# Risky Business: Quantitative Risk Assessments as Enabling Devices in Cybersecurity

Colette Alexander MSc | LUND UNIVERSITY

# Risky Business:

# Quantitative Risk Assessments as Enabling Devices in Cybersecurity

Thesis work submitted in partial fulfilment of the requirements for the MSc in
Human Factors and System Safety

**Colette Alexander MSc**

Under supervision of Dr Roel van Winsen

**Lund 2024**

# Risky Business:
# Quantitative Risk Assessments as Enabling Devices in Cybersecurity

Thesis work submitted in partial fulfilment of the requirements for the MSc in Human Factors and System Safety.

Colette Alexander MSc

Number of pages: 52
Illustrations: 2

Keywords
Quantitative risk assessment, cybersecurity, enabling device.

## Abstract

Quantitative risk assessment (QRA) is a growing practice in the cybersecurity field. This paper examines QRA the use in various industries and the problems with its use. The focus of the qualitative research is to understand why cybersecurity organizations might want to use QRA even if it produces untrue and potentially problematic results. It draws from other bodies of work that view QRA as a type of fantasy document and enabling device and posits that this could also be true within cybersecurity organizations. Interviews with Chief Information Security Officers (CISOs) and risk managers revealed that QRA clearly operates as an enabling device by aiding in budget approval with executives. Interviewees valued QRA for the perception of objectivity that it gave to others, even while understanding themselves that it was subjective. CISOs were more pragmatic about this tension, while risk managers who were more involved in the creation of the QRAs were more likely to want to have them continuously improved in the hope that they would eventually represent an objective truth. Even though it is often touted as a value of producing QRA, organizational learning was not an objective for any of the interviewees, and the method of collecting data for their QRAs was not always conducive to sharing information for broader learning. Overall, QRA clearly functions as an enabling device for the cybersecurity professionals interviewed, allowing them to advocate and receive crucial funding for cybersecurity projects.

Riskhantering och samhällssäkerhet
Lunds tekniska högskola
Lunds universitet
Box 118
221 00 Lund

http://www.risk.lth.se

Telefon: 046 - 222 73 60

Division of Risk Management and Societal Safety
Faculty of Engineering
Lund University
P.O. Box 118
SE-221 00 Lund
Sweden

http://www.risk.lth.se

Telephone: +46 46 222 73 60

## Dedication

This thesis is dedicated to the late Dr. Richard I. Cook, who inspired me to pursue the wonderful, messy world of human factors and systems safety by being a fantastic human and mentor.

# Acknowledgements

A whole village of people made it possible for me to complete this degree. First and foremost, my husband Eric, who took on the role of primary caregiver to allow me to spend hours at nights and weekends completing coursework and my thesis. My children, Gwen and Elvis, who brightened every one of my days whether I had a paper due or not. My parents and in-laws, who stepped in to help with childcare while I was at Lund Learning Labs and completing coursework. And especially Elvis' nanny, Jocelyn Vierzen, who made the daunting task of completing Course 3 and beginning my thesis with a brand-new baby in tow completely and totally possible.

Carsten Busch, Anne-Mette Petri and Noel Hengelbrook were phenomenal tutors. My peers in the 2024 cohort provided valuable encouragement and feedback during each of my courses, with a special thanks to Gareth Lock, my peer reviewer for all of course 2, and a steady source of help and inspiration during the entire course. Roel van Winsen guided me with pragmatism and generosity every step of the way during the thesis writing process. Anthony Smoker and Johan Bergström were unendingly patient, kind, and encouraging under every question and circumstance that I threw their way. Lastly, I feel very lucky that I can call John Allspaw a friend and mentor. He never stopped reminding me that I could complete this course and research even when I doubted myself, and he will, I hope, harangue me about how right he was forever.

# Contents

# List of Figures

# Abstract

Quantitative risk assessment (QRA) is a growing practice in the cybersecurity field. This paper examines QRA the use in various industries and the problems with its use. The focus of the qualitative research is to understand why cybersecurity organizations might want to use QRA even if it produces untrue and potentially problematic results. It draws from other bodies of work that view QRA as a type of fantasy document and enabling device and posits that this could also be true within cybersecurity organizations. Interviews with Chief Information Security Officers (CISOs) and risk managers revealed that QRA clearly operates as an enabling device by aiding in budget approval with executives. Interviewees valued QRA for the perception of objectivity that it gave to others, even while understanding themselves that it was subjective. CISOs were more pragmatic about this tension, while risk managers who were more involved in the creation of the QRAs were more likely to want to have them continuously improved in the hope that they would eventually represent an objective truth. Even though it is often touted as a value of producing QRA, organizational learning was not an objective for any of the interviewees, and the method of collecting data for their QRAs was not always conducive to sharing information for broader learning. Overall, QRA clearly functions as an enabling device for the cybersecurity professionals interviewed, allowing them to advocate and receive crucial funding for cybersecurity projects.

# 1. Introduction

Every day around the world, organizations of all kinds are faced with managing risk. Risk is the possibility that an undesired event or loss could occur and prevent the organization from achieving its objectives. Organizations attempt to manage risk by naming, analyzing and prioritizing it, and assigning resources to address it (Hubbard, 2020; Peace, 2017). During these steps, risk assessments are often created and referenced by safety experts along with subject matter experts (SMEs) at various levels of work (Peace, 2017).

Risk assessments can take many forms: some are structured to represent various risks as green, yellow, and red - color coded to communicate their importance or severity (Leveson, 2019). Managing organizational risk is at the heart of every cybersecurity department's mission, so it is not surprising that the skill of estimating, understanding, and communicating risk is fundamental to those who work in the field (Barnum, 2021; Sutton, 2017). In cybersecurity circles, quantitative risk assessment (QRA) is proposed as a better solution to risk management when compared with more 'traditional' ways of estimating risk, usually in the form of a qualitative risk matrix (Hubbard & Seierson, 2016; Freund, 2014). Probabilistic risk assessment (PRA) is a specific kind of QRA that assigns numerical, statistical probability to risks, and a numerical monetary range for the losses associated with those risks (Hansson & Aven, 2014; Hubbard, 2020). Some broader forms of QRA include only estimated losses and are not concerned with likelihood. For the purposes of this paper examining the reasons behind using quantitative prediction of any kind, QRA will be used to refer to both PRA and QRA.

In software engineering organizations, risk is present in the form of vulnerabilities in code, or architectural issues that can potentially expose information, or that simply threaten the overall quality and availability of their products. Cybersecurity is explicitly concerned with the risk of cybercrime: external and internal threat actors that can breach systems and steal data, attack the systems of customers, or shut down systems and hold them for ransom (Sutton, 2017). Risk assessments in software engineering organizations are often completed, or heavily informed by managers or individual contributors closer to the front lines of operation and surfaced to higher level management such as vice presidents or Chief Information Security Officers (CISOs) or Chief Technology Officers (CTOs) who have the power to ask for, approve and prioritize resources to address the risks laid out in them. While there are some data privacy laws that impact elements of risk in software products, when compared to more regulated industries like aviation or the nuclear industry, there are relatively few laws that govern the actions or create requirements for software companies around how they manage risk. Instead of relying on regulation, customers concerned with risk in the software that they buy rely primarily on third party auditing firms who assess whether an organization has sufficiently managed certain kinds of risk, including cybersecurity risk (Peace, 2017). Even though there are significant questions about the efficacy of QRA in the nuclear industry (Downer, 2013; Downer & Ramana, 2020), it is beginning to be adopted and advocated for in the realm of cybersecurity (Hubbard, 2020; Hubbard & Seierson, 2016).

## 1.1 The Use of QRA in Cybersecurity

Many organizations that create software or own digital assets of some kind have a cybersecurity department to help protect those assets. Cybersecurity activity is largely comprised of five activities that are well defined in the cybersecurity framework published by the National Institute of Standards and Technology (NIST) (2018):

**Identify**: identify assets that need to be protected, and the context that they exist in

**Protect:** implement protective capabilities for those assets (can take the form of tooling or training)

**Detect:** put monitoring in place to detect if a security breach has occurred

**Respond:** Define an incident process for responding to breaches

**Recover:** Develop plans for resilience against a breach to assist in incident breach

Risk, specifically the risk associated with an attack on an organization's digital assets, is at the core of the NIST framework. The official framework documentation recommends running a risk assessment to establish or improve any cybersecurity program of work (NIST, 2018). Risk assessments, then, are embedded in the processes of any cybersecurity team or department.

Perceived risk of a breach in the organization's operations and business will determine the size, nature, and functions of its cybersecurity group. Investment in cybersecurity is commensurate with the expected level of risk mitigation from such an investment. Because of this, financial and healthcare firms who have a great deal to lose from breaches tend to invest much more of their budget in cybersecurity than a small retail business does (Columbus, 2023).

For understanding levels of risk, QRA is perceived as more accurate than the qualitative risk matrix, and examples are often used where the process of doing QRA in some format helps an organization re-classify risks from higher to lower categories. The result of risks being lowered in severity saves teams time and allows them to focus on risks that are categorized as more likely or costly. For communicating risk, especially to non-technical decision makers, QRA is viewed as a much stronger tool than the qualitative risk matrix. Especially when discussing with executives who oversee budgeting and organizational management, the language of probabilities and financial costs is viewed as far more impactful than the simple high/medium/low categorization in the qualitative risk matrix (Freund, 2014; Hubbard & Seierson, 2016).

Advocates for QRA in cybersecurity suggest first analyzing historical data to gather probability and cost likelihoods (Freund, 2014; Hubbard & Seierson, 2016). If historical data is unavailable, they suggest gaining more accurate probabilities in prediction by "calibrating" SMEs, asking them about probabilities for discrete components or events that have been decomposed from more complex systems and incidents. This type of calibration is viewed as a best practice for generating probabilities in QRAs across industries (Winkler, 1986). The last step in generating a QRA involves averaging decomposed probabilities across multiple experts to reach a more accurate set of probabilities. Uncertainty can be expressed in results if desired by asking SMEs to rate their levels of uncertainty and injecting those values into various statistical simulators to generate ranges of estimates, rather than single numbered outputs (Hubbard & Seierson, 2016).

From these estimates, statistical probabilities and potential losses are calculated on a per-incident or scenario basis. Various methods like fault trees help to map dependencies between various components or scenarios. Some QRAs move from individual scenario probabilities into mapping the interaction between various failure modes using probability generators like Monte Carlo simulators and Continuous Time Markov Chains to inject randomness and simulate uncertainty in statistical probability ranges (Frenkel et al., 2014; Hubbard & Seierson, 2016). The outputs of QRA are usually in the form of matrices describing risk and costs for a single threat, potentially translating them in the form of a color-coded heat map as in Figure 1. Losses can be estimated as primary (direct fines, ransom paid) or secondary (customer communication costs, new customer loss through reputation damage).

## Figure 1

## An example of QRA

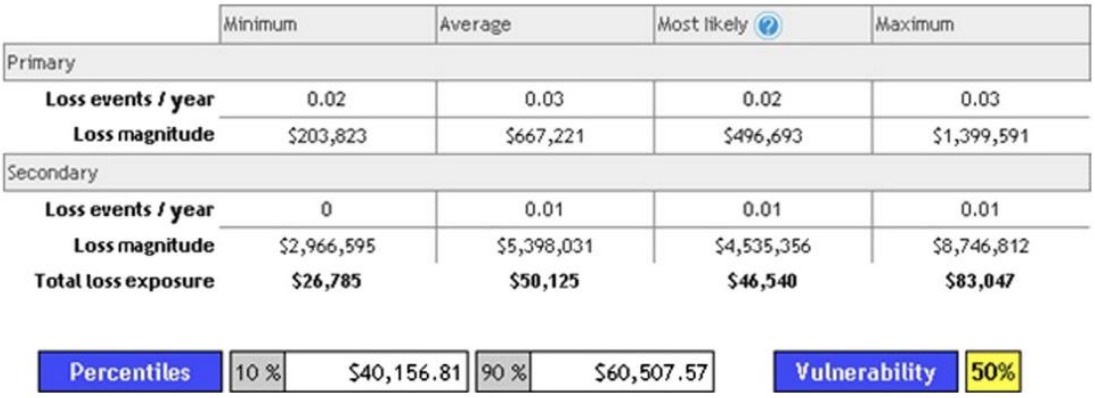

| | Minimum | Average | Most likely | Maximum |
|---|---|---|---|---|
| Primary | | | | |
| **Loss events / year** | 0.02 | 0.03 | 0.02 | 0.03 |
| **Loss magnitude** | $203,823 | $667,221 | $496,693 | $1,399,591 |
| Secondary | | | | |
| **Loss events / year** | 0 | 0.01 | 0.01 | 0.01 |
| **Loss magnitude** | $2,966,595 | $5,398,031 | $4,535,356 | $8,746,812 |
| **Total loss exposure** | **$26,785** | **$50,125** | **$46,540** | **$83,047** |

Percentiles 10 % $40,156.81 90 % $60,507.57 Vulnerability 50%

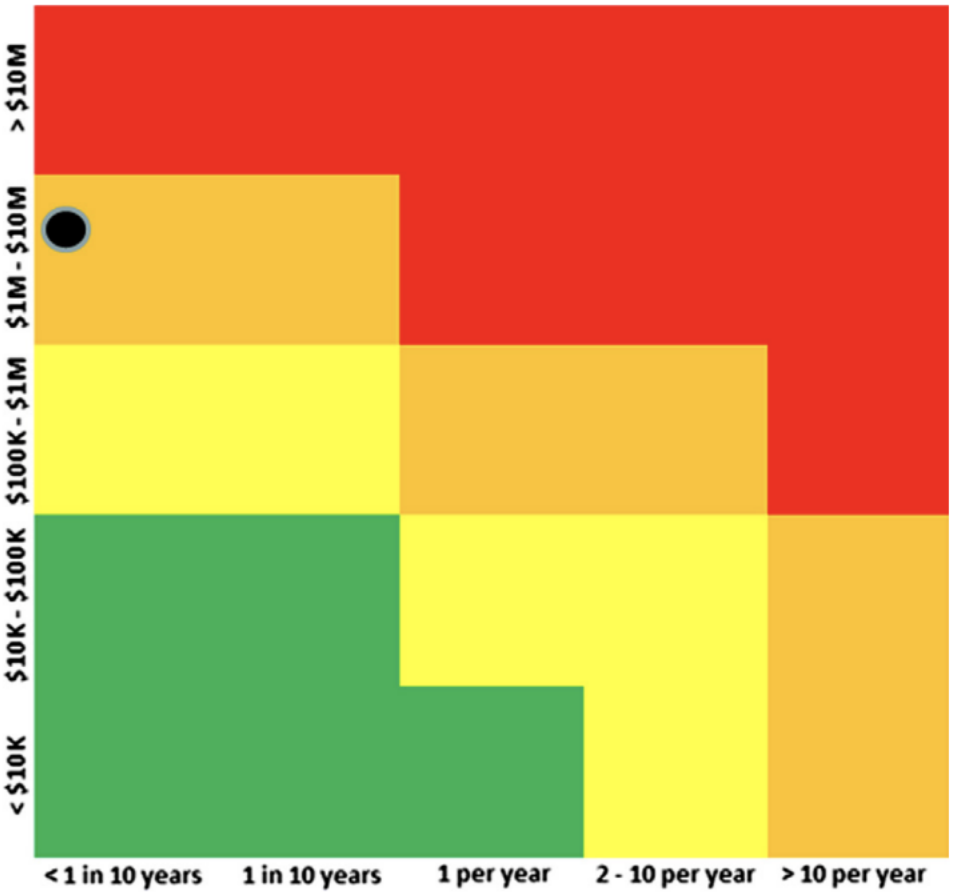

*Figure 1:* From *Measuring and Managing Information Risk – a FAIR Approach* (p. 107, 113), Freund & Jones, 2015, Butterworth-Heinemann. Copyright 2015, Elsevier, Inc.

## 1.2 Complexity and Modern Software

The words "complex" and "complexity" are used frequently in discussions around risk assessments and safety in socio-technical systems such as nuclear power plants, airplanes, and software systems. Charles Perrow first discussed complexity and coupling in 1984 in his book *Normal Accidents* (1999). Inspired by the nuclear accident at Three Mile Island, Perrow uses the accident to frame his broader theory of high-risk systems, pointing out that nuclear power plants are both highly complex and tightly coupled. Complexity is a characteristic of any system where components can interact in varied and potentially unanticipated ways. The opposite of complexity in his definition is linear: where a system has a set, ordered operation that it cannot deviate from. Coupling refers to how components relate to one another to accomplish a goal: tight coupling means that there are required sequences that are needed to meet a goal (often within a particular time frame). Loose coupling might be indicated by parts that are substitutable, or processes that are repeatable to correct mistakes. The riskiest technologies are systems that have both tightly coupled and highly complex interactions such as nuclear power plants, aircraft, space missions, and even DNA technologies (Perrow, 1999). For the sake of brevity, "complexity" in the rest of this paper will refer to both the concepts of complexity and coupling that are present in Perrow's work.

Modern cloud[1] software has many hallmarks of complexity that evoke Perrow's definitions. Allspaw (2015) outlines some main characteristics of that complexity in *Trade-Offs Under Pressure*:

- It is highly opaque with many layers of functionality and dependencies embedded in it.
- It relies on a dynamically distributed network (the internet) that is constantly changing and totally decentralized.
- The teams that manage it are often globally distributed. They make changes to shared code and dependencies at different times, often unbeknownst to one another.
- It is an open network, and its content is both consumed and produced by its participants.

[1] The "cloud" essentially means the internet, though public cloud and private cloud are distinguished in their meaning. Public cloud consists of hardware and software infrastructure that a provider (such as Amazon, Microsoft, or Google) leases to customers. Private cloud refers to that same type of infrastructure being owned and operated by a company themselves.

These elements make the diagnosis of problems, the coordination of technical resources and human expertise, and control of content and information flow very difficult problems.

One example of the complexity of modern cloud software are the dependencies that it contains, and how those are managed. Modern software applications often contain dependencies on multiple modules or packages that are open source which are written and maintained by groups of software engineers outside of the companies that are using them (Butler et al., 2022). Each of these packages can have multiple updates or releases per year for security patches, bug fixes, or adjusting functionality to manage other dependencies that the software relies on to work smoothly (Butler et al., 2022). Keeping up with these constant updates on many components means that in the span of a few months the software that a single company is running may process hundreds of changes authored by engineers that it does not even employ (Butler et al., 2022; Tapas et al., 2019). These changes are in addition to any features or bug fixes that the company itself desires to add to its source code. The dependencies within these systems and the precariousness that they can cause is well illustrated in a cartoon titled *Dependency* by Randall Munroe (Figure 2, Munroe).

**Figure 2**

**Dependency**

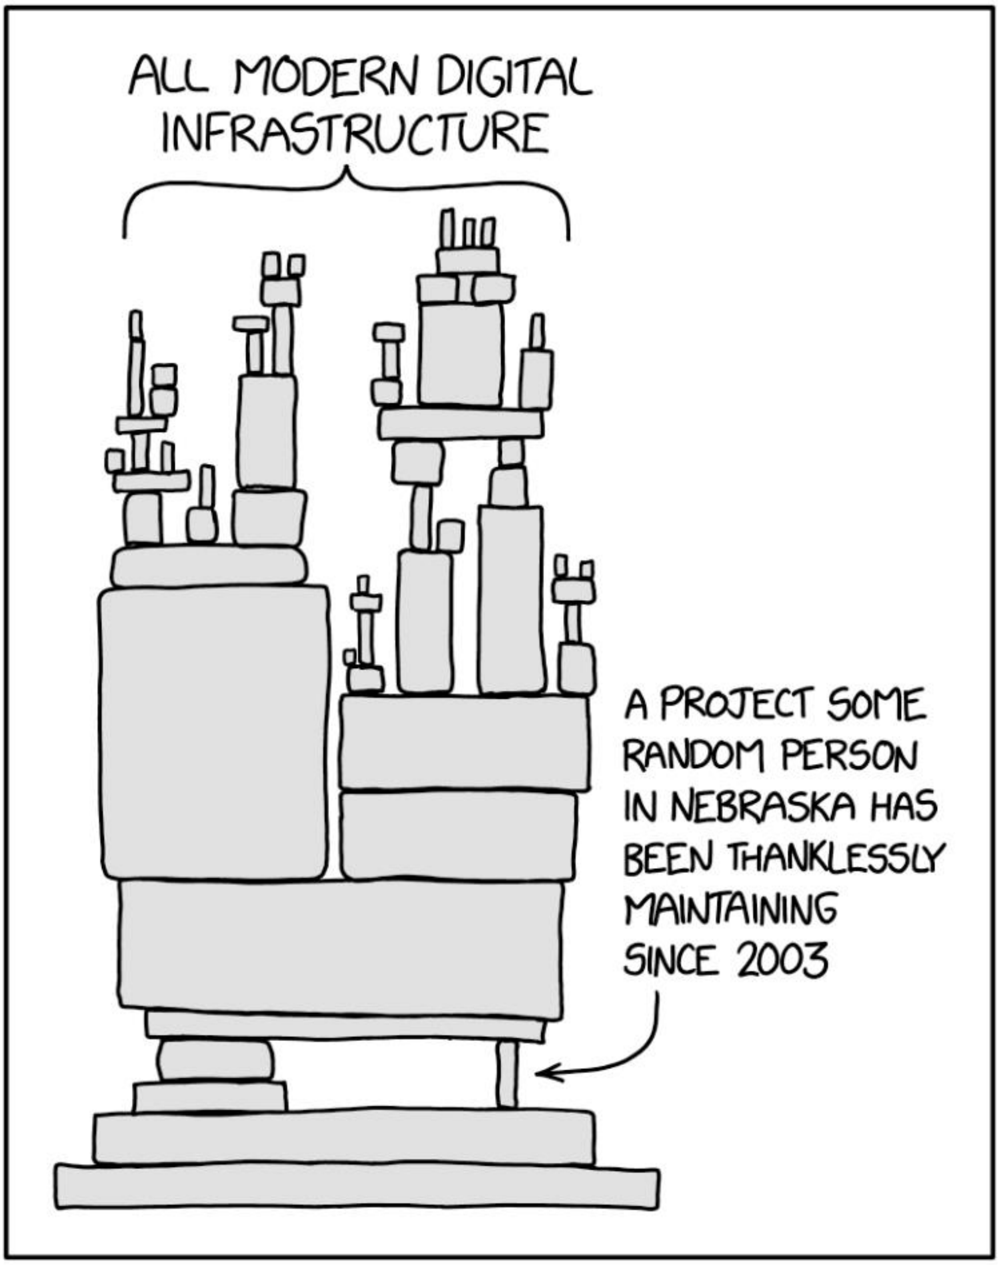

ALL MODERN DIGITAL INFRASTRUCTURE
A PROJECT SOME RANDOM PERSON IN NEBRASKA HAS BEEN THANKLESSLY MAINTAINING SINCE 2003

*Figure 2: Dependency* [Cartoon], Munroe.

## 1.3 The Problems with the Use of QRA in Cybersecurity

*"The future seems implausible; the past seems incredible"* (Woods & Cook, 2002, pg. 329)

*"Things that have never happened before happen all the time"* (Sagan, 1993, pg. 12)

There are many mathematical equations and algorithms that can be used to generate numbers and ranges of predictions for QRAs. There are so many possible algorithms that contribute to them that there are entire textbooks devoted to modeling QRAs for a variety of industries and environments (Frenkel et al., 2014; Kumamoto, 2007; Bedford and Cooke, 2001). On their surface, QRAs appear to be purely quantitative, but they are based on two decidedly qualitative inputs and assumptions: that historical events can predict the future, and that SMEs can predict the probability and the cost of incidents. The failure of QRA to accurately predict the future lies with these practices and the downfalls that accompany them.

Relying on historical data to project probabilities into the future has a high degree of failure within complex systems because they and the environment they operate in change and introduce novel risk over time (Leveson, 2019). A poignant example of historical data guiding a risk assessment to a disastrous result is the Challenger shuttle explosion. A launch decision was given go ahead not despite previous evidence of O-ring damage, but because that previous evidence in the presence of no major accidents proved to those making decisions that there was no real risk involved (Vaughan, 1996; Wynne, 1988).

Humans are notoriously bad at estimating risk, falling victim to many biases in the process of determining frequency probabilities. While some heuristic biases can be countered by calibration practices with SMEs, others are unrelated to statistics per se and are more difficult to manage. Some examples of heuristic biases in prediction can be found in Nancy Leveson's *Improving the Standard Risk Matrix: Part 1* (2019) where she summarizes a body of research from Kahneman and Tversky (1973) and Kahneman, Slovic and Tversky (1982): confirmation bias, availability heuristic, ease of scenario generation, difficulty in predicting cumulative causes, incomplete search for probable causes, and defensive avoidance all play a part in efforts to predict risks in complex systems. Each of these biases may or may not be at play in any given QRA, but collectively contribute to making its content both incomplete and highly uncertain.

A best practice when composing a QRA and gathering SME input is to decompose complex systems into smaller components (Winkler, 1986). From there, estimates for smaller components can be recombined into larger, interdependent models that can run multiple probabilities against each other for a holistic result (Hubbard & Seierson, 2016). The fundamental assumption behind these practices is that the complex systems they are assessing can be appropriately decomposed into smaller pieces by SMEs for the purposes of predicting anomalies or accidents. There are two problems with this assumption, though: SMEs in software engineering cannot know or understand the full set of interactivity or risk present in the systems that they work in, and these complex systems tend to create emergent phenomena that cause unique failure modes that are not able to be predicted (Cook, 2020; Dekker, 2011).

Cook (2020) establishes that in complex modern-day software systems, no single person can know or understand the system in its entirety: "As the complexity of a system increases, the accuracy of any single agent's own model of that system decreases rapidly" (p. 47). Because of this degraded understanding, models of the system or predictions of risk inherent in the system are also flawed.

Dekker points out that the assumption that a system can be decomposed into smaller pieces to make predictions about is Newtonian thinking - that the system contains linear cause and effect chains, and that the system can be fully known at a singular point in time. But, in complex systems this knowledge does not exist. No one can know the full state of a complex system, and no one truly understands the laws that govern a complex system (Dekker, 2011).

The exponential dimensionality of highly complex and tightly coupled systems allows for more frequent, less consequential failures to combine all at once and create new, emergent, catastrophic failure that cannot be anticipated (Leveson, 2019; Perrow, 1999). A single snapshot in time of the system (as QRAs represent) also tends to ignore the messy realities of work on the ground such as operator behaviors over time, and management or political changes that influence the systems (Leveson, 2019; Perrow, 2011). The adaptation of people and social systems in real time can change the trajectory of a technical risk dramatically, sometimes within seconds.

An example of how difficult it is to model the exponentially high probabilities in a complex system involves the very first QRA report created, WASH-1400 (a nuclear reactor safety study). One of the scenarios modeled was a pressurizer relief valve failing to close, which is exactly what kicked off the partial meltdown at Three Mile Island (TMI) one year after the report was released (Keller and Modarres, 2005). To some, the prescient inclusion of this failure mode of this pointed to the success of QRA as a technique in the nuclear industry. But WASH-1400 failed to anticipate the operator reaction to this failure - operators turning off the ECCS - which followed and exacerbated the incident (Keller and Modarres, 2005). The complexity of the socio-technical system on the ground was not represented meaningfully in the report, and so did not effectively model the whole system at play during the TMI incident.

Systems complexity can also lead to a common misinterpretation of quantitative risk reports: a highly complex system might have potential catastrophic errors with an assessed likelihood of one in 100 million. This very small likelihood may soothe someone reading about them in a document into thinking that these events are so unlikely as to be impossible. But because the system is so complex, there are a long tail of these types of errors that are possible. With 100 million catastrophic scenarios that have a one in 100 million probability, failure in a system is much more likely than these individual statistical probabilities initially indicate to a layperson (Woods, 2022).

*"Essentially, all models are wrong; but some are useful. However, the approximate nature of the model must always be borne in mind"* (Box & Draper, 1987, p 424).

The oft quoted maxim from George Box, that all models are wrong, but some are useful, is rarely combined with the second sentence that follows it: a warning to interpreters of statistical models and their results to be wary of the room for error within them. Within a QRA, this room for error is often expressed as two distinct types of uncertainty: aleatoric, true randomness, especially as it relates to mechanical failure, and epistemic, the uncertainty of SMEs who are providing subjective probabilities for modeling (Parry, 1996; Winkler, 1996). Hubbard and Seierson (2016) find that the predictive capabilities of SMEs trained in QRA are no better than those who are untrained. But they claim that those trained or calibrated to proper statistics mindsets with decomposition methods and rigorous questioning are more likely to have a better understanding of the uncertainty inherent in their answers. Inserting this uncertainty into various algorithms is meant to calculate windows of probabilities instead of single, static numbers allows for QRA to incorporate this dimension of the prediction and communicate it to others. But what about scenarios where SMEs are highly uncertain? Or are unaware (because of a cognitive bias or the complexity of the system they are assessing) of their actual levels of uncertainty? The practice of QRA becomes understandably more difficult to extract value from when the complexity of the system being measured (and thus, the envelope of uncertainty) increases dramatically and dynamically as modern software systems do.

It is also important to think of the consumer of QRA documents when discussing uncertainty. As Downer and Ramana (2020) point out, "consider the dilemma that would arise if experts were to assert they were 99.9999999 percent certain a reactor will not fail over a given period, but only 80 percent certain that this number is correct. Such a statement would be almost nonsensical, as the second variable should already be implicit in the first" (p. 6). While it can be dizzying to consider, critique, and examine the massive, complex models and algorithms designed to incorporate all these probabilities into one document, it is perhaps more important to ask whether a model with so many approximations and uncertainties baked into it can provide value at all.

With an examination of its methods, and with evidence post-hoc especially of various nuclear accidents, it has been argued that QRA cannot measure risk in an objectively true way (Downer, 2013; Rae et al., 2020). QRA has seen many doubts about its efficacy raised after high profile nuclear accidents such as Three Mile Island, Chernobyl, and Fukushima (Downer, 2013). Fukushima stands out as a more recent example of the failure of QRA to adequately predict failure and help organizations adapt to the possibilities of catastrophe. Using QRA prevented the organization from preparing for the unexpected successfully: the failure to build a seawall to contain a larger tsunami than had happened in the area before was based on "probabilistic thinking, not thinking about what was possible" (Perrow, 2011, p. 47).

There is much critique on the empirical truth of QRA but far less literature examines the forces in organizations that compel its use. Current literature questioning the use and validity of QRA has not caused it to decline in use: it is currently advocated for in many industries, including the software domain (Kumamoto, 2007; Mosleh, 2014; Rao et al., 2018; Hubbard & Seierson, 2016; Hubbard, 2020). QRA is woven into safety requirements in nuclear regulatory organizations, which lends it perpetual credibility as a tool within the safety world (IAEA, 2002; USNRC, 2020).

QRA might be particularly ill-suited as a tool to measure risk in cybersecurity given the complex and ever-changing landscape of code and infrastructure that internet enabled computer software entails. Despite this, it is seeing adoption within the field (Hubbard & Seierson, 2016). This research seeks to understand why QRA is used in cybersecurity including what kinds of organizational factors might be at play in its adoption.

# 2. Literature

Examining the literature of why an organization might use QRA led to research that examines QRA as an artefact of organizational culture. This body of work suggests various sociological factors might be at play for those who adopt the practice. It examines why organizations may use QRA and other types of objectively untrue documents and practices despite their shortcomings. In doing so, our focus turns from QRA and its inner workings towards a more sociological perspective of QRA: how does QRA function in an organization? What are some of the sociological forces that might compel its use? The literature explores these questions and provides a set of insights that frames this original research within cybersecurity organizations.

## 2.1 QRA as an Enabling Device

When viewed as an artifact in culture, QRA can be classified as a type of 'fantasy document', a term coined by sociologist Lee Clarke who defined a class of documents that were created with the intention of mitigating risk in case of disaster, but which often in practice contain unrealistic assumptions and beliefs that make them useless when disaster actually strikes (Clarke, 1999). Fantasy documents can have both positive and negative impacts in organizations and contribute to difficulties and even disasters. On the positive side, they can reveal information about how people learn in an organization and uncover and spread knowledge previously hidden within organizational silos (Clarke, 1999; Hutchinson et al., 2018). On the negative side, if people within organizations come to believe these documents are objectively true, they can ignore their own experience and evidence in front of them which can lead to disaster (Clarke and Perrow, 1996).

More specifically, QRA is an 'enabling device' (Hutchinson et al., 2022). Like a fantasy document, an enabling device creates risk in organizations when people attribute objective truth to it that it does not actually contain. Where fantasy documents focus on mitigating or managing potentially high-risk scenarios, enabling devices are more a part of everyday work. The purpose of an enabling device is to facilitate work - as in, greenlight a construction project or move to the next phase of a plan, by decreasing uncertainty with various partners (internal or external customers, leaders, government institutions, the public) around safety concerns (Hutchinson et al., 2022).

QRA is an enabling device within the nuclear industry. As Downer (2013) notes, probabilistic assessments "legitimate almost all contemporary public pronouncements and policy discourse about nuclear power" (p. 6). The goal of QRA in this space is to allow a particular nuclear power plant to achieve regulatory and public approval.

The framing of QRA as an enabling device confronts its subjectivity: creators of QRA often have an objective in mind when they create them, whether that is to advocate for or against a particular investment or project. As employees in an organization, they may have a stake in preserving certain business activities or goals that justify their continued employment or are simply more interesting to them than other activities. Interpreters of QRAs have similar biases: a set of business targets to meet or a belief that a specific goal or project should be invested in.

## 2.2 QRA and Expertise

QRA can provide value to organizations by uncovering and expressing expertise within the group. This gathering of expertise can also potentially spread knowledge around the organization that was previously hidden (Hutchinson et al., 2018). In the case of software engineering, SMEs sharing expertise about the systems they work with could serve as a moment of building awareness of the current state of software systems in the organization both with one another and with higher level management who are traditionally more removed from on-the-ground operational realities. This is a typical gap found between the sharp end of practitioners and blunt end of management (Cook et al., 1998). With QRA, the combination of risk scenarios that are generated when calculating various probabilities together can potentially identify important and previously unrecognized risky events (Keller and Modarres, 2005).

The communication of expertise outside of the specialist group and to a broader community within or outside of an organization can be fraught with misunderstanding and miscommunication, though (Hutchinson et al., 2018; Vaughan, 1999). The group with less expertise that is consuming a QRA usually has a great deal of power within an organization – they might be managers or members of the c-suite or board (Hutchinson et al., 2022). They will not necessarily be able to observe when people and systems deviate from the equations and system depicted in the QRA. New risks can emerge quickly in complex systems, and if those with decision making or budgetary power believe the important risks are managed, that could set the stage for catastrophe (DMAIB, 2017; Hutchinson et al., 2022).

## 2.3 QRA to Project Control and Soothe Uncertainty

QRA utilizes expertise to project control over complex and even dangerous technology (Downer, 2013). The transformation of subjective knowledge into numbers serves to soothe uncertainty through its appearance of objectivity, which allows business activity to commence (Hutchinson et al., 2018).

Numbers have an appearance of objectivity and can rise above qualitative arguments when advocating for a particular position or activity (Hutchinson et al., 2018). The appearance of objectivity is all that matters - the fact that inputs to QRA are highly subjective and unable to accurately predict calamity does not alter the impact that the appearance of numbers in a QRA has on public perception (Downer, 2013).

Using numbers allows QRA to be viewed as an objective measure, positioning facts against the feeling of fear around nuclear reactors within the public (Downer, 2013). QRAs use numbers to convince partners (the public, government officials, shareholders) that nuclear power is not as dangerous as other activities such as driving or flying in a commercial airplane (USNRC, 1975). These low probabilities for catastrophe clear a path for business activity. The goal of the US Atomic Energy Commission introducing QRA was to shift the conversation in the public to one around the very small likelihood of catastrophic failure occurring, rather than possible catastrophic scenarios that would result from nuclear power plant operation (Downer, 2013). In the nuclear use case, QRA was not meant to exhaustively model all possible disasters and address risks within these models, it was to change the conversation started in the first assessments of nuclear power plant risks that Congress commissioned which highlighted worst possible outcomes and assuage the public's fears of nuclear meltdown in their own neighborhoods (Downer, 2013; Rip, 1986). In this way, QRA transforms the risk of catastrophic possibility into acceptable risk through asserting probabilities.

A QRA document can also serve the function of soothing people into thinking they have properly controlled for any risks in their business. Lee Clarke discusses fantasy planning as a way for organizations to project competence: "When uncertainty about key aspects of a task is high, rationalistic plans and rational-looking planning processes become rationality badges, labels proclaiming that organizations and experts can control things that are, most likely, outside the range of their expertise" (Clarke, 1999, p. 4). This soothing can make accidents more likely to happen. If the risk assessment and subsequent mitigations put in place do not actually mitigate certain risks effectively, people in various roles in an organization may struggle to recognize signals or indications of serious issues, believing them to be improbable (Clarke & Perrow, 2012; DMAIB, 2017; Hutchinson et al., 2022).

In an examination of literature that addresses why QRA is used despite its lack of objectivity, the case for QRA as a fantasy document and enabling device in organizations is strong. That QRA can serve to uncover expertise within an organization, and that it can be used to project control over risk and thus soothe partners and shareholders is also apparent. This thesis seeks to draw on these sociological foundations of enabling devices and fantasy documents to understand there are any other themes or motivations for the use of QRA when it is adopted by cybersecurity practitioners.

# 3. Research Design

## 3.1 Epistemology and Methodology

This research seeks to discover what value QRA provides people creating them, and how people who are presented with QRA use and understand them. It aims to study this specifically in a cybersecurity context. It is concerned with constructed meaning around a particular object, the QRA, and so the research will take a constructionist approach (Crotty, 1998). One point of exploration is if QRA in an organization holds different meanings for different people, which is directly related to a constructivist viewpoint (Creswell, 2009).

Qualitative studies are well-suited to examining constructed meaning (Crotty, 1998; Creswell, 2009). The ethnographic methodology, including a combination of observation, interviews, and document analysis are the most advantageous ways to collect data for exploring meaning within a group setting (Crotty, 1998). Unfortunately, the direct observation aspect of ethnography is not a great fit in the current day software engineering world. Teams are often distributed around the world and more difficult to observe directly. With people consuming documents and collaborating asynchronously, observation does not contain the insight that it does in a more traditional single office team location. Documentation related to cybersecurity is also difficult to anonymize effectively to protect proprietary information, which would be necessary to do in order to publish research. Therefore, this research primarily relies on interviews with individual subjects.

## 3.2 Interviews and Participants

Interviewees included people who interacted with QRA documents: the creators of them as well as consumers within an organization. While it has been established in the introduction that QRA is a poor tool for predicting risk in an objectively true manner within complex socio-technical systems, research subjects who are using it in the field may have different beliefs. Therefore, the interviews were approached in as open-ended way as possible to understand the context that QRA is used in the organization, the relationships between the people who use it, as well as the meaning they ascribe to the QRA document. Questions focused on the events surrounding the creation or consumption of QRA, but allowed for free form, open-ended discussion related to those subjects.

Research participants were people who interact with QRAs as an artifact in an organization. They have specific roles related to the QRA and varying degrees and types of influence and power in an organization and over each other. They may have created QRAs or consumed them (or both). A note on sourcing interviewees: there was no filter placed on participation based on skepticism of the value of QRA within their organizations. The requirement for sourcing was only that participants had interacted with a QRA (created or consumed it) in their work and consent to be interviewed. Interviewee selection was not biased for any size of organization or type of organization. Cybersecurity is a common practice within many different types of businesses, but there was no anticipation that those various types or their size would determine attitudes towards QRA. In total, the following roles were sourced for interviews:

**Two (2) Chief information security officers (CISOs)**

- CISOs lead the security organizations in software organizations. They manage all sorts of security functions within a company, including security engineering (who create software engineering security solutions for company products and services), physical security (for office buildings, manufacturing plants, etc.), IT security (managing employee devices, access, etc.), compliance teams and engineers (managing audits and other compliance work) (Shayo and Lin, 2019). They are sometimes tasked with presenting QRAs to board members. They also might consume QRAs that are constructed by members of their organization for other reasons. CISOs are a powerful role in an organization, reporting directly to the CEO, sometimes to a lesser c-suite position (CIO, head legal counsel, etc.) (Williams, 2007)
- **Sourcing:** CISOs are very busy professionals who may not be initially interested in participating in safety research. Personal professional networks as well as professional networks that support the use of QRA (like SiRA, the Society of Information Risk Analysts) to source participants here.

**Three (3) Creators of QRAs**

- Many different roles within a security organization might create a QRA, and for a variety of reasons. This is the largest cohort of interviewees, partly because they are the easiest to access, and partly because as creators, they lie at the center of the generated meaning behind QRAs.

- **Sourcing:** Personal professional networks were relied on as well as public networks that promote the use of QRA in cybersecurity (like SiRA, the Society of Information Risk Analysts).

Research participants interviewed had different industry sizes, types, and backgrounds where they performed QRA. Interviews were conducted online, and an informed consent form (see Appendix A) was presented to participants who signed it. The content of the consent form, with specific emphasis on data storage and confidentiality guidelines was discussed at the beginning of the interview, with any concerns or questions being addressed at this point. Interviews lasted from 60-90 minutes and were recorded when possible. If consent for recording was not given, note taking was relied on.

## 3.3 Analysis

Transcripts of interviews were coded with themes related to the use of QRA. Specific attention was paid to patterns or discrepancies across the types of roles (risk manager vs. CISO) interviewed. Themes were also related back to the themes that existed in the literature review. Thematic analysis occurred in between multiple rounds of interviews, something that Braun and Clarke (2006) describe as a grounded theory 'lite' approach. Abductive reasoning was used to develop theory when comparing additional rounds of interview data and relevant scientific literature.

# 4. Results and Findings

Recruiting interviewees proved difficult at first, though not because of a lack of interest in participation. Many people signed up to participate in interviews, but when put through an initial screen to ensure they had, in fact, used QRA in their work, it was discovered that some had not. A small part of the group who initially signed up were not able to successfully schedule interviews in the three-month time frame that they were performed. The final subjects that made up my interview cohort all had experience with QRAs used in a cybersecurity context. They used them in various industries: insurance, health care, not-for-profit, government and higher education institutions were represented among the cohort.

The QRA types used by interviewees fell into two distinctive categories: the first is Factor Analysis of Informational Risk (FAIR), a method of quantifying cybersecurity risk that categorizes that risk and uses estimated probabilities for events combined with projected losses (Freund and Jones, 2014). The second is a more general use of quantitative data to help measure risk: some used probable monetary losses for types of events (ransomware, email compromise, etc.), some used scoring on cybersecurity audits, and they sometimes combined this information with historical events inside or outside of the company as added information. I've set context throughout these themes on which type of QRA is being referenced by practitioners.

Patterns that were discussed broadly in the literature review on QRA, including its function as an enabling device, to soothe uncertainty, and to uncover knowledge or to influence decision making provided initial themes for directing questioning in interviews (Hutchinson et al., 2018; Hutchinson et al., 2022). Open-ended discussion allowed exploration of the "why" behind using QRA for each person. From that discussion emerged a more nuanced take on QRA's perceived objectivity emerged from the interviews, though, which emphasized the ability of QRA to function as an enabling device with those more removed from knowledge and expertise in the organization.

## 4.1 Advocating for Budget with Executive Partners

The most consistent reason brought up for using QRA in discussions with CISOs and risk practitioners alike was the need to convince executives to either continue to fund or increase funding for cybersecurity work at the firm. Every single interviewee had used QRA to convince executives how to direct organizational budget. People who used QRA for this purpose believed that the perceived objectivity of numbers helped make QRA more effective in advocating for budget or project work than a qualitative risk matrix. CISOs consistently pointed out that the budget in a company was a zero-sum game, played against other areas of the business that were competing for money. One CISO described the scenario for budget decision makers:

> You're making investment cases, you have a limited pool of budget. I have $100 million to spend[...] What do I invest that in? Do I invest it in reducing risks? Do I invest it in marketing campaigns to go and get new customers? […] Do I invest it in hiring new staff? Do I buy a company with it? With my $100 million, everything there is quantitatively measured. The entire business case is money. (CISO 2)

They went on to emphasize the competitive nature of asking for budget as a cybersecurity leader:

> Convincing people of the merits of the risks and the value of the investments is a quantitative discussion. And QRA is an absolutely critical tool within that. Because without it, I'm going in there against somebody saying, let's put in Salesforce, automate this, get rid of 20 people in our environment and add $5 million in revenue. Okay, I need to go and spend $5 million dollars this year to do it. But next year we'll generate $10 million more revenue, we will have saved a million dollars a year in staffing costs […] If I'm going in there with 'this one's red' […] I'm on the backfoot. I'm in a losing environment. And I do not have an equal seat in that discussion. (CISO 2)

The CISO is convinced that without numbers for their budget proposal, and specifically dollars at risk, they will not be able to successfully advocate for investment into their cybersecurity programs. QRA is perfectly set up to provide support for their advocacy.

The effectiveness of QRA in cybersecurity to allow for comparison to other programs at a company stood out for a Risk Manager who was interviewed as well. Speaking about the presentation of the FAIR modeled QRA that was used, he mentioned its value to some executives: "I think it really helped illustrate for them how cybersecurity risks were on par with other risks that they typically understood better and were maybe more front of mind" (Risk Manager 2). They went on to discuss those other risks: pay outs to customers for not meeting contractual obligations around reliability were much more tangible for these leaders than the potential monetary losses for a cybersecurity event.

Another CISO discussed using NIST audit scores, and pitching the level of investment necessary to get to a better score for their organization to executives who controlled the budget.

> I said, "if you invest like $50,000, then we're going to get to the blue line, if you invest $100,000, and [...] if you invest 200,000 [...] we can get to this, this red line." And I showed our like, the height of where we were at. And I showed where I would be able to take us with each, you know, with the quantities, but they immediately invested the 200,000. (CISO 1)

By using the quantification of a score and predicting that the dollar amount investment would increase scores a certain amount, the CISO was able to use numbers to predict a lowering of cybersecurity risk (symbolized by a higher score on the NIST audit) for the organization. This use of quantification is more abstract and less tangible than the probabilities and dollar losses that the FAIR method advocates for, but the use of numbers to advocate for budget still stands out as a consistent practice.

In one case, a risk manager (3) worked with a Chief Financial Officer (CFO) to convince them to not invest in a particular tool due to the risk it would cause the organization. A colleague had suggested the use of a particular third-party tool that the risk manager understood to be extremely vulnerable to various forms of cyber-attack. Initially, they investigated doing a qualitative risk matrix but quickly concluded that it would not necessarily convince their CFO to ditch the tool: "When I sit down with my CFO and tell her the risk is 'medium' she is going to say, 'what the f*** does that mean?'" (Risk Manager 3) Compiling a FAIR model, the risk manager presented the risks in a quantitative format, emphasizing the many scenarios where the tool could be hacked and the probabilities they would occur, and what the result could be - including catastrophic financial loss for the organization. In the end, the CFO was responsive to the assessment and nixed the investment in the tool. "The monetary loss piece really landed it for her [the CFO]" (Risk Manager 3). Using QRA, the risk manager was able to successfully appeal to the CFO's risk appetite, and advocate for a different direction in selecting a third-party vendor.

Not all executives buy into the analyses they're presented, though. A Risk Manager (2) described presenting their FAIR assessment with probabilities and expected costs and finding that the CFO of their company was the most skeptical about projected possible losses. The CFO believed that the projected cost of reputational loss measured in the QRA was too big. The practitioner acknowledged that this could be based on the CFOs previous experiences:

> The industry that he came from was not heavily regulated. Quite honestly, no one would even care that much […] if all the data was breached. And so he just had a very different perspective on how much should be invested in security and how damaging security breaches could be […] I do think that that made him a skeptic upfront. (Risk Manager, 2)

This CFO's skepticism that losses would be high is borne out by the market, historically: data breaches do not always permanently damage a firm's reputation or have negative financial consequences (Spanos & Angelis, 2016).

A risk manager (1) noted the focus on budget that comes from discussing cybersecurity risk with finance executives: "Our finance director - they're all the same. Their opening question is 'good morning.' The next question is 'how much is that going to cost me?'" When asked to elaborate on what information is given to the finance director, they go on:

> I'm really focusing on getting the control model that's going to be based on NIST, then it can speak to risk. So, this is our new risk question from the board. And go "Okay, these are the key controls [related to the risk question from the board], these are the ones I know about, these ones I don't know about." So I know what I don't know. (Risk Manager 1)

This risk manager provides quantitative ratings for various controls in the NIST risk framework - they're identifying risks to the business and adding up/averaging the elements in place to help mitigate for those risks (also known as controls) in a score. The score is displayed, but with color coding that is similar to a risk matrix. Color coding data like this has been shown to ease understanding by consumers of information (Liu et al., 2021).

> Our non-technical management quite liked the NIST dashboard. So it's literally just the main columns, each section with a percentage in it. And because they're senior executives, you've got to make it a little bit more [...] executives want something more visual. They don't want to be reading a lot of text. They just want something nice and simple. It's easy to understand. If there's a nice style, say nought to 100. And we've color coded in say, well, above 33%. You're at your Amber. Then it gives them something to look at and focus on. (Risk Manager 1)

The scores and dials are presented to executive leadership by the CTO, highlighting what risk the organization is taking on and what risk they're undertaking in operations, which is meant to be approved by the board as 'risk appetite': "we report through the CTO route, that goes to our board and council [...] you kind of got to express that so they can understand. But it has come afterwards: 'Well, this is our risk appetite'" (Risk Manager 1).

Executives and boards signing off on risk appetite is at the center of any budgeting discussion for a cybersecurity organization: any unfunded programs or choosing to remain at incomplete remediations for risks will be viewed as being acceptable exposure for the company by the board. The risk appetite for an organization is never zero. There are remediations for cybersecurity risks that are unrealistically expensive, or impossible to implement for companies. In these cases, CISOs will highlight these risks to the board and ask for sign-off on not funding them, including them as part of the company's risk appetite. Setting a risk appetite with the board is part of the package of asking for funding for various projects. While funding a project purports to mitigate risk to the controller of budget, risk appetite states the opposite: what risks will not be mitigated. One CISO mentioned fielding questions from the board about a real-world case just before their interview:

> Look, yes, look on Wednesday. The announcement from Microsoft about the compromise of Federal Government Office 365 by Chinese state sponsored hackers. I get the question: "What have we done about this?" My answer: "you didn't set me as a target to keep state sponsored hackers out. That is not my target. That is with the risk appetite as specified by the board. Because the budgets to keep them out are so ridiculous that the NSA can't do it. Next!" (CISO 2).

A CISOs advocacy for a cybersecurity budget only goes so far: the risk appetite set by other leaders at the company (likely the CEO and board of directors) will determine whether a program will get funded or not. There is both a pragmatism at play when discussing setting risk appetite, and a real sense of needing to displace liability in the legal sense:

> I know I cannot protect everything. I explain to people what I'm working on protecting and say, 'Are you happy? Because you're the legally accountable executives?' Yeah. And if they say 'Yes', well I make sure I get it in writing. (CISO 2)

What the CISO means by saying "legally accountable" is that the board and C-suite of an organization hold a fiduciary duty to shareholders that can be called into question if they are seen to be negligent in enforcing standard cybersecurity protections for customers (Chaput et al., 2021). Setting risk appetite for cybersecurity investment is based on practicality and realism: CISOs know that they cannot invest limitless budgets in pursuit of stopping every possible breach. In some cases, though, setting risk appetite with sign off from the board is about accepting legal responsibility if that appetite is determined to be inappropriately high compared to the risks present for consumers and shareholders.

In all these cases, using quantitative assessments of risk with executives was viewed as a highly effective way of advocating for investments or specific decisions to be made about investments in cybersecurity. QRA is an enabling device for advocacy and funding between the presenter (whether a CISO or risk manager) and the executive group controlling the budget and major investment decision making. The users of QRA prefer it as an advocacy tool because they believe it projects a stronger case against other potential uses of the budget, and because it speaks in a monetary language that executives can comprehend. By expressing objective numbers, users of QRA believe that they're soothing the uncertainty of executives around what the right investments to make in the business are. QRA also functions as a mechanism to project control: by asserting that budgetary investment in some areas will protect the company from financial losses, users of QRA are asserting that they can control the future cyber risk of the company with monetary investment. QRA operates as a "rationality badge" (Clark, 1999, p. 4) to executives, assuring them that the objectivity of comparing numbers across budgets allows them to make the correct prioritization decisions.

## 4.2 Pragmatists vs. True Believers

Those who use QRA because they believe it is perceived as more objective do not necessarily believe that the tool is objectively true themselves. While practitioners might question the objectivity of the data they present, they still expected that the receiver of the information would find it objective and convincing. This theme adds nuance to the emphasis on the power of numbers and objectivity that much of the existing literature on QRA mentions (Hutchinson et al., 2018). Many interview subjects that I spoke with asserted that the data they were using was full of subjectivity and bias. The power of QRA lies with the perception of objectivity of the data by non-specialists including budget approvers and other peers that the cybersecurity experts interacted with. CISOs tended to balance the subjectivity of QRA against the power and impact of the data presented and pragmatically buy into what they viewed as a flawed model, whereas Risk Managers tended to want to make the data approach objective truth, even if they wrestled cognitively with the subjectivity that was present in it.

One CISO mentioned the power and influence data had, even while acknowledging its subjectivity:

> There are certain people who believe that it's not real unless you can measure it. And I don't have a lot of faith in the measurements that I'm doing in terms of their accuracy in terms of their repeatability. I do have a lot of faith in my ability to use something like that to coordinate behavior, though. (CISO 1)

Another CISO also reiterated that QRA was meant to be “accurate, but not precise” (CISO 2) and that their job was to invest in the appropriate things to try to mitigate risk, but that it was possible that this would not be enough. Following through on this thread with them, the CISO was asked what happened if the QRA turned out to be objectively false.

> **Interviewer**
>
> What happens to you? If one of those risks that you've assured an investment is enough for and controlled for or whatever for? What if that happens, and it is more catastrophic than..
>
> **CISO 2**
>
> I'm a CISO, you know what it stands for? Career Is Swiftly Over.

This CISO understood the consequences of a cybersecurity event occurring, despite their reassurances on controls and mitigations in place: they would be fired, and perhaps unhireable afterwards.

Unlike the CISOs above, risk managers desired QRA to be more objective, and expected the QRAs they produced to hold up as objectively true in the future. Risk managers had slightly different responses from the CISOs above when asked what would happen if the models proved to be not objectively true in the near future. All of them immediately suggested that they may have missed something or made a mistake with the model, or that their models were wrong and should be updated.

One risk manager responded that the end result of a catastrophic data breach previously deemed improbable would be similar to the CISOs assertion above: they would "pack up their desk and go home" (Risk Manager 3). But first, they asserted that their model was likely wrong and just needed updating: "You know, like with global warming, the one in 100-year floods are happening more often now, so the models need updating" (Risk Manager 3). This assertion, that an updated model could eventually approach objective truth, was a consistent response with the risk managers that were interviewed.

Another risk manager had a process-oriented response to an incident occurring that contradicts the model they present:

> If it's one of my processes, if something really went you thought it was low, it's caused a big problem. Hey, that's a security incident. So we log it as a security incident. And going through the process of okay, how do we fix it? But in our process for incidents, we have a root cause analysis. So we then go through and go, Why did we miss that? And that's that iterative process of going, Oh, we didn't think of that in our classification. Or we didn't, you know, we need to tune our model. So it's not do it once and forget about it. (Risk Manager 1)

Some risk managers had a sense of the subjectivity of their data, often expressing doubts of their own about its accuracy or of the accuracy of the SMEs who were tasked with coming up with data for the predictive models. One risk manager discussed this bias and their concern that SMEs who had a particular interest in a type of risk might overestimate its probability or cost:

> When they see a question about data exfiltration, and that's what they've been thinking about for the past month, and that's what keeps them up at night. They're just gonna respond quickly with their gut, with a, you know, a very high estimate of likelihood. I think there's probably more along those lines. I don't have, you know, evidence to support this. [...] In my perception, in glancing at some of the responses, I would say, oh, you know, that's so and so they really want to work on [a data exfiltration] project that looks like a high estimate to me. And I guess I could also see their estimate compared to you know, the other two or three people's estimates, and it was quite a bit higher in that case. (Risk Manager 2)

Are these SMEs overly concerned about data exfiltration, or are they simply so expert in it that they understand how much of a concern it should be to the organization? Are they overestimating the severity of these types of incidents, or are the other 2-3 estimates underestimating? These are questions that are impossible to know the answers to without the benefit of hindsight. This risk manager clearly understood that the data they were collecting from SMEs was subjective, but also has a desire to critique this subjectiveness and to transform their QRA into something more objective.

Ironically, the same risk manager responded with a decidedly subjective take when asked what the first step for the team that created a QRA was after getting initial results from SMEs:

> Actually, we, first of all, within the security team, went over the results together, initially, kind of as a sanity check, and see if things Yeah, just kind of felt right, if they jived with our general feelings, what keeps us up at night, things like that. And they did. (Risk Manager 2)

This reliance on a gut check when examining data directly contradicts the risk manager's previous concern about SMEs going more with their gut, subconsciously overweighting things that they care about in their estimates over others.

While exploring the tension in using QRA when it turns out to be objectively untrue, it is notable that CISOs who were interviewed were consistently pragmatic about its usefulness, even in the face of it turning out to be less than accurate. Those who created QRAs, however, were prone to wanting to update models, or question whether an SME might be overweighting elements of their predictions because of bias. They had a clear desire for QRA to represent some more objective truth than the CISOs did.

## 4.3 Lack of Emphasis on Organizational Learning

Some advocates of QRA emphasize that when many SMEs come together to create a QRA for an organization, the presentation and elicitation of knowledge is valuable and generates learning among peers (Rae et al., 2012). During interviews with research participants this benefit of QRA was not voluntarily mentioned as a motivating factor for using it, and the mechanisms that would provide for the most learning were not present in the practice of developing the data for QRAs.

For one risk manager utilizing FAIR, the collection of responses from SMEs happened via a survey: "we provided them with a survey, essentially, I think it was a Google form, to solicit their input on any risk scenarios within their purview" (Risk Manager 2). The surveying of individuals asynchronously made a discussion between SMEs about their predictions unlikely, and not facilitated by experts in estimation. The estimates submitted via the form were averaged across one another for a smoothing effect:

> The end result was we took the responses to those surveys, it went directly into Google Sheets. And then we did all kinds of Google Sheets foo and aggregation and math, to average out those estimates for both the likelihood and the impact. And we ended up with an overall [...] likelihood as a percentage of an event occurring over the next year. (Risk Manager 2)

Another risk manager pointed out the problem with interviewing at scale (in this case to scan for risks and aggregate a score), and how they solved the problem:

> It depends on the firm, because [company name], had 10,000 staff [...] so that's not going to happen. So but you kind of go by department go, who's the heads of department who the key teams are. Yeah, it'd be lovely to interview everybody, but you just don't have the resources to do it. So by doing it by key department, and department heads, take a sample of it tends not to be the very junior staff, it tends

> to be the staff that have been there a while and understand the business. (Risk Manager 1)

The emphasis here was on sampling interviewees but focusing on those with the most experience and knowledge. This technique makes clear that knowledge elicitation about the internal systems is the goal. But sharing that knowledge was not a result, rather they were immediately transformed into the risk manager's own product: a set of quantitative controls that are presented to leadership to let them know where their biggest risks lie.

When pressed about the possible value of sharing knowledge, a risk manager who had to do some interviews as well as surveys points at the value of interviewing people providing a richer picture of the problem space than simply surveying for numbers:

> So I did have conversations with maybe roughly half of the participants. [...] There were certainly good conversations. And whereas, or I guess backing up for a moment. In these surveys, you know, we asked people for their estimates. And then for each and every question, say something along the lines of You know, "please add any thoughts or notes here", just for my kind of edification, and then we could also incorporate those into future risk assessments, right, if someone brought something to light that we hadn't been aware of previously. Very few people who completed the survey independently added much in the way of notes. But anytime I would have a conversation with people, I would be filling out those notes myself. And we really did uncover some good information that could be incorporated. So I would say that was certainly a benefit of having conversations with people. (Risk Manager 2)

It is clear in this single case that having individual conversations with SMEs about the systems at play provided important context for the risk manager. Group conversations between SMEs could certainly lead to a similar outcome: increased organizational learning. But there are real organizational barriers to providing one-on-one interviews or facilitating SMEs in group interviews to spread learnings even further: scale. The risk managers interviewed for this research each collected SME contributions for their QRAs in ways that did not encourage or enable synchronous sharing and communication among participants. They spoke to both the scale of their organizations and in one case the fact that employees worked remotely as reasons for their asynchronous methods of participation. As the systemic knowledge is transformed into numbers for quantitative representation, and averaged across many participants, so is the nuanced and detailed information about the system behind those numbers.

# 5. Discussion

## 5.1 QRA as an Enabling Device in Cybersecurity

There is a great body of work that examines QRA, explains the processes behind it, or seeks to advocate for its use instead of a qualitative risk matrix. While the debate about whether QRA can be considered objectively true can be compelling to participate in, it does not seek to understand the motivations behind its use. Understanding the motivations behind the use of QRA are crucial first step to being able to replace it with a better tool. A smaller amount of literature examines the sociological forces that compel the use of QRA and viewing it as an enabling device that allows work to move forward highlights the objectives it achieves: soothing uncertainty and projecting control, and potentially allowing the distribution of expertise and knowledge throughout an organization. The goal of this research was to understand the adoption of QRA in cybersecurity, and to examine whether it functions as an enabling device with its attending themes.

QRA clearly functions as an enabling device in cybersecurity for every single one of the people who were interviewed. The most common activity QRA enabled for interviewees were budget approval for their cybersecurity organizations. QRA's enabling power lies in its perceived objectivity, which also reflects existing literature examining its use (Hutchinson et al., 2018). Even when interviewees were aware of the lack of objectivity underlying the production of QRA, they believed that the perception of objectivity by other key decision makers was what ultimately made it a powerful tool. They were convinced of its persuasiveness even if they were not all convinced it held objective truth. The use of QRA to enhance organizational learning, another possible reason for use explored in the literature review, was not volunteered as a reason for creating it by any of the interviewees. When questions approached the possible value of learning in the process, there was some belief that it could be valuable, but an acknowledgement that the methods often used to create QRAs, especially in highly distributed work environments like software engineering companies, did not make achieving that learning very likely.

Diving into the world of those who use QRA in cybersecurity was at times discombobulating: at one turn someone would discuss the difficulty of managing bias in their assessment participants, then minutes later discuss their own gut checking of results to assess the validity of their conclusions. Another participant would assert that QRA provided assurance to budget controlling partners that they were focused on the right cybersecurity projects and investments, but then insisted that executives must sign off on risk appetite to ensure that their choices would not find them legally culpable in case of an unfortunate cybersecurity incident (presumably deemed unlikely and too expensive to mitigate in the QRA). Throughout the process of interviewing people, this cognitive dissonance appeared to be a requirement not just for using QRAs, but for working in cybersecurity in the first place. CISOs spoke of limited budgets and being over-matched by cyber criminals, and how making tough choices with limited resources was a part of their job. There was nuance in their expectations of QRA to help them in this fight: they believed it presented a level of objectivity to those who controlled their budgets that was necessary, even if they understood that it was not objectively true in predicting the future.

## 5.2 CISOs vs. Risk Managers

There was a clear delineation between the CISOs and the risk managers that were interviewed and their expectations of the QRAs that they produced. CISOs had much more of a focus on the pragmatic use of a QRA to achieve an end. In contrast, risk managers had higher expectations of the objective truth of their QRAs, feeling that they should approach an accurate, predictive model. One reason for this might be how each role interacted with QRAs. CISOs are generally the roles who directly advocate for budget with their finance partners and the ones interviewed were using QRAs created by others (risk managers or external consultants). Risk managers, on the other hand, are the creators of QRAs and were observed expressing consistent desire to make them as objective as possible. As the creators of QRAs, risk managers might feel a sense of ownership over the contents of the document that CISOs are more abstracted away from.

## 5.3 The Alchemy of QRA

*"If we imagine the future in terms of probabilities, then risks look safe"* (Clarke, 2005, p. 42)

Clarke's warning is that the expression of risks as probabilities, quantified and given the appearance of objectivity, allows risk to be written off as safe, or in the case of cybersecurity: controlled. As a tool to advocate for budget and the acceptance of risk appetite, QRA operates as an enabling device, convincing the organization that approved projects can control against the complex, uncertain and ambiguous world of cybersecurity by following an objectively laid path of risks enumerated as probabilities. While QRA could potentially be viewed by an organization as a collection of expert information in its systems, once the information is gathered from SMEs and transformed into a QRA, the risk in question is legitimized, and any systems expertise communicated is obfuscated by numbers (Clarke, 1999).

It is notable that the experts interviewed expressed an understanding that QRAs were more subjective than objective. This known subjectivity by experts, but perceived objectivity by consumers could allow for QRA in cybersecurity to fall into organizational miscommunication issues mentioned by Vaughan (1999) and Hutchinson et al. (2018) in the literature review. In a situation where the nuance and lack of objective truth of the numbers is understood by practitioners but not by consumers of the fantasy document, organizations could be set up for dramatic failures that impact their customers and employees. This is a potential side effect of the alchemy of QRAs as enabling devices in cybersecurity: they turn risks, combined with systems knowledge into solid investments, controlled for by seemingly objective expressions of expertise.

## 5.4 Limitations and Opportunities for Future Work

As research grounded in a constructivist viewpoint, this thesis makes no claim to generalization of themes beyond the subjects interviewed. Despite the variety of industries represented among interviewees, it is possible that QRA for cybersecurity is used more in specific types of industries that might be more receptive to numeric and probabilistic analysis. Expanding this research to more participants across industries and potentially gaining insight in a quantitative way as to whether some industries are more receptive to QRA's use in cybersecurity would gather more valuable information that could contribute to answering the research question.

It is clear from this research that QRAs provide value to people who use them beyond any belief in their objective truth. Whether alternative tools (besides the qualitative risk matrix) that present risk and systems information could provide the same or better value would be a fascinating subject for further research. Many scholars advocate for investigations utilizing local rationality and human factors elements, especially in complex socio-technical systems, to help organizations understand more about their day-to-day operations and levels of risk (Rasmussen, 1997; Stoop & Dekker, 2012; Melsek, 2021; Weinger & Slagle, 2002). It might be valuable to compare the usefulness of QRA and these types of investigations, or even examine how they might work together to deepen knowledge of systems and risks for many partners, including those who control budgets.

Studying how a single QRA is perceived across its lifecycle by its creator or advocate (a risk manager or CISO), and then by its consumer (an executive budgetary decision maker or a practitioner in an organization) would produce profound insight into the potential for QRA to be reinterpreted or misinterpreted in a cybersecurity setting. Exploring the dichotomy between pragmatists and true believers who use QRA in cybersecurity would also be enlightening: why do some have an expectation that the model should reflect objective truth, and why do others seem to be content to let it serve a purpose despite its inability to accurately predict events and their costs?

It is difficult to recommend the use of QRAs in cybersecurity after assessing them. There could be many pitfalls to using them in an organization, especially when the consumers of QRAs do not understand their subjective and nuanced inputs. They take a large number of resources to create, and their output is lacking in objective truth. They achieve significant value for their users, though: a funded and prioritized cybersecurity program, in the face of budgetary pressures from all other parts of the business. Until there are better, viable alternative tools for practitioners to use to accomplish these goals, I imagine we will see the use of QRA for some time in the cybersecurity world.

## 6. Conclusion

This thesis set out to understand why QRA is used in a cybersecurity setting. The research findings suggest that QRA functions as an enabling device in cybersecurity by justifying budgets and advocating for investments in the organization. Increasing learning in the organization via QRA was not found to be a reason for its use. There was a difference in the perception of QRA between those who created them (risk managers), and those who interacted with them in leadership (CISOs). Risk managers desired QRAs to reflect objective truth. CISOs, on the other hand, were more pragmatic about its value, and did not have an expectation that it would reflect objective truth. More research is needed to examine the full lifecycle of a QRA and how each participant within the QRA process, including its consumers, perceive its value. Future studies may want to also examine alternative tools to QRA and compare their success in advocating for cybersecurity budgets.

## References

Allspaw, J. (2015). *Trade-offs Under Pressure: Heuristics and Observations of Teams Resolving Internet Service Outages* (A. Smoker & J. Bergström (eds.)) [MSc Human Factors and Systems Safety]. Lund University.

Barnum, T. (2021). *The Cybersecurity Manager's Guide.* O'Reilly.

Bedford, T., & Cooke, R. (2001). Probabilistic risk analysis: Foundations and methods. Cambridge University Press.

Box, G. E. P., & Draper, N. R. (1987). *Empirical Model-Building and Response Surfaces.* John Wiley & Sons.

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101.

Butler, S., Gamalielsson, J., Lundell, B., Brax, C., Mattsson, A., Gustavsson, T., Feist, J., Kvarnström, B., & Lönroth, E. (2022). On business adoption and use of reproducible builds for open and closed source software. *Software Quality Journal.*

Chaput, B., Peters, C. I., & Polsinelli, P. C. (2021). The Legal Liabilities of Enterprise Cyber Risk Management. *American Health Law Asssociation.*

Clarke, L. (1999). *Mission Improbable.* University of Chicago Press.

Clarke, L., & Perrow, C. (1996). Prosaic Organizational Failure. *The American Behavioral Scientist*, 39(8), 1040–1056.

Columbus, L. (2023, February 16). *Benchmarking your cybersecurity budget in 2023*. Venture Beat. https://venturebeat.com/security/benchmarking-your-cybersecurity-budget-in-2023/

Cook, R. I. (2020). Above the Line, Below the Line: The resilience of Internet-facing systems relies on what is below the line of representation. *Queue*, 17(6), 41–51.

Cook, R. I., Woods, D. D., & Miller, C. A. (1998). *A Tale of Two Stories: Contrasting Views of Patient Safety*.

Creswell, J. W. (2009). *Research Design, Third Edition*. Sage.

Crotty, M. (1998). *The Foundations of Social Research: Meaning and Perspective in the Research Process.* Sage.

Dekker, S., Cilliers, P., & Hofmeyr, J.-H. (2011). The complexity of failure: Implications of complexity theory for safety investigations. *Safety Science*, 49(6), 939–945.

DMAIB. (2017). *Marine Accident Report on Mærsk Battler's Loss of Tow on 21 and 22 December 2016*. Danish Maritime Accident Investigation Board.

Downer, J. (2013). Disowning Fukushima: Managing the credibility of nuclear reliability assessment in the wake of disaster. *Regulation & Governance.*

Downer, J., & Ramana, M. V. (2020). Empires built on sand: On the fundamental implausibility of reactor safety assessments and the implications for nuclear regulation. *Regulation and Governance.*

*Framework for Improving Critical Infrastructure Cybersecurity* (No. 1.1). (2018). National Institute of Standards and Technology.

Frenkel, I. B., Karagrigoriou, A., Lisnianski, A., & Kleyner, A. (Eds.). (2014). *Applied Reliability Engineering and Risk Analysis*. Wiley.

Freund, J. (2014). *Measuring and managing information risk: A FAIR approach*. Butterworth-Heinemann. https://doi.org/10.1016/c2013-0-09966-5

Hansson, S. O., & Aven, T. (2014). Is risk analysis scientific? *Risk Analysis: An Official Publication of the Society for Risk Analysis*, 34(7), 1173–1183.

Hubbard, D. W. (2020). *The Failure of Risk Management: Why It's Broken and How to Fix It.* Wiley.

Hubbard, D. W., & Seiersen, R. (2016). *How to measure anything in cybersecurity risk.* John Wiley & Sons. https://doi.org/10.1002/9781119162315

Hutchinson, B., Dekker, S., & Rae, A. (2022). Writing plans instead of eliminating risks: How can written safety artefacts reduce safety? *Safety Science*, 151, 105738.

Hutchinson, B., Dekker, S., & Rae, A. J. (2018, May 28). Fantasy planning: the gap between systems of safety and safety of systems. *Australian System Safety Conference 2018.*

International Atomic Energy Agency. (2002). *Procedures for conducting probabilistic safety assessment for non-reactor nuclear facilities* (IAEA-TECDOC-1267 ). IAEA.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment Under Uncertainty: Heuristics And Biases.* Cambridge University Press.

Kahneman, D., & Tversky, A. (1973). On the Psychology of Prediction. *Psychological Review,* 80(4), 237–251.

Keller, W., & Modarres, M. (2005). A historical overview of probabilistic risk assessment development and its use in the nuclear power industry: a tribute to the late Professor Norman Carl Rasmussen. *Reliability Engineering & System Safety*, 89(3), 271–285.

Kumamoto, H. (2007). *Satisfying safety goals by probabilistic risk assessment* (2007th ed.) Springer. https://doi.org/10.1007/978-1-84628-682-7

Leveson, N. (2019). *Improving the Standard Risk Matrix: Part 1.*

Liu, Y., Ma, W., Guo, X., Lin, X., Wu, C., & Zhu, T. (2021). Impacts of Color Coding on Programming Learning in Multimedia Learning: Moving Toward a Multimodal Methodology. *Frontiers in Psychology, 12*, 773328. https://doi.org/10.3389/fpsyg.2021.773328

Lund University. (2023, March 31). *Research ethics and animal testing ethics.*
http://www.staff.lu.se/research-and-education-0/research-support-0/research-ethics-and-animal-testing-ethics

Melsek, C. (2021). *Local Rationality and Frontline Child Welfare Workers' Decision Making* [MSc of Human Factors and Systems Safety]. Lund University.

Mosleh, A. (2014). PRA: A Perspective on Strengths, Current Limitations, and Possible Improvements. *Nuclear Engineering and Technology*, 46(1), 1–10.

Munroe, R. (n.d.). *Dependency*. https://xkcd.com/2347/

Parry, G. W. (1996). The characterization of uncertainty in Probabilistic Risk Assessments of complex systems. In *Reliability Engineering & System Safety* (Vol. 54, Issues 2-3, pp. 119–126). https://doi.org/10.1016/s0951-8320(96)00069-5

Peace, C. (2017). The risk matrix: uncertain results? *Policy and Practice in Health and Safety*, 15(2), 131–144. https://doi.org/10.1080/14773996.2017.1348571

Perrow, C. (1999). *Normal Accidents: Living with High Risk Technologies*. Princeton University Press.

Perrow, C. (2004). A Personal Note on Normal Accidents. *Organization & Environment*, 17(1), 9–14. https://doi.org/10.1177/1086026603262028

Perrow, C. (2011). Fukushima and the inevitability of accidents. *The Bulletin of the Atomic Scientists*, 67(6), 44–52. https://doi.org/10.1177/0096340211426395

Rae, A., Mc Dermid, J., & Alexander, R. (2012). *The Science and Superstition of Quantitative Risk Assessment.* https://www.researchgate.net/publication/256247163

Rae, A., Provan, D., Aboelssaad, H., & Alexander, R. (2020). A manifesto for Reality-based Safety Science. *Safety Science, 126,* 104654. https://doi.org/10.1016/j.ssci.2020.104654

Rao, A., Carreón, N., Lysecky, R., & Rozenblit, J. (2018). Probabilistic Threat Detection for Risk Management in Cyber-physical Medical Systems. *IEEE Software*, 35(1), 38–43.

*Reactor Safety Study: An Assessment of Risks in US Commercial Nuclear Power Plants* (No. WASH-1400). (1975). US Nuclear Regulatory Commission.

Rip, A. (1986). The Mutual Dependence of Risk Research and Political Context. Science & Technology Studies, 4(3/4), 3–15. http://www.jstor.org/stable/690407

Sagan, S. D. (1993). *The Limits of Safety*. Princeton University Press.

Shayo, C., & Lin, F. (n.d.). An Exploration of the Evolving Reporting Organizational Structure for the Chief Information Security Officer (CISO) Function.

Spanos, G., & Angelis, L. (2016). The impact of information security events to the stock market: A systematic literature review. *Computers & Security,* 58, 216–229. https://doi.org/10.1016/j.cose.2015.12.006

Stoop, J., & Dekker, S. (2012). Are safety investigations pro-active? *Safety Science*, 50(6), 1422–1430. https://doi.org/10.1016/j.ssci.2011.03.004

Sutton, D. (2017). *Cyber security: A practitioner's guide*. The Chartered Institute for IT.

Tapas, N., Longo, F., Merlino, G., & Puliafito, A. (2019). Transparent, Provenance-assured, and Secure Software-as-a-Service. *IEEE 18th International Symposium on Network Computing and Applications (NCA)*, 1–8. https://doi.org/10.1109/NCA.2019.8935014

Thomas, D. R. (2006). A general inductive approach for analyzing qualitative evaluation data. *The American Journal of Evaluation*, 27(2), 237–246. https://doi.org/10.1177/1098214005283748

USNRC. (2020, July 7). *Probabilistic Risk Assessment (PRA).* United States Nuclear Regulatory Commission. https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/probabilistic-risk-asses.html

Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press.

Weinger, M. B., & Slagle, J. (2002). Human factors research in anesthesia patient safety. *Proceedings / AMIA... Annual Symposium. AMIA Symposium*, 756–760.

Williams, P. (2007). Executive and board roles in information security. *Network Security*, 2007(8), 11–14. https://doi.org/10.1016/S1353-4858(07)70073-9

Winkler, R. L. (1996). Uncertainty in probabilistic risk assessment. In *Reliability Engineering & System Safety* (Vol. 54, Issues 2-3, pp. 127–132). https://doi.org/10.1016/s0951-8320(96)00070-1

Woods, D. D., & Cook, R. I. (2002). Nine steps to move forward from error. *Cognition, Technology & Work* , 4(2), 137–144. https://doi.org/10.1007/s101110200012

Woods, D. D., & Cook, R. I. (2006). Distancing Through Differencing: An Obstacle to Organizational Learning Following Accidents. In E. Hollnagel, D. D. Woods, & N. Leveson (Eds.), *Resilience Engineering: Concepts and Precepts* (pp. 329–338). Ashgate.

Woods, D. W. (2022-2023). On Being Prepared to be Surprised: 20 Key Insights from David Woods. *Hindsight*, 34, 11–13.

Wynne, B. (1988). Unruly Technology: Practical Rules, Impractical Discourses and Public Understanding. *Social Studies of Science*, 18, 147–167.

## Appendix A: Research Participant Consent Form

**Qualitative study on Quantitative Risk Analysis in Cybersecurity**

**Consent to take part in research**

**I……………………………………… voluntarily agree to participate in this research study.**

**I understand that even if I agree to participate now, I can withdraw at any time or refuse to answer any question without any consequences of any kind.**

**I have had the purpose and nature of the study explained to me in writing and I have had the opportunity to ask questions about the study.**

**I understand that I will not benefit directly from participating in this research.**

**____ I agree to my interview being audio-recorded.**

**____ If no audio recording consent is given, I agree to have notes taken during my interview.**

**I understand that all information I provide for this study will be treated confidentially.**

**I understand that in any report on the results of this research my identity will remain anonymous. This will be done by changing my name and disguising any details of my interview which may reveal my identity or the identity of people I speak about.**

**I understand that disguised extracts from my interview may be quoted in a MSc Thesis published by Lund University, and possible journal and conference papers that extend this thesis work.**

**If I agree to audio recording, I understand that any original audio recordings will be retained for up to two weeks in digital storage, to allow for transcription and encoding. Once transcription is complete, and transcripts are anonymized, original recordings will be deleted from cloud storage.**

**I understand that this consent form with identifying data will be retained for up to 2 years from granting of MSc (~June 2026)**

**I understand that I am free to contact any of the people involved in the research to seek further clarification and information.**

**Researcher: Colette Alexander, MSc Human Factors and Systems Safety, Lund University. Email: colettealexander@gmail.com**

**Thesis supervisor: Roel van Winsen, PhD**

**Email:  roel.van-winsen@lucram.lu.se**

**------------------------------------------ Signature of participant ---------------- Date**

**I believe the participant is giving informed consent to participate in this study**

**--------------------------- Signature of researcher ---------------------- Date**