RISCS

CYBER RISK QUANTIFICATION RESEARCH PROJECT

Overall Summary of Key Findings

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

VERSION 1.0
07/06/24

OXFORD BROOKES UNIVERSITY

DE MONTFORT UNIVERSITY LEICESTER

UNIVERSITY OF LEICESTER

UNIVERSITY OF LIMERICK OLLSCOIL LUIMNIGH

Table of Contents

Introduction

Cyber risk quantification provides a means to measure, and subsequently communicate and manage, the risk to an organisation from cyber attack or breach. Risk quantification requires evaluating the likelihood or probability of negative events and the loss that would be incurred as a result of those events. An overall measure of cyber risk to the organisation can then be estimated. Risk quantification can also be used to evaluate the expected impact of interventions such as enhanced cybersecurity controls (Orlando 2021). In doing so it provides a means to identify optimal investment opportunities. Given that risk can be communicated in terms that are familiar to boards it may help facilitate improved cyber security risk management.

Cyber risk quantification is, however, difficult to do well. The European Union Agency for Cyber Security (ENISA) recently criticised the lack of standard procedures for identifying, mitigating and quantifying cyber risk (ENISA 2023). Effective cyber risk quantification requires good, evidenced based models of cyber risk. This involves evaluating threats and events that are complex and about which, given the rapid evolution of cyber tactics, we may have relatively little quality data. Even a good model will inevitably result in large uncertainty around overall risk. Risk quantification also requires resources and expertise to inform, conduct and interpret the risk analysis as well as collect relevant data. Ideally, it should be a continuous, ongoing process with input from across the organisation. Demand for resource, expertise, and data can act as a barrier to the effective use of cyber risk quantification. It can also result in organisations adopting ‘flashy but ineffective’ risk quantification tools that can negatively inform strategic decision making. It is necessary, therefore, to question how organisations can be guided and supported to implement effective cyber risk quantification.

Cyber risk quantification can be seen as a subset of the broader notion of information security risk assessment (ISRA). It enables 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. In general there are three distinct phase: context establishment, risk identification and risk analysis (Shamala et al. 2013). Risk identification involves identifying assets, threats, existing controls, vulnerabilities and consequences (ISO/IEC 27005:2018). Risk analysis involves assessment of consequences, assessment of incident likelihood, and determination of the resultant risk level (ISO/IEC 27005:2018). As we discussed in more detail in the full reports, risk analysis can use a variety of different methods and approaches.

As part of a Research Institute for Sociotechnical Cyber Security (RISCS) project on Cyber Risk Quantification, we produced three reports summarising the current state of cyber risk quantification. The remit of the three reports can be briefly summarised as follows:

1. A systematic review of the academic literature on cyber risk quantification. The review includes: (a) an overview of quantification methodologies and their validity in addressing cyber risk challenges, (b) the use cases and arguments for and against cyber risk quantification, (c) the benefits and limits of cyber risk quantification, (d) prerequisites necessary for effective implementation and (e) gaps in the existing research literature.

2. A review of the cyber risk quantification literature in the context of specific challenge areas. The review includes: (a) the application of cyber risk quantification to aggregated risks, such as an entire sector or range of organisations, and comparing risk across sectors, (b) 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, and (c) the application of quantification to critical national infrastructure and/or safety critical settings.

3. A review on the feasibility of a standard model for costing cybersecurity incidents. The review includes: (a) categorising the costs (and benefits) of a cyber incident, (b) an overview of existing approaches for assessing the cost of incidents, (c) the pros and cons of quantifying incidents, and (d) recommendations on the use of standardised approaches for costing cybersecurity incidents.

In this summary report we provide an overview of the findings from each of the three reviews.

Photo by Risto Kokkonen on Unsplash

Definitions of Key Terms in Cyber Risk Quantification

Term	Definition
Framework	A foundational structure that provides concepts, principles, and best practices. Its role is to assist organisations in establishing a starting point for risk assessment, forming a basis for further development. Notable examples of frameworks include NIST CSF, OCTAVE, COBIT, the ISO 3100 Family, and FAIR.
Guideline	An instructional or recommendatory document designed to guide users. It suggests steps to aid in the risk assessment process. Frequently referenced guidelines encompass NIST SP 800-30 and the ISO 3100 family.
Standard	Denotes a universally agreed-upon risk assessment procedure adopted by an organisation or community. Standardized requirements or criteria enable users to compare their risk assessment results with those of other organisations. Prominent examples in this category encompass the ISO 27005 and ISO 27001 standards, along with NIST SP 800-30. For certain standards, like ISO 27005, there are certifications, which may influence the business case for adopting the standard.
Methods	A systematic, often mathematical, approach to assessing information security risks. This can encompass specific methodologies, examples include OCTAVE, MAGERIT, or Mehari, or the mathematical methods that are applied, examples including Fuzzy Theory, Bayesian Networks (BN), or Analytical Hierarchical Processes (AHP).
Tool	Can be a generic term for any cyber risk quantification method or framework. Often, though, means a specific process that can be used to conduct risk assessment. This can be, for example, documents to work through (e.g., PRAM) or a computer led process (e.g., MetricStream).

Table 1: Definitions of key terms used in this report

The Overview of the General Risk Cyber Quantification Landscape

A systematic literature review was conducted to encompass the cyber risk quantification landscape, identifying over 1,900 studies, with 713 deemed relevant to the review. In total, 137 frameworks, guidelines, and standards, along with 81 risk assessment methods, were identified to assist users in diverse ways in evaluating cyber threats. Our key findings from the comprehensive review of the academic literature on cyber risk quantification can be summarised as follows:

1. Large number of cyber risk quantification standards, guidelines, frameworks, and methods

We identified a total of 137 standards, guidelines, and frameworks discussed in the academic literature. The most frequently mentioned quantification approaches are (from most mentioned): ISO 27005, ISO 27001, OCTAVE, NIST SP800-30, CORAS, ISO 27002, COBIT, ISO 31000, FMEA, FAIR, and ISRAM.

#	Abbreviation	Name (long)	Price (approx.)	Freq.
1	ISO/IEC 27005	Guidance on managing information security risks	£177	126
2	ISO/IEC 27001	Information security management systems	£118	125
3	OCTAVE	Operationally Critical Threat, Asset, and Vulnerability Evaluation	Free	119
4	NIST SP 800-30	Guide for Conducting Risk Assessments	Free	89
5	CORAS	Consultative Objective Risk Analysis System	Free	69
6	ISO/IEC 27002	Information security controls	£197	50
7	COBIT	Control Objectives for Information Technology	Unknown	48
8	ISO 31000	Risk management - Guidelines	£88	46
9	FMEA	Failure mode and effect analysis	Free	43
10	FAIR	Factor Analysis of Information Risk	Unknown	32
11	ISRAM	Information Security Risk Analysis Method	Free	30

Table 1: Summary Table of Key Cyber Risk Quantification Standards and Methodologies

We also identified 81 cyber risk quantification methods, employing quantitative, qualitative, or semi-quantitative methodologies. Methods like Fuzzy Theory, Analytical Hierarchical Process, Common Vulnerability Scoring System (CVSS), Bayesian Networks, and Fault Tree Analysis were most frequently mentioned.

2. Cyber risk quantification approaches differ across a number of dimensions

There is wide variation in existing cyber risk quantification methods, frameworks, standards and guidelines. They can differ on the information needed as inputs into the process, the methods to identify, estimate, and evaluate risk, and the way the output is presented and communicated. Five specific dimensions of variability are (Vorster and Labuschagne 2005):

1. Whether approaches are self-directed and can be performed in-house, or require external expertise.

2. Whether they are best suited to risk assessment of individual assets, or provide a measure of overall organisational risk from specific risk scenarios.

3. The amount of information needed to inform the quantification process, and the trade-off between a quick approach or a more accurate approach that entails more resources.

4. The complexity of the assessment and output, and the trade-off between a simple and transparent approach or a more sophisticated but harder to interpret approach.

5. The type of output produced, e.g., whether risk is measured in absolute or relative terms.

There is little evidence on the comparative effectiveness of different methods and frameworks for cyber risk quantification. Moreover, there is no agreed benchmark for comparing methods. Many new methods being advertised in the private sector are not- disclosed (for IP reasons) and so are less open to scrutiny. The assumptions underlying open source methods are also difficult to compare and contrast.

3. The cyber risk quantification approach needs to be responsive to the needs of the organisation

There is no ‘one size fits all’ best method of cyber risk quantification. The optimal method or framework for a particular organisation will depend on the priorities and needs of that organisation, reflecting, e.g., the organisation’s risk profile, in-house expertise, financial resources and time availability. Organisations need guidance on how to choose a cyber risk quantification approach that is appropriate for their needs.

4. Cyber risk quantification should be a continuously ongoing process with input from across the organisation

Risk quantification assessment should be viewed as a continuous and routine process with input from across the organisation. This requires expertise (either in-house or external) and a commitment from across different departments.

New threats (e.g., deep fakes facilitated by AI) can make recent costly cyber security interventions ineffective while opening up new vulnerabilities. Hence, any cyber risk quantification exercise can quickly become out of date. This means that cyber risk quantification needs to be constantly under review to recognise the changing threat landscape and developments within the organisation (e.g., adoption of new software). The need for constant updating of knowledge is essential but can overwhelm organisations.

Effective cyber risk quantification needs accurate data and that means relevant departments (e.g., human resources, finance, research and development etc.) need to actively be part of the cyber risk quantification exercise. Moreover, output from the cyber risk quantification needs to be fed back to senior managers to inform policy and strategy. Risk quantification needs, therefore, commitment from across the organisation and should not be seen as solely a domain for the IT or cyber security team.

Continuous adaptation to reflect new threats and commitment from across the organisation requires dedicated resources to support risk quantification.

5. Organisations should look to develop in-house expertise to facilitate and/or implement cyber risk quantification

There exist tested and trusted cyber risk quantification frameworks, guidelines, standards, and methods that are ‘freely available’ to use and designed to be relatively accessible. While time and expertise is needed to implement such approaches it can be beneficial for organisations to develop in-house capabilities rather than solely rely on external providers. This can be facilitated by starting with ‘simpler’ approaches and building up capability over time. Specifically, given that cyber risk quantification needs expertise and significant resources it makes sense for organisations to build up experience, learn from positive and negative feedback, and evolve over time to more complex and tailored approaches.

6. Potential bias in cyber risk quantification approaches

Risk quantification methods can be systematically biased towards certain types of assets and risks, e.g., technical information, that are more ‘easily’ measured. This can neglect less tangible assets such as organisational knowledge. Cyber risk quantification approaches should also be based on a threat analysis that reflects the specific threats for the organisation.

Cyber Risk Quantification in the Context of Specific Challenge Areas

Our key findings on how cyber risk quantification approaches can be effectively extended or revised to specific challenge areas can be summarised as follows:

1. Aggregate risk

There is no agreed upon approach 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 (Tagarev et al. 2020, Welburn and Strong 2022). Evidence on aggregate risk tends to focus on systemic risk, particularly in the financial sector (e.g., Bouveret 2018, Orlando 2021) or risk to insurers (e.g., Zeller and Scherer 2022). Most work on aggregate risk has focussed on a value at risk (VaR) approach (e.g., Pal et al. 2021).

Value at Risk (VaR) Informally, VaR provides a probability distribution over potential cyber loses. 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.

In principal, 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. Measuring aggregate risk is complex because of two related considerations:

a) ***Economic and technological spillovers.*** An attack on one organisation may indirectly harm or benefit other organisations. For example, an attack may have negative spillovers on firms in their supply chain who face business disruption. In the case of negative spillovers a cyber risk quantification of one organization will underestimate the aggregate sector risk. The same attack may, however, have positive spillovers on competitors who can capture market share as a consequence of the attacked organisation’s loss of reputation. In the case of a positive spillover the cyber risk quantification of one organization will **overestimate** the aggregate sector risk. An aggregate cyber risk quantification needs to adequately account for both negative and positive spillovers. This requires data on market and sector networks of supply and competition.

b) ***Correlated risk.*** Cyber security incidents often entail correlated risks in which an attack on one organisation increases the likelihood of attacks on similar organisations (Amin et al. 2013, Pal et al. 2020, Welburn and Strong 2022). For instance, a supply side attack or zero day threat will impact multiple organisations simultaneously. 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. Most approaches to VaR and cyber risk quantification assume independent risks. This is a strong assumption in the case of aggregated cyber risk (Woods and Böhme 2021).

We remark that spillovers and correlated risk often interact. A supply side attack, for example, can simultaneously impact multiple organisations who also trade and compete with each other.

2. Top-down approach to modelling aggregate risk

We distinguish two broad frameworks to modelling aggregate risk:

i) ***A 'bottom-up' methodology*** 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 the most widely used to measure aggregate cyber risk. However, the independence assumptions underlying cyber risk quantification approaches mean they are arguably not well suited to modelling aggregate cyber risks. In particular, economic and technological spillovers as well as correlation of risk across organisations are not typically accounted for.

ii) ***A 'top-down' methodology*** can be used to quantify cyber risk by analysing aggregate level threats and contagion networks. This methodology looks to directly account for network interaction so as to measure spillovers and correlation of risk. Methods to estimate systemic risk tend to adopt a top-down methodology. Welburn and Strong (2022) provide one example. They study the NotPetya attack on Maersk and estimate upstream losses of between $663 to $773 million and downstream losses of between $16 billion to $19 billion compared to a direct loss to Maersk of between $250 to $300 million. The difficulty in adopting a top-down methodology is obtaining reliable information to inform the quantification exercise. Novel approaches such as prediction markets could be trialled to provide the needed information.

3. Markets in cyber risk quantification and aggregate risk

Quantifying cyber-risk is complicated by the fact that an organisation'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. Adverse selection can arise because of an informational asymmetry between the buyers of cyber-risk quantification tools, who cannot perfectly observe the quality of the product, and the vendors of those products. 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, in 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 tool is systematically biased in underestimating the loss from a zero day ransomware attack. Then there is an unaccounted for correlated risk that results from the cyber risk quantification market itself.

Information and transparency about the effectiveness of cyber risk quantification and mitigation tools are therefore crucial. A top-down methodology to modelling aggregate risk can take this into account by questioning bias in the cyber risk quantification market and trying to explore the market power of vendors and any consequences that may have for aggregate risk. To the best of our knowledge, however, there are no studies measuring market concentration in cyber risk quantification. Ideally the market would support diverse cyber risk quantification approaches that are well suited to the needs of organisations. It is unclear to what extent the market will deliver this ideal.

4. Critical national infrastructure

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. Existing risk quantification methods in these sectors primarily focus on internal threats and/or the failure of components about which there is reliable historical data.

For instance, risk analysis in the nuclear industry has primarily focussed on the reliability of physical components, an area where good historical data exists, and 'accidents'. Existing approaches and mindsets therefore, not easily adapted to consider cyber risk quantification with a continuously changing threat environment. In a comprehensive review of cyber risk quantification approaches for the nuclear industry, Eggers and le Blanc (2021) identify a wide range of gaps in existing approaches. Zhang and Kelly (2023) come to a similar conclusion and suggest a hybrid dynamic risk assessment model using Bayesian Networks.

Similarly, conclusions can be drawn about other sectors. For example, in the finance sector, banks' risk mitigation approaches are more effective in countering internal than external risks, leaving them vulnerable in the cybersecurity context (Pollmeier et al. 2023). In the health sector, Ksibi et al. (2023) provide a comprehensive overview of cyber security risk within e-health systems and argue current cyber risk quantification approaches are not well suited to the sector given the sector's diversity of behavior and capability, and heterogeneity of IT, coupled with a complex attack surface and continuously changing threat model.

5. Measuring costs and benefits of cyber security programmes and interventions

There are various ways, in principle, to use cyber risk quantification to analyse the cost-benefit implications of cyber security programmes. For instance, Erola et al. (2022) discuss how Value at Risk 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 cyber security interventions. Their approach is primarily based on using data of past events to extrapolate to future threats.

There are, though, challenges in using past data can inform on future threats. Woods and Böhme (2021) 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 and Böhme (2021) argue, however, that current studies are failing to do this and so creating unreliable results.

Approaches that combine past data with expert input show promise (e.g., Sheehan et al. 2021). Expert input can help overcome the problems with an over-reliance on sampling data from past incidents. FAIR, for example, has been criticised for an over-reliance on sampling and inflexible assumptions (Wang et al. 2020).

6. Measuring cyber security controls in context

The returns to the implementation of cyber security controls cannot be measured in isolation from other factors of the organisation. Factors to consider when estimating the financial returns to cyber security investment include:

a) ***Complementarity of Controls.*** The returns to cyber security controls will depend heavily on complementary controls. For example, the returns to implementing multi-factor authentication (MFA) will depend on the strength of other controls; if other controls are weak then MFA is likely to be ineffective.

b) ***Moral Hazard.*** Firms may exert less effort to prevent cyber-attacks if they believe their implemented method offers sufficient protection. 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.

c) ***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 and 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.

A Standard Model for Costing Cyber Security Incidents

Our key findings on the costing of cyber security incidents can be summarised:

1. Taxonomy of harms

There exist evidence based taxonomies of cyber harms that provide a good framework for estimating the cost of a cyber incident. Such taxonomies cover a wide range of costs including psychological and societal harms. Agrafiotis et al. (2018) distinguish the following five harm types (see also MacColl et al. 2024):

- Physical or Digital harm: Costs incurred from, e.g., damage, destruction, theft, modification, or corruption of assets. For instance, systems or data encrypted, data exfiltrated, IT infrastructure not functioning and backups being corrupted or destroyed.
- Economic harm: Costs incurred from disruption of operations, regulatory fines, investigation and PR costs. It also includes any financial consequences in terms of e.g., customer loyalty, reduced investment, loss of employment and loss of intellectual property, and increase in insurance premiums.
- Psychological and physical harm: Psychological costs incurred from, e.g., panic, anxiety, depression, guilt, anger, negative changes in perception. Physical costs incurred from exhaustion, sleep deprivation, burnout, and serious illness.
- Reputational harm: Costs incurred from e.g. damaged public perception, reduced corporate good will, negative impact on relationships with clients, suppliers and regulators, and reduced business opportunities. It also includes loss of staff and difficulty in recruiting new staff.
- Social and Societal harm: Costs incurred from services in society not functioning as expected as negative changes in public perception. Also costs from strained relationships and family life.

2. Established methods to estimate the costs of a cyber incident

There are well established methods to estimate the potential costs of a cyber incident. For instance, Monte Carlo simulation (Linsmeier and Pearson, 2000; Kavak et al., 2021) is a widely used probability-based technique that generates a range of possible outcomes for a given cyber threat by simulating the random processes that could lead to that threat materialising and, subsequently, estimating the potential costs. Woods et al. (2021) highlight the use of such statistical models in cybersecurity economics (see also Woods and Simpson 2018).

Erola et al. (2022) proposed a conceptual framework to calculate cyber value at risk (CVaR) using Monte Carlo simulations. Their paper details the successful application of this CVaR model in a real-world case study where there were millions of dollars of costs from a cyber incident. More recently, Franco et al. (2024) proposed another approach, the Real Cyber Value at Risk (RCVaR), which offers an economic estimation of potential financial losses from cyber incidents affecting various companies. This estimation leverages data drawn from public cybersecurity industry reports.

Classical approaches, including Monte Carlo simulations and other probability-based methods, have established the foundational principles for the financial quantification of cyber risks. However, these models are increasingly challenged by the rapid evolution and unique complexities of the cyber threat landscape. There are two main challenges:

a) Over reliance on historical data can lead to bias and misleading estimates of costs. It fails to adequately capture changes in the threat landscape, changes in controls that potential victims have implemented, and also changes in the cost landscape. The progression from Value at Risk (VaR) to Cyber Value at Risk (CVaR), and further to Real Cyber Value at Risk (RCVaR) signify a notable progression towards more sophisticated, forward-looking models that aim to encapsulate the full spectrum of cyber risk factors.

b) Evaluation of the full costs of an incident. Current methods tend to focus on specific costs that are ‘more easily’ measured in monetary terms, such as business disruption. They fail to take account of other costs, such as psychological costs, and may also not fully account for opportunity costs. The sequential nature of calculations within CVaR models (such as FAIR) mean inaccuracies or omissions can multiply through subsequent computations.

3. A standard model of costing a cyber incident

Building on existing approaches, particularly in terms of costing of intangible loss, we believe it is feasible to develop a standard model of costing a cyber incident that would be applicable to most organisations and types of cyber incident. It is, however, important that this standard model is developed based on a set of core principles that allow for transparency and an accurate estimation of costs. In the following we outline our key recommendations for a standard model of costing cyber incidents.

Photo by Risto Kokkonen on Unsplash

Recommendations for Costing of Cyber Security Incidents

1. Scope and purpose

A standard model needs to take into account the desired output of the user and scope of the costing exercise. For instance, the output could be a spot estimate of the cost of, say, a ransomware incident or a probability distribution over costs. Similarly, it could be an estimate of the total cost incurred or a timeline of costs as they occur over the short, medium and long run (Agrafiotis et al. 2018). Also, the costing exercise can be more or less inclusive in terms of the costs accounted for, e.g., whether social or spillover costs are included. The needs of a business or insurer in costing an incident may be very different to those of a government department or regulator and so the scope and purpose of the costing exercise need to be clear.

2. Transparency and flexibility of assumptions

The model should be clear and transparent in the assumptions used to cost the incident in order to yield meaningful, interpretable and robust outputs. Many current costing models are 'black box' making it difficult to verify the robustness of conclusions. Ideally, parameters in the model, e.g., discount rate, calculation of downtime and calculation of staff time, as well as quantification methods, e.g., monte-carlo techniques, and underlying assumptions, e.g., accounting versus economic cost, need to be clearly clarified. Moreover, the model should be flexible to allow for different underlying assumptions, e.g., estimating the cost of the incident for different values of the discount rate.

3. Clear guidance and reproducibility

The model should provide clear guidance on how to calculate the cost and, crucially, outline the data required to input into the calculation exercise. The organisation should perform mock exercises to verify that they would be in a position to capture the data required to calculate the cost of an attack. Crucially, the costing exercise should be reproducible in the sense that independent teams undertaking the exercise within an organisation would obtain similar estimates of cost.

4. Input from across the organisation

An evaluation of costs needs to have input from across the organisation, including IT, finance, HR and senior management. This is essential to guarantee the correct data is being fed into the costing exercise and so that all relevant costs are included. It also facilitates increased data availability and sharing across the organisation, e.g., data on staff absence and retention.

5. Reflects organisation's resource capacity

The model should reflect the resources and capabilities of the organisation to perform the exercise.
An effective costing of a cyber incident requires significant resource and expertise. Due consideration, therefore, needs to be taken of the organisations ability to perform the costing. For example, a small businesses will have less capability to perform a complex costing exercise than a government department. Crucially, the model should not seek simplification by narrowing scope. For instance, intangible losses (e.g., psychological stress) are likely to be particularly important for small organisations, like a micro business, and so simplification by omission is not appropriate. Instead the model needs to provide sufficient guidance and support to be simple enough to implement in the desired context.

6. Inclusion of intangible costs

The model needs to take account of intangible as well as tangible costs, e.g., psychological costs to staff impacted by the incident and loss of reputation to the firm. Current approaches primarily focus on tangible costs such as consultancy fees and replacement of physical assets. It is understandable that tangible costs can appear easier to source because they typically appear in financial accounts. There are, however, various tested, economic approaches that can be used to also quantify intangibles. For example, loss of productivity and/or loss of staff time due to sickness can be used as measures of psychological impact (Wolvetang et al. 2022). Similarly, reputational loss can be measured by looking at reactions in share prices (Cannas et al. 2009, Tweneboah-Kodua et al. 2019) or increased interest rate from debt. Such methods are currently under-developed in the analysis of cyber security incidents and so there is scope for more research exploring ways to measure intangible costs. Put differently, there is a need to move beyond a taxonomy of cyber harms and look to develop best practice methods of measuring the identified harms.

7. Dynamic flow of costs

The model needs to recognise that the costs of a cyber incident occur over time and there are often significant long run costs (Agrafiotis et al. 2018). Moreover, the long run costs of an incident (e.g., reputation loss) are difficult to quantify in the short run. The model needs, therefore, to allow for updates over time as new information becomes available. Moreover, costing needs to account for total cost, accumulated over time, recognising that cumulative costs can be considerable. Suppose, for instance, that a ransomware attack results in a loss of confidence and reputation with the victim’s customers and investors. This can change the entire future course of the organisation. Current costing models primarily focus on short to medium run costs that are more readily identifiable as a direct consequence of the attack.

8. Appropriate measures of economic opportunity cost

Current modelling of cyber incidents focus on tangible accounting costs. The appropriate metric, however, is economic cost taking into account opportunity cost. Specifically, outcomes following the attack should be judged relative to a forecast of outcomes without attack. Suppose, for instance, that service delivery falls following a ransomware incident. The appropriate metric is service delivery relative to that forecast. Costing of a cyber incident, therefore, needs to take account of forecasted outcomes. For a larger organisation one would expect that there are routine forecasts of key indicators that allow ready comparison. In a smaller organisation it may be necessary to produce ex-post forecasts based on extrapolating pre-incident data.

A cyber incident can have ‘benefits’ for an organisation, particularly in terms of ‘learning from experience’ and improving resilience and processes. Thus, at face value, some key performance indicators may improve as a consequence of an incident. For instance, investment in new IT infrastructure, post attack, may improve productivity in the medium run. However, such benefits should typically not be included in the costing of a cyber incident. Instead, the notion of opportunity cost should be invoked with a focus on attributable costs. In particular, it should be questioned whether the organisation could have ‘learnt from experience’ without the attack. In most instances, there are resources and information that could have enabled the organisation to improve their cyber security and would have been less costly than ‘waiting’ for a cyber attack. If, for instance, an organisation should have invested in new IT infrastructure then the cost of the new infrastructure should not be counted as a cost of the attack; similarly, the increase in productivity from the new infrastructure should not be counted as a ‘benefit’ of the attack.

Where to Next in Terms of Research?

Our review of the academic literature identified significant gaps in the existing research on cyber risk quantification. First and foremost we need more evidence on whether different cyber risk quantification approaches are effective at a technical level (i.e. the assumptions underlying the methods are sound) and at an implementation level (i.e. organisations can effectively implement the approach). We summarise as follows:

1. Comparative study of assumptions underlying cyber risk quantification approaches

There needs to be more research to analyse, compare and contrast the assumptions underlying different cyber risk quantification frameworks, standards, guidelines and methods. Approaches can differ across the implementation of the management process, including context establishment, risk identification and risk analysis, as well as communication. Current work tends to focus on comparing differences between approaches in terms of implementation (e.g., Wangen 2018, Sanchez-Garcia et al. 2022). There is a lack of current research analysing the modelling assumptions of different approaches and their advantages and limitations.

In recognising the need for more research analysing the assumptions underlying different cyber risk quantification approaches, we also highlight the difficulty of doing so. Our literature review identified several challenges in evaluating cyber risk quantification approaches. The assumptions underlying models are often not made explicit (or indeed set out in any transparent format). Moreover, different models are each framed in very different ways making it difficult to compare one model or approach with another.

We suggest first identifying an appropriate comparative framework, or starting point, with which to compare cyber risk quantification approaches. This could be based on a particular scenario of interest, e.g., ransomware attack, or asset of interest, e.g., loss of personal data. Different cyber risk quantification approaches can then be ‘run through’ that framework to identify underlying assumptions made and the robustness of those assumptions. Different comparative frameworks will likely favour different cyber risk quantification approaches because some approaches are better designed for scenario based and some for asset based analysis etc. By considering a range of different comparative frameworks it should be possible to obtain a better understanding of the strengths and limitations of different cyber risk quantification approaches.

A comparative study could be performed as desk based research. It would, however, ideally be combined with ethnographic studies of cyber risk quantification (see point 2 below). This is because the assumptions underlying different cyber risk quantification approaches may only become truly apparent when seeing how organisations implement the approach.

2. Studying the implementation of cyber risk quantification within organisations

The vast majority of the academic literature on cyber risk quantification has focussed on the theoretical properties of different models and approaches. There is a lack of work analysing how cyber risk quantification is practically implemented within organisations. Implementation is, however, key to effective cyber risk quantification and so more research is needed to study practical implementation. Existing case studies (e.g., Shedden et al. 2011 and Sheehan et al. 2021) have identified limitations with current approaches including potential bias towards assets that are more easily quantified.

An effective cyber risk quantification approach needs to be robust in terms of context establishment, risk identification, risk analysis and communication. One key measurement of robustness is that different people or teams (or the same person/team at different points in time) reach similar conclusions from the same underlying risk scenario. Anecdotal evidence suggests that this is not the case with current cyber risk quantification approaches. Research is needed to probe this issue in more depth. That research can take two complimentary forms:

a) An ethnographic study of cyber risk quantification in practice. Cyber risk quantification could be observed and studied in organisations that have experience of using cyber risk quantification. The study could map the cyber risk quantification process and identify strengths and weaknesses of different approaches.

b) Case study exercises in cyber risk quantification. Professionals with management and/or cyber responsibilities could be taken through exercises that capture particular aspects of the cyber risk quantification process (context establishment, risk identification, risk analysis or communication). The exercise could be run in, say, a two hour workshop. This controlled environment allows comparison of different cyber risk quantification approaches, and of different teams using the same approach.

3. Analysing and tracking cyber risk quantification implementation

There is a need for more research documenting the approaches currently being used by organisations. Our review focuses on the academic literature with the implicit assumption that approaches discussed more frequently in the academic literature are used more frequently in the wild. There is, though, a need for more work exploring the approaches that organisations are adopting and how approaches are evolving over time.

We were able to provide the results of a pilot study of 205 large and medium organisations in the UK. These showed that around two thirds of those surveyed are performing some form of cyber risk quantification. Future work could explore in more detail the type of approaches used (e.g., quantitative or qualitative), the purpose of the quantification (e.g., does it inform cyber security investment), how the results are used (e.g., are they reported to the board), what resources are available for cyber risk quantification, as well as the barriers and enablers to effective cyber risk quantification. It would also be desirable to compare across sectors, critical national infrastructure, size of organisation and whether they have cyber insurance or are subject to regulations.

4. Data for effective costing of cyber incidents

Cyber risk quantification relies on accurate and relevant data. This includes data on threat analysis, probability of breach/attack, and cost or loss in the event of a breach/attack. While improved data is needed for all three of these core elements we believe that there is particular potential for improved data on the costs of breach/attack. Costing of cyber security incidents currently relies on extrapolation from previous incidents. It is vital, therefore, we make sure that the costing of previous incidents is as accurate as possible and that the extrapolation exercise takes into account as many factors as possible, including the changing threat landscape.

Past attacks provide case studies with which to inform estimates of costs for both the organisation and wider supply chains and society (e.g., Welburn and Strong 2022). There is, though, a lot of ‘missing data’ and estimates tend to focus on readily available data. This can lead to systematic bias and also difficulty in accurately interpreting findings (Woods and Böhme 2021). Estimates also tend to focus on the aggregate cost while extrapolation requires more detailed understanding of the individual costs that contributed to the aggregate cost.

Future research could identify data that would be useful for costing cyber incidents but is currently not being collected or not being used. As a specific example, consider staff absence, staff turnover/retention and staff overtime during and after cyber incidents. Most large organisations should have such data that can inform a cyber costing exercise including the psychological costs of the incident. Similarly, costing of a cyber incident should not include costs that would be incurred anyway (e.g., IT upgrade) or were forecast (e.g., drop in sales because of an economic shock). Again, large organisations should have forecasts that can be fed into costing exercises.

We suggest identifying particular case studies of interest and analysing in depth the costing and extrapolation exercise. It should be recognised that the study of costing of cyber incidents is still in its infancy and more accurate models need to be developed over time.

The Data Challenge

A critical challenge in the field of cybersecurity is the prevalent lack of real-world data that authentically captures the dynamics of cyber risk and the effectiveness of security interventions. The necessity for field studies that involve actual organisations is paramount to establish the ecological validity of research findings. These studies are vital for reflecting true security scenarios as opposed to controlled environments that may not fully simulate real-world complexities. By studying cyber risk quantification in the field we can get a much clearer understanding of existing data gaps. Crucially, we then need to identify how those data gaps can be filled. We highlight three challenges with current data:

- The diversity and relevance of existing data sources require significant enhancement. Current research predominantly depends on a narrow range of data sources, mainly outdated industry reports. This reliance on limited data underlines the urgent need for more comprehensive datasets that encompass a broader spectrum of cyber incidents and their financial repercussions on firms.

- There are notable inconsistencies in how key variables are measured and reported across different studies and reports. This variability hinders the ability to compare and generalise results effectively, indicating a pressing need for more standardised data collection and reporting methodologies.

- There is a scarcity of empirical studies that directly link security controls with cyber risk outcomes. It is essential to conduct more empirical research to validate the efficacy of specific security measures and to uncover potential confounding variables that might skew these findings.

Advancing the science of cybersecurity and supporting evidence-based policy requires a concerted effort to gather more detailed, consistent, and real-world data. Such data is crucial for accurately quantifying risks and prioritising security controls as discussed by Woods and Seymour (2024). A theme throughout our work is that existing data is under-utilised. Cyber risk quantification requires a joined up approach in which data from across an organisation, and across organisations, is leveraged to inform the quantification approach.

Conclusion

This report distils the comprehensive findings and strategic recommendations from three detailed studies within the cyber risk quantification domain. It offers a high-level summary designed to encapsulate the insights and methodologies validated and utilised in addressing the multifaceted challenges of cyber risk.

Our work under the RISCS Cyber Risk Quantification Research Project delves into the quantification of cyber risks not only within individual organisations but also across aggregated risks spanning entire sectors, benchmarking risk across these sectors and underscoring the dynamic nature of the threat landscape. This project specifically considered the quantification's applicability to critical national infrastructure and safety-critical settings, reflecting the growing complexity and interconnectivity of these environments.

The research advocates for a versatile, multi-pronged approach to cyber risk quantification, taking into account the intricate fabric of economic and technological spillovers as well as correlated risks. The vast majority of extant cyber risk quantification research centres on intra-organisational assessments. However, the reports highlight an imperative for scalable methodologies that extend to aggregate measures. Our analysis suggests that systemic risk studies and prediction markets could offer valuable, complementary perspectives to enhance existing methods. We acknowledge the particularly volatile nature of cyber threats in sectors with complex risk profiles—like nuclear, finance, and health—and the insufficiency of historical data in developing future-facing cyber risk quantification. It is posited that a tapestry of methodologies, integrating expert insights, may yield the most comprehensive and precise outcomes.

In conclusion, the synthesised findings underscore that cyber risk quantification is not a one-size-fits-all endeavour. The adoption of any particular method is influenced by a myriad of organisational factors, from the size of the enterprise and the sensitivity of data to the expertise of users and the specific industry in question. No universally superior approach emerged from the studies; instead, a nuanced understanding of organisational context and a fusion of expert judgement with quantifiable data stand out as cornerstones for successful cyber risk quantification. This summary report encapsulates the essence of the research, aiming to furnish organisations with a foundational perspective to navigate the complexities of cyber risk in an increasingly digital world.

Photo by Shubham Dhage on Unsplash

References

Agrafiotis, I., Nurse, J. R., Goldsmith, M., Creese, S., & Upton, D. (2018). A taxonomy of cyber-harms: Defining the impacts of cyber-attacks and understanding how they propagate. Journal of Cybersecurity, 4(1), tyy006.

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.

Bouveret, A. (2018). Cyber risk for the financial sector: A framework for quantitative assessment. International Monetary Fund.

Cannas, G., Masala, G., & Micocci, M. (2009). Quantifying reputational effects for publicly traded financial institutions. Journal of Financial Transformation, 27, 76-81.

Eggers, S., & Le Blanc, K. (2021). Survey of cyber risk analysis techniques for use in the nuclear industry. Progress in Nuclear Energy, 140, 103908.

ENISA (2023). Demand Side of Cyber Insurance in the EU. Available at: https://www.enisa.europa.eu/publications/demand-side-of-cyber-insurance-in-the-eu [Accessed 05.07.2023].

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.

Franco, M. F., Künzler, F., von der Assen, J., Feng, C., & Stiller, B. (2024). RCVaR: an economic approach to estimate cyberattacks costs using data from industry reports. Computers & Security, 139, 103737.

ISO 2018. International Organization for Standardization ISO 31000:2018 Risk management – Guidelines. Available at: https://www.iso.org/standard/65694.html [Accessed on 26.01.2024].

Kavak, H., Padilla, J. J., Vernon-Bido, D., Diallo, S. Y., Gore, R., & Shetty, S. (2021). Simulation for cybersecurity: state of the art and future directions. Journal of Cybersecurity, 7(1), tyab005.

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., & Seifert, J. (2023). Secure Hardware Adoption in the Open Data Context.

Lam, W. M. W., & Seifert, J. (2023). Regulating data privacy and cybersecurity. The Journal of Industrial Economics, 71(1), 143-175.

Linsmeier, T. J., & Pearson, N. D. (2000). Value at risk. Financial analysts journal, 56(2), 47-67.

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.

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

Pal, R., Huang, Z., Lototsky, S., Yin, X., Liu, M., Crowcroft, J., ... & Nag, B. (2021). Will catastrophic cyber-risk aggregation thrive in the IoT age? A cautionary economics tale for (re-) insurers and likes. ACM Transactions on Management Information Systems (TMIS), 12(2), 1-36.

Pollmeier, S., Bongiovanni, I., & Slapničar, S. (2023). Designing a financial quantification model for cyber risk: A case study in a bank. Safety Science, 159, 106022.

Sánchez-García, I.D., Mejía, J. and San Feliu Gilabert, T. (2022) 'Cybersecurity Risk Assessment: A Systematic Mapping Review, Proposal, and Validation', Applied Sciences, 13(1).

Shamala, P., Ahmad, R. and Yusoff, M. (2013) 'A Conceptual Framework of Info Structure for Information Security Risk Assessment (ISRA)'. In Journal of Information Security and Applications, 18(1), 45-52.

Shedden, P., Scheepers, R., Smith, W. and Ahmad, A. (2011), 'Incorporating a Knowledge Perspective Into Security Risk Assessments', VINE, Vol. 41 No. 2, pp. 152-166.

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.

Tagarev, T., Pappalardo, S. M., & Stoianov, N. (2020). A logical model for multi-sector cyber risk management. Information & Security, 47(1), 13-26.

Tweneboah-Kodua, S., Atsu, F., & Buchanan, W. (2018). Impact of cyberattacks on stock performance: a comparative study. Information & Computer Security, 26(5), 637-652.

Vorster, A., & Labuschagne, L. E. S. (2005, July). A framework for comparing different information security risk analysis methodologies. In Proceedings of the 2005 annual research conference of the South African institute of computer scientists and information technologists on IT research in developing countries (pp. 95-103).

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.

Wangen, G., Hallstensen, C. and Snekkenes, E. (2018). 'A Framework for Estimating Information Security Risk Assessment Method Completeness: Core Unified Risk Framework, CURF', International Journal of Information Security, 17, 681-699.

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

Wolvetang, S., van Dongen, J. M., Speklé, E., Coenen, P., & Schaafsma, F. (2022). Sick leave due to stress, what are the costs for Dutch employers?. Journal of Occupational Rehabilitation, 32(4), 764-772.

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

Woods, D. W., Moore, T., & Simpson, A. C. (2021). The county fair cyber loss distribution: Drawing inferences from insurance prices. Digital Threats: Research and Practice, 2(2), 1-21.

Woods, D., & Simpson, A. C. (2018). Monte carlo methods to investigate how aggregated cyber insurance claims data impacts security investments.

Woods, D.W. and Seymour, S., (2024). Evidence-based Cybersecurity Policy? A Meta-review of Security Control Effectiveness. Journal of Cyber Policy, pp.1-19.

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.