The Economics of Cyber Risk Transfer

Daniel Woods

University College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2019

Acknowledgements

Thanks above all go to Andrew Simpson for making his students a priority. I enjoyed regular meetings and email responses, thorough feedback, and genuine care for my well-being throughout the DPhil. I am one of the few students who was encouraged to take *more* time off.

Beyond the big man's supervision, I am also grateful to Tyler Moore for his evangelism about empirical research and for patiently explaining the rules of American football as TU lost yet another game. Jason Nurse and Ioannis Agrafiotis supervised my first forays into research and suffered the worst of my transition away from writing like a mathematician. Chad Heitzenrater acted as an academic big brother in introducing me to his WEIS friends and sharing his models (toys).

Thanks to everyone involved in the CDT. The bond forged over six months of 9–5 lectures lives on. Aaron was a fountain of all round sagely wisdom. The revolutionary side of my family should thank Dennis for being a counter-balance to the economics literature. Any and all emotional development over the course of the DPhil was down to my partner in falafel and decaf tea, Adam.

Love to all the friends who brought balance to my life, often against my will. The Cruce continued to cruce, albeit less frequently. It was a joy to live with all of the inhabitants of Magpie Lane, Merton Street, Percy Street, and East 4th Street—the Percy Street strollers in particular. I had pretty good housemates growing up as well.

Abstract

Risk transfer plays an increasing role in information security risk management as organisations purchase cyber insurance and vendors offer cyber warranties. These cyber risk transfer products affect how risk managers make decisions. An archetypal example is insurers offering discounts on cyber insurance contingent on information security controls being in place. Alternatively, vendors offering cyber warranties incur relatively less cost if they produce more effective products, increasing the information risk managers possess when purchasing security products.

This dissertation uses mixed methods to ask *how might cyber risk transfer products increase information about security decisions?* Focusing on the incentives and strategies of market participants situates this dissertation within the Economics of Information Security. We collect empirical data in order to make realistic modelling decisions. We then introduce two decision-theoretic models to explore how mechanisms like cyber insurance and cyber warranties can increase information about the effectiveness of security controls. One of the resulting insights is operationalised by introducing a novel method to infer loss distributions from insurance prices.

Our first contribution collects data about cyber insurance risk assessment and how it feeds into pricing. A qualitative study involving nine insurance firms in the UK provides insights into market processes. We identify disparities between how an area of information security is valued by underwriters and how much information is collected in application forms. Additionally, we extract 26 regulatory filings describing how US insurers price cyber insurance, providing one of the first quantitative empirical studies of the cyber insurance market.

Our second contribution extends an existing model to consider multiple policyholders with an insurer coordinating information. Monte Carlo simulations are used to explore different strategies for the insurer. The results describe how the rate of attack, security spending, variance of losses, and gross return relate to the insurer's choice of strategy and number of insureds.

Our third contribution considers how consumers can use cyber warranties to increase information about the effectiveness of security products. We analyse 10 warranties attached to information security products to understand what they typically cover. We then introduce a simple model and derive four different inferences to be made, depending on the information held by the consumer. Numerical illustrations suggest vendors voluntarily offering warranties can force a separating equilibria. Finally, we discuss barriers to making these inferences in reality.

Our final contribution introduces a novel method to infer cyber loss distributions from insurance prices. We apply this to a set of $6,218$ cyber insurance prices extracted in our first contribution. This allows us to derive what we term the *County Fair Cyber Loss Distribution*, which aggregates the inferred loss models from 26 separate pricing schemes. The results provide real estimates that organisations can use to quantify cyber risk.

Contents

List of Figures

1 Introduction

1.1 Overview

The security of information systems and the data they hold depends on the decisions of the actors managing the risk, which is encapsulated in the declaration "information security is risk management" [1]. Risk managers are tasked with assigning limited resources across a range of possible security controls and procedures. The optimal investment level is intimately linked to both the likelihood and the impact of potential attacks, as well as the effectiveness of different security controls in mitigating risk. Put simply, "risks cannot be managed better until they can be measured better" [2].

This view suggests information security depends on measuring cyber risk. Yet quantifying potential losses is limited by the data available. Organisations face disincentives to report incidents because of the resulting reputation damage. Even where mandatory reporting laws force organisations to disclose data breaches, the data does not relate to financial losses and the population from which it is drawn is unknown.

Uncertainty about potential cyber losses is compounded by widespread ignorance about the effectiveness of different security controls. This can be illustrated by considering that buyers cannot easily determine the quality of information security

products. Performance must be evaluated against active adversaries who update their attack strategies in response to the actions of defenders. The resulting information asymmetry leads to a market for lemons [3] in which vendors face little incentive to provide higher quality products because buyers struggle to identify quality. Akerlof [3] identified mechanisms to address this including brand reputation, certification, liability laws, and warranties.

Establishing a reputation for effectiveness is undermined by the infrequency of incidents undermines this process. Romanosky [4] found incident rates of less than 0.3% per year in the majority of industries considered in his study. Even if a product halved the attack frequency, a randomised control trial should expect to need over $8,600$ organisations to participate to reject the null hypothesis that the security control has no effect at a 95% confidence level. Informal reputation systems and consumer reports would need an even larger sample size to establish effectiveness, especially given many incidents remain private.

Instead external experts could certify the effectiveness of the product. Yet experts often face incentives to conduct less rigorous on assessment. For example, the Common Criteria framework "motivated the vendor to shop around for the evaluation contractor who would give his product the easiest ride" [5]. There are also difficulties in using laboratory experiments to establish real world security.

Liability laws could shift the costs of an ineffective product back onto the vendor. This might incentivise vendors to create more effective products and even force firms selling ineffective products out of the market. However, the resistance to software liability is well documented [6, 7] and establishing proximate cause is difficult given the constellation of security controls employed by firms. This leaves consumers without traditional mechanisms or institutions to address information asymmetry about effectiveness.

Risk transfer products could help identify effective security products. Security vendors have started offering cyber warranties — voluntary obligations to indemnify the customer in the event of a cyber attack — to function as a quality signal. Much like how consumer protection laws are relatively more costly to firms offering

low quality products, cyber warranties are more costly for firms developing low quality enterprise security products. These warranties can increase information about the effectiveness of security products.

Cyber insurance is another form of risk transfer that may identify effective security controls. For example, Bruce Schneier [8] described a world in which the "computer security industry will be run by the insurance industry". Schneier's vision was motivated (at least in part) by an analogy with property insurance where businesses install alarms in their warehouses to "get a break in their insurance rates". Importantly, insurers will recognise "'snake-oil' peddlers" [8] and prevent policyholders from wasting resources on useless products. The nascent cyber insurance industry's failure to deliver on initial growth predictions means that such a world is not yet a reality. However, the underlying principle that the insurance industry will improve information security management has caught the attention of the policy-making community.

Policy makers have begun to seriously consider the impact and viability of such a partnership. At the 2003 Homeland Security Council[1] Paul B. Kurtz argued that "The Insurance industry has a pivotal role in play [in protecting our national infrastructure], particularly by developing cyberinsurance policies." Less than ten years later, the US Department of Homeland Security [9, 10, 11, 12] convened a 'Cybersecurity Insurance Workshop' and a series of other round-tables and working groups. The European Union Agency for Network and Information Security (ENISA) [13, 14] published an analysis of the barriers to a European cyber insurance market in 2012, with a follow-up report in 2016. In 2015, the UK Cabinet Office [15] published a policy paper titled 'UK cyber security: The role of insurance'. These discussions illustrate the progression of the market from one security expert's vision to an issue worthy of the attention of governments.

We have argued that cyber insurance and cyber warranties can increase information about effective controls. The insurance industry can also provide insight into the likelihood and impact of cyber losses. Actuarial science applies statistical

[1] Cyber-Insurance Metrics and Impact on Cyber-Security, Internet Society Alliance: `http://bit.ly/2oFIQIi`

methods to assess risk. The industry's expertise could help quantify potential cyber losses, despite the guesswork involved [16].

Our over-arching research question is *how might cyber risk transfer products increase information about security decisions?* A comprehensive answer is beyond our scope (and possibly unknowable), so we instead consider specific decisions in isolation. Chapter 5 investigates insurer strategies for aggregating claims information to increase knowledge about which security controls policyholders should employ. Chapter 6 explores how consumers can use cyber warranties to identify effective security controls. Chapter 7 introduces a method to infer cyber loss distributions from cyber insurance prices.

The dissertation adopts mixed methods. First we collect empirical evidence about the cyber insurance market. Later chapters will use these observations to justify modelling decisions. These models describe abstract processes by which risk transfer products increase information. Chapter 7 operationalises one such process to provide insights of interest to real world risk managers. By considering the economics of cyber risk transfer, this dissertation contributes to the emerging and interdisciplinary field of information security economics [2].

1.2 Contributions

This section identifies which parts of this DPhil dissertation represent novel contributions. Figure 1.1 describes how the chapters map to contributions and publications.

C1: consists of an empirical analysis of cyber insurance risk assessment (**C1a**) and a quantitative study of pricing (**C1b**). The risk assessment section is inspired by our paper published in the Journal of Internet Services and Applications in 2017 [17]. We add to this analysis of insurance documents by involving insurance professionals. Our 2017 paper [17] was the (joint) first empirical analysis of the cyber insurance assessment process. The pricing section describes how risk assessment information is used to price cyber insurance, these results formed part of [18].

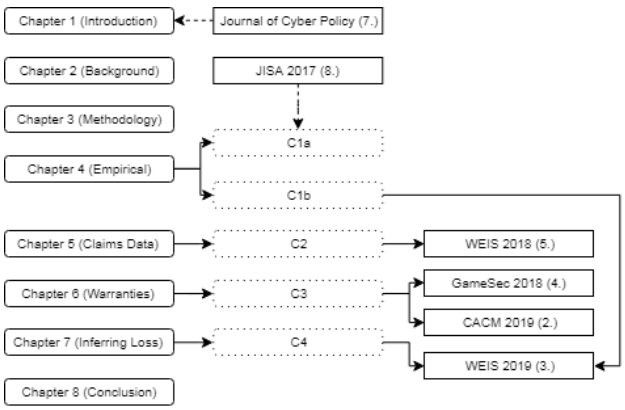

Figure 1.1: How chapters map to contributions and publications. Dotted arrows indicate the paper influenced that chapter or contribution, whereas solid arrows indicate the paper was directly imported. The paper numbers correspond to the list of publications in Section 1.4

C2: This is the first cyber insurance model in which the insurer has uncertainty regarding which security measures are effective. We believed this to be an important consideration after speaking to insurers who reported uncertainty about how to assess applicant's security level. Furthermore, it models insurers aggregating claims information, which is often discussed without formal study. It was published at the Workshop on the Economics of Information Security (WEIS) in 2018 [19].

C3: This contribution is the first academic consideration of cyber warranties. It focuses on how consumers can use warranties to identify effective security controls. One of the insights inspired C4. It was published in the proceedings of the 2018 Conference on Decision and Game Theory for Security [20]. This contribution also includes two sections from a viewpoint article submitted to the Communications of the ACM [21].

C4: We infer cyber loss distributions from cyber insurance prices. This line of research has not been considered in actuarial science and it represents a new technique. It provides insights into attack types not previously investigated. It was published at the Workshop on the Economics of Information Security (WEIS) in 2019 [18].

Some of the research conducted during the course of the DPhil was not directly included. This includes an article in the Journal of Cyber Policy [22] and a workshop paper at Cyber Science 2018 [23].

1.3 Structure

Figure 1.1 provides a high-level view of how chapters relate to the contributions described in the previous subsection. This subsection provides an overview of the content of each chapter.

Chapter 2 provides the necessary background for the reader. We introduce theories about how insurance impacts behaviour, as well as critical views of the insurance industry. We then turn to specific considerations of cyber insurance in the economics of information security. This predominantly consists of investigations grounded in economic modelling, but we also survey the limited number of empirical studies. We then explore existing approaches to quantifying cyber losses, as Chapter 7 introduces an alternative approach.

Chapter 3 justifies approaching this problem using both economic models and empirical studies. It also describes the data sources and data collection decisions for our empirical work. We describe a method for analysis used in Chapter 4 and a way to simulate the results of Chapter 5.

Chapter 4 describes cyber insurance risk assessment and pricing based on our empirical observations. The first part investigates Lasswell's problem [24] of "who says what in which channel to whom with what effect?" in the context of the cyber insurance application process. The second part of the chapter asks how the collected information determines the pricing of cyber insurance. These findings will be used to justify our modelling assumptions in Chapter 5 and Chapter 6, as well as providing the data to operationalise the method we introduce in Chapter 7.

Chapter 5 considers insurer strategies to share claims information and influence the insureds' security postures. We extend an information security investment model to approach this problem. Monte Carlo methods are used to explore three possible insurer strategies in guiding the policyholder's investments. We discuss

how realistic our modelling assumptions are and how to translate the simulation results into actions insurers can take. However, this consideration leaves consumers dependent on insurers.

Chapter 6 asks how consumers can identify effective controls using cyber warranties — voluntary obligations to indemnify the customer in the event of a cyber attack — as a quality signal for information security products. We collect 10 warranties and provide a preliminary analysis of what they typically cover. We then introduce a decision-theoretic model to explore how consumers might use cyber warranties to increase information when purchasing security products. Our analysis derives four inferences that consumers can make about a security product. The discussion reflects on the difficulties of making these inferences in the real world.

Chapter 7 introduces a novel technique to derive cyber loss distributions from cyber insurance prices. This chapter is inspired by a theoretical result from Chapter 6. By deriving fully parameterised loss distributions for various cyber incidents, we demonstrate how cyber risk transfer products can increase information about cyber risks.

Chapter 8 concludes the dissertation by reflecting on the disparity between theoretical predictions and empirical observations in the context of cyber risk transfer. We link our contributions back to our over-arching question of *how might cyber risk transfer products increase information about security decisions?*

1.4 Personal Contribution to Publications

This section describes my personal contribution to different publications and which parts are included in this document. I am an author on the following publications [17, 18, 19, 20, 21, 22, 23, 25, 26, 27]:

1. Sakshyam Panda, Daniel W Woods, Aron Laszka, Andrew Fielder, Emmanouil Panaousis. Post-Incident Audits on Cyber Insurance Discounts. *Computers & Security*, to appear.

2. Daniel W Woods and Tyler Moore. Cyber warranties: Market fix or marketing trick. *Communications of the ACM*, to appear.

3. Daniel W Woods, Tyler Moore, and Andrew C Simpson. The county fair cyber loss distribution: Drawing inference from insurance prices. In *Proceedings of The 18th Workshop on the Economics of Information Security (WEIS 2019)*, 2019.

4. Daniel W Woods and Andrew C Simpson. Cyber-warranties as a quality signal for information security products. In *International Conference on Decision and Game Theory for Security*, pages 22–37. Springer, 2018.

5. Daniel W Woods and Andrew C Simpson. Monte carlo methods to investigate how aggregated cyber insurance claims data impacts security investments. In *Proceedings of The 17th Workshop on the Economics of Information Security (WEIS 2018)*, 2018.

6. Daniel W Woods and Andrew C Simpson. Towards integrating insurance data into information security investment decision making. In *2018 International Conference On Cyber Situational Awareness, Data Analytics And Assessment(Cyber SA)*, pages 1–6. IEEE, 2018.

7. Daniel W Woods and Andrew C. Simpson. Policy measures and cyber insurance: A framework. *Journal of Cyber Policy*, 2(2):209–226, 2017.

8. Daniel W Woods, Ioannis Agrafiotis, Jason RC Nurse, and Sadie Creese. Mapping the coverage of security controls in cyber insurance proposal forms. *Journal of Internet Services and Applications*, 8(1):8, 2017.

My main contribution to (1) was framing the problem. I wrote the introduction and related work alone, only receiving comments and suggestions from the other authors. My ideas and experience informed the modelling choices. I made no contribution to the theoretical analysis or conducting the simulations, although I did

identify parameter values and wrote the justification for them. A few paragraphs from the related work were included in Chapter 2.

For all of the other papers I am the only student author. Publication (2) and (3) were completed while visiting the University of Tulsa. Tyler and I spent time discussing ideas. I collected all the data, wrote all of the analysis code, and wrote a number of first drafts alone, again only receiving comments and suggestions from both Andy and Tyler. However, for the final drafts I provided the document to Tyler and he edited a draft. I retained the document before Tyler's final edit.

(4) through (7) followed a common work pattern with Andrew Simpson. We discuss ideas and he provided weekly feedback on an ongoing draft document. However, I collected all of the data and wrote all of the code. The final draft was often turned over to Andy for copy-editing.

Finally, (8) was the result of a 10 week mini-project working directly with Jason Nurse and Ioannis Agrafiotis. Sadie Creese suggested investigating security controls. I conducted all of the analysis and wrote most of the report, possibly receiving some editing from Jason and Ioannis. However, none of this writing or analysis will be directly imported into this dissertation.

Chapter 2 contains many words imported from the related work of (1–8). Chapter 3 contains words from the data collection section of (3). Chapter 4 contains many figures from a part of (3). Chapter 5, Chapter 6 and Chapter 7 are imported from (5), (4) and (2), and (3) respectively. This involves directly copying sections from the papers, with the introduction and related work removed. The discussions were re-written to reflect being part of a wider document. Chapter 8 consists of new text.

2
Background

This chapter introduces the reader to the theory and prior results necessary to situate our contributions. In Section 2.1 we define the terminology we will use and identify some important processes in the cyber insurance market. Section 2.2 motivates this project by considering a theory of insurance and how it interacts with risk mitigation. Section 2.3 identifies relevant theoretical work on cyber insurance and Section 2.4 describes related empirical studies. Finally, we conclude by linking related work to the subsequent chapters in Section 2.5.

2.1 Terminology

Insurance is a financial contract between a legal person (*policyholder* or *insured*) and an insurance company (*insurer*). Although these terms suggest individual persons, throughout they will refer to organisations and their representatives unless otherwise stated. For example, the *insured* is used to refer to the purchasing organisation including the collection of individuals responsible for preparing the application, complying with the terms of the contract, and other such tasks.

The contract, known as an *insurance policy*, typically states that the insured party will pay a regular *insurance premium* in exchange for a promise of *indemnification*. The *insurance premium* is a defined amount of money the insured party pays to the

insurer. Meanwhile, *indemnification* is a financial payment intended to compensate the *insured party* in the event of a loss defined in the *insurance policy*. The fact that the event may or may not occur means that the indemnity payment is uncertain, unlike the insurance premium which is paid to the insurer independent of events.

Before a policy is bought, the individual will be referred to as an *applicant*. The insurer must decide whether to offer the insurance policy to the applicant. The process by which an insurer assesses a risk and then offers the contract is known as the *underwriting process*. In this role, the insurer may be referred to as the *underwriter*, which reflects the insurer's expertise in assessing risk. In the event of a loss, the insurer may send a *loss adjuster* to investigate the causes and impact of a loss.

Warranties represent an alternative form of risk transfer. A *vendor* might decide (or be forced) to attach a warranty to products. In general there is no additional cost to the *consumer*, although extended warranties may be offered for an additional fee. Much like an insurance contract, a warranty constitutes a promise of indemnification. *Cyber warranties* — voluntary obligations to indemnify the customer in the event of a cyber attack — will be considered in Chapter 6.

2.1.1 Defining Cyber Insurance

Cyber insurance is used as an umbrella term to describe a range of insurance policies. Defining cyber insurance is no simple task, given it is an evolving product. Tyler Moore offered an *intensional definition* (which tries to give the essence of a term) in a workshop organised by the US Department for Homeland Security [9] by defining cyber insurance to mean

> "a contract between an insurance carrier and a company that covers financial losses to the company resulting from damages caused by computer or network-based incidents."

The main objection to this definition is the lack of specificity. It is unclear which financial losses and incidents are covered by cyber insurance.

The collection of coverages outlined in Table 2.1 offers an *extensional definition* (a list of objects that a term describes) of cyber insurance. Siegel et al.[28] identify

Coverage	What is typically covered
First-Party Coverage	Coverage for the cost of replacing or restoring lost data. Excludes intellectual property.
Data Privacy and Network Security Liability	Coverage for liability claims of a third party (e.g. a data breach or unintentional transmission of a computer virus).
Business Interruption	Covers revenues lost as a result of network down time.
Cyber-Extortion	Cover for investigation costs, sometimes the extortion demand.
Public Relations	Fees for Public Relations firm to manage reputation in the event of a breach.
Multi-Media Liability	Costs relating to the content of a firm's website like copyright infringement.
Professional Services	Liability relating to a service offer such as web hosting or internet service.

Table 2.1: The range of coverage available.

the coverage outlined in Table 2.1. The results were corroborated by Baer and Parkinson who examined six insurance policies in [29] and Marotta et al. who analysed 14 policies in [30].

An exhaustive list of coverages cannot be expected as insurers regularly offer policy extensions in response to the dynamic nature of cyber risks. For example, one study by Majuca et al. [31] looked at seven different policies concurrently offered by one insurer, with optional coverage including physical theft of hardware, criminal rewards and crisis communication. The diversity of coverage illustrates the variation in policies offered by different insurers and even by the same insurer. A historical perspective sheds light on how cyber insurance has evolved.

2.1.2 A History of the Cyber Insurance Market

The history of cyber insurance is disputed. Marotta et al. [32] identify computer crime as the driver of the first insurance policies relevant to cyber risk in the late 1970s. Majuca et al. [31] claim that the first standalone Internet-based insurance policies were the "hacker insurance policies" of the late 1990s, in which an insurer partnered with a technology company to offer a policy covering the insured firm's first-party loss. They both agree that the coverage that traditional policies offered left significant gaps as firms outside the technology industry became increasingly dependent on their networks in the early 2000s.

Most business insurance policies covering tangible property excluded liability related to electronic data loss [33]. In response to this, insurance companies started to offer standalone cyber insurance policies. These policies are broken down into a number of sub-policies, with coverage offered for a specific set of risks. Though these sub-policies have evolved over time as the specific risks change, this model of standalone insurance has stayed constant.

In the DHS Workshop [9] alluded to in Chapter 1, Emily Freeman identified a number of key pieces of legislation in driving demand for cyber insurance. These were the passage of the Health Insurance Portability and Accountability Act of 1996 (HIPAA)[1] and California's data breach law in 2003[2]. She then mentioned that recent data breaches suffered by retailers and the ensuing court cases have caught boardroom attention and demand for insurance has increased accordingly. Common here is the effect of new legislation and court cases in increasing demand for cyber insurance. Note the General Data Protection Regulation (GDPR)[3] came into force in the European Union in 2018 and could boost demand for cyber insurance in Europe, which has traditionally lagged behind the US in this respect.

The disparity between uptake of cyber insurance in Europe and the US is illustrated by the following surveys. In the US, a 2015 market update[4] reveals that 26% of companies with a revenue of US$5 billion or more have cyber insurance, with 3% of those who return less than US$500k carrying insurance. In the UK, a 2015 report revealed that 2% of large companies use standalone cyber insurance while cyber insurance penetration is 'negligible' for smaller firms [15]. These surveys are corroborated by the UK Government report [15] revealing that "majority of exposure is concentrated in US (up to 95%)".

Today, cyber insurance is among the fastest growing areas of insurance. A report by PWC[5] estimated that global cyber insurance premiums were at US$2.5

[1] `https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html`

[2] `http://bit.ly/2tMfvQ7`

[3] `http://www.eugdpr.org/`

[4] Advisen Insight Cyber Insurance Market Update (`http://bit.ly/2sokjx8`)

[5] PWC, Insurance 2020 & beyond (`http://pwc.to/2ppT17P`)

billion in 2015, which they predicted would rise to at least US$7.5 billion by 2020. This trend is unlikely to slow as the world becomes increasingly dependent on information systems. Indeed, the Institute of Risk Management[6] identified cyber risk as one of the leading risks for businesses in 2017.

The next section describes theories about the societal impact of the insurance industry.

2.2 Insurance Theory

Section 2.2.1 describes a framework for how insurance affects risk management. We identify a number of critics of the insurance industry in Section 2.2.2. We apply the framework [34] to consider how cyber insurance affects information security in Section 2.2.3.

2.2.1 Insurance Theory

Zweifel and Eisen [34] identify six functions of insurance:

1. Efficient allocation of risk
2. Protecting existing wealth
3. Capital accumulation
4. Mobilisation of capital
5. Fostering governance
6. Financial relief to the government

Each will be described in terms of general insurance theory, before turning towards their relevance to cyber insurance in Section 2.2.3.

Efficient allocation of risk (1) sees the enlightened insurer support risk management. This may take the form of "inspection, auditing, and consulting" [34] and the subsequent risk treatment reducing risk ex-ante. Alternatively, insurance may

[6]IRM Risk Predictions 2017 (`http://bit.ly/2r8YotD`)

limit the size of losses ex-post by immediately providing funds to prevent further damage. The source of this expertise can be traced to "research into the causes of losses with the goal of limiting or even eliminating them" [34].

An insurance contract *protects existing wealth* (2) by financially compensating the policyholder for losses suffered as a result of a pre-defined event. It has been suggested this provides a more stable basis for planning. Indeed, Arrow's landmark paper [35] describes how an individual might rationally purchase medical insurance in the presence of uncertainty. However, the applicability of this concern to a publicly traded company whose shareholders have other diversification options has been called into question [34].

Capital accumulation (3) results from a combination of premiums being collected before any losses occur and insurers charging more in premium than is expected to be lost as claims. The resulting reserves are particularly relevant in the event of a catastrophic event which may affect a firm or individual's ability to raise funds. For example, a home owner may struggle to secure a loan against the value of a home in an area recently affected by flooding. The home owner should be able to access capital set aside by the insurer for this eventuality. Further, these funds may be invested productively to the benefit of the economy, contributing to the *mobilisation of capital* (4).

Insurance *fosters governance* (5) by financially quantifying decisions related to risk management. Insurance theory suggests that less risky decisions will result in lower premiums relative to high risk decisions. As a result, the enlightened insurer provides risk managers with financial "incentives to avoid excessive risks".

Finally, governments are sometimes expected to support individuals and organisations if a disaster occurs. For example, solidarity may lead taxpayers to contribute towards government led flood relief initiatives. The need for such contributions would be decreased if private insurance is purchased, providing *financial relief to the government* (6).

The final function reframes the beneficiaries of the insurance contract. Whereas the first five functions benefit the policyholder, the final function relates to the

benefits accrued to parties beyond the policyholder. These benefits may accrue to both the policyholder and the wider public. For example, income protection insurance both protects the individual from variable income and also lessens the load on the taxpayer in terms of welfare provision. Governments believe this function justifies their involvement in insurance markets, which we considered in [22] but exclude from this dissertation.

Describing these functions in the abstract obscures how they interact with the risks covered by a particular line of insurance. For example, financial relief to the government is far more relevant for health or life insurance than commercial property insurance because governments may otherwise be expected to provide healthcare, whereas shareholders usually absorb losses related to commercial property damage.

The first five functions can be broadly separated into improving risk mitigation and managing capital, which is out of scope. *Efficient allocation of risk* involves the insurer directly interfering in risk decisions, whereas the insurer provides indirect incentives for better risk management in *fostering governance*. The *protecting existing wealth* function seeks to make the insured's demand for capital more certain, *capital accumulation* concerns how much capital is available to cover losses, and *mobilisation of capital* concerns how the capital is invested.

Although the capital management functions determine the availability of funds for ex-post incident recovery, this dissertation focuses on how the risk mitigation functions affect security decisions.

2.2.2 Critics of the Insurance Industry

The previous subsection provides an idealised view of insurance. Zweifel and Eisen [34] predominantly focus on economic models. Models are often based on assumptions that abstract away from the intricacies of the market. We will reflect on the limits of economic modelling further in Chapter 3. But for now we will identify sociological works that can shed light on how insurance works in practice.

Giddens [36] contrasted manufactured risks with external risks in 1999, before cyber risks were widely identified as troubling. External risk has been stable enough

over time to be broadly predictable, whereas manufactured risk is created by the "progression of science and technology" [36]. It intrudes directly into personal life so that "one can no longer simply rely on tradition" [36]. The importance of this distinction is made clear by considering Beck's "risk society" [37].

Beck [37] observed how pollution risks are often produced without people or organisations being held responsible. This leads to the claim "smog is democratic" [37] in the sense that everybody suffers from air pollution. Yet some cannot afford to mitigate the consequences, by, for example, moving to the suburbs. A similar dynamic can be seen in the Equifax breach [38] in which millions of financial records were lost. Some can mitigate the consequences with private insurance. For example, Equifax were indemnified millions of dollars by a cyber insurance policy. In this way, cyber insurance may enable some parties to *produce risk* that other parties suffer the consequences of.

Giddens suggest we cannot calculate manufactured risks "accurately in terms of probability tables" [36], Ericson and Doyle [16] go further by suggesting this is the case for all insurance. Life insurance is partly "based on guesswork in support of gambling, with both unexpected windfalls and catastrophic losses" [16] despite the deference shown towards actuarial science. This work justifies skepticism about the expert authority of insurers.

The processes of insurance have problematic aspects. Zweifel and Eisen [34] suggest insurers audit policyholders to ensure a more efficient allocation of risk. Yet external audits can have perverse consequences. Audits incentivise making processes "externally verifiable via acts of certification" [39], which may run counter to the stated goal of creating internal improvement. Power also shows how best practice can be "taken for granted" [39] and evade meta-evaluation to probe the efficacy of the audit itself. Power's critical lens reminds us to be wary of unintended consequences of insurance as governance and to question audit efficacy.

Hilgartner focuses on sociotechnical networks in which different actors contribute to the "construction and control of risk objects" [40]. He identifies how Ralph Nader's Unsafe at Any Speed [41] pushed risk controls like built-in over-steer into

the public discussion. Analogously, the insurance industry constructs risk objects and the resulting controls. For example, the application process collects information about security controls, which may be used by risk managers to justify investing in those controls even if it is against the organisation's best interest.

Finally, Power [42] described how insurance companies developed "risk management consultancy and advisory practices" in response to organisations self-insuring. Quigley et al. [43] suggest that "cyber gurus" misrepresent their expertise to confer legitimacy [43]. This ties into an existing body of theory about how actors "profit from selling solutions to complex organizational issues" by persuading "audiences of the usefulness of their ideas" [43]. In this fashion, insurers may face incentives to "over-dramatize and over-simplify cybersecurity risks" [43].

These ideas will be invoked in later chapters to push back against the idealistic enthusiasm that often accompanies economic modelling. For example, the role of guesswork in pricing life insurance [16] undermines the assumption about accurate pricing that Chapter 7 relies upon.

2.2.3 Cyber Insurance Theory

Cyber insurance was predicted to impact security behaviour before any evidence was collected. In 2001, Schneier [8] predicted a world where every organisation would purchase network security insurance with discounts offered for security-enhancing decisions like replacing an "insecure Windows system" with a more "secure version of Linux". The optimistic essay ends by suggesting "good security [will be] rewarded in the marketplace" as insurers recognise "computer security *snake-oil* peddlers" (Schneier's emphasis).

A 2008 report [44] echoed Schneier's vision of premium discounts as incentives [8], while also suggesting discounts could be used as security metrics. The rule of thumb being an investment with a bigger discount is more effective. These discounts would be based on knowledge generated by aggregating claims data or even commissioning primary research. Despite the report being commissioned by a policy making

institution, the authors offered a *wait-and-see* conclusion regarding possible policy measures in support of the market.

A review of theoretical models of cyber insurance in 2010 revealed that "positive expectations about cyber-insurance have not been analyzed rigorously" [45]. Informal conjectures claimed the market would "improve information about security levels", "affect agents' choices of network products" and "aggregate information" without any corresponding parameters or model features. The authors observed further "we are not aware of any quantitative empirical work on cyber-insurance markets" [45].

Yet policymakers in the United States and the European Union latched onto the idea insurers can incentivise better risk management. Reports, collated in [22], released from 2012 onward by public institutions in the US and the EU suggest insurance contracts will contain "prescribed security controls and procedures" that the policyholder must implement for coverage to be valid. Insurers would additionally provide access to a network of incident response professionals and conduct forensic investigations of losses, which could be aggregated for analytical purposes.

Talesh [46] explored the topic by collecting evidence from industry conferences, educational webinars, and marketing documents. This self-reporting by insurers supports the conclusion that "insurer-sponsored help is greatly appreciated by organizations", which can be justified by revealed preference, but suggesting this leads to a "net benefit" (p.437) is premature. This is especially true given Talesh admits that "the value of these insurer-sponsored risk management services" is an open question. He claims that security incentives are better aligned than in directors and officers insurance where organisations resist "insurers interfering with their day-to-day decision making". This belies a mistaken belief that a lack of knowledge drives security failures, a view that has long since been debunked [2].

We collect together a list of mechanisms by which insurers are claimed to affect security levels:

C1 Assess Security Levels: "measure the organization's practices and make sure they are consistent with the prevailing security standards" [46].

C2 Incentives for Investment: choice of systems and security controls "will be strongly influenced" [8] by premium discounts.

C3 Create Security Obligations: contracts will contain "prescribed security controls and procedures" [22].

C4 Access to Response Services: provide "a menu of services that an organization can quickly access in the event of a data breach" [46].

C5 Generate Knowledge: aggregate claims information and conduct primary research to develop an understanding of cyber risk [44].

This dissertation will collect empirical evidence about **C1** and **C2** in Chapter 4. We introduce an economic model to consider strategies for generating knowledge (**C5**) in Chapter 5. The next two sections identify related theoretical models and empirical studies.

2.3 Economic Models

Researchers have introduced many economic models of the cyber insurance market. Section 2.3.1 uses an existing framework to classify approaches to modelling the cyber insurance market. We then contrast how two investment models represent uncertainty in Section 2.3.2 as we will adopt one of these approaches in Chapter 5.

2.3.1 Modelling Cyber Insurance

Böhme and Schwartz [45] proposed a framework to classify all cyber insurance market models, introducing a common terminology for an increasingly disparate body of research. In the following we discuss each of the five components of the modelling decisions that the framework identifies: (1) network environment, (2) demand side, (3) supply side, (4) information structure, and (5) organisational environment. We identify research illustrating these modelling choices where it is available. Doing so allows us to identify aspects of the market that have not yet been modelled.

Network Environment

The network environment is generally taken to be an *individual risk model*, which involves a series of nodes and the connections between them. Generally the models abstract from the particular systems, threats and connections they describe.

The *risk arrival* relates the *defense function* (mapping security level to probability of breach) and the *network topology* (connections in the network) to the probability of an event, which is then mapped to a loss event with a corresponding impact. This allows interdependent security and correlated risk to be modelled in a unified way, as each depends on the connections between nodes.

Examples Much of the early research focused on how interdependent security and correlated risk impact the viability of a cyber insurance market. Interdependent security occurs when the risk "depends on the actions of others" [47]. For example, all firms benefit if one firm's security control prevents it being infected and re-transmitting a virus or contributing to a denial-of-service attack. Optimists argued insurers could coordinate the resulting collective action problem [48], leading to a net social welfare gain and a viable market. Although later models [49, 50] introduced non-trivial network structures, the general argument suggests only insurers can internalise the positive externalities from interdependent security.

Skeptics instead focused on the "high correlation in failure of information systems" [51], citing it as a major impediment to the supply of cyber insurance. Böhme considers how the choice of platform impacts correlated risks in [52]. The author suggests government policies in support of a "diversity of systems" should be pursued. Böhme and Kataria [51] considered the "systemic interdependence of loss events within and between firms", suggesting insurance is best suited for risks with high internal and low global correlation. However, the "data sources at our disposal lack basic statistical requirements", making it difficult to identify whether cyber insurance risks fit these characteristics.

The relevance of interdependent security and correlated risk is intimately linked to how claims costs arise. Axon et al. [53] analyse 70 cyber insurance claims and

discover that data breaches, ransomware and non-compliance with legislation are the most common triggers of claims. Data breaches rarely correlate across companies. Interdependent security can be seen by considering that Target's 2013 data breach was traced back to "network credentials that were stolen from a third party vendor"[7].

Ransomware incidents provide an example of correlated risk because many claims can result from the same underlying cause. For example, the NotPetya attack is claimed [54] to have resulted in losses of over $3 billion when aggregated across multiple policyholders. Interdependent security results from firm A's compromise depending on firm B's compromise when firm B is the organisation from which firm A was exposed to the NotPetya malware. In the counter-factual that firm B avoided compromise by adopting security measures, firm A would also avoid compromise providing it was not otherwise exposed to the NotPetya malware.

Demand Side: Agents

In the framework of [45] agents make and invest in security choices for a given node, as well as bearing the financial consequences. It is these financial consequences that the insurance policy will indemnify against. Modelling choices include the heterogeneity of each agent in terms of the number of nodes controlled, size of loss, defense function and risk aversion. These define the characteristics of the agent and determine their *action space*. The *action space* defines the agent's ability to purchase full or partial insurance, invest in self-protection (reduces the likelihood of an attack succeeding) or self-insurance (reduces the impact of a successful attack on a given node), and ability to choose network formation. The authors [45] note that the final option has been under explored. The final consideration is whether the model is formulated as a *single-shot*, *repeated* or *continuous* game.

Examples Models often allow the policyholder to invest in risk mitigation, as well as insurance. Risk mitigation and risk transfer are alternately characterised as a complement or a substitute for mitigation. For example, Kesan et al. [55] show that cyber insurance and risk mitigation are complements provided that

[7] `https://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company/`

premiums are linked to security. This suggests that the cyber insurance market will increase security levels.

This is in contrast to an action space of self-insurance and risk mitigation considered by Grossklags et al. [56], in which self-protection and risk mitigation are substitutes. They showed that "there is almost always too little protection effort exerted compared to the social optimum" in a total effort game without a social planner. But surely the insurer can solve the collective action problem? Grossklags et al. show that "the common wisdom that having a central planner who decides upon security implementation always yields higher protection" [56] does not always hold. This highlights the importance of the choice of action space, network topology, and defense function.

Supply Side: Insurers

If the previous component of the market relates to the buyers of insurance products, this one relates to the suppliers. The first choice is whether the market consists of one insurer (*monopoly*), a few insurers (*oligopoly*) or many insurers (*competition*). The risk aversion and *markup* of the insurer can be chosen homogeneously (to be the same between insurers) or heterogeneously (to vary between insurers). The *markup* consists of the operating costs and profit of the insurer.

The *contract design* links the premium level to the security investment level and introduces the possibility of fines if this is observed to be lower than a certain level ex-post. Finally, insurers can transfer risk via financial products including *reinsurance*, *catastrophe bonds* or *exploit derivatives*. *Reinsurance* involves an insurer acting as an agent to purchase an indemnity contract that is triggered if the losses exceed a threshold. *Catastrophe bonds* and *exploit derivatives* are financial instruments that are triggered by the occurrence of a catastrophic event or the discovery of a vulnerability in a given system.

Examples In [52], Böhme investigates cyber insurance from the perspective of the insurer, as opposed to that of the individual insurance holders, which was previously the predominant perspective of study. The correlated nature of cyber risk is identified as an important consideration; this is contrasted to traditional indemnity insurance risks.

In a subsequent paper [51], Böhme and Kataria provide a framework for understanding the correlated nature of cyber risks both within a firm and at a global level. A variety of different cyber risks are classified in terms of each. For example, insider threat has a high correlation of risks to systems within a firm but low correlation to systems across the global market. In addition, they provide an empirical insight into the correlated nature of cyber attacks through the Leurre.com [57] honeynet project. The paper concludes that cyber insurance is most suitable for risks with high internal and low global correlation. From this the authors conclude systems managers should emphasise platform diversity. However, they do not consider how different security postures might affect the correlation of risks.

A copula-based model for pricing cyber-insurance is used in [58]; importantly it "considers nonlinear dependencies for correlated risks". The authors concede the approach is limited by the availability of historic data, as previously outlined in [58]. However, the paper establishes more detail regarding "specific security postures" as a future direction. In a 2011 journal paper [59], the authors further develop their copula-based model. In this paper, the number of computers is used as a proxy for π, despite their admission that π is more likely to depend on the number of vulnerabilities, as determined by a firm's "security posture". Incorporating security posture was suggested as a direction for future research, but to our knowledge this has yet to be addressed.

Information Structure

Many of the modelling parameters discussed above assume the existence of a measurable security level. Böhme and Schwartz define *perfect information* to be the situation where no uncertainty is present and *information asymmetry* to be

the situation where some players have private information not available to others. The authors [45] introduce *timing* to model *when* information arrives. Though the insurer may not be aware of the agent's behaviour before the event occurred, they can discover it afterwards in the loss adjustment process.

Two instances of information asymmetry that have been identified across all lines of insurance industry are *adverse selection* and *moral hazard*, known as hidden characteristics and hidden action respectively. Adverse selection occurs when an actor is more likely to seek insurance as a result of knowledge about their own risk profile. Moral hazard occurs when a policyholder engages in riskier behaviour in the knowledge that they are covered by an insurance contract. These are problematic when the insurer cannot establish an agent's risk profile before the contract is signed or when they cannot monitor a policyholder's behaviour.

Information asymmetry affects how the agent manages the nodes under its control and the insurer in distinguishing different types of agent (*adverse selection*) and their behaviour following signing a contract (*moral hazard*). Further, the uncertainty resulting from imperfect information about cyber risk motivates the purchase of cyber risk. Finally, both agents and insurers have imperfect information related to the effectiveness of risk controls.

Examples Timing tends to split into ex-ante risk assessment during the application process and ex-post investigation following an incident. Shetty et al. [60] investigated an insurer who could assess security levels perfectly or not at all, concluding that the latter cannot support a functioning market. Majuca et al. [31] showed that ex-ante assessments in combination with discounts for adopting security controls can lead to an increase in social welfare. A more recent model introduces stochastic uncertainty about the policyholder's security level [61].

None of these adverse selection studies consider the potential for insureds to misrepresent their security posture. Allowing malicious insureds to "subvert insurer monitoring" in both the application process and over the policy period was studied

in [62]. The analysis showed that no cyber insurance market could exist. However, we know from [63] that insurers audit insureds and refuse coverage for fraudulent claims.

Beyond ex-ante assessment, insurers make decisions regarding ex-post claims management. These decisions have received less attention. The impact of secondary losses on the policyholder's incentive to claim could lead to over-priced products [64]. Empirically it has been suggested insurers will refuse "claims arising from the insured's failure to maintain" [65] security levels, but the strategic aspects of insurers investigating the incidents leading to claims has not been considered[8].

Böhme and Schwartz [45] identified how "researchers write about how insurers will aggregate information about security (obtained from claims)" but do not directly model this. Nine years on from [45], the ability for insurers to aggregate claims information to reduce uncertainty about cyber risk has not been modelled. Chapter 5 introduces a model to address this gap in the literature.

Organisational Environment

Böhme and Schwartz [45] identify four stakeholders who may be included in cyber insurance models: regulator, ICT manufacturer, network intermediaries and security service providers. *Regulators* may enter the market to establish disclosure requirements, mandatory security controls, fines and subsidies for agents and insurers, mandatory insurance coverage requirements, and prudential supervision of the insurer. *ICT manufacturers* impact the defense function D and the system diversity. *Network intermediaries* such as ISPs can share information about observed threats and shape the network topology. Finally, *security service providers* may overcome information asymmetry by collecting and aggregating information from many agents.

Examples Modelling this aspect tends to make the defense function dependent on the actions of some third party. For example, models were introduced to explore the insurer's role in assessing the security of service providers and whether the insurer should invest in software quality. Khalili et al. [66] show that underwriting service providers improves both insurer profit and social welfare. Laszka et al. [67] found

[8] I contributed to a paper considering this [25] but did not include this analysis in this dissertation

that the insurer directly investing in software quality can "reduce non-diversifiable risks and can lead to a more profitable cyber-insurance market".

Zhao et al. [68] contrast the social benefits of cyber insurance, managed security services (MSSs) and risk pooling agreements (RPAs). MSSPs manage the security of multiple firms and "transfer their security risks" via service level agreements. The analysis concludes "MSSPs induce more efficient allocation of security resources [than cyber insurance and risk pooling agreements] because the MSSP, when serving all firms, internalizes the externalities of security investments between the member firms" [68]. This finding motivates considering how other security providers offering risk transfer might change the incentive structure — Chapter 6 will consider the impact of security vendors offering warranties with their products.

2.3.2 Security Investment Models

This subsection introduces two economic models that will be extended in Chapter 5 and Chapter 6. We first provide an overview of each model and then conclude by explaining what role they will play in this dissertation.

Rather than extend an insurance model to include security posture, we could instead extend an existing model of cyber risk in the context of the insurance industry. We look to investment models that ask: *how much should an organisation invest in information security?* Gordon and Loeb [69] conclude that the answer depends on the shape of the 'security breach function'.

The security breach function $S(\cdot, \cdot)$ relates the organisation's investment in security z_i and initial vulnerability v_0 to the probability of successful attack $S(z_i, v_0)$. The organisation's expected loss for a given investment level is calculated by multiplying the probability of successful attack by a fixed loss amount L. They established three core assumptions that a security breach function should fulfill:

A1: $S(z_i, 0) = 0$ for all $z_i \in \mathbb{R}$

A2: $S(0, v_0) = v_0$ for all $v_0 \in [0, 1]$

A3: $\frac{\delta S}{\delta z}(z_i, v_0) < 0$ and $\frac{\delta^2 S}{\delta z^2}(z_i, v_0) > 0$ for all $v_0 \in [0, 1]$ and $z_i \in \mathbb{R}$. Furthermore, for all $v_0 \in [0, 1]$ we have,

$$\lim S(z_i, v_0) \to 0 \text{ as } z_i \to \infty$$

The third assumption ensures that further investment always reduces the probability of attack, but does so at a diminishing rate. Further, no finite investment results in perfect security. In [69], Gordon and Loeb propose two classes to which the security breach probability function may belong. These will be used in Chapter 6 and may be expressed in the form

$$S^I(z_i, v_0) = \frac{v}{(\alpha z_i + 1)^\beta} \tag{2.1}$$

$$S^{II}(z_i, v_0) = v^{\alpha z_i + 1} \tag{2.2}$$

Unfortunately, the paper does not contribute a simple method for identifying the security breach function of a given vulnerability.

Indeed, Böhme [70] suggests that the problem with the Gordon and Loeb model is the direct mapping of security investment to vulnerability. Böhme instead suggests a direct mapping from investment to 'security level', which then stochastically maps to 'benefits of security'. The stochastic properties are a result of the indeterminacy of the attacker behaviour, which is assumed to be constant in Gordon and Loeb's model.

However, Gordon and Loeb [69] consider a one-time investment that cannot change in response to observed attacks. Böhme and Moore's Iterated Weakest Link (IWL) model [71] allows a defender to adopt a reactive investment strategy. The attacker's behaviour is modelled stochastically and the defender reduces uncertainty as this behaviour is observed. By introducing a parameter for uncertainty, the model provides a rational explanation for perceived under-investment in security — the defender is adopting a "wait and see" approach.

The IWL model provides three useful features for modelling the cyber insurance market: (i) an appropriate format for security posture; (ii) the uncertainty parameter; and (iii) a temporal dimension. For (i), the binary choice between protecting a given vulnerability or not is in keeping with the cyber insurance application process [17],

which predominantly consists of yes–no questions. This can be contrasted with models that assume continuous investment such as [69]. For (ii), the interviews [72] and pricing decisions [73] have identified that insurers lack information about how security controls relate to cyber risk. Indeed, uncertainty is one of the main challenges insurers face [74]. This can be contrasted with Gordon and Loeb's model, which assumes perfect information about the threat and vulnerability. Finally, (iii) allows us to model the insurer's ability to observe attacks in claims data and share this information with policyholders.

Chapter 6 will extend the Gordon and Loeb model to consider cyber warranties because we will be considering private information, not uncertainty. We will assume the security vendor controls the amount of investment in the product and this is hidden from the buyer. The model suggests offering a warranty is a signal about private investment.

We will consider uncertainty about which controls are effective in Chapter 5. The chapter (based on [19]) will develop the IWL to consider how an insurer can aggregate claims data from multiple insureds to guide investment. Fitting this into Böhme and Schwartz's framework [45], this allows us to model a new information structure that has not been previously considered. The next subsection provides an overview of the IWL.

2.3.3 Iterated Weakest Link

There are three aspects to the Iterated Weakest Link model [71]: the rules, the strategy adopted and the computation. The rules define whether an attack will take place and the defender's utility conditional on that attack. The strategy determines the defender's choice of defensive configuration across multiple rounds. Finally, the computation involves calculating the expected utility for adopting each strategy. We consider each in turn.

The Rules

Utility is optimised by balancing the cost of security investments with the likelihood of suffering an attack. The defender has perfect knowledge of the value of the asset a, the rate of return r on the asset and the cost of implementing a given defensive configuration. The loss if any of the n vulnerabilities are exploited is fixed at za, where z determines the proportion of the asset that is lost. The index i runs over the set of all possible defensive configurations. The probability of facing attack p_i for a defensive configuration, along with its cost c_i, determines the expected revenue as

$$R_i = ra - p_i za - c_i \tag{2.3}$$

The probability of facing attack p_i for each defensive configuration is determined by the true costs of attack, which are drawn from $\mathbf{x}_j \in \mathbb{R}^n$. The index j runs over the true costs of attack associated with n vulnerabilities. The true cost of attack x_j for each vulnerability is normally distributed (truncated at zero) around the expected cost of attack $\bar{x}_j$:

$$x_j = \sup(0, \chi_j) \text{ where } \chi_j \sim \mathcal{N}(\bar{x}_j, \frac{\sigma}{\Delta x}) \text{ for } j = 1, ..., n \tag{2.4}$$

The defensive configuration $\mathbf{d}_i \in \{0, 1\}^n$ describes whether the true cost associated with each vulnerability $x_1, ..., x_n$ is protected or not. We denote the k-th defence by d_k. The cost of employing the defensive configuration $\mathbf{d}_i$ is determined by an $n \times n$ matrix $\mathbf{C}$ that reflects "possible interdependent defenses" [71], so that $c_i = \mathbf{d}_i \mathbf{C} \mathbf{d}_i$. The matrix is set such that

$$c_i = \frac{\rho}{2} k^2 + (1 - \frac{\rho}{2}) k \quad \text{where } k = \sum_{i=0}^{n} d_i \tag{2.5}$$

The interaction between the defensive configuration $\mathbf{d}_i$ and the true cost of attack x_j determines whether an attack will place. The k-th vulnerability is defined to be *economically viable* if the true cost of attack x_k falls below the 'loot value' (za) that the attacker gains. The IWL derives its name from the assumption that the attacker will iteratively exploit the so-called "weakest link", which is the unguarded vulnerability with the lowest cost of attack.

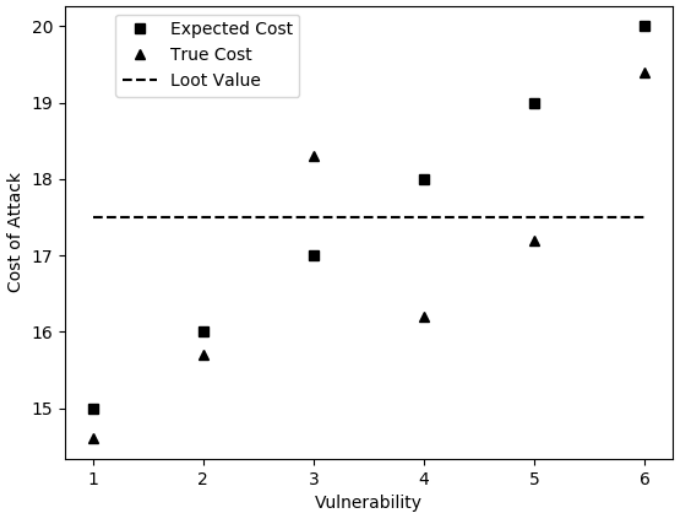

Figure 2.1: The true costs of exploiting the first, second, fourth and fifth vulnerabilities fall below the loot value, za. As a result, there will be an attack unless these defences are in place.

Figure 2.1 describes a situation in which the defender following the expected cost would result in the first three vulnerabilities being defended. Yet, the fourth and fifth would be unguarded and economically exploitable, contrary to the defender's expectations.

In the multiple round case, the cost of attack x_j remains constant but the defender can choose a different defensive configuration in each round. Changing configuration incurs a sunk cost, λa. The defender can choose a defensive configuration based on information regarding attacks suffered in the previous rounds.

The Strategy

Böhme and Moore [71] assume that the defender is rational and will employ the defensive configuration that maximises revenue. The defender is assumed to be risk-neutral. For a one-round game, this reduces to choosing the defensive configuration with the highest expected utility. For a multiple-round game, we must consider how the defensive configuration will be adapted in light of information about attacks.

Böhme and Moore [71] suggest that d_k should not be "tinkered with" once a defence is in place. Otherwise "the direct and indirect ($\rho > 0$) cost of the l-th

defence has to be borne for all intermediate rounds" [71]. If a vulnerability x_l is attacked, then the defender gains the information that it is economically viable. The optimal strategy chooses an initial defensive configuration and then places reactive defences in place only when a vulnerability is exploited.

If the cost of putting additional defences in place becomes too large, it may be rational to accept future attacks and dis-invest to zero defences. Additionally, if sunk costs are particularly high, the optimal strategy may involve accepting future attacks without changing the defensive configuration.

Calculation

If $\sigma = 0$, there is no uncertainty and the defender can directly compute the utility for each defensive configuration. With $\sigma > 0$, the probability of attack when $t_{max} = 1$ is calculated as the probability that at least one unprotected vulnerability x_k falls below z. As the true costs of attack are independent of each other, we can calculate this as

$$P(\text{Attack}) = 1 - \prod_{x_k \in B} P(x_k > z) \tag{2.6}$$

where B is the set of all unguarded vulnerabilities. As both the probability of attack and the cost of defence are determined by the defensive configuration, the optimal starting configuration is the defensive configurations $\mathbf{d}_i$ with the highest expected utility.

In the one-round case, an attack takes place if the true cost of attack x_j of at least one unguarded vulnerability falls below the loot value z, whereas the total number of economically viable unguarded vulnerabilities is important in the multi-round or dynamic case. Böhme and Moore [71] model each unguarded defence x_j as a Bernoulli random variable. Each has probability of being exploited given by

$$P(za > x_j) = \Phi(za; x_1 + (j-1)\Delta x, \frac{\sigma}{\Delta x}) \tag{2.7}$$

The number of economically viable unguarded vulnerabilities t_{att} is modelled as a Poisson binomial distribution with

$$\mu = \sum_{i=k+1}^{n} p_i \text{ and } \sigma = \sum_{i=k+1}^{n} p_i(1-p_i) \tag{2.8}$$

Böhme and Moore approximate this using $\mathcal{N}(\mu, \sigma)$, which is justified "for suitable parameter choices" [71]. Considering t_{max} rounds, there will be $t_{att} \leq t_{max}$ rounds in which the attacker is successful. The total revenue, for $t_{att} = i$ and initial defensive configuration k, is determined by

$$R(t_{att}, k) = \sum_{t=1}^{t_{att}} (ra - z - c_t) + \sum_{t=t_{att}+1}^{t_{max}} (ra - c_{t_{att}}) \tag{2.9}$$

As t_{att} is determined by the starting configuration $\mathbf{d}_1$, we can calculate the expected utility for a given initial defensive configuration $\mathbf{d}_1$ with k defences in place by:

$$U(k) = \sum_{i=0}^{t_{max}} P(t_{att} = i) R(t_{att} = i, k) \tag{2.10}$$

This involves a contribution for each possible value $0 \leq t_{att} \leq t_{max}$. The optimal initial defensive configuration $\mathbf{d}_1$ is the choice of $k = 1, ..., n$ that maximises $U(k)$. Each round contains an additional sunk cost λa each time the defensive configuration is changed, but this does not change the calculation.

To summarise, the IWL is a stochastic model that captures a defender interacting under uncertainty with an attacker over multiple rounds. The best strategy for a defender is to only deviate from an initial defensive configuration in response to observed attacks. These provide information related to the true cost of defence of the exploited vulnerability. The true cost of attack is modelled by random realisations of a normal distribution; to simplify calculating expected revenue, an approximation is used to calculate the likelihood of each realisation of these costs [71].

The next section turns to related empirical studies.

2.4 Empirical Studies

This section describes empirical work relevant to Chapter 4 and Chapter 7. Section 2.4.1 surveys the limited number of empirical investigations of the cyber insurance market. Section 2.4.2 identifies different approaches to quantifying cyber losses.

2.4.1 Empirical Cyber Insurance Research

In 2010, Böhme and Schwartz observed a lack of "quantitative empirical work on cyber insurance markets" [45]. This remains true with a few possible exceptions, depending on what is meant by quantitative. In this section, we introduce empirical studies and comment on whether they fulfil all three criteria of: quantitative, empirical and *on cyber insurance markets*. This frames Chapter 4 by arguing it is among the first empirical quantitative studies of the cyber insurance market.

Of the work covered so far, all of the cyber insurance models introduced in Section 2.3.1 fail to fulfil the empirical criterion. The study of insurability [75] fulfils all three criteria by analysing 994 operational risk incidents and analysing them against an insurability framework. Although, it can be argued it does not directly investigate the cyber insurance market.

One approach involves engaging cyber insurance professionals. Franke [72] conducted interviews to investigate the Swedish cyber insurance market. They found that insurers "impose information and IT security requirements on their customers", fulfilling the governance role described in Section 2.2.3. Further, premiums were found to be less than 1% of the limit ($10k for a $1M limit). This study relied on reports from qualitative interviews.

Cyber insurance policies, which were available from around 2000 [31], provide a different method of study. A number of qualitative studies analysed insurance policies to explore what is covered by cyber insurance. For example, Baer et al. [29] examined six policies and Marotta et al. [30] analysed 14 policies to ask which incidents cyber insurance covers. Kesan et al. [65] provided a comparative analysis of three policies. Majuca et al. [31] studied seven policies offered by the same insurer to understand how they change over time. These studies fail to meet the quantitative criterion.

Romanosky et al. [73] analysed documents that US-admitted insurers are required to submit (this dissertation will make use of these documents too). These include insurance policies, application forms and pricing schemes. This dataset was introduced in 2017 and allowed Romanosky et al. to analyse an order of

magnitude (over 100) more policies than previous studies. Their policy analysis revealed coverage is more consistent across policies than is commonly thought.

Analysing application forms allowed Romanosky et al. [73] to consider what information is collected for the risk assessment. They focused on the US market and used an inductive approach to identify 97 unique lines of questioning, which were classified into four broad sections: Organizational, Technical, Policies and Procedures, and Legal and Compliance. However, the presentation of results does not provide quantitative insights into which sections received the *most* focus.

An ENISA report [76] provides a cautionary tale in analysing proposal forms. The authors tried to map questions in the application forms to an information security framework, the Center for Internet Security Critical Security Controls [77]. There were a number of questionable mappings. For example, the question *'Do you have a group-wide privacy policy'* was considered to be related to 'Controlled Use of Administrative Privileges'. We had already conducted this exact analysis, which we will describe in Chapter 4.

Analysis of the pricing schemes provided the first insights into cyber insurance pricing. The results indicate that only 33% of cyber insurance premiums are modified based on the security posture of the insured, which is strange given how much information is collected about information security, as identified in [73, 17]. However, the results only provided qualitative insights into how prices are formulated even though this topic is seemingly ripe for quantitative research.

Chapter 4 contributes the first quantitative analysis of cyber insurance prices to this stream of literature. We also analyse the relative focus of the application process on different areas of information security (based on [17]), which cannot be answered by the analysis in [73]. We are the first to include both analyse documents and engage insurance industry participants in the same study, which provides insights into how the results from these two approaches relate.

Study	Year	Data Source	n	Frequency	Impact
[78]	2008	Identity Theft Resource Center	899		
[79]	2010	Open Security Foundation DatalossDB	956		Power law*
[80]	2016	Privacy Rights Clearinghouse (PRC)	2253		Lognormal*
[81]	2016	PRC & OPS DatalossDB &	8574		Pareto*
[82]	2017	Privacy Rights Clearinghouse	2266		Log-skew-normal*
[83]	2018	Privacy Rights Clearinghouse	600		Stochastic process*
[84]	2010	Research institute's event log	23,000		Lognormal*
[85]	2014	Nordic bank's event log	1,800		Lognormal*
[86]	2018	Proprietary data	53,308		
[87]	2018	Interviews	2200		μ
[88]	2015	Information security breaches survey	664	μ	μ
[89]	2018	Bank of Italy survey data	4,209		μ
[90]	2014	US court dockets	230		
[91]	2017	UK ICO regulatory actions	118		μ
[75]	2015	Operational risk database	994		μ, σ
[92]	2003	Event (DoS) window study	23		AR
[93]	2003	Event (Data Breach) window study	43		AR
[94]	2004	Event (Data Breach) window study	66		AR
[4]	2016	Proprietary data	12,585	μ	MR

Table 2.2: Research into the cost and frequency of cyber losses. * = Not financial impact, μ = mean/mode/median, σ = variation, MR = Multivariate regression, AR = Abnormal returns

2.4.2 Quantifying Cyber Losses

Table 2.2 provides an overview of the studies we have considered, including the data source of the study. The frequency and impact columns concern whether the study provides insights for individual organisation, not statistics aggregated across multiple firms within an industry or economy.

Studying data breach repositories like Privacy Rights Clearinghouse provides insights into the size of breaches and aggregate frequency. Aggregate frequency is found to be stable over time [80, 81] and is distributed according to a "Poisson or negative binomial" [82]. Breach size was shown to be best described by a power law [79], Lognormal [80], Pareto [81], and log-skew-normal [82] distribution in successive publications. A recent paper has suggested both aggregate frequency and breach size are better modelled by stochastic processes [83].

Governments are the only entities who can realistically operationalise insights into aggregate frequency. Unfortunately they have a limited number of levers to pull in response. A notable exception is mandatory breach reporting laws, which

have been shown to reduce identity theft by 6.1% [95]. Organisations have more risk management tools available, but how can they estimate potential losses?

Schroeder et al. [84] investigated 23,000 failures recorded on more than 20 different systems at a research institute. The results suggest time between failures is modelled by a Weibull distribution and "repair times are well modeled by a lognormal distribution". Franke et al. [85] investigated 1800 incidents in a large Nordic bank and also found that the "lognormal distribution offers the best fit of IT service time to recovery". These studies do not provide incident costs. It is not clear how much we can generalise from atypical organisations like the Los Alamos National Laboratory.

The Data Breach Investigations Report [86] describes the relative frequency of different types of incidents by industry. The Ponemon Institute Cost of a Data Breach [87] report estimates the average cost per record in a data breach. However, [86] does not provide absolute frequencies and [87] only surveys firms who have detected a breach. There are questions about the compatibility of commercial sponsors and scientific integrity, as evidenced by the $1 trillion cyber crime figure[9].

Surveys commissioned by governments provide an alternative. One study of security investments [88] used a survey commissioned by the UK Government [96] providing point-estimates of both frequency and impact distributed according to a Bernoulli distribution. Piggy-backing security-related questions on the Bank of Italy's annual survey provided a larger set of responses [89]. Neither of these sets of losses were fitted to a distribution. Furthermore, self-reported surveys are complicated by response biases [97].

Court dockets provide rich information about legal cases. Romanosky et al. [90] investigated factors affecting the data breach litigation. Freedom of information requests have been used to understand regulatory actions in the UK [91]. But they only provide insights into cases contested in the courts, leading to small sample sizes.

Biener et al. [75] extracted 994 cyber incidents from an operational risk database. The mean and mode loss are $41M and $1.9M respectively, and "the distribution of

[9] https://www.forbes.com/sites/andygreenberg/2012/08/03/mcafee-explains-the-dubious-math-behind-its-unscientific-1-trillion-data-loss-claim/#294dfc4525a3

the non-cyber risk sample is much heavier tailed than that of the cyber risk". The mean loss figure of $41M is 10 times larger than a similar figure in [4]. The types of attack included in their dataset depends on the "comprehensive set of keywords" comprising their definition of a cyber incident, so it is not clear what is included.

Event window studies provide insights into how events, such as denial of service attacks [92] and data breaches [93, 94], impact the stock market. By extracting information from the stock prices, event window studies represent an intellectual forefather of the technique we introduce in Chapter 7, which extracts information from insurance prices. However, event windows studies are limited to publicly reported events. As a result, they tend to have small sample sizes and questionable relevance beyond listed companies.

The most comprehensive study quantifying cyber losses is based on a proprietary dataset with 15,000 incidents [4]. Attack frequencies for each industry are calculated using census data on the number of firms in each industry. Their multivariate regression, based on 265 observations, describes how total cost of incident varies according to an organisation's revenue, industry, number of records and past incidents.

Table 2.2 summarises a research programme with strong evidence regarding how the number of records in a data breach is distributed and how frequently they occur across the United States. We have estimates [4] of breach frequency for an organisation, although this figure groups all firms in an industry regardless of their size[10], as well as point estimates of financial cost of incidents via survey data [88, 89] and an insightful multivariate regression [4].

However, quantifying the costs of attacks beyond data breach and regulatory fines is a challenge. In particular, the studies described in Table 2.2 provide no insights into the potential cost of business email compromise, business interruption incidents, or ransomware attacks. A possible exception is [75] but it is unclear what is included under operational loss.

Chapter 7 will introduce a method to derive distributions of dollar losses advancing the state-of-the-art in quantifying cyber losses. Using insurance prices

[10]Grouping corporations like Target with independent book stores under "retail trade" might explain the particularly low frequency for retail firms.

for specific firms suggests these distributions provide a granularity not available in alternative loss estimates. Further, the differences in coverage types allow us to delineate types of losses, which can provide estimates for previously unquantified loss types like ransomware.

2.5 Summary

Section 2.2.1 identified six functions of insurance of which efficient allocation of risk and fostering governance are most relevant to this dissertation. Section 2.2.3 introduced the idea that cyber insurance will improve information security levels through a number of processes; this dissertation will focus on how insurers assess applicants' security levels (Chapter 4), aggregate claims data to identify effective security controls (Chapter 5), and their expertise in quantifying losses (Chapter 7).

Section 2.4 explored empirical studies of the cyber insurance market. We found that the majority were qualitative and there was no quantitative study of pricing. Chapter 4 contains our contributions to the empirical cyber insurance literature. It provides observations of how premiums vary by coverage type, amount, policyholder type, and over time. Then we use mixed methods to build a richer picture of the application process in the UK.

In Section 2.3, we identified how different aspects of the cyber insurance market had been modelled, including the information structure and the wider ecosystem. Chapter 5 will introduce the first model to consider how insurers might aggregate claims information to identify effective security controls. Chapter 6 investigates how security product vendors offering warranties, an insurance like product, will impact investment decisions.

Finally, Section 2.4.2 surveyed approaches to quantifying cyber losses. There are few estimates of the distribution of financial losses available. Chapter 7 introduces a technique to estimate distributions of losses using the insurance prices collected in Chapter 4. Our insights provide granular insights for specific firm sizes and revenues, as well as for incident types not yet considered in the literature.

The next chapter discusses this dissertation's methodology and data collection.

3
Methodology

Collecting evidence about the effects of cyber insurance on how organisations make security decisions is difficult. Observing decisions in the real world is problematic because of selection effects. Organisations purchasing cyber insurance differ from those who do not. The decision to seek risk transfer solutions requires a mature risk management approach. Consequently, we cannot attribute the difference in observed effects to cyber insurance because it could result from pre-existing maturity or any other attribute that purchasing cyber insurance selects for.

Traditionally scientists have solved this issue by designing controlled experiments. The independent variable, in this case whether there is a relationship to an insurance provider, is randomly varied and the effects are measured. Randomly assigning the independent variable allows its causal role to be identified, ignoring the possibility of confounding variables.

External validity — the extent to which the results generalise beyond the study — is determined by experimental design. Experiments cannot hope to perfectly simulate the world. The causal mechanism may not generalise beyond the conditions of the experiment. This is particularly true when the world is represented as an abstract economic model[1], complete with simplifying assumptions.

[1]Maki [98] makes the argument that models are experiments. In this case, external validity is determined by how well the model's mathematical relations represent the world.

This dissertation adopts both theoretical modelling and empirical observations in an attempt to reconcile these problems. Empirical work describes the cyber insurance market. Economic models identify causal mechanisms in our mathematical representation of the market. The dissertation will tie these two strands together by reflecting on the discrepancies between modelling assumptions and empirical observations. While this cannot identify causal mechanisms in the real world with certainty, it can at least reduce uncertainty about them.

The rest of this chapter is split into an empirical section and a modelling section. Section 3.1 describes how we collected our empirical data. Section 3.2 justifies how we analyse this data. Section 3.3 explains our specific choice of economic model. Section 3.4 introduces an alternative way to estimate the results of the Iterated Weakest Link model [71], which will allow us to extend the model in Chapter 5.

3.1 Collecting Empirical Data

Speaking to insurance professionals [72] and analysing the documents they use [73, 46] are the primary approaches to collecting evidence about the cyber insurance market. Doing both allows us to comment on the relation between the two. For example, Chapter 4 will investigate how well security controls in application forms align with the controls valued by underwriters. Section 3.1.1 describes how we collected data about underwriting in the UK. Section 3.1.2 does the same for pricing in the US.

3.1.1 Underwriting Data

We employ a mixed methodology consisting of interviews, a content analysis and a questionnaire. Our mixed methods aim to achieve completeness [99] by engaging insurance professionals and analysing the forms they use; doing both allowed us to explore how these findings relate to each other. Such comparisons were strengthened by using the same sample for each research method.

We constructed a sample by making contact with an insurance broking firm who acted as a gatekeeper to the insurance industry. Using their experience and market

[A] European, target smaller risks	[B] US, target smaller risks
[C] European, target sector-specific	[D] European, target smaller risks
[E] European, target complex risks	[F] Australian, target complex risks
[G] Bermudan, target complex risks	[H] European, target complex risks
[I] European start-up, target sector-specific	

Figure 3.1: Brief description of the profile of each insurance company.

access, we established contact with 11 different insurance firms chosen because they represent a range of different market strategies: from offering 'off-the-shelf' insurance products to small businesses through to offering bespoke policies to large multinational companies. A brief profile of each is included in Figure 3.1.

All of our participants are members of the Lloyd's of London insurance market. The global nature of the Lloyd's market means the results may have international relevance. We acknowledge that this could result in sample bias, however there is no alternative specialist market in the UK.

All 11 firms provided the proposal form their organisation uses, 9 insurers completed a questionnaire, and 9 participated in an hour-long interview. Non-response was an issue given the time commitments insurance professionals face. Two firms did not complete an interview or questionnaire despite email reminders.

We conducted an hour-long semi-structured interview with an underwriter from each of the 9 insurance firms in our sample. The questions began descriptively asking about the role of different stages of the application process, and how the different stages were used to form a final decision. We then evaluated the application process, asking which areas are assessed well and which the insurer cannot assess with current techniques. Finally, we asked questions about how the assessment process improves over time.

The questionnaire provides an opportunity to directly compare participant responses to the content analysis (described in Section 3.2). This was achieved by asking each participant to rate the importance on a Likert scale [100] of each theme in the content analysis. These will be introduced in Section 3.2. In addition we asked the four additional questions outlined in Table 3.1. The multiple choice

Question	Type
Which of the following sources of information do you use to improve your risk assessment:	Multiple Choice
Which of the following do you regularly use to assess an applicant:	Multiple Choice
What are the most important aspects of an applicant's information security that you cannot effectively assess?	Open-ended
What do you consider the most important data points when assessing an applicant?	Open-ended

Table 3.1: Questionnaire questions that are not based on themes.

questions provide an 'Other' box to input free-form responses. The purpose of these questions is to capture responses we failed to anticipate.

The questionnaire was distributed electronically directly to the insurers. There were two non-responses from the individual assigned by that company, which is natural given the time constraints of the targeted professionals. We outline how we analysed this data in Section 3.2.

3.1.2 Pricing Data

Insurers in the United States operating in admitted markets are required to "file their policies and rate schedules with the state insurance departments" [73]. Rate schedules describe the formulas and tables used to calculate the premium for a given applicant. These documents are made publicly available by the National Association of Insurance Commissioners (NAIC).

Documents are organised by state and made available in the SERFF Filing Access system. We chose to focus on the state of California, which is the largest in the US, rather than partially search multiple states in order to collect the widest range of policies. Even if the set of policies skews towards firms unique to California, the insights from a given policy should have general applicability given there are no differences between states "that would materially bias any results or conclusions" [73].

Our objective was to quantify how cyber insurance premiums are adjusted according to coverage type, limit, deductible, industry, revenue, and security

infrastructure. This will allow us to generate sets of insurance prices for a hypothetical firm with specific characteristics, such as revenue or industry.

We extracted rate schedules related to cyber insurance by searching with keywords "cyber", "security" and "privacy", as in [73]. We only downloaded documents with either a "new program" or "rate" component in the filing type because these relate to pricing, whereas some filings concern the policy wording and have a "form" or "rule" filing type.

This resulted in 131 unique filings that were narrowed down to 26 rate schedules. Romanosky et al. [73] identified three types of pricing: flat rate pricing, base rate modification, and information security pricing. There were 40 filings related to the first type and they were excluded because the inflexible pricing structure makes inferences difficult. Further, being sold alongside existing products means these endorsements may not be commercially viable alone.

A further five filings relate to the price of excess layers in which insurers offer additional *layers* of coverage to supplement an existing policy. These tend to be priced as a percentage of the original policy, which makes inferences difficult. We also excluded filings for policies with pricing structures for specific industries. For example, one policy was priced according to the number of doctors employed by a healthcare provider.

Some filings introduced new coverage areas without updating the original rates. We considered these to be additions to the original filing. A small number of insurers updated prices for the same coverage and these were considered to be new filings. This resulted in 26 unique filings, which correspond to sets of prices.

A rate schedule does not provide prices directly. We had to read the document explaining how prices are calculated and extract the corresponding tables of multiplicative factors. The prices from the 26 rate schedules are determined by the product of a base rate (in USD) and many multiplicative factors, which increase or decrease the price.

After extracting the tables, we can calculate the price-limit-deductible triples for a hypothetical firm. The hypothetical firm would have characteristics, such

Revenue	Base rate	Deductible	Factor	Limit	Factor	Hazard Class	Factor
\$10m or less	\$1,914	\$10,000	1	\$500,000	0.809	1	.804
\$10m–\$20m	\$2,603	\$25,000	.95	\$1,000,000	1	2	1
\$20m–\$50m	\$3,502	\$50,000	.89	\$2,000,000	1.132	3	1.497
\$50m–\$100m	\$5,225	\$100,000	0.82	\$3,000,000	1.245	4	1.905

Table 3.2: A subset of the baserates and multiplicative factors found in a rate schedule, which can be downloaded from `https://tylermoore.utulsa.edu/cyberschedex.pdf` Price = (Base rate) × (Deductible Factor) × (Limit Factor) × (Hazard Class Factor)

as revenue or industry, corresponding to each multiplicative factor. Using the factors in Table 3.2, one of the triples for a retail firm with revenue \$50M would be $(4964, 1000000, 25000)$ because $4964 = 5225 \times 1 \times 0.95 \times 1$ where 1, 0.95 and 1 are the factors for the limit of $1M$, deductible of 25K, and hazard class of 2 respectively.

Generating the data set consists of taking all combinations of limits and deductibles offered by an insurer then computing the corresponding premium. If a rate schedule provided a choice of 8 deductible amounts and 15 limits, then there would be $8 \times 15 = 120$ triples. This leads to a total data set of $6,828$ price-limit-deductible triples across all 26 rate schedules, with price varying based on the hypothetical firm's characteristics. Chapter 7 will introduce a method to make inferences from these triples.

3.2 Analysing Empirical Data

This section describes how we analyse the risk assessment data predominantly. We provide simple observations (plots extracted directly from the data) of the pricing data in Chapter 4 and conduct no statistical tests. More complex analysis is delayed until Chapter 7 in which we introduce a new method to infer loss distributions from this data.

A deductive approach is used in the content analysis and the questionnaire, which allows us to relate our results to existing theory. The themes are derived from the information security standard ISO/IEC 27002:2013 [101]. Deductive content analyses are reliant on choosing themes that are mutually exclusive and exhaustive [99].

Woods et al. [17] identified that ISO/IEC 27002 covered the security controls found in cyber insurance proposal forms demonstrating it to be exhaustive in this respect. Further, Humphreys [102] describes it as the "de facto 'common-language' for information security". Although some contend that appeals to authority and popularity are not a sufficient basis for international standards [103], a widely agreed upon structure is important for our purposes.

Table 3.3 outlines the sections of ISO/IEC 27002 and the 35 sub-sections from which we derive the themes. For example, *Section 7: Human resource security* is split into three sub-sections with controls relating to *7.1 Prior to employment*, *7.2 During employment* and *7.3 Termination and change of employment* respectively. We chose not to extend the themes as the scope of ISO/IEC 27002 is so broad but we did collect together the units of analysis not mapped to a theme. We now turn to each of our methods in turn.

Content analysis "can be used to identify critical processes" [104], such as the factors involved in a risk assessment. Keyword analysis would not have captured the latent content of each question, yet a qualitative analysis may have been overwhelmed by the volume of questions.

A content analysis breaks the data into *units of analysis*, which are mapped to a theme. We define a *unit of analysis* to be a piece of text prompting a singular answer from the applicant. At times we departed from this to ensure the over-arching goal that units of analysis contain a similar amount of information, broadly defined.

We employed a similar rationale in mapping units of analysis to themes as was used in [17]. Of the 1043 questions across all 11 forms, only 243 were open-ended. The other questions were all closed and consequently it is easier to map each question to exactly one theme. The results of the content analysis, questionnaire and interviews will be presented in Chapter 4.

3.3 Modelling Choices

Economic models have been the primary method for investigating the cyber insurance market in the field of the economics of information security, as evidenced by

ISO/IEC 27002 Control	Contents of control
Section 5: Information security principles	Define a set of policies which direct and support information security.
Section 6: Organization of information security	Organisational roles and responsibilities for information security, controls for mobile devices and remote access.
Section 7: Human resource security	Policies to manage pre-employment screening, training and awareness during and post employment.
Section 8: Asset management	Maintaining an asset inventory, classifying information according to protection needed and managing information storage media.
Section 9: Access control	Defining an access control policy, controlling access rights throughout employment and ensuring users are aware of their responsibilities.
Section 10: Cryptography	Define a policy for the technical application of encryption.
Section 11: Physical and environmental	Enforce physical entry controls and protect equipment on and off site.
Section 12: Operations management	Control changes to IT facilities, install malware controls, take regular back-ups, maintain and monitor logs, control software installation, and manage vulnerabilities.
Section 13: Communications security	Maintain secure network architecture and control information transfer with third parties.
Section 14: System acquisition, development and maintenance	Specify security requirements and employ secure software engineering principles during development.
Section 15: Supplier relationships	Define policies and procedures for external supplier contracts and agreements. Monitor and audit the service provided.
Section 16: Info security incident management	Report, assess, respond to and learn from security events. Collect forensic evidence.
Section 17: Business continuity management	Define, update and review the organisation's business continuity plan.
Section 18: Compliance	Document and uphold security obligations to external authorities and third parties. Address external audis.

Table 3.3: Overview of the contents of each section of ISO/IEC 27002.

Chapter 2. Yet it is widely held that "all models are wrong" [105]. We make a case for the extended aphorism "all models are wrong but some are useful". Although Box's statement refers to statistical models, we will argue that it applies to economic modelling. We do this by sketching out a number of positions in the philosophy of economics, upon which we can draw to make our own argument.

In an influential 1953 essay [106], Milton Friedman argued that the realism of an economic theory is irrelevant; rather, the theory's predictive success is all that matters. He states that economic theory is "to be judged by its predictive power for the class of phenomena which it is intended to explain". He acknowledges the difficulties in collecting evidence to test economic predictions.

Moreover, assumptions need not conform with evidence. He suggests "in general,

the more significant the theory, the more unrealistic the assumptions" [106]. Instead, assumptions need only be "sufficiently good approximations for the purpose in hand" meaning that the model "yields sufficiently accurate predictions". Friedman's view might be summarised as: the wrongness of a model does not matter, providing it is useful in predicting reality.

Friedman's position both supports the importance of Popperian falsification [107] while also excusing economists from actually collecting data to falsify theories. Hausman [108] suggests economic methodologists are more attracted to Lakatos' [109] focus on "theoretical progress" rather than outright empirical success.

Frigg moves away from empirical testing by conceiving of scientific models as fictional worlds [110]. Just as fiction must be internally consistent (intrafictional truth), mathematical models must be logically consistent. Justifying how a model behaves like a real-world system (analogous to transfictional truth) involves arguing that "certain relevant properties are similar in relevant respects" [110]. The importance of similarity contradicts Friedman's antipathy towards realistic assumptions [106]. In this view, success looks like "a sequence of ever more realistic mathematical models' [111]'.

So why work towards realistic economic models? Hausman suggests models are an exploratory tool to conduct "investigations of possibilities" [112]. More realistic models are more useful as exploratory tools. Maki [98] goes further by arguing theoretical "models=experiments" in that they are both "representations that are manipulated" to isolate major processes. Maki identifies two differences: isolations are theoretical in models but causal in experiments, and theoretical models provide much "tighter isolations". This stance suggests theoretical models should be built to investigate processes that we cannot easily observe in the world.

Finally, McCloskey [113] argues against "methodological authoritarians" who only permit discussion within the "official rhetoric" of formal economic models. This view argues reasoning with economic models alone is insufficient. McCloskey instead suggests the whole "range of persuasive discourse in economics" should be used, like arguing by metaphor for example. This view reminds us to draw on

the sociology of insurance identified in Chapter 2 and the empirical results from Chapter 4 when interpreting the implications of model results.

To conclude, there is little hope of collecting direct evidence about the kind of questions we are interested in answering. This moves us away from Friedman's predictive success criteria when evaluating models. Instead, our modelling choices aim at realism. Realism is not pursued for its own sake, but to improve the model as an exploratory tool. Staying true to McCloskey [113], results will be interpreted by reflecting on the gap between models and empirical data.

3.4 Simulating the Iterated Weakest Link

This section simulates the original IWL using Monte Carlo methods and validates the results against those found using the approximation described in Section 2.3.3. The approximation used for calculation purposes in the original IWL [71] would break down if the ordering of vulnerabilities was important (like in our extended model in Chapter 5). The reader will have to take this on faith until we introduce the model in Chapter 5 and can explain why the ordering is important.

Knowing the number of economically viable vulnerabilities was sufficient to calculate returns on the original IWL [71]. Whether vulnerability i is exploited in the (t)-th or $(t+1)$-th round does not affect the revenue. The order in which vulnerabilities are exploited is unimportant. Consequently, expected revenue is calculated by approximating a series of Bernoulli trials that represent the number of economically viable vulnerabilities.

Simulating the IWL using Monte Carlo methods takes into account the ordering of vulnerabilities. To validate this approach, we simulate one defender following the rules and strategy employed in the original IWL (we consider multiple defenders in Chapter 5). If we were simulating the natural world, we would have to choose an appropriate distribution. However, we are simulating the IWL in which the true costs are assumed to be normally distributed.

Figure 3.2 compares the results of the IWL for two different methods of calculation: the "Original" approach using the approximation used in [71] and

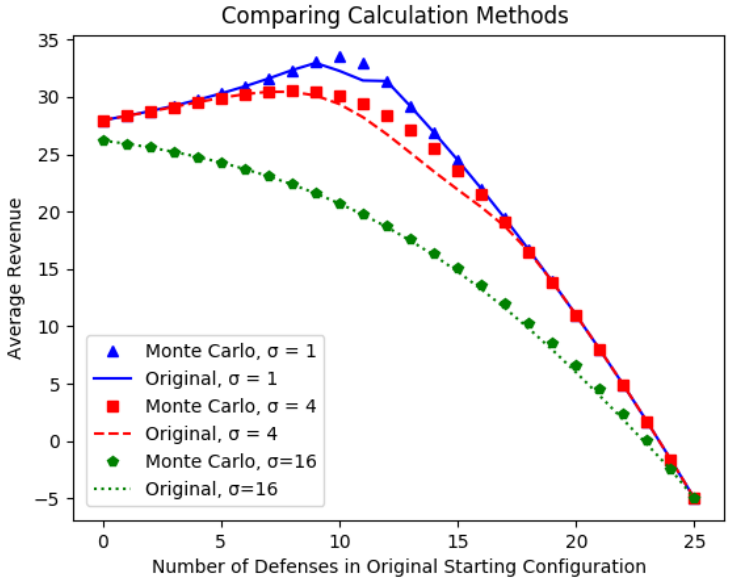

Figure 3.2: Average revenue of 10^7 Monte Carlo simulations compared to the expected revenue using the method of [71].

our "Monte Carlo" approach using 10^7 simulations of the IWL. The two methods broadly agree; as uncertainty increases, both expected revenue and the optimal number of starting defences fall. This suggests simulating the IWL is a reliable way of estimating the expected return on a given strategy.

For low values of σ, the methods diverge around the optimal point. Indeed, for $\sigma = 1$ the simulations suggest that the optimal starting configuration is different to the results of [71]. The source of divergence is most likely to be Böhme and Moore's decision to approximate the number of economically viable unguarded vulnerabilities as a Poisson binomial distribution [71]. Appendix B presents the standard error of the mean (SEM) for all of the calculations in Figure 3.2. Even when uncertainty is at its highest, the SEM results do not exceed 0.01.

However, we should not lose the forest for the trees. The IWL is not designed to be a predictive model. Rather, it is designed to gain insights into the strategies different parties might adopt under certain assumptions. Both the original computation and our simulations suggest that the optimal number of starting defences decreases as uncertainty increases; the original insight of [71] that sometimes it is rational to 'wait and see' with security investments holds true.

3.5 Summary

This chapter justifies the use of mixed methods in this dissertation. Empirical data is collected to understand the cyber insurance market. Speaking to insurers and analysing the documents they use provides a richer picture of the risk assessment process. ISO/IEC27002 [101] provides themes for our analysis, allowing us to see how the risk assessment aligns with what is generally held to be best practice.

The empirical results inform our modelling choices with the aim of providing realism. Economics models are used to explore processes and mechanisms that we cannot probe with observational data. We do not subscribe to the "official rhetoric" of economics [113] by taking modelling results as definitive answers. Insights from modelling will be interpreted in the context of empirical evidence and sociological critiques of insurance.

4

Empirical Observations of Risk Assessment and Pricing

Section 4.1 presents our results on how insurers assess risk. Section 4.2 describes the pricing schemes we obtained. This work is based on [18]. Section 4.3 discusses both sets of results individually and then how they relate to each other. We offer conclusions and link the results to the forthcoming chapters in Section 4.4

4.1 Risk Assessment

This section will follow the structure of the interviews by first looking at how and what information is collected in the application process, then evaluating the quality of the assessment, before finally looking at how the assessment process evolves over time. To avoid clumsy references to each interview, we will reference Table 4.1, which provides a brief description of the profile of each participant's firm. For clarity, we will refer to: the *underwriter* who receives the premium and pays out in the event of a loss; the *broker* who acts as an intermediary between policyholder and underwriter; the *actuary* who is concerned with the financial security of the insurance company; and the *reinsurer* who underwrites insurance for other insurance companies.

The application process begins with initial contact being made, which can involve submitting a proposal form directly to the underwriter or making an informal inquiry

[A] European, target smaller risks
[B] US, target smaller risks
[C] European, target sector-specific
[D] European, target smaller risks
[E] European, target complex risks
[F] Australian, target complex risks
[G] Bermudan, target complex risks
[H] European, target complex risks
[I] European start-up, target sector-specific

Figure 4.1: Brief description of the profile of each insurance company.

Method of Assessment	Number of insurers using it
Proposal forms	8
Telephone interview with IT team	7
External intelligence	4
Automatic risk assessment tools	3
Telephone interview with board	2
Other	2
Onsite Audit	1

Figure 4.2: The results of nine questionnaire responses to the question *Which of the following do you regularly use to assess an applicant?*

via the broker. Information about the applicant is collected directly using some combination of an applicant presentation or telephone interview with staff from the applicant (usually the IT team and/or the board) or indirectly collecting information via a method like open source intelligence or external scans of the applicant's network. The underwriter usually assesses how offering coverage will impact the underwriter's portfolio of risk, which may involve communication with an actuary within the company or an external reinsurer.

4.1.1 Methods of Assessment

Participants identified many different stages of the application process, which are outlined in Figure 4.2. Each method of assessment serves a different purpose and the balance between proposal forms, telephone interviews and other factors differed for each application. We now explore the role of each method.

Applicants are expected to undergo different methods of assessment depending on their size and industry. The underwriter at I reported that "our expectation for a large bank or healthcare facility would be very different from an SME tech

company". Smaller organisations are more likely to fill out a proposal form, which may contain fewer questions.

Proposal forms provide a standard assessment that can identify specific questions to be asked later in a telephone interview. The underwriter at D praised the "efficiency" of proposal forms in terms of time spent gathering information, but did identify their lack of flexibility. Proposal forms provide an analytical function in enabling the underwriter to "comparatively assess one insured against another with the same set of questions and information" (G).

Security calls provide an opportunity to "talk to people who are in charge" (C) and ask specific questions and address inconsistencies within the proposal form. Many participants (B, D, F, H, and I) stated security calls provide insights into the "security culture" of the applicant. The utility of a call depends on the role and competency of the staff member on the call. A suggested that "for it to be any good, you need the IT security person" on the call, while D stated they are "often [a] waste of time because they won't have the answers or they'll talk about something completely irrelevant".

On site audits are seen as both expensive and intrusive, particularly by I in reference to the manufacturing sector. E stated that automated tools were not considered useful in "pricing risk". External intelligence was used to estimate the size of a breach or even to understand the reputation of key security staff, such as the CISO (H).

To receive large coverage limits, organisations tend to seek co-insurance from multiple underwriters. In this situation, the underwriters rely on a security call in which the client presents to all of the underwriters and answers questions from each. However, the applicant will not complete a proposal form for each underwriter; instead they send the completed proposal form of the lead underwriter, or possibly none at all. This results in underwriters collecting different proposal forms, which undermines analytical opportunities.

Wider market conditions were also important in terms of the depth of assessment. Underwriter G noted that an over-supply of insurance meant that "people don't

want to ask the question because they're desperate to get the revenue in". There was consensus that this might change if big losses "harden the market" (B).

Many responses highlighted how the broker supports the applicant in accurately filling out a proposal form. Further, brokers play a role in determining how the information is collected. Underwriter D reported that

> "more and more we're seeing applications done by 'here's a financial statement and we're going to have an hour phone call and you can ask them what you want' versus 'here's a proposal form that says everything you need to know'."

Many participants worried that prevailing market conditions impact the "amount of risk assessment or risk management you can do with a company" because brokers will take "the path of least resistance" (C).

Underwriter F noted that other underwriters might be "more flexible with risk" because they have "the capital to deal with it". This illustrates the concept of insurance cycles, whereby an over supply of capital drives down premiums and risk assessment standards [114].

4.1.2 Collected Information

The content analysis identified what information is collected in proposal forms. Figure 4.3 details the total number of questions classified under each section of ISO/IEC 27002 and the average number of questions per sub-section. Together they provide an idea of the volume of information that is being collected per section (absolute score) and the volume of information collected relative to the size of that section (average score).

The absolute score suggests that the bulk of information collected relates to organisation structure, operational controls, information transfer, business continuity and compliance considerations. Controls related to cryptography, physical security and systems development are relatively less well-covered by this collection of forms.

Table 4.4 displays the number of questions corresponding to each sub-section of *Section 12: Operations security* and a typical example question for each. The majority of these questions are centred around security controls, back-up processes,

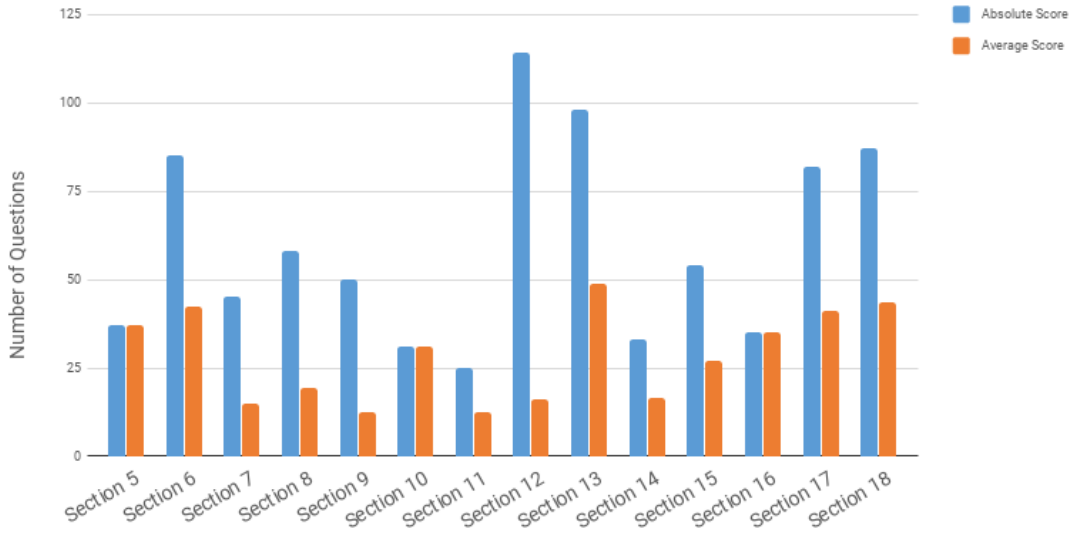

Figure 4.3: Absolute and average number of questions mapped to each section of ISO/IEC 27002 across all of the analysed proposal forms. See Table 3.3 for an outline of each section.

network monitoring and patch management. For example, for *12.2 Protection from malware* there are seven questions related to running anti-virus, five related to updating anti-virus and six related to phishing awareness training. Regarding *12.3 Backup*, only six forms asked whether a backup system was in place, with 10 questions asking about the specifics of implementation.

Questions corresponding to *12.4 Logging and monitoring* covered a variety of controls including monitoring access to sensitive data, collecting Point-of-Sale (POS) log data, installing intrusion detection systems (IDS) on devices, and monitoring user activity. Information regarding patch management policies and vulnerability scanning corresponded to *12.6 Technical vulnerability management*; six forms asked to describe patch management policy and provide details regarding frequency, while only two framed this as a yes–no question. The majority of questions regarding vulnerability scanning relate to specifics such as whether they scan POS systems, whether the process is regularly reviewed, and whether all web pages are scanned.

The average score provides an insight into which areas of ISO/IEC 27002 are not being fully explored by the proposal forms. For example, *Section 12* has a low average score in Figure 4.3 because *12.5 Control of operational software* and

Sub-Control	Sample question	# of questions
12.1 Operational procedures	Do you have a formal change control policy?	13
12.2 Protection from malware	Does the Applicant utilise Anti-Virus software?	22
12.3 Backup	Is all critical data backed-up at least weekly?	16
12.4 Logging and monitoring	Do you utilise Intrusion detection or prevention systems?	34
12.5 Control of operational software	Do you restrict staff's ability to install software?	2
12.6 Technical vulnerability management	What is your security patch management policy?	22
12.7 Info systems audit considerations	Do you periodically test or audit security controls?	5

Figure 4.4: The number of questions asked in relation to each sub-section of *Section 12: Operations security*.

12.7 Info systems audit considerations are barely addressed. Similarly, *Section 9: Access Control* has one of the lowest average scores.

The majority of the questions that were not mapped to a theme relate to an organisation's profile such as revenue, employee numbers, IT budget, previous insurance cover, mergers and acquisitions, and contact information. The answers to many of these questions are publicly available.

4.1.3 Making a Decision

Although the proposal forms provide insights into what information is collected, they do not show how this information is used to make a decision. The questionnaire results shows the relative importance of different areas of information security in the application process. Figure 4.5 plots the average score in the proposal forms alongside the average score in the questionnaire, for each section of ISO/IEC 27002. We converted the Likert scale responses to an index (because we are interested in the average of multiple responses) and then normalised for ease of comparison with the content analysis.

Using ISO/IEC 27002 for both the content analysis and the questionnaire allows us to cross-validate the results. Figure 4.5 shows disparities between the amount of information collected and the underwriters' priorities. For example, low-level technological controls are deemed to be important yet relatively little information is collected about them. The converse is true of high-level strategic

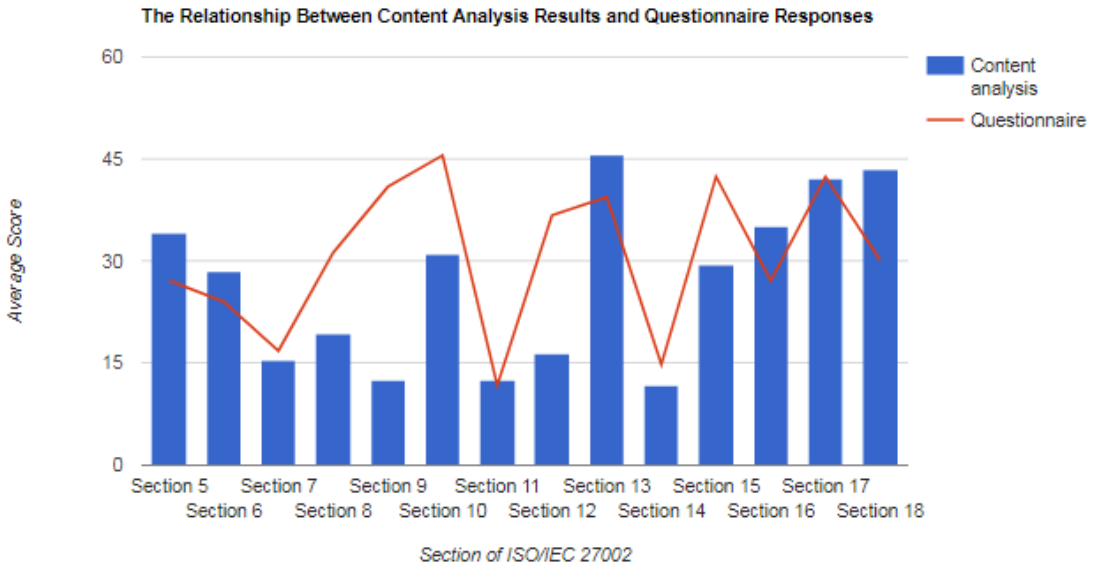

Figure 4.5: Showing the relationship between the number of questions per sub-control and the sub-control's importance as reported in the questionnaire averaged across each section and across all respondents.

policies. *Section 8: Asset management*, *Section 9: Access control* and *Section 10: Cryptography* are examples of the former, while examples of the strategic sections include *Section 5: Information security policies*, which is concerned with the overarching information security policy, and *Section 6: Organization of information security*, which predominantly relates to roles and responsibilities.

The way that participants classified security controls does not perfectly correspond with the ISO/IEC 27002 sections, which supports the discrepancy of the previous paragraph. For example, data security was often used to describe a combination of access control, cryptography and communications security, which suggests the conceptual level of ISO/IEC 27002 is more granular than that of the insurers when it comes to technical controls.

Many of these technical controls "wouldn't necessarily affect price" (D), although they might make the underwriter more likely to offer coverage (F). Pricing in general was not seen as consistent. Underwriter A reported surprise at the "varying prices that come out for cyber". Underwriter E believed his "technical background" meant that he was more concerned by information security.

The theme of "organisational culture" or "security culture" emerged as an important part of the final decision (B, D, F, H and I). Yet, *human resources security* was deemed of relatively low importance in both the questionnaire and the proposal forms, which suggests there is some aspect of organisational culture that is not covered by the controls outlined by ISO/IEC 27002.

The organisation's leadership, particularly within security, is seen to influence "culture". For example, D asks for the firm's CISO to attach a CV to the application, stating you would "want to have a look at the pilot's CV" when insuring a single plane. Underwriter H does "open source research on a CISO" to see if they speak at industry events and participate in information sharing schemes. However, E observed they were blind as to whether this was replicated "within that team".

There were other concerns beyond the applicant's information security posture. Underwriter D stated that "number of records" is the most important piece of underwriting information, with other respondents stating that it was among the most important (D, F, G and I).

There was consensus that assessing the aggregation of risks presents a challenge. Underwriters A, C, D, G, H and I all reported aggregation as an area they were actively seeking a better understanding of, while B, E and F seemed to be more confident in their understanding. Common service providers were identified as a major source of aggregation by (A, B, D, F, G and I); Underwriter D believed there is "no cyber equivalent" to how natural catastrophe risk are limited to one geographical location, which means there is no way to diversify a portfolio by insuring different cyber risks.

4.1.4 Evaluation and Improvements

We asked how the assessment process evolved over time. Rather than ask what has changed, we focused on what drives the change. Figure 4.6 identifies the individual drivers, with the 'Other' response relating to bench-marking against industry peers.

Figure 4.6 shows that claims data is shared between underwriters. Underwriter F highlighted how claims data improves the risk assessment process:

Method of evaluation	Frequency
Claims history from your own organisation	7
News reports	7
Security industry threat reports	6
Legal rulings and regulation	6
Publicly available data sets	5
Claims history from other organisations	4
Government level information sharing	2
Other	1

Figure 4.6: Responses to the question *Which of the following sources of information do you use to improve your risk assessment?*

> "when I first started under writing ... it was all about IT. If you look at our claims trends at my previous job and this one, the claims trends is saying 30 to 40 percent are human error."

Underwriter A said information sharing between underwriters provided "a sense [of] emerging trends" but nothing more detailed because of competitive reasons. Underwriter C identified the "fear that if you say too much to someone, they [might] steal business off you". Further, B was concerned that information sharing might constitute collusion "if we all get in the room together and say we're not happy about this type of procedure".

There was a tension between remaining ahead of emerging threats and not overlooking long term trends. Underwriter D cautioned that the Dyn attack[1] "only happened once" while there has been "hundreds of ransomware events", illustrating the danger of relying on news reports.

There are opportunities for underwriters to share expertise. Underwriter E identified informal sharing of assessment practices

> "with these conference calls everyone's on the same call. If someone asks a good question, you'll think 'that's a good question, I'll ask that in the future' "

Discussing the future of the assessment process revealed a tension between long term analytics and short term expediency; underwriter C hoped that completed proposal forms could lead to an "analytic database". However, he also believed that

[1] `https://dyn.com/blog/dyn-statement-on-10212016-ddos-attack/`

large companies "are going to move more towards meetings and calls". Underwriter D suggested that the future may involve "security architecture on a flash drive" and "security score card science".

4.1.5 Summary

The cyber insurance assessment process is provisional on factors related to the applicant and wider market conditions. Larger companies are asked to provide more information. Underwriters are limited in the amount of information they can collect by market dynamics. Proposal forms collect information about security controls, while telephone interviews provide insights into 'security culture'. External intelligence, automatic risk assessment tools and audits are infrequently used.

Areas with non-security aspects like compliance and business interruption received the most focus, as well as procedures for encrypting data at rest and in transit. Areas like access control, asset management and supplier relationships showed the most disparity between number of questions asked and how they were valued by underwriters. Despite early visions of an "insurer's checklist" [8], there is no monolithic list of security controls that applicants are assessed against.

Many technical controls would not "necessarily affect price" but might determine whether coverage was offered. Factors affecting the potential impact, like number and type of records stored, were deemed more important than security considerations. Participants reported varying prices and pricing methods for cyber insurance. Further, the method to improve this over time are ad-hoc with each underwriter managing their own education.

The next section describes the pricing data.

4.2 Pricing

This section describes how prices are adjusted for coverage, limit and deductible, revenue and industry in turn. We then provide a longitudinal perspective on cyber insurance prices.

Coverage	n	μ	σ
Cyber liability	26	1	0
Data breach/first party costs	15	0.95	0.97
Regulatory proceedings	11	0.23	0.17
Website multimedia	14	0.44	0.44
Business interruption	15	0.36	0.34
Contingent business interruption	4	0.06	0.05
Ransomware	15	0.15	0.13
Wire transfer	11	0.12	0.07
Notification Costs	7	0.21	0.08
Crisis management	9	0.1	0.11
PCI costs	5	0.24	0.15
Forensics	3	0.18	0.15
Data recovery	11	0.17	0.22

Table 4.1: Available endorsements to cyber liability coverage. The value of μ is calculated by taking the mean of the price of each coverage as a fraction of the price for cyber liability.

Pricing by coverage type Table 4.1 lists the coverage categories identified, along with its frequency of occurrence (**n**), the average price as a fraction of the cyber liability premium (μ) and standard deviation (σ). Prices for different coverages are typically expressed as a fraction of the cyber liability premium, often in a range (e.g., 0.05–0.15), in which case we used the mid-point.

The prominence of cyber liability coverage is not surprising given most filings fall under insurance lines "related [to] corporate liability policies" [73]. Market entrants began offering coverage including first party, business interruption and ransomware in later years.

Data breach and first party costs vary across insurers in terms of what is covered, often including some combination of the last five entries in Table 4.1, which explains the high standard deviation. Further, policyholders can combine notification costs, public relations and forensics to build the equivalent of a comprehensive data breach policy. This may explain the seemingly small amount of insurers offering first-party data breach coverage. Although only three and five insurers offer a forensics and PCI costs endorsement respectively, coverage may still be offered under a general "first-party" costs endorsement.

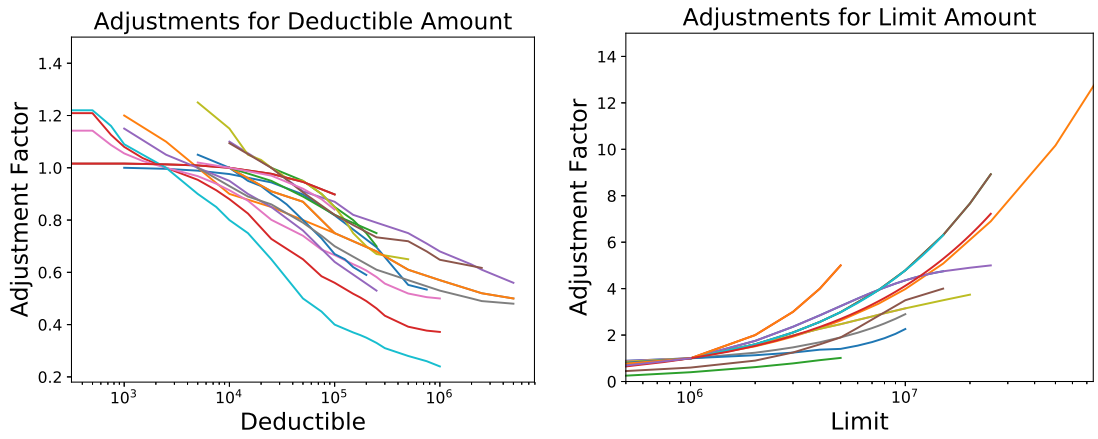

Figure 4.7: Insurers apply an adjustment factor depending on the selected deductible (left) and limit (right).

Wire transfer fraud coverage is consistently priced at around 5–15% of cyber liability. Ransomware coverage tends to be somewhat more expensive. Business interruption coverage is even more expensive still. Variance in the price of multimedia coverage could result from it being a type of third-party cover, which is traditionally difficult to price.

Pricing by coverage amount Figure 4.7 (left) shows that increasing the deductible leads to a decrease in price as we would expect. The absolute adjustment depends on the specific insurer's baseline, but we can see exponential increases in the deductible lead to linear decreases in the adjustment factor in general. Insurers use different baseline deductibles (corresponding to an adjustment factor of 1). Whereas most insurers use \$1,000,000 as the baseline limit, hence why most of the lines in Figure 4.7 (right) pass through $(1000000, 1.0)$.

The most obvious difference between insurers occurs when they adjust for higher limits. Remarkably one insurer began including an option for a \$1 billion limit in 2017, which is 56 times the price of a \$1 million limit. Figure 4.7 (right) shows that some insurers offer sub-linear increases in price for an exponential increase in revenue, which suggests losses exceeding the limit become increasingly less likely.

Some insurers use the applicant's revenue to set the adjustment for the deductible – when this was the case we used the adjustment for the smallest revenue band

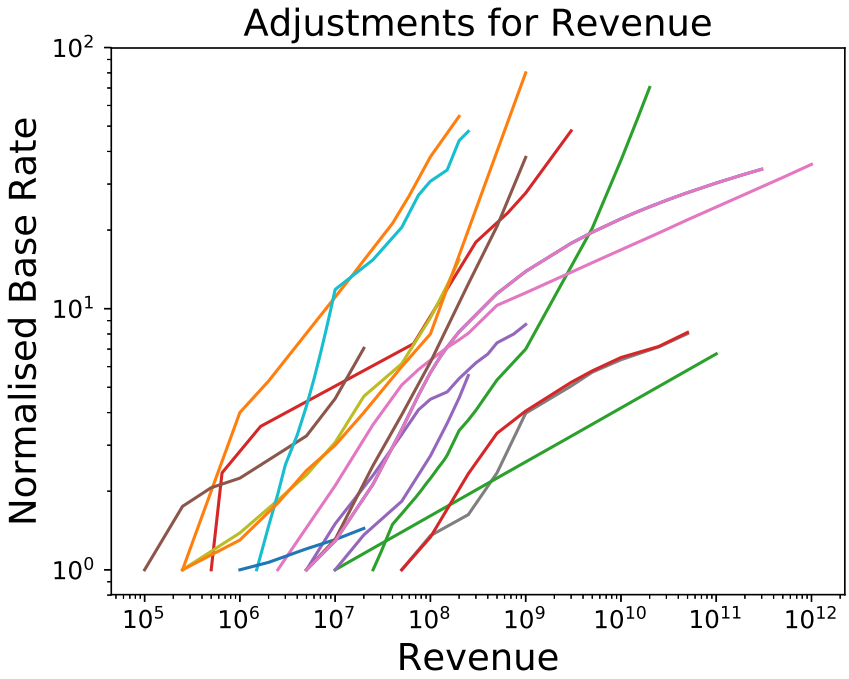

Figure 4.8: Insurers apply an adjustment factor depending on the policyholder's revenue.

available for these figures. Other than these insurers, adjustments for the deductible and limit do not depend on the coverage type or firm characteristics. This suggests the coverage type or industry has little effect on the shape of the distribution of losses, which will become important in Chapter 7.

Pricing by policyholder characteristics We cannot hope to describe in aggregate the range of adjustments for policyholder characteristics. Many of these adjustments are at the underwriter's discretion, often ranging from 0.75 to 1.25. This is further complicated by insurers describing what they are adjusting for differently, which impedes cross-comparison.

One might expect to be able to compare across adjustments for the policyholder's industry, especially given there are popular standards for defining industry. As Romanosky et al. [73] observed, "there was no consistency regarding approach [to industry definitions], or any consensus on what the insurance industry would consider the 'most' risky". As a result, we cannot compare across these adjustments because there is no consistent definition.

Table 13. Operational Modification Factor

	Range of Factors		
	Low	**Medium**	**High**
Risk	**0.75**	**1.00**	**1.25**
Risk Management Department	Yes and has Chief Info Officer or Equivalent	Yes	None
Website Sophistication Level	Informational Only; no publishing or customer info stored	Some customer info stored; no blogs or publishing or transactional/on-line sales	One or more of the following – blogs, transactional sales, publishing, web hosting
Risk Control – Section IV of App Q1-Q10	"Yes" to all questions	"Yes" to 6-8 of these questions	"Yes" to 5 or less
Info Security or Privacy evaluation performed,	"Yes" within past 24 mos. and favorable	"Yes" with most deficiencies corrected	No

SECTION III – RISK CONTROL

1.	Do you have a firewall? If yes, identify the hardware / software used:	Yes	No
2.	Do you enforce a software update process, including software patches and anti-virus software definition upgrades?	Yes	No
3.	Do you have a virus protection program in place? If yes, identify the software used:	Yes	No
4.	Do you use a standard configuration for firewalls, routers, and operating systems?	Yes	No
5.	Do you have a process for managing computer accounts, including the removal of outdated access accounts in a timely fashion?	Yes	No
6.	Do you have physical security controls in place to control access to your **computer systems**?	Yes	No
7.	Do your access control procedures address access to critical and sensitive **computer systems**?	Yes	No
8.	Do you have a written business continuity/disaster recovery plan that includes procedures to be followed in the event of a disruptive computer incident?	Yes	No
9.	Do you have a designated individual or group responsible for information security?	Yes	No
10.	How long would it take to restore your operations after a computer attack or other loss/corruption of data?		Hours

Figure 4.9: How one company adjusts the price according to security controls.

Fortunately we *can* examine how prices are adjusted for revenue. Larger companies are charged higher premiums by assigning tiered base rates according to revenue. Figure 4.8 displays the adjustments that are made according to revenue. This also shows the maximum revenue that insurers are willing to price without further consultation. It should be noted most rate schedules offer rates for larger organisations by request.

Pricing for security posture First we describe two examples at opposite ends of the objective–subjective scale. Figure 4.9 displays a pricing schema with relatively

E. Security Infrastructure Factors

The Security Infrastructure factors describe the underwriter's assessment of the technical safeguards utilized to protect a computer network and/or data. The factors reflect the level of underwriting confidence that the current security infrastructure can effectively mitigate a potential security incident. The factors are based on industry standard protocols and best practices for network infrastructure security as well as peer based analysis.

Security Infrastructure	Ranges
Excellent	0.75 to 0.84
Good	0.85 to 0.99
Average	1.00 to 1.14
Fair	1.15 to 1.24
Poor	1.25 to 1.50

Determining factors when considering Security Infrastructure:

- Are commercially available software and hardware tools utilized, including but not limited to firewalls, anti-virus and data loss prevention tools?
- Are software patches implemented on a regular basis?

Figure 4.10: How one company adjusts the price according to security controls.

objective criteria for how to apply discounts. Not having a risk management department leads to a 25% price increase, whereas having a Chief Information Officer leads to a 25% decrease. Similarly, the answers to those 10 questions translate to price adjustments of around 5% per question. This can be contrasted against Figure 4.10 which relies heavily on the underwriter's judgement of the security infrastructure.

For the fabled discounts for security controls, we cannot easily compare across discount factors because they are defined differently by each insurer. We would have to justify why adjustments, for example "privacy controls", "encryption" and "encryption for network data, laptops and mobile devices", refer to the same set of controls. Even if we did, the results would simply show the underwriter's subjective assessment of these controls leads to a discount or price increase of up to 25%.

Pricing over time Base rates are typically given for $1,000,000 of cyber liability coverage, with different rates depending on the organisation's revenue band. Figure 4.11 shows the base rate for a retailer with $50 million of revenue, along with the date the policy was introduced.

Three insurers updated their rate schedules with new prices, providing the only *apples to apples* comparison because the prices relate to the same underlying

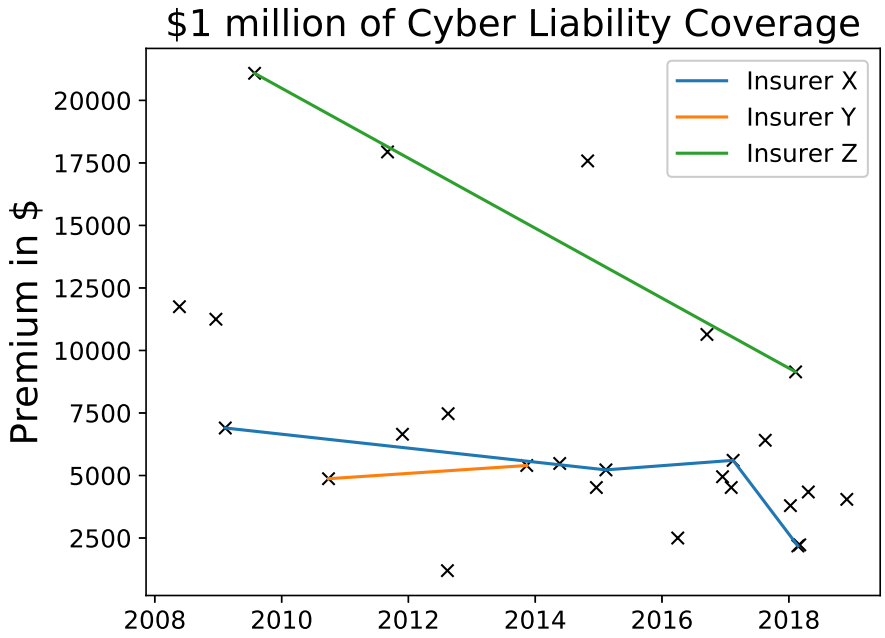

Figure 4.11: Cyber liability insurance premiums over time for selected insurers.

policy. Insurer X updated their prices four times with a notably sharp decrease in 2018. Meanwhile, Insurer Y had a minor increase in price. Insurer Z reduced their price by over 40% after 9 years. It is unclear whether prices for new and existing customers gradually fell during this period.

The trend in prices among market entrants is more difficult to assess due to subtle differences in the products being sold (although these differences are often overstated [73]). The cheapest policy was introduced in 2012 and the maximum limit was $5M, which suggests it was targeted at smaller companies. We would need to control for changes in coverage to make reliable claims about how prices change over time[2]. Although, it is fair to say market entrants and price updates have become more frequent.

[2]Anecdotally, industry insiders suggest coverage has become broader. If this is true, then Figure 4.8 suggests prices are falling, which would be consistent with greater competition.

4.3 Discussion

The results related to risk assessment and pricing are discussed in Section 4.3.1 and Section 4.3.2 respectively. We then discuss how they relate to each other in Section 4.3.3.

4.3.1 Risk Assessment Discussion

The results depend on a sample that was limited by the total number of insurers selling cyber insurance in the UK and their availability. While psychology studies can recruit from the local undergraduate pool, cyber insurance professionals may not be as motivated by, for example, entry into a raffle to win an Amazon voucher. In spite of this, we conducted a similar number of interviews as [72], analysed half as many proposal forms as [73, 22] and collected questionnaire responses.

Using mixed methods allowed us to directly compare the results from analysing proposal forms to engaging insurers; it is difficult to compare qualitative studies such as [72] to quantitative studies like [73, 17]. There seem to be disparities between the two casting doubt over whether document analysis can provide insights into how underwriters assess risk — a previously unexplored premise upon which [76], [73] and [17] were based. However, we would need to collect more data to make any concrete claims.

Disparities may result from: (i) proposal forms are not the only method by which information is collected; (ii) the sub-sections may differ in size, meaning they require different amounts of information to be collected; and (iii) it may require fewer questions to establish the security measures of certain areas of information security (*informational efficiency*).

Alternatively, the questionnaire responses could provide an imperfect measure of how the insurance industry assesses cyber risk. Yet asking insurance professionals how they assess risk faces its own challenges. For example, acquiescence bias meant no participants deemed a section of ISO/IEC 27002:2013 [101] unimportant.

A critical view of each methodology suggest the "triangulation" [99] of mixed methods might be more appropriate. For instance, the content analysis identified all

of the risk assessment considerations identified in [73]. However, portfolio concerns and the importance of organisational culture as assessed through security calls was overlooked. Similarly, although the qualitative approach in [72] identified the former two concerns, it could provide little insight into which areas of information security are assessed.

Turning to how our analysis relates to ISO/IEC 27002, the results should not motivate generic recommendations for improving proposal forms; doing so would be based on a normative assumption that ISO/IEC 27002 provides an authoritative view of appropriate information security. This denies insurance professionals' experience and expertise in assessing risk.

The results might instead be seen as a way to think about the cyber insurance risk assessment process. As [101] states,

> "Organizations that adopt ISO/IEC 27002 must assess their own information risks, clarify their control objectives and apply suitable controls using the standard for guidance."

In this vein, the insurance industry could use these results to reflect on why they ask the questions they do — paying particular attention to whether technical controls are assessed at the right conceptual level.

However, the application process faces similar problems to auditing in determining efficacy. Power [39] identifies how regulators and auditors avoid probing the ability of audits to achieve their stated goal. Underwriters earn a living by assessing risk and cannot be reasonably expected to question their ability to accurately do so.

The information security community could reflect on insurers' unique perspective in seeing the losses from many more organisations than a risk manager within a single organisation might. Further, they have expertise in financially quantifying the impact of cyber attacks, which IT managers have historically failed to articulate within their organisation. Perhaps the community could reflect on whether they are overly focused on technical controls at the expense of understanding the culture of the organisation implementing them.

A more critical lens would question the insurance industry's focus on internal controls, especially given the "very short period of time over which it has come about" [42]. For example, insurance professionals may be incentivised to represent themselves as cyber gurus [43] to confer legitimacy.

Insurance application forms can be used by internal actors to justify investments [40]. This becomes problematic when these forms are influenced by the vendors and consultants who advise insurers on security. The objectivity associated with actuarial expertise could disguise the commercial interest in providing advice.

Portfolio concerns are leading the insurance industry to consider the systemic risk associated with common service providers, as evidenced in the interviews. An efficient market would price insurance accordingly, potentially offering discounts if they increase the diversity of service provision. This is a new consideration compared with traditional information security risk assessment, although the insurance industry is still seeking a better understanding.

4.3.2 Pricing Discussion

The insights gained from the regulatory filings depend on the extent to which insurers comply with them. While it is possible that insurers file one thing and do another, it is more likely that the scope for subjective judgement introduces uncertainty into how pricing actually occurs. On the other hand, regulators have the power to revoke licenses. This incentive to accurately report pricing schemes is far greater than any incentive that researchers can provide for accurate reporting in interviews and surveys. We will now turn to the findings themselves.

Pricing different coverage types by multiplying the price of cyber liability coverage by a constant suggests each type of cyber loss is driven by a similarly shaped distribution. Two separate studies [84, 85] that find the recovery time for IT systems to be well modelled by a lognormal distribution, much like number of records in [80]. Although we cannot map these distributions to financial costs, this evidence supports the viability of multiplicative pricing.

The adjustments for policyholder revenue cast further doubt over aggregate statistics regarding expected losses that group together companies regardless of revenues, for example [75]. Some insurers force large companies to pay 100 times the premium that smaller companies do, suggesting the expected loss is up to 100 times as large. This increases the effect of the composition of firms in the sample and ties into some of the problems with surveying cyber crime cost estimates [97].

There is very little to learn about adjustments for security because the underwriter has so much scope for case-by-case judgement. However, Figure 4.9 shows that a 25% discount requires all 10 questions to be answered satisfactorily. This means the marginal benefit of an additional security control can be expected to be some fraction of 25%. This calls into question how much of an incentive insurance actually provides for firms to invest given the relative cost of security controls. This contradicts a common theme in the literature [8] about insurers offering significant incentives for better security.

The scope for subjective judgement is in keeping with the finding that life insurance is "based on guesswork in support of gambling, with both unexpected windfalls and catastrophic losses" [16]. If life insurance falls short of an exact science despite mortality tables going back hundreds of years, there is very little hope for cyber insurance. The next subsection asks about how the pricing and risk assessment results relate to each other.

4.3.3 Relating Risk Assessment to Pricing

What can we learn from comparing US pricing with UK risk assessment? One could argue corroborating evidence can be generalised to both markets, whereas contradictions might only result from the difference in markets. As such we will focus on discussing what both studies find evidence for, as this is more insightful. For example, there is a lot of variation between insurers in terms of how they conduct risk assessment in the UK and how they price insurance in the US; this suggests neither market has reached the maturity of a stable industry best practice.

First, comparing the results of studies of the US [73] and Swedish [72] markets with our study of risk assessment suggest cyber insurance is assessed similarly to the UK market. Or at least, the coarse research methods cannot identify how the markets differ.

In Chapter 2 we identified how the majority of insurers do not use security information to price risk [73] while still collecting security information. Some underwriters explained that security information was instead used to determine whether coverage would be offered, providing a partial explanation; others believed that, as the market matured, the link between security and losses will become clear. The latter suggests that security information may be collected for analytical purposes in the future.

The interview responses corroborate the finding from studying pricing schemes that discounts for individual security controls are possibly over-exaggerated. Respondents suggest discounts are based on a holistic view of an applicant's security and they cannot quantify the effect of a single control. One interviewee suggested accreditation to standards leads to a "not necessarily massive discount", while another said "it wouldn't be like 10% discount for if you take good logs".

The longitudinal analysis in Figure 4.11 reveals remarkable variance in pricing. This is corroborated by reports about the "varying prices that come out for cyber" (p.60). Further, the scope for subjective judgement may provide space for underwriters to make ad hoc adjustments for "organisational culture". This was identified as an important part of the final decision by many participants (B, D, F, H and I).

There is an interesting discrepancy between individual underwriters' efforts to understand information security, as exemplified by formal education or time invested in news reports, and the infrequency by which insurers file new pricing schemes. This suggests there could be ongoing informal innovation in terms of risk assessment and pricing that is not reflected in regulatory filings. Focusing on filings alone led Romanosky et al. [73] to fail to "observe any substantial changes in policy length, style, or composition over time".

4.4 Conclusion

We used mixed methods to explore the actors and artifacts involved in the cyber insurance assessment process. Our findings suggest there is no consistent assessment process; insurers vary the depth and methods of assessment according to the applicant. The assessment process balances the applicant's information security practices and with the underwriter's portfolio and wider market forces. The complex and contextual nature of the assessment process presents challenges for researchers seeking to model cyber insurance.

A number of important factors did emerge. Insurers tend to collect information using yes–no questions. This suggests they conceive of security as a discrete set of variables, as in the IWL [71], rather than as a continuous security level as in [69]. In Chapter 5, our model extends the IWL [71] to more accurately represent the information that insurers collect.

There was widespread uncertainty among insurers about what constitutes effective security, as evidenced by underwriters' efforts to educate themselves about information security. This is mirrored by the inconsistent prices identified in Section 4.2. This seems to be more realistically modelled by the uncertainty parameter in the IWL [71], which we can contrast against the Gordon and Loeb model's assumption that what constitutes security is known and the only question is how much to invest in it [69].

The survey showed that "claims history from your own organisation" is used by the majority of insurers to improve their risk assessment. This suggests the model should allow historical information to influence decisions. The multiple rounds of the IWL [71] capture this aspect.

The pricing data will be most relevant to Chapter 7. In particular, the adjustments for industry and revenue will allow us to gain granular insights for specific firm types. The disparity between how insurers adjust premiums for the policy limit will drive disparities between our inferences, which we address by aggregating inferences in the hope random noise will cancel itself out.

5

Insurer Strategies for Sharing Information

The preceding chapter presented observational data about the cyber insurance market. This provided a description of the world but no insights into the mechanisms and processes driving security outcomes. For example, there is evidence that insurers assess security and adjust the insurance price based on this. But we cannot tell whether this changes security decisions as compared to the counter-factual of no insurance market existing.

In this chapter we investigate how cyber insurance changes the information structure. Chapter 2 showed that information asymmetry, in which the insured has private information, has received a lot of attention [31, 60, 61] but there has been no consideration of insurers lacking knowledge about which controls are effective. Yet, Chapter 4 showed that insurers are uncertain about assessing and pricing cyber risk.

This chapter's contribution consists of modelling this uncertainty. We explore different strategies by which insurers can share information to reduce uncertainty. The model allows us to measure security budgets, attacks suffered and social welfare.

Section 5.1 begins by justifying our modelling choices, arguing they improve the realism as compared to previous models. We then simulate three strategies for sharing information and present the results in Section 5.2. Finally, we discuss the implications of the results in Section 5.3. We offer a chapter summary in Section 5.4. The work of this paper appeared previously in [19].

5.1 Model

Our extension, the Iterated Weakest Link – Cyber Security Insurance (IWL-CSI) introduces new rules, new strategies and a new method of computation in extending the original model to consider m policyholders purchasing insurance from a single insurer. In Section 5.1.1, we introduce those new rules. We outline the passive, active and diverse strategies for the insurer in Section 5.1.2. In Section 5.1.3, we explain why Monte Carlo methods are used to simulate the IWL-CSI.

5.1.1 New rules

The IWL-CSI introduces the parameter m to represent the number of policyholders that an insurer may share claims information with. To model the insights gained from aggregating claims data, information about attacks against one policyholder must be relevant to the attacks other policyholders might face in future rounds.

The IWL-CSI makes two assumptions: (i) the true cost of exploiting a given vulnerability is the same for each defender; and (ii) the defenders can adopt different defensive configurations, which the insurer influences. Both (i) and (ii) are strong assumptions that increase the value of claims information and allow the insurer to mandate security controls respectively (although the passive insurer will not use this power).

To understand (i), consider that network effects empower software monopolies [2], which leads to a lack of so-called "cyber diversity" [115]. The result of policyholders adopting similar information systems is that they share vulnerabilities. Consequently, if the insurer learns that the i-th vulnerability is used to attack one insured, then that same vulnerability will be economically viable in other insured' systems.

We assume (ii) because diverse defensive configurations allow information about the attacker to be shared and collective uncertainty reduced. For example, different organisations protect the same operating system with different security products. The insurer is assumed to have control over which defensive configurations are in place for each policyholder. Insurers have expressed a "desire" to recommend security controls [22] and some insurers even offer "a list of actions to be taken" [72]

Description	Symbol	Default Value
Busines Model		
Number of insureds	m	4
Asset value	a	1000
Total number of rounds	t_{max}	25
Return on asset per round	r	0.025
Attacker		
Number of threats	n	25
Loss given attack (as a fraction of asset value)	z	0.025
Expected minimum attack cost	$\bar{x}_0$	15
Attack gradient	Δx	1
Level of uncertainty	σ	1
Defender		
The i-th defense of the j-th defender	d_{i_j}	0 or 1
Cost of each defense	1	1
Defence interdependence	ρ	0.1
Sunk cost (as a fraction of asset value)	λ	0

Table 5.1: Describing each of the parameters in the motel

to improve security. Of course, this ability will be dependent on the wider market; an under-supply of insurance allows the insurer greater freedom to select with whom they enter into contract.

To translate these assumptions into the model, we assume homogeneous defenders so that for the j-th and k-th defender we have: $a_j = a_k$, $t_{max_j} = t_{max_k}$, $r_j = r_k$, $n_j = n_k$, $z_j = z_k$, $x_{1_j} = x_{1_k}$, $\Delta x_j = \Delta x_k$, $\sigma_j = \sigma_k$, $\rho_j = \rho_k$, and $\lambda_j = \lambda_k$. We will drop the index for each defender to ease notation, unless using it provides clarity. Table 5.1 describes each parameter, along with the default value.

Defensive costs related to the interdependence of controls employed by an organisation are determined by an $n \times n$ matrix C_{int} (as in the original IWL [71]). Additionally, we introduce an $m \times m$ matrix C_{ext_i} for each of the n possible defenses. These matrices reflect the extent to which defenders can co-operate to take advantage of returns to scale, which may vary depending on the particular defensive investment d_{i_j}. The cost to the j-th defender of employing the defensive configuration $\mathbf{d}_{i_j}$ is

$$C_j = \mathbf{d}_{i_j} C_{int} \mathbf{d}_{i_j} + \sum_{k=0}^{n} \mathbf{b}_k C_{ext_k} \mathbf{b}_k \tag{5.1}$$

where $\mathbf{b}_k \in \{0,1\}^m$ with $b_{k_t} = d_{k_t}$. The j-th element of $\mathbf{b}_k$ represents whether the j-th defender employs controls k. To remain in the scope of this study, the matrices C_{ext_i} are chosen so that defensive costs scale linearly in the number of insureds, calculated using Equation 2.5. In Section 5.3, we discuss how future work might modify this assumption to explore non-linear scaling.

As in [71], the attacker will exploit each defender's unguarded vulnerability with the lowest cost of attack, unless this is greater than the loot value, za. The true costs of attack are modelled as follows:

$$x_{i_j} = \sup(0, \chi_i) \text{ where } \chi_i \sim \mathcal{N}(\bar{x}_i, \frac{\sigma}{\Delta x}) \text{ for } j = 1, ..., m \tag{5.2}$$

Consequently, $x_{i_j} = x_{i_k}$ for all $i \in \{1, ..., n\}$ and $j, k \in \{1, ..., m\}$, where n and m are the number of vulnerabilities and insureds respectively. To identify the "weakest link", we must maintain an ordering on the set of vulnerabilities and use the untruncated values χ_i to do so. If more than one vulnerability had a true cost of 0, it is not clear which would be exploited first. For $x_{i_l} = 0 = x_{k_l}$, we say that i is the weakest link if $\chi_i < \chi_k$, where the values of χ_i and χ_k are determined by Equation 5.2.

5.1.2 Novel strategies

In this subsection we characterise the passive, active and diverse strategies. These explore different approaches the insurer might employ to discover which controls are effective in mitigating losses. At a high level we can say that the passive insurer is hands-off, the active approach ensures the policyholders have different levels of security, and the diverse insurer maintains diverse security configurations similar to a controlled experiment. The active and diverse approaches assume different levels of confidence in terms of where the attacks might land. Each strategy is illustrated in Table 5.2.

As the passive insurer does not share claims information or mandate controls, we can assume that each policyholder acts rationally by adopting the optimal defensive configuration. The game reduces to the original IWL with m different policyholders facing the same realised true costs with no ability to communicate

	Passive			**Active**			**Diverse**		
	x_1	x_2	x_3	x_1	x_2	x_3	x_1	x_2	x_3
Policyholder A	o			o				o	o
Policyholder B	o			o	o		o		o
Policyholder C	o			o	o	o	o	o	

Table 5.2: Illustration of the different strategies with $m = 3$.

with each other. Consequently, the policyholders can be expected to adopt the same strategy as in the original IWL.

For the active insurer to share claims information, its policyholders must adopt different defensive postures. Otherwise they will each gain identical information about the attacker. We assume the m policyholders have different security levels driven by some combination of premium incentives, mandated controls or differing risk aversion levels. These assumptions are supported by evidence from Chapter 4 that insurance coverage is offered to applicants with different security posture (possibly at different prices). This can be modelled by assuming that the i-th policyholder guards one more vulnerability than the $(i-1)$-th policyholder. If a given vulnerability is exploited in any policyholder, it will be protected by every policyholder in the following round. This part of the strategy is facilitated by the active insurer sharing claims information.

Finally, the diverse approach ensures the choice of defensive configurations maximises the amount of information gained. Doing so would require the insurer to offer premium incentives or mandate security controls as no rational defender would protect a vulnerability while one that was more likely to be exploited was left unguarded. For the n most likely vulnerabilities to be exploited, the i-th policyholder guards all n vulnerabilities apart from the i-th. Thus, if the i-th policyholder is attacked, then the diverse insurer knows that the i-th vulnerability is economically viable to the attacker. Consequently, the insurer can gain up to m pieces of information per round (one for each policyholder), at the cost of each policyholder implementing different, and possibly more expensive, defensive configurations.

We now turn to the calculation.

	x_1	x_2
Policyholder A	x	
Policyholder B	o	x

x_1'	x_2'
	x
o	x

Table 5.3: An illustration of how the ordering of x_l and x_l' impacts the realisation of attacks.

5.1.3 New Method of Calculation

We demonstrated that simulations can be used to calculate returns on the original IWL in Chapter 3. Recall from Chapter 3 that the return on the original IWL [71] depended only on the number of economically viable vulnerabilities. It was independent of the ordering of vulnerabilities. We now explain why this is not true for the IWL-CSI.

In the IWL-CSI, the ordering of vulnerabilities determines how much information is revealed per round. For example, consider the two separate realisations of true costs of defence x_l and x_l' in Table 5.3. The defender expects the true cost of exploiting the first vulnerability to be lower than the true cost of exploiting the second vulnerability. In both cases the first policyholder protects neither and the second policyholder protects the first vulnerability.

If the expectation is realised ($x_1 < x_2$) as in the first case, then the insurer discovers that both vulnerabilities are economically viable. If the ordering is not as expected ($x_1' < x_2'$), then the attacker exploits the second vulnerability in both policyholders. The insurer discovers that the second vulnerability is economically viable but there is still uncertainty about whether the first vulnerability is exploitable. Consequently, the ordering of the true costs of attack impacts the IWL-CSI. However, the approximation used in the original IWL does not account for this.

5.1.4 Summary

The original IWL consists of the rules, strategy and method of computation to consider the strategic interaction between one defender and an attacker. In Section 5.1.1, we extended the model to consider m defenders. Based on the new rules, we identified three separate strategies an insurer might employ in Section 5.1.2.

Then, in Section 5.1.3, we provided some validation for using Monte Carlo methods to compute the original IWL and explained why simulations are better suited to computing the new rules. In Section 5.2, we will use the new rules and Monte Carlo methods to consider strategies the insurer of m policyholders might employ.

5.2 Results

We compare the average revenue of the three strategies in Section 5.2.1 and consider how these might be priced in Section 5.2.2. In Section 5.2.3, we look at the variance of claims. We then explore the impact of varying the number of insureds in Section 5.2.4. Apart from the parameters for uncertainty, sunk costs and numbers of insureds, all of the simulations use the same values as in [71].

5.2.1 Comparing the passive, active and diverse strategies

The three strategies can be compared in terms of attacks suffered, defensive spending, or revenue. Defensive spending is likely incurred by the policyholder while attacks suffered may result in claims paid out by the insurer. How these costs are divided between insurer and insured will depend on the particular insurance contract, which we will look at in the next subsection. This subsection compares the average revenue for each strategy because it reflects attacks suffered and defensive spending without assuming a particular contract.

All of the results in this subsection are based on simulations with four insureds. With this parameter value, the insurer can gain information related to a maximum of 16% of the vulnerabilities per round (often the number will be less). The choice of four illustrates the differences between the three strategies by choosing a middle 'road'. Section 5.2.4 shows the effect of changing the number of insureds.

Figure 5.1 shows the result of 10^7 simulations without sunk costs. Displaying results like this highlights the optimal initial defensive configuration for a given strategy. For high uncertainty ($\sigma = 16$), it is still rational to under-invest and "wait and see" [71], but sharing claims information does result in a higher revenue.

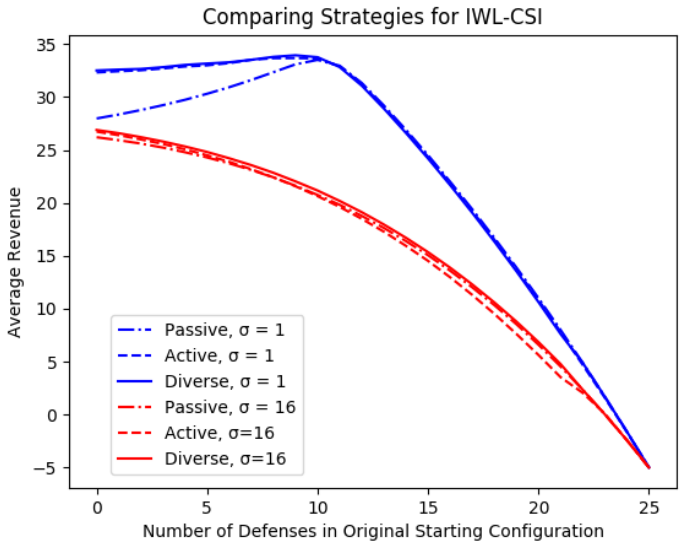

Figure 5.1: Average revenue for each of the three strategies based on 10^7 simulations with four insureds and no sunk costs.

For low uncertainty ($\sigma = 1$), investing less initially but sharing claims information will result in higher average returns.

For risk-averse strategies involving high initial investment, the passive strategy becomes superior. When the number of initial defenses (k) is high, observed attacks are less likely and there will be less value in a strategy that uses information gained from observing attacks. However, these are strictly sub-optimal and we will stop displaying these initial configurations as they result in lower revenue than just accepting an attack every round. These diagrams are not well-suited to comparing between the passive, active and diverse strategies.

Figure 5.2 and Figure 5.3 provide a comparison of active against passive and diverse against active respectively. Figure 5.2 shows that using claims data is most beneficial when the policyholders begin significantly under-invested and uncertainty is low. A policyholder not sharing claims information suffers many attacks in determining which vulnerabilities are economically viable.

Sharing claims information is least beneficial when the defender invests heavily in defence and uncertainty is low. The cost of maintaining diverse defences is greater than the value gained by observing attacks because so few attacks are observed.

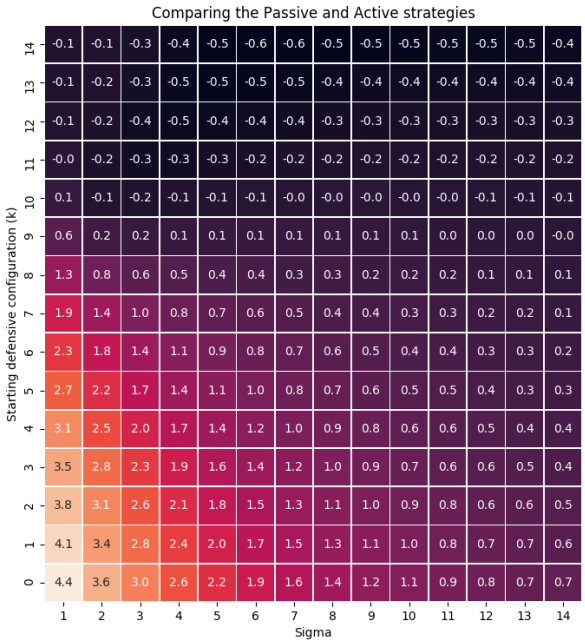

Figure 5.2: Improvement in revenue gained by adopting the active strategy as opposed to the passive strategy.

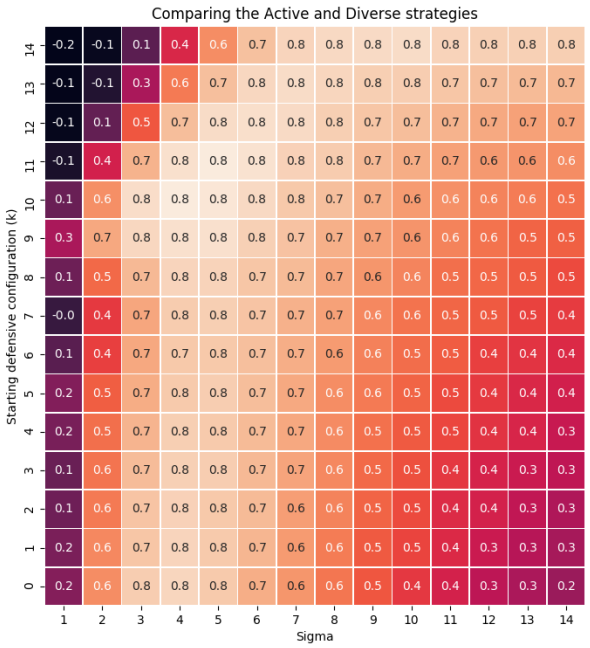

Figure 5.3: Improvement in revenue gained by adopting the diverse strategy as opposed to the active strategy.

	Level of uncertainty								
	$\sigma = 1$			$\sigma = 4$			$\sigma = 16$		
Strategy	Passive	Active	Diverse	Passive	Active	Diverse	Passive	Active	Diverse
No sunk costs ($\lambda = 0$)									
Optimal defence k^*	10	9	9	8	6	6	0	0	0
Attack intensity I (% rounds)	2.7	1.7	1.5	10.0	5.4	4.6	45.5	33.6	30.0
SD of attack intensity σ_I	0.026	0.009	0.007	0.030	0.013	0.008	0.100	0.079	0.075
Avg. gross return (% asset)	33.5	33.7	33.9	32.3	33.2	33.3	26.2	26.7	26.9
Avg. security spending (% asset)	15.8	15.9	15.7	15.7	15.5	15.6	12.4	15.9	15.6
ALE_0	0.68	0.43	0.38	2.50	1.35	1.14	11.4	8.4	7.5
ROSI (% security spending)	53.8	54.5	56.9	48.4	53.1	53.1	9.9	4.4	12.2
Loading factor limit l^*	1	1.38	2.13	1	1.63	1.82	1	1.77	1.95
Total load limit L^*	0	0.17	0.43	0	0.86	0.94	0	6.45	7.13
Sunk costs ($\lambda = 0.025$)									
Optimal defence k^*	10	9	9	10	9	9	0	0	0
Attack intensity I (% rounds)	2.7	1.74	1.53	7.1	4.9	3.7	45.5	33.7	30.0
Avg. gross return (% asset)	32.8	33.6	33.9	28.2	29.7	30.8	14.8	18.5	19.5
Avg. security spending (% asset)	16.5	15.9	15.7	19.9	19.1	18.3	23.8	23.1	23.0
ALE_0	0.68	0.44	0.38	1.9	1.2	0.9	11.4	8.4	7.5
ROSI (% security spending)	47.4	54.3	56.9	16.2	24.7	33.6	-42.8	-28.2	-23.9
Loading factor limit l^*	1	2.89	3.91	1	2.21	3.81	1	1.44	1.63
Total load limit L^*	0	0.82	1.11	0	1.48	2.56	0	3.70	4.73
Memo item: No defence									
Avg. gross return (% asset)	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.5
ALE_1	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0

Table 5.4: The results of 10^7 simulations with and without sunk costs.

However, Figure 5.1 shows that high initial investment will result in lower revenues and can only be justified by extreme risk-aversion.

Figure 5.3 compares the active and diverse strategies, which offer two different approaches to gathering claims information. The diverse strategy is most effective when uncertainty is moderate; the active strategy requires accurately forecasting the ordering of the true costs of exploitation. When uncertainty is low, the active strategy accurately predicts the ordering and so there is little benefit in adopting the diverse approach. The diverse strategy can tolerate a threshold of inaccuracy. It performs similarly to the active strategy when the threshold is exceeded, which occurs when uncertainty is high and there is a low starting configuration.

Table 5.4 displays a number of security investment metrics. Broadly speaking, the active and diverse strategies result in adopting a less secure initial defensive configuration with a higher gross return. In every case, the passive strategy results in a higher attack intensity I, which describes the proportion of rounds in which a policyholder suffers an attack.

The relative spend on security depends on the trade-off between the costs incurred making predictions about which vulnerabilities might be exploited and the

ability to adopt a less costly initial defensive configuration, which must be incurred for all future rounds. When uncertainty is high, the passive strategy spends less on security because it only defends the vulnerabilities that are exploited. However, the passive strategy defends vulnerabilities that may not even be economically viable when uncertainty is moderate. On the other hand, the diverse strategy under-invests initially but gains information quickly, which results in a lower average security spend.

Different parameter values reward the balance between making no predictions as to which vulnerabilities might be exploited (passive strategy), relying on predictions as to which vulnerabilities are important to defend (active strategy) and accepting these predictions are likely to be misguided (diverse strategy). Without sunk costs, the active and diverse strategy achieve the same return on security investment (ROSI) for moderate uncertainty ($\sigma = 4$), whereas, the diverse strategy ROSI is three times that of the active strategy for high uncertainty ($\sigma = 16$). In fact, the passive strategy achieves a higher ROSI than the active strategy in this case. The active strategy is being punished for the poor predictions as to which vulnerabilities are important to defend; these predictions are poor because uncertainty is so high.

When sunk costs are introduced, the active strategy begins to achieve a higher ROSI than the passive strategy even when uncertainty is high. Despite the active strategy's predictions being poor, the few that do gain new information mean the insureds need to make less costly defensive configurations. Indeed, this is even more true for the diverse strategy in the case of moderate uncertainty, which achieves twice the ROSI of the passive strategy. This suggests sharing claims data is especially valuable in guiding investments when it is costly to change defensive configuration.

It is rational to abandon all three strategies and accept a loss each round when sunk costs are present and uncertainty is high. This is evidenced by the negative ROSI values for all three strategies. The implication is that there are values of σ and λ for which the only rational way to invest in security is if aggregated claims data can guide the investments.

5.2.2 Pricing

The pricing of each policy determines how the revenue is divided between policyholder and insured. For each policyholder, the insurer should expect an annual loss expectancy of

$$ALE = (za)I \tag{5.3}$$

per round. Consequently, the insurance premium will be priced at $l.ALE$ where $l > 1$ is a loading factor to account for the insurer's profit and operating costs.

If the policyholder purchases insurance, they will receive a return on the asset of ra minus the insurance premium $l.ALE$ and the cost of defence C, equal to

$$R = ra - l.ALE - C \tag{5.4}$$

This is to be contrasted to the case without risk transfer where the policyholder's revenue is dependent on the random realisation of attacks. A risk-averse policyholder might prefer to lose a guaranteed $l.ALE$ rather than risk losing za, even if purchasing insurance results in lower expected revenue than can be expected without purchasing insurance.

The insurer sharing claims information and guiding security investments may reduce the attack intensity I sufficiently to offset the loading factor l. We define the loading factor limit l^* to be the point at which purchasing insurance does not change the expected revenue. If ALE_1/C_1 and ALE_0/C_0 are the expected annual loss expectancy/cost of defence with and without the insurer guiding investments respectively, then

$$\lambda a - l^* ALE_1 - C_1 = \lambda a - ALE_0 - C_0 \tag{5.5}$$

Rearranging, the loading factor limit is given by

$$l^* = \frac{ALE_0 + (C_0 - C_1)}{ALE_1} \tag{5.6}$$

The total load limit L^* is defined to be the amount of premium that remains after the expected number of claims are paid,

$$L^* = (l^* - 1)ALE_1 \tag{5.7}$$

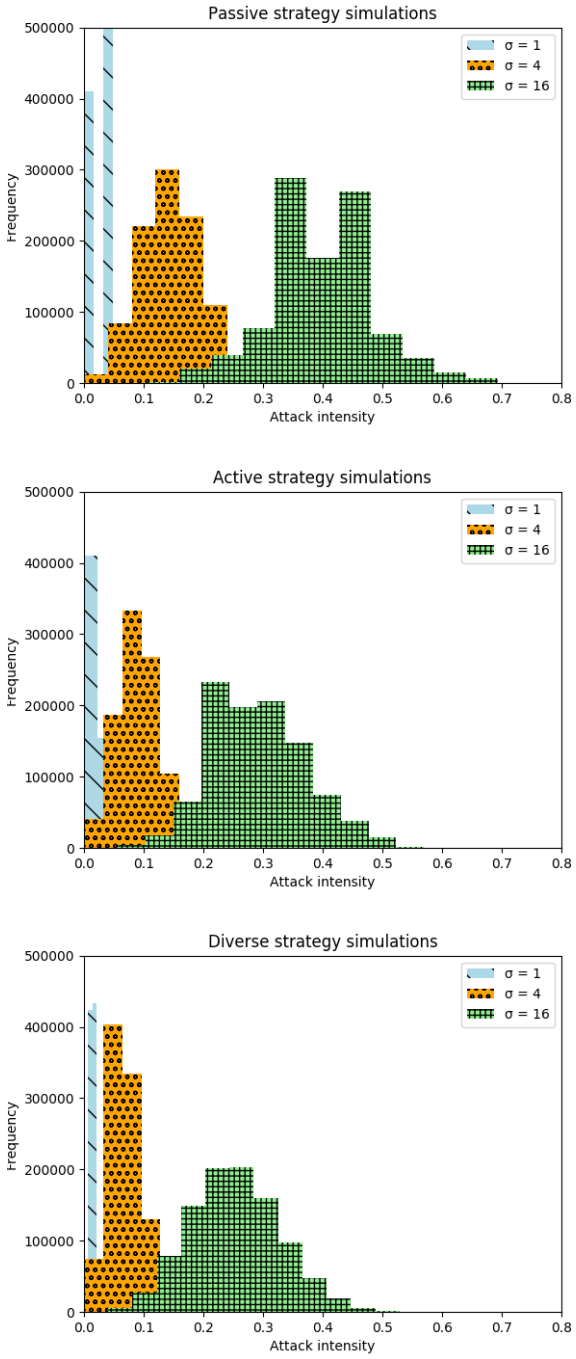

Figure 5.4: Distribution of claims for the passive, active and diverse strategies.

Table 5.4 contains the calculations of l^* and L^* based on the simulations. The results suggest that the insurer can charge a higher percentage of the premium to cover operating costs and profits when sunk costs are present. The value to the policyholder in sharing claims information is higher when it is costly to change defensive configuration frequently. However, the size of the premium is far higher in absolute terms when uncertainty is high, which is driven by the higher attack intensity. However, the insurer must hold capital reserves to cover potential claims, which contributes an additional cost. The next subsection investigates this aspect.

5.2.3 Measures of dispersion

Using Monte Carlo methods allows us to explore the dispersion of the results. We focus on the distribution of the attack intensity (I) because it determines the variability of claims the insurer faces.

Figure 5.4 highlights how the distribution of claims differs for the optimal starting configuration for each strategy. Comparing the three figures shows that the passive strategy suffers more attacks for each uncertainty value. As uncertainty increases, both the mean and variance of the attack intensity increases. The insurer must hold more capital to account for the higher mean attack intensity and additional capital must be held to account for the variance of the attack intensity. Alternatively, the insurer may choose to purchase reinsurance to cover the tail of these attack intensities.

The passive strategy simulations illustrate how ordering is not important for the original IWL; when $\sigma = 1$ the two bars capture the two cases when 0, 1 or 2 economically viable exploits are left unguarded. Recalling that the passive strategy is equivalent to the original IWL, these results suggest that the "wait and see" approach that was found to be optimal for high uncertainty in [71] leads to a concerning variability in the number of attacks faced. Table 5.4 shows that both the mean and the standard deviation of the attack intensity are lower than for the passive and active strategies for all uncertainty levels. This suggests that less capital would need to be held in reserve for the strategies that involve sharing claims data.

5.2.4 Varying the number of insureds

The prior results show that sharing claims information is superior to not sharing it. However, the results have, with a few exceptions, not revealed a significant difference between the active and diverse strategies. This subsection does not display the results for the passive insurer because the passive insurer does not share information. Consequently, the average revenue is the same for all number of insureds.

Figure 5.5 varies the number of insureds, displaying the average revenue for the optimal starting configuration for each parameter choice. The active strategy is optimal for between two and four policyholders for all the uncertainty levels we consider. For the diverse strategy, the optimal number of policyholders increases with the level of uncertainty. Figure 5.5 shows that for $\sigma > 4$, 16 policyholders leads to the highest revenue. In fact, the passive strategy outperforms the active strategy when m and σ are high.

When m is high, the active insurer diversifies across 16 different security levels. High uncertainty σ means that less information is collected because the ordering of true costs of attack is not as the insurer expects. For example, the insurer gains 16 pieces of information if the ordering of true costs is as expected. Suppose the true cost of attack x_j, which is expected to be j-th in the ordering, was actually the i-th lowest. Then, for defenders i through to j, x_j would be left undefended and would be exploited by the attacker. Thus only one piece of information is gained by the $j - i$ defenders. The more the ordering diverges from expectations, the less information is gained — yet the policyholders expend a high cost in implementing the different security levels.

Figure 5.5 allows us to consider the marginal benefit of insuring another policyholder. When uncertainty is low, insuring an additional policyholder will not significantly affect the average revenue for either of the active or diverse approaches. The absolute difference between all numbers of policyholders is less than 0.5 when $\sigma = 1$. As uncertainty increases, the marginal benefit diverges for each strategy. For the active approach, there is a small gain in moving from 2 to 4 policyholders and the marginal benefit of another policyholder is negative. However, for the diverse

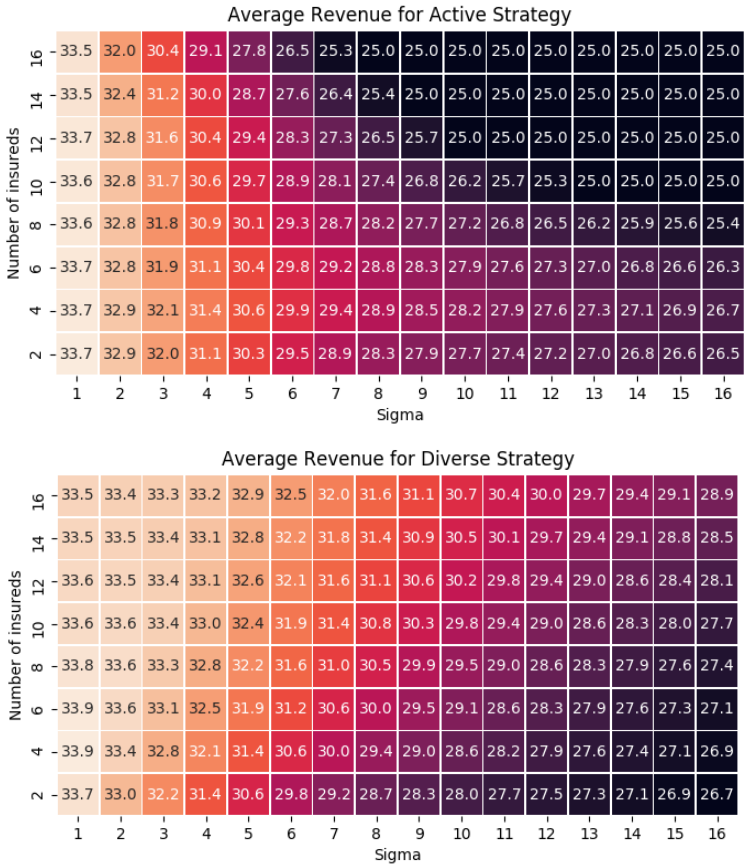

Figure 5.5: Average revenue for 10^7 simulations with different numbers of policyholders and uncertainty levels for the active and diverse strategies.

strategy, this benefit is always positive for $\sigma > 8$ and for $\sigma = 16$ the difference between 2 and 16 policyholders is four times that for $\sigma = 1$. The implications of these results in relation to market composition are discussed in the next section.

5.3 Discussion

Our motivating concern is the ability of the insurer to aggregate claims information across numerous multiple policyholders. In Section 5.3.1, we discuss the validity of the assumptions in the IWL-CSI that characterise this ability. We then discuss how these results inform the optimal strategy for the insurer to adopt in Section 5.3.2.

5.3.1 Assumptions

IWL-CSI makes strong assumptions in order to gain insight into the ability of the insurer to aggregate claims data. The assumptions break down into: (i) defenders are homogeneous and the cost to the attacker of exploiting the same vulnerability in two different defenders is equal; and (ii) the insurer influence the security posture of its policyholders. These assumptions may only be relevant in particular contexts.

Assumption (i) might be relevant in the context of organisations in the same industry who rely on similar service providers or off-the-shelf software. It is inappropriate if the policyholders have different organisational profiles or deploy different information systems. For example, an attack on an organisation holding financial data may be economically viable, but not economically viable on an organisation holding less valuable data — rendering knowledge about attacks on the other as useless.

Assumption (ii) depends on market dynamics. It is less appropriate in a 'soft market' characterised by an over-supply of insurance. Insurers may not have the freedom to discriminate between applicants based on their security posture, nor be able to mandate security controls, when the insured can easily find coverage elsewhere. The active approach requires less insurer influence on security posture than the diverse approach. The interviews upon which Chapter 4 is based revealed that the cyber insurance market could be described as a 'soft market' in 2017. Although the strategies may be difficult to implement at present, they could be saved until market conditions are right.

Beyond the assumptions we introduced, there is a question about what constitutes the vulnerabilities and controls that protect them. The model was used to explain payment card fraud and online crime in the original IWL [71]. In payment card fraud, banks implementing chip and pin security shifted attackers towards exploiting card-not-present vulnerabilities.

It is an open question whether the controls about which insurers collect information can have the same effect. For example, whether an organisation has an anti-virus system in place is unlikely to shift the attacker's strategy given they

have existed in various forms for thirty years. What is arguably more likely is attackers exploit particular vulnerabilities and shift to new vulnerabilities as these are patched. This suggests insurers might reflect on the security information they collect and how valuable it will be for increasing information.

The next subsection discusses strategies the insurer might adopt to take advantage of this market power.

5.3.2 Insurer strategy

The results suggest that an insurer sharing claims information will pay less in claims, as evidenced by the lower attack intensity in Table 5.4. The ROSI calculations show this increased revenue results from more efficient security investments. Further, sharing claims information reduces the variance of claims meaning that insurers need to hold less capital in reserve.

However, achieving higher revenues relies on security investments being influenced by claims data. Empirical work has shown that insurers do not seem certain about which security controls are important, as two thirds of insurers in one study did not adapt premium prices based on security controls [73]. Our results suggest that this may even be rational — when uncertainty is high, the optimal starting configuration is no defensive investment. Perhaps insurers will begin to require security controls as they observe attacks on their policyholders and discover which defensive measures are effective.

Such a strategy is reliant on dynamically adapting the defensive configuration. Woods and Simpson [22] identified a number of mechanisms by which the insurer can influence security controls of their insureds, which we call *risk selection*, *incentivisation* and *integration*.

Risk selection involves the insurer offering coverage based on the security level of the applicant. These levels may result from risk-aversion and the premium may be dependent on the security level. However, this approach cannot respond to attacks because the insurer cannot offer incentives to invest in security until the insurance policy lapses.

Incentivisation involves the insurer offering premium discounts for adopting certain security controls and these will indirectly lead to a market composed according to the insurer's optimal strategy. Again, security controls must be adopted dynamically but policies lapse after a year. One solution that provides responsiveness is dynamic risk management in which the price of the policy is frequently updated according to the insured's behaviour and external events, such as attacks on other policyholders.

Integration involves the insurer managing the insured's security directly and is the most intrusive mechanism. At present, we are not aware of anyone utilising integration to offer insurance and security[1]. It could be achieved either by insurers partnering with security companies or by security companies offering insurance.

Each mechanism achieves a different level of responsiveness, highlighting the tension between the insurer's desire for adaptive security mechanisms and the insured's distaste for intrusive obligations. In the IWL-CSI the defender can adapt their defensive configuration every round. The corresponding length of time in the real world is not clear; does the defender have to update every year, every month or even every day?

This question depends on the change to defensive configuration. Unlike in the stylised IWL-CSI model, security controls are not homogeneous and some may be more difficult to implement than others. For example, secure software engineering investments may be more efficient in pre-deployment [116] — long before the insurance policy has been purchased. This raises the concern that the insurance industry's incentives do not align with the insured's long term interest — a concern that was borne out in the security controls that cyber insurance application forms focus on [17].

Insurers even discovering which controls are effective is by no means guaranteed. Collecting information about the security posture of the insured presents problems in terms of the granularity of the information collected and the reliability of self-reported answers [73]. An application form may report that a firewall is in place

[1]One of the cyber warranties analysed in Chapter 6 is only valid if the vendor controls certain aspects of the product, such as the white list.

without specifying the monitoring processes, let alone whether it is configured correctly. These issues are exacerbated by difficulties linking controls to attacks suffered in post-event reporting. To our knowledge, no empirical work has established what information is collected in the claims process, nor how this is used by insurers.

Another concern is whether all the policyholders achieve the same return. As the IWL-CSI strategy relies on defenders adopting different postures, they will suffer different attacks and costs of defence. This raises the question of which defender should adopt the most costly defensive configuration or, equally, the one most likely to suffer attack. These concerns are somewhat similar to ethical concerns in medical research where the control group does *not* receive a potentially life saving medicine. In theory this can be handled by the market; the insurer should offer premium discounts to defenders adopting more costly defences. However, it would not be the first market failure in information security if this were not the case [2].

In spite of these concerns, cyber insurance is a desirable product, as evidenced by the growing market. The active and diverse insurer can improve the return on security investments and achieve greater expected revenue. But this does not consider the insurer's operating costs in doing so. It could be that, for low uncertainty, the small improvement in revenues by adopting the active or diverse strategy is not worth the increased operating costs associated with doing so.

Market composition affects the total revenue across all the insureds. Under conditions of high uncertainty, insurers can increase average revenue by taking on an increasing number of policyholders, which raises competition concerns. This is consistent with the accounts stating that large insurers see their claims data as a competitive advantage. Indeed, "insurers with a small amount of claims data were (perhaps understandably) far more enthusiastic about data being shared" [22]. The results suggest that under conditions of high uncertainty, the value of claims data may be such that the marginal benefit of an additional insured drives the market composition towards monopoly. Again, it would not be the first time incentives have tended towards monopoly [2]. However, this analysis does not take into account the risk of aggregated losses which may push the market towards many providers.

In the IWL-CSI, uncertainty and costs of defence are fixed for the entire game. However, in reality these will change over time: organisations will adopt new information systems changing the cost of attack and uncertainty regarding the attacker may rise or fall. The "technical flux of change" undermines the utility of actuarial data over time [117]. There are two competing factors here: (i) changing costs of exploit make old claims data less valuable, and (ii) old claims data may reduce uncertainty about the attacker in future games. The former erodes the competitive advantage of claims data over time while the latter strengthens it. For example, if one insurer holding a larger share of the market leads them to operate with lower uncertainty than another, they can expect to gain higher revenues. There could be a tendency towards monopoly in the market, depending on the balance of (i) and (ii).

5.4 Summary

The IWL-CSI explores the interaction between security investment under uncertainty and the insurer's ability to aggregate claims data. Our main theoretical contribution is providing a more realistic representation of a number of aspects of the cyber insurance market that we observed in Chapter 4: (1) uncertainty about which controls are effective, (2) conceiving of security as a discrete set of variables, and (3) insurers aggregating claims information to increase information.

The simulations show that sharing claims information can increase the average revenue of the policyholders in the IWL-CSI model, particularly when uncertainty is high and initial defensive investment is low. Exploring pricing structures reveals that both the insurer and the insured can benefit from this increased revenue, as well as the reduction in the variability and the rate of claims.

Translating the strategies that are defined in terms of the IWL-CSI into business strategies will be difficult. An insurer might reflect on how claims information could help insureds make more effective information security investments, which may translate into a lower volume of more predictable claims. Further, claims data from policyholders employing diverse defences could be more valuable, even

accounting for the additional cost incurred as a result of maintaining this diversity. In particular, we recommend revising what information is collected in the application process with analytics in mind.

In this chapter we explored how insurers can identify more effective security controls. In Chapter 6, we explore how consumers can use risk transfer products to identify more effective security controls.

6

Cyber Warranties as a Quality Signal

Enterprise security firms have begun to offer so-called *cyber warranties* as an alternative signal of quality. For example, a managed security provider[1] has offered a $100,000 warranty and an end-point protection firm[2] offers a $1,000,000 warranty. In this chapter, we consider cyber warranties to be voluntary obligations in which enterprise security providers offer to indemnify consumers in the event of a successful attack. Security firms are seeking to unilaterally shape market dynamics by accepting these obligations.

This chapter presents the first economic consideration of how cyber warranties affect the market for enterprise information security products. The contents are based on [20] and [21]. Section 6.1 analyses 10 warranties to understand what is typically covered. In Section 6.2, we introduce a decision-theoretic model that captures both the vendor's short-run decision of setting the warranty level while investment is fixed and the long-run decision in which investment can vary. Section 6.3 derives four inferences the consumer can make based on the cyber warranty level. Section 6.4 illustrates how these inferences depend on the information structure between the consumer and the vendor. We discuss how

[1] `https://www.armor.com/cyber-warranty/`

[2] `https://www.theregister.co.uk/2016/07/29/sentinelone_ransomware_guarantee/`

Description	Contract	Amount	Coverage	Notification Window	Modification	Required Procedures	Physical Defects
Network architecture	Y	Product	T	?	Y	Y	
Various products	N	Product	T	?			
Mobile security	Y	Product	T	10 Days			
Routers	Y	Product	T	?	Y	Y	Y
Firewall	Y	Product	T	?	Y	Y	Y
Source code review	Y	$5,000,000	C	48 hours	Y		
Back-up services	Y	$10,000/$50,000	C	30 Days	Y	Y	Y
End-point protection	N	$100,000/$500,000 /$1,000,000	C	?	Y	Y*	
Monitoring	Y	$1,000,000 or 2x License fee	C	45 Days	Y	Y	Y
End-point protection	N	$1,000,000 or $1,000 per device	C^{RWP}	?	Y	Y	

Table 6.1: Columns refer to: a description of the vendor, whether we have a copy of the contract, the amount of coverage offered, the type of coverage, and whether there are exclusions for modifying the product, not following procedures or physical defects.

applicable these inferences are in real-world decisions in Section 6.5. Section 6.6 offers a summary of the chapter.

6.1 What do cyber warranties cover?

This section contains a preliminary analysis of what is covered by cyber warranties. We searched for cyber warranties using a popular search engine with terms using a combination of *warranty*, *cyber*, *information security*, *cybersecurity*, and *security product*. We excluded any physical security products, such as burglar alarms. In seven of the cases we could get the actual contract and in three (the products with an N in the Contract column in Table 6.1) we had to rely on the vendor's description of the product. The analysis consisted of inductively identifying common components across the different warranties.

Consumers should first ask whether the product comes with a traditional or cyber warranty. Of the 10 warranties attached to information security products we collected, half covered only defective hardware or software, denying coverage for first- or third-party costs resulting from an attack. There was little variation in the coverage among these products, which we will call *traditional warranties* from now on and denote by T in Table 6.1.

Cyber warranties represent more meaningful coverage, which we denote by C in Table 6.1. Our analysis identified *coverage* and *trigger conditions* as the

main components of the warranty. *Coverage* describes which costs the vendor will indemnify and the total indemnification limit. *Trigger conditions* describe the conditions that (in)validate the coverage, which includes the security procedures the consumer must follow.

Four of the five cyber warranties in our sample covered first-party costs like notifying customers and hiring consultants for forensic investigation, public relations or legal review. One vendor explicitly covered ransomware payments and nothing else (denoted C^{RWP}), while others explicitly excluded ransomware payments.

None of the warranties (T or C) cover regulatory fines or third-party liability. The amount of coverage ranged from \$10,000 to \$5,000,000 with one vendor offering \$1,000 per end-point and another offering the minimum of twice the subscription fee or \$1,000,000.

Trigger conditions contain most of the variation in warranties. Product modifications and physical defects invalidated both traditional and cyber warranties alike. Common procedures that must be followed in cyber warranties include proper installation, configuration, maintenance, application of updates, and operation. The lack of concrete definitions for what *proper* entails may provide vendors with discretion when it comes to paying out. One vendor adopted a different approach, requiring the client to relinquish write access to the product and forgo control of security functions like the white list, for the warranty to be valid.

Exclusions varied based on the product type. It is hard to evaluate the relative frequency of losses that will and will not be refused because of these exclusions. A back-up provider excluded "any breach due to weak or stolen credentials" or "denial of service". A monitoring product excluded breaches that were "not a result of APT activity". A firm offering source code review excluded coverage if the attack resulted from unknown vulnerabilities, defined elaborately using the CVE database and a list of 122 known vulnerabilities.

This preliminary analysis suggests cyber warranties do not transfer much liability. In terms of *coverage*, they explicitly exclude third party liability and income lost to business interruption, reputation damage or intellectual property loss. The

trigger conditions exclude even more losses via ambiguous definitions of appropriate maintenance and operation of the product, not to mention all of the events that are explicitly excluded.

The next section introduces an economic model to understand how warranties affect the market. This allows us to abstract away from the current state of cyber warranties. We hope our analysis will remain useful as these warranties mature, potentially transferring more liability in the future.

6.2 Model

A one-round decision-theoretic model is used because we are primarily interested in what a consumer can infer about a security product based on the cyber warranties offered. We do not investigate the strategy of the vendor in response to the consumer (which might lead to a game theoretic approach) nor do we consider multiple procurement cycles (which might lead to a multiple round model). Our model takes inspiration from how cyber insurance has been modelled.

The model considers a number of vendors $V_1, .., V_n$, with each V_i selling a single security product S_i. Each vendor sets the amount of development investment z_i that represents costs, including developer time, training costs, participation in threat intelligence schemes and purchasing development tools.

We assume a Bertrand model of competition [118] in which a vendor chooses a price P_i and a warranty $\Psi_i \in [0, 1]$; how this choice interacts with market demand determines the quantity supplied. The Bertrand model is relevant to software markets where quantity supplied can dynamically meet market demand [119] — unlike, for example, car manufacturers who must forecast market demand in order to begin a production process that may take months to complete.

To model the random nature of cyber attacks, we consider a Bernoulli trial in which the consumer faces a set loss λ with probability of occurrence p_i when the consumer purchases product S_i and a probability of v_0 if no purchase is made. This realisation of losses is in line with the common approach to modelling other forms of risk transfer [45, 120, 121, 122]. As the set loss is fixed, the security

Symbol	Description
V_i	The i-th vendor
S_i	The product offered by the i-th vendor
c_{f_i}	The fixed costs incurred in offering product S_i
z_i	The amount of investment into security during development of S_i
P_i	The price of S_i
Ψ_i	The proportion of realised losses the i-th vendor will indemnify
λ	The set loss resulting from a successful attack
v_0	The consumer's vulnerability before employing a security product
$S(v_0, z_i)$	The probability of successful attack given an investment of z_i
$R_c(R_v)$	The revenue of the consumer (vendor)
Π_i	The profit accruing to the i-th vendor

Table 6.2: Descriptions of each parameter in the model.

products mitigate the probability of successful attack without affecting the impact of the attack. Consequently, our analysis will be less relevant to security products that seek to reduce the impact of losses.

There are many functions relating p_i to the investment z_i in the security product S_i. For example, Fultz et al. [120], Pal et al. [122] and Johnson et al. [121] each use a different form. We adopt the functions introduced in [69] given it has empirical support [123] and wide acceptance. There is a more detailed overview in Section 2.3.2. For the purposes of this chapter, we only need to know that they may be expressed in the form

$$S^I(z_i, v_0) = \frac{v}{(\alpha z_i + 1)^\beta} \tag{6.1}$$

$$S^{II}(z_i, v_0) = v^{\alpha z_i + 1} \tag{6.2}$$

The vendor incurs total cost $c_i = c_{f_i} + z_i$, where c_{f_i} represents the fixed cost of offering the product and z_i is the variable describing investment in product development. Each vendor seeks to maximise their profit Π_i by setting P_i, z_i and Ψ_i:

$$\Pi_i = P_i - S(z_i, v_0)(\lambda \cdot \Psi_i) - c_i \tag{6.3}$$

Table 6.2 describes Ψ_i as a proportion of losses that the vendor will indemnify. However, the previous section showed that cyber-warranties actually place a limit

or ceiling on the amount that can be claimed. This translates into a proportion of realised losses because we chose a model in which the loss amount λ is fixed. In a model with variable losses (and in reality), a given warranty would not indemnify a fixed proportion of losses.

The consumer only has knowledge of the price P_i, warranty Ψ_i and set loss λ. The investment z_i is assumed to be unobservable due to information asymmetry. The consumer chooses the security product S_i that minimises

$$R_c = P_i + S(z_i, v_0) \cdot \lambda(1 - \Psi_i \cdot H_i) \tag{6.4}$$

When making supply-side comments, we assume that customers are homogeneous and all demand the same product leading to winner-takes-all market dynamics, which have been observed in many other software markets [2, 119]. We are more concerned by the demand-side and the majority of our comments relate to the consumer's purchasing decision.

6.3 Analysis

We consider a market without security warranties ($\Psi_i = 0$) to illustrate the market for lemons. Using Equation 6.3, the vendor receives

$$P_i - (c_{f_i} + z_i)$$

while the consumer's expected security expenditure is

$$P_i + S(z_i, v_0) \cdot \lambda$$

The vendor has no incentive to increase the development investment beyond $z_i = 0$ because the consumer cannot observe ex-ante the resulting decrease in vulnerability. In a competitive market without warranties, the market equilibrium is $P_i = c_{f_i}$ with $z_i = 0$. Clearly vendors still invest in product development without offering warranties in spite of this result and we discuss why they might do so in Section 6.5. The rest of this section identifies four inferences consumers can make regarding security products, as well as the information they need to do so.

First, we consider a vendor V_i with a fixed investment of $c_{f_i} + z_i'$ in the product. Each vendor can offer the product at a price P_i with warranty Ψ_i. Equation 6.3 shows that the vendor's profit at the price P_i is as follows.

$$\Pi_i(P_i) = P_i - S(z_i', v_0)(\lambda \cdot \Psi_i) - (c_{f_i} + z_i') \tag{6.5}$$

Microeconomists [124] often assume that the vendor can only incur losses up to the value of the fixed costs of operation $(c_{f_i} + z_i')$ in the short-run. This observation leads to the constraint

$$\Pi_i(P_i) \geq -(c_{f_i} + z_i')$$

from which we derive Inference 1.

Inference 1 *Vendor V_i can offer S_i in the short-run with a warranty level of $\Psi_i \in [0, 1]$ at any price*

$$P_i \geq S(z_i', v_0)\lambda \cdot \Psi_i$$

The left-hand side represents the expected value of the indemnification payment to the consumer. It is reasonable to assume that no risk-neutral vendor would offer the warranty unless they receive at least this value as an up-front payment. This provides an upper bound of $\frac{P_i}{\Psi_i}$ for the expected loss a consumer faces — dividing by Ψ_i adjusts for the proportion of the loss that the vendor pays. The inference can be made in the presence of information asymmetry regarding the vendor's security efficiency (α, β), the shape of the probability breach function $S(\cdot, \cdot)$, or their security investment during development z_i.

The consumer seeks to minimise Equation 6.4 despite having incomplete information about z_i. The consumer can use Inference 1 to calculate a lower bound z_{min_i}, which represents the smallest investment value such that vendor i can break even in offering offer a product with warranty Ψ_i at price P_i. This value may be used to calculate the worst-case expected loss $R_{c_{min}}$ resulting from purchasing the product S_i:

$$R_{c_{min}}(S_i) = P_i + S(z_{min_i}, v_0) \cdot \lambda(1 - \Psi_i) \tag{6.6}$$

The consumer is assumed to be indifferent between purchasing the product S_i and the product S_j if

$$R_{c_{min}}(S_i) = R_{c_{min}}(S_j) \tag{6.7}$$

From this, we can construct a (worst-case) indifference curve for the consumer.

Calculating the (worst-case) indifference curve involves finding the smallest z_i such that

$$\Pi(S_i) \geq -(c_{f_i} + z_i') \tag{6.8}$$

Using Equation 6.3 and the formulae for each class of probability breach function, we derive Inference 2.

Inference 2 *If the product S_i has been offered at price P_i and the warranty level is Ψ_i we have that*

$$z_{min_i} = \begin{cases} \frac{(\frac{\lambda\Psi' v_0}{P_i})^{\frac{1}{\beta}} - 1}{\alpha} & \text{if } S(\cdot,\cdot) \text{ is Class I} \\ \frac{ln(P_i) - ln(\Psi\lambda v_0)}{\alpha ln(v_0)} & \text{if } S(\cdot,\cdot) \text{ is Class II} \end{cases}$$

It is worth noting that Inference 2 may provide an under-estimate of the product investment. A profit-making vendor analysed as if the vendor was breaking even would appear to have invested less than they did in actuality.

The long-run decision reduces to first selecting a warranty level and then determining the optimal investment as setting the investment first would reduce to the short-run analysis.

Suppose that the vendor unilaterally sets the warranty level at $\Psi_i' > 0$. The vendor will make the long-run investment z_i^* that optimises profit $\Pi_i(P_i)$ for all values of P_i. The marginal net benefit of investment is given by

$$\frac{\partial \Pi_i}{\partial z_i} = -\frac{\delta S}{\delta z}(z_i, v_0)(\lambda \cdot \Psi') - 1 \tag{6.9}$$

Using the convention that investment is non-negative ($0 \leq z_i$), we can derive the following.

Inference 3 *If the vendor has committed to the warranty level Ψ_i', the optimal choice of product investment is*

$$z_i^* = \begin{cases} \frac{(\alpha\beta\lambda\Psi' v_0)^{\frac{1}{\beta+1}-1}}{\alpha} & \text{if } S(\cdot,\cdot) \text{ is Class I and } \alpha\beta\lambda\Psi' v_0 > 1 \\ \frac{-ln(-\alpha\lambda\Psi' v ln(v))}{\alpha ln(v)} & \text{if } S(\cdot,\cdot) \text{ is Class II and } \alpha\lambda\Psi' v ln(v) > -1 \\ 0 & \text{otherwise} \end{cases}$$

Inference 3 allows the consumer to infer the exact level of investment providing the warranty level was decided in the long-run and investment was optimised for this decision. If the investment z_1^* can be inferred, the consumer can expect revenue $R_c(S_i)$ if they purchase S_i, where

$$R_c(S_i) = P_1 + S(z_i^*, v_0)(1 - \Psi_i)\lambda \tag{6.10}$$

In a fully competitive market, we can expect that

$$P_i = S(z_i^*, v_0)\Psi'\lambda + c_{f_i} + z_i^* \tag{6.11}$$

However, Inference 2 and Inference 3 both rely on the consumer knowing the shape of the probability breach function and the vendor's security productivity.

In both the short-run and the long-run, the vendor can vary the warranty level. This is analogous to the vendor changing the indemnification limit observed in Section 6.1. In this sense, Ψ_i is a variable and so we can differentiate with respect to it. Again, the real world analogy sees the consumer observe the change in the price of the product for a given change in warranty level. If profits are constant across warranty levels, the price P_i must increase to compensate for any increase in the warranty Ψ_i at a rate equal to the risk-transfer rate of substitution (RTRS) $\frac{\partial \Pi_i}{\partial \Psi_i}$ in order to keep profits constant.

Inference 4 *The risk-transfer rate of substitution for the vendor V_i is equal to the consumer's expected loss when the security product S_i is in place.*

$$\frac{\partial \Pi_i}{\partial \Psi_i} = S(z_i', v_0)\lambda$$

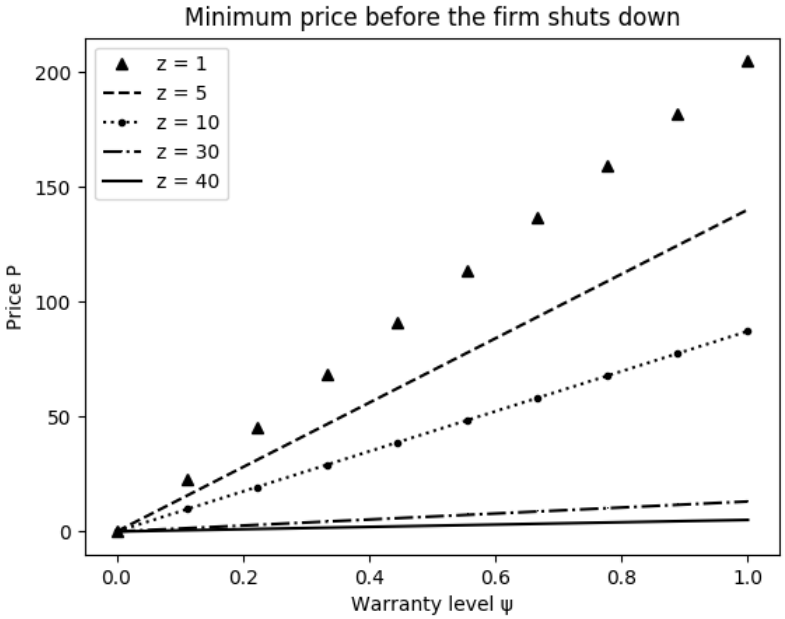

Figure 6.1: The price at which the vendor would shut down if price fell any further, for different investment levels z and a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.5$ and $c_f = 5$.

The consumer can discover the expected loss if the risk-transfer rate of substitution is observed. This inference can be made with knowledge of only the price and warranty level regardless of whether the warranty has been offered in the short-run or the long-run. This inference might be considered the most powerful as it can be made with information asymmetry regarding the vendor's technological constraints.

The price and warranty offered by each vendor will depend on the market environment in both the short-run and the long-run. If the vendors have perfect information about the competitors' investments in product development, there may exist one vendor who can extract a supplier surplus by setting (P_i, Ψ_i) such that any competitor would suffer an economic loss in offering a competing product. However, this will depend on the particular values of both investments z_i, existing vulnerability v_0 and breach probability function $S(z_i, v_0)$. The relative risk aversion of the vendor and the consumer will determine the optimal pair (P_i, Ψ_i).

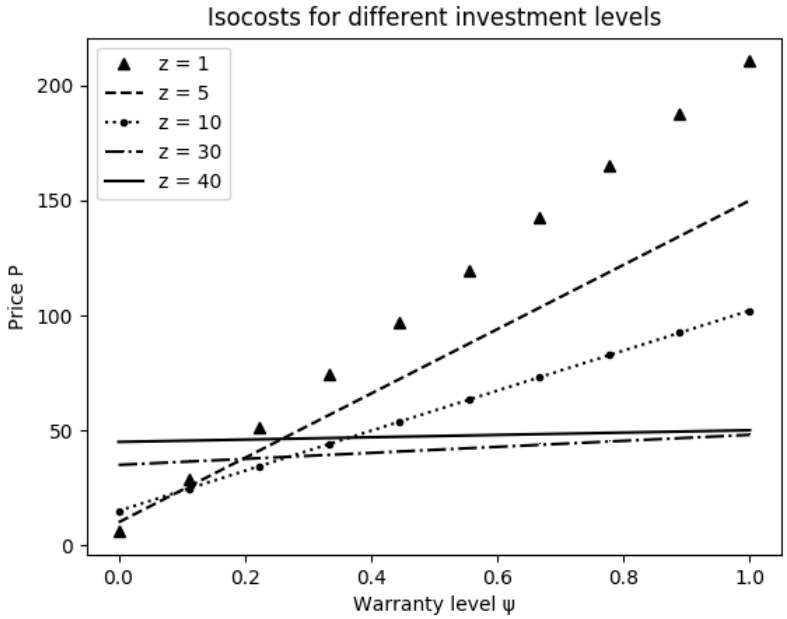

Figure 6.2: The price at which the vendor would shut down if price fell any further, for different levels of security investment z and a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.5$ and $c_f = 5$.

6.4 Numerical Illustrations

This section illustrates each of the inferences in turn. Firms will only shut down in the short-run if price exceeds average cost. Figure 6.1 shows how the minimum price is determined by the warranty level and the investment. We define the shutdown-isoprofits to be the lines with a loss equal to fixed costs; the curve for $z = z'_j$ represents the possible pairs (P_j, Ψ_j) for which the j-th vendor's profit $(\Pi_j(S_j))$ is equal to $c_{f_j} + z'_j$. Although vendors V_1 and V_5 have invested $z_1 = 1$ and $z_5 = 40$ respectively, both accept a minimum price of 0 when no warranty is offered. The difference between the size of their losses will be given by

$$\Pi_5(0) - \Pi_1(0) = (c_{f_5} + z_5) - (c_{f_1} + z_1) = -39$$

because V_5 has larger fixed costs as a result of higher fixed investment z_5.

Figure 6.2 illustrates the isoprofits when the firms break even. For a given warranty level Ψ, the vendor with the isoprofit curve intersecting $x = \Psi$ at the lowest point can offer the most competitive product. This provides a graphical

Figure 6.3: The minimum investment value z_{min_i} and worst-case loss $R_{c_{min}}$ for a given price P_i and warranty level Ψ_i, for a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.9$ and $c_f = 5$.

illustration of the market for lemons as the product with investment $z = 1$ is most competitive when no warranty is offered. The downside of over-investment can be seen by considering that the vendor who invested $z = 30$ is more competitive at every warranty level than the vendor who invested $z = 40$.

If the consumer had knowledge about the shape of the probability breach function and the vendor's security efficiency, Inference 2 can provide information about the vendor's minimum investment z_{min_i} in the short-run. Figure 6.3 highlights the points at which the strongest inference can be made; more information is contained in a warranty as the price it is offered at decreases. Consider a duopoly with vendors V_1 and V_2 who have made investments of $z_1 = 5$ and $z_2 = 10$ respectively. If vendor V_2 sets (P_2, Ψ_2) to be equal to any pair of Figure 6.3 with $z_m in > 5$, then V_1 would sooner shut-down than offer the same contract. Offering such a contract functions as a reliable signal of product quality in this scenario.

For each contract (P_i, Ψ_i), Inference 2 may also be understood graphically as the smallest value of z_i for which the associated isoprofit curve intersects (P_i, Ψ_i) or falls beneath it. For z_{min_i} to be the worst case, we have to assume that the isoprofit corresponded to the points where the loss is equal to the fixed costs. Inference 2 might lead to different conclusions if we used the isoprofit curves corresponding to a different profit condition, such as breaking even as in Figure 6.2. If a functional

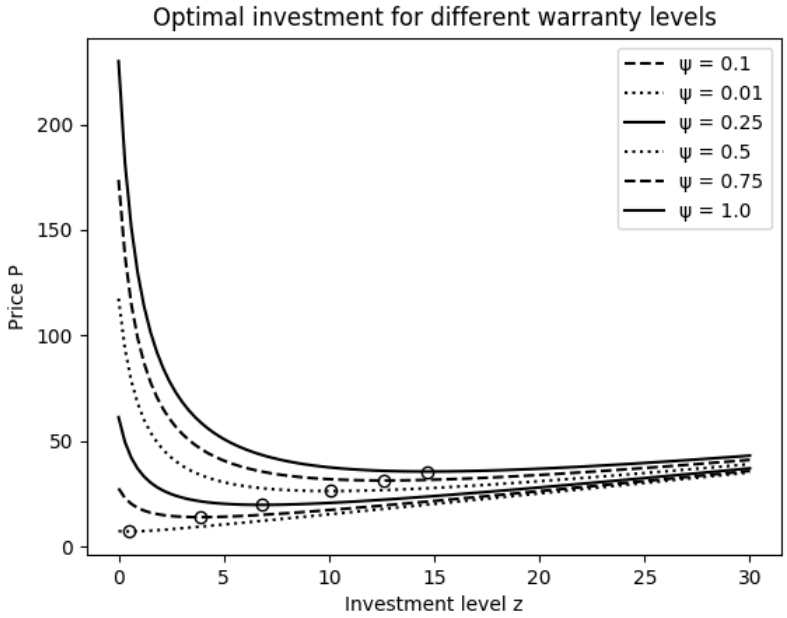

Figure 6.4: The choices of price and investment level that lead a vendor to make zero profit, for different warranty levels Ψ, for a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.9$ and $c_f = 5$.

form is difficult to obtain for a given profit condition, the graphical interpretation of Inference 2 may be used instead.

Turning to long-run investments, Figure 6.4 shows the price a vendor must charge for a given warranty level in order to break even. We have circled the optimal investment for each warranty level and can see that it is increasing in Ψ_i. Consumers may use Inference 3 to discover the optimal investment level z^* and use it to calculate their expected loss. When investment costs are fixed, the consumer can only infer a lower bound for investment whereas the consumer can now infer the optimal investment level for a given warranty level.

Figure 6.5 describes the expected loss (R_c) for each customer if they purchase a product with warranty Ψ_i assuming the optimal investment has occurred. As the curve is always downward-sloping we must have

$$\frac{\partial R_c}{\partial \Psi_i} < 0 \text{ for all } \Psi_i \in [0, 1]$$

However, the customer's expected loss falls at a diminishing rate so that

$$\frac{\partial^2 R_c}{\partial \Psi_i^2} > 0 \text{ for all } \Psi_i \in [0,1]$$

These results, derived via Inference 3, suggest that greater risk transfer to the agent deciding amount of security investment leads to a more efficient allocation of resources. As such, consumers should push to increase the warranty level over time, which can be seen in the decreasing expected loss for greater warranty levels in Figure 6.5. Inference 3 requires knowledge about the vendor's technological constraints, much like Inference 2.

Inference 4 states that the risk-transfer rate of substitution (RTRS) of the i-th vendor is equal to the expected loss when employing the i-th security product. The RTRS for a given vendor is equal to the slope of that vendor's isocost curve (in Figure 6.2 and Figure 6.1). The lowest investment has the steepest isoprofit curve and hence the highest expected loss. Negotiating with the vendor might reveal the RTRS if the vendor stated how much the price would have to rise for a given increase in warranty level.

In summary, Inference 1 provides a lower-bound on expected losses and Inference 2 provides a lower bound on product investment. Both of these are valid in the short-run. Inference 3 provides an exact value of the optimal investment for a given warranty but it is only valid in the long-run. However, Inference 2 and Inference 3 require knowledge about the vendor's technological constraints. Inference 4 provides an exact value of the expected loss. It is valid in both the long-run and the short-run, and requires no knowledge beyond the RTRS. The next section discusses some of the real world limitations of these inferences.

6.5 Discussion

This section first discusses the modelling results in Section 6.5.1 as they are the primary contribution of this chapter. We then turn to the current state of cyber warranties in Section 6.5.2.

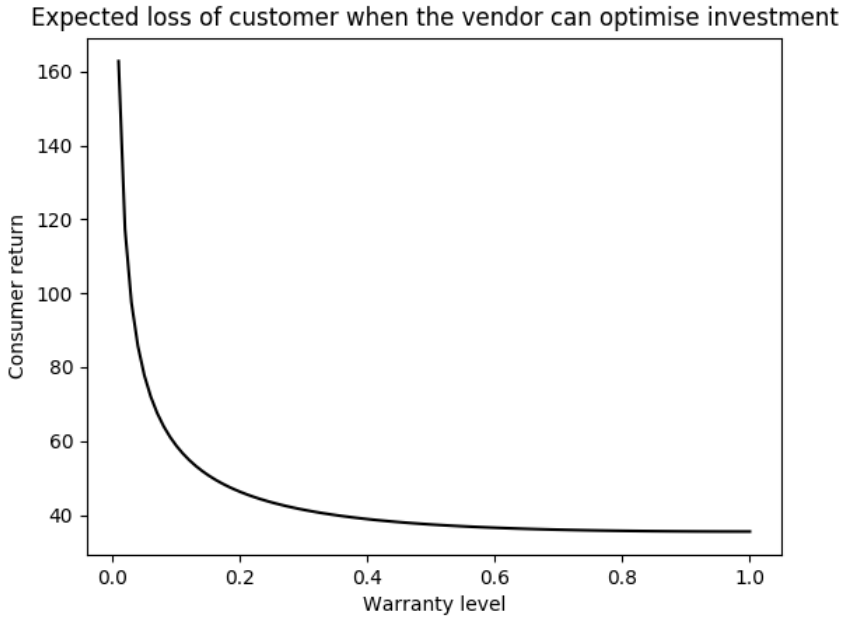

Figure 6.5: The choices of price and investment level that lead a vendor to make zero profit, for different warranty levels Ψ, for a Class I probability breach function with: $\alpha = 0.9, \beta = 1, \lambda = 500, v_0 = 0.9$ and $c_f = 5$.

6.5.1 Modelling results

The consumer can use inferences 1–4 to estimate the expected loss when implementing the security product S_i, which can be compared against the expected loss without any security product or some other security product S_j. This allows the consumer to estimate the expected benefit from the security product. More knowledge about the vendor or the risk transfer rate of substitution may allow the consumer to make stronger inferences and increase confidence in these estimates. Unfortunately, these estimates will be weakened by many complicating factors in the real world.

Inference 1 and 2 are only as reliable as the underlying assumption that firms will not accept losing more than their fixed costs. It breaks down when firms are willing to incur significant losses, which can occur when they seek market dominance [125]. Such strategies can be underwritten by venture capital funding, which often prioritises long-term scale over short-term profit. Whether this assumption holds depends on the vendor under consideration. For example, older firms are less likely to be incurring large losses to achieve scale.

The warranty level will likely take the form of a contract that will not stipulate a proportion of risk the vendor will cover. The contract might instead define a selection of events for which the warranty is valid. Estimating the proportion of the expected loss that these events represent requires that risk managers understand their organisation's risk profile. It may even be difficult to compare two warranties and decide which is more meaningful.

Although the model suggests that full risk transfer achieves the optimal solution for the consumer, it may not be possible in the real world. Cyber insurance policies do not cover intangible losses such as reputation damage and intellectual property loss precisely because it is difficult to quantify such losses. There is no reason why vendors are better suited than insurers to offer warranties covering these risks.

A further complicating factor is that prices must reflect principal–agent problems such as adverse selection and moral hazard. These problems have presented a major problem for cyber insurance, as we observed in Chapter 2. Solutions to these problems, such as monitoring the consumer's security practices to prevent moral hazard or performing an in-depth assessment to prevent adverse selection, come at a cost. Chapter 4 revealed how policyholders pushed back against security obligations and intrusive monitoring. The same would likely be true with cyber warranties.

The risk-transfer rate of substitution is only equal to the expected loss if the vendor is risk neutral. Otherwise the consumer would have to correct for the vendor's discomfort with holding greater liability associated with a higher warranty level. Vendors should also be concerned by the possibility of a 'cyber hurricane' in which interdependent events trigger multiple indemnification claims [52, 126].

Furthermore, the insolvency risk to vendors grows as they hold more liability. The risk may be managed via self-insurance or market insurance to ensure that vendors have funds available for indemnification. An equilibrium between consumer demand and the cost of these insolvency prevention measures will determine the warranty level available to the consumer at a given price.

6.5.2 The current state of cyber warranties

Warranties must transfer non-negligible amounts of liability to vendors in order to meaningfully overcome the market for lemons. Consumers should question whether warranties can function as a costly signal when narrow coverage means vendors accept little risk. Cynics characterising warranties as marketing tricks cannot be dismissed based on our analysis.

Buyers cannot compare across warranty contracts when so many events are excluded and the *trigger conditions* are so ambiguous. Moving towards standardised terms and conditions may help consumers, as has been pursued in cyber insurance [22]. Alternatively, the managed security provider's warranty conditioned on the vendor taking full control of security decisions might be more meaningful than a warranty conditioned on ambiguous definitions of appropriate operation.

Limited warranties at least communicate the attached product's limitations. For example, the source code review firm cannot over-exaggerate their product when they only indemnify losses resulting from known vulnerabilities with a corresponding CVE number. Consumers might look to the vendors with the fewest exclusions as warranties force the emperor to admit he is wearing no clothes.

6.6 Summary

Customers face information asymmetry when deciding which information security product to purchase. The results discussed in this chapter suggest cyber warranties can overcome this information asymmetry by creating a separating equilibrium in which the vendor reveals the level of product investment. Vendors selling information security products face lower costs in offering cyber warranties if they invest in developing more effective products.

Our model identifies four inferences that customers can make about a potential information security purchase based on the warranty offered. In general, more information is gained when there is more prior knowledge about the vendor. However, these inferences are likely to be weaker in the real world. Consumers must adjust

for factors including the extent of the vendor's risk aversion, costs incurred to mitigate principal–agent problems, and risk-loading to deal with the variability of (potentially correlated) losses.

Chapter 5 asked how insurers can use claims data to identify effective controls; this chapter asked how consumers might use cyber warranties to identify effective security products. We have not yet asked how much to invest. Inference 4 suggests consumers could use warranties to infer loss estimates. Operationalising this is impossible due the novelty of cyber warranties. However, Chapter 4 showed that the relative maturity of the cyber insurance market, along with some enabling regulation, means we can collect cyber insurance prices. The next chapter introduces a technique to make inferences about loss estimates based on insurance prices.

7

Inferring Loss Distributions from Insurance Prices

Having identified how risk transfer products might help identify effective controls, the question of how much to invest has not been considered. A first step in doing so is estimating the potential losses an organisation might face. This chapter introduces a method to infer loss distributions from insurance prices. This operationalises Inference 4 from Chapter 6 using real prices extracted in Chapter 4.

Chapter 2 (in particular Section 2.4.2) identified weaknesses in existing approaches to quantifying cyber risk including: (1) estimates not in terms of financial losses, (2) the granularity—sample size trade-off, (3) reporting biases, and (4) limited number of incident types. Weakness (1) is driven by measuring impact in terms of number of records [81] or business interruption duration [72]. Weakness (2) results from asking how relevant the losses suffered by a $300 billion energy firm are to an independent book store, as both are considered in aggregate data breach studies [80]. Weakness (3) relates to reporting biases described in [97] and (4) reflects the fact that there is little empirical data regarding incidents like ransomware and business email compromise.

A novel data source for cyber loss distributions could be valuable to complement existing information sources, especially given their limitations. Hayek identified the

price system as a mechanism to communicate information about "how to secure the best use of resources" [127], even when the "knowledge of the relevant facts is dispersed among many people". This link between prices and dispersed information underlies faith in the efficient markets hypothesis stating that asset prices reflect all available information; although, it has recently come under attack [128].

Prices can be operationalised for decision support. For example, prediction markets have been created to predict election outcomes [129], influenza outbreaks [130], scientific results [131], Oscar nominees and winners [132], and government policy outcomes [133]. In the security context, it has been suggested that exploit derivatives can be used to estimate the probability a software product will be compromised [134]. However, it is important to remember "markets incorporate falsehood as well as truth" [132].

This chapter introduces a method to infer loss distributions from cyber insurance prices, much like how probabilities are extracted from prediction markets. If the efficient markets hypothesis holds for insurance markets, the price of insurance for a single retail firm with a revenue of \$50M reflects all available information about retail firms with a revenue of \$50M. Admittedly, factors like operational costs, transaction costs, and solvency risk add noise to the inferences.

This method complements existing approaches to quantifying cyber losses in a number of ways. Insurance prices are provided based on firm characteristics like revenue and industry, which leads to granular insights. Insurance is offered for many cyber incidents that currently lack related data like business interruption, business email compromise, and ransomware. This method provides parameterised distributions of dollar losses, whereas previous studies report point-estimates of dollar losses or distributions of the number of records breached.

We introduce the theoretical framework behind our inferences in Section 7.1. Section 7.2 applies the inference framework to evaluate how well these parameterised distributions explain observed data. We discuss limitations and compare our inferences to other empirical work related to loss events in Section 7.3. Finally, we offer conclusions in Section 7.4.

7.1 Method for Inferring Loss Distributions

This section describes a method to make inferences from the pricing data introduced in Chapter 4. In particular, we extracted 26 pricing algorithms (as filed to state regulators) used to price cyber insurance. This provided quantitative insights into how prices are adjusted for coverage type, limit and deductible, revenue and industry. This allows us to generate sets of prices for each pricing algorithm. If we keep the firm-specific characteristics constant and only adjust the amount of coverage (change the limit and deductible), then we observe how the price of cyber insurance changes as the amount of coverage is varied.

All of our inferences will be based on this relationship between the price and amount of cyber insurance coverage. In keeping the firm-specific and incident type constant, the inferred distribution will be relevant to a specific firm and incident type. We are not aware of any existing method to infer loss distributions from insurance prices and so we had to create one[1].

Figure 7.1 describes our iterative method for inferring loss distributions. The method begins by choosing parameters for a given loss distribution. The parameterised loss distribution is used to generate the predicted price for an insurance policy with a given limit and deductible. The set of predicted prices is compared to the observed prices. The result of this comparison is used to improve the next parameter choice. The cycle repeats until some termination condition is met.

We describe how we generate predicted prices from a parameterised loss distribution in Section 7.1.1. A metric to compare predicted prices with observed prices is defined in Section 7.1.2. The heuristic governing parameter choices and the termination condition is described in Section 7.1.3. Finally, Section 7.1.4 explains how to translate between loss distributions for different coverage types and revenues.

7.1.1 Generating Price Predictions

This subsection assumes we already have a parameterised distribution for all loss events covered by a given insurance policy. This gives rise to the universe of possible

[1] We became aware of [135] after the dissertation had been written and the exam took place.

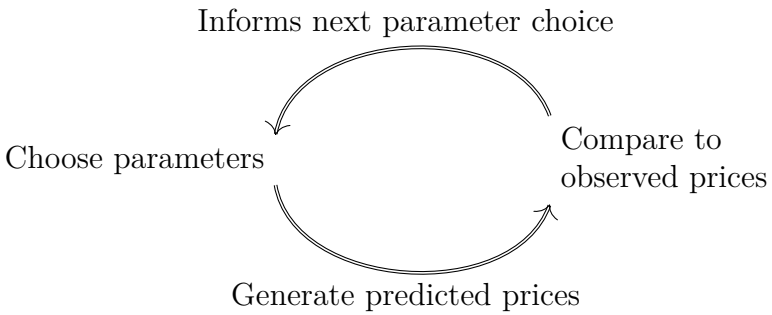

Figure 7.1: High level description of our iterative method for inferring loss distributions from insurance prices.

losses $\Omega = \mathbb{R}_+$, which is distributed according to a random variable X. Each possible loss value x is some non-negative real number occurring with probability determined by the probability density function

$$f(x) : \mathbb{R}_+ \to \mathbb{R}_+$$

The probability that the loss amount is greater than a and less than b is equal to

$$P(a < X < b) = \int_a^b f(x)dx \tag{7.1}$$

The first axiom of probability follows provided f is non-negative. The second follows by normalising f so that

$$P(\Omega) = \int_0^\infty f(x)dx = 1 \tag{7.2}$$

The third follows from the definition of an integral.

For our purposes, an insurance contract is a promise that the insured will be indemnified for any losses greater than the deductible D up to a limit L. This gives rise to an indemnity function $I_{D,L}(x) : \mathbb{R}_+ \to \mathbb{R}_+$ with

$$I_{D,L}(x) = \begin{cases} 0, & \text{for } x < D \\ x - D, & \text{for } D \leq x \leq L + D \\ L, & \text{for } L + D < x \end{cases} \tag{7.3}$$

which describes the size of the indemnity for a given loss x.

As a result, the insurer's expected loss is equal to

$$\begin{aligned} E(X)_{D,L} &= \int_{-\infty}^{\infty} I_{D,L}(x) f(x) dx \\ &= \int_D^{L+D} (x - D) f(x) dx + \int_{L+D}^{\infty} L f(x) \end{aligned} \tag{7.4}$$

Our method assumes insurance prices are determined by the value of $E(X)_{D,L}$. Algebraically we assume that a premium $p_{D,L}$ with limit L and deductible D is equal to

$$p_{D,L} = \Lambda E(X)_{D,L} \tag{7.5}$$

where $\Lambda \geq 1$ is the loading factor. The loading factor determines how much the insurer sets aside for non-claims related expenses. An actuarially fair policy occurs if $\Lambda = 1$.

7.1.2 Evaluating Loss Distributions via Predicted Prices

This subsection assumes the loss distribution has generated a set of predicted prices. We want to evaluate how the predicted prices compare to observed prices. We define the following loss functions.

Definition 1 (Observed Price Cost Functions) *For a set of premium–limit–deductible triples $S = \{(p_{D_1,L_1}, L_1, D_1), ..., (p_{D_n,L_n}, L_n, D_n)\}$, denote the predicted price and observed price respectively as*

$$P_i = E(X)_{D_i,L_i} \text{ and } O_i = p_{D_i,L_i}$$

Then define the following cost functions to evaluate a parameterised loss distribution X with an associated probability density function $f(x) : \mathbb{R} \to \mathbb{R}_+$ over the set of prices S:

$$C_T(f|S) = \frac{1}{n}\sum_{i=1}^{n}\frac{P_i - O_i}{O_i}$$

$$C_A(f|S) = \frac{1}{n}\sum_{i=1}^{n}\frac{|P_i - O_i|}{O_i}$$

The *absolute* function C_A sums the absolute differences between predicted and observed prices. A value of 0 occurs if and only if expected prices perfectly predict observed prices. Meanwhile, the *total* function C_T allows under-predictions to cancel out over-predictions in the aggregate. A score of zero can be achieved by balancing over- and under-predictions.

The parameterised distribution X with the lowest cost function value *best explains* the observed prices $(p_{D_1,L_1}, L_1, D_1), ..., (p_{D_n,L_n}, L_n, D_n)$.

7.1.3 Choosing Parameters and Termination

Our method can be considered as an iterative optimisation problem. For a distribution with n parameters, this involves searching through $\mathbb{R}^n$ for the n-tuple of parameters minimising C_A. The problem is non-convex with no way of distinguishing local minima from the globally optimal solution. Further, the objective function f is costly because it consists of $2m$ integrals where m is the number of price-limit-deductible triples offered by the insurer. These integrals will be computed numerically when an analytic solution is not easily obtained (as will often be the case).

Many optimisation techniques could have been applied and may still prove fruitful in future work. The particle swarm optimisation (PSO) heuristic [136] consists of multiple candidate solutions moving around the search space. The direction and speed of movement of each candidate is determined by a weighting of each candidate's best known position and the best known position from the entire swarm. The search stops when the candidates converge or after a certain number of iterations. The process is repeated multiple times to decrease the chance of selecting an unsatisfactory solution.

PSO was chosen because it does not require calculating the gradient of the objective function, unlike gradient descent. It requires few assumptions about the search space, which is ideal for exploratory research. Further, multiple particles in the search space reduce the chance of terminating at a local minima. The Pyswarm [137] package was used for implementation.

7.1.4 Accounting for Multiplicative Pricing

This subsection provides a partial answer as to the effect of changing the coverage type, a characteristic about the firm, or even the profit loading factor Λ. It could save re-running the optimisation analysis. We first provide the intuition behind the adjustment.

Suppose that we found that $f(x)$ best explains a set of prices $S = \{p_1, ..., p_n\}$. Then let $aS = \{ap_1, ..., ap_n\}$ be a set of prices for a new coverage type with $0 < a$,

then intuitively we might say that all losses are a times as likely. The following reasoning suggests $af(x)$ explains aS as well as $f(x)$ explains S.

$$C_A(af|aS) = \frac{1}{n}\sum_{i=1}^{n}\frac{|P_i - ap_i|}{ap_i} = \frac{1}{n}\sum_{i=1}^{n}\frac{|E(af)_{D_i,L_i} - ap_i|}{ap_i} \tag{7.6}$$

Then

$$\begin{aligned} E(af)_{D,L} &= \int_D^{L+D}(x-D)af(x)dx + \int_{L+D}^{\infty} Laf(x) \\ &= aE(f)_{D,L} \end{aligned} \tag{7.7}$$

substituted into (7.6) reveals

$$C_A(af|aS) = \frac{1}{n}\sum_{i=1}^{n}\frac{a|E(f)_{D,L} - p_i|}{ap_i} = C_A(f|S) \tag{7.8}$$

For example, if the price of business interruption is 36% of cyber liability across all limits and deductibles, then denoting the cyber liability prices as S_{CL} we have

$$C_A(0.36f(x)|0.36S_{CL}) = C_A(f(x)|S_{CL})$$

where $C_A(f|S_{CL})$ is the cost function value of $f(x)$, the cyber liability loss distribution.

Unfortunately, $af(x)$ is not a probability density function since

$$\int_0^{\infty} af(x)dx = a\int_0^{\infty} f(x)dx = a \neq 1$$

Providing $a < 1$, we can adjust $af(x)$ so that a loss of 0 occurs with probability $(1-a)$. Constructing "mixed discrete/continuous distributions" has precedent in actuarial science [138]. This translates into a business interruption loss of 0 occurring with frequency 0.64. If $a > 1$, re-running the method on the new prices might be a better option.

7.2 Empirical Results

Our method begins by selecting a loss distribution with unknown parameters to be inferred. The polynomial distribution was chosen for analytic tractability. The Lognormal, Pareto, Burr and Gamma distributions were selected because they

are commonly used by insurers [139]. Also, they were identified in Chapter 2. The Weibull distribution was included because it is mentioned in one of the documents analysed in Chapter 4.

We generated a set of $6,828$ price-limit-deductible triples for cyber liability insurance for a hypothetical retail firm with a revenue of \$50M. We selected a multiplicative factor of 1 for all factors other than limit, deductible, industry, revenue and coverage type. Even if we assumed certain security controls were in place, it would be difficult to assign the corresponding multiplicative factors because the value depends on the underwriter's subjective judgement.

Section 7.2.1 contains analysis of a set of prices from one insurer, enabling a more detailed description of the cost functions introduced previously. In Section 7.2.2, we extend our analysis to consider all of the data collected in Chapter 4. The *County Fair Cyber Loss Distribution* is derived in Section 7.2.3.

7.2.1 Analysis of One Insurer

Before considering the prices from all 26 rate schedules, we focus on the insurer with the most extensive pricing set. This pricing schedule contained $2,211$ prices for cyber liability coverage, up to a limit of $\$50,000,000$.

Table 7.1 identifies parameter values that minimise C_A across all of the observed prices. The Gamma distribution does the best job of predicting these prices. The average error (C_A) is less than 20% of the observed premium and the net error (C_T) is less than 5%. We will achieve as low as 5% absolute error for some insurers in the next subsection.

The relationship between C_T and C_A reveals the balance of over- and under-predictions. The Lognormal is the worst offender for consistently under-predicting, as C_T is close in value to C_A. This is likely to occur when the distribution is not sufficiently heavy tailed to predict the prices. Indeed, Burnecki et al. [139] suggest the Pareto is more appropriate than Lognormal "where exceptionally large claims may occur" and the Burr distribution provides a more flexible heavy tailed distribution.

Distribution	$f(x)$	**Parameter Values**	C_T	C_A
Polynomial	$-\frac{1+a}{c^{a+1}}x^a$	a=-1.369 c= 1.0	-0.261	0.371
Lognormal	$\frac{1}{\sigma\sqrt{2\pi}}e^{-(log(x)-\mu)^2/2\sigma^2}$	μ = 1e-10, σ =5.418	-0.312	0.39
Pareto	$\frac{\alpha {x_m}^\alpha}{x^{\alpha+1}}$	$\alpha = 0.3692$, $x_m = 0.6712$	-0.266	0.37
Burr	$ck\frac{x^{c-1}}{(1+x^c)^{k+1}}$	c=1.219, k=0.303	-0.264	0.371
Gamma	$\frac{\beta^\alpha}{\Gamma(\alpha)}x^{\alpha-1}e^{-\beta x}$	$\alpha = 0.001579$, β = 5.02e-08	-0.058	0.185
Weibull	$\frac{k}{\lambda}(\frac{x}{\lambda})^{k-1}e^{-(\frac{x}{\lambda})^k}$	$k = 0.3039$, $\lambda = 0.0001574$	-0.263	0.41

Table 7.1: The distributions and parameters values minimising C_A across a set of $2,211$ cyber liability prices offered by one insurer.

Both loss functions have an unbounded punishment for over-predictions but a bounded punishment for under-predictions (unless negative prices are predicted). This shifts all of the distributions towards under-prediction. We experimented with a loss function bounding the punishment for over-prediction at +1. For the best performing distributions, there was little difference in the optimal parameter values minimising C_A and the bounded loss function, and so we omitted this analysis. But it's important to note the parameter choices consistently tilt towards under-prediction.

Analysing the inferred loss distribution from just one insurer provides little more than a candidate distribution used in that specific insurer's pricing algorithm. Inferring loss distributions across all 26 pricing algorithms extracted in Chapter 4 is more insightful and is the task of the next subsection.

7.2.2 Market Analysis

We ran the same analysis across all $6,828$ of the insurance prices and show the results in Table 7.2, with each set of prices corresponding to a row. The number of prices (n) in a given set is determined by the number of limits and deductibles choices offered by the insurer. The maximum limit (L_{max}) ranges from \$1M to one thousand times that.

Set	Date	n	L_{max}	Poly	Lognrm	Pareto	Burr	Gamma	Weibull
1	3/2/17	39	2,000,000	0.322	0.31	0.188	0.324	**0.174**	0.334
2	16/09/16	55	5,000,000	**0.1**	0.101	0.276	0.287	0.183	0.385
3	28/11/11	56	5,000,000	0.112	0.102	0.22	0.22	**0.082**	0.282
4	15/12/16	396	15,000,000	0.174	**0.172**	0.313	0.317	0.544	0.465
5	17/08/12	74	5,000,000	0.194	0.182	0.171	0.168	**0.156**	0.182
6	22/05/14	458	25,000,000	0.17	**0.163**	0.333	0.342	0.277	0.516
7	29/07/09	168	25,000,000	0.327	0.327	0.389	0.391	**0.21**	0.384
8	1/4/16	90	25,000,000	0.19	0.237	0.19	0.19	**0.156**	0.305
9	10/02/09	132	1,000,0000	0.064	**0.055**	0.213	0.213	0.125	0.326
10	05/03/18	374	10,000,000	0.286	**0.269**	0.398	0.417	0.324	0.507
11	01/12/18	120	5,000,000	0.713	0.766	0.293	0.689	0.568	**0.179**
12	03/09/11	234	20,000,000	**0.162**	0.167	0.439	0.446	0.229	0.531
13	22/05/10	135	15,000,000	0.178	0.203	0.178	0.178	**0.113**	0.242
14	12/2/15	55	10,000,000	0.149	**0.147**	0.203	0.203	0.153	0.314
15	22/02/18	325	10,000,000	0.305	0.336	0.26	0.303	**0.208**	0.487
16	18/12/08	290	25,000,000	0.358	0.349	0.387	0.393	**0.235**	0.418
17	28/09/10	2,211	50,000,000	0.371	0.39	0.37	0.371	**0.185**	0.41
18	19/12/14	54	2,000,000	0.46	0.45	0.336	0.46	**0.172**	0.35
19	18/08/17	168	15,000,000	0.359	0.367	0.341	0.359	**0.281**	0.375
20	30/10/14	435	50,000,000	0.306	0.28	0.354	0.358	**0.227**	0.502
21	13/08/12	81	5,000,000	0.35	0.45	**0.15**	0.339	0.243	0.423
22	22/05/08	30	1,000,000	0.369	0.341	0.235	0.292	**0.175**	0.229
23	02/14/17	98	10,000,000	0.213	0.21	0.3	0.314	**0.194**	0.378
24	23/04/18	490	25,000,000	0.251	0.25	0.258	0.267	**0.213**	0.367
25	13/11/13	30	1,000,000	0.369	0.362	0.217	0.312	**0.128**	0.189
26	09/02/18	230	1,000,000,000	**0.267**	0.326	0.436	0.455	0.45	0.72
Mean				0.274	0.279	0.28	0.326	0.222	0.363
Variance				0.138	0.148	0.083	0.111	0.115	0.106

Table 7.2: The minimal C_A value across for each set of cyber liability prices offered by the insurers in our dataset. The best score is in bold text. The corresponding parameter values are described in Appendix C.

No loss distribution consistently outperformed all of the others. This could results from heterogeneity in the sets of observed prices. Factors might include the range of maximum limits and deductibles, or differences in the policy wording. Equally, heterogeneity may result from differing expectations among insurers about losses. In short, the best loss distribution is contingent on the set of prices it aims to predict. So what can we say?

Aggregating the scores suggests the Gamma distribution is the best candidate for predicting cyber liability prices, with the Burr and Weibull performing poorly. Price set 11 is curiously better explained by a Weibull parameterisation with a heavy tail. The actuarial model behind the 11th set of prices might expect a relatively more heavy tailed distribution of losses.

	Polynomial	Lognormal	Pareto	Burr	Gamma	Weibull
Polynomial	-	0.973	0.28	0.595	0.34	-0.017
Lognormal	0.981	-	0.258	0.607	0.389	0.048
Pareto	0.188	0.146	-	0.817	0.656	0.707
Burr	0.756	0.75	0.64	-	0.736	0.567
Gamma	0.422	0.458	0.449	0.685	-	0.668
Weibull	-0.194	-0.146	0.688	0.283	0.374	-

Table 7.3: Describing how the optimal C_A score for one distribution correlates with the optimal C_A score of another distribution across all of the rate schedules ($n = 26$). The top triangle shows the Spearman's rank correlation coefficient and the bottom triangle shows the Pearson correlation coefficient.

Seeing how these scores correlate with each other sheds some light. The scores in Table 7.3 tilting towards positive suggests that some sets of prices are easier to predict than others. The Lognormal's performance on a given set of prices is a remarkably good indicator of the polynomial distribution's performance. If the Weibull predicts a set of prices relatively well, we can expect the Polynomial/Lognormal to do relatively poorly. This provides more evidence of differing expectations regarding how heavy-tailed losses are.

7.2.3 The County Fair Cyber Loss Distribution

The previous subsection provides 26 potential candidates for the parameterised distribution of cyber losses. An argument could be made that the set of prices with the highest maximum limit is most useful, as this leads to the least extrapolation. Alternatively, we might want to select the distribution achieving the lowest C_A score, even though this discards all but 132 of our data points. Instead we take our lead from Francis Galton's [140] method for estimating the size of an ox by aggregating guesses from attendees at a county fair. A similar method was applied to estimate the number of jelly beans in a jar [133].

The County Fair Cyber Loss Distribution (CFCLD) is derived by averaging the optimal parameterised loss distribution for each set of prices for a given firm and coverage type. The rest of this section will illustrate the CFCLD by considering cyber

	County Fair	**Poly**	**Lognorm**	**Pareto**	**Gamma**	**Weibull**
$P(\$0 < X < \$10\text{K})$	0.9146	0.7894	0.9021	0.9968	0.9612	0.6374
$P(\$10\text{K} < X < \$50\text{K})$	0.0386	0.1422	0.0581	0.0017	0.0124	0.0007
$P(\$50K < X < \$100\text{K})$	0.0094	0.0257	0.0143	0.0004	0.0052	0.0003
$P(\$100\text{K} < X < \$250\text{K})$	0.0089	0.0195	0.0121	0.0004	0.0065	0.0004
$P(\$250\text{K} < X < \$500\text{K})$	0.0048	0.0084	0.0055	0.0002	0.0043	0.0003
$P(\$500\text{K} < X < \$1\text{M})$	0.0035	0.0053	0.0034	0.0001	0.0035	0.0003
$P(\$1\text{M} < X < \$2.5\text{M})$	0.0031	0.0041	0.0025	0.0001	0.0034	0.0004
$P(\$2.5\text{M} < X < \$5\text{M})$	0.0015	0.0018	0.001	0.0001	0.0017	0.0003
$P(\$5\text{M} < X < \$10\text{M})$	0.0009	0.0012	0.0005	0	0.0011	0.0003
$P(\$10\text{M} < X < \$50\text{M})$	0.0007	0.0014	0.0005	0.0001	0.0006	0.0007
$P(\$50\text{M} < X < \$100\text{M})$	0.0001	0.0003	0.0001	0	0	0.0003
$P(\$100\text{M} < X < \$500\text{M})$	0.0001	0.0004	0	0	0	0.0007
$P(\$500\text{M} < X < \$1\text{B})$	0	0.0001	0	0	0	0.0003
$P(\$1\text{B} < X < \$10\text{B})$	0.0001	0.0001	0	0	0	0.0011
$E(X)_{min}$ (\$)	107,328	242,592	28,360	17,449	23,660	1,3107,18
$E(X)_{mid}$ (\$)	428,261	914,258	70,540	75,655	51,214	6,300,519
$E(X)_{max}$ (\$)	749,194	1,585,924	112,719	133,860	78,768	11,290,318

Table 7.4: An overview of the distribution of losses for the County Fair Cyber Loss Distribution, as well as the average contribution from each of the distributions.

liability losses for a retail firm with revenue of \$50M across all 26 pricing schemes. We could easily derive the CFCLD for a different revenue, firm size or coverage type.

Table 7.4 displays the probability of different loss amounts according to this CFCLD. Treating 0 as part of a continuous distribution means these loss distributions underestimate the proportion of firms facing no losses. Interpreting smaller losses (such as those less than \$50K) as 0 may correct for this. This adjustment suggests an incident rate of 0.0468 per year since 95% of losses are less than \$50K.

The probability of a loss of between \$100K and \$250K is 0.0089. This falls to 0.0083 for losses between $\$250,000$ and \$1M. The probability of a loss of \$1M–\$10M is 0.0045, falling to 0.0009 for losses of \$10–100M.

Inferences about losses beyond the maximum limit are essentially extrapolation. The indemnity payment for a loss exceeding the limit by a dollar is the same as for a loss exceeding the limit by a billion dollars. This leads to a question regarding losses exceeding \$50M as only one firm provides limits beyond this. For the CFCLD, this amounts to 1.4% of the distribution and it is not clear how it should be interpreted.

Table 7.4 also describes the average contribution to the CFCLD from each distribution. The contributions of each distribution are weighted according to how often

each distribution led to the highest C_A score for a given set of prices. Consequently the Gamma, Lognormal, Polynomial, Pareto and Weibull distributions contribute 16, 5, 3, 1 and 1 respectively, with no contribution from the Burr distribution.

The implied Pareto distribution is notable for it's vanishing tail and it corresponds to the set of prices including the remarkably cheap policy in Figure 4.8. The Weibull distribution (at least for this parameterisation) is a notable outlier and provides most of the support for the CFCLD's tail. In fact, the singular Weibull parameterisation contributes 99.2% of the losses greater than $100M despite comprising one 26th of the CFCLD.

Table 7.4 also includes approximations of the expected loss of each distribution using the buckets in the table. $E(X)_{min}$, $E(X)_{mid}$ and $E(X)_{max}$ are calculated by summing the $P(a < X < b)$ weighted by the minimum (a), midpoint ($\frac{b-a}{2}$), and maximum b of the range respectively. This presentation allows the reader to understand which sections of each distribution are contributing the most. For example, the over-extrapolated section of the Weibull distribution representing losses of between $1 billion and 10 billion contributes 54% of $E(X)_{mid}$.

Readers might instead calculate expected losses by counting all losses above $100M as $100M. They might also decide to omit the contribution of the Weibull distribution. Extracting a concrete expected loss from these distributions is so challenging because any expected loss relies on some degree of extrapolation, unless the maximum limit exceeds plausible losses.

Section 7.1.4 provides a simple way of converting between different coverage types. Table 7.4 describe the cyber liability losses of a retail firm with revenue of $50M. Section 4.2 showed that business interruption is on average priced at 38% of cyber liability. Section 7.1.4 suggests we can multiply all of the probabilities by 0.38 to estimate the distribution of business interruption incidents for the same hypothetical firm. This results in an expected loss of $154K for business interruption events. The same process for ransomware and wire transfer incidents leads to expected losses of $64K and $51K respectively.

7.3 Discussion

We discuss how the results relates to other attempts to quantify cyber losses in Section 7.3.1. We critique our method in Section 7.3.2.

7.3.1 Quantifying Cyber Losses

The results allows us to comment on the shape of the distribution of losses and point-estimates of losses in dollars. Not identifying a single best distribution for cyber losses is in line with data breach studies, which have found no consensus; subsequent studies concluded breaches were best described by power law [79], Lognormal [80], Pareto [81], and log-skew-normal [82] distributions. However, the Gamma distribution being the most likely candidate differs from such studies.

Liability dollar losses could plausibly be less heavy tailed than distributions of the number of records in a data breach. Liability costs are assigned by courts, which are unlikely to follow the distributions found in data breaches. For example, an equivalent legal ruling to Yahoo!'s[2] breach affecting up to 3 billion accounts is unlikely as it is an order of magnitude bigger than the next largest. Further evidence can be found in a study [75] of operational losses, which found that "the distribution of the non-cyber risk sample is much heavier tailed than that of the cyber risk".

Pricing different coverage types by multiplying the price of cyber liability coverage by a constant suggests each type of cyber loss is driven by a similarly shaped distribution. We have two separate studies [84, 85] that find the recovery time for IT systems to be well modelled by a Lognormal distribution, much like number of records in [80]. Although we cannot map these distributions to financial costs, this evidence supports the viability of multiplicative pricing. Given the different coverage types that continue to be bought and sold, we can make estimates about previously under-studied types of cyber losses like ransomware.

The lower-bound, mid-point and upper-bound for the expected loss from the CFCLD are \$107k, \$428K, and \$749K respectively. For comparison, one study of 921 event costs has mean of \$7.84M, median of \$250K and maximum of

[2] `https://reut.rs/2xSpiXr`

$750M [4]. Another study found mean and median cyber losses of $41M and $1.9M respectively [75].

Why are the point estimates from the County Fair Loss Distribution smaller than related studies? Our inferences are based on the insurance prices for a retail firm with a revenue of $50 million. Meanwhile we do not know the size or industry of the firms constituting the samples from which [4, 75] are based. The firms likely have revenues exceeding $50M. This highlights the problem of granularity in cyber security data.

7.3.2 Reflecting on the Method

If our aim is to uncover the true loss distribution, there are two types of error: insurer error and method error. Insurer error is driven by the insurer's uncertainty around the loss distribution faced by an insured party. Method error is introduced by making flawed inferences from the observed prices. We can only seek to reduce method error, but it is worth discussing both.

Insurer Error Insurer error results from uncertainty in quantifying cyber risk (to be contrasted against certainty in quantifying the risk of a coin toss). This results in Knightian "unmeasurable" uncertainty [141] as evidenced by discussions in which insurers bemoan the lack of actuarial data in cyber insurance [22]. Even worse than random uncertainty, insurers may exhibit systemic bias resulting from tight professional networks. The resulting group-think could prevent aggregated information from mitigating random noise [142]. For example, the solicitors' professional indemnity market was systemically under priced in 2010 [17]. Finally, the phenomenon of the underwriting cycle, in which prices across all lines of insurance rise and fall cyclically [143], illustrates how actual prices differ from the actuarially fair price. This introduces noise into our inferences.

Although the causes of cyber losses are new to insurers, costs are often realised like traditional insurance lines. For example, cyber liability is still assigned by courts and cyber business interruption is calculated via lost revenues. Insurers have amassed much experience quantifying realised losses; the industry dates to

at least 14th Century Italy [144] when insurance could be purchased to cover voyages into uncharted territory. Professional qualifications should help with probability estimates given they are improved by even light-weight calibration training interventions [133]. Aggregating the inferred distributions is intended to reduce random noise, much like how different political polls are combined into a *poll of polls* [132].

Prediction markets provide incentives for deviating from group-think. An analogous argument would go: if a line of insurance is systemically over-priced, then individual insurers have an incentive to defect by offering lower prices to increase market share. Figure 4.8 shows how more market entrants drove some insurers to reduce prices. Although losses should discipline generous underwriting, competitive pressures can temporarily lead to systemic under pricing. Cass Sunstein's aphorism "markets incorporate falsehood as well as truth" [132] summarises this discussion.

Method Error Even if the insurer priced risk perfectly, we still might make flawed inferences based on those prices. Results are influenced by the sample of rate schedules and prices. Section 7.2 showed that the inclusion of rate schedule 11 more than doubled the expected loss from the County Fair Cyber Loss Distribution. For our purposes, choosing the largest data set was justified because any anomalies illustrate challenges for future work to overcome. Further, we presented the results so that the reader can extract loss estimates excluding the contribution of anomalous distributions.

Our model assumes risks are considered independently. In reality, insurers must consider how losses correlate with each other and implement costly solvency risk mitigation measures like reinsurance or holding more capital. Not considering insurer risk aversion or wealth endowment could lead to flawed inferences. These factors complicate extracting probabilities from prediction markets [145]. Our method is likely to face even greater challenges given the relative complexity of insurance markets as compared to the binary events considered by conventional prediction markets.

Given the novelty of the method, is is unclear whether the prediction errors are acceptable. For example, the best performing distribution (the Gamma) across all of the rate schedules has an average error of 22%. In our defence, there is heterogeneity in the expected losses of each insurer and the observed prices do not account for the underwriter's subjective adjustment. Further, we are predicting as many as $2,200$ data points with just two parameters and one distribution, which are commonly used by insurers [139]. More flexible distributions would improve predictions but increase the risk of over-fitting.

Making the best inference relies on solving an optimisation problem in a non-convex search space. The possibility of terminating at a local minima is ever present. Understanding the search space better and selecting the appropriate optimisation heuristic is a must for future work.

7.4 Summary

We introduced a method to infer loss distributions from insurance prices. The method uses particle swarm optimisation to iterate through candidate parameter values to identify the parameterised loss distribution which best explains the observed prices. The Gamma, Lognormal, Polynomial, Pareto and Weibull distributions best predicted 16, 5, 3, 1 and 1 respectively of the 26 sets of premiums.

The County Fair Cyber Loss Distribution aggregates each of the 26 parameterised distributions to provide estimates about cyber losses for a retail firm with a revenue of $50M. The results suggest the expected loss resulting from cyber liability incidents is $428K with a 0.0055 probability of a loss of between $1M and $10M. Expected losses for business interruption, ransomware and wire transfer incidents are $154K, $64K and $51K respectively.

These estimates complement existing approaches to quantifying cyber losses by providing distributions of dollar losses, cost estimates for novel incident types, and granular insights for a specific revenue and industry. This first attempt at inferring losses from insurance prices can be improved by speaking to insurance professionals to understand how to construct a better sample of prices, including

more flexible distributions to improve predictions, and by analysing performance on prices generated by a known distribution.

8
Conclusion

The starting point of this dissertation is uncertainty about potential cyber losses and the effectiveness of risk controls. Unfortunately there is no single source or repository of related information; it is dispersed across many actors. For example, security vendors understand the functionality of security products. Insurers can draw on expertise in quantifying risks for hundreds of years, as well as insights from aggregating cyber insurance claims data. But these actors hold the information privately. This dissertation has explored three mechanisms by which risk transfer markets can increase information about cyber risk to improve security decisions.

Each contribution belies a faith in the ability of markets to aggregate dispersed information. A richer picture is necessary to understand the gap between processes in an economic model and the corresponding processes in the real world. We tie together some of the implications of the preceding chapters by reflecting on this disparity.

Studying documents used by insurers in Chapter 4 suggests a sophisticated operation that collects reams of security data to precisely calculate prices. Yet the variance in pricing across insurers suggests the majority of insurers are pricing inaccurately. This is in keeping with life insurance, which involves a surprising amount of guesswork despite the reputation of actuarial science [16]. These findings cast doubt over the quality of the information we hope to extract.

Chapter 5 introduced a multi-round model to consider insurer strategies for aggregating claims information. This was motivated by survey results showing that this is the primary means by which insurers improve their understanding of cyber risk. The main recommendation is that insurers should ensure insureds maintain a diverse set of defensive configurations when there is uncertainty about which risk controls are effective. However, these results are premised on the information extracted from claims data. It is unclear how insurers will learn about security controls given how information is collected at present.

Cyber warranties provide an alternative way of identifying effective security controls, at least if the model introduced in Chapter 6 describes reality. Unfortunately many of the inferences require far more information than consumers generally have available, such as the vendor's security efficiency. This is exacerbated by a lack of clarity regarding what warranty contracts actually cover. Both Chapter 5 and Chapter 6 suggest processes by which risk transfer products increase information, which are more complicated in the real world.

Chapter 7 introduced a method to infer cyber loss distributions from a collection of insurance prices. This is the clearest illustration of how risk transfer products increase information about cyber risk. However, predictive errors are often greater than 10% and there are question marks over how well cyber insurance prices reflect the risk.

It is tempting to question why this dissertation includes contributions based on models that fail to accurately describe reality. Rather than wade into an open debate about the value of (unrealistic) economic models, Section 8.1 answers this question by outlining how different stakeholders might interpret and use the insights from this dissertation for their own purposes.

8.1 Benefits and Impact

Insurance professionals can look to Chapter 4 to benchmark the application process and pricing schemes of their own organisation. The diversity of approaches to risk assessment means we cannot make concrete recommendations. Insurers will have

to interpret the results individually. For example, insurers might ask why they are not collecting information about an area of information security that others value. Justifications are conceivable, such as targeting a certain industry, but the exercise should strengthen understanding regardless. Quantitative comparisons with competitors' prices could help similarly.

Describing the model introduced in Chapter 5 should prompt insurers to consider how they collect information. For example, insurers are unlikely to link the absence of an intrusion detection system to the cause of a breach. Instead insurers might consult with forensics specialists to identify which application questions could be directly linked to incidents. Over time this linking might provide analytical insights.

This model has implications for policymakers too. The market tilts towards monopoly when uncertainty is high. The answer is not to break up dominant firms, as this will increase uncertainty and reduce social welfare. Rather, regulators must remain vigilant that dominant firms do not exploit this. The system of regulatory filings in the US might make this task easier. Policymakers might also consider the collective action problem leading insurers away from collecting information in standardised form, undermining future analytical insights.

Risk managers can use our empirical results when considering cyber risk transfer products. For example, our descriptions of risk assessment and pricing in Chapter 4 can help set expectations regarding purchasing cyber insurance. The range of questions in proposal forms means many departments from their organisation would have to be involved. The preliminary analysis of cyber warranties in Chapter 6 provides a coarse description of what might be available.

The distribution of losses introduced in Chapter 7 could be useful to both insurers and risk managers. Risk managers can quantify the financial risk beyond mere point estimates and for incident types not previously studied, such as ransomware and business email compromise. Insurers can again benchmark their own risk estimates against the CFCLD.

The academic community can benefit from new research directions. The next section suggests how our contributions might be taken forward in future work.

8.2 Future Work

The empirical studies in Chapter 4 would benefit from replications in new markets and jurisdictions. For example, there has not been a study of cyber insurance in European markets like Germany, France, Italy or Spain. Similarly, we only collected pricing data about the Californian market. There would also be value in repeat studies to determine whether the market is maturing over time.

Future work related to Chapter 5 could introduce uncertainty into risk assessment and post-breach forensics. Reasoning about the trade-offs between information gained by the insurer and the burden placed on the applicant may provide insights into phenomena like "race-to-the-bottom" cyber risk assessment standards [22] or information asymmetries [60]. Future work could consider an insurer with a limited amount of capital depleted each time a policyholder is attacked and build upon the variance of claims analysis seen in Section 5.2.3, providing insights into aggregated risk.

The model introduced in Chapter 6 could be taken forward by exploring how the balance of risk aversion between vendors and consumers affects the supply and demand for cyber warranties. Another factor to consider is the vendor's costs in terms of mitigating (via self insurance or market insurance) the insolvency risk when increasing the warranty level. Identifying an empirical basis for the parameter choices may increase relevance for practitioners.

Finally, the method from Chapter 7 could be advanced by considering more flexible distributions to improve predictions of observed prices. Combined distributions, in which the tail is parameterised separately, might help given we have assumed the area around 0 and the tail are one distribution so far. But we must remain conscious of the bias–variance dilemma [146]. Insurance industry professionals should be engaged to choose a more representative sample of rate schedules or prices. Further, the technique could be applied to non-cyber lines of insurance.

8.3 Conclusion

To conclude, we return to the question *how might cyber risk transfer products increase information about security decisions?* Claims data can be aggregated to identify effective controls providing information is appropriately collected, is shared and acted upon quickly enough, and there is sufficient diversity in the security postures of the insureds. However, the devil is in the details of what appropriate and quickly enough entail; at present insurers collect information about high-level controls and procedures that cannot be rapidly implemented or linked to losses.

Warranties allow security vendors who believe in the effectiveness of their product to send a costly signal to the market. Consumers should look to the policies that affirmatively cover losses and contain few exclusions. However, we discovered there is too much contractual diversity to compare across warranties at present. We should wait and see how these warranties develop naturally before suggesting some other intervention — although consumers should be cautioned against relying on the cover provided.

Inferring loss distributions from insurance prices provides the most hope at present. Investment decisions can be supported by evaluating their efficacy in protecting potential losses, which we argue the County Fair Loss Distribution can estimate. There is a lot of noise present in our inferences *but* a measurement need not eliminate uncertainty, rather it can reduce uncertainty [133] and still be of value. Our method can reduce uncertainty about losses related to under-explored incidents like business email compromise and business interruption.

Bibliography

[1] Bob Blakley, Ellen McDermott, and Dan Geer. Information security is information risk management. In *Proceedings of the 2001 Workshop on New Security Paradigms (NSPW 2001)*, pages 97–104. ACM, 2001.

[2] Ross Anderson and Tyler Moore. The economics of information security. *Science*, 314(5799):610–613, 2006.

[3] George A Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. In Peter Diamond and Andrew Rothschild, editors, *Uncertainty in Economics*, pages 235–251. Elsevier, 1978.

[4] Sasha Romanosky. Examining the costs and causes of cyber incidents. *Journal of Cybersecurity*, 2(2):121–135, 2016.

[5] Ross Anderson. Why information security is hard-an economic perspective. In *Computer security applications conference, 2001. acsac 2001. proceedings 17th annual*, pages 358–365. IEEE, 2001.

[6] Michael L Rustad and Thomas H Koenig. The tort of negligent enablement of cybercrime. *Berkeley Tech. LJ*, 20:1553, 2005.

[7] Dirk van der Linden and Awais Rashid. The effect of software warranties on cybersecurity. *ACM SIGSOFT Software Engineering Notes*, 43(4):31–35, 2019.

[8] Bruce Schneier. Insurance and the computer industry. *Communications of the ACM*, 44(3):114–114, 2001.

[9] US Department of Homeland Security. Cybersecurity Insurance Workshop Readout Report Available: `http://bit.ly/2qY8MUW`, 2012. [Online; accessed 27-February-2017].

[10] US Department of Homeland Security. Cyber Risk Culture Roundtable Readout Report Available: `http://bit.ly/2qdfraY`. 2013.

[11] US Department of Homeland Security. Healthcare and Cyber Risk Management: Cost/Benefit Approaches Available: `http://bit.ly/2pFjhI0`. 2014.

[12] US Department of Homeland Security. Insurance for Cyber-Related Critical Infrastructure Loss: Key Issues Available: `http://bit.ly/1nZC0wM`. 2014.

[13] ENISA. Incentives and barriers of the cyber insurance market in Europe Available: `http://bit.ly/2qXUUd7`. 2012.

[14] ENISA. Cyber insurance: Recent advances, good practices and challenges available: `http://bit.ly/2fNiQIC`. 2016.

[15] UK Cabinet Office. UK cyber security: the role of insurance Available: `http://bit.ly/1Htk2XB`. 2015.

[16] Richard Victor Ericson and Aaron Doyle. *Uncertain business: Risk, insurance and the limits of knowledge*. University of Toronto Press, 2004.

[17] Daniel W Woods, Ioannis Agrafiotis, Jason RC Nurse, and Sadie Creese. Mapping the coverage of security controls in cyber insurance proposal forms. *Journal of Internet Services and Applications*, 8(1):8, 2017.

[18] Daniel W Woods, Tyler Moore, and Andrew C Simpson. The county fair cyber loss distribution: Drawing inference from insurance prices. In *Proceedings of The 18th Workshop on the Economics of Information Security (WEIS 2019)*, 2019.

[19] Daniel W Woods and Andrew C Simpson. Monte carlo methods to investigate how aggregated cyber insurance claims data impacts security investments. In *Proceedings of The 17th Workshop on the Economics of Information Security (WEIS 2018)*, 2018.

[20] Daniel W Woods and Andrew C Simpson. Cyber-warranties as a quality signal for information security products. In *International Conference on Decision and Game Theory for Security*, pages 22–37. Springer, 2018.

[21] Daniel W Woods and Tyler Moore. Cyber warranties: Market fix or marketing trick. *Communications of the ACM*, to appear, 2019.

[22] Daniel W Woods and Andrew C. Simpson. Policy measures and cyber insurance: A framework. *Journal of Cyber Policy*, 2(2):209–226, 2017.

[23] Daniel W Woods and Andrew C Simpson. Towards integrating insurance data into information security investment decision making. In *2018 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (Cyber SA)*, pages 1–6. IEEE, 2018.

[24] Harold D Lasswell. The structure and function of communication in society. *The communication of ideas*, 37(1):136–39, 1948.

[25] Sakshyam Panda, Daniel W Woods, Aron Laszka, Andrew Fielder, and Emmanouil Panaousis. (submitted) post-incident audits on cyber insurance discounts. *Computers & Security*, 2019.

[26] Daniel W Woods and Tyler Moore. Does insurance have a future in governing cybersecurity? *IEEE Security and Privacy Magazine*, to appear, 2019.

[27] Daniel W Woods and Jessica Weinkle. Market definitions of cyber war. In *Proceedings of the 2019 Conference on Cyber Norms: Dealing with Uncertainty.* The Hague, Netherlands, 2019.

[28] Carol A Siegel, Ty R Sagalow, and Paul Serritella. Cyber-risk management: technical and insurance controls for enterprise-level security. *Information Systems Security*, 11(4):33–49, 2002.

[29] Walter S Baer and Andrew Parkinson. Cyberinsurance in IT security management. *IEEE Security & Privacy*, 5(3):50–56, 2007.

[30] A Marotta, F Martinelli, S Nanni, and A Yautsiukhin. A survey on cyber-insurance. available: `http://www.iit.cnr.it/en/node/36039`, 2015.

[31] Ruperto P Majuca, William Yurcik, and Jay P Kesan. The evolution of cyberinsurance. *arXiv preprint cs/0601020*, 2006.

[32] Angelica Marotta, Fabio Martinelli, Stefano Nanni, Albina Orlando, and Artsiom Yautsiukhin. Cyber-insurance survey. *Computer Science Review*, 2017.

[33] Hazel Glenn Beh. Physicial losses in cyberspace. *Connecticut Insurance Law Journal*, 8:55, 2001.

[34] Peter Zweifel and Roland Eisen. *Insurance Economics*. Springer Science & Business Media, 2012.

[35] Kenneth Joseph Arrow. Uncertainty and the welfare economics of medical care (American Economic Review, 1963). *Journal of Health Politics, Policy and Law*, 26(5):851–883, 2001.

[36] Anthony Giddens. Risk and responsibility. *The Modern Law Review*, 62(1):1–10, 1999.

[37] Ulrich Beck. *Risk society: Towards a new modernity*, volume 17. Sage, 1992.

[38] Tyler Moore. On the harms arising from the equifax data breach of 2017. *International Journal of Critical Infrastructure Protection*, 19(C):47–48, 2017.

[39] Michael Power. *The audit society: Rituals of verification*. OUP Oxford, 1997.

[40] Stephen Hilgartner. The social construction of risk objects: Or, how to pry open networks of risk. *Organizations, uncertainties, and risk*, pages 39–53, 1992.

[41] Ralph Nader. *Unsafe at any speed. The designed-in dangers of the American automobile*. Grossman Publishers, 1965.

[42] Michael Power. *Organized uncertainty: Designing a world of risk management*. Oxford University Press on Demand, 2007.

[43] Kevin Quigley, Calvin Burns, and Kristen Stallard. 'cyber gurus': A rhetorical analysis of the language of cybersecurity specialists and the implications for security policy and critical infrastructure protection. *Government Information Quarterly*, 32(2):108–117, 2015.

[44] Ross Anderson, Rainer Böhme, Richard Clayton, and Tyler Moore. Security economics and the internal market. *Report to the European Network and Information Security Agency*, 2008.

[45] Rainer Böhme and Galina Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Proceedings of The 9th Workshop on the Economics of Information Security (WEIS 2010)*, 2010.

[46] Shauhin A Talesh. Data breach, privacy, and cyber insurance: How insurance companies act as "compliance managers" for businesses. *Law & Social Inquiry*, 43(2):417–440, 2018.

[47] Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of risk and uncertainty*, 26(2-3):231–249, 2003.

[48] Hulisi Ogut, Nirup Menon, and Srinivasan Raghunathan. Cyber insurance and IT security investment: Impact of interdependence risk. In *Proceedings of The 4th Workshop on the Economics of Information Security (WEIS 2005)*, 2005.

[49] Jean-Chrysostome Bolot and Marc Lelarge. A new perspective on internet security using insurance. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 1948–1956. IEEE, 2008.

[50] Marc Lelarge and Jean Bolot. Economic incentives to increase security in the internet: The case for insurance. In *INFOCOM 2009, IEEE*, pages 1494–1502. IEEE, 2009.

[51] Rainer Böhme and Gaurav Kataria. Models and measures for correlation in cyber-insurance. In *Proceedings of The 5th Workshop on the Economics of Information Security (WEIS 2006)*, 2006.

[52] Rainer Böhme. Cyber-insurance revisited. In *Proceedings of The 4th Workshop on the Economics of Information Security (WEIS 2005)*, 2005.

[53] Louise Axon, Arnau Erola, Ioannis Agrafiotis, Michael Goldsmith, and Sadie Creese. Analysing cyber-insurance claims to design harm-propagation trees. In *2019 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (Cyber SA)*, June 2019.

[54] Reinsurance News. Petya cyber industry loss passes $3bn driven by Merck silent cyber, url = `https://bit.ly/2u1dui9`, note = Accessed: 2019-08-24.

[55] Jay P Kesan, Rupterto P Majuca, and William J Yurcik. The economic case for cyberinsurance: for securing privacy in the internet age, 2005.

[56] Jens Grossklags, Nicolas Christin, and John Chuang. Secure or insure?: a game-theoretic analysis of information security games. In *Proceedings of the 17th international conference on World Wide Web*, pages 209–218. ACM, 2008.

[57] F Pouget, M Dacier, VH Pham, et al. on the advantages of deploying a large scale distributed honeypot platform. In *Proceedings of the E-Crime and Computer Evidence Conference*, 2005.

[58] Hemantha SB Herath and Tejaswini C Herath. Cyber-insurance: Copula pricing framework and implication for risk management. In *Proceedings of The 6th Workshop on the Economics of Information Security (WEIS 2007)*, 2007.

[59] Hemantha SB Herath and Tejaswini C Herath. Copula-based actuarial model for pricing cyber-insurance policies. *Insurance Markets and Companies: Analyses and Actuarial Computations*, 2(1):7–20, 2011.

[60] Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. Competitive cyber-insurance and internet security. In *Economics of information security and privacy*, pages 229–247. Springer, 2010.

[61] Aron Laszka, Emmanouil Panaousis, and Jens Grossklags. Cyber-insurance as a signaling game: Self-reporting and external security audits. In *Proceedings of the 9th Conference on Decision and Game Theory for Security (GameSec 2018)*. Springer, 2018.

[62] Galina Schwartz, Nikhil Shetty, and Jean Walrand. Why cyber-insurance contracts fail to reflect cyber-risks. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 781–787. IEEE, 2013.

[63] Complaint in Columbia Casualty Company v. Cottage Health System, No. 2:16-cv-03759 (Circuit District California), 2016.

[64] Tridib Bandyopadhyay, Vijay S Mookerjee, and Ram C Rao. Why IT managers don't go for cyber-insurance products. *Communications of the ACM*, 52(11):68–73, 2009.

[65] Jay Kesan, Ruperto Majuca, and William Yurcik. Cyberinsurance as a market-based solution to the problem of cybersecurity: a case study. In *Proceedings of The 4th Workshop on the Economics of Information Security (WEIS 2005)*, 2005.

[66] Mohammad Mahdi Khalili, Mingyan Liu, and Sasha Romanosky. Embracing and controlling risk dependency in cyber-insurance policy underwriting. In *Proceedings of The 17th Workshop on the Economics of Information Security (WEIS 2018)*, 2018.

[67] Aron Laszka and Jens Grossklags. Should cyber-insurance providers invest in software security? In *European Symposium on Research in Computer Security*, pages 483–502. Springer, 2015.

[68] Xia Zhao, Ling Xue, and Andrew B Whinston. Managing interdependent information security risks: Cyberinsurance, managed security services, and risk pooling arrangements. *Journal of Management Information Systems*, 30(1):123–152, 2013.

[69] Lawrence A Gordon and Martin P Loeb. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)*, 5(4):438–457, 2002.

[70] Rainer Böhme. Security metrics and security investment models. In *International Workshop on Security*, pages 10–24. Springer, 2010.

[71] Rainer Böhme and Tyler Moore. The "iterated weakest link" model of adaptive security investment. *Journal of Information Security*, 7(2):81–102, 2016.

[72] Ulrik Franke. The cyber insurance market in Sweden. *Computers & Security*, 68:130–144, 2017.

[73] Sasha Romanosky, Lillian Ablon, Andreas Kuehn, and Therese Jones. Content analysis of cyber insurance policies: How do carriers write policies and price cyber risk? In *Proceedings of The 16th Workshop on the Economics of Information Security (WEIS 2017)*, 2017.

[74] Rob Thoyts. *Insurance theory and practice*. Routledge, 2010.

[75] Christian Biener, Martin Eling, and Jan Hendrik Wirfs. Insurability of cyber risk: An empirical analysis. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 40(1):131–158, 2015.

[76] European Union Agency for Network and Information Security (ENISA). Commonality of risk assessment language in cyber insurance available: `http://bit.ly/2fNiQIC`. 2017.

[77] Center for Internet Security. CIS Critical Security Controls - Version 6.0 Available: `https://www.sans.org/critical-security-controls`, 2015.

[78] Matthew Curtin and Lee T Ayres. Using science to combat data loss: Analyzing breaches by type and industry. *A Journal of Law and Policy for the Information Society*, 4:569, 2008.

[79] Thomas Maillart and Didier Sornette. Heavy-tailed distribution of cyber-risks. *The European Physical Journal B*, 75(3):357–364, 2010.

[80] Benjamin Edwards, Steven Hofmeyr, and Stephanie Forrest. Hype and heavy tails: A closer look at data breaches. *Journal of Cybersecurity*, 2(1):3–14, 2016.

[81] Spencer Wheatley, Thomas Maillart, and Didier Sornette. The extreme risk of personal data breaches and the erosion of privacy. *The European Physical Journal B*, 89(1):7, 2016.

[82] Martin Eling and Nicola Loperfido. Data breaches: Goodness of fit, pricing, and risk measurement. *Insurance: Mathematics and Economics*, 75:126–136, 2017.

[83] Maochao Xu, Kristin M Schweitzer, Raymond M Bateman, and Shouhuai Xu. Modeling and predicting cyber hacking breaches. *IEEE Transactions on Information Forensics and Security*, 13(11):2856–2871, 2018.

[84] Bianca Schroeder and Garth Gibson. A large-scale study of failures in high-performance computing systems. *IEEE Transactions on Dependable and Secure Computing*, 7(4):337–350, 2010.

[85] Ulrik Franke, Hannes Holm, and Johan König. The distribution of time to recovery of enterprise it services. *IEEE Transactions on Reliability*, 63(4):858–867, 2014.

[86] Verizon LLC. 2018 Data Breach Investigations Report available at `https://enterprise.verizon.com/en-gb/resources/reports/dbir/`, 2018.

[87] Ponemon Institute. Cost of a data breach study available at `https://www.ibm.com/security/data-breach`, 2018.

[88] Chad D Heitzenrater and Andrew C Simpson. Policy, statistics and questions: Reflections on UK cyber security disclosures. *Journal of Cybersecurity*, 2(1):43–56, 2016.

[89] Claudia Biancotti. The price of cyber (in) security: evidence from the italian private sector. In *Proceedings of The 17th Workshop on the Economics of Information Security (WEIS 2018)*, 2018.

[90] Sasha Romanosky, David Hoffman, and Alessandro Acquisti. Empirical analysis of data breach litigation. *Journal of Empirical Legal Studies*, 11(1):74–104, 2014.

[91] Aaron Ceross and Andrew Simpson. The use of data protection regulatory actions as a data source for privacy economics. In *International Conference on Computer Safety, Reliability, and Security*, pages 350–360. Springer, 2017.

[92] Anat Hovav and John D'Arcy. The impact of denial-of-service attack announcements on the market value of firms. *Risk Management and Insurance Review*, 6(2):97–121, 2003.

[93] Katherine Campbell, Lawrence A Gordon, Martin P Loeb, and Lei Zhou. The economic cost of publicly announced information security breaches: empirical evidence from the stock market. *Journal of Computer Security*, 11(3):431–448, 2003.

[94] Huseyin Cavusoglu, Birendra Mishra, and Srinivasan 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.

[95] Sasha Romanosky, Rahul Telang, and Alessandro Acquisti. Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management*, 30(2):256–286, 2011.

[96] Department for Business, Innovation Skills. Information security breaches survey, 2015. [Online; accessed 27-July-2016].

[97] Dinei Florêncio and Cormac Herley. Sex, lies and cyber-crime surveys. In *Economics of information security and privacy III*, pages 35–53. Springer, 2013.

[98] Uskali Mäki. Models are experiments, experiments are models. *Journal of Economic Methodology*, 12(2):303–315, 2005.

[99] Alan Bryman. *Social research methods*. Oxford University Press, 2015.

[100] I Elaine Allen and Christopher A Seaman. Likert scales and data analyses. *Quality progress*, 40(7):64, 2007.

[101] IsecT Ltd. ISO/IEC 27002:2013 information technology — security techniques — code of practice for information security controls (second edition) available: `http://www.iso27001security.com/html/27002.html`, 2016.

[102] Edward Humphreys. Information security management standards: Compliance, governance and risk management. *Information Security Technical Report*, 13(4):247–255, 2008.

[103] Mikko Siponen and Robert Willison. Information security management standards: Problems and solutions. *Information & Management*, 46(5):267–270, 2009.

[104] Satu Elo and Helvi Kyngäs. The qualitative content analysis process. *Journal of advanced nursing*, 62(1):107–115, 2008.

[105] George EP Box. Robustness in the strategy of scientific model building. In *Robustness in statistics*, pages 201–236. Elsevier, 1979.

[106] Milton Friedman and Marilyn Friedman. *Essays in positive economics*. University of Chicago Press, 1953.

[107] Karl Popper. *The logic of scientific discovery*. Routledge, 2005.

[108] Daniel M Hausman. Philosophy of economics. *The Stanford Encyclopedia of Philosophy*, 2018.

[109] Imre Lakatos. Falsification and the methodology of scientific research programmes. In *Can theories be refuted?*, pages 205–259. Springer, 1976.

[110] Roman Frigg. Models and fiction. *Synthese*, 172(2):251, 2010.

[111] Mary S Morgan, Margaret Morrison, and Quentin Skinner. *Models as mediators: Perspectives on natural and social science*, volume 52. Cambridge University Press, 1999.

[112] Daniel M Hausman. *The inexact and separate science of economics*. Cambridge University Press, 1992.

[113] Donald N McCloskey. The rhetoric of economics. *Journal of Economic Literature*, 21(2):481–517, 1983.

[114] Lawrence A Berger. A model of the underwriting cycle in the property/liability insurance industry. *The Journal of Risk and Insurance*, 55(2):298–306, 1988.

[115] Mark Stamp. Risks of monoculture. *Communications of the ACM*, 47(3):120, 2004.

[116] Chad Heitzenrater, Rainer Böhme, and Andrew Simpson. The days before zero day: Investment models for secure software engineering. In *Proceedings of the 15th Workshop on the Economics of Information Security (WEIS 2016)*, 2016.

[117] Daniel Geer, Kevin Soo Hoo, and Andrew Jaquith. Information security: Why the future belongs to the quants. *IEEE Security & Privacy*, 99(4):24–32, 2003.

[118] Jospeh Bertrand. Theorie mathematique de la richesse sociale. *Journal des Savants*, pages 499–508, 1883.

[119] Carl Shapiro and Hal R Varian. *Information rules: a strategic guide to the network economy*. Harvard Business Press, 1998.

[120] Neal Fultz and Jens Grossklags. Blue versus red: Towards a model of distributed security attacks. In Roger Dingledine and Philippe Golle, editors, *International Conference on Financial Cryptography and Data Security*, pages 167–183. Springer, 2009.

[121] Benjamin Johnson, Rainer Böhme, and Jens Grossklags. Security games with market insurance. In *Proceedings of the 2nd Conference on Decision and Game Theory for Security*, pages 117–130. Springer, 2011.

[122] Ranjan Pal and Leana Golubchik. Analyzing self-defense investments in internet security under cyber-insurance coverage. In *Proceedings of the IEEE 30th International Conference on Distributed Computing Systems (ICDCS2010)*, pages 339–347. IEEE, 2010.

[123] Hideyuki Tanaka, Kanta Matsuura, and Osamu Sudoh. Vulnerability and information security investment: An empirical analysis of e-local government in japan. *Journal of Accounting and Public Policy*, 24(1):37–59, 2005.

[124] Hal R. Varian. *Intermediate Microeconomics: A Modern Approach*. W.W. Norton Co., New York, eighth edition, 2010.

[125] Hal Varian. Managing online security risks available: `http://people.ischool.berkeley.edu/~hal/people/hal/NYTimes/2000-06-01.html`. *New York Times*, 1, 2000.

[126] Christian Biener, Martin Eling, and Jan Hendrik Wirfs. Insurability of cyber risk: An empirical analysis. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 40(1):131–158, 2015.

[127] Friedrich August Hayek. The use of knowledge in society. *The American economic review*, 35(4):519–530, 1945.

[128] Burton G Malkiel. The efficient market hypothesis and its critics. *Journal of economic perspectives*, 17(1):59–82, 2003.

[129] David Rothschild. Forecasting elections: Comparing prediction markets, polls, and their biases. *Public Opinion Quarterly*, 73(5):895–916, 2009.

[130] Philip M Polgreen, Forrest D Nelson, George R Neumann, and Robert A Weinstein. Use of prediction markets to forecast infectious disease activity. *Clinical Infectious Diseases*, 44(2):272–279, 2007.

[131] Adam Mann. The power of prediction markets. *Nature News*, 538(7625):308, 2016.

[132] Cass R Sunstein. *Infotopia: How many minds produce knowledge*. Oxford University Press, 2006.

[133] Douglas W Hubbard. *How to measure anything: Finding the value of intangibles in business*. John Wiley & Sons, 2014.

[134] Rainer Böhme. A comparison of market approaches to software vulnerability disclosure. In *International Conference on Emerging Trends in Information and Communication Security*, pages 298–311. Springer, 2006.

[135] Moshe A. Milevsky, Thomas S. Salisbury, and Alexander Chigodaev. The implied longevity curve: How long does the market think you are going to live?, 2018.

[136] James Kennedy. Particle swarm optimization. In *Encyclopedia of Machine Learning*, pages 760–766. Springer, 2011.

[137] Lester James V. Miranda. Pyswarms, a research-toolkit for particle swarm optimization in python, 2017. 10.5281/zenodo.986300.

[138] David Bahnemann. Distributions for actuaries. *CAS Monograph Series*, (2), 2015.

[139] Krzysztof Burnecki, Grzegorz Kukla, and Rafał Weron. Property insurance loss distributions. *Physica A: Statistical Mechanics and its Applications*, 287(1-2):269–278, 2000.

[140] Francis Galton. Vox populi (the wisdom of crowds). *Nature*, 75(7):450–451, 1907.

[141] Frank H Knight. *Risk, uncertainty and profit*. Courier Corporation, 2012.

[142] Scott E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies.* Princeton University Presss, 2007.

[143] Greg Niehaus and Andy Terry. Evidence on the time series properties of insurance premiums and causes of the underwriting cycle: new support for the capital market imperfection hypothesis. *Journal of Risk and Insurance*, pages 466–479, 1993.

[144] Humbert O Nelli. The earliest insurance contract. a new discovery. *Journal of Risk and Insurance*, pages 215–220, 1972.

[145] Charles F Manski. Interpreting the predictions of prediction markets. *economics letters*, 91(3):425–429, 2006.

[146] Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992.

Appendices

A

Qualitative Research Design and Analysis

This appendix describes the designs of the qualitative work contained in Chapter 4. We share the questionnaire design and anonymised responses as a comma-separated values file in the electronic appendix. This file was exported from the online distribution platform and the column with company identifiers was deleted. The code book used to conduct and document the content analysis is included in the electronic appendix. However, we will not share coded data from the proposal form analysis because this could be used to de-anonymise the company and possibly the respondent.

The interviews were conducted under the agreement that only "selected quotes" would be published. The terms were drafted like this so participants were open and did not revert to the public position of their employer. Consequently, we cannot share full transcripts or extended extracts. However, we share the interview design:

> Opening question: "Which trends do you see emerging in cyber insurance and which trends do you see emerging in the future?"
>
> Q1: "What stages would an applicant go through in applying for cyber insurance?" Follow up questions to clarify whether certain methods (proposal form, onsite audit, telephone interview) are used and what their role in the process is.
>
> Q2: "How do you balance the information gained from each? Follow up question of what happens if there are discrepancies?"
>
> Q3: "Which areas of information security do you feel the assessment process provides insights into? Which areas are relatively less well covered?" Follow up questions around how they assess certain areas like human resources, processes, organisational structure.
>
> Q4: "So you've said these are the areas that can and can't be assessed, what would you say are the most important areas to assess, regardless of the ability to do so?"
>
> Q5: "How do you keep up-to-date on what the important areas of information security are?"

Q6: “In 10 years time what will cyber insurance assessment look like?”

After the previous six questions, the interviews became far less structured and would vary between participants, we covered topics like:

- Whether insurers directly recommend security controls
- What coverage will look like in the future
- The role of IoT
- Definitions like "cyber war" and "cyber terrorism"
- The challenge of assessing aggregated risk

B

Error Bars for the Simulations

Standard Error of the Mean (10^{-3})

k	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$
0	1.59	3.17	4.86
1	1.65	3.28	4.95
2	1.71	3.39	5.03
3	1.77	3.49	5.09
4	1.83	3.59	5.15
5	1.89	3.65	5.20
6	1.95	3.68	5.22
7	2.01	3.67	5.23
8	2.07	3.60	5.24
9	2.09	3.48	5.22
10	1.86	3.28	5.19
11	1.20	3.03	5.13
12	0.48	2.72	5.06
13	0.12	2.38	4.98
14	0.02	2.02	4.86
15	0.00	1.67	4.72
16	0.00	1.34	4.56
17	0.00	1.04	4.37
18	0.00	0.79	4.16
19	0.00	0.58	3.90
20	0.00	0.41	3.61
21	0.00	0.29	3.26
22	0.00	0.19	2.85
23	0.00	0.12	2.35
24	0.00	0.06	1.68
25	0.00	0.00	0.00

Table B.1: Standard error of the mean when simulating the original IWL for various values of σ. Section 3.4 refers to this table.

C
Parameter Values

Set	Poly	Lognrm	Pareto	Burr	Gamma	Weibull
1	-1.415, 6.39	0.6926, 5.1	0.329, 0.03881	0.2066, 1.576	0.007277, 1.147e-06	0.08108, 0.001172
2	-1.62, 564.9	5.643, 3.364	0.2652, 2.877	1.133, 0.2318	0.02423, 1.238e-06	0.06136, 0.002266
3	-1.66, 548.6	4.612, 3.692	0.3022, 2.511	0.6558, 0.4608	0.008544, 6.02e-07	0.5245, 0.001431
4	-1.655, 308.7	4.411, 3.593	0.3517, 2.501	1.2, 0.2931	0.02773, 3.106e-06	0.06657, 0.0005691
5	-1.286, 2.255	0.3091, 6.4	0.2558, 0.1918	0.1512, 1.494	0.008674, 9.961e-08	0.06174, 0.002714
6	-1.655, 340.9	4.694, 3.515	0.3517, 3.0	0.7278, 0.4834	0.008751, 5.452e-07	0.6507, 0.0005137
7	-1.396, 255.8	4.974, 4.838	0.2031, 0.1322	0.7801, 0.2604	0.01345, 1.018e-07	0.06178, 0.002765
8	-1.398, 1.19	1e-10, 5.126	0.3922, 0.3753	0.7442, 0.5269	0.001502, 9.592e-08	0.05443, 0.0002993
9	-1.722, 710.6	4.859, 3.475	0.3338, 2.073	0.8557, 0.3901	0.007281, 4.453e-07	0.3396, 0.0008387
10	-1.708, 494.9	4.722, 3.441	0.3449, 3.0	1.032, 0.3336	0.01308, 1.094e-06	0.05877, 0.0006117
11	-1.489, 1.0	1e-10, 4.215	0.3377, 5e-10	0.2231, 2.094	0.003038, 1.892e-07	0.1066, 0.001243
12	-1.796, 4023.0	7.622, 2.728	0.2625, 3.0	1.036, 0.2534	0.03272, 8.414e-07	0.0433, 0.001744
13	-1.367, 1.371	1e-10, 5.452	0.3571, 0.2254	0.5841, 0.6112	0.002467, 1.147e-07	0.491, 0.0004963
14	-1.706, 262.3	2.881, 3.949	0.3924, 2.878	1.165, 0.3368	0.008694, 1.168e-06	0.0536, 0.0004344
15	-1.42, 1.0	1e-10, 4.828	0.4142, 0.1765	0.3608, 1.159	0.001997, 2.887e-07	0.8208, 0.0002809
16	-1.468, 172.3	4.228, 4.426	0.2697, 0.01978	0.6569, 0.4132	0.01359, 2.172e-07	0.0687, 0.001438
17	-1.369, 1.0	1e-10, 5.418	0.3692, 0.6712	1.219, 0.303	0.001579, 5.02e-08	0.3039, 0.0001574
18	-1.438, 10.51	0.5467, 5.178	0.325, 0.03928	0.1881, 1.707	0.006347, 9.875e-07	0.08164, 0.001064
19	-1.307, 1.042	1e-07, 6.028	0.3072, 0.3652	0.6888, 0.446	0.004438, 1.012e-07	0.07249, 0.001049
20	-1.417, 82.3	3.978, 4.738	0.2743, 2.014	0.699, 0.3914	0.008581, 9.422e-08	0.04431, 0.0007566
21	-1.469, 1.0	1e-07, 4.339	0.4641, 0.04153	0.1861, 2.393	0.001529, 6.166e-07	0.6957, 0.0002709
22	-1.44, 124.6	2.485, 5.876	0.1988, 3.06e-05	0.1073, 1.118	0.01903, 9.225e-07	0.08372, 0.004
23	-1.709, 311.7	4.34, 3.406	0.3818, 3.0	0.6891, 0.5541	0.01925, 2.315e-06	0.0584, 0.0005017
24	-1.477, 20.38	1e-13, 5.427	0.3622, 0.1995	0.6908, 0.5254	0.002696, 1.385e-07	0.7393, 0.0004358
25	-1.439, 20.87	5.28e-07, 5.914	0.2942, 0.004664	1.043, 0.007721	0.01292, 2.127e-06	0.09196, 0.002301
26	-1.478, 104.4	4.655e-07, 6.201	0.3179, 9.962	0.8672, 0.3542	0.008126, 2.027e-07	0.03755, 0.0002793

Table C.1: The parameter values corresponding to Table 7.2.