



**University of
Zurich^{UZH}**

Real Cyber Value at Risk: An Approach to Estimate Economic Impacts of Cyberattacks on Businesses

*Fabian Künzler
Zürich, Switzerland
Student ID: 15-941-842*

Supervisor: Dr. Muriel Franco, Chao Feng
Date of Submission: January 18, 2023

Zusammenfassung

Um in der heutigen digitalisierten Wirtschaft Prozesse effizient zu gestalten und Dienstleistungen an die Kunden zu bringen, sind viele Unternehmen auf Computerprogramme angewiesen. Jedoch erhöhen die digitalen Werkzeuge nicht nur die Geschäftsmöglichkeiten, sondern auch das Risiko Opfer von Cyberangriffen zu werden. Um dieses digitale Risiko zu managen, existieren verschiedenste Ansätze in der akademischen Literatur als auch in der Beratungsindustrie. Allerdings beinhalten die meisten dieser Lösungen keine individualisierten, auf empirischen Daten beruhenden, quantitativen Angriffskostenschätzungen. Insbesondere für Kleine und Mittlere Unternehmen (KMU) stellt eine Kostenschätzung, aufgrund begrenzter Ressourcen und mangelnder IT-Kenntnisse, eine Herausforderung dar. Die vorliegende Arbeit schließt diese Lücke in der aktuellen Literatur, indem sie den neuartigen Real Cyber Value at Risk (RCVaR) präsentiert. Bestehend aus drei Komponenten, liefert der RCVaR unternehmensspezifische, monetäre Kosten- und Risikovorhersagen für ein Jahr. Die numerische Angabe von Risiko und Kosten lässt eine individuelle Interpretation nach eigenen Risikopräferenzen zu und erlaubt gleichzeitig einen bereichsübergreifenden Risikovergleich. Die Auswertung der Kostenvorhersagen auf der Grundlage von zuvor “ungesehenen” realen Vorfällen zeigt, dass der RCVaR einen Absoluten Prozentualen Fehler (APE) von 2 % erreicht. Weiter beweist die Auswertung, dass das Modell reale, quantitative Kostenmuster von Angriffen widerspiegelt. Um die Risikokomponente des RCVaRs abzubilden, wurde der neuartige Cyber Value at Risk (CVaR) in das Modell integriert. Im Gegensatz zu früheren Forschungsansätzen wird der CVaR nicht durch Monte-Carlo-Simulationen berechnet, sondern mit tatsächlichen historischen Daten. Die Risiko- als auch die Kostenschätzung des RCVaR sind zudem auf die Bedürfnisse von KMU zugeschnitten und über eine Webanwendung leicht zugänglich. Der letzte Beitrag dieser Arbeit adressiert, mit der Einführung einer Federated-Learning (FL)-Methode, den frappanten Mangel an Kostendaten im Bereich der Cybersicherheitsökonomie. Ein Vergleich der Performance-Resultate von verschiedenen FL-Modellen mit dem Output traditioneller, zentralisierter, neuronaler Netzwerke zeigt, dass FL erfolgreich Kostenvorhersage-Funktionen erlernen kann. Folglich stellt Federated Learning eine praktikable Lösung für das Problem der Datenknappheit dar. Zusammenfassend lässt sich sagen, dass der Real Cyber Value at Risk einen neuartigen und kosteneffizienten Ansatz bietet, um quantitative Kosten- und Risikomaße für den Budgetplanungsprozess zu erstellen.

Abstract

To compete in today's digitized economy, companies rely on computer programs to manage processes efficiently and bring their services directly to customers. However, these tools increase not only business opportunities but also the risk of falling victim to cyber attacks. Consulting firms and academic literature provide several approaches to manage this risk exposure. Nonetheless, most solutions fail to provide individualized, quantitative attack cost estimates based on real-world empirical data. Especially Small and Middle-Sized Enterprises (SME) struggle to quantify their attack exposure due to limited resources and a lack of IT knowledge. This thesis addresses this gap in the current literature by proposing the novel Real Cyber Value at Risk (RCVaR) framework. Consisting of three components, the RCVaR provides a monetary, annualized cost and risk prediction for an individual firm. Thus, addressing the issue of individual risk perception and allowing cross-domain risk comparisons. Evaluating the cost predictions on previously "unseen" data from real-world incidents shows that the RCVaR achieves an Absolute Percentage Error (APE) of 2%. The evaluation further proves that the model reflects quantitative real-world attack cost behavior. To portray the risk component of the RCVaR, the newly proposed Cyber Value at Risk (CVaR) is integrated into the model. In contrast to previous research, the CVaR is not computed with Monte Carlo simulations but on the basis of actual historical quantitative data. Both, cost and risk predictions, are tailored towards SMEs and are easily accessible over a web application. The last contribution of this thesis is a Federated Learning (FL) methodology to address the prevalent lack of real-world cost incident data in cyber security economics. Comparing the performance of different FL models against traditional centralized networks suggests that the process can successfully learn cost prediction functions. Consequently, Federated Learning presents a viable solution to the data scarcity issue. In conclusion, the Real Cyber Value at Risk provides a novel and cost-effective approach to obtain quantitative cost and risk measures that integrate seamlessly into the company's overall budget planning process.

Keywords - Cyber Security, Cyber Attack Cost, Cyber Risk, Cyber Value at Risk, Federated Learning

Acknowledgments

First of all, I would like to thank Dr. Muriel Franco for the great collaboration and his feedback which contributed immensely to the quality of this work. I am also glad for the Communication Systems Research Group (CSG) and its head Prof. Dr. Burkhard Stiller, for providing me with such an exciting topic.

Besides the persons mentioned above, I would also like to thank Dennis Arend and Simon Künzler for proofreading this thesis and correcting the few spelling mistakes.

Last but not least, I would like to thank all others who supported me while writing this thesis with their encouragement and advice.

Contents

Zusammenfassung	i
Abstract	iii
Acknowledgments	v
1 Introduction	1
1.1 Description of Work	2
1.2 Thesis Outline	3
2 Background	5
2.1 Attack Costs	5
2.2 Cyber Value at Risk (CVaR)	6
2.3 Machine Learning	8
3 Related Work	11
3.1 Economic Approaches for Cyber Security	11
3.2 Industry Reports on Cyber Attack Costs	14
3.3 AI-Based Approaches for Risk Assessment	17
3.4 Discussion	19
3.4.1 Cost estimation	19
3.4.2 Risk Estimation	20
3.4.3 ML and FL for Information Sharing	22

4	The Real Cyber Value at Risk Model	23
4.1	Data Gathering	24
4.1.1	Primary Data Source	24
4.1.2	Data Extraction	26
4.2	Cost Estimation	29
4.2.1	Size Scaler	29
4.2.2	Time Scaler	31
4.2.3	Factor Selection	35
4.2.4	Factor Scaler	38
4.2.5	Complete Model	44
4.3	Risk Measure	45
4.3.1	Distribution of Cost	45
4.3.2	Discussion on RCVaR Application	49
4.4	Model Development	52
4.4.1	Data Generation	52
4.4.2	Data Preprocessing	56
4.4.3	Neural Network with FL	57
4.5	Web-solution	61
4.5.1	Architecture	61
4.5.2	User Interface	62
5	Evaluation	67
5.1	Cost Estimator	67
5.1.1	Qualitative Evaluation	67
5.1.2	Quantitative Evaluation	70
5.2	Risk Measure	72
5.3	Federated Learning	74
5.3.1	Quantitative Evaluation of the Neural Network	74

<i>CONTENTS</i>	ix
5.3.2 Comparison of Different Data Splits	76
5.3.3 Comparison to Centralized Models	79
5.4 Limitations and Discussion	82
6 Summary and Conclusions	85
6.1 Future Work	87
Bibliography	87
Abbreviations	103
List of Figures	104
List of Tables	106
A Complete Evaluation Data	109
B Complete Limitation List	115
C Documentation on Website	117
D Installation Guidelines	123
E GitLab	127

Chapter 1

Introduction

As the world becomes increasingly digitized, companies are adopting new digital tools to keep up with the rapid pace of change. The COVID-19 crisis has accelerated this trend [48], with 70% of Small and Middle-Sized Enterprises (SME) reporting to have intensified their use of digital technologies [111]. However, digitization comes with challenges, such as migrations of traditional services, training of employees, and exposure to cyber attacks. Today, cyber security is a top concern for C-suite executives. A survey conducted in 2020 by PricewaterhouseCoopers (PWC) [131] found that 49% of CEOs were highly concerned about cyber attacks negatively influencing their business in the following year. This result is astonishing, especially considering that cyber risk outranked the second-placed health risk during a global pandemic [131]. The fear of falling victim is justified, as data from the Federal Bureau of Investigation (FBI) [43] and Accenture [3] shows that criminal cyber activity has been steadily increasing over the past five years.

SMEs, in particular, face difficulties in quantifying the individual risk and impact of a cyber attack. These companies often lack in-house cyber security personnel and financial resources. Due to this absence of expertise, SMEs frequently struggle to understand the current plethora of cyber risk assessment frameworks. Furthermore, these frameworks tend to focus more on the threat environment than on the economic impact of cyber attacks [133, 47]. Leaving out economic indicators during the analysis can lead to inefficient protection measures for businesses that must consider a trade-off between security, risk, and cost [135]. To optimize between these factors in the risk assessment process, Artificial Intelligence (AI) has been employed successfully. Results from studies in the financial [50], engineering [59], and cyber security [77, 91] fields present promising evidence for the use of AI.

Nevertheless, the scientific landscape for approaches that explore Artificial Intelligence and, more specifically, Machine Learning (ML) to investigate the economy of cyber security is scarce. This leaves room for novel methodologies that support SMEs to elicit their economic costs and risks due to cyber attacks.

Such new approaches are needed as SMEs are still unprepared to face criminals in cyberspace. A survey by European Network and Information Security Agency (ENISA) [40] uncovered that most SMEs in the European Union (EU) use information technologies for communication and bank transactions, with 80% claiming to process critical information

in their systems. This means that there is a lot at stake for Small and Middle-Sized Enterprises. Despite the high stakes for these companies, only 70% of SMEs have basic cyber security products deployed [40]. Additionally, SMEs are of utmost importance on a macroeconomic level, as they make up 99% of all companies operating in the European Union (EU), according to the ENISA [40]. A similar situation exists in Switzerland, where SMEs represent 95% of the economy [145]. With their jobs and economic output, SMEs are undoubtedly crucial to European societies. Therefore, cyber risk assessment of SMEs is a priority for many countries and even a matter of national security [156, 112]. Consequently, there is a need for cyber risk assessment solutions, which can be operated without specific knowledge of cyber security concepts and produce easily interpretable results.

1.1 Description of Work

This Master thesis focuses on the research, design, development, and evaluation of a new ML-supported approach for estimating the economic impacts of cyber attacks. The approach is specifically designed for Small and Middle-Sized Enterprises and features an intuitively understandable output metric.

To achieve this, the business characteristics and their relationship with the costs of attacks are thoroughly analyzed and described. Based on the analysis, two different models are developed to estimate risk and expected cost. The risk is estimated by applying the Cyber Value at Risk (CVaR) measure to the cost distribution, while costs are predicted using a model that reflects current quantitative cyber loss behavior in related works. Given these two models, a semi-synthetic dataset of hypothetical firms and their associated costs is generated. A ML model is then trained and evaluated on the dataset. Privacy concerns regarding the sharing of attack information are addressed using Federated Learning (FL), which means only the trained model is shared among participants, not the data itself [95]. The risk and the cost estimating approach, including the ML solution, are summarized in the Real Cyber Value at Risk (RCVaR) model, which can be accessed through a web application.

The proposed RCVaR has several benefits over existing risk assessment frameworks in the cyberspace field. First, the novel approach focuses heavily on short and long-term economic impact estimation, while current frameworks only touch very shallowly on costs occurring during or after an attack. Leaving a recognizable gap in the literature, as different researchers [133, 47] point out. Secondly, the technical solution, which evolved from this thesis' theoretical approach, can be operated without any knowledge of cyber security or even informatics concepts, consequently *(i)* lowering the need for cyber security specialists. It also *(ii)* enables non-cyber technicians to engage with cyber security and facilitates *(iii)* a shared understanding of cyber topics in budget discussions. Therefore, the developed approach solves three of seven cyber security challenges identified by ENISA [40].

Additionally, the cost and risk estimation are both based on quantitative data from real-world surveys rather than theoretical concepts. The numerical risk value also addresses

the challenge of individual risk perception, which is a drawback of current risk-ranking approaches [49, 83]. Both measures cover a wide range of attack vectors and are presented on an annualized basis to facilitate integration into annual business planning.

Last but not least, the developed dataset provides a starting point for future research in cyber economics, which chronically suffers from data scarcity as the Congressional Research Service (CRS) points out. As identified by the CRS, one reason for this is companies' incentive not to report incidents [17, 101]. To tackle this problem at its root, Federated Learning techniques are employed to allow an anonymous sharing of knowledge from attacks among SMEs and enable an evolving model.

1.2 Thesis Outline

This Master thesis begins by establishing a knowledge foundation in Chapter 2, with a particular focus on the concept of Value at Risk (VaR). In the following Chapter 3, the relevant academic literature and industry reports that focus on cyber attack costs are reviewed. The chapter also discusses risk and cost estimation in the literature and examines the ability of Federated Learning (FL) to learn risk and cost estimation functions.

The main chapter of this thesis (*cf.* Section 4) addresses three main goals:

First, Section 4.1 presents the primary cost data source before explaining the extraction process. Next, Chapter 4.2 illustrates the development process of the cost prediction model. This process includes scaling cost over time (*cf.* Section 4.2.2) and specific company sizes (*cf.* Section 4.2.1). Heavy emphasis also lies on the customization of the cost through individual company characteristics (*cf.* Section 4.2.3 and 4.2.4).

Secondly, to address the lack of individually interpretable risk measures, a Value at Risk based measure is developed in Section 4.3.

The final contribution of this thesis is presented in Section 4.4, where a Federated Learning (FL) model (*cf.* Section 4.4.3) is trained on a generated dataset (*cf.* Section 4.4.1).

The thesis closes by evaluating all three solutions in Chapter 5. First, the cost estimation model is tested against unseen data. Then, Section 5.2 compares the risk measure's distribution against the results from related work. In the end, the Federated Learning model's performance (*cf.* Section 5.3) is evaluated before stating the limitations (*cf.* Section 5.4) of the presented approach. Finally, the thesis concludes with an outlook on possible future work (*cf.* Section 6.1).

Chapter 2

Background

This chapter builds the knowledge foundation for the concepts used later in the main Chapter 4. First, the scope of the cyber attack costs is defined. Next, the concept of the Value at Risk (VaR) is introduced with a particular focus on cyber economics. Finally, the theory of Federated Learning (FL) with Deep Learning (DL) models is explained in detail.

2.1 Attack Costs

This thesis focuses on costs associated with cyber attacks, but understanding what contributes to the total costs of an incident can be difficult. Therefore, this chapter provides an overview of the multiple pillars that make up the total cost. In the end, the definition of cyber attack costs used in this thesis is provided.

The two main dimensions by which incident costs are typically classified in the current literature are shown in Table 2.1. The first dimension, used in [51, 62], divides the cost into *Anticipation*, *Consequence* and *Response* costs. These dimensions can be seen in the top row of Table 2.1. The second dimension, often seen in industry reports [2, 3], distributes the costs into categories: *Direct*, *Indirect*, and *Opportunity* costs. This second dimension can be viewed along the y-axis of Table 2.1.

Table 2.1 illustrates exemplary costs for each category commonly used in the related literature [51, 51, 2, 3]. *Direct* costs are those related to immediate cash outflows, such as money stolen by criminals or the loss in revenue. On the other hand, *Indirect* costs refer to losses incurred as a result of additional resources being expended, such as the cost of vetting security software before deployment within a firm. Such a process, for instance, requires an invitation to tender for the contract, which claims time from the management and other company intern entities. Finally, *Opportunity* costs stem from lost business opportunities [2]. For example, if a financial firm experiences an attack, it is likely to lose the trust of its customers, leading to potential future loss of business. These costs are especially difficult to measure, as they can have a long time horizon.

The first cost contributor along the x-axis, called *Anticipation*, covers all costs prior to an attack. These include spending positions on additional security measures and missed opportunities due to these actions. The next category, *Consequence*, is typically related to the immediate aftermath of an attack. The final category covers costs occurring post-incident. Depending on the work, this category includes the public’s loss, such as the investigation of law enforcement authorities and the justice system, or it focuses on the firms’ post-attack response. A combination of these two views is presented in Table 2.1 to provide a comprehensive picture. As a result, costs in the *Response* column include the investigation efforts of law enforcement as well as the possible business disruption of the firm due to these investigations.

	Anticipation	Consequence	Response
Direct	Cyber Insurance Premium Price for Security Measures	Revenue Loss Theft of Funds	Lawyer Fees Regulation Fines
Indirect	Vetting of Security Software Employer Cyber Training	Repair of Infrastructure Monitoring of the Attack	Investigation Documentation
Opportunity	Product R&D Investment in New Project	Intellectual Property Loss Reputation Loss	Business Disruption Higher Cost of Capital

Table 2.1: Common Cyber Cost Dimensions in Literature

Since this work’s model is based on real-world numbers from industry reports, their definition of costs has been adopted. Therefore, the model output covers the firm specific *Direct*, *Indirect*, and *Opportunity* costs of the two last columns in Table 2.1. Furthermore, it is noteworthy that the two steps, *Consequence* and *Response*, can be divided further into subcategories, which are discussed in detail in Section 4.1.1. It is also important to remember that the costs are not solely focused on a single attack vector but rather cover all types of cyber attacks included in the survey reports. Since assigning costs to attack vectors is a complex endeavor, the reports do not provide an exhaustive list of covered attack vectors.

Ultimately, the costs reflect the estimated annualized costs in United States dollars (USD). Conversion of the amount to local currency is subject to currency risk and is outside the scope of this thesis. While analyzing the cost output of the RCVaR, it is also important to understand that not all expenses are covered. More specifically, costs occurring in the *Anticipation* stage, costs due to unsuccessful attacks, expenses related to compliance, and investments in new security measures are not included in the cost prediction of the model.

2.2 Cyber Value at Risk (CVaR)

Cyber attack costs have significant impact on organizations considering that the total loss due to attacks is estimated to be one percent of global GDP according to Allianz, a leader in corporate insurance [5]. Hence, it is not surprising that cyber risk management is crucial for risk officers. In recent years, the scientific community has paid particular

attention to the Cyber Value at Risk (CVaR) measurement. The CVaR is a quantile-based risk measure that originated from the Value at Risk (VaR), which was outlined by Markowitz in 1952 and gained widespread use in finance in the 1990s [61, 119, 37]. VaR has three dimensions: Confidence, also known as probability, worst-case loss, and a time frame [119].




Figure 2.1: Value at Risk Representations Based on [160]

According to the efficient market hypothesis (EMH) [41], the returns of assets in efficient markets follow a random walk, which morphs into a Brownian motion in continuous time [84]. Therefore, it can be argued that returns are normally distributed, as depicted in Figure 2.1. If the distribution in Figure 2.1 represents monthly returns, the monthly return at quantile α can be calculated. The result is the Value at Risk with confidence $1 - \alpha$ over a time frame of one month with a worst-case loss of the return at the respective quantile.

In the year 2015, the cyber resilience initiative of the World Economic Forum (WEF) [159] proposed the idea of using the Cyber Value at Risk as a risk benchmark across different domains and industries [1, 119]. Compared to traditional “scoring” approaches in the field of cyber security (*cf.* Section 3.1), CVaR offers several advantages. First, the complexity of the risk can be represented by a single, individually interpretable number, allowing C-suite executives to scale risk in the cyber domain to their risk appetite [1]. Residual risk, which managers are unwilling to take on, can then be outsourced to cyber insurance companies [119]. Secondly, quantile measurements are well established and therefore allow for comparison of risks across different domains, *e.g.* financial, operational, and cyber risk. Additionally, due to the experiences in the financial sector with the VaR, there exist multiple frameworks and models that build upon the theory of quantile risk measurements. The CVaR, as a quantile risk quantifier, is compatible with these frameworks, allowing a user to conduct further risk analysis.

However, the CVaR model inherits not only VaR’s benefits but also its drawbacks. The most severe limitation imposed is that the CVaR is a backward-looking measure, relying on historical data and assumptions, which do not necessarily hold in the future [1]. For instance, just because an event has never occurred in the past does not mean it has a probability of zero in the future. Another often overlooked drawback is that “Black

Swan” events, which are infrequent events with a massive impact, are not represented well by past data. Meaning fat tails of the distribution, signalized through the kurtosis, are disregarded by the Cyber Value at Risk [161]. A solution to this problem could be the expected shortfall, also known as the conditional VaR, depicted in Figure 2.1. The expected shortfall represents the expected loss under the condition that the α -quantile is pierced; hence it takes losses in the tails into account.

When the World Economic Forum presented the concept of the CVaR, they did not provide a methodology for calculating it. Instead, they identified specific components the model should encompass, such as the *Vulnerability* of assets, *Assets* under threat, and *Profiles of attackers* to which the assets are exposed [1, 159]. In a later report [157], Deloitte calculated the expected loss as well as the CVaR for different sectors in the dutch economy, but with several underlying limitations. For instance, the authors explicitly excluded Small and Middle-Sized Enterprises and certain sectors such as construction or real estate. Their approach uses a combination of real-world data and estimates to determine the threat profiles for information assets. Based on these likelihood estimations and the identified asset values, the expected value and the Cyber Value at Risk for the industry sectors are calculated. Nevertheless, they fail to provide a general model which is applicable to individual organizations. To close this gap, a recent study from the Oxford university [37] developed a model which calculates the CVaR based on four pillars: asset values, harm probability, threat probability, and effectiveness of protection. Once these values are determined, mainly through estimations, a Monte Carlo simulation is run to establish a distribution, and the CVaR is then derived from that distribution. Both Deloitte and Oxford emphasize in their contributions that there is an enormous lack of data to calculate probabilities and distributions [37, 157], resulting in many probability estimates that can be skewed due to, for example, behavioral bias [74]. Because of this complexity, the calculation of the CVaR is costly and complicated, which is why it is not commonly used in the industry [119]. One company that has the CVaR in its product palette is MARSH [93]. They use a similar approach to the Oxford study by collecting relevant data at the client’s location and estimating the missing probabilities before deriving the loss distribution through a Monte Carlo simulation. This process requires a high level of trust from the client, as he must provide MARSH with critical company data.

To summarize, the Cyber Value at Risk is a relatively new measure in risk management. Its complexity and the lack of data hinder it from being used broadly and in a continuous controlling setting [119]. A particular gap in CVaR solutions exists for Small and Middle-Sized Enterprises.

2.3 Machine Learning

Artificial Intelligence (AI) is a field comprising various subfields, one of which is Deep Learning (DL). Deep Learning (DL) has gained attention in the past decade due to its performance accuracy, the availability of data, and the increased computing power [7]. Deep Learning models can have a multitude of architectures but on the smallest level consist of neurons. These are connected to one another through weights and organized into layers. The term Deep Learning refers to the fact that these layers are often stacked on top

of each other to achieve a “deep” model. During the forward propagation of the learning process each neuron’s output is weighted, summed, and passed through an activation function to introduce nonlinearity [7]. The forward pass results in a predicted value, which is compared to the actual value using a loss function. To achieve the smallest difference between the actual and predicted value, the loss function is minimized by updating the weights in the opposite direction of the gradient. This is done by subtracting the gradient from the weights.

Traditionally, Deep Learning models were fed with vast amounts of data to achieve optimal performance. This involved collecting data in a central entity and then training the Deep Learning model on this data. However, this approach can fuel privacy concerns due to multiple reasons. First, the data might be sensitive and, therefore, should not leave the personal device. Secondly, the server, with all data located there, presents a single point of failure for attacks. One data breach might compromise all observations. To address this issue, Google introduced FL in 2016 [99, 64]. Figure 2.2 demonstrates the Federated Learning process using the example of cyber attack costs (*cf.* Section 4.4.3).

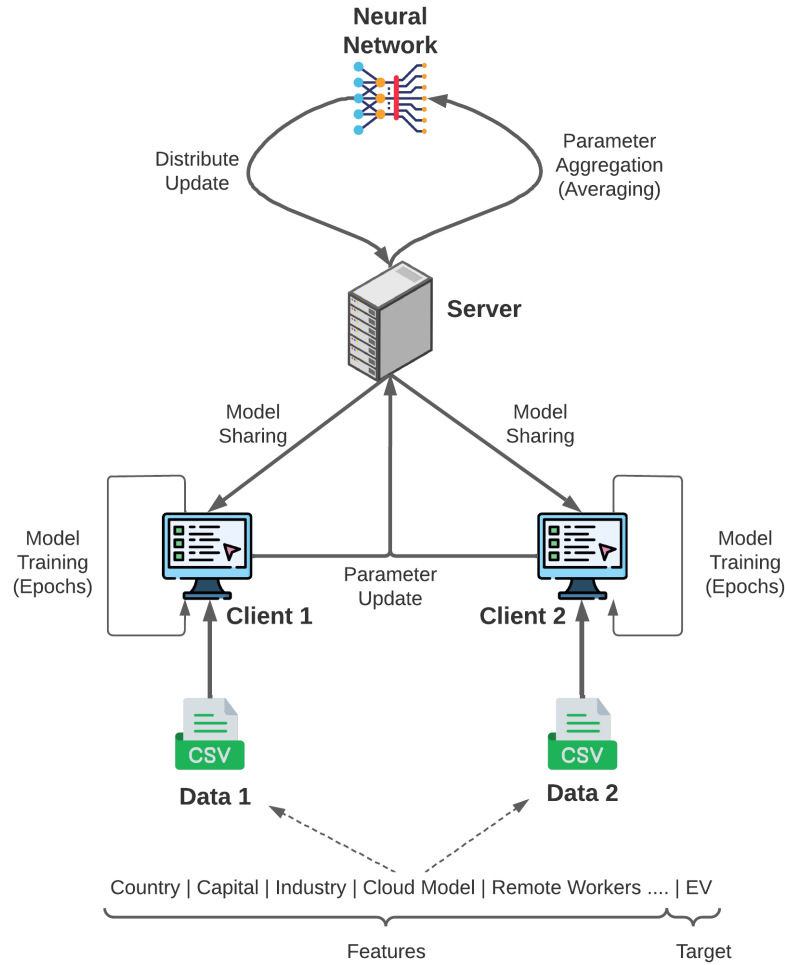


Figure 2.2: General Federated Learning Process

In this thesis, Federated Learning was conducted with only two clients as a proof of concept. Each client has its own dataset of business characteristics, which serve as input features, and expected cost, which serve as the target variable. The process of generating data in this format is described in Section 4.4.1. The data is stored in a predefined format in a CSV document and is kept on the client’s device. Since both clients have the same type of data, the Federated Learning process can be classified as horizontal Federated Learning [64]. The process begins by sharing a pre-trained model (*cf.* Section 4.4.3), which is located on the server, with all clients. They then perform the DL algorithm described above on their local devices for several epochs. After the clients have trained the model, the weights are encrypted and sent back to the server, while the training data remains on the client’s device. On the server, the weights are aggregated by averaging before redistributing the model to the clients. Such an iteration is called “round” in the Federated Learning process. After several rounds, the final model, trained in a decentralized manner, is available on the server [64, 99].

Federated Learning has the ultimate benefit of achieving accurate models while maintaining privacy during the learning process. Additionally, it opens new opportunities in the Internet of Things (IoT) field. However, some drawbacks of this new learning method are higher bandwidth requirements for sharing the model. It is also challenging to ensure compliance with data protection regulations. Specifically, how to delete a client’s contribution upon his withdrawal from the FL process mid-training.

Chapter 3

Related Work

This chapter overviews the related literature on attack cost estimation, Artificial Intelligence for risk assessment, and Federated Learning. First, the research of cost estimating methodologies is split into two chapters, with one section focusing on economic-research papers while the other prioritizes industry reports. Next, AI-based methods' role in risk and cost prediction are presented. Finally, the chapter closes by pointing out the gap in the affiliated literature and how this thesis will address it.

3.1 Economic Approaches for Cyber Security

Current literature provides a flood of risk assessment and cyber security investment frameworks. Two of the most well-known cyber security frameworks are from the National Institute for Standards and Technology (NIST) and the International Organization for Standardization (ISO). The guidelines specified in ISO 27005 [70] focus on information security risk management. Its concepts are based on other reports from the ISO 27000 family (ISO 27001 [68], ISO 27003 [69], etc.). The NIST framework [108] comprises of three subparts: Core, Implementation Tiers, and Profiles. Implementation Tiers represent a methodology to rank a company's security processes, whereas the Profiles represent different states, *e.g.* current-state or target states. The Core consists of different cyber security activities, which help organize the risk assessment. Nevertheless, these frameworks are hard to apply to SMEs due to their high level and complexity. None of these models provide a practical approach to assign a numerical measure to risk. To address this need, the Information Assurance for Small and Medium Enterprises Consortium (IASME), in association with the National Cyber Security Center (NCSC), developed a practical questionnaire [67] for SMEs to evaluate their cyber risk. Another more businesslike methodology is presented by the SecRiskAI solution [49]. Its neural network estimates risk based on business and cyber security characteristics, such as security investments, revenue, or known vulnerabilities. One issue with this model is that it is trained on a dataset that was artificially generated using a theoretical formula with no quantitative evidence to back it up. Furthermore, the formula makes several assumptions that may not hold in reality, such as a linear relationship between input and risk. In the end, the model classifies risk,

similar to the visual tool for investment analysis [66], into one of three abstract categories: Low, Medium and High. This benchmark gives a baseline for decision-making but is still too abstract to be used in a corporate risk analysis.

In comparison to the cyber risk assessment frameworks, far fewer out-of-the-box approaches to estimate the cost of cyber incidents exist. Two frameworks to measure the monetary loss due to cyber attacks are presented in SEconomy [135] and [51]. In the SEconomy paper [135], the overall economic impact is estimated in the fifth step of the framework based on inputs such as threat cost, mitigation cost, and initial security investments. Detailed formulas for each of these inputs are also provided through the framework. Nevertheless, to apply the framework, the investigating entity needs to estimate the threat cost for each role and system, which is a challenge in itself. A similar approach was proposed in [51], where the authors identified cost contributors such as insurance, repair, and intellectual property fraud and classified them among two dimensions: The time the cost occurred and the type of the cost. The ensuing extensive list of cost types aims to determine a company's overall cost by estimating the individual contributors and summing them up. However, numerical values or guidelines for obtaining monetary values are not provided.

Both risk and cost are essential input factors for cyber security planning and investment. CyberTEA [47] is an approach consisting of a methodology, a framework, and a set of solutions for cyber security planning. In the *Decision Layer* of the proposed framework, the costs are estimated before finally recommending a resilience-improving action. The economic impact is calculated using the SECAdvisor proposed in [113]. The SECAdvisor calculates the overall cost by multiplying the monetary value per data record, determined by a survey from IBM, with the number of records. Furthermore, the value of a web server is estimated as the 30-fold of the monthly profit generated over the web store. Compared to the previous frameworks, the SECAdvisor provides a numerical cost number rooted in the real world. A shortfall is that the economic costs are only tailored to a company by altering one parameter: The number of records. A similar approach was followed by Li et al. [83] in developing their visual tool for cyber security/investment planning. They assumed the monetary value of a record, respectively a loss, to be the monetary penalty defined in policy regulations such as the EU's General Data Protection Regulation (GDPR) [38]. A particularly significant part of cyber security planning is deciding on how much a company should invest in cyber security. The most established frameworks for this process are the Return on Security Investment (ROSI) measure [39] and the Gordon-Loeb model [54]. The ROSI illustrates the monetary gain of the security measure relative to the invested capital. To get the nominator in the formula, one has to estimate the annual monetary loss (ALE). Unfortunately, the ROSI framework does not provide ways to estimate it for cyber incidents. The other major approach concerned with cyber security investment is from Gordon and Loeb [54]. The model suggests that a company should only invest 37% of the expected loss in security measures. The expected loss is defined by their paper as the vulnerability multiplied by the potential loss. However, no guidelines are provided on how to estimate these numbers in a corporate environment. Similarly vague on how to estimate costs are the frameworks presented in [18] and [66].

A more detailed look into the cost of cyber crime is provided by Anderson et al. [6], whose paper can be considered the first systematic study of costs. In their work, they distinguish

between *Direct*, *Indirect*, and defense costs as well as between different types of cyber crimes. The heart of the research is a table of collected costs from multiple sources such as CSI, ENISA, and Symantec statistics. But even with extensive research some numbers had to be approximated, with the report stating that 9 of the 28 cyber-crime-type costs include highly uncertain estimates. It is also noteworthy that the authors focused on the combined public and firm costs for the United Kingdom, meaning no numbers regarding individual company losses are displayed.

Despite the difficulty of estimating the cost for corporate entities, some solutions which provide customized numerical estimates exist. Two approaches researched in the scope of this thesis are the ReCIst [57] and the questionnaire of the cyber security osservatorio [26]. The ReCIst is a project that provides a visual tool to investigate cyber security investments by comparing a measure's benefits against its cost. The system calculates costs based on a range of complex inputs, including the impact of an attack on revenue and factors related to the likelihood of the attack. Estimating these inputs accurately, as acknowledged by the authors, can be cumbersome. The Cyber Security Osservatorio [26], on the other hand, uses a questionnaire, which does not require estimates of cyber attack costs¹.

The main problem in researching the cost of cyber incidents is the lack of publicly available data. There are several reasons for this. First, victims are incentivized not to reveal their breaches [101, 17]. Secondly, it is hard to quantify *Indirect* and *Opportunity* costs [17]. To address the scarcity of data, the Department of Homeland Security (DHS) has started an initiative [155] to collect data on cyber security incidents, but the current status of this database is unknown.

An alternative approach to estimate costs uses stock prices. According to the efficient market hypothesis [41], new information is immediately priced upon release. Furthermore, stock prices, in theory, reflect not only the company's current value but also the overall market sentiment of future business development. Therefore, market expectations can be a good proxy for cyber incident costs since they incorporate the market's expectations of *Indirect* and long-term costs. Following this theoretical approach, Cavusoglu et al. [19] discovered that stock prices drop 2.1% within two days after an attack. Tweneboah-Kodua et al. [153] found in their paper that statistically significant price fluctuations of sector-specific stocks can be observed after a cyber incident. Nevertheless, the study also points out that these changes can only be detected in particular industries. The most robust evidence of a price impact due to a cyber attack is found in financial industry stocks.

Throughout the research of related papers, it becomes clear that none of the papers above presents numerical data on cost influential factors. Furthermore, most papers base their cost estimations on something other than real-world data or leave the estimation to the company itself. This thesis addresses this gap and provides a solution based on real-world data to assess and predict a company's incident cost.

¹The website of the Cyber Security Osservatorio [26] was very unreliable during the test run conducted in the scope of this thesis. The questionnaire, did not provide results instead https-errors were thrown.

3.2 Industry Reports on Cyber Attack Costs

As pointed out in the previous section, current academic research struggles to classify the cost or risk numerically. However, corporate risk management requires an understandable numerical metric to estimate the impact of an attack. To address this need, major consulting and technology firms conducted surveys to develop cost estimations on which companies can rely during their cyber security decision-making process. Most of these surveys were conducted with the help of the Ponemon Institute [85], an independent organization specializing in data privacy and information security. It is widely considered the leading institute regarding economic cost estimation following cyber incidents [89]. In the scope of this thesis, all publicly available reports [126, 127, 128, 129, 130, 139, 85] related to attack costs of the past 12 years of the public Ponemon library are investigated. Nevertheless, the main focus lies on IBM’s 2022 [27] and Accenture’s 2017 [2] reports due to their recency and the number of cost-influential factors (*e.g.*, country or organization size) described numerically.

To determine these cost factors, Accenture surveyed 2182 individuals from 254 companies across seven countries². Their research focused on costs related to “cyber crime” incidents. They estimated the average cost per annum for a company of the sample to be \$ 11.7 million in the year 2017. Their sample group only included larger companies from various sectors. As can be observed in Table 3.1, the survey identified the organization size, industry sector, region, year, and security measures as possible influences on the average cost. But Accenture also states that they omitted other essential variables. However, they do not explicitly mention which variables were omitted. Additionally, IBM’s survey confirms some influential factors mentioned by Accenture and adds new variables that impact the cost.

Regarding the influence of operational processes on costs, the IBM work is more detailed. It shows that risk management, remote work, cloud services, and their security, sufficiently staffed teams, as well as breaches in the supply chain, influence the average cost. All costs in the IBM report are expressed as the average per-company breach cost, as opposed to Accenture’s per annum incident cost-unit. A data breach defines an event that puts an individual’s sensitive data at risk. IBM estimated the average per-breach cost for 2022 to be \$ 4.35 million. The estimation is based on their survey of roughly 3600 individuals from 550 enterprises. Their sample included companies from 17 regions³ and 17 sectors. Most sectors are similar in both studies, with IBM having the additional sectors Entertainment, Media, Research, and Pharmaceuticals. Table 3.1 summarizes the numerical cost-influencing factors provided in both reports.

Accenture [2] concludes its 2017 report by identifying a “value gap” between the spending on a security measure and its cost-saving ability. The report locates the highest potential in the usage of Security Intelligence Systems, which collect system data, calculate risk-related measures and display them [144]. IBM, on the contrary, sees remote work and cloud environments as the biggest challenges for the company’s data security. To address

²United States, Germany, Japan, United Kingdom, Australia, France and Italy

³United States, India, United Kingdom, Brazil, Germany, Japan, France, Middle East, South Korea, Australia, Canada, Italy, ASEAN, Latin America, South Africa, Scandinavia, Turkey

these challenges, the report recommends, among others, Identity and Access Management (IAM) products and Data Encryption Technologies. Furthermore, the report recognizes the potential to reduce costs by implementing Security AI and Automation products, Incident Response Plans, and Zero-Trust frameworks [27].

Both studies are very similar regarding their methodology. Both base their results solely on survey information. Additionally, both claim to capture the *Direct*, *Indirect*, and *Opportunity* costs separately. Furthermore, the two investigate only cost attributed to activities in the two categories *Consequence* and *Response* (cf. Section 2.1). These activities are further split into “cost activity centers” and roughly include *Detection*, *Investigation*, *Containment*, *Post Breach Activities*, and *Lost Business Revenue*. Both reports have many similarities due to their Accenture’ and IBM’s partnership with the Ponemon Institute during the development.

In 2019, Accenture partnered again with the Ponemon Institute to release the cybercrime report of 2019 [3]. The research involved extensive data collection, with 2600 individual interviews from 355 companies operating in 11 countries⁴. However, the report presents fewer cost-influencing facts as the previous report from 2017. Overall, the cost contributors country, industry, sector, attack type, and security measures on the costs are identified. Additionally, a composition of costs per attack type into the categories: *Business Disruption*, *Information Loss*, *Revenue Loss*, and *Equipment Damage* are shown. The report also states clearly that *Information Loss*, due to theft or destruction is the most significant contributor to the average annualized per-company cost of \$ 13 million. Accenture further discovered that attackers tend to gain access to IT systems through supply chain partners. This poses new challenges, especially for large, interconnected companies. The report concludes that the most prevalent measure to keep attacks and costs low is actively living a security culture, including having individual accountability and cyber security training. Additionally, the authors recommend investing in information protection, such as Cryptography and Detection Automation Technologies, in order to reduce monetary consequences of attacks.

The Accenture report from 2019 was the last cost of cyber crime survey conducted by Accenture. The successor reports focus on capturing trends in the threat landscape and investigating the spending on different security measures. For instance, Accenture’s 2021 cyber resilience report [10] emphasizes the need to include security considerations in strategic planning by continuously measuring threats and security, particularly concerning cloud-based solutions. Other Accenture reports shed light on federal agencies’ cyber security [42], and business leadership in the current IT security environment [30].

⁴United States, United Kingdom, Germany, Japan, France, Brazil, Canada, Australia, Spain, Italy, Singapore

Table 3.1: Summary and Comparison Between the Two Most Established Cost Reports from Accenture [2] and IBM [27]

Business Characteristic	Specifications of Information form	Accenture (2017) [2]			IBM Report (2022) [27]		
		Data type	2nd variable	Unit	Data type	2nd variable	Unit
Security Measure	Effectness of specific Security Measures	Cat.	Security spending	Rank			
	Cost savings per security measure	Cat.		ann. USD	Cat.	Time comparison (2020-2022)	avg. USD per breach
	Security effectiveness Score (SES)	Cat.	Country	ann. USD			
Time	Increase of avg. cost (2013-2017)	Num.		ann. USD	Num.		avg. USD per breach
Region	Avg. Cost per region	Cat.	Time comparison (2016-2017)	ann. USD	Cat.	Time comparison (2021-2022)	avg. USD per breach
	Costs per attack per country	Cat.	Threats	ann. USD threat cost per country			
Organization size	Influence of Organization size on cost	Cont. & Cat.	Time comparison (2013-2017)	ann. USD			
	% spend on threats for two sizes	Cat.	Threats	% of total cost per threat			
Industry Sector	Cost per Sector	Cat.		ann. USD	Cat.		avg. USD per breach
Operation	General factors influencing cost				Num.		avg. USD per breach
	Risk management influencing cost				Cat.		avg. USD per breach
	Influence of Supply Chain data breach				Cat.	Time to resolve in days	avg. USD per breach
	Cloud model and its security measures				Cat.	Time to resolve in days	avg. USD per breach
	% working of employees working remotely				Num.		avg. USD per breach
	Sufficiently staffed security teams				Cat.		avg. USD per breach
Threats	Attack type experienced	Cat.	Time comparison (2016-2017)	% of sample has experienced it			
	Costs per threat	Cat.		ann. USD	Cat.	Frequency of attack in %	avg. USD per breach
	Avg. time to resolve threats	Cat.	Time comparison (2016-2017)	Days	Cat.	Time comparison (2016-2022) & Cost & attack type	Days & avg.USD per resolve period
Cost composition	Composition of cost	Num.	Time comparison (2015-2017)	% of total cost & category	Cat.	Time comparison (2017-2012)	avg. USD per breach and category

Ann. = Annualized, Avg. = Average, Cat. = Categorical (Presented as Histogram)

Num. = Numerical (Presented as single Numbers) , Cont. = Continuous (Presented as continuous Line Graph)

A report conducted by Kaspersky [72] independently of the Ponemon Institute in 2013 supports the findings of Accenture and IBM. Kaspersky's main goal was to summarize the types of data breaches companies experience and how they protect themselves against them. Additional emphasis lies on the company's perception of the current cyber security environment. The work only presents numerical results for two influencing factors of cost: Region and attack type. The numbers for those factors were collected by interviewing 2895 IT professionals among 24 countries⁵. Based on this inquiry, the average per security breach impact was estimated to be \$ 649'000 for larger enterprises. An essential difference from the other reports is that Kaspersky gives a separate cost approximation for SMEs (\$ 50'000 per incident). It is also the only report researched in this thesis, which looked at the company's incident costs in China and Russia. In the end, Kaspersky identifies two major causes of incidents: Bring-your-own-device policies (BYOD) and internal leaks due to insider attacks or poor training. Solutions to address these problems include employee security training and targeted investments in professional IT security applications.

The last industry report investigated in the scope of this thesis was conducted by Deloitte [102] in 2016. The main focus of the report lies on researching *Indirect* cost of cyber attacks, *e.g.* impact on customer relationships, insurance premium increases, and devaluation of the brand name. Unlike the previously presented reports, Deloitte bases its monetary valuation of these impacts not on surveys but on estimations of its in-house industry experts. They further support their claims with data from the Ponemon Institute [85]. To present their cost estimations, Deloitte demonstrates the cost calculation of a cyber attack through two fictional scenarios: A data breach on a healthcare provider and an attack on a manufacturer. The total scenario costs approximate to \$ 1.679 and \$ 3.258 million. To improve cyber resilience, the report recommends more in depth cyber risk management, which includes scenario stress tests [102]. Furthermore, investment in Detection Technologies and the development of Incident Response Plans can shorten an attack's life cycle and decrease costs.

Other reports which deal with costs of cyber security include the Ponemon reports from 2014 [129], 2016 [130], 2017 [139], as well as Zurich Insurance's 2015 report [89]. The information of these reports is not as recent as the data introduced by Accenture [2, 3, 10, 42, 30] or IBM [27]; however they help gain a picture of cost development over time.

3.3 AI-Based Approaches for Risk Assessment

Due to the lack of cost-related data, using machine learning to estimate the expected costs of cyber attacks is difficult. Therefore, this thesis could not find a ML model in recent literature that predicts cyber incidents' costs. Nevertheless, there are countless examples of machine learning algorithms in risk assessment and cyber threat detection. One solution which uses ML to determine the cyber risk exposure of a company is the SecRiskAI [49]. During the development four machine learning algorithms: Decision Trees, Support

⁵United States, Canada, Mexico, Columbia, Peru, Chile, Brazil, United Kingdom, Germany, France, Spain, Italy, Greece, Czech Republic, Hungary, Saudi Arabia, United Arab Emirates, South Africa, Russia, China, Japan, India, Australia, Kazakhstan.

Vector Machines, K-Nearest Neighbor, and a Multi-layer Perceptron Neural Network were evaluated. The paper concludes that all algorithms achieve an accuracy of over 95%. However, it is worth noting that all the models were trained on synthetic data, so the function learned by the models was also artificially constructed and, therefore, previously known. In their paper, Abishek and Kumar [140] used a similar approach to identify risk factors and apply machine learning models such as Decision Trees or Randomizable Filter Classifier to investigate threats specific to the cloud computing domain.

Another paper by Chih-Hung et al. [63] argues that current cyber risk assessment heavily relies on cyber security specialists and their domain knowledge. To lower this dependency, which is also pointed out by ENISA [40], they propose a fuzzy scoring system that returns a risk ranking. The authors achieved promising results when evaluating their novel system on a real-world cyber risk dataset. Besides a lower dependency on rare cyber security specialists, introducing AI into cyber risk management has additional advantages, as a Deloitte report [31] from 2018 points out. Deloitte analyzes that when Artificial Intelligence is included in the risk management process, the procedure becomes more forward-looking than without AI.

Supplementary to cyber risk management, ML is also used in other sectors to manage different risks. A summary of engineering industries and their respective algorithms can be found in Table 3.2. It is noteworthy that only the most frequent algorithms used in risk assessment, according to a meta-study conducted in 2020 [59], are listed. Besides engineering industries, ML techniques for risk management are also central in other economic sectors, such as the financial industry [79, 50, 97].

Industry	Algorithm	Paper
Construction	ANN	[33]
	SVM	[80]
	DT	[53]
	RF	[152]
Railways	ANN	[44]
	SVM	[81]
	DT	[142]
	RF	[15]
Mining	ANN	[143]
	SVM	[143]
	DT	[143]
	RF	[52]

ANN = Artificial Neural Network, SVM = Support Vector Machine

DT = Decision Tree , RF = Random Forest

Table 3.2: Risk Assessment ML-Algorithms in Different Industries [59]

Since Machine Learning has multiple applications throughout different domains, it is not surprising that Federated Learning (FL) is also commonly used to protect privacy during model training. With its sensitive but volume-rich data, the healthcare sector is destined

for Federated Learning. Brisimi et al. [14] developed a decentralized optimization framework, which allows different data holders to jointly train and predict hospitalizations for cardiac events without sharing their data. Further, research by Kim et al. [75] demonstrates on medical data that Federated Learning is competitive with centralized trained models in terms of accuracy while still adhering to privacy restrictions. Other applications for FL include visual inspection of products in engineering processes [58], credit risk assessment [73], and protection of drones [103]. A more extensive list of Federated Learning applications can be found in the paper by Li et al. [82].

3.4 Discussion

The gap in the literature discussed in Sections 3.1, 3.2, and 3.3 serves as the foundation for the development of this thesis. Therefore, this section discusses the most significant findings from related literature. The main objective is to develop a comparable and interpretable **Cost Estimation Model** based on real-world data. The model should be tailored towards companies without specific cyber security knowledge, more specifically SMEs. The second goal of this thesis is to provide companies with a **Risk Measure** that allows them to evaluate risk in consensus with the principles for risk measures established by the World Economic Forum [159]. Finally, this thesis addresses the lack of data in the field of cyber security economics by introducing a **Federated Learning (FL) Model** that allows for sharing knowledge about economic impacts without revealing sensitive information.

3.4.1 Cost estimation

As shown in Section 3.1, very little research about attack costs is available in the current academic literature. Examples of work that provide discussions on cyber security economics are Anderson et al. [6], SEconomy [135], and SECAdvisor [113]. From these three, only the SECAdvisor provides the option to tailor costs to the company by specifying the data type and the number of records [113]. However, the approach is limited since it focuses solely on data breaches, which excludes important attack vectors such as ransomware, which primarily aims to disrupt ongoing operations. To account for more factors such as residence, industry sector, and a multitude of attack types, a company has to conduct complex analysis of current industry reports from IBM [27] or Accenture [2, 3, 10].

In these reports, the relationships between average costs and influencing factors (*cf.* Table 3.1) are numerically described. One issue with the reports is that they do not provide guidance on how to combine the average costs for various factors, units of measurement, and dates. For example, suppose an American company has an average cost of \$ 21.22 million, and the average cost for companies in the financial sector is \$ 18.18 million. In that case, it is unclear how an estimator should combine this information. Furthermore, how should one scale the cost over time and different company sizes?

This thesis addresses this gap by developing a parameterizable cost estimator based on real-world data from various sources [2, 3, 10, 27, 126, 127, 128, 129]. The prediction can be established without specific knowledge in the cyber security domain, which benefits Small and Middle-Sized Enterprises. Furthermore, the estimator’s results are scaled to the company size at a certain point in time, allowing predictions of future costs.

3.4.2 Risk Estimation

Although the gap in economic cyber impact estimation is clear, the expected cost value can not be analyzed isolated from risk as current approaches or consultant reports do. Cyber risk is often expressed in some score variants, such as SecRiskAI [49] classifying risk into the categories low, medium, and high. Another example is Li et al. [83], which uses a score between 1 to 100 to categorize risk. The issue with such approaches is that they do not account for individual risk perception and assume a global score. As stated in [16], risk perception is a highly personal decision-making process based on an individual’s frame of reference developed over a lifetime. Therefore, a company with a high-risk score in the SecRiskAI approach might perceive this score as desirable, while another company views medium risk as intolerable. Comparing these two companies with the same benchmark, which assumes low risk as the best possible outcome, can lead to wrong conclusions. This situation could also be shown using Figure 3.1: The absolute variance of \$ 3650 in Figure 3.1a might be considered high cyber risk by the scoring benchmark. However, if the company’s overall revenue fluctuates by an absolute variance of \$ 100,000 due to risk, the company may perceive a cyber risk of \$ 3650 as relatively small. On the other hand, a company with an overall risk in the form of an absolute variance of \$ 36 might consider the cyber risk of Figure 3.1a as unbearable. It is, therefore, unpractical to use risk measurements from related work researched in the scope of this thesis. Another downside of risk scoring is that quantitative categories are complicated to compare to risks of other domains. Hence, it violates an essential principle of the cyber risk concept established by the cyber resilience initiative [159].

The reason why the expected value of costs cannot be viewed independently from risk is shown in Figure 3.1, which highlights two hypothetical probability density distributions of annualized cyber costs of a company. In the context of this example, cost includes investment costs into security measures and the expected costs of cyber attacks, thereby the problem consists of two variables (risk and cost) with an assumed partly independent, possibly inverse, relationship. Since the prices for security measures are company-individual and often publicly available, this thesis focuses solely on the expected cost and risk estimation. In this simplified demonstration, the distribution on the left 3.1a has the lower mean, hence the lower expected value of costs. However, what can also be observed is that the variance of the highly right-skewed distribution is more significant than in Figure 3.1b. In this scenario, the company can now decide which distribution it prefers based on the cyber security investments for the coming year. A more risk-averse company will invest more into security, leading to higher costs but much less spread around the mean. Whereas another company is willing to take the higher risk with the benefit of lower total costs. This hypothetical scenario depicts that a company cannot opt for a solution based only on risk or cost but must optimize between them.




Figure 3.1: Hypothetical Cyber Cost Distributions

One risk measure that solves the issue of scoring-based risk measures is the variance. The variance is individually interpretable and can be compared across different domains. Nevertheless, variance is a rather abstract statistical measurement, which is challenging to interpret depending on the specifics of a distribution. It further is a two-sided risk measure, which means it also accounts for upside (lower cost) risk. Since a company would not classify lower cost than expected as risk, the variance is not an ideal risk measure. Therefore, it contradicts the requirements for cyber risk measures set forth by the World Economic Forum [159], particularly regarding ease of interpretation and transparency. An alternative, which originated in the financial sector, is the Value at Risk (*cf.* Section 2.2). As one observes in Figure 3.1, the VaR scales similarly to the variance. Nevertheless, with its time, value, and probability dimensions, it is an easily understandable risk measure with a lot of research background in other fields, mainly the financial sector.

Currently, there is very little research on the Cyber Value at Risk, with the most recent research using Monte Carlo simulations [37, 93] to determine the distribution and the desired quantile. The process of estimating specific input probabilities to the Monte Carlo simulation is cumbersome and expensive [37, 157]. In addition, they require intimate knowledge of the company's IT systems and extensive experience in the field of cyber security. These factors make it difficult for companies with low cyber security budgets, such as Small and Middle-Sized Enterprises, to attain CVaR numbers.

To address this gap, this thesis produces the Real Cyber Value at Risk (RCVaR) as a risk measurement unit. The RCVaR is based on the theory of the VaR, hence is comparable across domains and individually interpretable. It further leverages real-world data from industry reports (*cf.* Section 3.2) to develop probability distributions from which the α quantile is calculated. In consequence, the RCVaR does not rely on harm, security effectiveness, or threat probability estimations from entities with cyber security expertise. Using total costs, consisting of the expected attack cost estimation and the publicly available prices for security measures, and risk estimation allow companies to determine their optimized security plan. Therefore this thesis, with its expected cost and risk estimation, creates the foundation to conduct risk-cost considerations in Small and Middle-Sized Enterprises, allowing them to be prepared for a new digitized economy and to face their adversaries in cyberspace.

3.4.3 ML and FL for Information Sharing

The proposed approach in this thesis heavily relies on the accuracy and availability of data. Even though Section 3.2 presents an extensive summary of available data from different sources, data scarcity remains. The undersupply of data in cyber security economics is a well-established problem [17, 101]. The lack of data stems from the incentive of companies to hide data breaches and cyber incidents since their publication could trigger reputation effects, financial market impacts, or lawsuits [17].

To address this issue, this thesis also develops a neural network with Federated Learning. To initialize the network, it is trained to produce the same results as the cyber cost estimator model, which is also developed as part of this thesis. The advantage of the Federated Learning approach is that it allows sharing of conclusions while keeping anonymity (*i.e.*, without revealing any information besides the trained model). Consequently, the model will allow for improved cost estimation over time due to more data. Promising results in using centralized trained ML models in other industries provide evidence that networks can learn cost-estimating functions [33, 44, 143, 79]. Furthermore, related research shows that FL is used successfully in other privacy-sensitive sectors to achieve similar performance as centralized models while maintaining privacy [75]. Therefore, this thesis lays the groundwork for an anonymous data sharing tool for financial cyber security data to improve future models.

Chapter 4

The Real Cyber Value at Risk Model

This chapter proposes the novel model named Real Cyber Value at Risk (RCVaR) and its methodology in detail. The complete RCVaR model consists of three parts: The cost predictor, the risk estimation, and a Federated Learning (FL) solution. Since the risk measure in the RCVaR is a quantile-based metric, which follows the logic of the Value at Risk presented earlier, the term Real Cyber Value at Risk is sometimes referring only to the risk output. Nevertheless, the whole model is considered to be the Real Cyber Value at Risk. When using the model, both cost and risk estimation outputs underlie multiple assumptions. These assumptions are briefly mentioned at the end of each subchapter, and a more detailed discussion of them can be found in the evaluation section of this thesis. An exhaustive list can also be found in Appendix B.




Figure 4.1: Overview of the RCVaR Development Process

This chapter is structured as shown in Figure 4.1. First, a light is shone on the data sources used to compute the RCVaR. This topic includes an extensive discussion on how the data is extracted from the reports. The following section then gives an in depth look at the development of the cost estimator by deriving all sub-parts of the cost prediction. Once the financial impact estimation is covered, the focus of this chapter shifts toward risk prediction. There, the distribution of costs is derived statistically before presenting the risk model with its assumptions in detail. Besides risk and cost estimation, this thesis introduces a Federated Learning approach in Chapter 4.4 to address the issue of data scarcity in cyber security economics. In the end, a web solution that incorporates all models and makes them accessible to a broad audience is presented.

4.1 Data Gathering

As the Congressional Research Service (CRS) already pointed out in their 2004 report [17], a scarcity of data related to attacks and breaches exists. This has different reasons: Firstly, companies have strong incentives not to report breaches. The incentives stem from multiple fears. Most prevalent is the anxiety of reputation loss and increased costs to raise capital. Companies usually do not want to expose themselves to liability lawsuits or signalize that they are soft targets. Secondly, it is tough to assign a monetary value to the damage of an attack. Particularly, *Indirect* and *Opportunity* costs are complicated to measure. For example, it is hard to quantify the overall value lost due to costumers' decision not to buy more services and goods because of the attack [18].

During this thesis, to the best of the author's knowledge, there is no dataset consisting of monetary costs of cyber attacks publicly available. However, sparse resources are accessible in reports of major consulting firms such as Accenture [2, 3, 10], IBM [27], Ponemon Institute [126, 127, 128] and Kaspersky [72]. Nevertheless, due to different measurement units, survey years, and survey regions, it is impossible to effectively compare and merge the cost-related data between reports. Moreover, these articles only state the mean cost for their respective sample pool. As discussed in Section 3.4.2, the expected cost alone is not helpful due to the connection between total cost and risk. The only report which presents an array of monetary consequences of cyber attacks, from which risk can be derived, is the Accenture report from 2017 [2]. As a consequence, the Accenture reports [2, 3] act as a primary data source. Compared to other reports researched in the scope of this thesis, Accenture is the only consulting firm that included multiple types of cyber incidents in its research. Furthermore, the report covers the *Direct*, *Indirect*, and *Opportunity* economic losses of attacks. Moreover, they connect the costs to a time dimension, which eases the interpretation of the cost values. In contrast, the IBM handout [27] demonstrates the cost per breach but does not provide any information about their frequency. This complicates cost predictions since even minor impacts can have cumulatively high costs depending on how often they occur. Therefore, the Accenture reports [2, 3] are the primary data sources, whereas the other consultant papers are secondary sources of information.

4.1.1 Primary Data Source

The only publicly and freely available dataset of costs of multiple attacks was found in the Accenture report from 2017 [2]. Figure 4.2 shows how the anonymized data is presented in the report. The figure shows the annualized costs in US-Dollar (USD) of 254 companies surveyed, represented by dots.




Figure 4.2: Company Individual Cost Presented in the Accenture 2017 Report [2]

To obtain the information in the graph, Accenture surveyed over 2100 individuals of 254 larger enterprises, which boils down to an average of 8.5 interviews per company. The majority of the participating individuals (36%) were tasked with IT security or IT operations in their company. During the interviews, the participants were asked to report *Direct*, *Indirect*, and *Opportunity* costs for each incident response step over four consecutive weeks. The incident steps included subcategories of the *Consequence* and *Response* stages (*cf.* Section 2.1). These subcategories are *Detection*, *Investigation of the incident*, *Containment of the impact*, *Recovery*, and *Ex-post response*. Besides these data-gathering process, external consequences such as business disruption or revenue loss were additionally calculated using shadow-costing methods. After the data was collected in the research period of four weeks, the cost numbers were annualized [2] and converted to USD according to the exchange rates provided by the Wall Street Journal on the 16. August 2017 [34]. The complete methodology is confidential and, therefore, not explicitly explained in the study.

The sample of companies represents a diverse mix of 15 sectors. Most companies operate in the financial (16%) and the industrial sector (12%) [2]. The granularity of the sectors roughly corresponds to the 11 sectors of the S&P 500 [45]. Regarding the company size, only large enterprises, with a minimum of 1'050 enterprise seats, were selected. The number of enterprise seats corresponds to the number of employees with access to internal IT systems. Regional-wise, only companies from specific western economies, such as Germany, Australia, or the United States, were interviewed in the 2017 report [2].

It is crucial to outline that the study is not based on actual accounting information but on the statements of multiple senior officials per company. These statements might not accurately depict the reality, even though checks were implemented in the survey to assess the correctness [2]. The study further does not cover preventive expenditures for information security and company policy measures [2]. The study also does not consider the number of stopped attacks for the attack expenditure calculation. Furthermore, the cost estimations may suffer from non-response- or sampling bias. Additional bias could

have been introduced by the non-disclosed shadow-costing method. While Accenture does use statistical inference to calculate confidence intervals for the mean in their research, they state that it is not advisable to use the results for statistical tests.

Assumptions

- *Bias free dataset.*
- *The report from senior officials is a good representation of the reality.*
- *The currency exchange rates are not abnormal during conversion.*

4.1.2 Data Extraction


Unfortunately, the data shown in Figure 4.2 of the Accenture report [2] cannot be easily processed because it is presented in a graphical format [13]. To address this issue, multiple approaches from computer vision [65, 115, 116] were explored in this work to identify the actual cost numbers. All approaches work similarly by estimating the y-coordinates of the dots as well as the top and bottom lines. Once all these coordinates are known, the actual cost value can be calculated using Equation 4.1. For all approaches, the OpenCV library [132] was used, which sets the origin per default in the top left corner. Consequently, the y-coordinate of the bottom line is higher than the top line.

$$Cost = 90 \times \frac{Bottom_line(y, -) - Dot(y, -)}{Bottom_line(y, -) - Top_line(y, -)} \quad (4.1)$$

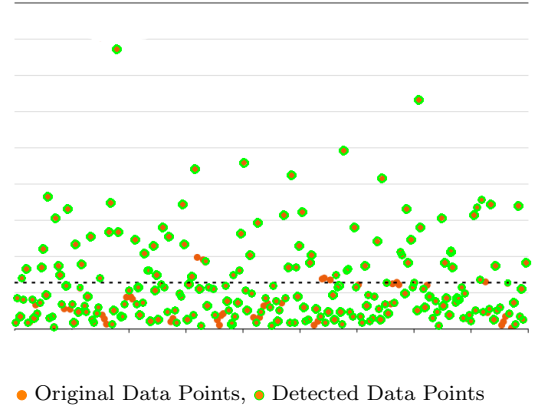
When inspecting Figure 4.2, one can observe visually two colors: Orange and Black. These characteristics are used in the first approach to identify the center coordinates of the data points [65]. More concretely, a simple comparison of each pixel color to the desired color orange (255, 153, 51) is conducted. This lookup results in a cluster of pixels with the color orange. In the next step, the center coordinates are derived using a depth-first search among the clusters of orange pixels. Unfortunately, this results in very few data points since the orange color observed by humans can have multiple slightly different Red, Green, and Blue (RGB) values. In total, the computer could observe roughly 2900 RGB combinations in Figure 4.2. Therefore, the acceptable range for the color orange is expanded. Through trial and error, it is determined that the best detection results are achieved by taking the 14 colors most similar to orange out of the array of unique colors detected by the computer. The color lookup approach resulted in 252 recognized dots, representing 99% of all company data points. However, upon visual inspection of Figure 4.3a, it becomes clear that the method detected some points twice and many not at all. Therefore, it was determined that searching for colors in the image to derive the coordinates and, finally, the cost numbers is not a reliable approach.

As a second approach, the OpenCV's Hugh Circle transform [115] is considered. This algorithm uses a 3D accumulator to estimate the three parameters of a possible circle,

namely the radius and the two center coordinates. Before applying the Hugh Circle transform, an edge detection convolution is applied to the image to detect edge coordinates, which are the foundation to fill values in the accumulator. The default algorithm used for the Hugh Circle transform is the Canny Edge detector [114]. The detected edges (white) in Figure 4.4b illustrate that edges are remarkably accurately detected. In addition to the default edge detection algorithm, the Sobel algorithm [117] is applied to compare the results. As one can see, comparing Figures 4.4b and 4.4a visually, the Canny Edge detection algorithm returns thinner lines and, therefore, more accurate edges.




(a) Color-Dot-Detection




(b) Dot-Detection With the Hugh Circle Transform Algorithm in Combination With Canny-Edge Detection.

Figure 4.3: Results of Dot-Detection Approaches Applied to Figure 4.2

When applying the Hugh Circle transform algorithm with the Canny Edge detection mechanism, 85% of the 254 actual points are detected, as seen in Figure 4.3b, where the green encircled dots mark the correctly located dots. Compared to the results in 4.3a, where the blue dots are the estimated dots and the orange the actual ones, there are fewer false positives.



(a) Detected Edges Using the Sobel Edge Detection Algorithm



(b) Detected Edges Using Canny Edge Detection Algorithm

Figure 4.4: Results of Edge Detection Algorithms Applied to Figure 4.2

In the last approach, the graph is visually examined and the dots are selected manually through a mouse click [116]. The program then remembers the mouse click coordinates and calculates the costs with Equation 4.1 in the next step. With this approach, 250 dots could be identified, which resembles 98% of all the data available. The visual inspection of Figure 4.5 also reveals that there are no false positives and that the estimated points capture the distribution well.




Figure 4.5: Manual-Detection of Data Points Through Mouse Clicks

All of the above-described methods, the color-detection, circle-detection, and the manual-detection, introduce some error to the actual truth. Since the development of the risk metric requires a correct distribution but not an accurate estimation of every single data point, the inaccuracy in the data is not critical. If all extraction methods have a similar distribution with a certain confidence, the overall probability of choosing the wrong distribution in Section 4.3.1, based on the data provided, is small. Therefore, a Kolmogorov-Smirnov 2-sample test [105] is performed to evaluate the parity of the data series extracted through the three methods from Figure 4.2. When comparing all three distributions, no null hypothesis can be rejected. This means that for none of the data distribution comparisons, one can reject the parity hypothesis. For instance, the data gathered by the circle-detection and manual-detection methods is not significantly different with a confidence level of 85%. Additionally, the hypothesis that the distributions between the manual-detection and color-detection methods are the same cannot be rejected with a high confidence level of 95%.

Assumptions

- *The extracted data is representative of the distribution in Graphic 4.2.*

4.2 Cost Estimation

This section outlines the development process of the economic loss estimation model to address the identified gap of customizable cost assessment models. Each component of the model is explained individually before merging the building blocks together in the Section 4.2.5. The initial section focuses on the process of scaling the cost based on the monetary size of a company. Next, the scaling of both costs and valuation over time is introduced. Finally, the design of the customization factors is shown, thus, allowing the model to tailor the cost to an individual company.

4.2.1 Size Scaler

Based on the data extracted in Section 4.1.2, a firm's considered average annualized costs in 2017 were \$ 11.7 million. The reason why the 2017 average is chosen over IBM's average cost per data breach for 2022 [27] is that Accenture covers all incidents, not only data breaches. Furthermore, to give an accurate cost estimate based on IBM's numbers, one must to determine the frequency of data breaches to estimate costs per time period. Therefore, out-of-the-box annualized costs from Accenture fit the role better and allow companies to integrate expected costs into their yearly financial reporting.

The first challenge when estimating the average cost per company is how to scale the cost to the company's size. To achieve this, the estimator uses market capitalization or, for unlisted companies, their equity valuation to scale the costs. Since the survey results are anonymized, it is impossible to determine the study participants' average valuation and match them to the average cost. The only information available is that the companies operate in 15 different sectors, which roughly correspond to the 11 sectors of the S&P 500. Furthermore, Accenture's survey only included larger organizations with enterprise seats (*i.e.*, the number of direct connections to the enterprise systems) ranging from 1'050 to 259'000.

Since the S&P 500 is comprised of the 500 largest US-based businesses, its average market capitalization is not a good approximation of the average valuation of an enterprise in the survey. An alternative is the Russell 1000 [87], which includes the large stocks of the S&P 500 but is more representative of the US economy due to an exposure of 93% to the entire US equity market. It is further unlikely that tech giants (*e.g.*, Apple, Amazon or Microsoft) were part of the Accenture survey since those large companies often have in-house consultants and their own cyber security experts. Furthermore, Accenture states in its report that the largest enterprise in the survey had 259'000 enterprise seats. Due to these restrictions, it can be assumed that a subset of the Russell 1000 should be utilized to approximate the average market capitalization of the sample.

The subset should consist of a large spectrum of company sizes from different sectors. The ideal match is the Russell Mid-Cap Index. It consists of large companies but ultra-large businesses, that are part of the Russell Mega-Cap Index, are excluded. In other words, enormous multinational companies, such as Apple or Amazon, are not part of the Index. At the same time, the Mid-Cap Index also excludes listed Micro-Cap enterprises, which

aligns with the minimum requirement imposed by Accenture. Furthermore, the Index includes 800 companies from a wide variety of industry sectors, which covers all sectors listed in the sample description. Due to all these reasons, the Russell Mid-Cap Index presents an excellent fit to investigate the sample's characteristics.

Figure 4.6 depicts the historical market capitalization of the Russel Mid-Cap Index according to the Bloomberg database [11]. The total market capitalization at the end of 2017 was \$ 7.7 trillion. Since the membership list of the Index is publicly available at [86], the average market capitalization for the year 2017 can be calculated as approximately \$ 9.63 billion. This indicates that the average valuation of a listed enterprise in the US economy, not including ultra-large and extremely small companies, was \$ 9 billion in 2017. From this information, it can be inferred that the average company in the survey sample had a similar market capitalization. Although the study also included companies from non-US economies, this conclusion remains valid since European companies in the study have similar characteristics as their American counterparts.




Figure 4.6: Market Cap of the Russell Mid-Cap Index (RMC) Based on Data From [11]

The conversion ratio can then be computed based on the average market capitalization for 2017 and the average annualized costs for cyber incidents. The conversion ratio or cost valuation ratio (*cv_ratio*), as shown in Equation 4.2, was developed in the scope of this thesis. The idea is to use a firm's 2017 valuation to estimate its expected costs for that year. Using the company's valuation or market capitalization enables a straightforward calculation of costs without having to understand the specific cyber security cost behavior of the company. To determine a company's equity value, known methods from economics can be used, such as stock market capitalization, peer group comparison, discounted cashflow or liquidation value [71].

$$cost_valuation_ratio(cv_ratio) = \frac{total_avg_cost_{2017}}{MarketCap_{RussellMidcapIndex}} \quad (4.2)$$

Assumptions

- *The average market cap of the Russell Mid Cap Index approximates the average market cap of the companies in the data samples.*

4.2.2 Time Scaler

As introduced above, scaling the average costs to the size of a company is only valid if the valuation is based on data from 2017. Additionally, the output only provides an estimate of cost for 2017. Therefore, this section addresses the challenge of accurately scaling the valuation and costs to the relevant year.

Numerous factors exist that influence an increase or decrease in annualized incident costs. For instance, the frequency and type of attacks might vary over time. Also, new law regulations might require additional spending in case of an incident. Another driver of cost increase over the years is undoubtedly inflation. When products and goods become more expensive, so become ransom, consultants, and protection solutions.

Figure 4.7 depicts the yearly inflation from 2010 to 2020 in the US on the y-axis. The inflation data was obtained from the Bloomberg database [12]. To determine the inflation rate, one must compute the increase of price of the Consumer Price Index (CPI). The CPI is an Index that measures the current price for a basket of goods. The price change of this basket over time reflects inflation [154]. On the x-axis, the average annualized financial impact of cyber incidents per company can be observed. The data is based on three Ponemon and three Accenture reports [126, 127, 128, 2, 3, 10]. Only these reports are used, because they use the same measurement units and methodology for collecting data, but not the same region. Including additional reports, such as those from IBM, would significantly decrease the explanatory power of the regression.




Figure 4.7: The Influence of Inflation on Cost

Since Accenture shifted the focus of the annualized reports from cost to threats, the Accenture 2019 report is the last cost of cyber crime study available. Therefore, the cost

impact for 2020 needs to be computed based on the cyber resilience report of 2021 [10]. The report categorizes companies into four categories according to how the company utilizes cyber security in terms of business strategy or resilience. The largest subgroup by far, with 55% of survey-participating companies fall into the “Vulnerable” category. Companies in that category have immature cyber security and protect the bare minimum of their infrastructure. To determine the expected cost for the year 2020, under the assumption that the affiliation of a company to a subgroup is unknown, one has to calculate the weighted average of costs of all categories. The weights in this computation represent the frequency of successful attacks of the respective subgroup. Even though the frequency is given for each category, the expected cost is only provided for two of the four. Therefore, the cost needs to be approximated with the expected cost of the two remaining subgroups. In the end, the weighted average cost of the subgroup yields annualized costs per company of \$ 13.8 million for 2020. Through interpolation, the remaining value for 2019 was deduced.

It is essential to point out that the number representing 2020 is based on research published in March 2020. Respectively before the full scale of the pandemic revealed itself. Since cyber crime saw a strong surge in activity during the pandemic [10], the effective expected cost for the year 2020 is considerably higher. However, this thesis considers the circumstances during 2020 and 2021 as rare abnormalities. Therefore, the regression result would be heavily skewed if later costs were included. Consequently, the adjusted pandemic cyber activity numbers were purposefully left out during the construction of time scaling variables to improve the predictability of the cost estimator model in regular times. To deal with cyber crime influencing black swan events, such as pandemics or wars, one has to view the cost estimator together with the risk measure developed in Section 4.3.2. More specifically, in terms of company individual black swan events, one would have to look at the tail beyond the α quantile. Risk measures for these areas exist, such as the expected shortfall. Nevertheless, it is tough to quantify them in the cyber domain because since the broad adaption of the internet; there have been very few global long-lasting crises. And even fewer that measurably impacted the financial cost of cyber attacks. Hence the lack of data is highly prevalent at the tail of the distribution.

Based on the inflation numbers from Bloomberg [12] and the annualized per-company average cost from Accenture and Ponemon, several regressions can be computed to determine relationships. Figure 4.7 shows the relationship between inflation and the economic cost of attacks. It can be observed that there is a positive relationship between both time series, meaning higher inflation also leads to higher costs. However, the result, with a p-value of 34%, can also be stated as not statistically significant. Furthermore, even if the relationship would show significance, its results must be consumed with caution due to very few data points and the selection bias introduced by leaving out 2021 numbers.

As inferred from Figure 4.7, the correlation between inflation and cost could be better. Therefore, the cost estimator model needs to scale cost and market capitalization independently over time. The intention behind it is that the valuation will be scaled to the year 2017, where it can be converted to cost after multiplying with the *cv_ratio*. Afterward, the resulting costs are re-scaled to the desired year.

$$company_valuation_{2017} = \frac{company_valuation_{2017+T}}{discount^T} \quad (4.3)$$

Since the discount model should work with simple multiplication (*cf.* Equation 4.3), respectively, by extrapolating the discount factor, a regression against the cumulative inflation is conducted. The beta of the regression in Figure 4.8 yields an average inflation of 1.81% from 2010 to 2020. An inflation discount factor of 1.81% is reasonably close to the real annualized inflation of 1.72%. The difference partly stems from the fact that a linear model was applied to only 11 data points. Furthermore, the limited scope of the analysis to the past 11 years means that the results cannot be generalized well beyond this time period. If, for example, current valuations are discounted with an inflation rate of 1.81%, the cost output would be heavily skewed due to current inflation numbers, which are close to double digits. Furthermore, it is indispensable to mention that this thesis does not consider asset growth. Meaning the cost of a company for the year 2014 is most likely lower than the output suggests if the enterprise experienced higher than inflation growth due to business activity. Nevertheless, the current estimator gives a reasonably accurate cost estimation for a company for the past few years. To increase the accuracy of the scaling process, the model could use the exact inflation numbers and inflation predictions. Furthermore, other discount values, such as the Weighted Average Cost of Capital could be used. However, using such measures would add more complexity to the model and, in the case of the exact inflation numbers, reduce generalizability over a considerable time horizon.




Figure 4.8: Evolution of the Cumulative Inflation Over Time

Since inflation is determined to be the valuation discount factor, the cost increase represents the discount factor to scale the cost to the desired year. Similar to the inflation discount factor, it should work by simply extrapolating the average cost increase by years (*cf.* Equation 4.4). Consequently, the costs were computed as a percentage of the initial cost in 2010. Meaning the cost of \$ 6.5 million in 2010 is marked as 100%, whereas the cost of \$ 8.4 million in 2011 is represented in percentage of the 2010 cost (129%).

$$company_cost_{2017+T} = company_cost_{2017} \times discount^T \quad (4.4)$$

The result of running a regression on the cost percentages is depicted in Figure 4.9. Similar to the regression where inflation is the dependent variable, the regression's slope represents the annualized increase. Hence it can be assumed that cyber cost rise on average by 11.2% (not inflation-adjusted). There are two significant differences to the inflation discount factor. First of all, the T in Equation 4.4 can be negative. A negative T means a down-scaling of cost for years prior to 2017. The second difference is that the cost data stems from various sources, which comes with challenges. The earliest data regarding average annualized cyber incident data is available in [126]. In the scope of this thesis, no data was found regarding cyber security costs on a company level before 2010. Moreover, data before 2013 often relies on samples that consist of only few companies. The smallest amount of 45 participating enterprises occurred in Ponemon's first study in 2010 [126]. Another relevant detail is that the studies before 2013 only researched US-based enterprises. This could explain the sharp drop in costs in 2013 since the study in that year specifically included a few European companies. Nevertheless, the regression is significant, with a p-value of 2%. Nevertheless, the significance must be viewed cautiously due to few data points. Overall, the beta of the regression is 11.2%, which, as stated before, reflects the yearly non inflation-adjusted cost increase. It can be expected that the increase is even a bit higher since the drop in the year 2013 smooths out the overall trend.




Figure 4.9: Evolution of Cumulative Cost Over Time

Assumptions

- *The Costs of the “Vulnerable” category in the 2021 report [10] can be reasonable approximated.*
- *Inflation is a reasonable discount factor for different assets across various industries.*
- *Both regression approximate the annualized increase of cost, respectively inflation, relatively well.*

4.2.3 Factor Selection

After applying the proposed scaling approaches, the result will resemble the expected cost in a particular year. But as several studies from IBM [27], Kaspersky [72], and Accenture [2, 3] demonstrate, costs deviate significantly from the overall average when more specific business characteristics are considered. Exemplary characteristics, among others, include the location or the industry. These company specifications lead to customized threat and vulnerability profiles and, consequently, individual cyber costs per company. This section defines the business characteristics from research (*cf.* Table 3.1), which are most important for tailoring costs to a business and, thus, are used in this thesis.

The different business characteristics, also referred to as factors in this thesis, are chosen based on data availability and relevance in other cyber security awareness assessment tools. The first assessment tool, which is investigated in the scope of this research, is a questionnaire developed by the Information Assurance for Small and Medium Enterprises Consortium (IASME) in association with the UK National Cyber Security Center [67]. The questionnaire is tailored explicitly towards SMEs and should enable them to identify vulnerabilities and how to act on them. In combination with the questionnaire, an official government certificate can be attained, which should signal to customers that the company places heavy emphasis on cyber security, according to the NCSC [104]. The second self-assessment tool, which serves as an orientation in the factor selection, is provided by the Cyber Security Osservatorio [26]. The Osservatorio is part of Italy's largest public research institution and receives funding from the European Commission and the Italian government [25]. Its services are specifically tailored to inform SMEs on cyber security topics [29]. Both self-assessment tools do not provide any cost estimations for cyber attacks and are solely qualitative tools to elicit the vulnerability of Small and Middle-Sized Enterprises.

Combining the data in reports and the requirements from both assessment tools leads to the selection of factors in Table 4.1. Each factor has several Parameters. For instance, a company can select between the USA, Germany, and many others within the category of business residence. It is noteworthy that even though the factor selection is based on related questionnaires, it goes far beyond the capabilities of both assessment reports. This is also illustrated in detail in Table 4.1. Furthermore, the estimator developed in this thesis covers multiple countries and a wider variety of security measures and more parameters than both assessment tools, which have a limited parameter selection.

As illustrated in Table 4.1, the *Country* is the first factor. The *Country* in this context refers to each country where the company has an office that has an operating computer. Even though the *Country* is in neither self-assessment tool, it is an integral factor since regulations, security pricing, and threat environments vary across regions. The same is true for the industry in which a company is operating.

The third factor shines a light on the IT system connections of company suppliers. Multiple reports [10, 3, 27] underline the importance of attacks on companies through supplier networks. For example, according to IBM [27], 19% of all data breaches in 2022 occurred due to infiltration over the supply chain. Therefore it is necessary to evaluate if the sup-

plier can be trusted. Unfortunately, due to the lack of data, the estimating tool requires the user to input whether he believes the supplier has trustworthy IT security or not.

As already established in Section 4.2.1, the size of a company plays a vital role in determining a firm's threat profile. Since the size of a company can be not only defined by its valuation but also by the number of employees, the headcount of an organization is a key input factor. To ease the use of the estimator tool, the parameters provided in this section align with the Organisation for Economic Co-operation and Development (OECD) classification of SMEs [118]. However, not only is the number of employees a crucial input factor, but also what cyber awareness they bring to the table. This circumstance is reflected in the *Training* factor, which asks the user to specify if the workforce in the company has received any cyber security training. For instance, training in detecting spam mails.

The fifth factor pays tribute to the fact that most businesses in today's digitized society use a cloud solution to sell their products or operate on a daily basis. Different network architectures in this context can lead to different vulnerabilities, which in consequence, need to be reflected ultimately in the cost. Hence the cloud type must be considered when determining the financial impact of cyber attacks.

Another major contributor to a company's costs is how many employees access the company's IT systems remotely. IBM discovered that the correlation between the number of remote workers and the cost of attacks is significant [27]. In 2021, for instance, companies with more than 80% of remote personnel experienced roughly 1.5 million higher costs on average than companies with more than 80% of the workforce on-site. To assess this factor, one needs to specify how much of a company's workforce is located remotely, by entering the exact percentage number. Internally, the program will assign the amount of remote workers to one of the categories listed in Table 4.1.

The last four categories resemble the cyber security measures which can be taken to secure an organization. These characteristics can be modified by investing in technology or by purchasing a cyber insurance contract. As all reports researched in the scope of this thesis declare, security measures and defense-enabling systems improve a company's cyber resilience and consequently reduce the expected cost. Hence, *Multi-Factor Authentication*, *Identity Access Management* systems, *Insurance*, and other measures are major factors to the estimator. The discussion regarding security measures often centers around which defense measure gives the biggest bang for the smallest buck. In other words, which action delivers the most effective defense for the lowest amount of money.

Table 4.1 summarizes all factors and their parameters used in the estimator tool of this thesis. Furthermore, the last two columns indicate whether the factor is mentioned by the respective self-assessment tool represented at the beginning of this chapter. Overall, 11 factors for customization were selected. Even though the valuation itself is a business characteristic needed to shape the cost to an individual company, it is omitted in the table due to its extensive discussion in Section 4.2.1.

Cost Factor	Parameters	IASME	Osservatorio
<i>Country</i>	USA, UK, Germany, France, Italy, Canada, Spain, Scandinavia, Turkey	No	No
<i>Industry</i>	Banking, Utilities, Aerospace and Defence, Software, Health, US Federal, Consumer Goods, Retail, Life Sciences, Communications and Media, Travel, Education, Automotive, Insurance, High Tech, Capital Markets, Energy, Public Sector, Pharmaceuticals, Industrial	Yes	Yes
<i>Supplier</i>	Supplier Safe, Supplier Not Safe	No	Yes
<i>Number of Employers</i>	Micro, Small, Medium, Large	Yes	Yes
<i>Cloud Model</i>	Public, Private, Hybrid	Yes	No
<i>Employer Training</i>	Training Received, No Training Received	Yes	Yes
<i>Percentage of Remote Employers</i>	0-20%, 21-40%, 41-60%, 61-80%, 81-100%	Yes	Yes
<i>Cyber Insurance</i>	Insurance, No Insurance	No	No
<i>Multi-factor Authentication</i>	Multi-factor Auth., No Multi-factor Auth.	Yes	Yes
<i>Identity Access Management</i>	Identity Access, No Identity Access	Yes	Yes
<i>Deployed Security Measures</i>	Automated Checks & ML and AI, Cyber Analytics and Behavior Analytics, Encryption Technologies, Risk management, Sufficient Security Staff, Incident response plan testing, Security intelligence systems, Advanced identity and access, Advanced Perimeter Controls, Data Loss prevention measures	Yes	Yes

Table 4.1: Cost-Influencing Factors and Their Parameters

4.2.4 Factor Scaler

This section explains the mathematical foundation of calculating the factors in detail. Before the calculation can be performed, the data must be preprocessed. This step comes with its own set of assumptions that also need to be discussed. Next, the data sources for every factor are presented before displaying the final parameters ratios graphically.

Preprocessing of the Data

After defining the significant factors for cost customization in Section 4.2.3, the factors need to be calculated from the available data. Table 4.2 illustrates from which sources the data was taken to compute the factors. It becomes clear that not all factors are available in every report. For instance, the influence of remote workers on the cost is described for two years in the IBM report [27] but not in the Accenture report [2, 3]. On the other hand, the number of employees is only stated for the year 2017 in the Accenture report [2]. Further challenges arise from the fact that even if the same factor is present in two reports, the parameters might not coincide. This issue is best demonstrated on the country factor, where the parameters Canada and Italy are present in the data for 2018 but not in 2017 [2, 3].

Further adjustments due to unclean data were necessary. The following paragraph describes them in detail:

- The cost estimator is tailored towards European and North American SMEs due to sparse data available for businesses outside these regions. The limited information available for Asian, Middle Easter, or South American companies is therefore purposefully omitted to achieve higher accuracy.
- Data for the parameter *Health* in the industry factor for the years 2021 and 2022 is also omitted intentionally. This corresponds to the argumentation presented in Section 4.2.2. More specifically, this thesis views the pandemic as a black swan event, which occurs relatively rarely. If the attack cost of the health industry, which spiked during these two years, is included in the factor calculation, the ultimate factor would be highly skewed. The inclusion would be especially grave since the data for 2021 and 2022 would make up 50% of all data related to the health industry.
- The last adaption of data needed to compute the factors is the unification of parameter names. For instance, the industry named “technology” in the IBM review [27] has a matching equivalent in the Accenture reports [2, 3] called “high tech”. To effectively use both industries as one parameter, the names were adjusted accordingly where necessary. In the case of the “technology” industry, this meant renaming it to “high tech”.

Cost Factor	Years	Source
<i>Country</i>	2017, 2018, 2021, 2022	[2, 3, 27]
<i>Industry</i>	2017, 2018, 2021, 2022	[2, 3, 27]
<i>Supplier</i>	2022	[27]
<i>Number of Employees</i>	2017	[2]
<i>Cloud Model</i>	2021	[27]
<i>Employer Training</i>	2022	[27]
<i>Percentage of Remote Employees</i>	2021, 2022	[27]
<i>Cyber Insurance</i>	2022	[27]
<i>Multi-factor Authentication</i>	2022	[27]
<i>Identity Access Management</i>	2022	[27]
<i>Deployed Security Measures</i>	2017, 2018, 2022	[2, 3, 27]

Table 4.2: Data Sources of Factors and When They Are Available

Computation of Factors

Once the data is cleaned according to Section 4.2.4, the issue of different time horizons and measurement units remains. To resolve this issue, the relative cost of companies with a certain business characteristic compared to the overall average for n reports is computed using Equation 4.5.



$$parameter_ratio = \frac{1}{n} \sum_{i=1}^n \frac{cost_parameter_i - avg_cost_factor_i}{avg_cost_report_i} \quad (4.5)$$

The Equation 4.5 cancels the respective measurement unit due to the fraction. This calculation is demonstrated with the concrete example of the *Banking* parameter. In the year 2017, the expected incident cost for companies in the banking sector was measured to be \$ 18.28 million (*cost_paramter*) [2]. The average across all k industry sectors (parameters) is calculated with Equation 4.6. Resulting in a sample average of \$ 10.348 million (*avg_cost_factor*) if companies were equally distributed across the different industry buckets. As one can imagine, the actual average of the sample deviates from the *avg_cost_factor* due to a non-uniform distribution across categories. However, the sample average is luckily stated by each report. For the year 2017, the overall expected cost for the sample is \$ 11.7 million (*avg_cost_report*). Then, with the help of Equation 4.5, the deviation of the banking sector's average company for the year 2017 can be calculated. The resulting ratio is 67.79%, which is also visible in Figure 4.10b. Expressed in other words: A company in the banking sector, on average, faces costs that are roughly 67% higher than the overall cost signalized in the report for the respective year. This parameter ratio is then calculated for each year where data for the parameter is available. For the banking sector, this analysis shows that additional costs were incurred in the following amounts: 67% in 2017, 45% in 2018, 37% in 2021, and 42% in 2022 (*cf.* Figure 4.10b). To ultimately obtain a single ratio per parameter, the average of all these numbers across the n reports in which data for that parameter is available is calculated. In the case of the banking category, the rounded average would be 48%. This is the final result of the Equation 4.5, which is further used in Equation 4.7 to tailor costs toward a company.

As stated in the previous example, the report provides two of the three inputs for Equation 4.5. The one parameter that needs to be computed before inserting it into the Equation 4.5 is the *avg_cost_factor*. This variable represents the expected value of costs if the industry is unknown. It deviates slightly from the mean sample cost stated in the report since not all industries are represented equally in the data set. The average cost over the factors (*avg_cost_factor*) is used because the deviation from the average industry, rather than the deviation from an average company, more accurately reflects the additional costs associated with a company's sector affiliation. To compute the mean cost per factor, one calculates the average over all k parameters within the respective factor. This computation is executed according to Equation 4.6.

$$avg_cost_factor = \frac{1}{k} \sum_{i=1}^k cost_parameters_i \quad (4.6)$$

Selected parameters for the *Country* and *Industry* factors can be viewed in Figure 4.10a and 4.10b, respectively. The ratios depicted there reveal only minor fluctuations over the years. In other words, the parameter ratios occur to be relatively stable. This situation can be observed in Figure 4.10a, where companies located in the US generally have expected financial costs of above 86% while France remains negative the whole time. The ranking among these three countries further stays the same for the duration of the data. A similar conclusion can be drawn from Figure 4.10b. The ranking among the four industries stays the same over time, except for the travel industry in 2021. Remarkably, even the tiny distance between Communications and Media and Consumer Goods is persistent over time. This observation is even more astonishing, considering that the parameters were extracted from different reports [2, 3, 27] with different samples, measurement units, and regions. This realization supports the hypothesis that the relative additional costs are persistent. Hence, taking the average of ratios over different years is a good approximation.

(a) Parameter Ratios of *USA, France, and Germany*(b) Parameter Ratios of *Banking, Communications and Media, Travel, and Consumer Goods*Figure 4.10: Parameter Ratios for *Country* and *Industry* Factors

After computing the parameter ratios, one can observe which characteristics influence the cost the most. Figure 4.11 shows each factor's maximum range in the positive and negative direction. Additionally, the parameter responsible for the ratio is displayed next to the bar. As one can observe in the bar chart, the business location can have a substantial impact. For instance, US-American companies have double the cost of an average company. The second highest influential factor is the organization's size. The more employees work in a company, the higher the cost due to cyber incidents. Regarding the organization size, the highest impact has to be expected in SMEs with more than 250 employees. Another interesting observation from Figure 4.11 is that the most effective

security action is the introduction of a Security Intelligence system [144], since it lowers cost by 20% on average. However, business characteristics that define a company's core, such as employee amount, location, or industry have a far more significant impact on costs than security systems. When looking at the graph, the most efficient action to reduce the cost with a minor investment might be reducing the number of remote workers. Providing an on-site workplace can reduce the financial impact by roughly 14%. This might be due to many reasons, such as fewer connections over insecure Wireless Local Area Networks (WLAN) or a more limited use of personal devices. Especially Bring Your Own Device (BYOD) policies which are strongly connected to remote work, can have a detrimental effect on the vulnerability of a business [72].




Figure 4.11: Maximum and Minimum Impact Parameters per Factor

Compared to Figure 4.11, Illustration 4.12 paints a more granular image. It can be observed that the *Country* factor consists of a few countries with very high costs, whereas the majority of business locations have a negative influence. This realization suggests that some countries suffer more from attacks than others. This could be due to a higher frequency of attacks or simply because of a more costly resolve. In the case of the US, the relatively higher cost compared to other countries could stem from an increased amount of targeted attacks of state-side actors. This hypothesis is further supported by data from the American think tank called Council on Foreign Relations, which blames US rivals

such as China, Russia, Iran, and North Korea for 77% of all globally state-side sponsored attacks [28]. Another verdict from Figure 4.12 is that there is an indication of normally distributed additional cost in the *Industry* factor. Similar conclusions can be drawn from the *Organization Size* factor.




Figure 4.12: Distributions Within Factors

As mentioned above, the industry reports only state the relationship of a single business characteristic to the average cost of all firms with this characteristic. This data structuring prevents an in depth investigation into correlation effects among factors. There are two types of correlation effects that may skew the results of the model. First, two business characteristics might often appear together, but only one substantially influences costs. For example, it could be that a majority of banking firms are located in the US. Moreover, since the parameter *Banking* increases cost, the average of a US-based company is also increased even though the location *US* might be negligible. Secondly, business characteristics can have different effects depending on whether they appear individually or in combination with each other. Expressed alternatively, if the parameters are not independent, their joint influence cannot be expressed by multiplication of factors. For instance, the parameters *Large* and *Industrial* increase costs when appearing individually. However, in combination, the increase might be less severe. Both correlation-introduced errors are very unlikely to heavily skew the data due to a well-diversified dataset and a factor selection based on reasoning and related work instead of quantitative reports. Nevertheless, the correlation effects should be kept in mind while analyzing the output of the RCVaR. Another assumption to remember is that the parameter ratios change over time. Figure 4.10 highlights that the changes over time are minor; however, they exist and therefore introduce an error to the model.

Assumptions

- *Cross-Correlation between Factors is assumed to be zero.*
- *Parameter ratios are constant over time.*

4.2.5 Complete Model

This section merges all the building blocks discussed in the previous chapters to construct the final estimator model. The economic impact estimation of a cyber attack for a company can be determined with Equation 4.7. The first step in this process is to determine the valuation of the company for which the costs are being estimated. The valuation should resemble the equity value of a company for the current year. Then, the valuation is discounted with the respective discount factor described in Section 4.2.2, where T_1 represents the number of years that have passed since 2017. The outcome of this operation is the valuation of a company for the year 2017. This interim result is then converted by the *cv_ratio* to costs (*cf.* Section 4.2.1). In the next step, the costs are scaled to the year for which the estimation is required. For instance, if an approximation for 2025 is required, the *discount_cost* is multiplied by itself eight times. In the end the estimation is customized by computing the product of *parameter_ratios* of all factors for which an input was provided. If no specification for a factor was entered upon execution of the estimator, the ratio is set to zero, which results in a multiplication of value one.

$$company_cost_{year} = \frac{valuation_{2017+T_1}}{discount_{valuation}^{T_1}} \times cv_ratio \times discount_{cost}^{T_2-2017} \times \prod_{i=1}^{11} (1 + param_ratio_i) \quad (4.7)$$

Both discount factors in Equation 4.7 can be retrieved from Table 4.3. The numbers in the table resemble the numbers presented in Section 4.2.2. More specifically, the cost discount factor of 11% was determined by the beta of the regression in Figure 4.9. The same is true for the valuation discount, which was identified by running the regression visible in Figure 4.8. The resulting beta of 1.8% is then defined as the discount factor for the market capitalization input.

Discount Factor	Slope	Sources
<i>Cost Discount</i>	$1 + 0.110$	[2, 3, 10, 126, 127, 128]
<i>Valuation Discount</i>	$1 + 0.018$	[12]

Table 4.3: RCVaR Discount Factors for Time and Size Scaling

4.3 Risk Measure

As discussed in Section 3.4.2, more than the expected value is needed for developing a plan of action during the cyber security management process. Therefore, this section presents the risk part of the RCVaR model. In the first step, the Generalized Inversion Gaussian distribution is determined to be the best fit of the spread for the sample. The development of a representative distribution underlies some strong assumptions, which are further discussed in detail in the following sub-chapters, before finally displaying the CVaR risk measure of the RCVaR model.

4.3.1 Distribution of Cost

The data extraction, described in Section 4.1.2, results in a data series of costs in US-Dollar. These costs stem from different large companies operating in 15 different sectors. In Figure 4.13, the distribution of extracted costs is displayed, with the y-axis representing the number of occurrences in the sample and the x-axis indicating the cost value in USD. It becomes clear that the distribution is heavily skewed to the left side, with a long tail on the right side. The heavy tail in Figure 4.13 is caused by a few companies with very high annualized costs. These findings are consistent with the 2020 IBM report [27], which suggests that “mega breaches” - incidents with exceptionally high costs - are relatively rare events. Of the 550 companies in the IBM sample, only 2.3% experienced these significant events, roughly corresponding to 2.7% of companies that experience extreme values (above 40 Million) in the Accenture sample (*cf.* Figure 4.13) [2]. Based on these findings, one cannot state any globally applicable relationship or characteristics. Nonetheless, the similarity of these findings suggests that despite differences in years, methodologies, and companies, the distribution of extracted costs in the sample is representative of the actual distribution of costs.




Figure 4.13: Cost Density Distribution of the Accenture Sample [2]

Further reports, such as the survey paper from Kaspersky [72], support the hypothesis of a heavily skewed distribution. The Kaspersky report states that a non-insignificant part of the costs can be attributed to professional services required after a breach. These services include external cyber security consultants, specialized lawyers, or PR consultants. The report indicates that consultants are required for 87% of incidents and, on average, make up 26% of the overall financial impact. Upon investigation of this sizeable cost contributor, it becomes clear that it approximates a similar distribution as Figure 4.13.

These empirical indicators can further be supported by theoretical reasoning: Most cyber criminals do not particularly choose a specific target [141]. Instead, the victim's reaction to bait or the presence of a vulnerability often attracts the attention of a cyber criminal. Exemplary of this circumstance are Phishing attacks. Phishing is a social engineering attack in which the attacker contacts the victim under a fake identity to infiltrate a network or to deceive the victim into revealing sensitive information [56]. According to the IBM 2022 cost report [27], Phishing is the most expensive attack vector and the second most frequent one. An attacker aiming to maximize profit will likely not target one specific email account. Instead, he will send the Phishing mail to as many email addresses as possible. The majority of these emails will be caught by spam filters, or the attack will be detected in an early stage. As IBM points out, the earlier an attack is detected, the fewer financial consequences the company will suffer. Following this argument, hypothetically only a tiny fraction of the sent Phishing emails will result in a successful attack and therefore in financial gain for the attacker. If one plotted the emails sent on the y-axis and the monetary gain of the criminal per email on the x-axis, the distribution would look very similar to Figure 4.13. Since the criminals' gains are the victims' cost, the distribution must be mirrored to get the firm's cyber cost. The mirroring is achieved in Figure 4.2 by changing the x-axis label from gains to cost. Finally, one arrives at the cost-per-incident distribution displayed in Figure 4.13. If all other attack vectors are subject to the same distribution, the overall distribution is a linear transformation of the previously established Phishing cost distribution.

In the next step, a goodness-of-fit test is conducted to determine which distribution fits best the data sample from the report. The test is necessary to determine the continuous function from which the α quantile can be computed. The Kolmogorov-Smirnov 1-sample test is used to perform the goodness-of-fit test. The KS goodness-of-fit test works as follows: First, the data is sorted by increasing order. In the next step, the Empirical Cumulative Distribution Function (ECDF) is computed and then compared to the Cumulative Distribution Function (CDF) of the target distribution. Finally, the vertical distance of each data point to the reference CDF is calculated. The best-fitting distribution is the one with the lowest single maximum distance [98].

Before applying the test to the extracted data sample, it is necessary to discuss the assumptions of the KS goodness-of-fit test:

- **Continuous Distributions:** The test is only applicable to non-discrete data. If this requirement is disregarded, the p-values are not exact, meaning the confidence threshold is most likely to be lower. The fewer data points are in the sample, the more severe this effect is. Since we search for extreme p-values in a sizeable sample, this assumption is relaxed. Furthermore, the KS test can be extended to categorical

discrete ordered data as stated on the NIST website [106]. Since the data could be ordered into buckets, this requirement could be full-filled. However, the application of the extended discrete KS test goes beyond the scope of this thesis. To conclude, due to the sample size and the expectation of very high p-values we relax the continuous assumption.

- **Sensitive at the Center:** The Kolmogorov-Smirnov test is more sensitive at the center of the ECDF than at the tails. This circumstance is not a severe issue for the extracted sample because more data is in the center, which increases confidence in the distribution. Sensitivity at the tails could be more challenging because of very few indications of the actual distribution shape in that specific region.
- **Fully specified Distribution:** When a distribution's parameters, such as mean and standard deviation, are determined from the data sample itself, it could lead to biased p-values. This means that a sample would not be rejected due to a high enough p-value, even though it should have been [146]. Similar limitations apply to the Anderson-Darling test, an alternative to KS [106]. For testing the hypothesis of normal distribution the Lilliefors test, which allows estimations of parameters from the sample, is applicable [146]. Unfortunately, testing in this thesis's scope requires a test which is applicable to multiple distributions. Since hypothesis testing aims to determine the distribution which is most unlikely to be rejected by the true sample, the assumption can be relaxed. The expected p-values should be well above the significant 5% threshold; therefore, a more conservative Kolmogorov-Smirnov shall not affect the results. Furthermore, every graph is visually assessed, to ensure the best possible fit (*cf.* Figure 4.14). Nevertheless, it is vital to remember while reading that the p-values in this thesis are approximations and should be considered as such.

After considering the limitations of the KS 1-sample test, the test was conducted using the Python library `scipy.stats` [151], as shown in Listing 4.1. In the first step, the code iterates over all 101 continuous distributions available in the `scipy` package and fits the distribution parameters to the data sample (*cf.* 3rd KS assumption). After storing the parameters in a dictionary, the Kolmogorov-Smirnov test is performed. For each test, the maximum distance and the p-value are returned. Finally, the distribution with the lowest likelihood of rejection, based on the p-value, is selected from the dictionary. The results can be observed graphically in Figure 4.14. Besides displaying the distributions with the highest p-values, the figure additionally shows the adjusted normal distribution. Figure 4.14 shows that all four best-fitting distributions are an appropriate continuous representation of the discrete data sample. The only difference is whether the density at the origin is zero or not. If the example regarding Phishing emails from earlier is revisited, the argumentation can be made that even when the spam filter detects the mail, costs occur. These costs can be categorized into *Detection*, *Investigation and Escalation*, or *Ex-post Response* cost [2]. For instance, costs occur during the activities and deployment of technologies for early detection. Moreover, a spam email might trigger an investigation if it is a part of a large-scale attack and if other employees have received identical emails. And last but not least, the company will draw conclusions on how to protect against spam emails in the future. The analysis and implementation of these post-attack responses are spending points that should not be neglected. To summarize, there is a theoretical argument that

zero costs cannot occur hence the density is zero at the origin. Consequently, the Pareto distribution is deemed a poor match. A lower p-value of 46.9%, compared to bell curve distributions, supports this conclusion.

```


1   for dist_name in dist_names:
2
3       # get the distribution of the stats module in scipy
4       dist = getattr(stats, dist_name)
5
6       # fits the distribution to the sample
7       try:
8           param = dist.fit(data)
9       except:
10          continue
11
12      # store the parameters in the summary dictionary (params)
13      params[dist_name] = param
14
15      # Applying the Kolmogorov-Smirnov test
16      D, p = stats.kstest(data, dist_name, args=param)
17      # The letter "D" stands for maximum vertical "distance"
18      # p stands for the p-value
19
20      print("p value for " + dist_name + " = " + str(p))
21
22      # select the best fitted distribution (with the highest p-value
23      best_dist_name, best_p = (max(dist_results, key=lambda item: item
24      [1]))
25
26      print("Best fitting distribution: " + str(best_dist_name))
27      print("Best p value: " + str(best_p))

```

Listing 4.1: KS 1-Sample Test in Python

Another finding of the KS 1-sample test that can be discovered in Figure 4.14 and Table 4.4 is that the normal distribution is not a good continuous representation of the data sample. Even after adjusting for multiple testing with the Bonferroni method, which reduces the 5% threshold to $4.95e-4$, the normal distribution as a hypothesis can still be rejected. The multiple testing adjustment is necessary since, with every iteration in Listing 4.1, the chance of significance due to luck increases [90]. The graph in Figure 4.14 further supports this rejection since negative costs due to cyber attacks, which would represent gains for the company, are unlikely to occur.

The test concludes that the Generalized Inversion Gaussian (Geninvgauss) distribution is most likely not to be rejected. Important to point out is that the two statements: “distribution which is most likely not to be rejected” and “most probable distribution” are not equal. Therefore, it is not possible to determine “the one” solution. This circumstance can also be observed in Figure 4.14, where only minor differences between the distribution with the highest p-value and the subsequent ones are visible. Nevertheless, in the scope of this thesis, the Generalized Inversion Gaussian is chosen as the best fit due to the highest p-value.



Distribution	P-Value
Geninvgauss	0.98921
Exponnorm	0.96195
Norm	4.89e−6
Recipinvgauss	0.97147
Pareto	0.46933

Figure 4.14: Distributions Fit to Sample

Table 4.4: P-Values of Distributions

Assumptions

- *The assumptions of the Kolmogorov-Smirnov test can be relaxed.*

4.3.2 Discussion on RCVaR Application

Before applying the RCVaR, the assumptions regarding the risk distribution must be questioned to avoid wrong conclusions. As described in Section 4.3.1, the distribution of costs is derived from the loss of multiple companies. Meaning the data does not represent a time series of a single company. Instead, the distribution is determined by a data snapshot of multiple businesses with different characteristics, which is problematic when applying the distribution to a single company to evaluate its risk. In other words, the distribution of costs of a single company does not necessarily converge with the distribution of costs of multiple companies. Therefore the assumption is made that risk does not vary over different industries, countries, and other characteristics. Consequently, this allows viewing all data points as costs from the same company. Subsequently, the model can make statements on a single company's Value at Risk.

The assumption regarding risk similarities among different business characteristics is rather strict. Nevertheless, one could argue that the companies in the sample are somewhat similar to each other due to the selection criteria used by Accenture [2]. The sample only includes companies from certain sizes and industries, and the business processes of large companies tend to have many similarities because of regulatory requirements, market conditions, and best practices recommended by consulting firms. This hypothesis is supported by empirical data from an IBM study [4], which found that 39% of cross-industry processes have commonalities. This number is even higher within the industries themselves. Furthermore, the study looked at the degree of similarities and discovered that 20% of processes have more than 50% similarity according to their scoring.

This result indicates that, although the assumption of equivalency is a simplification, the distribution should still yield accurate risk predictions to conduct a risk-reward analysis. In other words, the outcome of the trade off between the estimated costs from Section 4.2 and the risk from the distribution should have a neglectable error rate.

Another challenge that arises due to the unavailability of data is that the distribution change over time is impossible to elicit. If more data were available, research into seasonality and long-term effects could be conducted. In the scope of this thesis, only one study from Stanford [78] was identified that investigates the evolution. It shows that enterprises manage to eliminate the significant tail risks over time. Hence the expected average cost also decreases. Calculating the cyber risk and expected cost with a heavy tail distribution, as seen in Figure 4.13, leads most likely to higher cost estimates than what should be expected in the future. The further in the future the prediction lies, the more negatively skewed the risk. Therefore, the risk as well as the expected cost of the model can be viewed as conservative estimations. Due to the results from Stanford [78], it becomes clear that the distribution does not become worse over time. It is therefore allowed to assume constant variance of the determined risk distribution over time since it most likely leads to over-estimation of the risk rather than under-estimation.

The last assumption which needs to be discussed in the context of the probability density distribution of costs is the effect of company size. The final model aims to estimate risk and cost for Small and Middle-Sized Enterprises. Since SMEs are not represented in the extracted data sample due to the enterprise seat restriction, one has to assume that the exact cost distribution also applies to smaller firms. This assumption is supported by the Accenture report, which suggests a linear relationship between costs and the number of employees. A study by the university of St. Gallen [36] from March 2022 supports the heavy influence of business size on cost and risk numbers. However, in their work, the authors contradict the findings from Accenture [2] in saying that the relationship might not be linear. They found higher per-data breach costs for smaller incidents. Under the assumption that smaller incidents are mostly affiliated with SMEs, they hypothesize that SMEs might have higher relative costs than larger corporations. Similarly to Accenture [2], the conclusion of this study [36] is based on a limited number of observations (324) and is restricted to US companies (unlike the Accenture report). As the *Country* is an essential factor, as shown in Figure 4.12, and the relationship between factors has not been investigated, it is necessary to exercise caution when considering this general statement. Overall, no empirically significant statements regarding the relative higher costs in smaller-sized companies are formulated in the study. Therefore, the constraint of different cost behavior across the organization size factor is relaxed. Meaning the same behavior as in the sample is assumed for SMEs. Nevertheless, it is worth noting that scaling also introduces an error that may result in an underestimation of costs for smaller firms, according to the university of St. Gallen [36].

After being aware of the assumptions that go into the risk distribution computation, one can apply the model. In the first step, the expected costs are predicted with Equation 4.7 in Section 4.2.5. In the next step, the continuous distribution of the Generalized Inverse Gaussian function is created with the help of the `scipy.stats` library in Python [151]. Following these steps, the Percentage Probability Function (PPF) has to be computed. The PPF, which returns a discrete value for a given probability [107], defined by four

parameters is described in [134]. Besides the scale and the location, the input parameters are the same as the ones derived from the distribution of the sample data. The exception regarding scale and loc stems from the fact that scaling of the distribution requires modification of these two parameters, as stated by the library [150]. To derive the new location the original variable is scaled proportionally to the size of the expected value. Meaning the distribution is shifted on the x-axis until the correct position is found. Next, the scale, which resembles the spread is adapted similarly. Following these instructions will result in the same distribution, which is just linearly scaled to the size of the expected value. Ultimately, setting the variance, and with it the risk, constant.




Figure 4.15: RCVaR of an Average Company in 2019 With 95% Confidence

The result of the scaling can be viewed in Figure 4.15. It shows the distribution of costs of a company with a valuation \$ 1'000'000 for the year 2019. Based on this distribution, the discrete value to which the realized costs are less or equal, with a certain probability (95%), can be calculated [107]. This discrete value can be found at the 95th quantile in this specific example. To reiterate, the green area under the curve in Figure 4.15 marks the range of costs up to the 95th quantile. While the x-axis of the blue area shows costs that occur with a probability of 5% or less. In summary, the expected costs are determined to be roughly \$ 8'000 with a Real Cyber Value at Risk of \$ 23'448. This number can also be observed in the graphic. Additionally, it is noteworthy that the mean of the distribution and the expected value, which resembles the average, do not match. This is because the continuous PPF is only an approximation fit of the sample. And even though it is the best possible fit, its mean cannot match the expected value since that variable is based on exact numbers from the data. Overall, based on the expected value and the risk represented by the RCVaR, a business can start the cyber risk management discussion and determine the best-suited risk-cost optimum for the individual company.

Assumptions

- *Companies in the sample are similar.*
- *Distribution is the same over time and factors.*
- *Costs of SME have the same distribution.*

4.4 Model Development

Throughout this thesis, the issue of data scarcity has come up on multiple occasions. This dissertation addresses this issue by proposing a Federated Learning (FL) approach. Federated Learning allows training models in a decentralized manner, meaning the data stays on the user's local device, and only the model weights are shared with a central entity. This procedure allows all clients, in this case, companies, to profit from the experiences of others while keeping the exposure to cyber attacks private. In the first step, the model is initialized by learning the proposed function in Section 4.2, building a foundation for further improvements through the involvement of companies' participation in the FL process. However, a synthetic dataset must be created before the neural network can learn the cost prediction Equation 4.7. Therefore, Section 4.4.1 describes the development of the semi-synthetic data in detail. In the next and final step, a neural network architecture is proposed and trained using a FL pipeline.

4.4.1 Data Generation

A synthetic dataset is an artificially created dataset that incorporates real-world characteristics while keeping data anonymity. If the synthetic data closely resembles the actual observations, no performance loss should be observed when training the model on the artificial data [60]. In the case of the Real Cyber Value at Risk, no real observations are provided by the reports [27, 2, 3, 10]. For example, no report presents an actual company with its business characteristics and associated costs. Instead, only the costs of anonymous organizations are depicted in graphical form. After the data extraction (*cf.* Section 4.1.2), relations between a single characteristic and costs can be evaluated (*cf.* Section 4.2.4). Nevertheless, the connection of detailed real-world entities to costs is still unfathomable. Therefore, this thesis generates a dataset that adds several business characteristics to a theoretical firm and then determines that firm's cost based on the known cost patterns. This approach assumes that cross-correlation effects among factors can be disregarded. In this case, a cross-correlation example would be if a large percentage of financial companies were situated in the US, and their expected cost was higher not due to the industry but the country affiliation. As mentioned before, the publicly available data only states the relationship of one business characteristic to costs, which prevents research into cross-correlation effects. Nevertheless, the observed cost patterns in the consultant reports allow a rough assessment of where the costs lie. Consequently,

the generated dataset is marked as semi-synthetic since it consists of hypothetical firms but reflects real-world cost patterns.

Listing 4.2 shows a simplified version of the code responsible for generating data. As one can observe, artificial companies are generated from 2012 till 2022. Within each year-loop, the capital associated with an observation has to be determined. This is achieved by selecting a random value of the normal distribution around the mean of the Venture Capital (VC) valuation for that specific year. VC is a form of private equity which aims to find startups or small companies with potential. Upon finding, a Venture Capital fund assesses the value of the enterprise before seeking participation in the company in exchange for money. The companies' Venture Capitalists invest in can further be divided into early and late-stage companies. Early-stage companies can be described as companies with a functioning business model and a usable product. They further might already generate revenue [121]. A study by Deloitte [32] in 2020 further elaborates that the average employee amount of VC firms is around 14, which corresponds to the organization size of SMEs. Due to all these reasons, an early-stage Venture Capital firm is a good representation of a SME. This similarity in characteristics is essential since there is more data available regarding the market capitalization of Venture Capital firms than SMEs. Therefore, in the scope of this thesis, the market capitalization of VC companies is assumed to be a good approximation of the average SME equity value. Report data [123, 124] of Pitchbook, a large financial data supplier company [125], provides quarterly equity value averages of early-stage Venture Capital firms for 2012 to 2022. Then, as stated at the beginning of this paragraph, a value is randomly chosen from a normal distribution around the mean for the respective year. The *select_random_capital()* function represents this process in the code.

Once the year and the capital are determined, the program code in Listing 4.2 elicits the single factors in line 27. From all parameters available for a particular factor (*cf.* column parameters in Table 4.1), one is chosen arbitrarily. It is noteworthy that the term *None* was added to each factor. Hence, when the value for the *Supplier* factor is set to *None*, it is representing the state where the user does not specify an input. There are four exceptions for which a more complicated procedure is applied. These exceptions include the following factors: *Industry*, *Country*, *Security Measure* and *Remote*. In the *Remote* case, a random number is generated with a probability of 66%, or the factor is set to *None* in the remaining 34%. Since internally, the model later categorizes these numbers into one of the five categories visible in Table 4.1, the categorization process is already performed in the preprocessing. The *get_random_value_remote()* function symbolizes this step.

```

1 for year in range(10):
2
3     input_year = 2012 + year
4
5     # for each year generate n datapoints
6     for i in range(nr_data_points):
7         capital = select_random_capital()
8
9         row_dict = {"year": input_year, "capital": capital[i]}
10
11        # for each datapoint iterate over all parameters
12        for ind, row in df_param.iterrows():

```

```

13
14     # default
15     selected = ['None']
16
17     # special case remote (random number between 0-100)
18     if row['param_name'] == "remote":
19
20         # in 30% of the cases it should be None
21         if random.uniform(0, 1) < 0.66:
22
23             selected = [get_random_value_remote('param_name')]
24
25         # alternatively select one of the options (None included)
26         else:
27             selected = [get_random_value('param_name')]
28
29         # in 50 % of cases add additional security measures
30         # or countries or industries
31         if row['param_name'] == 'industry' or
32             row['param_name'] == 'country' or
33             row['param_name'] == 'security_measure':
34
35             # with a probablility of 50% and not all parameters
36             while random.uniform(0, 1) > 0.5 and
37                 len(selected) < row['nr_param']-1:
38
39                 additional_param = get_random_value('param_name')
40
41                 # add the new parameter if it is not None
42                 if additional_param != "None":
43                     # add additional parameters
44                     selected.append(additional_param)
45                     # remove duplicates
46                     selected = list(set(selected))
47
48                 # if None was originally in the array
49                 if len(selected) > 1 and "None" in selected:
50                     selected = list(filter(
51                         lambda item: item != "None", selected)
52                     )
53
54
55     row_dict[row['param_name']] = selected
56
57     # calculate the expected value based on business characteristics
58     row_dict['EV'] = data_generator.get_expected_value(...)

```

Listing 4.2: Abstract Data Generation

The other exceptions are due to multi-selection. Since Small and Middle-Sized Enterprises can operate in multiple countries and industries, this should also be reflected in the data. Furthermore, companies might employ multiple security measures at the same time. Therefore, the data generator keeps adding parameters to the factors *Country*, *Industry* and *Security Measure* with a probability of 50%. During these operations, it is crucial to ensure that the final array of parameters only consists of unique parameters

and no *None* term since a user cannot select parameters of a factor while not providing input information simultaneously.




Figure 4.16: Semi-Synthetic Data Along Capital and Cost Dimensions

In the end, the expected costs are calculated with the function developed in Section 4.2.5. These costs are the y, or target value, in the generated dataset. Since the goal of the later developed neural network model is only to predict the expected costs but not the risk measure, no other variable has to be provided for an observation in the dataset. If a quantile-based risk measure, namely the RCVaR, should also be determined by a Machine Learning model, the CVaR must be added to each observation. After collecting all variables for an observation, the factors, the expected values, the year, and the capital are merged into a row vector. By repeating the above-described process for *nr_data_points* required per year, the final semi-synthetic data set is created.

The result of the data-generating process can be observed in Figure 4.16. The figure depicts the data plotted along the two non-restricted axes. Meaning that besides the capital and the predicted cost, all variables have a finite number of possibilities (*c.f* Section 4.5). As one can observe in Figure 4.16, the data is well distributed among the two dimensions. It further becomes clear that no point is negative as it should be. Moreover, the expected cost increases the more the company is worth. This observed characteristic reflects the relationship manifested in Equation 4.2. Additionally, the breadth in the logarithmic y-direction stems plausibly from the customizing factors, including the time scaler.

Assumptions

- *Venture Capital market capitalization approximates the valuation of SME.*

4.4.2 Data Preprocessing

Since the data generated in Section 4.4.1 is semi-synthetic, only very little preprocessing is needed before feeding it to the neural network for training. Nevertheless, some steps are required to ensure a smooth training process. This section lists the steps sequentially and explains them in detail.

- **Special Column Naming:** All the data except the equity value of a firm and its expected financial loss are categorical variables. Thus, they have to be one-hot-encoded to be used as input for the neural network. One-hot-encoding is a process where each possible categorical value is represented by a binary vector in the dataset [7]. Problematic in this context is that two variables with the same name exist: The insurance industry and the *Insurance* factor. Since this would lead to a single one-hot-encoded vector instead of two, the insurance sector is renamed to *insurance_sector* in the preprocessing.
- **Introducing None:** Another difficulty that arises in the context of one-hot-encoding is due to the existence of *None* values. Every non-mandatory input can be *None*, which then results in one binary vector for all *None* values during the one-hot-encoding process. This result is undesirable since a non-specified *Country* factor has a different impact on the expected cost than a non-specified *Industry* factor. Therefore, to achieve multiple *None* one-hot-encoded columns, one for each factor, the factor as a suffix was added to each occurring *None* value. For instance, if no input is given for the *Industry* factor, the value *None* is replaced with *None_industry*. This procedure leads to factor-many one-hot-encoding vectors for the value *None*.
- **One-Hot-Encoding:** After the previous adjustments had been made, the data was one-hot-encoded with the help of the MultiLabelBinarizer [138] of the sklearn package [122]. This complex binarizer was necessary because a single factor might have a list of parameters as input, *e.g.* a firm could have offices in Canada and the US, which would lead to an input of type: *[CAN,US]*. Since a one-hot-encoded for both *CAN* and *US* was desired instead of one vector for each combination, a more complex encoder had to be used.
- **Remove Binary Input Variables:** Some factors in Table 4.1 have only two possible states. For instance, a firm can either have insurance or not. Technically, a third state exists with the above-introduced *None* case. After one-hot-encoding these states, there is redundant information in an observation data point. Since a zero in the insurance column indicates a one in the *No Insurance* parameter column. The redundancy of information does not bring any benefit. On the contrary, it increases the issue of vanishing gradient because it adds a zero as input for any observation.

To address this issue, redundant information is removed. More concretely, the one-hot-encoded vector of the negative category of all binary factors is removed.

- **Scaling:** Besides the categorical data presented above, the input variables also consist of numerical inputs in the form of year and capital as well as a numerical cost output value. Even though these variables are essential, their higher value compared to the ones and zeros of the one-hot-encoded vectors over-emphasizes their contribution to the output. In other words, their influence on the outcome of the neural network is too strong. Therefore, they must be brought in reasonable distance to the one-hot-encoded scale. In the case of the *Year* parameter, this was done by subtracting the year 2009 from the specified input year. This scales the years to a scale from 1 to 16. For the other two variables, capital and expected cost, the MinMaxScaler [137] available in the sklearn package [122] was used. The scaler adjusts the numerical values according to Equation 4.8 to a range of 1 to 20. The relatively higher values of time and monetary valuation compared to the binary values of the one-hot-encoded vectors should signal the neural net the importance of the different inputs. The reason for the higher importance of numerical variables is the fact that they impact the output more heavily than the customizing factors (*cf.* Section 4.2.4). The most significant customizing component by far, for example, is the size scaler presented in Section 4.2.1.
- **Shuffle and Split:** In the final step, the dataset is shuffled to remove the chronological ordering introduced during data generation in Listing 4.2. After shuffling, the data is split into train and test sets according to the ratio of 80:20.

$$X_{scaled} = \frac{X - X_{min}}{X_{max} - X_{min}} \times (20 - 1) + 1 \quad (4.8)$$

4.4.3 Neural Network with FL

After the preprocessing step in the previous chapter, a neural network is trained with Federated Learning as described in Section 2.3. The model architecture chosen for this approach is a simple feed-forward neural network since it is the most common network architecture used in the researched literature in Section 3.3. Furthermore, it is chosen because a neural network can technically approximate any function [7], and therefore is capable of learning the cost-estimating function. The exact architecture can be viewed in Figure 4.17. In total, the network consists of 4 hidden layers, with the first two layers having the same amount of neurons as the input layer. On the other side, the output layer of the dense network consists of only one neuron, which contains the predicted cost of a company. Due to a high probability of having multiple zeros in the 66 input features, the danger of vanishing gradient is prevalent. To address this issue and achieve a smooth training process, batch normalization [136] is introduced after every layer. Batch normalization specifically means that each layer's output is normalized to a distribution with a mean of zero and a standard deviation of one [20]. Then before passing the output further to the next layer, they are put through an activation function. In all the layers, the same activation function was used: Relu [22].


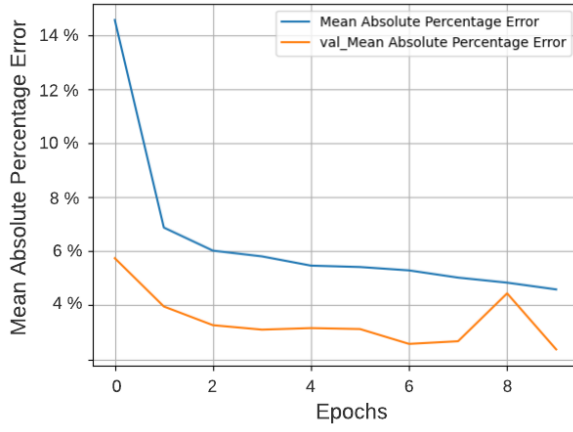


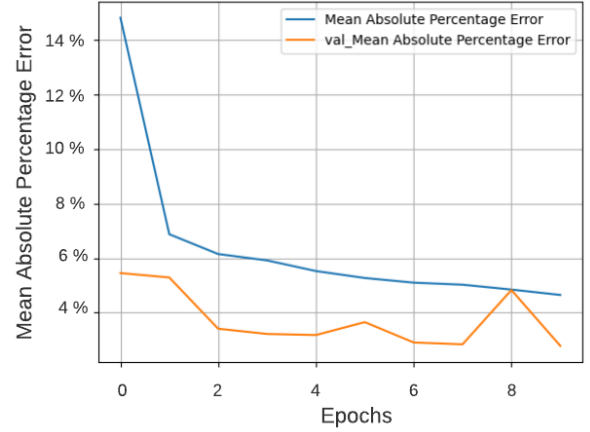
Figure 4.17: Neural Network Architecture for Federated Learning

The architecture in Figure 4.17 is trained in a decentralized manner. Meaning the model is initially located on the server before it is shared among the clients participating in the learning process. In the scope of this thesis, the Federated Learning process consists of two clients and a server. All of these instances were running on the same machine and were communicating the parameters of the model over the ports of the host machine. After the model is shared with the clients, they train in parallel on their local dataset. The randomly shuffled original training data, which stems from the train-test split conducted in the preprocessing step, is divided into two equally sized training datasets. Other partition strategies, such that specific parameters only appear in one client but not the other, are also evaluated (*cf.* Section 5.3.2). The reason for these separation strategies is to ultimately achieve two clients with different training data. Within each client, a second split is performed to divide the data into actual training and valuation sets according to a 90:10 ratio.

Each client then trains the model depicted in Figure 4.17 individually. Both use the Adam optimizer [76, 23] with the Mean Absolute Percentage Error (MAPE) loss function [24]. The MAPE function computes the average deviation of the prediction to the true label in percentage of the true label. This measure is chosen because the relative difference between the result from the cost estimator model in Section 4.2.5 and the neural network should be as small as possible. The Adam optimizer [76], on the other hand, updates the weight parameters of the neural network according to the gradient. Adam uses two momentum factors, namely the exponential moving average of the gradient and squared gradient, to adjust the learning rate. These two momentum factors further decrease over time by β_1 and β_2 . The model weights are updated after each batch propagation using the gradient, the two momentum factors, and a learning rate of 0.001. The exact numerical values for the constants of the optimizer, batch size, and the learning rate can also be found in Table 4.5.



(a) MAPE Loss Round 1 Client 1



(b) MAPE Loss Round 1 Client 2

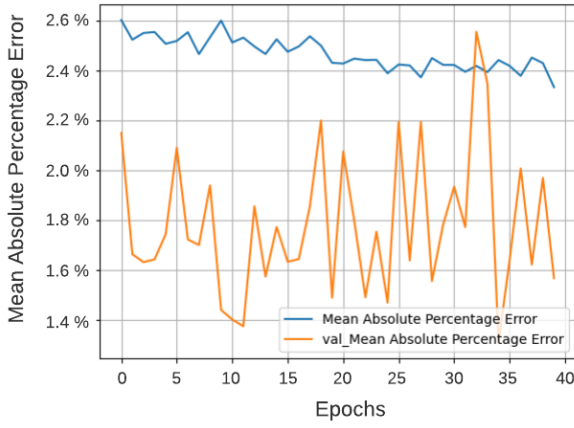
— Training MAPE — Validation MAPE

Figure 4.18: Mean Absolute Percentage Error (MAPE) of Clients in the First Round

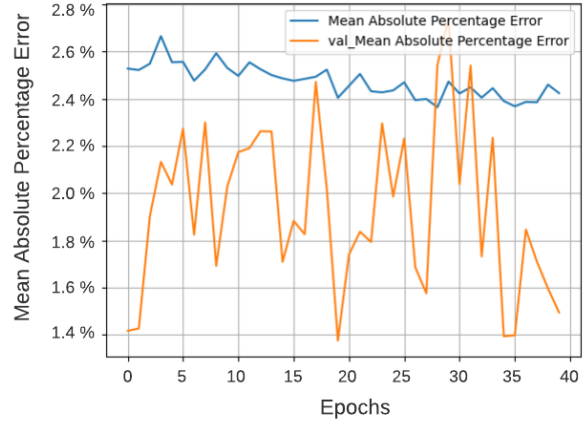
For the first three server-side aggregation rounds, the models are trained on the client side with ten epochs. This means the model will have seen each client's data ten times before it is sent back to the server. The relatively small number of epochs is chosen because of a fast loss decrease in the first few epochs. Therefore, a small epoch number reduces training time and allows for a more smooth model improvement on the server side since no extreme weight deviations are expected in the beginning. Figure 4.18a and 4.18b show both clients' Mean Absolute Percentage Error of the first round. As expected, the loss decreases relatively fast in the first two epochs in both clients. Visually, only minor differences in the plot can be perceived.

In the last two of the five rounds, the model is trained with 40 epochs on the client side. Figure 4.19a and 4.19b display both clients' loss and validation loss for the fifth round. As highlighted in both figures, the training loss settles somewhere in the interval of 2.4% and 2.6%. Furthermore, one can observe that the validation loss only passes the training Mean Absolute Percentage Error between epochs 30 and 35. Consequently, the model does not generally overfit during training. As in the loss graphs of the first round, it is evident that both clients' loss histories are very similar. This likely stems from similar observations since the generated data was divided into two halves before the actual training. Therefore other partition strategies are utilized in Section 5.3.2 to investigate if the model could still learn the loss function with clients whose observations do not contain all cost parameters.

After each round, the model weights are sent back to the server, where they are aggregated. To aggregate them, the default function is used, which calculates the average among the returned parameters [92, 8]. After aggregating the weights, the model is evaluated on the server with the test dataset. Therefore, the loss is not only tracked on the client side, but also on the server side. Figure 4.20a shows the loss of the server-side model after each round. It highlights the improvement of the MAPE from 80% to lower than 10% over the training period. Furthermore, the Mean Squared Error (MSE) is also tracked throughout the rounds. The MSE development can be noted in Figure 4.20b. Analogous



(a) MAPE Loss Round 5 Client 1

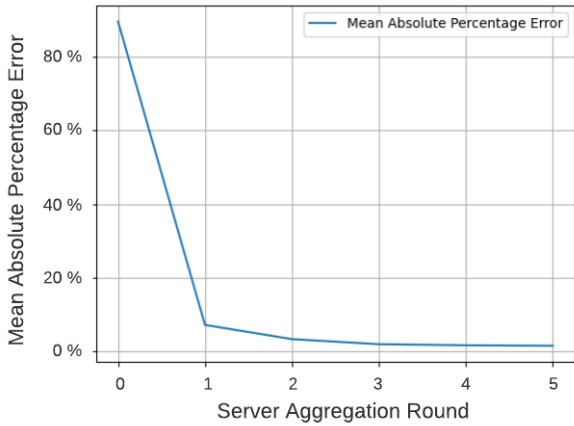


(b) MAPE Loss Round 5 Client 2


— Training MAPE — Validation MAPE

Figure 4.19: Mean Absolute Percentage Error (MAPE) of Clients in the Fifth Round

to the decrease of the Mean Absolute Percentage Error, the Mean Squared Error decreases similarly due to their likeness in computation. The final Mean Squared Error loss of the server-side model settles below 0.5 in round 5.



(a) MAPE Server-Side Loss per Round



(b) MSE Server-Side Loss per Round

Figure 4.20: Mean Absolute Percentage Error (MAPE) and Mean Squared Error (MSE)

The complete Federated Learning process is conducted using the Flower framework [8]. Flower is a framework developed by the universities of Cambridge and Oxford. The framework aims to enable researchers to conduct large-scale FL experiments. Additionally, Flower supports, as one of the few frameworks, both TensorFlow [96] and PyTorch [149] models. The Flower developers further claim that once the model is implemented, the migration from simulated to real devices works seamlessly. To build the neural network shown in Figure 4.17, TensorFlow with Keras [21] is used. The hyperparameters chosen for the model are discussed in the text above or can be viewed in Table 4.5.

Parameter Name	Parameter Value
<i>Train-Test Split</i>	80:20
<i>Train-Validation Split</i>	90:10
<i>Number of Clients</i>	2
<i>Number of Features</i>	66
<i>Optimizer</i>	Adam
<i>Learning Rate</i>	0.001
<i>Beta 1</i>	0.9
<i>Beta 2</i>	0.999
<i>Loss</i>	MAPE
<i>Batch Size</i>	32
<i>Epochs</i>	40 (round $< 3 \rightarrow 10$)
<i>Rounds</i>	5

MAPE = Mean Absolute Percentage Error

Table 4.5: Training Parameters for Federated Learning

4.5 Web-solution

The last contribution of this thesis is a web-based interface through which companies can interact with the models developed above. This section depicts the components of the website before taking a closer look at the output produced.

4.5.1 Architecture

Figure 4.21 shows the general components of the mobile responsive website. The left box represents the application’s Frontend, developed using Meta’s React Framework [100]. Instead of Javascript, Typescript is used to enforce type specifications on functions and variables [9]. The user interface with which a user interacts consists of three web pages. The landing page is the *Homepage*, where the user receives general information about cyber security. The second page depicted in Figure 4.21 is the *RCVaR* page. It is where the user can interact with the models to compute the expected cost and the Cyber Value at Risk. Upon completion of the input specifications, a GET Request is sent to the Backend, which returns all necessary information to display the cost in the evaluation part of the *RCVaR* page. The graphical output of the cost and Cyber Value at Risk information is explained in Section 4.5.2. The last view of the application is the *Documentation* page. Here, the user finds a short description of how to use the tool as well as how to interpret the risk measure and the predicted cost. Furthermore, the semi-synthetic data set covered in Section 4.4.1 is available for download. Last but not least, the *Documentation* page

gives an overview of the most critical industry reports presented in the related work part of this thesis (*cf.* Section 3).

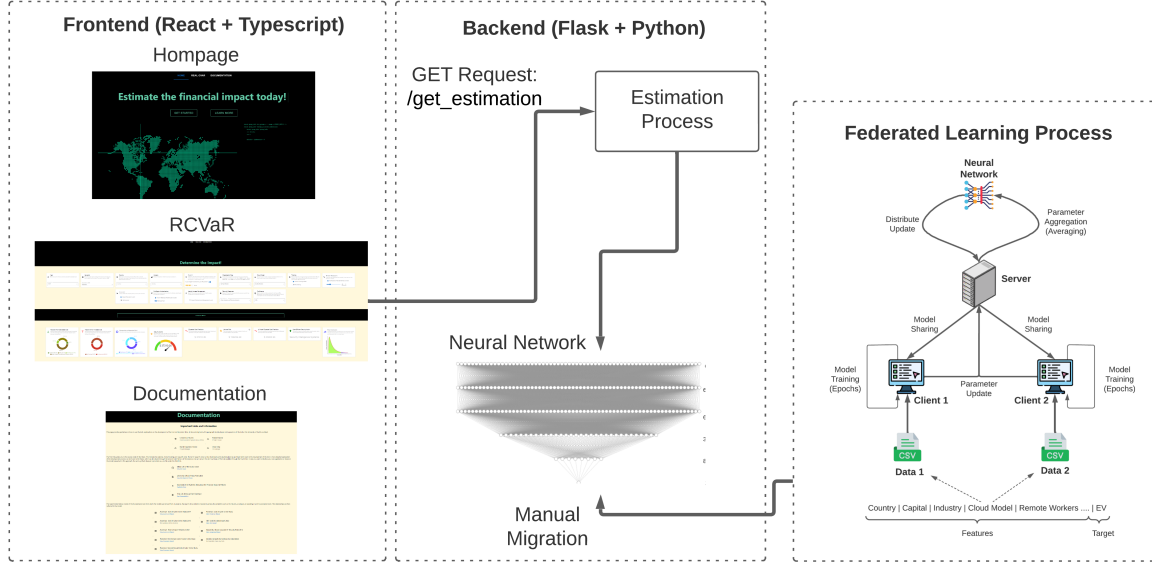


Figure 4.21: High-Level Components of the RCVaR Website

The Backend of the application is developed in Python [46] using the Flask [55] framework. It is a straightforward Backend architecture with only one API-Endpoint, which receives the business characteristics as input and returns the results of the models. The *Estimation Process* in the Backend calculates the predicted costs on the fly using the information from the industry reports as well as the neural network. To predict the expected cost based on the neural network, one has to one-hot-encode (*cf.* Section 4.4.2) the categorical values and bring them in the correct order before a forward propagation through the network is conducted.

The neural network was developed using the Federated Learning approach explained in Section 4.4.3 and 2.3. For instance, the figure representing the Federated Learning process is described in detail in Section 2.3. Overall, the key takeaway of the graph is that the training process is separate from the website infrastructure. Therefore, the final model from the simulation process from Section 4.4.3 needs to be migrated manually to the website's Backend. This process could be automated in a future version; respectively the Federated Learning process could be included in the main web application (*cf.* Section 6.1).

4.5.2 User Interface

This section focuses on the *RCVaR* page presented in the previous chapter. In the web page's first section, the user must enter company-specific characteristics. These charac-

teristics include the equity valuation, the location, or the industry of a company. Furthermore, the year for which the costs should be predicted needs to be inserted as input. The user has the option to select any year from 2012 to 2025. This was done to avoid wrong conclusions on years too far off from the range 2017-2020 since most data for the prediction stem from this time frame. In addition to the company-specific values, a confidence probability must be chosen, which represents the α percentile of the Cyber Value at Risk. The default value for confidence is set to 95%. For example, a confidence of 97.5% means that the probability of experiencing costs higher than the Cyber Value at Risk is 2.5%.

A complete list of all inputs, excluding the year, the capital, and the confidence, can be found in Table 4.1 in Section 4.2.3. In Figure 4.22 below, one can view four examples of how to enter business characteristics. The card in the top left corner shows the selection of security measures of a company. Since a company can have multiple cyber defense mechanisms, multiple instances can be selected. The *Country* or *Industry* factor inputs behave similarly. On the top right corner, one can input the trust or the defense capabilities of a company's supplier by rating the security within a range of 0 to 5 with a step size of 0.5. If the company has no connection to the supplier's IT systems, then the switch can be disabled, which signals the website to use the default value *None*.

Besides the year, capital and confidence, no input is required. Therefore these non-mandatory values have the default value of *None* initially. Using the *None* value means that the factor will not influence the average expected cost. Or expressed mathematically, the unspecified factors are set to zero in Equation 4.7. Therefore, binary inputs such as the existence of a multifactor authentication or the presence of cyber insurance require a check box. This can be viewed in the bottom right corner of Figure 4.22. When the check box is checked, the user can choose whether the *Multifactor Auth* factor should be true or false, respectively, if multifactor authentication is deployed or not. When the box is unchecked, the factor is set to *None* internally and the switch button disappears.

The last input card visible in Figure 4.22 shows how the user can enter a firm's percentage of remote workers into the model. To do so, the slider has to be moved to the correct position. Like other input cards, the default value *None* is set for the *Remote* factor upon disabling the check box in the middle of the input card.

The screenshot displays the RCVaR Web-Interface for four input factors, organized in a 2x2 grid:

- Security Measures:** A section titled "Security Measures" with a checkmark icon. It asks to "Select all IT security measures which the company currently has fully deployed/operational". A search bar contains "securitymeasure" and "Cyber Analytics and Behavior Analytics". Below it, a list shows "Automated Checks & ML and AI" (unchecked) and "Cyber Analytics and Behavior Analytics" (checked).
- Supplier:** A section titled "Supplier" with a shopping cart icon. It asks to "Rate the IT security and the trust in the company's suppliers. If the company has no supplier or their IT systems are not connected to the supplier IT: Set the switch to False". A toggle switch for "Is your Supplier connected to your IT systems?" is turned on. Below it, a rating of four yellow checkmarks and one grey shield is shown, labeled "Good".
- Remote Employees:** A section titled "Remote Employees" with a person icon. It asks to "Enter how many employees work remotely. Check the box if the amount is unknown." A checkbox "The Number of Remote Workers is known" is checked. Below it, a slider shows a value of 40, with a percentage sign and "40 %" indicated.
- Multifactor Authentication:** A section titled "Multifactor Authentication" with a QR code icon. It asks to "Specify if the company uses some kind of multifactor authentication to protect their IT systems." A checkbox "I know if Multifactor Authentication is used" is checked. Below it, a toggle switch for "Multifactor Auth" is turned on.

Figure 4.22: RCVaR Web-Interface for Four Input Factors

After all the inputs are specified, the user clicks the *Evaluate* button, which prompts the model to predict the cost for the company. The first few outputs are visible in Figure 4.23. It shows how the customization parameters listed in Table 4.1 influence the average cost. As one can observe on the output card on the top left in Figure 4.23: The supplier safety, Cyber and Behavior Analytics systems, multifactor authentication, and the percentage of remote workers had a positive influence on the cost. Positive in this sense means a reduction of the predicted cost. The percentage sign associated with the parameters reflects how much of the reduction is caused by that parameter. For instance, the Cyber and Behavior Analytics system contributed most towards the total reduction of all the customizing parameters. However, it does not mean that the costs were reduced by 50.7% due to Cyber and Behavior Analytics systems. The output card on the right top works similarly, with the slight difference of focusing on the cost increase. These two outputs allow firms to compare characteristics and prioritize business or security actions. For example, a company should focus first on training its employees before considering a cyber insurance since an untrained staff increases the expected cost more than incomplete insurance coverage. In the next step, a consideration of the benefits of these two actions against their costs needs to be conducted. The customized costs of cyber insurance and cyber awareness training are publicly available and are therefore not included in this thesis.

Furthermore, a firm can compare its cyber resilience against a customized benchmark. The output card in Figure 4.23 compares the current cost against the cost if no or all security measures are implemented. Security measures in this context mean all non-core business characteristics. More specifically, the following factors are considered in the benchmark: *Training*, *Insurance*, *Multifactor Auth*, *Identity Access Management*, and *Security Measures*. All their parameters can be viewed in Table 4.1. It is crucial to remember that this scoring does not represent risk. Instead, the benchmark is the same

company but with different security measures. It gives a firm an indication if costs have the potential to be reduced. It further gives them a reference point whether the expected costs returned by the system are rather high or low.




Figure 4.23: Web-Interface for Cost Decomposition and Factor Analysis

The last card in the bottom left of Figure 4.23 depicts how costs are composed. These numbers stem from the IBM report [27] and do not fluctuate significantly over the years. Nevertheless, it helps companies to locate where in the company the highest costs occur. Upon locating the problem, one can use this information to increase resilience in a particular area. For example, it becomes clear that the highest cost contributor is lost business revenue (42%). As known from the output card on the top right, the primary increasing cost driver is the lack of training (50.7%). Therefore, a reasonable action to increase cyber resilience is to provide training to the staff closest connected to the IT systems which enable daily operations. In a retail setting, this could apply to the salesmen and women who gather customer information upon checkout.

Figure 4.24 highlights the numerical output of the different models. The card in the

upper left corner shows the predicted cost by Equation 4.7 and therefore represents the “exact” cost. On the other hand the value in the bottom left corner shows the expected loss estimated by the AI model. As one can observe, the two values are at a reasonable distance from each other, indicating a good-performing Federated Learning process (*cf.* Section 5.3). Figure 4.24 furthermore depicts the Value at Risk for the company. The VaR represents the costs not exceeded by the realized costs with a certain confidence. It is noteworthy that the RCVaR has a constant variance as described in detail in Section 4.3.2. Consequently, the Value at Risk represents a linear transformation, namely a multiplication of 2.9 with the expected value in the case of 95% confidence. Nevertheless, the risk measure is a foundation for a security-investment-risk analysis during the planning management process (*cf.* Section 3.4.2). The last card in Figure 4.24 indicates the parameter from the factors *Training*, *Insurance*, *Multifactor Auth*, *Identity Access Management*, and *Security Measures* that is not yet implemented and that would have the highest effect on reducing the average cost.

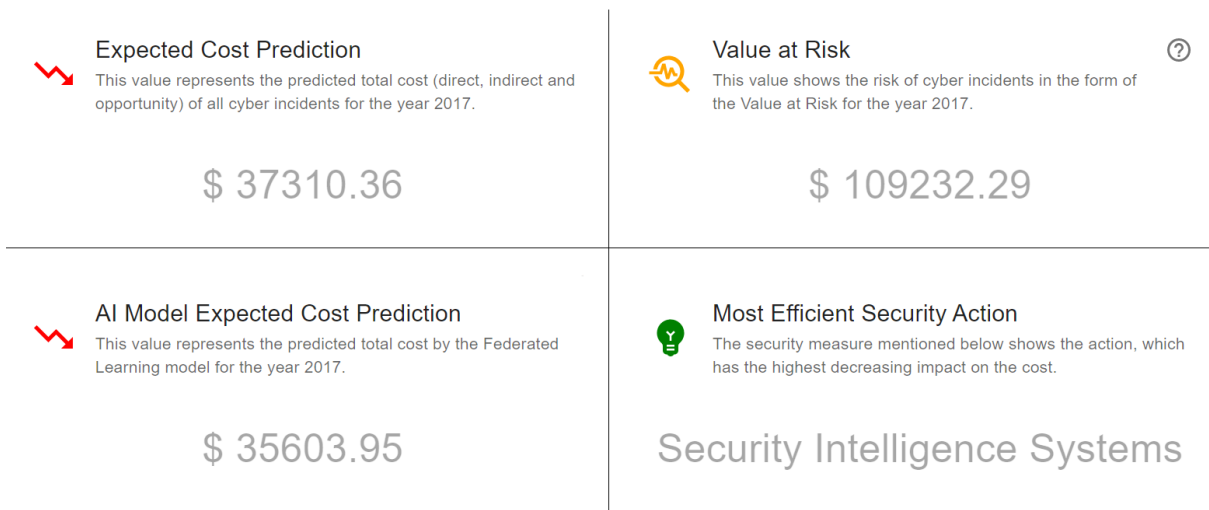


Figure 4.24: Web-Interface for Numerical Output of Cost Analysis

Besides the outputs presented here, the website further displays the customized cost distribution with appropriate coloring for interpretation (*cf.* Figure 4.15). All the outputs are presented in an easily interpretable manner, with a short explanation available on each card. This provides Small and Middle-Sized Enterprises with the opportunity to assess their possible cost together with their risk. The website further provides insights on how to reduce the expected cost. The following steps after the RCVaR analysis are to research the cost of the most efficient security action. These costs vary across countries and industries and are not included in this thesis. After reaching a good pricing overview of the offered products and actions one can come back to the website and assess the total costs (investment costs + cost predictions) against the risks and the cost-benefit of the action. When using the results with careful consideration of the model’s assumptions, the Real Cyber Value at Risk (RCVaR) can be a valuable addition to the cyber risk management process (*cf.* Section 5).

Chapter 5

Evaluation

This chapter assesses whether the three main objectives of this thesis have been achieved. To investigate the extent to which the cost estimation model reflects real-world behavior, a quantitative and qualitative evaluation is conducted. To assess the accuracy of the risk measure, the cost distribution is compared to that of related research. Finally, the Federated Learning model is assessed by comparing the performance of different models.

5.1 Cost Estimator

First, the model is evaluated by applying it to multiple theoretical companies and comparing their predicted losses. This approach allows assessing the model's ability to replicate real-world behavior as indicated by industry reports. In the second part of the chapter, the estimated financial loss obtained using the RCVaR is compared with actual costs in two use cases.

5.1.1 Qualitative Evaluation

To evaluate the RCVaR qualitatively, the cost behavior of its results is compared to the loss patterns found in the industry reports (*cf.* Section 3.2). Table 5.1 illustrates a selection of hypothetical firms from the data generation process (*cf.* Section 4.4.1). The columns show the firm's specifications that were used as input into the Real Cyber Value at Risk model. The output of the RCVaR can then be viewed in the last two columns to the right. The colors in the cells indicate whether the parameter influences the average cost for a particular capital and year positively (reduces the cost → Green) or negatively (increases the cost → Red). If no specification is provided to the Real Cyber Value at Risk, the cell is colored orange.

Upon examining Table 5.1, it is clear that all firms, except Firm 13, have the same equity value (*i.e.*, capital). This uniformity in market capitalization allows for the independent study of the impact of cost factors (*cf.* Section 4.2.3). Additionally, a comparison between

Firm 13 and Firm 2 illustrates that the RCVaR effectively adjusts costs based on a company's financial capital (*cf.* Section 4.2.1). This phenomenon has been well documented in previous research [36, 102, 153].

Similar to the financial size of a company, costs also tend to increase over time. When inspecting Table 5.1, one notices that all firms predict the cost for the year 2022, except Firm 14. This company can be contrasted with Firm 5, which has the same business characteristics but much higher costs. This stark difference in costs for different years aligns with the findings of the Accenture [2, 3] and IBM [27] reports, which both state double-digit cost growth during their research period.

As previously discussed in Chapter 4.2.4, the *Country* factor has a significant impact on costs. Examining the two banking firms, Firm 3 and Firm 15, it becomes apparent that the German company has approximately \$ 30'000 higher costs than the French company. This trend is also reflected in real-world data according to IBM [27] and Accenture [2, 3]. A similar pattern can be observed when comparing Firm 7 and Firm 16, which both operate in the healthcare sector: The US-German pharmaceutical company has significantly higher costs, as estimated by the RCVaR model, compared to the German enterprise. In fact, related literature [27, 2, 3, 72, 43] supports the notion that North American companies tend to have higher cyber costs.

The second highest cost contributor is the number of employees. A comparison of Firm 6 and 7 highlights the difference between a *Large* and *Medium* sized company. If other factors such as *Country* also positively influence average cost, the predicted loss can be reduced to a tiny amount, as can be seen in the case of Firm 4. The size of a company has also been identified as an essential cost contributor in multiple research papers [36, 102, 2].

Similarly, the *Industry* factor has also been well-researched in related literature [36, 102, 2, 3, 27, 126, 128, 139]. In particular, financial institutions tend to have high cyber attack costs [36, 153]. This situation is also reflected in the output of the RCVaR model for Firm 16 and Firm 15 in Table 5.1. Other relationships, such as the higher cost of industrial manufacturers compared to retailers [102] (Firm 1 vs. Firm 2) or the significant impact of attacks on healthcare providers, are also reflected in the comparison provided in the table.

In general, the companies that have better security measures tend to have lower expected costs [36, 27, 2, 31]. For instance, most companies with three security measures deployed have costs below \$ 40'000. Additionally, it is noteworthy that Firm 1 is the only *Large* company in the table with costs below \$ 100'000. Comparing Firm 7 and Firm 5 also illustrates the influence of the cloud model and the percentage of remote workers on the cost estimation by the RCVaR model. The output suggests that firms with a smaller remote workforce and a hybrid cloud model tend to sustain lower costs due to cyber attacks. This realization aligns with the argument presented in the IBM [27] report, which shows this exact correlation between the cloud model, remote workers, and costs.

	Year	Capital	Country	Industry	Number of Employees	Cloud Model	Percentage of Remote	Supplier	Employee Training	Cyber Insurance	Multi-factor Authentication	Identity Access Management	Deployed Security Measures	Expected Cost	AI Expected Cost
Firm 1	2022	\$ 70 M	GER	Retail	Large	Public	40%	0	1	1	1	-	Risk Management Advanced Perimeter Controls Data Loss Prevention Measures	\$ 98'209	\$ 96'623
Firm 2	2022	\$ 70 M	CAN GER	Industrial	Large	-	20%	0	-	0	1	-	-	\$ 137'646	\$ 128'690
Firm 3	2022	\$ 70 M	FRA	Banking	Large	-	20%	1	1	1	1	1	Encryption Technologies	\$ 129'780	\$ 122'398
Firm 4	2022	\$ 70 M	ESP	Life Sciences	Micro	Private	20%	0	0	0	0	0	-	\$ 6'390	\$ 6548
Firm 5	2022	\$ 70 M	US	Insurance	Medium	Hybrid	20%	1	1	1	1	1	Cyber Analytics and Behavior Analytics Advanced Identity and Access	\$117'680	\$109'639
Firm 6	2022	\$ 70 M	US CAN	Health High Tech	Large	-	100%	1	0	0	1	1	Data Loss Prevention Measures	\$ 473'101	\$ 288'054
Firm 7	2022	\$ 70 M	US GER	Health Pharmaceuticals	Medium	Public	80%	1	0	0	1	1	Security Intelligence System	\$ 280'245	\$ 220'435
Firm 8	2022	\$ 70 M	UK SCA TUR	Energy Industrial	Large	Public	20%	0	1	1	0	-	Risk Management Security Intelligence System Data Loss Prevention Measures	\$ 15'547	\$ 16'168
Firm 9	2022	\$ 70 M	ITA	Consumer Goods	Medium	Private	20%	0	0	1	0	-	Sufficient Security Staff Security Intelligence Systems Incident Response Plan Testing	\$ 36'116	\$ 34'939
Firm 10	2022	\$ 70 M	CAN	Retail	Small	-	20%	0	-	0	1	1	-	\$ 44'556	\$ 43890
Firm 11	2022	\$ 70 M	GER	Education	Small	-	20%	-	1	0	1	0	-	\$ 28'676	\$ 28'475
Firm 12	2022	\$ 70 M	US	Public Sector	Medium	Private	20%	-	1	0	1	1	Automated Checks & ML and AI	\$ 79'337	\$ 75'155
Firm 13	2022	\$ 34 M	CAN GER	Industrial	Large	-	20%	0	-	0	1	-	-	\$ 71'290	\$ 64'772
Firm 14	2014	\$ 70 M	US	Insurance	Medium	Hybrid	20%	1	1	1	1	1	Cyber Analytics and Behavior Analytics Advanced Identity and Access	\$ 51'064	\$ 47'349
Firm 15	2022	\$ 70 M	GER	Banking	Large	-	20%	1	1	1	1	1	Encryption Technologies	\$ 157'214	\$ 144'155
Firm 16	2022	\$ 70 M	GER	Health	Large	Public	40%	0	1	1	1	-	Incident Response Plan Testing Data Loss Prevention Measures	\$ 149'360	\$ 140'951

● Cost Increasing Factors, ● Cost Decreasing Factors, ● No Input Provided and Therefore No Influence on Cost

Table 5.1: Hypothetical Companies and Their Associated Costs

5.1.2 Quantitative Evaluation

As part of the evaluation, the output of the cost estimator model is compared to cost predictions found in related literature. When evaluating the model, the most prevalent problem one faces is the lack of data. The few public cost numbers are anonymized, thus, the link between business characteristics and loss prediction is difficult to establish. In the scope of this quantitative evaluation, the focus lies on two reports. Important to point out is that both reports have not contributed any number to the development of the cost estimator. Therefore, these numbers can be viewed as Out-of-Sample test data, *i.e.*, the model has not seen the data previous to this test run.

The first cost estimation number found in the literature stems from the Kaspersky report [72] of the year 2013. Very similar to the methodology of the Accenture reports [2, 3], Kaspersky conducted roughly 3000 interviews with IT specialists familiar with both IT-security as well as the business process in their companies. The total sample consisted of 2364 companies, which were divided into two groups depending on their amount of computerized workplaces. Companies with less than 1500 workplaces are labeled as SMEs, while the others belong to the group of large corporations. One conclusion based on the survey was that the average loss of Small and Middle-Sized Enterprises due to a cyber incident is \$ 50'000 (*cf.* Section 3.2).

Given the Venture Capital data from Pitchbook used to generate data, it is possible to approximate the average equity value of an SME for the year 2022. Combining this information with the year the Kaspersky study was conducted, all the information necessary to calculate the expected yearly cost is given. Upon entering the target year 2013 with a market capitalization of \$ 168 million, the model discounts the equity to the year 2017 (*cf.* Section 4.2.2) before converting it to cost (*cf.* Section 4.2.1) and further discounting them to the desired year. The output of the RCVaR model results in costs of roughly \$ 70'000. This amount can also be observed in Table 5.2. There, it is highlighted that the estimation deviates from the true value (Absolute Percentage Error) by 37%.

It is important to be cautious when interpreting this number. An error of 37% can be viewed as inaccurate. Nevertheless, it is crucial to point out that the \$ 50'000 of the Kaspersky report portray the average bill per serious incident, while the cost prediction of the RCVaR model represents the total annualized incidents costs. Given that companies have multiple severe incidents per year on rare occasions, annualized costs of 70'000 seem plausible. Furthermore, it is noteworthy that the VC equity input might lead to overestimating the cost in this specific instance since Venture Capital valuations spiked in the first quarter of 2022 and reached a record high. If the value is smoothed with the average over the past 12 months, the equity approximation of SMEs is equal to \$ 134 million. After re-running the estimation with this market capitalization value, the RCVaR results in a yearly cost estimation of \$ 55'728, which is very close to the value of the Kaspersky report [72].

In the past two years, multiple studies [158, 36] focusing on econometrics have researched the cost of cyber attacks based on insurance risk premium data. Theoretically, the pre-

Cost Source	Cost Estimation of Source	Cost Predicted using RCVaR	Absolute Percentage Error
<i>Kaspersky Report [72]</i>	\$ 50'000	\$ 68'786	37 %
<i>Woods et al. [158]</i>	\$ 428'000	\$ 420'444	2 %

Table 5.2: Unseen Cost Estimation Vs. RCVaR Cost Prediction

mium paid by the insuring company to hedge against cyber incidents reflects the firm's belief about future costs. The study conducted by Woods et al. [158] used pricing data from roughly 7000 observations across 26 insurance firms to derive a real-world cost distribution. Based on this distribution, the authors predict the price of a hypothetical retail company. Additionally, it is known that the insurance data consists of only US, primarily Californian, companies. Therefore, it can be assumed that the insured company is also based in the US. Another piece of information provided about the retail company is its yearly revenue of \$ 50 million. Given the Return on Equity (ROE) ratio, the revenue can be converted to equity. The Return on Equity (ROE) is a percentage measure representing the ratio of revenue in terms of the money invested to achieve this income [35].

Using this approach to determine market capitalization has two advantages. First of all, the Return on Equity is deeply anchored in the economic literature [35, 110], and therefore there exists a lot of researched statistical data for different regions and industries. Secondly, not determining the equity value through Venture Capital data gives insights into the robustness of the RCVaR model due to the strong dependence on the capital input in the model. Applying the ROE data from the prestigious NYU [110] to convert the revenue to equity, one receives a market capitalization of \$ 249 million for US-based general retail companies. Adjusting for inflation, today's equity value equals \$ 253 million. Providing the industry (Retail), the equity value (\$ 253 million), the location (USA) and the desired year (2021), which is the year the study was conducted. All the information needed to analyze the firm with the RCVaR is given.

The comparison of the result from the study presented in [158] and the output of the Real Cyber Value at Risk shows that they are very close together. In contrast to the Kaspersky report [72], both numbers resemble the annualized total cyber incident costs. It can be inferred that the more information is provided to the RCVaR, the more accurate the prediction is. Furthermore, it shows that the output has a real-world connection and can approximate the cost numbers of two reports with different methodologies, regions, years, and information sources.


In summary, the RCVaR yields accurate predictions when tested on "unseen" data. Even if the input information is scarce, the Real Cyber Value at Risk still provides the user with an estimation in the correct area. The more information is entered, the more accurate the model becomes. Furthermore, both evaluation experiments show that the Real Cyber Value at Risk is rooted in the real world and is not just of theoretical nature. Based on the information extracted from seven industry reports [27, 2, 3, 10, 126, 127, 128], the proposed Real Cyber Value at Risk model achieves highly accurate cost predictions on theoretical and real-world firms.

5.2 Risk Measure


In this chapter, the risk measure developed within the context of this thesis is examined. While there is little empirical data available on the Cyber Value at Risk (CVaR), one source that exists is a report by Deloitte [157] which studied the Value at Risk of cyber incidents in the Netherlands for various industries. However, due to limited information on business characteristics and the inability of the RCVaR to make predictions for Dutch companies, a different approach has been chosen.

The loss distribution is evaluated by comparing the cost distribution of the General Inverse Gaussian (as proposed in this thesis) with the cost distributions of related research. According to [37], both Log-Normal and Skew-Normal distributions provide a good fit for the actual distribution, though other distributions may also perform well. But it is considered inevitable that the financial impact of cyber incidents follows a heavy-tailed distribution, as supported by [78, 158].

Other distributions mentioned in the literature include Power-Law, Pareto, and Weibull. It is worth noting that none of these distributions could be rejected through the Kolmogorov-Smirnov test conducted in the scope of this thesis (*cf.* Section 4.3.2). The Pareto distribution even had the fourth highest probability after the General Inverse Gaussian. Given the similarity between these distributions, it can be assumed that the risk measure of the RCVaR is a good approximation of real-world circumstances. The main difference between distributions in cyber economics research is whether or not the cost density at the origin is zero. Since this thesis considers all incident costs, including *Indirect* and *Opportunity* costs, it is assumed that the costs are never zero, even with no successful attack occurring. As a result, bell-shaped curves such as the Weibull or Skewed-Norm distributions, which are mentioned in [78, 158], are considered to be a better fit.



(a) Cost Distribution Based on MARSH [93]



(b) Cost Distribution of the RCVaR

Figure 5.1: Cyber Cost Density Distributions MARSH [93] Vs. RCVaR

In Figure 5.1a, one can see the cost distribution generated by MARSH [93] using Monte Carlo Simulations (*cf.* Section 2.2). MARSH already utilizes the CVaR derived from this distribution to advise their clients on cyber security planning. On the hand, Figure 5.1b depicts the distribution of the RCVaR derived in the scope of this thesis.

Upon comparing these two graphs, it becomes clear that they have similar shapes. The only difference is that the RCVaR has density zero at the origin, which can be explained by the different definitions of cyber attack costs. Based on the matching distributions in the literature and the cost distribution of the RCVaR, it can be concluded that the RCVaR is relatively closely approximating the distribution of real-world costs.

5.3 Federated Learning

This section focuses on the evaluation of the Federated Learning model introduced in Section 4.4.3. First, the performance numbers of the deployed model are explained. Then, the same model is trained using different data splits to simulate real-world training among clients and compare the performance of these models to the deployed model. Finally, the results of the original FL are shown in contrast to the output of two centralized model architectures. The overall conclusion drawn from this chapter is that neither the different data splits nor the centralized training approach result in significantly different performance measures compared to the original model.

5.3.1 Quantitative Evaluation of the Neural Network

The center of attention in this chapter lies on the originally trained Federated Learning model. It is trained on two halves of the dataset generated in Section 4.4.1. To establish the two halves, each observation is assigned randomly to a client until both have the same amount of observations. After distributing the data among the clients, the FL is conducted as follows: The server shares the central model with the clients, which train the model on their half of the data. Then, the parameters are sent back to the server, where they are aggregated before sharing the updated model again. This circular process is repeated multiple times until the model converges. The final result is the model with the aggregated weights on the server.

Figure 5.2 depicts the test subset of the generated data. Each observation in the figure consists of 14 dimensions: The expected value, capital, year, and all other 11 factors. When plotting the observations by the capital dimension on the x-axis and the target variable, namely the true expected cost, in logarithmic scale, on the y-axis, one receives Figure 5.2.

The coloring indicates the Absolute Percentage Error (APE) of each observation. The Absolute Percentage Error, in this case, is the absolute deviation of the model estimation in percentage points of the actual value. More concretely, blue means that the predicted cost deviates more than 10% from the cost predicted by the function in Section 4.2.5. Green means that the deviation is less than 10%, and an orange coloring indicates that the APE is less than 5%. As Figure 5.2 highlights, the model is more accurate in predicting higher expected costs than smaller impacts. This circumstance is likely attributed to an imbalance in the total generated data since less than 5% of the generated data has an expected financial loss of less than \$ 1'000. This imbalance is not problematic since costs of less than \$ 1'000 are a rare occurrence because it would require a firm to deploy many security measures while having shallow equity reserves. Consequently, it is unlikely that poorly capitalized firms will invest in many IT-security products simultaneously. Related research, specifically the Accenture report [10], supports this hypothesis. Nevertheless, one could improve the model's accuracy in the low cost area by producing more semi-synthetic data with the existing process or by developing a new data generation process with less imbalanced data.





Figure 5.2: Differences Between the Predicted and the True Values on the Test Sample

The overall percentage of test observations with a deviation of less than 10% is 78.11% in Figure 5.2. By achieving a training Mean Absolute Percentage Error of roughly 2.5% and a test APE of below 10% for the majority of observations the model can be considered as good [148]. Upon investigating Figure 5.2 further, it becomes evident that no observation accomplished an Absolute Percentage Error on the test set of below 5%.

When observing Figures 5.3a and 5.3b, one can also view the ability of the model to learn the parameter ratios from Equation 4.5. In both figures, one line represents the true target costs from Equation 4.7, while the neural network model creates the other cost line. The lines of the model are generated by varying one factor of an observation while keeping the others the same. For instance, Figure 5.3a is produced by inputting nine observations, which are all the same except for their *Country* factor, into the FL model. If the neural network would exactly predict the correct costs for the different countries, the two lines would lie perfectly on top of each other. Or formulated differently, the distance between the lines represents the error, while the similarity in the line shape reflects if the correct relative importance between the parameter ratios is learned. Despite the error, both Figures highlight that the model learns the relative relation of the parameter ratios well due to their similar shapes. The only constant that is not learned well is the *Insurance* parameter in Figure 5.3b.

To increase the model's overall performance, it must be trained on larger datasets. A preliminary experiment conducted in the scope of this thesis shows that the set of observations with a deviation of less than 10% increases to 91.02% (*cf.* Appendix A.2) when training the model on a data set of 500'000 observations. Nevertheless, even when training

on a larger dataset, no test observation achieves an APE of below 5%. It can be assumed that the training set needs to be increased even more to achieve such test errors. Since the main focus of this thesis is not on developing the best neural network to approximate the RCVaR equation, training the model on more data is out of the scope for this thesis. Furthermore, the training time increases dramatically with larger data samples. Hence this thesis's primary model is trained on 100'000 observations.



(a) Cost Deviations Along Country Dimension (b) Cost Deviations Along Industry Dimension

—●— Cost Predicted by RCVaR Cost Predictor (Equation 4.7) —▲— Cost Predicted by Machine Learning Model (ANN)

Figure 5.3: Cost Accuracy Along Factor Dimensions (FL/Random Split/40 epochs)

To conclude this section, the Federated Learning approach manages to learn the cost function established in Section 4.2.5 well. An Absolute Percentage Error of less than 10% for the majority of the test samples is considered good [148] even though no observation in the test sample broke the 5% mark. The overall performance can be increased when training the model on more data. However, it remains questionable if test observations with a low MAPE can be achieved consistently.

5.3.2 Comparison of Different Data Splits

As shown in Section 5.3.1, the model, trained with two clients, performs well. Previous to the training, the observations are assigned randomly to the clients to create the different datasets. Since the data consists of 100'000 observations and the combinations of factors are limited, both clients likely have similar combinations in their data halves. Hence, if the data between the clients is not vastly different, the whole model could also be trained just with one client, which would be equal to centralized learning. To prove that a genuinely Federated Learning process can learn the cost function defined in Section 4.2.5, one has to

show that similar performance can be achieved in a training environment where specific parameters only occur in a single client.

The new model’s performance with the same architecture and hyperparameters, as described in Section 4.4.3, can be viewed in Figure 5.4. The only difference to the previous training process is that the data split is not conducted randomly but according to simple rules to ensure heterogeneous data halves. Figure 5.4a shows the model’s performance trained on data that is split according to the *Country* factor. The reasoning behind this approach is that Small and Middle-Sized Enterprises often are not multi-national firms. Therefore, they will not have local cost data available for different locations to train the model. To simulate this behavior, the data is divided so that the first client only has observations from the USA, UK, France, and Italy. While the second client contains firm data for Canada, Spain, Turkey, and Scandinavia.

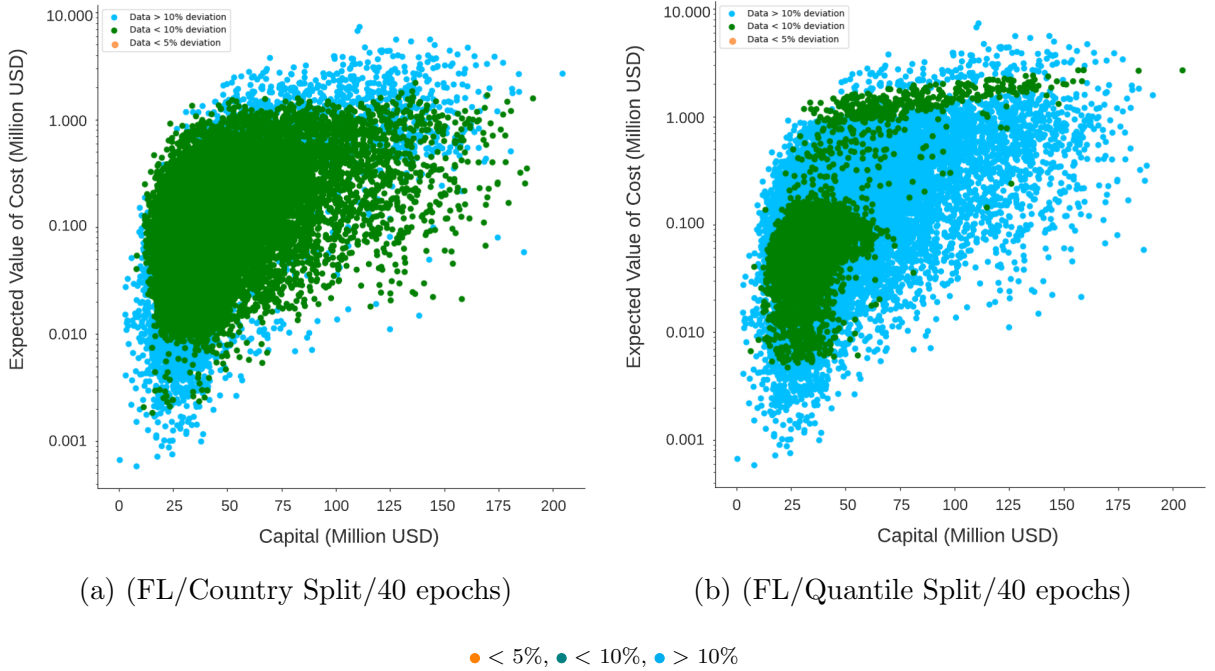



Figure 5.4: Differences Between the Predicted and the True Values on the Test Sample

As expected, the model trained on the country-split data achieves a very similar performance as the original FL model from Section 5.3.1. With roughly 85% of all test observations having a deviation of less than 10% from the true value, the model’s prediction of Figure 5.4a outperforms the original model. The country-split model can especially achieve higher accuracy in the low-cost environment in the lower left part of the graphic. This difference in performance compared to the original model is likely caused by the way the data is split. When looking at the parameters of each client in more detail, one recognizes that client 2 possesses countries associated with relatively lower costs. Therefore, client 2 focuses specifically on the low-cost area during training, whereas in the previous data split, these observations were divided among the clients. A second explanation could be that, despite the expectation previous to the country-split experiment, the clients’ data distributions after the country split are “more” similar to each other and the true

data distribution. As suggested by [109, 162, 120], the performance of models could vary based on how accurately the clients' distributions reflect the overall data structure when training on data that is not Independent and Identically Distributed (IID). Since the data generation process in Section 4.4.1 generates the capital based on a restricted normal distribution and the factors with varying probabilities, the observations do not have the same probability of appearing in the generated data set. Hence, they can not be described as Independent and Identically Distributed. Therefore, it can be argued that the original split creates distributions in the clients that differ more from the overall density structure than the country-split distributions. Such a minor difference in data distribution could lead to reduced performance compared to the centralized model [120]. In other words, despite the previous expectation, the client's distributions from the country split could resemble the actual distribution of the data more closely and, therefore, could lead to higher accuracy, as visible in Table 5.3. Nevertheless, the higher performance can also be due to the randomness of shuffling the training data. Therefore, more simulations are needed to paint a complete picture.

The influence of IID data in Federated Learning can also be observed in Figure 5.4b. It depicts the performance of the FL model trained on data split along the capital dimension. Meaning the observations are either assigned to client 1 if their capital is in the lower half of all equity values or to client 2 if their value is higher than the 50% quantile. Because the data is split across two clients, the 50% quantile is used as the dividing threshold. Alternatively, in other words, the threshold for two available clients is the median of all equity values. When looking at the training data of both clients in Figures 5.5a and 5.5b, it becomes evident that each distribution does not follow the shape of the proper dataset, as observed in Figure 4.16.



(a) Observations With Below Median Equity (b) Observations With Above Median Equity

Figure 5.5: Training Data Distribution of Capital Quantile Split

Based on the different local distributions in Figure 5.5, the aforementioned issue of Non-IID data becomes apparent. When looking at the model's performance on the data split along the capital dimension, it can be observed that the FL network predicts the cost

in the low capital area well. This is most likely because one client solely consists of low capital observation. Thus the FL model is more focused on this region during training. However, it is apparent from the numbers in Table 5.3 that there are only 75% green observations. Meaning that the model trained on the quantile split achieves the lowest amount of observations with a deviation of less than 10% from the true cost value of all FL models. When comparing the two models regarding the fixed factors *Country* and *Industry* (cf. Appendix A), it can be observed that the lines are very similar to the ones depicted in Figure 5.3. Depending on their accuracy, they only differ in the distance between the true and the predicted lines. Interesting in this regard is that the “quantile” model’s predictability of the *Country* parameter is much better than in all previous models. This could stem from the fact that the parameter ratios have a relatively higher weight when training on lower capital data. The complete line graphs can be found in Appendix A.

To summarize this section: The data split can make the difference between a good and an excellent performance. Hence Federated Learning should be able to reproduce the cost estimation function sufficiently no matter the data distribution among clients. Nevertheless, enough clients must participate so that the model can achieve outstanding results.

Model	> 10 %	< 10 %	< 5 %
<i>Federated Learning (Random)</i>	21.88 %	78.11 %	0.0 %
<i>Federated Learning (Country)</i>	15.02 %	84.98 %	0.0 %
<i>Federated Learning (Capital)</i>	24.86 %	75.14 %	0.0 %

Table 5.3: Absolute Percentage Error (Accuracy) Comparison of Different Data Splits

5.3.3 Comparison to Centralized Models

In the final evaluation step, the performance of the “decentralized” model is compared to a centralized approach. Meaning the same model is trained on the complete dataset on the server. Besides the epoch, all parameters stayed the same during the training process to allow an effective comparison between the two architectures. The epoch was altered minimally to elicit a good epoch number to ultimately compare the Federated Learning model to the “best” possible model trained with the complete dataset.

Figure 5.6a shows the predictions against the true target values of the centralized model trained with 50 epochs. The observations colored green indicate that their prediction deviates less than 10% from the actual value, whereas the light blue represents observations with a prediction which is more than 10% off. As can be observed, the performance is vastly the same compared to the Federated Learning model. A slight difference is that the centralized model has higher accuracy in the low-cost area. Since less data is available for this area, the FL model might lose some of its predictability during the aggregation in the data-poor region. As with the different data splits (cf. Section 5.3.2), the small difference in performance could also stem from the fact that the distribution of data among clients differs heavily. This circumstance is documented in the Federated Learning literature [109, 120, 162], which states that the performance of a “decentralized” model might be

worse if the local data distributions are not similar. The results from Section 5.3.2 further support this fact since different data split approaches lead to higher accuracy over the complete dataset.

Table 5.4 shows the accuracy of different models compared to the original model (top row) from Section 5.3.1. The best-performing Federated Learning model, which is the one trained on the country split, is also included in the second row. It becomes clear that there are only minor differences in performance between the centralized Artificial Neural Network and the Federated Learning models. For instance, none of these models achieves a deviation of less than 5% for any observation. Overall, the deviation percentages of less than 10% fluctuate within the range of 75% to 85%. Considering that a Mean Absolute Percentage Error of less than 10% is considered good to very good [148], all models are deemed capable of approximating the cost function of Section 4.2.5.

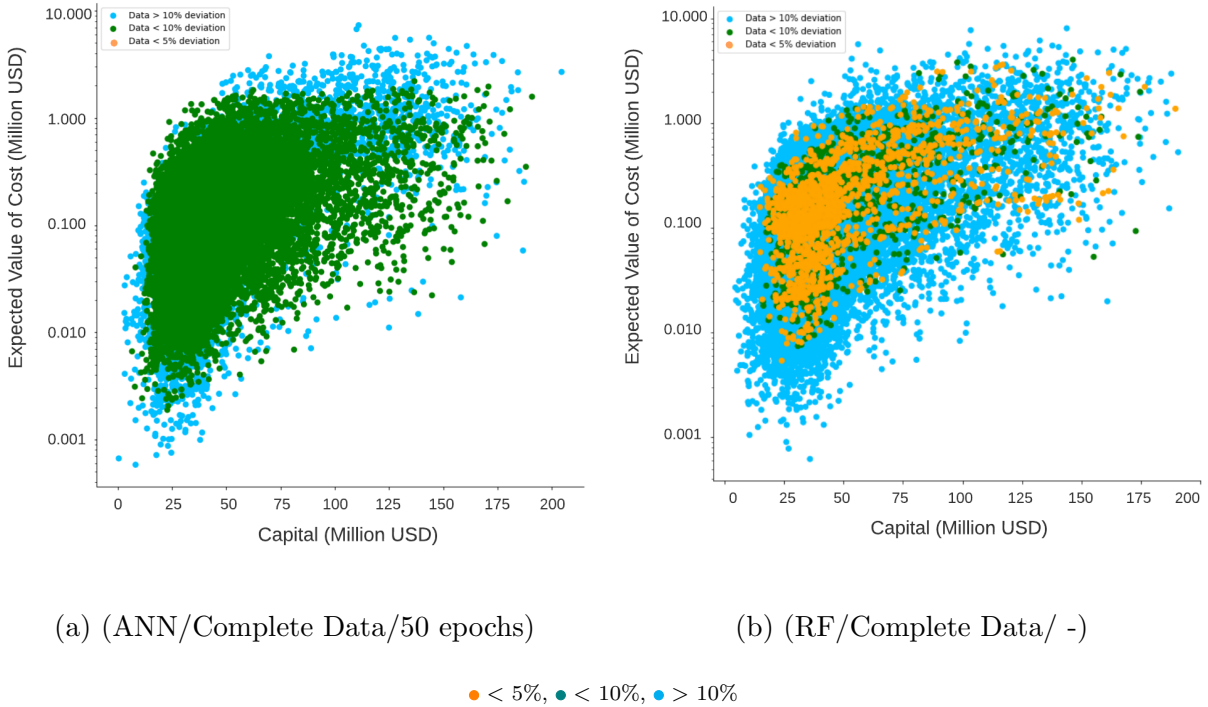


Figure 5.6: Differences Between the Predicted and the True Values on the Test Sample

The literature review conducted in Section 3.3 shows that the most promising and most frequently used ML architecture type to approximate risk or cost functions is the Artificial Neural Network. The second place takes the Random Forest (RF) architecture. A Random Forest consists of multiple decision trees whose output is aggregated [147]. In the case of cost estimation, the predictions of each single decision tree are aggregated by taking the overall average. In this thesis's scope, an RF model is built using Google's TensorFlow framework [94]. The tree-based model is configured with 500 individual trees with a depth of ten. When applying the Random Forest model on the semi-synthetic dataset, it yields the test performance portrayed in Figure 5.6b. The figure signals that the RF achieves very high performance on a few observations, while having an enormous deviation from the true value in the majority of cases (84%). The exact performance

statistics can be observed in Table 5.4. The differences in performance between the ANN and the RF architecture are likely related to the incapability of the Random Forest to capture trends and handle sparse data [88]. Since the data consists of many binary inputs and can be described as sparse, the RF model is not the best fit. Overall the RF model achieves a deviation of less than 5% for 8% of all test observations depicted in Figure 5.6b.

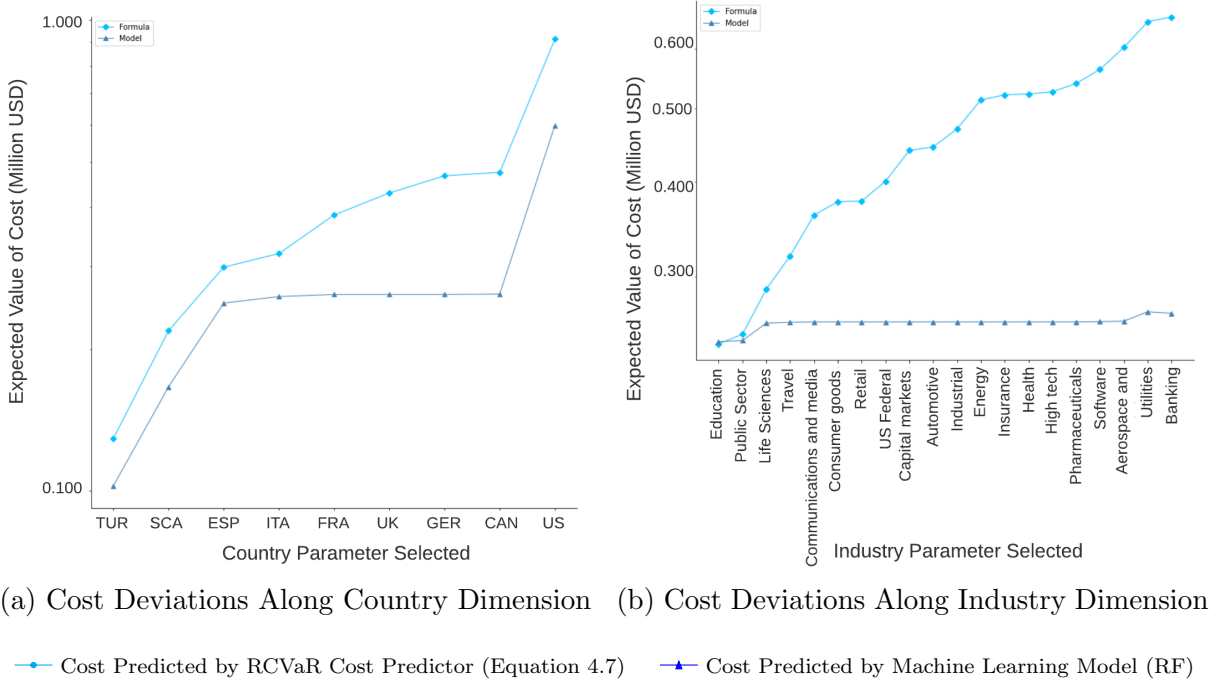


Figure 5.7: Cost Accuracy Along Factor Dimensions (RF/Complete Data/-)

The centralized ANN and RF performance can also be evaluated according to Figure 5.7. Since the lines for the neural network trained on the complete dataset look very similar to the ones of the Federated Learning model, they are omitted in this section but can be viewed in Appendix A. This section focuses on the RF's line charts, which are visible in Figure 5.7. It shows the predicted versus the true cost along either the country or industry dimension. To generate both plots, all other inputs, such as capital and year, are fixed. As can be seen in Figure 5.7a, the Random Forest manages to capture the relation between parameter ratios of the locations Turkey, Scandinavia, Spain, Canada, and the US relatively well. While it fails to reproduce the shape of the line between Italy and the UK. Nevertheless, the closest estimation with an APE of 21% (Turkey) is still far off from the true label. A worse picture is painted by Graph 5.7b, which shows the inability of the Random Forest to estimate the costs of firms among most industries correctly. Although most of the predictions are far off, there exist a few industries that have a very low error. For instance, the prediction for the firm for the education industry has only a deviation of 1% from the true label. However, starting from the parameter Life Science the APE starts to grow from 10% to almost 60%.

Overall, it can be concluded that the centralized Artificial Neural Network has a very similar performance to the neural network trained with Federated Learning. A look at

Table 5.4 even reveals that the FL model trained on data divided along the country dimensions achieves higher accuracy than the centralized model. Nevertheless, the differences between the central and the FL are minor, which indicates that Federated Learning is a viable solution to tackle the data scarcity problem in the cyber security economics field. Additionally, a centralized Random Forest model was evaluated. It achieved superior performance than all other models in a few observations while failing to perform well on the broader dataset. Hence it can also be concluded that the RF model is not well suited to be trained on the semi-synthetic dataset available on the thesis' website.

Model	> 10 %	< 10 %	< 5 %
<i>Federated Learning (original, 40)</i>	21.88 %	78.11 %	0.0 %
<i>Federated Learning (country, 40)</i>	15.02 %	84.98 %	0.0 %
<i>Random Forest</i>	84.22 %	8.0 %	8.0 %
<i>Neural Network (40)</i>	24.31 %	75.69 %	0.0 %
<i>Neural Network (50)</i>	16.79 %	83.21 %	0.0 %
<i>Neural Network (60)</i>	23.26 %	76.74 %	0.0 %

Table 5.4: Absolute Percentage Error (Accuracy) Comparison With Centralized Models

5.4 Limitations and Discussion

Both the risk measure and the cost estimation model underlie certain limitations. Therefore, the predictions in terms of risk and cost need to be consumed with caution. The base assumption that the later developments build upon is that survey data reflects the actual numbers accurately. Even though Accenture and IBM introduced checks and balances during their development, it cannot be ensured that the survey results are perfectly accurate. Especially since selection bias could have been introduced. Further data pollution could have occurred during the data extraction phase with OpenCV [132].

Due to the anonymization of data, only the relationship between one factor and costs can be observed at a time. However, cross-correlation effects can occur. For instance, the Industrial and the No Identity Access Management parameter increase the costs when they occur individually. This does not mean that their combinatorial appearance automatically has the same effect. However, the diverse data sample and careful factor selection suggest that the model's output is unlikely to be significantly skewed due to cross-correlation effects. Nevertheless, every cost prediction should be subject to a common sense review to ensure reasonable results.

Another assumption on which the model output relies heavily is that the average market capitalization of the Accenture sample can be approximated with the equity values of firms in the Russell Mid-Cap Index. Additionally, if the costs or the market capitalization needs to be scaled over multiple years, the discount factors strongly influence the results. Due to these two reasons, the predicted costs are relatively sensitive to capital input.

Regarding the risk measure of the RCVaR, many assumptions needed to be taken to develop the probability distribution. First of all, the assumptions of the Kolmogorov-Smirnov test are partly relaxed to conduct the test on discrete data. Additionally, it is assumed that all companies in the Accenture data can be treated as a single entity and have cost behavior similar to SMEs. While evidence in the literature (*cf.* Section 4.3.2) suggests that some of these assumptions may be relaxed, it is essential to keep them in mind when analyzing the risk output. Lastly, it is assumed that risk stays the same over time and in different business configurations. This assumption is certainly not valid [78] but was necessary due to limited data about the influence of business characteristics on variance over time. However, future work can provide more insights with the help of the Federated Learning (FL) approach investigated in this thesis (*cf.* Section 6.1). The complete list of assumptions for each chapter can be found in Appendix B.

Chapter 6

Summary and Conclusions

Cyber security is a critical concern for businesses due to the potential costs associated with attacks. In order to determine appropriate levels of investment in cyber security systems, it is necessary to consider the trade-off between the potential costs of a cyber attack, its risks, and the cost of security measures. This thesis proposes the Real Cyber Value at Risk (RCVaR) as a solution to address this issue.

The RCVaR addresses three issues in the field of cyber security economics. First, it provides a model for estimating the potential costs based only on quantitative data from the industry. It further offers more customization options than current solutions offered by the United Kingdom’s National Cyber Security Center or Italy’s Cyber Security Osservatorio. Moreover, the model is specifically tailored for use by SMEs. Through a simple web interface with extensive documentation, firms can access the proposed model and estimate the potential loss due to cyber attacks without requiring specialized knowledge in cyber security or informatics. This is particularly beneficial for SMEs, which often lack in-house computer specialists. The output provided by the model reflects the yearly instead of per-incident costs, allowing firms to consider these costs in their yearly investment planning. Unlike most reports and frameworks, the model covers all attack vectors and includes estimates of complex long-term costs, such as reputation damage or increased cost of capital. Additionally, the model allows reviewing past estimations as well as future predictions up to the year 2025.

Evaluating the cost estimation model on “unseen” work from industry and academic research suggests that the estimates are accurate. For example, the model accurately predicts the costs for an average US-based retail company in 2021 at \$ 420’444, deviating by only 2% from the estimate produced by Woods et al. [158] using insurance premiums from 26 insurance companies. In another case, the RCVaR could approximate the costs reported in Kaspersky’s 2013 cyber cost survey [72] with relatively little information. Besides that, the thesis demonstrates qualitatively that the cost behaviors of Accenture’s [2, 3, 10], IBM’s [27], and Ponemon’s [126, 127, 128, 129, 130, 139, 85] industry reports are incorporated successfully in the model. In addition to providing an expected cost estimate, the web application also provides insights into which business characteristics have the most significant impact on potential costs. Together with the most effective security

measure stated as an output, the company can start the investment process to decrease its costs.

The second contribution of this thesis is an individually interpretable, numerical risk measure. While current literature [49, 83] in the field has failed to provide organizations with a comprehensive, cross-domain risk measure that takes individual risk perceptions into account, the World Economic Forum proposed the Cyber Value at Risk (CVaR) measure in 2015. The first fully-developed model was presented in the scope of research at a study in Oxford [37] in the year 2021. However, previous studies on CVaR have relied on simulations based on threat and harm probability estimates. Therefore, new approaches are needed to empirically model the attack loss distribution, as related literature suggests [119]. Consequently, this thesis advances the work on the CVaR by developing a measure derived entirely from real-world empirical data and does not require probability estimates. This risk measure is easily accessible through the web interface, along with extensive documentation on the CVaR measure. A significant benefit of using the CVaR is that firms can compare cyber risk with risks of other domains, such as operational or financial risk, where the VaR is already used. By considering the expected cost and risk of cyber attacks, as well as the cost of security investments, companies can optimize their overall risk profile based on empirical performance indicators.

The CVaR measure established in this thesis was evaluated against the CVaR measure currently employed by the consulting firm MARSH [93]. Comparing the cost distribution shows that both distributions are very similar. Further examination of the limited literature [78, 158, 37] about attack cost distributions confirms the derived distribution in this thesis and consequently validates the CVaR risk measure.

The third contribution of this thesis addresses the prevalent lack of data in the field of cyber security economics. Most companies are hesitant to share information about cyber incidents for fear of additional costs [17, 101]. This thesis proposes a Federated Learning neural network, which allows firms to share their findings about cyber costs while maintaining their privacy. Evaluation of the FL model against a “centralized” learned model shows that the performance is comparable, even when some input features only appeared in a single client. Overall, a model which achieves an Absolute Percentage Error of less than 10% for more than 85% of observations was developed. Preliminary experiments in this thesis’s scope further indicate that with more semi-synthetic data, even better results can be achieved.

All three contributions show promising results and can provide fundamental information to the management of a company in the budget or risk planning process. The expected cost and risk of cyber attacks can be easily estimated without specialized knowledge, enabling informed discussions and recommendations for improving security most efficiently, based on quantitative data. Furthermore, Federated Learning provides a solution to extend the empirical basis to improve results in the future even further.

6.1 Future Work

As data scarcity is the most prevalent issue in cyber security economics, future work could deploy the Federated Learning process developed in the scope of this thesis. Results from a model in the production stage would allow the generation of synthetic data which resembles real-world circumstances more closely. It would give further insights into research fields that are currently covered insufficiently. Especially the cost distribution's evolution over time, different industries and countries is of enormous interest when conducting statistical cost predictions. Similarly, the results from the FL model could provide insights into the correlation effects between the factors developed in this thesis. Utilizing this information could lead to more accurate predictions in the future. Last but not least, this thesis used the most significant 13 factors from current industry reports. However, IBM's report [27] displays additional factors, such as the deployment stage of security measures and more detailed information on security systems. Incorporating this information might make the Real Cyber Value at Risk more complex but also more accurate.

Bibliography

- [1] A. Raghavan, and A. Thomas. Deloitte Review: Quantifying Risk, 2016. Accessible online: <https://www2.deloitte.com/us/en/insights/deloitte-review/issue-19/quantifying-risk-lessons-from-financial-services-industry.html>, Last accessed Oct. 2022.
- [2] Accenture and Ponemon Institute LLC. 2017 Cost of Cyber Crime Study, 2017. Accessible online: https://www.accenture.com/_acnmedia/pdf-62/accenture-2017costcybercrime-us-final.pdf, Last accessed Nov. 2022.
- [3] Accenture and Ponemon Institute LLC. The Cost of Cybercrime, 2019. Accessible online: https://www.accenture.com/_acnmedia/pdf-96/accenture-2019-cost-of-cybercrime-study-final.pdf, Last accessed Aug. 2022.
- [4] R. Akkiraju and A. Ivan. Discovering Business Process Similarities: An Empirical Study With SAP Best Practice Business Processes. In *International Conference on Service-Oriented Computing*, pages 515–526, San Francisco, USA, December 2010. Springer.
- [5] Allianz. Cyber: The Changing Threat Landscape, 2022. Accessible online: <https://www.agcs.allianz.com/news-and-insights/reports/cyber-risk-trends-2022.html>, Last accessed Aug. 2022.
- [6] R. Anderson, C. Barton, R. Böhme, and et al. Measuring the Cost of Cybercrime. In R. Böhme, editor, *The Economics of Information Security and Privacy*, pages 265–300. Springer, Berlin, Germany, 2013.
- [7] Y. Bengio, I. Goodfellow, and A. Courville. *Deep learning*. MIT press Cambridge, USA, 2017.
- [8] D. Beutel, T. Topal, A. Mathur, and et al. Flower: A Friendly Federated Learning Framework, 2022. Accessible online: <https://flower.dev/>, Last accessed Nov. 2022.
- [9] G. Bierman, M. Abadi, and M. Torgersen. Understanding Typescript. In *European Conference on Object-Oriented Programming*, pages 257–281, Uppsala, Sweden, July 2014. Springer.
- [10] K. Bissell, J. Fox, R. M. LaSalle, and et al. How Aligning Security and the Business Creates Cyber Resilience, 2021. Accessible online: <https://www.accenture.co>

- m/_acnmedia/PDF-165/Accenture-State-Of-Cybersecurity-2021.pdf, Last accessed Nov. 2022.
- [11] Bloomberg L.P. Market Cap Data for Russell Mid-Cap (RMC) Index 01/01/12 to 17/11/22. Retrieved from Bloomberg Database, 2022.
 - [12] Bloomberg L.P. Prices of CPI YoY 01/01/10 to 01/12/21. Retrieved from Bloomberg Database, 2022.
 - [13] R. Blumberg and S. Atre. The Problem With Unstructured Data. *Dm Review*, 13(42-49):62, 2003. Powell Publishing Inc.
 - [14] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi. Federated Learning of Predictive Models from Federated Electronic Health Records. *International Journal of Medical Informatics*, 112:59–67, 2018. Elsevier.
 - [15] D. E. Brown. Text Mining the Contributors to Rail Accidents. *IEEE Transactions on Intelligent Transportation Systems*, 17(2):346–355, 2015. IEEE.
 - [16] V. Brown. Risk Perception: It’s Personal. *Environmental Health Perspectives*, 122:A276–A279, 2014. NLM-Export.
 - [17] B. Cashell, W. D. Jackson, M. Jickling, and B. Webel. The Economic Impact of Cyber-Attacks, 2004. Accessible online: https://archive.nyu.edu/bitstream/2451/14999/2/Infosec_ISR_Congress.pdf, Last accessed Oct. 2022.
 - [18] H. Cavusoglu, H. Cavusoglu, and S. Raghunathan. Economics of ITSecurity Management: Four Improvements to Current Security Practices. *Communications of the Association for Information Systems*, 14(1):3, 2004. AIS eLibrary.
 - [19] H. Cavusoglu, B. Mishra, and S. Raghunathan. The Effect of Internet Security Breach Announcements on Market Value: Capital Market Reactions for Breached Firms and Internet Security Developers. *International Journal of Electronic Commerce*, 9(1):70–104, 2004. Taylor & Francis.
 - [20] F. Chollet. BatchNormalization Layer, 2015. Accessible online: https://keras.io/api/layers/normalization_layers/batch_normalization/, Last accessed Nov. 2022.
 - [21] F. Chollet. Keras, 2015. <https://keras.io>, Last accessed Nov. 2022.
 - [22] F. Chollet. Layer Activation Functions, 2015. Accessible online: <https://keras.io/api/layers/activations/>, Last accessed Nov. 2022.
 - [23] F. Chollet. Optimizers, 2015. Accessible online: <https://keras.io/api/optimizers/>, Last accessed Nov. 2022.
 - [24] F. Chollet. Regression losses. Accessible online: https://keras.io/api/losses/regression_losses/, Last accessed Nov. 2022, 2015.
 - [25] Consiglio Nazionale delle Ricerche. About Us, 2022. Accessible online: <https://www.cnr.it/en/about-us>, Last accessed Nov. 2022.

- [26] Consiglio Nazionale delle Ricerche Institute of Informatics and Telematics. Self Assessment Tools, 2018. Accessible online: <https://www.cybersecurityosservatorio.it/it/Services/survey.jsp>, Last accessed Sep. 2022.
- [27] IBM Corporation. Cost of a Data Breach Report 2022, 2022. Accessible online: <https://www.ibm.com/security/data-breach>, Last accessed Dec. 2022.
- [28] Council on Foreign Relations. Cyber Operations Tracker, 2022. Accessible online: <https://www.cfr.org/media>, Last accessed Nov. 2022.
- [29] Cyber Osservatorio and Consiglio Nazionale delle Ricerche. About Us, 2018. Accessible online: <https://www.cybersecurityosservatorio.it/en/node/286>, Last accessed Nov. 2022.
- [30] P. Dal Cin, J. Fox, and R. M. LaSalle. Elevating the Cybersecurity Discussion, 2022. Accessible online: https://www.accenture.com/_acnmedia/PDF-177/Accenture-Elevating-the-Cybersecurity-Discussion.pdf, Last accessed Aug. 2022.
- [31] Deloitte. Managing Cyber Risk With Smart Cyber, 2019. Accessible online: <https://www2.deloitte.com/content/dam/Deloitte/global/Documents/Risk/gx-ra-smart-cyber.pdf>, Last accessed Dec. 2022.
- [32] Deloitte Development LLC. VC Human Capital Survey, 2021. Accessible online: <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/audit/vc-human-capital-survey-3rd-edition-2021.pdf>, Last accessed Dec. 2022.
- [33] L. Ding and C. Zhou. Development of Web-Based System for Safety Risk Early Warning in Urban Metro Construction. *Automation in Construction*, 34:45–55, 2013. Elsevier.
- [34] Dow Jones & Company, Inc. Exchange Rates, 2022. Accessible online: <https://www.wsj.com/market-data/currencies/exchangerates>, Last accessed Nov. 2022.
- [35] P. D. Easton. PE Ratios, PEG Ratios, and Estimating the Implied Expected Rate of Return on Equity Capital. *The Accounting Review*, 79(1):73–95, 2004. American Accounting Association.
- [36] M. Eling, K. Jung, and J. Shim. Unraveling Heterogeneity in Cyber Risks Using Quantile Regressions. *Insurance: Mathematics and Economics*, 104:222–242, 2022. Elsevier.
- [37] A. Erola, I. Agrafiotis, J. Nurse, L. Axon, M. Goldsmith, and S. Creese. A System to Calculate Cyber-Value-at-Risk. *Computers & Security*, 113:102545, 2021. Elsevier.
- [38] European Commission. Rules for Business and Organisations, 2016. Accessible online: https://commission.europa.eu/law/law-topic/data-protection/reform/rules-business-and-organisations_en, Last accessed Sep. 2022.

- [39] European Network and Information Security Agency (ENISA). Introduction to Return on Security Investment, 2012. Accessible online: <https://www.enisa.europa.eu/publications/introduction-to-return-on-security-investment>, Last accessed Sep. 2022.
- [40] European Union Agency for Cybersecurity (ENISA). Cybersecurity for SMEs, 2021. Accessible online: <https://www.enisa.europa.eu/publications/enisa-report-cybersecurity-for-smes>, Last accessed Aug. 2022.
- [41] E. F. Fama. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2):383–417, 1970. JSTOR.
- [42] A. Faulkner, J. Layman, Maj. Gen. G. Franz, and D. Dalling. Achieving Federal Cyber Resilience, 2020. Accessible online: https://www.accenture.com/_acnmedia/PDF-130/Accenture-Federal-Cyber-Resilience-2020-report.pdf, Last accessed Aug. 2022.
- [43] Federal Bureau of Investigation (FBI). 2020 Internet Crime Report, 2021. Accessible online: https://www.ic3.gov/Media/PDF/AnnualReport/2020_IC3Report.pdf, Last accessed Aug. 2022.
- [44] F. Feng, W. Li, and Q. Jiang. Railway Traffic Accident Forecast Based on an Optimized Deep Auto-Encoder. *Promet-Traffic&Transportation*, 30(4):379–394, 2018. Fakultet prometnih znanosti Sveučilišta u Zagrebu.
- [45] Fidelity Investment. Sectors & Industries Overview, 2022. Accessible online: https://eresearch.fidelity.com/eresearch/markets_sectors/sectors/sectors_in_market.jhtml, Last accessed Nov. 2022.
- [46] Python Software Foundation. The Python Language Reference, 2023. Available at <https://docs.python.org/3/reference/>, Last accessed Dec. 2022.
- [47] M. F. Franco. *CyberTEA: a Technical and Economic Approach for Cybersecurity Planning and Investment*. PhD thesis, Communication Systems Group (CSG), University of Zurich, February 2023.
- [48] M. F. Franco, B. Rodrigues, C. Killer, E. J. Scheid, A. De Carli, A. Gassmann, D. Schoenbaechler, and B. Stiller. WeTrace: a Privacy-Preserving Tracing Approach. *Journal of Communications and Networks*, 23(5):374–389, 2021. KICS.
- [49] M. F. Franco, E. Sula, A. Huertas, E. J. Scheid, L. Z. Granville, and B. Stiller. SecRiskAI: a Machine Learning-Based Approach for Cybersecurity Risk Prediction in Businesses. In *24th IEEE International Conference on Business Informatics (CBI 2022)*, pages 1–10, Amsterdam, Netherlands, June 2022. IEEE.
- [50] J. Galindo and P. Tamayo. Credit Risk Assessment Using Statistical and Machine Learning: Basic Methodology and Risk Modeling Applications. *Computational Economics*, 15(1):107–143, 2000. Springer.
- [51] C. H. Gañán, M. Ciere, and M. Van Eeten. Beyond the Pretty Penny: The Economic Impact of Cybercrime. In *2017 New Security Paradigms Workshop*, pages 35–45, Santa Cruz, USA, October 2017. Association for Computing Machinery.

- [52] J. Gernand. Evaluating the Effectiveness of Mine Safety Enforcement Actions in Forecasting the Lost-Days Rate at Specific Worksites. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems*, 2(4):041002, 2016. American Society of Mechanical Engineers Digital Collection.
- [53] Y. M. Goh, C. Ubeynarayana, K. L. X. Wong, and B. H. Guo. Factors Influencing Unsafe Behaviors: A Supervised Learning Approach. *Accident Analysis & Prevention*, 118:77–85, 2018. Elsevier.
- [54] L. A. Gordon and M. P. Loeb. The Economics of Information Security Investment. *Association for Computing Machinery Transactions on Information and System Security (TISSEC)*, 5(4):438–457, 2002. Association for Computing Machinery.
- [55] M. Grinberg. *Flask Web Development: Developing Web Applications With Python*. O’Reilly Media, Inc., 2018.
- [56] S. Gupta, A. Singhal, and A. Kapoor. A Literature Survey on Social Engineering Attacks: Phishing Attack. In *2016 International Conference on Computing, Communication and Automation (ICCCA)*, pages 537–540, Noida, India, April 2016. IEEE.
- [57] R. A. Hallman, M. Major, J. Romero-Mariona, R. Phipps, E. Romero, M. John, and S. Miguel. Return on Cybersecurity Investment in Operational Technology Systems: Quantifying the Value That Cybersecurity Technologies Provide after Integration. In *5th International Conference on Complexity, Future Information Systems and Risk*, pages 43–52, Prague, Czechia, May 2020. SciTePress.
- [58] X. Han, H. Yu, and H. Gu. Visual Inspection with Federated Learning. In *International Conference on Image Analysis and Recognition*, pages 52–64, Waterloo, Canada, August 2019. Springer.
- [59] J. Hegde and B. Rokseth. Applications of Machine Learning Methods for Engineering Risk Assessment—A Review. *Safety Science*, 122:104492, 2020. Elsevier.
- [60] M. Hittmeir, A. Ekelhart, and R. Mayer. On the Utility of Synthetic Data: An Empirical Evaluation on Machine Learning Tasks. In *14th International Conference on Availability, Reliability and Security*, pages 1–6, Kent, UK, August 2019. Association for Computing Machinery.
- [61] G. Holton. History of Value-at-Risk: 1922-1998, 2002. Accessible online: <https://econwpa.ub.uni-muenchen.de/econ-wp/mhet/papers/0207/0207001.pdf>, Last accessed Nov. 2022.
- [62] Home Office Science Advisory Council. Understanding the Costs of Cyber Crime, 2018. Available at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/674046/understanding-costs-of-cyber-crime-horr96.pdf, Last accessed Dec. 2022.
- [63] C. Hsieh, W. Chao, P. Liu, and C. Li. Cyber Security Risk Assessment Using an Interpretable Evolutionary Fuzzy Scoring System. In *International Carnahan Conference on Security Technology (ICCST)*, pages 153–158, Taipei, Taiwan, September 2015. IEEE.

- [64] IBM Research. What Is Federated Learning?, 2022. Accessible online: <https://research.ibm.com/blog/what-is-federated-learning>, Last accessed Dec. 2022.
- [65] Iconiq Inc. How to Detect Specific Colors from an Image Using OpenCV, 2022. Accessible online: <https://www.projectpro.io/recipes/detect-specific-colors-from-image-opencv>, Last accessed Nov. 2022.
- [66] C. Inan. *A Visual Tool for the Analysis of Cybersecurity Investments*. Bachelor's thesis, Communication Systems Group (CSG), University of Zurich, August 2020.
- [67] Information Assurance for Small and Medium Enterprises Consortium Consortium Ltd and National Cyber Security Center. Cyber Essentials Questions, 2021. Accessible online: <https://getreadyforcyberessentials.iasme.co.uk/questions/>, Last accessed Sep. 2022.
- [68] International Organization for Standardization. Information Technology - Security Techniques - Information Security Management Systems - Requirements, 2005. Accessible online: <https://www.sis.se/api/document/preview/916657/>, Last accessed Sep. 2022.
- [69] International Organization for Standardization. Information Technology - Security Techniques - Information Security Management Systems - Guidance, 2017. Accessible online: <https://www.sis.se/api/document/preview/921740/>, Last accessed Sep. 2022.
- [70] International Organization for Standardization. Information Technology - Security Techniques - Information Security Risk Management, 2018. Accessible online: <https://www.sis.se/api/document/preview/80005503/>, Last accessed Sep. 2022.
- [71] J. Schmidt and Corporate Finance Institute. Valuation, 2022. Accessible online: <https://corporatefinanceinstitute.com/resources/valuation/valuation/>, Last accessed Dec. 2022.
- [72] Kaspersky Lab ZAO. Global Corporate IT Security Risks: 2013, 2013. Accessible online: https://media.kaspersky.com/en/business-security/Kaspersky_Global_IT_Security_Risks_Survey_report_Eng_final.pdf, Last accessed Dec. 2022.
- [73] D. Kawa, S. Punyani, P. Nayak, A. Karkera, and V. Jyotinagar. Credit Risk Assessment from Combined Bank Records Using Federated Learning. *International Research Journal of Engineering and Technology (IRJET)*, 6(4):1355–1358, 2019.
- [74] M. W. Khaw, L. Stevens, and M. Woodford. Individual Differences in the Perception of Probability. *PLoS Computational Biology*, 17(4):e1008871, 2021. Public Library of Science San Francisco, USA.
- [75] Y. Kim, J. Sun, H. Yu, and X. Jiang. Federated Tensor Factorization for Computational Phenotyping. In *23rd Association for Computing Machinery SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 887–895, Halifax, Canada, August 2017. Association for Computing Machinery.

- [76] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *3rd International Conference for Learning Representations*, pages 1–13, 2015.
- [77] V. S. Kumar and V. L. Narasimhan. Using Deep Learning for Assessing Cybersecurity Economic Risks in Virtual Power Plants. In *7th International Conference on Electrical Energy Systems (ICEES)*, pages 530–537, Chennai, India, February 2021. IEEE.
- [78] Kuypers, M. and Maillart, T. and Paté-Cornell, E. An Empirical Analysis of Cyber Security Incidents at a Large Organization. *Department of Management Science and Engineering, Stanford University, School of Information, UC Berkley*, 30, 2016.
- [79] M. Leo, S. Sharma, and K. Maddulety. Machine Learning in Banking Risk Management: A Literature Review. *Risks*, 7:29, 2019. MDPI.
- [80] H. Li, Z. Lü, and Z. Yue. Support Vector Machine for Structural Reliability Analysis. *Applied Mathematics and Mechanics*, 27(10):1295–1303, 2006. Springer.
- [81] H. Li, D. Parikh, Q. He, B.e Qian, Z. Li, D. Fang, and A. Hampapur. Improving Rail Network Velocity: A Machine Learning Approach to Predictive Maintenance. *Transportation Research Part C: Emerging Technologies*, 45:17–26, 2014. Elsevier.
- [82] L. Li, Y. Fan, M. Tse, and K. Lin. A Review of Applications in Federated Learning. *Computers & Industrial Engineering*, 149:106854, 2020. Elsevier.
- [83] T. Li, G. Convertino, R. K. Tayi, and S. Kazerooni. What Data Should I Protect? Recommender and Planning Support for Data Security Analysts. In *24th International Conference on Intelligent User Interfaces*, pages 286–297, Marina del Ray, USA, March 2019. Association for Computing Machinery.
- [84] G. Liu, C. Yu, S. Shiu, and I. Shih. The Efficient Market Hypothesis and the Fractal Market Hypothesis: Interfluves, Fusions, and Evolutions. *SAGE Open*, 12(1):21582440221082137, 2022. SAGE Publications Sage CA: Los Angeles, CA.
- [85] Ponemon Institute LLC. Why We Are Unique, 2022. Accessible online: <https://www.ponemon.org/about/>, Last accessed Sep. 2022.
- [86] London Stock Exchange Group. Membership list, 2021. Accessible online: https://content.ftserussell.com/sites/default/files/rumidcap_membershiplis_t_20210628.pdf, Last accessed Nov. 2022.
- [87] London Stock Exchange Group. Russell 1000 Index, 2022. Accessible online: <https://content.ftserussell.com/sites/default/files/russell-1000-index-product-highlights.pdf>, Last accessed Nov. 2022.
- [88] G. Louppe. Understanding Random Forests: From Theory to Practice. *arXiv e-prints*, page arXiv:1407.7502, 2014. arXiv.
- [89] Zurich Insurance Company Ltd. The Good, the Bad and the Careless, 2015. Accessible online: https://www.zurichcanada.com/-/media/project/zwp/canada/knowledge/docs/english/security_privacy/good_bad_careless_eng.pdf, Last accessed Sep. 2022.

- [90] M. Goldman. Statistics for Bioinformatics, 2008. Accessible online: <https://www.stat.berkeley.edu/~mgoldman/Section0402.pdf>, Last accessed Nov. 2022.
- [91] V. S. Mai, R. J. La, and A. Battou. Optimal Cybersecurity Investments in Large Networks Using SIS Model: Algorithm Design. *IEEE/Association for Computing Machinery Transactions on Networking*, 29(6):2453–2466, 2021. IEEE.
- [92] A. B. Mansour, G. Carenini, A. Duplessis, and D. Naccache. Federated Learning Aggregation: New Robust Algorithms with Guarantees. *arXiv*, abs/2205.10864, 2022.
- [93] MARSH LLC. Cyber Value-at-Risk (Cyber VaR), 2017. Accessible online: <https://www.marsh.com/content/dam/marsh/Documents/PDF/US-en/Cyber%20Value-at-Risk.pdf>, Last accessed Dec. 2022.
- [94] A. Martín, A. Ashish, B. Paul, and et al. Build, Train and Evaluate Models With TensorFlow Decision Forests, 2015. Accessible online: https://www.tensorflow.org/decision_forests/tutorials/beginner_colab, Last accessed Dec. 2022.
- [95] A. Martín, A. Ashish, B. Paul, and et al. TensorFlow Federated: Machine Learning on Decentralized Data, 2015. Accessible online: <https://www.tensorflow.org/federated>, Last accessed Nov. 2022.
- [96] A. Martín, A. Ashish, B. Paul, and et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. Accessible online: <https://www.tensorflow.org/>, Last accessed Nov. 2022.
- [97] A. Mashrur, W. Luo, N. Zaidi, and A. Robles-Kelly. Machine Learning for Financial Risk Management: A Survey. *IEEE Access*, 8:203203–203223, 2020. IEEE.
- [98] J. Massey and J. Frank. The Kolmogorov-Smirnov Test for Goodness of Fit. *Journal of the American Statistical Association*, 46(253):68–78, 1951.
- [99] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Arcas. Communication-Efficient Learning of Deep Networks From Decentralized Data. In *International Conference on Artificial Intelligence and Statistics*, pages 1273–1282, Ft. Lauderdale, USA, April 2017. PMLR.
- [100] Meta. React, 2022. Available at <https://reactjs.org/>, Last accessed Dec. 2022.
- [101] T. Moore. The Economics of Cybersecurity: Principles and Policy Options. *International Journal of Critical Infrastructure Protection*, 3(3-4):103–117, 2010. Elsevier.
- [102] E. Mossburg, J. Gelinne, and H Calzada. Beneath the Surface of a Cyberattack, 2016. Accessible online: <https://www2.deloitte.com/global/en/pages/risk/cyber-strategic-risk/articles/beneath-the-surface-of-a-cyberattack.html>, Last accessed Aug. 2022.
- [103] N. I. Mowla, N. H. Tran, I. Doh, and K. Chae. Federated Learning-Based Cognitive Detection of Jamming Attack in Flying Ad-Hoc Network. *IEEE Access*, 8:4338–4350, 2019. IEEE.

- [104] National Cyber Security Center. About Cyber Essentials, 2021. Accessible online: <https://www.ncsc.gov.uk/cyberessentials/overview>, Last accessed Nov. 2022.
- [105] National Institute for Standardization and Technology. Kolmogorov Smirnov Two Sample, 2016. Accessible online: <https://www.itl.nist.gov/div898/software/dataplot/refman1/auxillar/ks2samp.htm>, Last accessed Nov. 2022.
- [106] National Institute for Standardization and Technology. Kolmogorov-Smirnov Goodness-of-Fit Test, 2021. Accessible online: <https://www.itl.nist.gov/div898/handbook/eda/section3/eda35g.htm>, Last accessed Nov. 2022.
- [107] National Institute for Standardization and Technology. Related Distributions, 2021. Accessible online: <https://www.itl.nist.gov/div898/handbook/eda/section3/eda362.htm>, Last accessed Nov. 2022.
- [108] National Institute of Standards and Technology. Framework for Improving Critical Infrastructure Cybersecurity, 2018. Accessible online: https://www.baltimorecityschools.org/sites/default/files/inline-files/NIST.CSWP_.04162018.pdf, Last accessed Sep. 2022.
- [109] A. Nilsson, S. Smith, G. Ulm, E. Gustavsson, and M. Jirstrand. A Performance Evaluation of Federated Learning Algorithms. In *Second Workshop on Distributed Infrastructures for Deep Learning*, pages 1–8, Rennes, France, December 2018. Association for Computing Machinery.
- [110] NYU Stern School of Business. Return on Equity by Sector (US), 2022. Accessible online: https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/roe.html, Last accessed Dec. 2022.
- [111] OECD. *The Digital Transformation of SMEs*. OECD Publishing, Paris, 2021.
- [112] Office of the Prime Minister. A Strong Britain in an Age of Uncertainty: The National Security Strategy, 2010. Accessible online: https://webarchive.nationalarchives.gov.uk/ukgwa/20121015000000/http://www.direct.gov.uk/prod_consum_dg/groups/dg_digitalassets/@dg/@en/documents/digitalasset/dg_191639.pdf, Last accessed Sep. 2022.
- [113] C. Omlin. *A Gordon-Loeb-based Visual Tool for Cybersecurity Investments*. Bachelor’s thesis, Communication Systems Group (CSG), University of Zurich, January 2022.
- [114] OpenCV. Canny Edge Detection, 2022. Accessible online: https://docs.opencv.org/4.x/da/d22/tutorial_py_canny.html, Last accessed Nov. 2022.
- [115] OpenCV. Hough Circle Transform, 2022. Accessible online: https://docs.opencv.org/3.4/d4/d70/tutorial_hough_circle.html, Last accessed Nov. 2022.
- [116] OpenCV. Mouse as a Paint-Brush, 2022. Accessible online: https://docs.opencv.org/3.4/db/d5b/tutorial_py_mouse_handling.html, Last accessed Nov. 2022.

- [117] OpenCV. Sobel Derivatives, 2022. Accessible online: https://docs.opencv.org/3.4/d2/d2c/tutorial_sobel_derivatives.html, Last accessed Nov. 2022.
- [118] Organisation for Economic Co-operation and Development. Enterprises by Business Size, 2022. Accessible online: <https://data.oecd.org/entrepreneur/enterprises-by-business-size.htm#:~:text=SMEs%20are%20further%20subdivided%20into,employ%20250%20or%20more%20people.>, Last accessed Nov. 2022.
- [119] A. Orlando. Cyber Risk Quantification: Investigating the Role of Cyber Value at Risk. *Risks*, 9(10):184, 2021.
- [120] M. S. Ozdayi, M. Kantarcioglu, and R. Iyer. Improving Accuracy of Federated Learning in Non-IID Settings. *arXiv e-prints*, page arXiv:2010.15582, 2020. arXiv.
- [121] Peak Capital. Understanding Early Stage Venture Capital, 2022. Accessible online: <https://peak.capital/understanding-early-stage-venture-capital/>, Last accessed Dec. 2022.
- [122] F. Pedregosa, G. Varoquaux, A. Gramfort, and et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011. JMLR.
- [123] PitchBook. US Venture Capital Valuation Trends in Four Charts, 2022. Accessible online: <https://pitchbook.com/news/articles/us-venture-capital-valuation-trends-charts>, Last accessed Dec. 2022.
- [124] Pitchbook. VC Valuations Report, 2022. Accessible online: https://files.pitchbook.com/website/files/pdf/Q1_2022_US_VC_Valuations_Report.pdf, Last accessed Dec. 2022.
- [125] PitchBook. Who We Are, 2022. Accessible online: <https://pitchbook.com/about>, Last accessed Dec. 2022.
- [126] Ponemon Institute LLC. First Annual Cost of Cyber Crime Study, 2010. Accessible online: https://citadel-information.com/wp-content/uploads/2010/12/2010-ponemon_Cost_of_Cyber_Crime_study.pdf, Last accessed Oct. 2022.
- [127] Ponemon Institute LLC. Second Annual Cost of Cyber Crime Study, 2011. Accessible online: https://www.ponemon.org/local/upload/file/2011_2nd_Annual_Cost_of_Cyber_Crime_Study%20.pdf, Last accessed Dec. 2022.
- [128] Ponemon Institute LLC. 2012 Cost of Cyber Crime Study: United States, 2012. Accessible online: https://www.ponemon.org/local/upload/file/2012_US_Cost_of_Cyber_Crime_Study_FINAL6%20.pdf, Last accessed Dec. 2022.
- [129] Ponemon Institute LLC and Experian Data Breach Resolution. The Aftermath of a Data Breach: Consumer Sentiment, 2014. Accessible online: <https://www.ponemon.org/news-updates/blog/security/the-aftermath-of-a-data-breach-consumer-sentiment.html>, Last accessed Aug. 2022.
- [130] Ponemon Institute LLC and IBM. 2016 Cost of Data Breach Study: Global Analysis, 2016. Accessible online: <https://www.cloudmask.com/hubfs/IBMstudy.pdf>, Last accessed Sep. 2022.

- [131] PricewaterhouseCoopers (PwC). Reimagining the Outcomes That Matter, 2022. Available at https://www.pwc.com/gx/en/ceo-survey/2022/main/content/downloads/25th_CEO_Survey_PDF_report.pdf, Last accessed Aug. 2022.
- [132] Python Software Foundation. OpenCv-Python 4.6.0.66, 2022. Accessible online: <https://pypi.org/project/opencv-python/>, Last accessed Sep. 2022.
- [133] M. Rea-Guaman, J. A. Calvo-Manzano, and T. San Feliu. A Prototype to Manage Cybersecurity in Small Companies. In *13th Iberian Conference on Information Systems and Technologies (CISTI)*, pages 1–6, Caceres, Spain, June 2018. IEEE.
- [134] Christian Robert. Generalized Inverse Normal Distributions. *Statistics & Probability Letters*, 11(1):37–41, 1991. Elsevier.
- [135] B. Rodrigues, M. Franco, G. Parangi, and B. Stiller. SEconomy: a Framework For the Economic Assessment of Cybersecurity. In *International Conference on the Economics of Grids, Clouds, Systems, and Services*, pages 154–166, Leeds, UK, September 2019. Springer.
- [136] S. Santurkar, D. Tsipras, A. Ilyas, and A. Madry. How Does Batch Normalization Help Optimization? *Advances in Neural Information Processing Systems*, 31:2488–2498, 2018. arXiv.
- [137] Scikit-Learn Developers (BSD License). Sklearn.Preprocessing.MinMaxScaler, 2022. Accessible online: <https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.MinMaxScaler.html>, Last accessed Nov. 2022.
- [138] Scikit-Learn Developers (BSD License). Sklearn.Preprocessing.MultiLabelBinarizer, 2022. Accessible online: <https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.MultiLabelBinarizer.html>, Last accessed Dec. 2022.
- [139] IBM Security and Ponemon Institute LLC. 2017 Cost of Data Breach Study, 2017. Accessible online: https://www.ncsl.org/documents/taskforces/IBM_Ponemon_2017CostofDataBreachStudy.pdf, Last accessed Sep. 2022.
- [140] A. Sharma and U. K. Singh. Modelling of Smart Risk Assessment Approach for Cloud Computing Environment Using AI & Supervised Machine Learning Algorithms. *Global Transitions Proceedings*, 3(1):243–250, 2022.
- [141] D. Shinder and M. Cross. *Scene of the Cybercrime*. Elsevier, 2008.
- [142] S. Y. Sohn and S. H. Lee. Data Fusion, Ensemble and Clustering to Improve the Classification Accuracy for the Severity of Road Traffic Accidents in Korea. *Safety Science*, 41(1):1–14, 2003. Elsevier.
- [143] L. Sousa, T. Miranda, R. Sousa, and J. Tinoco. The Use of Data Mining Techniques in Rockburst Risk Assessment. *Engineering*, 3(4):552–558, 2017. Elsevier.
- [144] S. Stack and IBM. The Definition of Security Intelligence, 2022. Accessible online: <https://securityintelligence.com/defintion-security-intelligence/>, Last accessed Sep. 2022.

- [145] State Secretariat for Economic Affairs (SECO). Digitalization of SMEs in Switzerland: a Key Factor, 2021. Accessible online: https://www.accenture.com/_acn/media/pdf-96/accenture-2019-cost-of-cybercrime-study-final.pdf, Last accessed Aug. 2022.
- [146] D. Steinskog, D. Tjøstheim, and N. Kvamstø. A Cautionary Note on the Use of the Kolmogorov–Smirnov Test for Normality. *Monthly Weather Review*, 135(3):1151–1157, 2007. AMS.
- [147] V. Svetnik, A. Liaw, C. Tong, J. C. Culberson, R. P. Sheridan, and B. P. Feuston. Random Forest: A Classification and Regression Tool for Compound Classification and QSAR Modeling. *Journal of Chemical Information and Computer Sciences*, 43(6):1947–1958, 2003. ACS.
- [148] D. A. Swanson. On the Relationship Among Values of the Same Summary Measure of Error When Used Across Multiple Characteristics at the Same Point in Time: An Examination of MALPE and MAPE. *Review of Economics and Finance*, 5(1):1–14, 2015. Better Advances Press, Canada.
- [149] The Linux Foundation. From Research to Production, 2015. Accessible online: <https://pytorch.org/>, Last accessed Nov. 2022.
- [150] The SciPy Community. Scipy.Stats.Geninvgauss, 2022. Accessible online: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.geninvgauss.html>, Last accessed Nov. 2022.
- [151] The SciPy Community. Statistical Functions, 2022. Accessible online: <https://docs.scipy.org/doc/scipy/reference/stats.html>, Last accessed Nov. 2022.
- [152] A. Tixier, M. R. Hallowell, B. Rajagopalan, and D. Bowman. Automated Content Analysis for Construction Safety: A Natural Language Processing System to Extract Precursors and Outcomes from Unstructured Injury Reports. *Automation in Construction*, 62:45–56, 2016. Elsevier.
- [153] S. Tweneboah-Kodua, F. Atsu, and W. Buchanan. Impact of Cyberattacks on Stock Performance: A Comparative Study. *Information & Computer Security*, 26(5):637–652, 2018. Emerald Publishing Limited.
- [154] U.S. Bureau of Labor Statistics. Consumer Price Index, 2022. Accessible online: <https://www.bls.gov/cpi/>, Last accessed Dec. 2022.
- [155] U.S. Department of Homeland Security. Enhancing Resilience Through Cyber Incident Data Sharing and Analysis, 2015. Accessible online: https://insidecybersecurity.com/sites/insidecybersecurity.com/files/documents/sep2015/cs09142015_Data_Categories_White_Paper.pdf, Last accessed Sep. 2022.
- [156] U.S. Department of Homeland Security. Cybersecurity, 2022. Accessible online: <https://www.dhs.gov/topics/cybersecurity>, Last accessed Sep. 2022.

- [157] M. Van Wieren, E. Van Luit, Estourgie R., V. Jacobs, and J Bulters. Cyber Value at Risk in the Netherlands, 2016. Accessible online: <https://securitydelta.nl/images/deloitte-nl-risk-cyber-value-at-Risk-in-the-Netherlands.pdf>, Last accessed Nov. 2022.
- [158] D. W. Woods, T. Moore, and A. C. Simpson. The County Fair Cyber Loss Distribution: Drawing Inferences from Insurance Prices. *Digital Threats: Research and Practice*, 2(2):1–21, 2021. Association for Computing Machinery.
- [159] World Economic Forum. Partnering for Cyber Resilience Towards the Quantification of Cyber Threats, 2015. Accessible online: https://www3.weforum.org/docs/WEFUSA_QuantificationofCyberThreats_Report2015.pdf, Last accessed Nov. 2022.
- [160] Y. Yamai and T. Yoshiba. On the Validity of Value-at-Risk: Comparative Analyses with Expected Shortfall. *Monetary and Economic Studies*, 20(1):57–85, 2002. Institute for Monetary and Economic Studies, Bank of Japan.
- [161] Y. Yamai and T. Yoshiba. Value-at-Risk Versus Expected Shortfall: A Practical Perspective. *Journal of Banking & Finance*, 29(4):997–1015, 2005. Elsevier.
- [162] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra. Federated Learning With Non-IID Data. *arXiv e-prints*, page arXiv:1806.00582, 2018. arXiv.

Abbreviations

AI	Artificial Intelligence
ANN	Artificial Neural Network
APE	Absolute Percentage Error
APF	Absolute Prozentuale Fehler
BYOD	Bring Your Own Device
CDF	Cumulative Distribution Function
CRS	Congressional Research Service
CSI	Computer Security Institute
CVaR	Cyber Value at Risk
CPI	Consumer Price Index
DHS	Department of Homeland Security
DCF	Discounted Cashflow
DL	Deep Learning
ECDF	Empirical Cumulative Distribution Function
ENISA	European Network and Information Security Agency
EU	European Union
FBI	Federal Bureau of Investigation
FL	Federated Learning
GDPR	General Data Protection Regulation
IASME	Information Assurance for Small and Medium Enterprises Consortium
IID	Independent and Identically Distributed
IoT	Internet of Things
ISO	International Organization for Standardization
KMU	Kleine und Mittlere Unternehmen
KS	Kolmogorov-Smirnov
MAPE	Mean Absolute Percentage Error
MSE	Mean Squared Error
ML	Machine Learning
NCSC	National Cyber Security Center
NIST	National Institute for Standards and Technology
NPV	Net Present Value
OECD	Organisation for Economic Co-operation and Development
PWC	PricewaterhouseCoopers
PPF	Percentage Probability Function
RF	Random Forest
RGB	Red, Green, and Blue

ROE	Return on Equity
ROSI	Return on Security Investment
RCVaR	Real Cyber Value at Risk
SME	Small and Middle-Sized Enterprises
USD	US-Dollar
VaR	Value at Risk
VC	Venture Capital
WEF	World Economic Forum
WLAN	Wireless Local Area Network

List of Figures

2.1	Value at Risk Representations Based on [160]	7
2.2	General Federated Learning Process	9
3.1	Hypothetical Cyber Cost Distributions	21
4.1	Overview of the RCVaR Development Process	23
4.2	Company Individual Cost Presented in the Accenture 2017 Report [2]	25
4.3	Results of Dot-Detection Approaches Applied to Figure 4.2	27
4.4	Results of Edge Detection Algorithms Applied to Figure 4.2	27
4.5	Manual-Detection of Data Points Through Mouse Clicks	28
4.6	Market Cap of the Russell Mid-Cap Index (RMC) Based on Data From [11]	30
4.7	The Influence of Inflation on Cost	31
4.8	Evolution of the Cumulative Inflation Over Time	33
4.9	Evolution of Cumulative Cost Over Time	34
4.10	Parameter Ratios for <i>Country</i> and <i>Industry</i> Factors	41
4.11	Maximum and Minimum Impact Parameters per Factor	42
4.12	Distributions Within Factors	43
4.13	Cost Density Distribution of the Accenture Sample [2]	45
4.14	Distributions Fit to Sample	49
4.15	RCVaR of an Average Company in 2019 With 95% Confidence	51
4.16	Semi-Synthetic Data Along Capital and Cost Dimensions	55
4.17	Neural Network Architecture for Federated Learning	58


4.18	Mean Absolute Percentage Error (MAPE) of Clients in the First Round . .	59
4.19	Mean Absolute Percentage Error (MAPE) of Clients in the Fifth Round . .	60
4.20	Mean Absolute Percentage Error (MAPE) and Mean Squared Error (MSE)	60
4.21	High-Level Components of the RCVaR Website	62
4.22	RCVaR Web-Interface for Four Input Factors	64
4.23	Web-Interface for Cost Decomposition and Factor Analysis	65
4.24	Web-Interface for Numerical Output of Cost Analysis	66
5.1	Cyber Cost Density Distributions MARSH [93] Vs. RCVaR	72
5.2	Model 1; Prediction Deviations (FL/Random Split/40 epochs)	75
5.3	Model 1; Cost Accuracy Along Factor Dimensions	76
5.4	Model 3 and 4; Prediction Deviations	77
5.5	Training Data Distribution of Capital Quantile Split	78
5.6	Model 6 and 8; Prediction Deviations	80
5.7	Model 8; Cost Accuracy Along Factor Dimensions	81
A.1	Model 1 (FL/100'000 Observations/Random Split/40 epochs)	109
A.2	Model 2 (FL/500'000 Observations/Random Split/40 epochs)	109
A.3	Model 3 (FL/100'000 Observations/Country Split/40 epochs)	110
A.4	Model 4 (FL/100'000 Observations/Quintile Split/40 epochs)	110
A.5	Model 5 (Centralized ANN/100'000 Observations/No Split/40 epochs) . . .	111
A.6	Model 6 (Centralized ANN/100'000 Observations/No Split/50 epochs) . . .	111
A.7	Model 7 (Centralized ANN/100'000 Observations/No Split/60 epochs) . . .	112
A.8	Model 8 (Centralized RF/100'000 Observations/No Split/ -)	112
C.1	Required Inputs	119
C.2	How the Factors Influence the Cost	120
C.3	Cost Composition and Security Rating	120
C.4	Numerical Outputs of the RCVaR	121

List of Tables


2.1	Common Cyber Cost Dimensions in Literature	6
3.1	Summary of the Two Most Established Cost Reports	16
3.2	Risk Assessment ML-Algorithms in Different Industries [59]	18
4.1	Cost-Influencing Factors and Their Parameters	37
4.2	Data Sources of Factors and When They Are Available	39
4.3	RCVaR Discount Factors for Time and Size Scaling	44
4.4	P-Values of Distributions	49
4.5	Training Parameters for Federated Learning	61
5.1	Hypothetical Companies and Their Associated Costs	69
5.2	Unseen Cost Estimation Vs. RCVaR Cost Prediction	71
5.3	Absolute Percentage Error (Accuracy) Comparison of Different Data Splits	79
5.4	Absolute Percentage Error (Accuracy) Comparison With Centralized Models	82
A.1	Prediction Vs. True Cost Along Country Dimension	112
A.2	Prediction Vs. True Cost Along Industry Dimension	113
D.1	Technical Requirements to Run the Code	123
D.2	Technical Requirements to Run the Code	124
D.3	Technical Requirements to Run the Code	125

Appendix A


Complete Evaluation Data



(a) Deviation in Test Sample
(● < 5%, ● < 10%, ● > 10%)




(b) Predicted and True Cost
along Country Dimension




(c) Predicted and True Cost
along Industry Dimension


Figure A.1: Model 1 (FL/100'000 Observations/Random Split/40 epochs)



(a) Deviation in Test Sample
(● < 5%, ● < 10%, ● > 10%)




(b) Predicted and True Cost
along Country Dimension



(c) Predicted and True Cost
along Industry Dimension

Figure A.2: Model 2 (FL/500'000 Observations/Random Split/40 epochs)




(a) Deviation in Test Sample
(● $< 5\%$, ● $< 10\%$, ● $> 10\%$)

(b) Predicted and True Cost
along Country Dimension

(c) Predicted and True Cost
along Industry Dimension

Figure A.3: Model 3 (FL/100'000 Observations/Country Split/40 epochs)




(a) Deviation in Test Sample
(● $< 5\%$, ● $< 10\%$, ● $> 10\%$)


(b) Predicted and True Cost
along Country Dimension

(c) Predicted and True Cost
along Industry Dimension


Figure A.4: Model 4 (FL/100'000 Observations/Quintile Split/40 epochs)



(a) Deviation in Test Sample
(● < 5%, ● < 10%, ● > 10%)




(b) Predicted and True Cost
along Country Dimension




(c) Predicted and True Cost
along Industry Dimension


Figure A.5: Model 5 (Centralized ANN/100'000 Observations/No Split/40 epochs)



(a) Deviation in Test Sample
(● < 5%, ● < 10%, ● > 10%)




(b) Predicted and True Cost
along Country Dimension



(c) Predicted and True Cost
along Industry Dimension

Figure A.6: Model 6 (Centralized ANN/100'000 Observations/No Split/50 epochs)




(a) Deviation in Test Sample
(● < 5%, ● < 10%, ● > 10%)

(b) Predicted and True Cost
along Country Dimension

(c) Predicted and True Cost
along Industry Dimension

Figure A.7: Model 7 (Centralized ANN/100'000 Observations/No Split/60 epochs)



(a) Deviation in Test Sample
(● < 5%, ● < 10%, ● > 10%)

(b) Predicted and True Cost
along Country Dimension

(c) Predicted and True Cost
along Industry Dimension

Figure A.8: Model 8 (Centralized RF/100'000 Observations/No Split/ -)

	<i>TUR</i>	<i>SCA</i>	<i>ESP</i>	<i>ITA</i>	<i>FRA</i>	<i>UK</i>	<i>GER</i>	<i>CAN</i>	<i>US</i>
<i>Model 1</i>	15%	18%	16%	15%	14%	14%	15%	16%	18%
<i>Model 2</i>	21%	20%	27%	28%	22%	18%	11%	10%	0%
<i>Model 3</i>	10%	10%	11%	10%	11%	12%	10%	10%	10%
<i>Model 4</i>	6%	3%	8%	9%	6%	6%	3%	4%	2%
<i>Model 5</i>	13%	13%	11%	15%	17%	14%	17%	14%	15%
<i>Model 6</i>	14%	14%	17%	16%	15%	15%	17%	18%	20%
<i>Model 7</i>	8%	15%	12%	12%	12%	10%	10%	10%	19%
<i>Model 8</i>	21%	24%	16%	19%	32%	39%	44%	45%	34%

Table A.1: Prediction Vs. True Cost Along Country Dimension

	<i>Education</i>	<i>Public Sector</i>	<i>Life Sciences</i>	<i>Travel</i>	<i>Communications and Media</i>	<i>Consumer Goods</i>	<i>Retail</i>	<i>US Federal</i>	<i>Capital Markets</i>	<i>Automotive</i>	<i>Industrial</i>	<i>Energy</i>	<i>Insurance</i>	<i>Health</i>	<i>High Tech</i>	<i>Pharmaceuticals</i>	<i>Software</i>	<i>Aerospace and Defence</i>	<i>Utilities</i>	<i>Banking</i>
<i>Model 1</i>	18%	18%	17%	16%	14%	14%	15%	13%	14%	15%	16%	17%	29%	17%	17%	17%	17%	16%	16%	17%
<i>Model 2</i>	30%	31%	35%	34%	27%	31%	26%	22%	16%	16%	12%	10%	32%	7%	7%	2%	2%	0%	4%	3%
<i>Model 3</i>	11%	11%	9%	9%	11%	10%	10%	12%	10%	12%	10%	10%	26%	9%	9%	9%	9%	8%	7%	7%
<i>Model 4</i>	11%	11%	10%	9%	7%	6%	7%	7%	5%	7%	5%	4%	12%	4%	7%	8%	5%	2%	1%	1%
<i>Model 5</i>	17%	13%	11%	12%	19%	15%	19%	18%	16%	16%	14%	14%	30%	13%	15%	16%	15%	17%	18%	15%
<i>Model 6</i>	18%	16%	17%	15%	13%	15%	15%	14%	14%	15%	18%	18%	29%	17%	19%	18%	19%	22%	24%	20%
<i>Model 7</i>	18%	19%	15%	12%	13%	13%	13%	10%	8%	10%	8%	11%	25%	10%	12%	11%	12%	14%	18%	16%
<i>Model 8</i>	1%	2%	10%	18%	28%	31%	31%	35%	41%	41%	44%	49%	50%	50%	50%	52%	54%	56%	59%	59%

Table A.2: Prediction Vs. True Cost Along Industry Dimension

Appendix B

Complete Limitation List

Assumptions

- *Bias free dataset.*
- *The report from senior officials is a good representation of the reality.*
- *The currency exchange rates are not abnormal during conversion.*
- *The extracted data is representative of the distribution in Graphic 4.2.*
- *The average market cap of the Russell Mid Cap Index approximates the average market cap of the companies in the data samples.*
- *The Costs of the “Vulnerable” category in the 2021 report [10] can be reasonable approximated.*
- *Inflation is a reasonable discount factor for different assets across various industries.*
- *Both regression approximate the annualized increase of cost, respectively inflation, relatively well.*
- *Cross-Correlation between Factors is assumed to be zero.*
- *Parameter ratios are constant over time.*
- *The assumptions of the Kolmogorov-Smirnov test can be relaxed.*
- *Companies in the sample are similar.*
- *Distribution is the same over time and factors.*
- *Costs of SME have the same distribution.*
- *Venture Capital market capitalization approximates the valuation of SME.*

Appendix C

Documentation on Website

This chapter portrays the short Summary file accessible through the *Documentation* page on the RCVaR website:

This short documentation should help you understand how to interact with the RCVaR interface and how to interpret the numbers in a meaningful way without requiring you to read the whole thesis.

The RCVaR aims to provide you with three primary pieces of information.

1. What are my expected costs due to cyber attacks for a specific year?

Costs in the scope of the RCVaR include Indirect and Opportunity costs. For instance, the RCVaR captures the potential revenue lost due to customers' unwillingness to use your services as a consequence of the decrease in trust. Another example would be the increased lending costs or follow-up lawsuits.

2. What is the Cyber Value at Risk for a specific year?

The Cyber Value at Risk is a numerical risk measure that reflects the highest possible attack costs for a specific year with a certain confidence. In other words, with a certain confidence, the costs for a year will not be higher than the Cyber Value at Risk.

3. What are the most efficient actions to lower my expected cyber attack cost for a specific year?


The output provides multiple insights into the composition of costs. It also shows which business characteristics have a positive or a negative influence on the expected costs. In addition, the RCVaR provides a security measure with the highest possible chance of decreasing the cost for your firm.

The steps below give more detailed instructions on how to receive these three pieces of information and how to interpret them.

1. First, you must navigate to the RCVaR web page by clicking on the RCVaR Symbol in the header.
2. Upon navigating to the RCVaR web page, you see multiple different input fields. In total, you can configure the output of the RCVaR by customizing over 14 different characteristics. For all of them, a short description is available above the input field itself. To evaluate your company's expected cost, a minimum of three characteristics must be specified. First, the year for which the expected loss should be computed must be specified in the *Year* Input section. Next, the company valuation must be provided to the interface in the *Capital* field. The valuation reflects the worth of the company, more specifically, the equity value of the company. A good starting point to elicit the company value is by taking a look at the tax-forms or the company's balance sheet. More advanced methods are Peer-Group comparisons or Discounted

Cashflow methods. The capital must be specified for the current year. The last mandatory input is the *Confidence*. The *Confidence* reflects how secure you want to be regarding the CVaR value displayed in the end.

Year


 Enter the year for which you want to calculate the expected cost

Required *

2017

(a) *Year* Input

Valuation

 Enter the equity value/ valuation of the company for which you want to determine the impact. The amount must be in USD.


Valuation in USD

70000000

\$

(b) *Capital* Input

Confidence

 Enter the level of confidence desired for the Value at Risk. Respectively, how confident do you want to be that your cost will not be higher than the cost displayed as Value at Risk?

Confidence

95%

(c) *Confidence* Input

Figure C.1: Required Inputs

After providing all these inputs, additional information, such as the amount of remote workers or the industry of the firm can be specified as well. This additional information ultimately helps the RCVaR provide you with more accurate and customized answers. It is therefore recommended to give as much input as possible.

3. Once the inputs are entered, you can click the *Evaluate Impact* button at the bottom.
4. The first two outputs show the influence of the characteristics on the average expected loss if only the mandatory inputs were provided. For instance, the chart in Figure C.2b shows that not having insurance negatively impacts the expected costs, meaning the loss increases without insurance. The percentage next to the cost-influencing factor weights the factor so that the different factors can be compared. Of all cost-decreasing factors, the retail sector has the most significant influence, with 67.4%. Based on this information, you can reduce your company's exposure to negative factors and introduce more cost-decreasing factors.




Figure C.2: How the Factors Influence the Cost

The subsequent output symbolizes how the costs are distributed among different actions in a specific year. For instance, the chart in Figure C.3a indicates that the biggest cost component in that particular year are lost business opportunities. Figure C.3b shows a comparison of your company’s current security state against the best and worst possible security state of your individual company. Hence, the best possible security state (if all security-related actions are implemented) of your personal company is the best possible state in this comparison.




Figure C.3: Cost Composition and Security Rating

When looking at the numerical outputs in Figure C.4, you can observe two estimates for the expected costs for the specified year. One estimate was produced using a

mathematical model, while the other value is based on an Artificial Neural Network. Due to limited data during the development of this thesis, the mathematical model is considered more accurate. Next, the security measure, which would have the highest impact of decreasing the expected cost, is depicted together with the aforementioned Cyber Value at Risk.

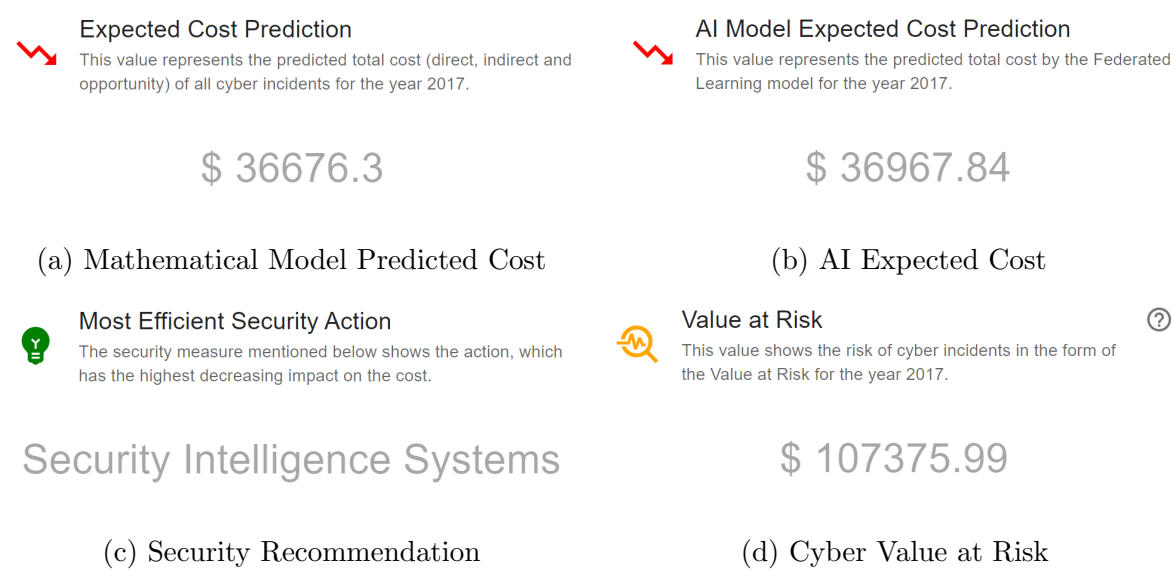


Figure C.4: Numerical Outputs of the RCVaR

This summary hopefully helps you to better understand the Real Cyber Value at Risk and improve your company’s cyber resilience. For a more detailed understanding, please consider reading the thesis.

Appendix D

Installation Guidelines

The code for this thesis is divided into two components. One component, which is referred to as *Research Code*, is used to conduct statistical tests, generate graphs and data. The other component is the code that is responsible for the website and is referred to as *Website Code*. The *Website Code* consists again of two parts: Frontend and Backend.

The overall requirements to run either component are listed in Table D.1.

Model	Version	Downloads
<i>Python</i>	3.9.13	https://www.python.org/downloads/
<i>Pip</i>	22.0.4	https://pip.pypa.io/en/stable/installation/
<i>Git</i>	2.39.0	https://git-scm.com/download/win
<i>Node.js</i>	16.14.0	https://nodejs.org/en/download/

Table D.1: Technical Requirements to Run the Code

After *Git* is installed the repository can be cloned into the desired location with the following command:

Command

```
git clone https://gitlab.com/FinanceLecture/cybervar.git
```

1. *Research Code*

To run the *Research Code* you have to navigate into the project and execute the file *Preparation/main.py*:

Command

```
cd cybervar/Preparation  
python main.py
```

If the centralized machine learning part should be executed the file *Preparation/-model_training/main_model_training.py* needs to be run. To train the model with Federated Learning, the file *Preparation/model_training/federated_learning/FL_main.py* needs to be executed.

Command

```
cd cybervar
python -m Preparation.model_training.main_model_training

OR

cd cybervar
python -m Preparation.model_training.federated_learning.FL_main
```

Before running any file from the *Research Code*, the modules in Table D.2 need to be installed.

Model	Version	Command
<i>NumPy</i>	1.23.2	pip install numpy
<i>OpenCV</i>	4.6.0	pip install opencv-python
<i>PiWin32</i>	302	pip install pypiwin32
<i>SciPy</i>	1.9.1	pip install scipy
<i>Pandas</i>	1.2.4	pip install pandas
<i>Statsmodels</i>	0.13.2	pip install statsmodels
<i>Matplotlib</i>	3.5.3	pip install matplotlib
<i>Seaborn</i>	0.11.2	pip install seaborn
<i>OpenPyXL</i>	3.0.9	pip install openpyxl
<i>Sklearn</i>	1.1.2	pip install scikit-learn
<i>TensorFlow</i>	2.10.0	pip install tensorflow
<i>Flower</i>	1.0.0	pip install flwr

Table D.2: Technical Requirements to Run the Code

During the execution of the *Preparation/main.py* file, graphs will be generated automatically. The program will continue the execution once the graphs are closed. Moreover, the program is interrupted by one input request and only continues after an input is provided. If the current data should not be changed enter: *No* or *N*.

If the *Research Code* is executed in an editor such as PyCharm, one has to make sure that the *Working Directory* is set to the root project folder (*cybervar*) for the execution of the machine learning files.

2. Website Code

To run only the *Website Code*, one has to navigate to *cybervar/Website/RCVAR* and install the libraries required for the Backend listed in Table D.3. Afterwards, the Backend can be run with the following command:

Command

```
cd cybervar/Website/RCVAR
python run_rcvar.py
```

Model	Version	Command
<i>PyYaml</i>	6.0	pip install pyyaml
<i>TensorFlow</i>	2.10.0	pip install tensorflow
<i>Flask-Cors</i>	3.0.10	pip install flask-cors
<i>Flask-SocketIO</i>	5.3.1	pip install Flask-SocketIO
<i>Pandas</i>	1.5.1	pip install pandas
<i>Matplotlib</i>	3.5.3	pip install matplotlib
<i>SciPy</i>	1.9.1	pip install scipy
<i>Sklearn</i>	1.1.2	pip install scikit-learn

Table D.3: Technical Requirements to Run the Code

Once the Backend is running, the Frontend, for which *Node.js* must be installed first, can be started. The following commands navigate to the Frontend folder, then install the necessary packages and start the program in the end.

Command

```
cd cybervar/Website/RCVAR_Frontend/rcvar
npm install - -force
npm start
```

The installation process was tested on both Windows and Linux. Thus, should also be applicable for MAC OS. The Tables D.3 and D.2 list all necessary packages which have to be installed via Pip-command. Nevertheless, it is possible that the user must install additional dependency packages.

Appendix E

GitLab

Both the *Research Code* and the *Website Code* are available on GitLab under the following link:

<https://gitlab.com/FinanceLecture/cybervar>