats. The difficulty with this approach is the availability of relevant data to inform the analysis and the question of whether past data can inform future threats.

There are also quantification products, most notably FAIR-CAM, that look to identify the effect of cyber security controls. This approach makes transparent the diverse channels through which a cyber security intervention can impact financial benefits (FAIR Institute 2021). For instance, the intervention may reduce the loss magnitude and/or reduce the probability of the loss being incurred. It can also take account of indirect effects (e.g. decision support controls) and direct effects. One crucial aspect emphasised in the FAIR-CAM documentation is that 'all controls have a relationship with, and dependencies upon, other controls, which is not accounted for in common control frameworks. As a result, weaknesses in some controls can diminish the efficacy of other controls' (FAIR Institute 2021, p. 3). Analysis of cyber security interventions needs, therefore, to be considered at a holistic level, taking into account the range of controls an organisation has in place and the specific needs and threats of the organisation. This, again, brings us back to the question of data.

Woods & Böhme (2021) demonstrate the difficulty of using data extrapolation to inform on cyber security controls. They show that a superficial look at data typically suggests cyber security interventions are associated with *higher* cyber losses. To explain and control for this perverse finding, it is necessary to take account of the threat level an organisation may face. In short, organisations at higher threat of cyber loss are likely to spend more on security interventions and suffer greater losses; thus, creating a spurious positive correlation between security and loss. Once risk is accounted for it is possible to evaluate the effectiveness of interventions. Woods & Böhme (2021) argue, however, that current studies are failing to do this and so create unreliable results. Research on methods to reliably estimate the impact of cyber security interventions is, therefore, lacking. This is crucial because it means we have limited data with which to evaluate the potential benefit of, say, extra controls.

A lack of reliable data means that cyber risk quantification approaches can provide inaccurate inferences. For instance, they could underestimate the monetary benefit of an intervention because they fail to correct for threat level. Or they could overestimate the monetary benefit because they fail to recognise a lack of complementary controls. The modelling assumptions underlying approaches such as FAIR can also be questioned in relying too heavily on sampling. For example, Wang et al. (2020) argue that FAIR makes inflexible assumptions that limit its accuracy. They propose an alternative Bayesian Network FAIR approach that allows modelling for, e.g. attack-defender causal processes in understanding threats. This can be crucial in evaluating potential cybersecurity interventions.

There are other approaches that allow complementing data with expert input. Sheehan et al. (2021), for example, propose an approach based on the bow-tie model and risk matrix to quantify cyber risks, applied the method to a hospital and identified interventions that could reduce risk. This approach allowed for expert input alongside data analysis. The study found that a lack of staff awareness training and no procedures for vetting prospective employees or monitoring employee activity were risk factors that could be addressed. Such work shows the potential to identify and measure the potential gains from cyber security interventions.

Different cyber risk quantification methods more naturally align with modelling cyber security interventions. For instance, Eggers & Le Blanc (2021) surveyed 36 different cyber risk quantification methods and distinguished them in terms of their end goals (see below for more information). For some methods, the goal is financial analysis to decide how best to allocate cyber security funds. These methods include mean failure cost, return on security investment (ROSI), prioritised attack trees with analytical hierarchy process, financial goal, and game theory (Brangetto & Aubyn 2015). For other methods, the goal is to determine which security controls would be the most effective in reducing the impact of cyber-attacks, e.g. intrusion modes and criticality analysis (IMECA). We discuss this more when looking at specific sectors below. As highlighted by Orlando (2021) there are also options in how to measure risk reduction. For instance, ROSI can focus on expected losses before and after a cyber security intervention or could look at, say, expected loss before the intervention compared to a reasonable worst case scenario after the intervention.

We finish this section by noting two further complications in measuring the returns to cyber security investment, which are not accounted for in current approaches. First, we consider moral hazard. Moral hazard captures the idea that agents may take less care to prevent negative events from occurring if they are insured against the consequences of those events. In the cyber-risk context, firms may exert less effort to prevent cyber-attacks if they believe their implemented method offers sufficient protection.[7] If firm effort to prevent cyber-attacks is reduced when new tools are employed, the effectiveness of those tools in mitigating cyber risks is undermined. Quantification of the effectiveness of cyber security interventions needs, therefore, to take account of human factors, reiterating the need to consider the complementarity of controls.

A final consideration concerns the interactions between firms' cybersecurity practices and their wider business decisions in digital markets. Incentives to invest in security and to adopt security-enhancing technologies interact with firms' data sharing practices and competitive decisions (Lam & Seifert 2023, 2024). As soon as firms share the consumer data they collect with third parties, they should take greater precautions to offset the incremental cyber risk associated with data sharing. Thus, when firms are quantifying their cyber-risk exposure, it is important for them to take the implications of their wider business operations into account. In complex digital markets, this further complicates the task of quantifying cyber risk and evaluating the return to security investment.

[7] This setting differs from the canonical moral hazard or 'hidden action' problem in economics insofar as the lack of precaution in this example arises from an over-estimation of the effectiveness of the tools rather than a shift in liability for negative events.

Applicability of Quantification to Critical National Infrastructure or Safety Critical Settings

In the UK, there are 13 national infrastructure sectors, namely:

- Chemicals
- Civil nuclear
- Communications
- Defence
- Emergency services
- Energy
- Finance
- Food
- Government
- Health
- Space
- Transport
- Water

In this report, we briefly overview the cyber risk quantification literature for nuclear, finance, and health care. We do so because these three sectors are some of the most studied in the cyber risk quantification literature. They also provide three contrasting risk environments: e.g., nuclear has a catastrophic low probability risk, health care has a very wide array of possible risks (e.g. data sharing to use of legacy systems), and the financial sector is likely to be relatively advanced in measuring financial risks.

NUCLEAR INDUSTRY

The nuclear industry is very well versed in the use of probabilistic risk assessments to measure safety risk. Moreover, quantitative risk analysis is considered to be more suited for risks associated with low-probability and high consequence events (Rausand 2013). One might expect, therefore, that the sector would be well placed to lead on cyber risk quantification. As Eggers and le Blanc (2021) point out, however, in their comprehensive review, this is not the case: 'cyber risk analysis of digital assets is still an immature field with unproven techniques due, in part, to the continuously changing threat environment and the challenge of digital assets failing in unexpected ways'. Safety risk assessments in the nuclear industry have primarily focused on areas where very good data exists on, e.g., failure rates. Moreover, the analysis has often focused on 'accidents' and internal threats rather than outsider threats. Existing approaches and mindsets are, therefore, not easily adapted to consider cyber risk quantification.

To expand on these arguments we summarise the review exercise undertaken by Eggers and Le Blanc (2021). They looked at the strengths and weaknesses of cyber risk analysis methods as applicable to the nuclear industry. They rated existing methods on 3 criteria: (a) scope: which measures whether all relevant hazards are considered (scored from 0-4), (b) adoptability: which measures whether the method can be implemented in the nuclear sector (0-3), and (c) repeatability: which measures whether the method allows for ongoing evaluation (with changing threats) and whether the technique gives consistent results for different analysts (0-3). In Table 3, we summarise their findings across 31 cyber risk quantification methods. The table details whether the method is formula based or not, quantitative or qualitative, the end goal, and then the score across the three criteria.

Method	Formula	Type	End goal	Scope	Adoptability	Repeatability	Total
MFC Mean Failure Cost	FB	QT/SQ	Financial analysis	1	1	2	4
ROSI Return on Security Investment	FB	QT	Financial analysis	1	1	2	4
Equation Based Risk Score RS (Kure et al.)	FB	SQ	Identify need for controls	2	2	2	6
RS (Papa et al.)	FB	SQ	Identify need for controls	1	1	1	3
RS (Caralli et al.)	FB	SQ	Identify need for controls	3	1	1	6
RS (Wu et al.)	FB	SQ	Identify need for controls	2	2	1	5
Risk Matrix RM (NIST 2012)	MBnG	SQ	Prioritise controls	3	2	1	6
RM (ANSI/API)	MBnG	SQ	Prioritise controls	3	2	1	6
RM (Braband)	MBnG	SQ	Prioritise controls	2	1	1	4
RM (Hutle et al.)	MBnG	SQ	Prioritise controls	3	2	1	6
RM (Moore)	MBnG	SQ	Prioritise controls	3	1	1	5
RM (Mohr)	MBnG	SQ	Cyber informed safety analysis	2	1	1	4
Game Theory GT	MBnG	SQ	Identify optimal strategy	2	1	2	5
RM + GT	MBnG	SQ	Identify optimal strategy	3	2	1	6
Intrusion Models + Criticality Analysis IMECA	MBnG	SQ	Prioritise controls	1	1	2	4
Security Argument Graph SAG	MBG	SQ	Identify need for controls	1	1	2	4
Petri-Net PN	MBG	SQ	Identify need for controls	1	1	1	3
Attack Trees + Vulnerability Tree AT + VT	MBG	SQ	Prioritise controls	1	1	1	3
AT + CD Chain Diagrams	MBG	QL	Identify need for controls	2	2	2	6
AT + AHP Analytical Hierarchy Process	MBG	SQ	Financial analysis	1	1	2	4
Failure Modes and Effects Analysis FMEA + AT	MBG	QL	Identify need for controls	1	1	1	3

FMEA + HHM ATs Hierarchical Holographic Modelling	MBG	SQ	Identify need for controls	2	1	1	4
Failure Modes, Vulnerabilities and Effects Analysis- + Systems Theoretic Process Analysis FMVEA + STPA	MBG	SQ	Cyber informed safety analysis	2	1	1	4
FTA + ETA + AT Fault Tree Analysis + Event Tree Analysis + Attack Trees	MBG	SQ	Cyber informed safety analysis	2	2	2	6
BDMP Boolean logic Driven Markov Processes	MBG	QT	Cyber informed safety analysis	3	2	1	6
HAZCADS = FTA + STPA	MBG	QL	Cyber informed safety analysis	2	2	2	6
FTA + STPA + Attack Graphs	MBG	SQ	Cyber informed safety analysis	2	2	1	5
UML + HAZOP Unified Markup Language + Hazards and Operability	MBG	QL	Cyber informed safety analysis	2	1	1	4
ETA + BN Bayesian Networks	MBG	SQ	Cyber informed safety analysis	3	1	1	5
BN	MBG	SQ	Identify need for controls	1	1	1	3
3D Risk Profiling	MBG	SQ	Identify need for controls	1	1	1	3

Table 3: Summary of findings from Eggers and Le Blanc (2021). FB = formula based, MBnG = model based non graphical, MBG = model based graphical. QT = quantitative, QL = qualitative, SQ = semi-quantitative.

The main take-away from Table 3 and the analysis of Eggers and Le Blanc (2021) is that, despite the wide range of methods analysed, there is no method that scores highly. Moreover, the following gaps were identified in existing methods: there is no focus on safety-related consequences, a lack of final risk determination, a focus on overall facility security, the method requires analysis of every asset or intensive process, requires modification to current probabilistic risk assessment, lacks detail or poor usage guidance, or relies on limited or subjective data (i.e. system relationships, threat, vulnerability, attack vectors, failure times).

The limitations of existing cyber risk quantification approaches for nuclear are recognised elsewhere (e.g. Son et al. 2023). For instance, in their recent review, Zhang and Kelly (2023) also argue that longstanding methods of risk quantification in the nuclear industry are not well adapted to cyber risk quantification, and new methods are needed. They conclude (p. 499) that 'Since an ideal solution is not currently available, cannibalising existing model elements to fit individual needs may prove to be the best current solution. A hybrid dynamic risk assessment model taking advantage of [Bayesian Networks] is a feasible and promising approach.'

FINANCE

Systemic risk plays an important role in the financial sector. Besides systemic concerns, the financial sector is particularly interesting in terms of cyber risk quantification for two reasons. Firstly, the business activities of banks involve quantifying and trading risks of various types. As such, banks may be expected to employ best practices with respect to the quantification of cyber risk. Secondly, financial markets provide a means of aggregating information that is dispersed among diverse market participants (Siga & Mihm 2021). Through observed market prices, this information can be translated into an estimate of financial risk. This approach motivates the application of market-based methods for the quantification of cyber risk beyond the financial sector (see Prediction Markets above).

Banks are exposed to a variety of risks, including (Bank of England 2020):

- Credit risk: the possibility that outstanding loans will not be repaid.
- Market risk: the possibility that movements in market prices will reduce the value of a bank's trading portfolio.
- Operational risk: the possibility of losses due to failures of internal processes, people and systems.

Cyber risks in the financial sector are mainly discussed in the context of operational risk (Aldasoro *et al*. 2023, Bank of International Settlements 2021). However, cyber-risks may also feed into both credit risk (when loan defaults are caused by cyber-attacks targeting banks' customers) and market risk (when cyber-attacks impair the settlement and clearing facilities that enable market transactions, for example).

It is helpful to consider how banks quantify cyber-risks arising in connection with cyber-attacks on the bank itself, which fall under operational risk, and risks arising in connection with counterparties or third-parties, which mainly fall under credit risk. In terms of the former, survey-based research highlights a perception on the part of bank managers that the consequences of cyber-attacks cannot be quantified financially (Pollmeier *et al*. 2023). This is somewhat surprising given banks' expertise in quantifying financial risks more generally, and the results of other surveys (albeit conducted by professional services firms) that suggest cyber-risks rank as the top concern of chief risk officers (EY and IIF 2024). Banks' risk mitigation approaches were found to be more effective in countering internal than external risks, leaving them vulnerable in the cybersecurity context (Pollmeier et al. 2023). Concerning counterparty risk, credit rating agencies such as Moody's now offer models that integrate cyber risk into their analysis of creditworthiness.[8] Although these are proprietary models, details of which are not freely accessible, publicly available descriptions of these tools tie in with the data-driven, top-down approach to cyber-risk quantification discussed above.[9]

The second particularity of the financial sector is that markets exist for the trading of financial instruments, the value of which is often linked to the risk of particular events. Credit default swaps (CDS) are a good example of this. In return for regular payments, a CDS promises to pay out in the event that an underlying credit instrument, such as a bond, defaults. As such, it serves as an insurance mechanism against adverse credit events. This insurance becomes more costly as the likelihood of default increases, allowing the probability of default to be inferred from market prices.

---

[8] See, for example, https://www.moodys.com/web/en/us/about/what-we-do/cyber.html

[9] See, in particular, https://www.bitsight.com/security-ratings/data-advantage

Markets typically offer rewards to better-informed traders, who can exploit the mispricing of financial instruments through arbitrage strategies that end up driving the market price of the instrument towards its fundamental (arbitrage-free) value. In the context of a CDS, this should translate into an accurate estimate of credit risk. The ability of markets to provide accurate measures of credit risk is supported by research showing that CDS prices can predict sovereign credit events, while credit ratings do not (Rodríguez et al. 2019).

There are, of course, limitations to the market pricing of financial risks. These limitations emerged clearly in the 2007-2008 financial crisis. An important element of this crisis arose in connection with sub-prime mortgage-backed assets, the credit risk of which was underestimated. While financial innovation played an important role in developing complex financial derivatives that markets were unable to price accurately, the ability of markets to price risk correctly in connection with simpler financial instruments is more widely supported (Rodríguez et al. 2019).

In summary, although risk is fundamental to the operations of the financial sector, banks appear not to be far ahead of other firms in quantifying cyber-risks. That said, financial markets have the potential to aggregate information that is widely dispersed and thereby to arrive at accurate assessments of risk. This motivates the potential use of prediction markets to quantify cyber risk discussed above.

HEALTH

With nuclear and finance we have some common themes emerge. In particular: (a) traditional and long established ways of modelling risk within these sectors are not well adapted to modelling cyber risk, and (b) existing cyber risk quantification methods are not ideal for the specific challenges in these complex sectors. While there is less research on cyber risk quantification in health (than nuclear or finance), existing studies suggest that findings (a) and (b) also apply to healthcare.

For instance, Ksibi et al. (2023) provide a comprehensive overview of cybersecurity risk within e-health systems. They highlight the complex environment: 'the e-health environment is mainly characterised by its ubiquity, heterogeneity of devices, diversity of behaviour and capability, scarcity of computing resources, changing infrastructures, etc. From the security perspective, it has a wide and complex attack vector/surface since it is under various and continuously changing threat models.' Ksibi (2023) evaluate existing cyber risk quantification methods (e.g. OCTAVE and FAIR) and identify drawbacks with using such approaches in the complex healthcare environment.

Kandasamy et al. (2020) provide a summary conclusion in their review of Internet of Things (IoT) services. They argue that existing cyber risk quantification methods (such as OCTAVE) are not well suited to study the complex risks of IoT. Moreover, they highlight health care as one sector where the need for new methods is pressing. Kandasamy et al. (2020) and Ksibi et al. (2023) outlined alternative approaches to cyber risk quantification in health care. See also they study by Sheehan et al. (2021) that we discussed earlier using a quantitative bow-tie method to assess risk in a hospital. These approaches, however, should be seen as promising possibilities to be explored in future work, rather than frameworks ready to be implemented at scale in healthcare settings.

Key Findings

Our key findings can be summarised:

1. There is no agreed upon method to quantify cyber risk at an aggregate level for sectors or industries. Cyber risk quantification methods are primarily designed for specific organisations and not easily adapted to quantify aggregate risk. Extending cyber risk quantification methods to measure aggregate risk is difficult.

2. A 'bottom-up' approach can be used to quantify cyber risk at an aggregate level in which risk is quantified for representative organisations and then scaled up to give an aggregate measure. This approach is complicated by economic and technological spillovers as well as correlation of risk across organisations.

3. A 'top-down' approach can be used to quantify cyber risk by analysing aggregate level threats and contagion networks. Methods to estimate systemic risk tend to adopt this approach. The difficulty in this approach is obtaining reliable information to inform the quantification exercise. Novel approaches such as prediction markets could be trialled to provide the needed information.

4. There are various ways, in principle, to use cyber risk quantification to analyse the cost-benefit implications of cyber security programmes. The challenge is to obtain accurate data. Data analysis needs to take account of relative threat levels, the overall combination of cyber security measures employed, and the human factors in cyber security implementation.

5. There is a need for improved cyber risk quantification methods in key sectors, such as nuclear, finance, and health. While these sectors have a long history of risk quantification, existing methods are not well adapted to cyber risk quantification. For instance, they tend to focus on internal threats about which there is reliable data. New cyber risk quantification methods are needed that are better suited to the challenges these sectors face.

6. There is a pressing need for more research on cyber risk quantification at an aggregate level and in complex environments. The lack of research likely reflects, in part, the limited commercial potential for aggregate cyber risk quantification. Future work could explore complementarities between different approaches, including top-down and bottom-up, and explore how they can be combined to produce robust risk estimates.

Conclusion

Cyber risk quantification is designed to enable organisations to identify security risks, outline risk scenarios, identify the consequences and associated costs of such scenarios, as well as their frequency or likelihood and possible interventions (Shamala et al. 2013). As part of a RISCS Cyber Risk Quantification Research Project, we were commissioned to produce a review of the cyber risk quantification in the context of specific challenge areas. This includes: the application of cyber risk quantification to aggregated risks, such as an entire sector or range of organisations, and comparing risk across sectors. The use of quantification to develop a business case for cyber security programmes and interventions, and evaluating their effectiveness, particularly in an environment of an ever changing threat landscape. The application of quantification to critical national infrastructure and/or safety-critical settings.

We have argued that the vast majority of research on cyber risk quantification focuses on quantification within an organisation. There is, therefore, a lack of methods to measure aggregate cyber risk quantification. It is highly challenging to scale up cyber risk quantification at an organisational level to obtain aggregate measures because of, e.g., economic and technological spillovers and correlated risks. Alternative approaches may, therefore, be needed. For instance, studies of systemic risk take as a starting point for network spillovers and so provide a complementary approach to bottom-up approaches. We also suggested prediction markets as a viable approach to obtain expert insight in a controlled, quantitative way to inform quantification estimates.

We also argued that current cyber risk quantification methods are inadequate to measure complex risks in critical sectors, including nuclear, finance and health, because of the more complex risk environment. The fact that such sectors have a long history of successful risk quantification is not translating into cyber risk quantification because of the different nature of cyber risk, particularly in terms of the fast changing threat landscape. More research is needed, therefore, to improve cyber risk quantification methods. The optimal approach to aggregate cyber risk quantification will likely involve an amalgam of different approaches that combine to produce robust results and relevant margins of error. In particular, the complexity of the cyber environment means that past data alone is unlikely to yield accurate cyber risk quantification. Some level of expert input is needed to inform the process.

References

Adelmann, F., Ergen, I., Gaidosch, T., Jenkinson, N., Khiaonarong, M. T., Morozova, A., Schwarz, N. & Wilson, C. (2020). Cyber Risk and Financial Stability: It's A Small World After All. International Monetary Fund

Aldasoro, I., Gambacorta, L., Giudici, P., and Leach, T. (2023). Operational and Cyber Risks in the Financial Sector. International Journal of Central Banking, 19(5), 341-402

Amin, S., Schwartz, G. A., & Hussain, A. (2013). In Quest of Benchmarking Security Risks to Cyber-Physical Systems. IEEE Network, 27(1), 19-24

Anderson, R., & Moore, T. (2006). The Economics of Information Security. Science, 314(5799), 610-613

Bagwell, K. and Riordan, M.H. (1991). High and Declining Prices Signal Product Quality. American Economic Review, 81(1), 224-239

Bank of England (2020). What Risks Do Banks Rake?. Available at: https://www.bankofengland.co.uk/explainers/what-risks-do-banks-take

Bank of International Settlements (2021). Newsletter on Cyber Security, 20 September 2021. Available at: https://www.bis.org/publ/bcbs_nl25.htm

Berg. J., Forsythe, R., Nelson, F. & Rietz, T. (2008). Chapter 80 Results from a Dozen Years of Election Futures Markets Research. In Handbook of Experimental Economic Results, Vol. 1, Plott, C. & Smith, V. (eds.). Amsterdam: Elsevier, 742-751. Available at: https://doi.org/10.1016/S1574-0722(07)00080-7

Böhme, R., Laube, S., & Riek, M. (2019). A Fundamental Approach to Cyber Risk Analysis. Variance, 12(2), 161-185

Bouveret, A. (2018). Cyber Risk for the Financial Sector: A Framework for Quantitative Assessment. International Monetary Fund

Brangetto, P., & Aubyn, M. K. S. (2015). Economic Aspects of National Cyber Security Strategies. Brangetto P., Aubyn MK-S. Economic Aspects of National Cyber Security Strategies: project report. Annex, 1(9-16), 86

Campello, M., Graham, J.R. & Harvey, C.R. (2010). The Real Effects of Financial Constraints: Evidence from a Financial Crisis. Journal of Financial Economics, 97(3), 470-487

Chen, K.-Y. & Plott, C. (2002). Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. California Institute of Technology Social Science Working Paper, No. 1131

Deck, C., Hao, L. & Porter, D. (2015). Do prediction markets aid defenders in a weak-link contest? Journal of Economic Behavior & Organization, 117, 248-258

Deck, C., Lin, S. & Porter, D. (2013). Affecting policy by manipulating prediction markets: Experimental evidence. Journal of Economic Behavior & Organization, 85, 48-62

Deck, C. & Porter, D. (2013). Prediction Markets in the Laboratory. Journal of Economic Surveys, 27, 589-603

Duncan, S. E., Reinhard, R., Williams, R. C., Ramsey, F., Thomason, W., Lee, K., ... & Murch, R. (2019). Cybersecurity: A New Perspective on Protecting US Food and Agricultural System. Frontiers in bioengineering and biotechnology, 7, 63

Eggers, S., & Le Blanc, K. (2021). Survey of Cyber Risk Analysis Techniques for Use in the Nuclear Industry. Progress in Nuclear Energy, 140, 103908

ESRB (2020). Systemic Cyber Risk, February 2020. Available at: https://www.esrb.europa.eu/pub/pdf/reports/esrb.report200219_systemiccyberrisk~101a09685e.en.pdf

Erola, A., Agrafiotis, I., Nurse, J. R., Axon, L., Goldsmith, M., & Creese, S. (2022). A System to Calculate Cyber Value-at-Risk. Computers & Security, 113, 102545

European Central Bank (2009). Financial Stability Review. December 2009. Available at: https://www.ecb.europa.eu/pub/pdf/fsr/financialstabilityreview200912en.pdf

Evans, A. (2019). Managing Cyber Risk. Routledge

EY and IIF (2024). 13th Annual EY/IIF Global Bank Risk Management Survey. Accessible at: https://www.iif.com/portals/0/Files/content/Regulatory/32370132_2312-4407639_eyiif-global-bank-risk-mgmt-survey_final2.pdf

FAIR Institute (2021). An Introduction to the FAIR Controls Analytics Model (FAIR-CAM)

Jamilov, R., Rey, H., & Tahoun, A. (2021). The Anatomy of Cyber Risk (No. w28906). National Bureau of Economic Research

Janssen, M.C.W. and Roy, S. (2010). Signalling Quality Through Prices in an Oligopoly. Games and Economic Behavior, 68(1), 192-207

Jones, P. and Hudson, J. (1996). Signalling Product Quality: When is Price Relevant? Journal of Economic Behavior & Organization, 30(2), 257-266.

Kandasamy, K., Srinivas, S., Achuthan, K., & Rangan, V. P. (2020). IoT Cyber Risk: A Holistic Analysis of Cyber Risk Assessment Frameworks, Risk Vectors, and Risk Ranking Process. EURASIP Journal on Information Security, 2020(1), 1-18

Kotidis, A., & Schreft, S. (2022). Cyberattacks and Financial Stability: Evidence from a Natural Experiment

Ksibi, S., Jaidi, F., & Bouhoula, A. (2023). A Comprehensive Study of Security and Cyber-security Risk Management within e-Health Systems: Synthesis, Analysis and a Novel Quantified Approach. Mobile Networks and Applications, 28(1), 107-127

Lam, W.M.W. and Seifert, J. (2023). Regulating Data Privacy and Cybersecurity. Journal of Industrial Economics, 71(1), 143-175

Lam, W.M.W. and Seifert, J. (2024). Competition, Data Sharing and Secure Hardware Adoption. Working Paper

MacColl, J., Hüsch, P., Mott, G., Sullivan, J., Nurse, J. R., Turner, S., & Pattnaik, N. (2024). Ransomware: Victim Insights on Harms to Individuals, Organisations and Society

Malavasi, M., Peters, G. W., Shevchenko, P. V., Trück, S., Jang, J., & Sofronov, G. (2022). Cyber Risk Frequency, Severity and Insurance Viability. Insurance: Mathematics and Economics, 106, 90-114

Milgrom, P. and Roberts, J. (1986). Price and Advertising Signals of Product Quality. Journal of Political Economy, 94(4), 796-821

Orlando, A. (2021). Cyber Risk Quantification: Investigating the Role of Cyber Value at Risk. Risks, 9(10), 184

Pal, R., Huang, Z., Yin, X., Lototsky, S., De, S., Tarkoma, S., ... & Sastry, N. (2020). Aggregate Cyber-Risk Management in the IoT Age: Cautionary Statistics for (re) Insurers and Likes. IEEE Internet of Things Journal, 8(9), 7360-7371

Pandey, P., & Snekkenes, E. A. (2014, September). Using Prediction Markets to Hedge Information Security Risks. In International Workshop on Security and Trust Management (pp. 129-145). Cham: Springer International Publishing

Pattnaik, N., Nurse, J. R., Turner, S., Mott, G., MacColl, J., Huesch, P., & Sullivan, J. (2023, July). It's More Than Just Money: the Real-World Harms from Ransomware Attacks. In International Symposium on Human Aspects of Information Security and Assurance (pp. 261-274). Cham: Springer Nature Switzerland

Pollmeier, S., Bongiovanni, I. and Slapničar, S. (2023). Designing a Financial Quantification Model for Cyber Risk: A Sase Study in a Bank. Safety Science, 159, 106022

Rausand, M. (2013). Risk Assessment: Theory, Methods, and Applications (Vol. 115). John Wiley & Sons

Rodríguez, I.M., Dandapani, K. and Lawrence, E.R. (2019). Measuring Sovereign Risk: Are CDS Spreads Better than Sovereign Credit Ratings? Financial Management, 48(1), 229-256

Siga, L. and Mihm, M. (2021). Information Aggregation in Competitive Markets. Theoretical Economics, 16(1), 161-196

Sheehan, B., Murphy, F., Kia, A. N., & Kiely, R. (2021). A Quantitative Bow-tie Cyber Risk Classification and Assessment Framework. Journal of Risk Research, 24(12), 1619-1638

Shevchenko, P. V., Jang, J., Malavasi, M., Peters, G. W., Sofronov, G., & Trück, S. (2023). The Nature of Losses from Cyber-related events: Risk Categories and Business Sectors. Journal of Cybersecurity, 9(1), tyac016

Son, K. S., Song, J. G., & Lee, J. W. (2023). Development of the Framework for Quantitative Cyber Risk Assessment in Nuclear Facilities. Nuclear Engineering and Technology, 55(6), 2034-2046

Tagarev, T., Pappalardo, S. M., & Stoianov, N. (2020). A Logical Model for Multi-sector Cyber Risk Management. Information & Security, 47(1), 13-26

Wang, J., Neil, M., & Fenton, N. (2020). A Bayesian Network Approach for Cybersecurity Risk Assessment implementing and extending the FAIR model. Computers & Security, 89, 101659

Welburn, J. W., & Strong, A. M. (2022). Systemic Cyber Risk and Aggregate Impacts. Risk Analysis, 42(8), 1606-1622

Wolfers, J. and Leigh, A. (2002). Three Tools for Forecasting Federal Elections: Lessons from 2001. Australian Journal of Political Science, 37(2), 223-240.

Wolfers, J., & Zitzewitz, E. (2004). Prediction Markets. Journal of Economic Perspectives, 18(2), 107-126

Wolinsky, A. (1983). Prices as Signals of Product Quality. Review of Economic Studies, 50(4), 647-658

Woods, D. W., & Böhme, R. (2021, May). SoK: Quantifying Cyber Risk. In 2021 IEEE Symposium on Security and Privacy (SP) (pp. 211-228). IEEE

Zeller, G., & Scherer, M. (2022). A Comprehensive Model for Cyber Risk Based on Marked Point Processes and Its Application to Insurance. European Actuarial Journal, 12(1), 33-85

Zhang, F., & Kelly, K. (2023). Overview and Recommendations for Cyber Risk Assessment in Nuclear Power Plants. Nuclear Technology, 209(3), 488-502

Zhang, Y., Wang, L., Liu, Z., & Wei, W. (2020). A Cyber-Insurance Scheme for Water Distribution Systems Considering Malicious Cyberattacks. IEEE Transactions on Information Forensics and Security, 16, 1855-1867