The Economics of Information Security: Investment, Insurance, and Evaluation

Henry Robert Keith Skeoch
University College London
Department of Computer Science

2024

Submitted to University College London (UCL) in partial fulfilment of the requirements for the award of the degree of Doctor of Philosophy.

Primary supervisor: Professor David J. Pym
Secondary supervisor: Professor Christos Ioannidis

Examining committee: Professor Rainer Böhme (Universität Innsbruck) and Professor Julian Williams (Durham University)

A strange game.

The only winning move is not to play.

— WOPR [1], *WarGames* (1983)

[1] War Operation Plan Response

Declaration

I, Henry Robert Keith Skeoch, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Henry Robert Keith Skeoch
London, United Kingdom
17th January 2024

Word count: 51,494

Abstract

The complexity and connectivity of modern information technology systems present a significant modelling challenge from a security perspective. The field of security economics aims to use the tools of economics to reason about security problems to understand and evaluate the different possible outcomes of varying security postures. This thesis is primarily focused on the economics of cyber-insurance, which provides indemnity for individuals and organizations against financial losses related to degradation in cyber-security risk parameters such as confidentiality, integrity, or availability. The research presented addresses several pertinent questions related to cyber-insurance. An existing influential model for cyber-security investment decisions, the Gordon-Loeb Model, is expanded to included cyber-insurance and the optimal combination of defensive security investment versus insurance is investigated. A modelling framework is then developed that combines a systems-focused descriptive approach to security modelling using entity relationship diagrams and security maturity models to demonstrate how their outputs might be used to provide or adjust parameters for an expected utility maximization insurance pricing approach. This model helps to provide a consistent methodology for pricing cyber-insurance, which then poses the question of scalability and insurance market capacity. If a market is not efficient, theory suggests financial imbalances will develop. Via simulations, it is demonstrated that better information sharing is a key condition to strengthen the sustainability of the cyber-insurance market. Finally, an economic model of a ransomware attack across a network is developed. Ransomware is a key concern for cyber-insurers as it has the potential to trigger immediate financial losses and accordingly is an important modelling target for the cyber-insurance industry.

Impact Statement

This thesis makes a number of contributions to both the academic field of security economics and to practitioners in cyber-insurance.

The literature review organizes the existing literature on cyber-insurance into categories, providing structure to a literature that is currently fairly disperse in terms of publication venue and content.

In terms of the specific scholastic contributions, Chapter 3 expands a well-established model in security economics, the Gordon-Loeb model, to cyber-insurance. This is of theoretical interest from a decision theoretic approach to security investment decisions but the model simulations developed could easily be used by a corporate decision-maker given an appropriate set of parameters. Chapter 4 proposes a novel modelling framework for combining a systems 'picture' of an organization and using that to develop a security maturity assessment of that organization which can then be used to provide parameters or adjustments to parameters in a utility function for pricing cyber-insurance. Chapter 5 tackles the significant and important problem of the conditions required for efficiency in cyber-insurance markets and the conditions needed to provide more risk capacity to that market. In a geopolitical climate where organizations face significant harms from cyber-threats, cyber-insurance can help mitigate the financial losses associated with cyber-incidents and provide access to external expertise to help organizations recover. The conclusion of the modelling work in this chapter is that better information sharing about cyber-incidents is needed to assist with convergence in beliefs on probability distributions of cyber-losses, which is a key condition for efficiency. Finally, Chapter 6 presents a partially observable Markov decision process (POMDP) model of a ransomware attack. This could be used by strategic decision-makers, such as a Chief Information Security Officer (CISO), to engage in discussions with managers and ultimate risk owners about the optimal defence strategy against such attacks.

At a more structural level, security economics has a long-standing problem in that much research is either very detailed and focused on specific systems, attacks or vulnerabilities. On the broader scale, there is a criticism that is valid in some cases that economic models are too abstract to be of practical use. This thesis aims to show that it is possible to combine both approaches to deliver useful insights, particularly within the insurance domain. Robust modelling of outcomes might be used to inform policy and regulation around cyber-insurance, response to cyber-threats, and providing rigour and structure to defensive security investment decisions.

Acknowledgements

The first, and most significant acknowledgement goes to my principal supervisor, Professor David Pym. I first met David on a rainy day in March 2019 knowing very little about cyber-security but left his office absolutely convinced that spending the next 4 years of my life studying for a PhD at UCL was a very exciting prospect. The experience has not disappointed. Particular thanks go to David for his guidance and support at all stages of the PhD and especially for his patience whilst dealing with my sometimes hair-brained ideas. Second, to Professor Christos Ioannidis, my second supervisor. A conversation with Christos is never dull and always concludes with at least three more books added to one's reading list. Chris's knowledge of the economics literature is encyclopedic and has saved me a great deal of research time, though the time saved has inevitably then been reallocated to discussions about politics, economics, sports, or some other subject. I have been extremely fortunate to have had the support of two supervisors whose company I have hugely enjoyed (even if our discussions ended up being more virtual than we initially expected).

I would like to thank all those at UCL who I have studied with. My fellow students on the inaugural Centre for Doctoral Training in Cyber-Security cohort — Antoine, Antonis, Arianna, Hawra, Niamh, Phil, and Sergi — and our 'honorary CDT member in the office', Sharad, have been great fun to study alongside and have made the necessary but sometimes dull aspects of doctoral research much easier to bear. Marie Vasek, Lorenzo Cavallaro, Philip Jovanovic, and Steven Murdoch proved to be very kind 'employers' of my services as a Postgraduate Teaching Assistant; I learnt a lot about security from all of them and greatly strengthed my technical understanding from our discussions. Finally, to the CDT Centre Managers, Fiona

Mannion and Pui Sin and the CDT Directors, David Pym, Madeline Carr and Shane Johnson for keeping the ship sailing and creating the opportunity to pursue doctoral research on an exciting multi-disciplinary programme.

I have greatly enjoyed research discussions with Daniel Woods and Martin Eling, both of whom produced research which has proved hugely inspiring to me in the course of this work. Lawrence Gordon and Martin Loeb were hugely helpful in response to questions about their model and the scope for expanding it to insurance. Tyler Moore and the organizing committee of the Workshop on the Economics of Information Security 2022 in Tulsa, Oklahoma provided a wonderful opportunity to interact with fellow researchers as the pandemic started to ease and helped inspire the later research in this thesis.

The thesis examiners, Professors Rainer Böhme and Julian Williams provided many helpful comments and suggestions that have greatly improved the quality of the thesis and the presentation of the ideas contained within it.

On a personal note, I would like to thank my parents and family for their support throughout my life so far. Professor Jonathan Goodman at the University of Cambridge generously responded to my endless requests for references back in 2019, having supervised my Part III Chemistry thesis over 10 years prior. Finally, to my former Barclays colleagues who encouraged me to take a massive risk and leave a comfortable job to go and start a PhD — the payoff was worth the premium.

Funding

This work was supported partly by the UK Engineering and Physical Sciences Research Council grant for Doctoral Training EP/R513143/1.

Contents

Table of Acronyms **15**

List of Figures **16**

List of Tables **18**

1 Introduction **21**

1.1 Background 21

1.1.1 The development of information security as a field of research 22

1.2 Systems and security 23

1.2.1 Core concepts in security 24

1.2.2 Describing systems 24

1.2.3 Measuring security 25

1.3 Economics 26

1.4 Cyber-insurance 27

1.4.1 Demand 28

1.4.2 Supply 28

1.5 Thesis structure 29

1.5.1 Relation to prior published work 30

1.6 Thesis Contributions 30

2 Literature Review **33**

2.1 Insurance economics 33

2.1.1 Expected utility and the theory of insurance demand 34

2.1.2 Information asymmetry, adverse selection and moral hazard . 36

2.1.3 Insurance supply and pricing 38
2.1.4 Reinsurance 39
2.1.5 Relevant contemporary work 40
2.2 Security investment models 41
2.2.1 The Gordon-Loeb Model and some alternatives 41
2.2.2 Criticisms of the Gordon-Loeb Model 42
2.2.3 Extensions of the Gordon-Loeb Model 43
2.3 Systematic cyber-insurance literature review 43
2.3.1 Economic modelling 44
2.3.2 Actuarial models 46
2.3.3 Frameworks and policy 47
2.3.4 Game theory 48
2.3.5 Law 49
2.3.6 Surveys and empirical analyses 50
2.3.7 Insurance-linked security pricing 52
2.3.8 Capital requirements 53

3 Expanding the Gordon-Loeb Model to Cyber-Insurance 54

3.1 Introducing the Gordon and Loeb Model 54
3.1.1 Key results from the Gordon-Loeb Model 55
3.2 Related work 57
3.2.1 Critique of the approach of Young et al (2016) to combining the Gordon-Loeb Model and Cyber-insurance 57
3.2.2 Mazzoccoli and Naldi on mixed insurance/investment cyber-risk management 57
3.2.3 Return on Security Investment (ROSI) 58
3.3 Economic utility 59
3.3.1 Utility functions 59
3.4 Incorporating cyber-insurance into the Gordon-Loeb Model 61
3.5 Simulation 64
3.5.1 Method 64

3.5.2 Optimal investment per the Gordon-Loeb Model, variable cover 65

3.5.3 Premium versus vulnerability under optimal security investment 67

3.6 Remarks 72

3.6.1 Model limitations 72

3.7 Summary 75

4 Pricing Cyber-Insurance Based on System Structure 76

4.1 Background 76

4.2 Problem statement 77

4.3 Theory 78

4.3.1 Grammar 78

4.3.2 Describing systems 78

4.3.3 Describing security 80

4.3.4 Parameters 80

4.3.5 Security Maturity Models 81

4.4 Related work 81

4.5 Chapter organization 82

4.6 Modelling framework overview 83

4.6.1 Preliminaries 84

4.6.2 Modelling framework structure and parameters 87

4.7 Setting up the modelling framework 90

4.7.1 Entity Relationship Diagrams 92

4.7.2 Security Maturity Models 99

4.7.3 Pricing cyber-insurance using utility functions 104

4.8 Using the modelling framework 111

4.8.1 Case studies 111

4.8.2 Online retail 113

4.8.3 Consumer Bank 120

4.8.4 Pharmaceuticals 124

4.8.5 Reflecting on the Case Studies 126

4.9 Discussion 127

4.9.1 Real world considerations for using the modelling framework 127
4.9.2 Weaknesses 129
4.9.3 Extensions and refinements 130

5 Modelling the Cyber-Insurance Market with Risk Transfer via Reinsurance 131

5.1 Background 132
5.1.1 Insurance market structure 132
5.1.2 Technological advancement, information deficiency and cyber-insurance 134
5.1.3 What claims might arise in relation to cyber-insurance? . . . 135
5.1.4 What is the motivation for reinsurance involvement in the cyber-insurance market? 136
5.1.5 Relation to existing literature 137
5.2 Economic Theory on Efficiency 137
5.2.1 Ex-ante efficiency 138
5.2.2 Ex-post efficiency 138
5.2.3 Rational belief equilibria 140
5.3 Model 140
5.3.1 Insurance buyer 141
5.3.2 Supply of insurance 142
5.3.3 Introducing reinsurance 144
5.3.4 Modelling cyber-risks 146
5.4 Simulations 151
5.4.1 Preliminaries 152
5.4.2 Simulation Strategy 154
5.4.3 Simulation 1: Benchmark simulation 155
5.4.4 Simulation 2: Reinsurance supply and price 158
5.4.5 Simulation 3: Insurance buyers of variable risk 171
5.5 Discussion 175
5.5.1 Information asymmetry 175

5.5.2 Cyber-insurer loss experience 176

5.5.3 Loss transparency 179

5.5.4 Consistency of reference 181

5.5.5 Supply and demand 181

5.5.6 Further Work 182

5.5.7 Conclusions 182

6 Threat Modelling Ransomware Attacks on Enterprise Networks 185

6.1 Background 186

6.1.1 Modelling ransomware requires a clear, disciplined approach . 187

6.1.2 Distinguishing different ransomware attacks 188

6.1.3 Ransomware may incur reparative costs 189

6.2 Theory 189

6.2.1 Solutions 190

6.3 Related work 191

6.3.1 Game theory models of ransomware infection 191

6.3.2 POMDP models of penetration testing 192

6.3.3 Incident response and recovery 193

6.3.4 Business continuity insurance 194

6.3.5 Network malware models 194

6.4 Model 194

6.4.1 Problem statement and economic considerations 194

6.4.2 POMDP Model Structure 196

6.4.3 Pricing ransomware insurance 200

6.5 Simulations 202

6.5.1 Sensitivity analysis 202

6.5.2 Ransomware infection 206

6.5.3 Insurance pricing 214

6.6 Further work 215

6.7 Summary 217

7 Conclusion **218**

7.1 Closing remarks 218

7.2 Further work 220

Bibliography **221**

Author Contribution Statement **255**

Table of Acronyms

Notation	Description
CARA	Constant Absolute Risk Aversion.
CIA	Confidentiality, Integrity, and Availability.
CISO	Chief Information Security Officer.
CMMC	Cybersecurity Maturity Model Certification.
CRRA	Constant Relative Risk Aversion.
ENBIS	Expected Net Benefit of Investment in Security.
ERD	Entity Relationship Diagram.
ILS	Insurance-linked Security.
ISO	Independent Organization for Standardization.
MDP	Markov Decision Process.
NAIC	US National Association of Insurance Commissioners.
NIST	US National Institute of Standards and Technology.
POMDP	Partially Observable Markov Decision Process.
RDS	Realistic Disaster Scenario.
ROSI	Return on Security Investment.
SCCS	Synchronous Calculus of Communicating Systems.
SMM	Security Maturity Model.
WEIS	Workshop on the Economics of Information Security.

List of Figures

3.1 Example Gordon-Loeb security breach functions 56

3.2 Utility as a function of cover assuming $z = z^*$ 66

3.3 Maximum available cover under the cash constraint at different levels of z 67

3.4 Utility as a function of investment with maximum insurance coverage purchased 68

3.5 Highest premium rate at which maximum cover may be obtained for vulnerability, v and variation of optimal GL investment with vulnerability, v 69

3.6 Utility functions for different vulnerabilities assuming investment at the Gordon-Loeb optimum, z^* and maximum coverage respecting the model cash constraint 70

4.1 Relationship between criticality/sensitivity and confidentiality/integrity/availability 85

4.2 Conceptual overview of model 88

4.3 A simple entity relationship diagram 94

4.4 Chen notation 96

4.5 Barker notation 96

4.6 Crow's foot notation 97

4.7 Payroll model with criticality and sensitivity 98

4.8 Procedure for estimating potential losses 107

4.9 Entity relationship diagram for online retailer 114

4.10 Entity relationship diagram for consumer bank 121

4.11 Entity relationship diagram for pharmaceutical manufacturing . . . 125

5.1 Categorisation of cyber losses by frequency and severity 135

5.2 Benchmark severity (LHS) and frequency (RHS) distributions 156

5.3 Simulated loss distributions 157

5.4 Reinsurer loss distributions 159

5.5 Loading versus target loss ratio for different insurance portfolios . . 162

5.6 Insurer Simulated Loss Distributions (Section 5.4.4) 164

5.7 Insurer profit with varying quota share proportions (ρ). 168

5.8 US cyber-insurer losses vs premium written 178

5.9 Epanechnikov kernel versus fitted log-normal distribution for NAIC reported cyber-insurance loss ratios, 2018-2021 180

6.1 Stylised transition probability structure for a single machine 198

6.2 Utility versus number of simulation steps. $p_{encrypt} = p$; all other probabilities set to 0.5 205

6.3 Utility versus number of simulation steps. $p_{comrpomise}, p_{spread\_low} = p$; all other probabilities set to 0.5 205

6.4 Utility versus number of simulation steps. $p_{spread\_medium}, p_{spread\_high} = p$; all other probabilities set to 0.5 206

6.5 Simulation step distribution for rare/sophisticated strain 210

6.6 Simulation step distribution for common/unsophisticated strain . . . 210

6.7 Simulation step distribution for 50/50 strain 210

6.8 Utility evolution for rare/sophisticated strain 211

6.9 Utility evolution for common/unsophisticated strain 211

6.10 Utility evolution for 50/50 strain 211

6.11 Varying probabilities of domain controller compromise, Common/Unsophisticated Strain 213

6.12 Varying probabilities of successful repair, Common/Unsophisticated Strain 213

List of Tables

2.1 Summary of systematic literature review on cyber-insurance 44

3.1 Sample parameters for different vulnerabilities, v. $\tau l = 400,000$. $P_{max} = z_{max} - z^*$, i.e. the maximum cash available to pay an insurance premium after investing at the optimal level given by the GL model. p is the highest premium rate at which coverage equal to the expected loss after investment, $\tau l S(v, z^*)$, can be achieved. 71

4.1 Mapping of ERD and natural language grammar 93

4.2 Currently used security maturity models 99

4.3 Objectives for payroll example 102

4.4 d_i parameters for payroll example 103

4.5 Number of practices for each Domain in the CMMC by Level 104

4.6 p_i parameters for payroll example 104

4.7 Objectives for online retailer maturity 116

4.8 d_i parameters for online retailer maturity 116

4.9 Summary of utility function parameters for online retail example . . 119

4.10 Utility function parameters for the consumer bank 123

4.11 Cloud provider and cyber-insurer collaborations 130

5.1 Expected versus simulated benchmark distribution losses 156

5.2 Reinsurance premia for excess-of-loss policies on the 50% chance of claim distribution. 157

5.3 Table of reinsurance distributions 158

5.4 Illustrative premium rates for target loss ratios under different distributions at cover fixed at \$500mn 160

5.5 Insurance contracts in the market 160

5.6 Insurance policies written by insurance panel 161

5.7 Calculating premium loading rates for insurance companies based on simulated losses. The loading rate is expressed in percentage points. 161

5.8 Simulated Losses 165

5.9 Reinsurance ceding fractions that maintain insurer solvency at different stress-test values 166

5.10 Excess of Loss Pricing Example 169

5.11 Insurance buyer premium ceilings 172

5.12 Distribution specification for insurance contracts offered to buyers . 172

5.13 Expected losses for policies 173

5.14 Technical premium for insurance contracts 173

5.15 Stress Test losses for policies, with Frequency set at $F^{-1}(0.995)$, Severity at $F^{-1}(0.5)$ 173

5.16 Optimal reinsurance purchase fraction for each contract implied by stress test values 174

5.17 Cyber-insurer loss experience in the US market (†denotes weighted average by DWP) Source: NAIC, Researcher calculations 177

6.1 Attack impact on utility 195

6.2 Model actions 196

6.3 Model reward structure 197

6.4 Machine Specification 199

6.5 Network transition probabilities 200

6.6 Sensitivity analysis: number of simulation steps until terminal state reached varying single probability variable, holding others constant at 0.5 203

6.7 Sensitivity analysis: number of simulation steps until terminal state reached varying all probabilities to p with different network sizes. . . 203

6.8 Sensitivity analysis: number of simulation steps until terminal state reached varying number of domain controllers in network, 5 fileservers, 100 workstations. 204

6.9 Simulation Specifications 207

6.10 Simulation Results - average wealth 209

6.11 POMDP specifications for insurer-specified ransomware strains . . . 214

6.12 Poisson parameters for ransomware strains 215

6.13 Results of insurance claim simulation 216

Introduction 1

The adoption of technology by society has transformed computers from curiosities to useful tools to essential staples of human existence. This 'digitization' of society has brought many benefits: information transmission and retrieval take mere fractions of a second in many cases; people can communicate seamlessly across the globe at barely any cost; and global commerce has been transformed courtesy of online shopping. However, while much of the technological innovation has been harnessed for good, its potential for misuse also presents many significant challenges. The field of information security (often also known as cybersecurity), which aims to articulate and mitigate against these threats, has seen rapid development since internet use expanded through the mid-to-late 1990s and now comprises many different disciplines. This thesis is concerned with the economic aspects of information security and in particular the financial costs of degradation in information security.

1.1 Background

In the early days of computer use, information security was a minor consideration outside highly sensitive sectors such as the military. In its infancy, the internet was largely geared towards collaboration and communication with only modest consideration of security[1]. Early users of the world wide web may recall that it was not uncommon for the individual hobbyist webmaster to publish their postal addresses on their website. Prior to the internet, systems communicated on closed or semi-public networks, providing a quasi-physical barrier. Maintaining physical

[1] Not all users were unconcerned about privacy; Chaum proposed a transaction system under which the user would not be identified in the mid-1980s [58] while Zimmerman developed Pretty Good Privacy (PGP) in the early 1990s as described in [308]

media confidentiality would ensure relative security and viruses were largely intended to cause a nuisance, or developed as a proof-of-concept, rather than for revenue generation. Of course, in the days of the modem and dial-up or dial-in access, there was always the risk of a social engineering attack[2], but this required persistent effort.

In the Web 2.0 era [217], systems may be distributed, hosted in 'clouds', and linked to or even directly accessible via the internet. A company anywhere in the world can be held to ransom by an agent located anywhere with an internet connection, with the ransom payments facilitated by cryptocurrencies, which allow pseudonymous monetary transfers to take place outside of the conventional financial system. Known as 'ransomware', these attacks have moved from being an academic concept introduced by Young and Yung in 1996 [304] to a serious threat to corporate profitability, to healthcare, and to education. Connectivity is now an operational requirement for many individuals and industries simply for daily life, but the defensive calculus has accordingly shifted from how to *stop* all attacks with certainty to how to *limit* the damage.

1.1.1 The development of information security as a field of research

There has been interest among economists in the economics of information since the 'dotcom boom' of the late 1990s [261]. In 2001, Ross Anderson argued at the Seventeenth Annual Computer Security Applications Conference that many information security problems can be explained more clearly and convincingly using the language of microeconomics than via technical measures [10]. Anderson's arguments arguably helped generate significant interest in the development of the field of information security economics. In 2002, Lawrence Gordon and Martin Loeb proposed a mathematical model for evaluating the optimal amount of investment in information security [128], that has become known as the Gordon-Loeb model and is now generally as a foundational contribution in the use of economic theory to capture security trade-offs and forms the focus of Chapter 3. If Anderson articulated the qualitative issues associated with the analysis of information security, Gordon

[2] Kevin Mitnick deployed this with great success, achieving significant notoriety among law enforcement [200]

and Loeb played a pivotal role in demonstrating how quantitative methods might be used to analyze security problems.

1.2 Systems and security

As terms such as cyber-incident, cyber-attack, and data breach enter the mainstream vernacular with increasing frequency and form news headlines, the need for a disciplined and sensible analysis of information security risks has never been greater. Generalization can lead to sensationalizing of threats, potentially clouding decision-making by those tasked with managing information security. Economics can help support decisions related to security by providing concepts and mathematical techniques for understanding and framing losses associated with a degradation in the intended security posture of a dataset or system. The tools of the economics discipline are most powerful when their application and range of parameters are clearly identified and specified and sources of potential uncertainty clearly defined.A particular challenge of cybersecurity modelling is that technology and systems are not homogenous and constantly evolve. This means that modelling certain systems and risks cogently and coherently is challenging, but by no means impossible.

When aiming to model systems at a large scale, it is rarely possible or advisable to capture every feature of every individual component — in essence, building a complete replica of the ecosystem of interest. Rather, the task of the modeller is to identify the necessary set of features that allow for the desired outcomes of the analysis to be generated. For applications such as insurance pricing, developing a depiction of systems that allows for probability distributions to be generated or fitted is a useful goal, particularly if the depiction allows for comparison between different companies and their security postures.

The intended contribution of this thesis is to demonstrate the application of economic techniques and methodology to a range of security problems. This section provides a brief overview of fundamental concepts in analyzing and describing systems and their security and how these related to an economic analysis of information security. This supplies the core foundations upon which the work in this thesis builds.

1.2.1 Core concepts in security

In security, it has become commonplace to frame problems in terms of *confidentiality*, *integrity* and *availability* since work by Anderson in the early 1970s [9] (see, for example, [135] for further historical background). These concepts may be explained as follows

- *Confidentiality* — only authorised or intended parties may view the information
- *Integrity* — the information must not be altered from its intended state
- *Availability* — the information must be accessible on demand

One can further reduce the set of parameters to *criticality* and *sensitivity* as described in, for example, [236]. Criticality maps to availability, sensitivity to confidentiality, and integrity is the intersection of criticality and sensitivity. The property of having as few parameters in a model is known as parsimony; using just the necessary set of parameters to describe a system may be viewed as a desirable modelling objective. These security parameters are useful for economic models as they may be defined across the interval $[0, 1]$ and be intuitively understood while aiding the use of models by enforcing clear bounding. This is helpful where the monetary value of these relatively abstract concepts is hard to instantiate.

1.2.2 Describing systems

As computer systems have become more complex, significant research endeavour has been devoted to producing mathematically rigorous descriptions of systems. In particular, there has been significant output in the field of logic. The literature is far too rich to cover in any depth within the scope of this thesis, but [199] is recommended as an excellent introduction to the use of logic to describe processes, while [67] gives a more contemporary overview of systems modelling with practical examples. Further suitable references are contained within Chapter 4 in Sections 4.4 and 4.7.1. In the context of this thesis, it is important to acknowledge the existence of rigorous systems modelling and that it can be compatible with economic modelling. However, detailed systems modelling is used only lightly in Chapter 4. Combining systems depictions and economic models for any of the other chapters would in itself be a significant, multi-year research endeavour and is consequently outside the scope of what might reasonably be accomplished within this work.

1.2.3 Measuring security

In a simple system, security may be readily defined and enforced in absolute terms. However, for most real world applications, security is not a binary. For some very sensitive applications, such as the military, security might be afforded the highest priority. For others, certain elements of an operation or data set may require stringent security, but for services where availability and ease of use is a greater concern, enforcing complex or prescriptive security procedures may not be a desirable objective.

The concept of measuring security might be construed as trying to form an assessment of how vulnerable a system is as measured against a set of criteria. This takes many different forms; examples might include penetration testing of an organizational network, formal methods and verification of computer code or configuration files, or behavioural sampling. Security standards such as ISO27001 try to provide a structured set of controls across different domains that must be fulfilled to achieve certification to the particular standard [49].

Measurement of security is important, as it allows those reliant on it being maintained to have confidence in the measures deployed to implement security and to guard against 'window dressing'. One of the most well-know explanations of the phenomenon of signalling security without actually enhancing security is 'security theater', a phrase coined by Bruce Schneier in 2003 [257]. The ability to distinguish policies and processes that genuinely improve security rather than merely giving the impression of so doing is particularly critical for industries such as cyber-insurance, which rely heavily on client declarations of security posture. Usually, companies will undergo some form of audit, which is designed to convey assurance of controls. However, as discussed in Chapter 28 of [11], auditors usually rely on declarations of 'good faith' or to the 'best of knowledge'. Almost every major 'cyber-incident' exposes some weakness in underlying processes, and it is arguably very important for company management and boards to take external advice to expose potential shortcomings in internal security assessments.

1.3 Economics

Economics provides a structured means of analyzing competing incentives between agents, evaluating the strategies and trade-offs in the case of game theory, and determining whether an equilibrium can exist. It also allows for rigorous mathematical solution of problems under certain circumstances where the preferences of agents may be modelled using so-called utility functions. A technical, systems-centric perspective of security need not compete with an economic perspective; the two analyses can be complementary and help explain different elements of the problem.

As systems and their connectivity become more complicated, building a 'bottom-up' model becomes inefficient and some level of approximation and/or abstraction is likely to yield a more tractable model. Technical controls are of course important and the first line of defence, but vulnerabilities (as they are known in the field of information security) are techniques that allow an attacker to circumvent the intended controls to manipulate a system away from its intended state. If security were perfect, then cyber-security practitioners need not exist. However, this is far from the case. It may not be possible to defend every attack surface in a system fully and then economic analysis can help order the allocation of resources to give the best defence possible subject to relevant constraints.

It is also worth considering the management structure and decision making process of organizations. In a typical large organization, the board of directors will have overall strategic oversight and responsibility for security. Increasingly, many organizations employ a chief information security officer (CISO). The CISO acts as a conduit between the organizational management team and the technical personnel responsible for monitoring and delivering operational information security. Therein lies a potential divergence of knowledge and experience; it is not guaranteed that organizational management will be well versed in the technical aspects of information security. A well-developed economic model may assist optimal decision making by structuring the information in a manner more familiar to executives and presenting a clear range of investment options, objectives, risks and trade-offs.

1.4 Cyber-insurance

As awareness of information security risks developed and companies appeared willing to spend money to protect against the financial losses associated with cyber-attacks, underwriters in the insurance market spotted an opportunity to launch a new product: cyber-insurance [132]. Cyber-insurance allows for individuals or companies to pay a premium to be reimbursed, either in money or services, for damage or losses associated with an information security event. Initially, cyber-insurance proved profitable for insurers as while the threat landscape developed, information security incidents did not cause severe economic losses for companies. Insurance typically uses a rigorous claim process, enforced by the rule of law. Accordingly, unless the contract provides otherwise, costs must be tangible. This leads to a perception by some that cyber-insurance is unlikely to ever pay out; this is not the case, but the information exchange between purchasers of cyber-insurance and the insurers who write the policies is certainly complex and asymmetric, which presents an interesting if not unique modelling challenge among insurance perils.

The cyber-insurance market has seen rapid growth in recent years, especially in the United States where the National Association of Insurance Commissioners estimates total premium written has grown from $1.4bn in 2015 to $6.5bn in 2021 [210][3]. Cyber-insurance is particularly interesting as a research topic as it is relatively new as a line of cover and past data on claims is not necessarily as comparable as for other, more established lines of business. Much of the content of this thesis is devoted to the analysis of problems relevant to cyber-insurance and how economics can help to suggest solutions to them. We now briefly sketch a few broad considerations as to the motives and incentives of buyers and sellers of cyber-insurance.

[3] Some caution is required with these figures as the premium income is a function of both the policy rate charged and the amount covered.

1.4.1 Demand

It is important to note that there is no universal cyber-insurance policy that covers all outcomes. Commercial cyber-insurance policies are highly tailored and specific contracts that cover a range of clearly stated perils up to a stated limit [88, 247]. This is no different to any form of insurance; even mass-market consumer insurance is tailored to the preferences and risk tolerance of the buyer [64, 86]. The diversity of insurance policies reflects the different requirements of the various buyers. A small firm with fairly simple requirements is unlikely to possess extensive in-house cybersecurity capabilities or have consultants on retainer. In the event of a suspected data breach, a cyber-insurance policy will connect the policyholder with appropriate assistance to diagnose, remedy, and prevent future recurrence of the breach [295]. Larger companies may have dedicated cybersecurity functions or outsource these to an external provider. For these entities, cyber-insurance largely serves as a means of balance sheet management — in the event of a serious breach of information security, coverage for investigations by specialist (expensive) consultants and subsequent remediative measures might help mitigate the loss of revenue the company would otherwise face. The extent to which a company assumes versus indemnifying risk will be commensurate with its level of risk aversion.

1.4.2 Supply

Supply of insurance comes from firms who believe that the activity will be profitable — the value of premiums written will exceed the value of claims and other expenses [63, 6, 72]. Usually, insurance pricing is determined via probability distributions derived from past experience, parameters or other data [229]. The field of risk quantification in this manner is known as actuarial science. For cyber-insurance, the field is relatively new and faces several unique challenges related to the evolution of technology and the inter-connectedness of the internet removing the geographical element of diversification that insurers usually rely on for lines such as flood insurance. The particular challenge for cyber-insurers is how to assess the likelihood of a cyber-insurance customer claiming on their policy, especially if the customer itself is not able to fully articulate its own security posture. This thesis will explore some possible approaches to answering this question using models and model frameworks.

When insurers are uncertain or uncomfortable with the totality of risks they have underwritten, they may rely on reinsurance — insurance on insurance [53]. This can act as a simple risk transfer mechanism or as a portfolio management tool, which provides cover for losses in excess of certain pre-defined characteristics. However, the dynamics of information exchange are important for a reinsurance market to be sustainable and the aforementioned difficulties for cyber-insurers relating to risk analysis apply equally to reinsurance providers.

1.5 Thesis structure

The remainder of this thesis is organized as follows. Chapter 2 presents a comprehensive Literature Review of the cyber-insurance market and the economic methods needed to support analysis of it.

Chapter 3 introduces the Gordon-Loeb (GL) model for investment in information security. The GL model is used as the basis for an expected utility maximization problem for deriving the optimal combination of cyber-insurance and defensive security investment under exogenously observed premia. This demonstrates how given a few parameters, the case for competing allocations of capital to investment versus insurance might be evaluated for a number of different outcomes.

One issue with modelling approaches such as the Gordon-Loeb model is the estimation of appropriate parameters in the absence of clear priors. Chapter 4 proposes a methodology for describing an organization and its systems using entity relationship diagrams, calculating the maturity of the security posture of the organization and using this assessment to deliver parameters that can be used to populate the utility function of an insurance company providing cyber-insurance. This addresses the criticism of excessive abstraction from real world technical considerations sometimes levied at economic approaches to information security modelling. Chapter 5 examines the difficulties of achieving efficient information exchange in the cyber-insurance market. The analysis concludes that only with statutory public reporting of losses from information security incidents could pricing of cyber-insurance be efficient. Inefficiency does not imply that transactions cannot or should not take place, but suggests that there is a risk of financial imbalances building. Chapter 6 illustrates how a information security specific risk might be modelled, describing a partially observable Markov decision process (POMDP) model of ransomware spreading through a model organizational network.

1.5.1 Relation to prior published work

Chapter 3 in based on [267] and Chapter 6 on [269]. Chapter 4 is adapted from [266] and Chapter 5 [265], both of which were under submission as separate articles to *The Journal of Cybersecurity*, at the time of thesis submission. A statement of contributions to these papers is provided in the Appendix. The literature reviews from all the aforementioned papers were combined to form Chapter 2.

1.6 Thesis Contributions

This thesis contributes to the academic literature on security economics and cyber-insurance, but the three models and single modelling framework contained within it also have relevance for insurance companies, managers of enterprise information security, and policymakers. It aims to bridge the gap between a systems-focused and economic view of security by showing how these disciplines can be complementary rather than competing in an analysis of security.

The literature review in Chapter 2 is intended to provide an overview of the body of work that should be read to gain a thorough grounding in the discipline of insurance economics. The review of the cyber-insurance literature is aimed to be comprehensive and the categorization imparts structure on a field that can often appear lacking in organization [101].

When published as [267], Chapter 3 was the first example of an expected utility model for cyber-insurance pricing using the Gordon-Loeb model. It provides a credible use of a well-established model for the vulnerability of an information set to breach to price cyber-insurance in a classical sense. One particularly useful feature of the Gordon-Loeb model is its flexibility. Parameters such as the inherent vulnerability of the information set to a breach or expected loss from a breach could be taken from other models. Consequently, the model imparts a degree of structure to framing decision problems that might otherwise become fairly complex.

Chapter 4 addresses a long-standing research challenge of taking serious account of system structure in pricing cyber-insurance policies. No single framework let alone model could ever price a cyber-insurance policy on its own. However, the modelling framework developed combines well-established modelling technologies in the form of entity relationship diagrams as a language, security maturity models as a measure, and utility functions as a pricing model. This could easily be mapped to insurance company pricing factors as described in [247]. Informal discussions with individuals from three insurance companies and one insurance broker have suggested that the modelling framework may be of use by practitioners. Currently, the insurance industry is heavily reliant on questionnaires and there is an apparent lack of standardization. At the very least, the modelling framework proposed acts as a template for discussions on how pricing of cyber-insurance might be improved. If deployed on real-world organizations, the modelling framework might be able to deliver outputs that could be used in the model developed in Chapter 3, acting as a support to decision-makers.

Chapter 5 represents a novel contribution in considering the interaction between reinsurance and cyber-insurance. According to regulatory sources, almost half of all cyber-insurance premium written is ceded to reinsurers [210]. Yet, to date, reinsurance has received barely any attention in the cyber-insurance focused literature. Reinsurance has been cited as an important factor in the liability insurance crisis of the 1980s [291, 28]. Given the growth in the cyber-insurance market [210], the interaction between cyber-insurance and reinsurance merits consideration. The nature of cyber-threats and technological development presented in the chapter suggests that, according to established economic theory, *ex post* efficiency will be almost impossible to achieve in the cyber-insurance market. This is supported by analysis using simulations of a fictional cyber-insurance market, which has been carefully constructed to be representative of the existing cyber-insurance market.

There have been a number of empirical studies and game theoretic analyses of the cyber-insurance market as reviewed in Chapter 2 but very few Monte Carlo-type simulations. This has broader relevance for insurance modelling beyond cyber-insurance alone. The conclusion of the analysis of the cyber-insurance market is that better co-ordination and public information sharing is needed. This demonstrates the need for the modelling framework presented in Chapter 4 and illustrates why it is not just of theoretical interest.

Chapter 6 is somewhat different in nature from the preceding chapters in modelling a specific rather than generic attack. Ransomware has been modelled using game theory (see Section 6.3.1) but the outputs of these models are not particular helpful for insurance loss estimation. The POMDP model introduced in the chapter can produce loss estimates for different ransomware strains on different networks as illustrated in Section 6.5.3. The output of the ERD modelling framework from Chapter 4 could be used to deliver parameters to the POMDP model and produce loss estimates that might be used in the GLCI model in Chapter 3. This in turn might be used to assess the need for reinsurance per the modelling in Chapter 5.

Literature Review 2

This chapter contains a review of the literature relevant to this thesis. It begins with an overview of the classical economic literature devoted to or useful for an analysis of insurance decisions. This is aimed to present the foundational economics used within the models that follow in subsequent chapters and should not be viewed as a comprehensive analysis of a rich field of work, particularly on asymmetric information.

The literature on security investment models is then reviewed, followed by a systematic review of the literature on cyber-insurance. This was initially conducted in late 2019 but has been updated with relevant subsequent contributions where possible. Finally, the literature on modelling ransomware attacks is reviewed. This should be regarded as a specific case study of a threat that is highly relevant to cyber-insurance.

2.1 Insurance economics

The benefits of insurance in spreading risk away from the individual towards society were described by Adam Smith in the late 18th century [272]. The modern discipline of insurance economics was arguably established by the work of Borch, Pratt, Arrow and Mossin in the 1960s, following von Neumann and Morgenstern's expected utility theory [211]. This was followed by key developments in the 1970s with regard to asymmetric information, particularly the celebrated contributions of Akerlof [5], Spence [274] and Rothschild & Stiglitz [248]. A literature detailing the theory of insurance supply also subsequently developed. The coverage of insurance economics in this thesis aims to summarise some notable contributions to the literature rather than representing a complete survey of a what is a diverse and well-established field.

2.1.1 Expected utility and the theory of insurance demand

The economic concept of utility, essentially the mathematical formulation of preferences or behaviours, is fundamental to a quantitative analysis of insurance markets. Expected utility was first introduced by Bernoulli in the 18th century [110]. In classical economics, expected utility is usually descriptive rather than normative[1]. Von Neumann and Morgenstern introduced an axiomatic version of expected utility theory [211]. The essence of their argument is that it is particularly hard to describe utility as a number and they assume that "the aim of all participants in the economic system... is money". In rudimentary terms, their proposition is similar to a notion in physics that while certain fundamental properties of nature such as mass and charge can be readily defined in theoretical terms, their properties are most apparent and readily understood in an experimental sense. The axioms they propose for a system of abstract utilities are shown to be interpretable as one of numbers up to a linear transformation. Von Neumann-Morgenstern utility functions form the basis of the theory of insurance demand.

It should be noted that the expected utility hypothesis is not universally accepted: the Allais and Ellsberg paradoxes provide noted counterexamples [7, 93]. One of the most famous critiques of expected utility theory known as prospect theory was introduced by Kahneman and Tversky [154]. The core idea of prospect theory is that "choices among risky prospects exhibit several pervasive effects that are inconsistent with the basic tenets of utility theory." In particular, Kahneman and Tversky argue that people underweight outcomes that are merely probable in comparison with those that are obtained in certainty; they develop a theory that assigns value to gains and losses rather than to final assets and in which probabilities are replaced by decision weights. The additional versatility of prospect theory is likely to prove important in modelling cyber-insurance, where the loss function is still primarily monetary but has an additional dimension in the form of loss of information. This adds additional complexity to the problem.

[1] A normative model is one which dictates rather than describes the behaviour of an agent

Pratt introduced $r(x) = -u''(x)/u'(x)$ as a measure of local risk aversion, where $u(x)$ is a utility function for money and $u'(x)$ and $u''(x)$ are the first and second derivatives of the utility function, respectively [234]. $r(x)$ is often known as Arrow-Pratt risk aversion given contemporaneous work by Arrow [14]. Mossin analyzed four different problems in terms of the wealth effect on the propensity to take insurance coverage: the maximum acceptable premium for full coverage, optimal reinsurance quota, the optimal coverage at given premium, and the optimal amount of deductible — an amount of loss below which no claim is paid by an insurer [203]. These are foundational to the theory of insurance demand and are collectively sometimes called the Mossin Theorem. Arrow considered optimal insurance and generalized deductibles [15]. He demonstrated that a risk averse buyer will prefer a policy offering complete coverage beyond a deductible. This form of contract effectively places a cap on the loss of wealth an individual may incur.

Borch addressed the issue of insurance pricing under incomplete information or analytical methods [40]. He argued that the ends or objectives of an economic analysis of insurance ought not to be subservient to the means of analysis available, in essence that data quality or methods need not preclude transactions. Borch argued initially that insurance should be considered using the principle of equivalence, from which the insurance premium an agent is willing to pay should be equal to the sum of expected claim payments and administrative costs. He then expands the simple principle of the equivalence model to multiple contracts, suggesting that the choice of market premium ultimately implies a choice of profit distribution. This choice of premium is a subjective decision and will depend on the insurance company's objectives. The problem may be reduced to the task to maximizing a mathematical expression after reformulation using expected utility. Finally, Borch discusses some of the issues involved in applying traditional economic analysis to insurance and proposes an equilibrium price in which total insurance supply would equate to total insurance demand. However, under classical economic theory, there is no natural unit of insurance cover from which to define a price.

2.1.2 Information asymmetry, adverse selection and moral hazard

A key development in modern economics is models incorporating asymmetric or imperfect information, which allow for a more realistic and versatile depiction of many real world problems. For insurance markets, adverse selection and moral hazard are two widely studied problems in this domain. In simple terms, adverse selection is the risk that an insurance buyer takes advantage of their personal knowledge of their circumstances to which the insurer is not privy; moral hazard is the risk that possessing insurance encourages risky behaviour.

One of the most important contributions in understanding asymmetric information,*The Market for Lemons*, was introduced by Akerlof in 1970 [5]. Akerlof introduced a structure for determining the economic costs of dishonesty, which provides the foundation for analysis of adverse selection in insurance. Akerlof's model relied upon linear utilities to avoid algebraic complication but also to allow clear focus on the asymmetry of information rather than endogenous factors such as the treatment of risk aversion inherent in a concave utility function. The analysis of the used car model Akerlof uses to illustrate his theory highlights the connection between price and quality: if a market contains sufficient inferior goods of lower price, and the buyer is willing only to pay the lower price for fear of being sold an inferior goods at a higher price, the inferior goods drive out the superior good. Akerlof uses the example of the over-65 health insurance market, arguing that this group has difficulty in buying health insurance, but that the price does not rise to match the risk. The reason given for this is that as the price rises, only those in need of the insurance will take it out; that is, the quality of the applicant moves in inverse proportion with the price. This has the potential result that no sale may take place at any price. This principle is readily applicable to many insurance markets, and has clear relevance for cyber-insurance. Spence introduced the idea of signalling within the context of the job market [274]. His idea was that job candidates will possess certain characteristics such as a college degree, which signals to employers that they have a capacity to learn. Any candidate could claim that capacity to learn, but there is then an information asymmetry between candidate and prospective employer; the college degree acts as a signal to resolve the information asymmetry. This concept is particularly valuable in the context of cyber-insurance where the poorly protected might claim otherwise to try to lower insurance costs; however, clear evidence of preventative measures such as firewalls or information security policies might act as a signal in this instance.

Rothschild and Stiglitz considered competitive markets in which the "characteristics of the commodities exchanged are not fully known to at least one of the parties to the transaction." [248]. The key insight from this paper is that when a competitive equilibrium does exist, they may have strange properties compared with a more traditional sense of equilibrium. In an insurance market, a consumer is not offered a price at which they can buy all the insurance they desire; rather, they are offered a quantity and a price. Rothschild and Stiglitz argue that high risk individuals cause an externality as low risk individuals are generally worse off as insurance consumers than they would be in the absence of the high-risk group. However, the high-risk group are indifferent to the existence of the low-risk group. Rothschild and Stiglitz are able to show that under some circumstances, a competitive insurance market may have no equilibrium. Wilson also found that no stationary equilibrium may exist if all firms have static expectations with regard to the policy offers of other firms [290]. However, under a different policy rule in which any policy is immediately withdrawn that become unprofitable after that firm makes its own policy offer, the equilibrium is found to exist.

Moral hazard as it relates to the improvement of contracts has been studied by Holmström [144]. He argues that by creating additional information systems or by using other available information about the agent's action or the state of nature, contracts can generally be improved. A particular relevant point for further analysis raised in this work is that in a long-term relationship, the propensity for moral hazard is decreased as if an agent repeatedly behave recklessly, their insurer will soon recognise this and their premiums will commensurately increase upon renewal. Lee considers how the problem of moral hazard might be solved by provision of a loss-preventative good [171]. For cyber-insurance, an example would be the government providing anti-malware software to the population. Moral hazard may also be addressed via the use of deductibles in insurance contracts [82].

2.1.3 Insurance supply and pricing

A key contribution in explaining the supply of insurance was made by Raviv, who noted that in the earlier model proposed by Arrow, it is unclear whether the optimal insurance policy with a deductible is due to risk neutrality of the insurer, non-negativity of insurance coverage or loading on the premium [239]. Raviv proposed a solution to this question via a general formulation of the insurance problem, which embedded previous models such as those proposed by Borch and Arrow. Raviv found that the cost of insurance could be shown to be the driving force behind the deductible results proposed by Arrow. He showed that the Pareto optimal insurance policy involves a deductible and coinsurance of losses above the deductible. If the cost of providing insurance is independent of the insurance contract, then the Pareto optimal contract does not have a deductible.

Borch developed a model to investigate regulation and supervision of insurance companies, finding that if a company is interested solely in making a short-term quick profit, then regulation is needed [41]. However, if the management of the company take a long-term view, no regulation should be necessary. Borch also shows there are limits to what a government can achieve by regulation of private insurance companies which operate in a free economy. Munch and Smallwood examine the case for solvency regulation in the property and casualty insurance industry, noting that the case for solvency regulation derives from the difficulty of a policyholder establishing the financial soundness of alternative firms [205] . However, firm owners are also at risk as they may lose their entire equity in a firm whereas the insurance buyer may just receive partial coverage. Finsinger and Pauly argue that beyond an assumption of consumer ignorance of risk of insurance company default two further assumptions are necessary to justify regulation: if not regulated, firms will hold reserves below the socially optimum level and regulators can determine and enforce a level of reserves that is closer to the social optimum than the unregulated level [109].

2.1.4 Reinsurance

The following series of works represent significant contributions on understanding reinsurance, but is not exhaustive. As a general overview, Dionne is an excellent collection of important papers related to reinsurance and includes many key contributions to the field [80]. Within this, Borch is of particular relevance to this work, focusing on describing the conditions required to achieve equilibrium in a reinsurance market via generalizing the classical theory of commodity markets to include uncertainty [97].

Kaluszka et al introduce a more contemporary example of optimal reinsurance treaty pricing derived from the classical literature on the subject [155]. The work suggests that stop loss and truncated stop loss contracts are the optimal structure for maximising utility, the stability and the survival probability of the cedent for a fixed reinsurance premium.

Schlesinger and Doherty provide a useful treatment of issues associated with incomplete insurance markets, in particular suggesting that focusing on correlation of risks is essential for making use of incomplete markets theory [256]. This is an argument as to why an insurer who does not currently offer cyber-insurance might enter the market should it believe that cyber-losses will not be highly correlated with areas in which it currently has exposure. Empirically, there is concern of hidden or 'silent' cyber-risks within existing lines, meaning that for many insurers offering cyber-insurance could be utility detracting.

Froot and O'Connell discuss the pricing of US catastrophe insurance with some illustrative data [113]. They find that price increases and quantity declines are more pervasive than they should be within catastrophe reinsurance based on fundamental data; this is strongly suggestive of historical inefficiency.

Aase discusses the Nash bargaining solution in relation to the competitive equilibrium allocation for a reinsurance syndicate and finds that in certain cases, a first order Taylor approximation of the two solutions may be equivalent [1]. This has relevance for comparing potential game theoretic representations of reinsurance.

Mata provides a theoretical methodology to calculate the distribution of total aggregate losses for two or more consecutive layers when there is a limited number of reinstatements [187]. In the reinsurance market, a layer is a specific proportion of losses a reinsurer will cover and reinstatements are the number (as oppose to magnitude) of loss events covered by a policy. This thesis only considers single or aggregate losses, but the pricing methodology proposed by Mata would be useful for pricing more complex structures such as those with specific triggers or event definitions.

2.1.5 Relevant contemporary work

The field of behavioural economics aims to combine psychology and economics to model the behaviour of agents endowed with human characteristics [147, 204, 12]. This has potentially interesting applications to cyber-security decision making, where machines behave according to a prescribed set of instructions but controls must be adhered to by humans who are either susceptible to errors [150, 237, 213] or may be malicious actors [202, 262, 178]. Within the area of contracts, Kőszegi, reviews the behavioural economic literature of relevance to contract theory and accordingly insurance. The issue of informed principals —- where the principal has superior information regarding a variable that affects both their incentive-design problem and the agent's willingness to exert effort — discussed by Kőszegi among other topics related to asymmetric information is relevant for cyber-insurance. While directed towards labour market concerns, the model of Fang and Moscarini [102] where the principal receives a private signal about each agent's ability and can decide whether to make different contract offers is closely aligned to the work of an underwriter in insurance (the insurance ecosystem is discussed in Chapter 5).

The issue of preference heterogeneity in insurance has been studied by Cutler et al [71]. The authors find that in annuity and acute health insurance markets, higher-risk individuals have more insurance, as classical theory would predict. However, in life insurance and long-term care insurance, they find that "advantageous selection" is evident where lower-risk individuals have higher coverage. This may suggest that lower-risk individuals have higher risk aversion and thus are more likely to seek insurance; in cyber-insurance, it is possible that high-risk entities may be restricted from full coverage due to capacity constraints [299].

Gollier et al comprehensively review the literature on decision making under uncertainty in the post-expected utility hypothesis era [126]. From this survey, the findings of Eeckhoudt & Gollier that the intuitive idea that risk aversion should always increase self-protection is mistaken [84] has implications for a decision-maker balancing security investment versus cyber-insurance. The survey also contains an interesting discussion on higher-order derivatives of utility functions that may have future relevance for cyber-insurance decision problems that might be modelled using more complex utility functions than the standard formulations deployed in this thesis.

2.2 Security investment models

2.2.1 The Gordon-Loeb Model and some alternatives

Gordon and Loeb introduced an economic model in 2002 that determines that optimal amount to invest to protect a given set of information [128]. The Gordon-Loeb (henceforth GL) model is discussed in full detail in Chapter 3, but its most important contributions are presented here for comparison with other relevant literature. The key result of the GL model is that investment should not exceed more than 37% of the expected loss. Gordon and Loeb introduce the concept of a security breach function with three key assumptions:

1. If the information set is completely invulnerable, it will remain perfectly protected for any security investment;

2. If there is no investment in information security, the probability of a security breach is the inherent vulnerability of the information set;

3. As investment in security increases the information is made more secure but at a decreasing rate.

The GL model is conditioned using security breach functions (SBFs) that are linear (class I) and exponential (class II) in the inherent vulnerability of the dataset. These security breach functions are a function of investment in information security, and provide a route to quantification of the risk reduction provided by such investment and as such the Gordon-Loeb model may be considered a security investment model.

The GL model laid the foundations for a rigorous quantitative structured analysis of information security investment problems. The two types of breach function introduced are intuitive to understand and fairly simple to manipulate, which is a distinct advantage of the model. Böhme gives a good summary of security investment models, their terminology and parameters [35]. Huang and Behara likewise provide an excellent summary of the various security models and derive similar security breach functions to Gordon and Loeb, albeit via a mathematically more sophisticated route [146]. While this approach might be regarded as superior by the more mathematically inclined, it is not necessarily superior to the approach taken by Gordon and Loeb as the GL model is arguably more intuitively accessible to a broader audience.

2.2.2 Criticisms of the Gordon-Loeb Model

There have been examples in the literature of attempts to disprove the Gordon-Loeb optimal security investment. The first of these is due to Hausken who provides a counter-example via the use of a logistic function but with quite a few changes to the original Gordon-Loeb assumptions [138]. Willemson disproves the conjecture by also showing investment up to 50%; upon relaxation of the original requirements he shows that with the Gordon-Loeb framework, levels of close to 100% investment can be achieved [288]. With any simple mathematical model, it is relatively straightforward to engineer a counter-example and these prove useful in understanding the limitations of the Gordon-Loeb model. The key advantage of the Gordon-Loeb model is the balance it strikes between rigour and simplicity while offering useful insights into how to consider security investment. Baryshnikov aims to counter the assertion made by some critiques of the GL model that the $1/e$ rule of investment does not hold in generality [25].

2.2.3 Extensions of the Gordon-Loeb Model

A body of literature has developed evaluating potential empirical uses of the Gordon-Loeb model and the calibration of its parameters. Matsuura proposes a productivity space of information security, specifically considering a productivity regarding threat reduction and a productivity involving vulnerability reduction [188]. In essence, this might be regarded as an extension of the original Gordon-Loeb model to a two-dimensional case. Tatsumi and Goto add a timing dimension to the original Gordon-Loeb model using a real options approach [280]. Gordon, Loeb, Lucyshyn and Zhou extend the original Gordon-Loeb model to include costs associated with the externalities of security breaches rather than just the firm's private costs [129]. The revised model is sometimes referred to in the literature as the GLLZ model.

Farrow and Szanton propose extensions to the Gordon-Loeb and GLLZ models, based on mathematical equivalency with a generalized homeland security model [104]. Gordon, Loeb and Zhou explain how the Gordon-Loeb model can be used in a practical setting and the intuition underpinning the model's parameters [130].

Young et al use the Gordon-Loeb as the foundation for setting up an insurance problem involving minimising a linear combination of expected loss, security investment and insurance premium [305]. Naldi and Flamini provide a thorough investigation of the productivity parameters in both classes of Gordon-Loeb security breach functions and propose estimators for these parameters [209]. Mazzoccoli and Naldi [189] expand upon the work of [305] by providing closed form solutions to the same problem.

2.3 Systematic cyber-insurance literature review

This section presents a systematic review of the literature on cyber-insurance. Papers are classified into the following categories: Economic Modelling, Frameworks and Policy, Game Theory, Law and Surveys and Literature Reviews. This classification is necessarily subjective - for the game theory category, papers are selected where the abstract explicitly references games or the paper structure clearly implies this was the main purpose of the analysis. Economic modelling papers are generally classified as such where the papers follow a risk management or optimization-type narrative. Actuarial models covers literature that identifies its contribution as in the domain of actuarial science, which is the rigorous mathematical modelling of financial risks. There is similarity between the actuarial/economic model classifications and there is not intended to be a strong distinction; classification in this literature review is informed in part by publication venue. Frameworks and policy aims to categorise those papers which provide either a qualitative means or risk classification or normative discussions on cyber-insurance relevant policy. Law is self-evident. Surveys and empirical analyses encompasses questionnaire based research, market analysis and 'state of the field' type analyses. Cyber-insurance linked securities and capital requirements are specific categories with direct relevance for insurance companies. Table 2.1 summarizes the surveyed literature by category.

Classification	**Relevant Literature**
Economic modelling	[39, 222, 263, 223, 253, 224, 232, 264, 116, 158, 215, 233, 24, 34, 100, 160, 185, 21, 300, 37]
Frameworks and Policy	[127, 36, 219, 218, 259, 191, 157, 176, 170, 95, 115, 159, 231, 287]
Game Theory	[151, 301, 225, 302, 139, 79, 186, 307, 107, 106, 245, 296, 226]
Law	[140, 212, 278]
Surveys and empirical analyses	[32, 89, 281, 112, 182, 193, 246, 294, 271, 87, 92, 91]
Actuarial models	[29, 141, 306]
Cyber-insurance linked securities	[161, 156, 43]
Capital requirements	[90]

Table 2.1: Summary of systematic literature review on cyber-insurance

2.3.1 Economic modelling

Bojanc et al outline a variety of different economic techniques that could be used for information security risk management; they discuss cyber-insurance as a potential solution to the problem [39] but note that prior literature raised a cause for concern that some cyber policies may not pay out [181]. Pal et al investigate the problem of self-defense investments in the Internet under full and partial insurance coverage models, finding that cooperation among users results in more efficient self-defence investments and that partial insurance motivates non-cooperative internet users to invest efficiently in self-defense mechanisms [222]. There is some agreement in the literature that cyber-insurance does not necessarily improve network security from a theoretical perspective, though user welfare generally improves [263, 224, 185]. Khalili et al suggest that an insurance company can increase profit by insuring both a primary and associated party and that this reduces collective risk [158]. This seems a counter-intuitive result unless the purchase of cyber-insurance encourages better security, which is at odds with the findings of other papers; this work is expanded in [160].

The literature on pure cyber-insurance modelling is relatively modest in quantity. Pal et al introduce a cyber-insurance model, Aegis, in which a user accepts a fraction of loss recovery on themself and transfers the rest of the loss recovery to a cyber-insurance agency [223]. Bodin et al provide a model for selecting the optimal set of cyber-security insurance policies by a firm, given a finite number of policies being offered by one or more insurance companies [34]. Bandyopadhyay et al build a model to capture the impact of secondary loss in structuring the use of cyber-insurance and then combine the backward analysis of myriad breach scenarios to derive the overall optimal decision to purchase cyber-insurance [21]. This appears an area where there is significant opportunity for further work.

Similarly, the literature on theoretical pricing of cyber-insurance appears particularly sparse and underdeveloped. Saini et al attempt to produce a model for deriving utility functions for cyber-insurance, using a university network as an example [253]. Determining the optimal utility function to describe insurance buyer and supplier behaviour is fundamental in developing a sound pricing model as it established a fair value for risk, making this a useful contribution. Fahrenwaldt et al introduce a polynomial approximation of claims together with a mean-field approach that allows to compute aggregate expected losses and prices of cyber-insurance [100]. However, the limited data publicly available around cyber-insurance would make such a model difficult to validate. Piromsopa et al propose a rudimentary cyber-insurance scoring model, which can incorporate existing security standards - this is most applicable to enterprise risk management [233]. Xu et al propose a three component model based on the epidemic mode, loss function and premium strategy and study the dynamic bounds for infection probability based on Markov and non-Markov models and propose a simulation approach to compute the premium for cybersecurity risk [300]. This is an interesting approach, although the cyberattack model is somewhat simplistic relative to the variety of overall threats.

Böhme et al [37] builds on and updates the analysis in an earlier paper by Böhme and Schwartz [38]. The newer paper develops a cascade model of cyber risk arrival factors, and links these to practical top-down and bottom-up risk management frameworks deployed by practitioners. A model is then developed, which accounts for information asymmetries. The model is then expanded to account for security interdependence — dependence not only on one's own security but also on the security of all connected nodes. The work concludes that a detailed understanding of cyber risk is required for cyber-insurance to be a sustainable line of insurance and that modelling must be predictive rather than reactive and grounded in scientific principles.

2.3.2 Actuarial models

Developing specific actuarial models for cyber-insurance has become a focus through the support of industry funding [17]. Some of the models that have been proposed thus far through this initiative are briefly summarised here, though for brevity technical details are omitted. Bessy-Roland et al introduce a multivariate Hawkes process for cyber-insurance and demonstrate how it can be calibrated using the Privacy Rights Clearinghouse database of data breaches to provide a full joint distribution of future cyber attacks [29] (see also [142] for an application of such modelling to cyber-insurance derivatives). Hillairet and Lopez propose a stochastic diffusion model for estimating the propogation of cyber-incidents within an insurance portfolio [141].

In developing a model for cyber-insurance claims, catastrophic claims are a significant concern. Dassios and Jang use the Cox process[2] to model the claim arrival process for catastrophic events [75]. Baldwin et al use the multi-variate Hawkes

[2] A doubly stochastic Poisson process

process as the basis of a model for estimating contagion in cyber attacks [20]. Bessy-Roland et al introduce a multi-variate Hawkes framework for modelling and predicting cyber attacks frequency across firms following successful cyber-attacks against a subset of the population [30]. Gollier suggests how to design optimal insurance under ambiguous risks [125]. An ambiguous risk is one where the uncertainty is not known and well-understood

Zeller and Scherer propose a statistically motivated model of marked point processes for capturing the dynamic nature of cyber-risk in insurance pricing [306]. While the proposed model of Zeller and Scherer is judged to be well-motivated by expert practioners [243], given the paucity of extant data, in practice the model would be very hard to calibrate and some paramaters may not be estimable. This appears to be a significant limitation at the moment in the cyber-insurance literature as the available data for modelling is many steps behind the current state of theory.

2.3.3 Frameworks and policy

The literature concerning frameworks and optimal cyber-insurance policy is rather better developed than that on the economics of cyber-insurance. The four step decision plan of Gordon, Loeb and Sohail [127] is one of the earliest contributions specifically on cyber-insurance we have identified in the literature. The difficulties surrounding potential risk correlation are well studied: Böhme and Gaurav investigate the potential limits of cyber-insurance in the context of the high correlation of potential risks [36] , which Bandyopadhyay et al develop arguing that cyber-insurance tends to be overpriced as insurers cannot estimate the potential secondary losses of customers [219]. Ogut et al find that firms invest less than the socially optimal level when risks are correlated but that the appropriate social intervention policy to induce a firm to invest at these levels depends on whether insurers can verify a firm's self-protection levels [218]. How the latter would be achieved in practice would depend on regulation.

Shackelford argues that firms should take a proactive stance toward managing cyber attacks implicitly cautioning against over-reliance on cyber-insurance [259]; this perspective is somewhat countered by Laszka and Grossklags who suggest that insurance providers taking a role in helping improve software security can lead to a more profitable cyber-insurance market [170].

In terms of papers describing the field, Linton et al and Keegan summarise research in the cyber-security chain [176, 157] while Elnagdy et al outline the taxonomy of cyber-risks for cyber-security insurance of the financial industry in cloud computing[95]. There are a few papers, which propose frameworks for cyber-insurance, specifically cloud based insurance for big data; organization insurance; and pre-screening and security interdependence [115, 231, 160]. As with economic modelling, contract and cyber-insurance design is an important field for developing a functioning market and likely merits further work. However, the field remains underdeveloped and as such critical evaluation of the existing output is difficult.

2.3.4 Game theory

A reasonable number of papers have attempted to analyse cyber-insurance from a game theoretic perspective. Players in the games in the simplest formulations are attackers and targets or defenders, while more complex games involve attackers, defenders, and a regulator or policy co-ordinator. Massacci et al critique the emergent narrative that insurance companies act as a clearing house for information and then provide guidance on appropriate security investment to firms seeking liability coverage [186]. Their modelling framework demonstrates that this view of cyber-insurance as a delegated policy tool is unlikely to yield the anticipated coordination benefits and may in fact erode the aggregate level of security investment undertaken by targets. This is a similar result to that identified within the economic modelling strand of the literature.

Johnson et al find that equilibria with a joint investment in protection and self-insurance may exist in a one-shot security game under a restricted case with a weakest-link externality [186]. The key conclusion of the analysis is that full market insurance should only be chosen when it is cheaper than an option involving a combination of protection and insurance or full protection against risks (though full protection is arguably unachievable in the real world).

Yang et al investigate cyber-insurance as part of a Bayesian network game analysis on security investment [301, 302]. They argue that when insurance is offered at the actuarially fair price (highly unlikely in practice) that the optimal insurance is full coverage. Pal et al propose Bonacich/Eigenvector centralities of network users as an appropriate parameter for differentiating insurance clients, highlighting the value of network topology as a defence mechanism [225]. Hayel & Zhu deploy a game-in-games framework where a zero-sum game is nested within a moral-hazard game problem to model cyber-insurance with the goal of enabling the systematic design of a robust insurance policy [139]. Martinelli et al investigate how a drop in security investments for non-competitive cyber-insurance markets might be prevented [184]. Zhang et al develop a bi-level game appraoch to attack-aware cyber insurance of computer networks [307]. Rios Insua et al model a number of cyber-insurance problems as a network and offer decision making models for cyber-insurance, although the models are not solved or analyzed in detail [245]. Woods et al investigate how aggregated claims data impacts investments in information security using Monte Carlo methods to simulate an extended iterative weakest link model [296].

The game theoretic literature suggests that in theory, there is benefit for defenders in holding cyber-insurance policies, but that the contracts require careful design to deter negligent behaviour by defenders. One challenge is that attackers are not uniform and may have different objectives, such as theft of data for one group but direct extortion for another. Network topology is a particularly interesting area for analysis, given the increasing use of cloud storage.

2.3.5 Law

The literature treating cyber-insurance from a legal perspective is surprisingly sparse considering that contracts are an integral to insurance. Economics informs the optimal pricing of and decision making around insurance, but whether the contract pays out or not on a claim is open to legal interpretation, especially in complex cases. Niewesteeg et al provide the first contemporary legal analysis of cyber-insurance contracts we are aware of focused on the Netherlands [212]. Their results suggest that there are two current options for insurers: a strategy of rigorous market penetration with easily accessible and attractive insurance products, or a strategy of significant hedging of correlated risks that reduces the potential of cyber-insurance. Talesh conducts an analysis contributing to two literatures on organizational compliance: new institutional organizational sociology studies of how organizations respond to legal regulation and sociolegal insurance research on how institutions govern through risk [278]. Talesh concludes that insurers act as de facto compliance managers for organizations dealing with cyber security threats via the provision of risk management services. Heath explores theories of torts and insurance in driving efficient management of risk and addresses the possibilities and limiations of both fields in developing effective deterrence of risk [140].

2.3.6 Surveys and empirical analyses

Surveys of individual corporate decision-makers and empirical analyses of the state of the cyber-insurance market represent a notable component of the body of cyber-insurance literature. Biener et al emphasise the distinct characteristics of cyber risks compared with other operational risks including highly interrelated losses, lack of data and severe information asymmetries based on an analysis of almost 1000 cyber risk incidents [32]. The lack of business and economics literature on cyber-insurance has been identified by Eling et al, who concur with this review in describing the lack of data and modelling approaches in the cyber-insurance literature [89]. Tondel et al explore the challenges insurance companies face in assessing risk including from interviews with insurers; they propose two options for improvement: basing analysis on reusable sector-specific risk models, and including managed security service providers in the value chain [281]. Marotta et al undertake a highly comprehensive survey of cyber-insurance, albeit analyzing only a small number of insurance firms [182]. Their characterization of risks via a 'heat map grid' type analysis is particularly pertinent and helps elucidate the range of technical challenges associated with and complexity of cyber-insurance.

Romanosky et al collect and analyse over 100 cyber-insurance policies filed with state insurance commissioners in the United States to fulfil regulatory requirements [246]. This is an important paper, as it represents deployed pricing schedules in the admitted US insurance market. The analysis is accordingly more objective than qualitative survey of market participants, which are susceptible to sampling bias or subjectivity of participant response. The authors find that policies were generally classified as property and casualty lines and that cyber-insurance is generally not covered under a single line of business. Regarding pricing, they found that the firm's asset value base rate rather than specific technology or governance controls, was the single most important factor used in policy pricing.

Regarding surveys, Woods et al present the first systematic analysis of cyber-insurance proposal forms, suggesting that to avoid adverse selection the number of controls that proposal forms include should be in alignment with two key information security controls: ISO/IEC27002 and the CIS Critical Security Controls [294]. De Smidt and Botzen provide an analysis of individual perceptions of cyber risks among professional decision-makers; they find that the probability of a successful cyber attack is overestimated in general and the financial impact underestimated [271]. A reluctance to insure cyber risks is noted compared against expected value-based decision making, which supports a notion that some may believe that cyber-insurance is unlikely to pay out. Eling and Zhu analyse the relationship between corporate characteristics and the writing of cyber-insurance in the US property and casualty insurance industry; a key finding is that insurers writing cyber-insurance policies use more reinsurance to transfer their risk [92]. Eling and Wirfs use extreme value theory to estimate cyber-risk costs based on an operational risk database [91].

Nurse et al investigate the types of data used in pricing cyber-insurance via a qualitative study of professional practitioners including underwriters and actuaries [216]. Their analysis sheds useful light on the trade-offs faced by insurance suppliers, though their interview sample size is relatively small and the paper acknowledges support from a sole insurer, meaning the analysis might not fully represent the views of the broader insurance market.

Dambra et al surveys past research into cyber-insurance and classifies the outputs into four areas: economic aspects, mathematical models, risk management methodologies, and cyber-event prediction [74]. The authors then identify areas of practical research endeavour where data-driven methodology and automated tools might be used to replace qualitative reporting in insurance pricing across risk prediction, data collection, catastrophe modelling, and forensic analysis.

2.3.7 Insurance-linked security pricing

Reinsurance is a common device used by insurance companies to manage their exposure, but in the event of needing to raise further external capital (especially to cover catastrophic events), there is a relatively active market for insurance-linked securities (ILS), which are also sometimes known as 'catastrophe bonds' [44].

Kolesnikov et al suggest how a cyber-catastrophe bond might be priced drawing on classical catastrophe bond pricing [161]. This work is a useful contribution on pricing a cyber bond, however, the use of an exponential distribution for modelling a cyber event is probably too simplistic a depiction of reality. It might have use in modelling a single catastrophic event, analagous to credit event, but this is unlikely to be helpful in valuing, for example, an aggregate XL structure triggered by multiple small claims.

Kasper provides a very thorough feasibility study into cyber-bonds, concluding that they could be sufficiently attractive to capital market investors and sufficiently utility enhancing for the issuers to be viable [156].

Braun gives a comprehensive overview of historical primary market catastrophe bond issuance [42]. Liu et al present a pricing model for an Insurance Linked Security that uses stochastic calculus in the tradition of financial dervative pricing theory [177]. The authors price the ILS by generating a stream of payments linked to an insurance risk process (a multi-dimensional compound Poisson process) and a reference rate process. This model has potential for use in pricing cyber-insurance linked securities, however, work on appropriate insurance processes remains in early stages and is arguably some way from being ready for deployment in a marketable security. Further, the nature of cyber-risk means that a process that well describes risks at a point in time may quickly become obsolete.

Braun et al conduct a feasibility study into cyber-insurance linked security facilitation of risk transfer [43]. They conclude that this is feasible but only if cyber-risk is better understood, highlighting the need for better modelling of the cyber-peril.

2.3.8 Capital requirements

Eling and Schnell provide an overview of the capital requirements for cyber-insurance, analyzing Solvency II, US risk-based capital standards and the Swiss solvency test [90]. This work thus captures many of the key jurisdictions in which insurance is commonly written. The authors conclude that the regulatory models surveyed underestimate the potential risks associated with cyber-threats due to the heaviness of the tail of distributions of potential losses.

Expanding the Gordon-Loeb Model to Cyber-Insurance 3

3.1 Introducing the Gordon and Loeb Model

Gordon and Loeb proposed a model[1] for decisions on information security investment in 2002, in which the probability of a security breach occurring reduces with investment according to a specified function [128]. Under such a framework, a rational decision-maker will aim to maximise the expected net benefit of investment in information security. Gordon and Loeb consider two classes of security breach function and show that for these functions, the optimum security investment will always be less than $(1/e)$ times the expected loss. The Gordon and Loeb model is well suited to the type of Marshallian cost-benefit analyses undertaken by decision-makers in firms as it is intuitive, adaptable and does not require advanced Mathematical knowledge. This work addresses the research question of whether the Gordon-Loeb model can form the foundations of a classical expected utility maximization problem to investigate some of the trade-offs between security investment and cyber-insurance. Following a literature review of the fields of insurance economics, security economics and cyber-insurance, we present a single period, two-state model where the utility of an insurance buyer is maximised subject to a number of constraints. We assume that decisions around information security may be framed solely based on economic considerations, which results in a model that is fairly abstract compared with reality.

[1] The original paper is titled *The economics of information security investment*, but the model contained in it is generally known in the literature as the 'Gordon and Loeb' or 'Gordon-Loeb' model

However, we believe that the model yields useful insights on cyber-insurance pricing and provides the foundations for further work and development in this field.

3.1.1 Key results from the Gordon-Loeb Model

Some key results and assumptions underpinning the Gordon-Loeb (henceforth GL) Model are briefly summarised here, which have relevance for the model developed in this research. Gordon and Loeb assume that an information set may be characterised by three parameters: l, τ and v which represent the loss conditioned on a breach occurring, the probability of a threat occurring and the vulnerability (the probability that a threat once realised would be successful). In the GL model, τ and l are assumed to be constant. The expected loss from a breach event if no investment is made is then $E[L] = \tau v l$. This loss may be reduced by an investment in security z, which the model accommodates via the introduction of a security breach probability function, $S(v, z)$. The GL model makes three assumptions about $S(v, z)$:

A1: $S(z, 0) = 0$ for all z

A2: For all v, $S(0, v) = v$

A3: For all $v \in (0, 1)$ and all z, $S_z(z, v) < 0$ and $S_{zz}(z, v) > 0$ where S_z and S_{zz} are the first and second partial derivatives of the security breach probability function with respect to z.

The expected benefit of investment in information security (EBIS) may be defined as:

$$EBIS(z) = [v - S(z, v)]\tau l \tag{3.1.1}$$

This is the reduction in the expected loss as a result of the investment z. Subtracting the investment, z, then yields the expected net benefit of investment in information security (ENBIS):

$$ENBIS(z) = [v - S(z, v)]\tau l - z \tag{3.1.2}$$

ENBIS neatly encapsulates the cost-benefit trade-off of security investment and should be strictly positive for a rational decision-maker investing in security measures.

As the security breach probability function is strictly convex in z by definition, $ENBIS$ is accordingly strictly concave in z, meaning that an interior maximum $z^* > 0$ is given by the first order condition:

$$-S_z(z^*, v)\tau l = 1 \tag{3.1.3}$$

The GL model proposes two classes of security breach function: $S^I(z,v) = \frac{v}{(az+1)^\beta}$ and $S^{II}(a,v) = v^{\alpha z+1}$. α and β are parameters for the productivity[2] of information security. The optimal level of investment in defence for a particular information set are then easily obtained for the two classes of security breach functions:

$$z^{I*}(v) = \frac{(v\beta\alpha l\tau)^{1/(\beta+1)} - 1}{\alpha} \tag{3.1.4}$$

$$z^{II*}(v) = \frac{\ln(1/-\alpha v l\tau(\ln v))}{\alpha \ln v} \tag{3.1.5}$$

Gordon and Loeb show that for either of these forms of S, $z^*(v) < (1/e)v\tau l$. The GL security breach functions are illustrated in Figure 3.1 for vulnerability $v = 0.65$ and the corresponding ENBIS for these breach functions.

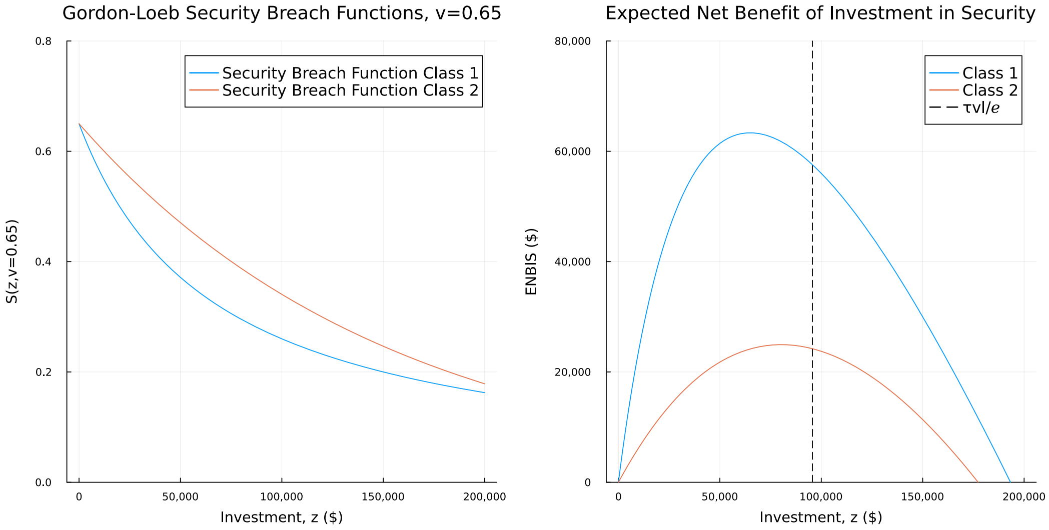

Figure 3.1: Example Gordon-Loeb security breach functions

[2] In economics, productivity is a measure of the efficiency of an input

3.2 Related work

3.2.1 Critique of the approach of Young et al (2016) to combining the Gordon-Loeb Model and Cyber-insurance

Young et al adopt a similar conceptual approach to the model introduced in this chapter in terms of setting up an optimization problem incorporating parameters from the Gordon-Loeb model [305]. They propose minimising the expression $S(z,v)\tau l + z + P$ with the constraints that the cost of security investment and insurance premium cannot exceed the security budget and that coverage should be fixed at l. For the premium, they assume a base premium rate of 8% which is then discounted in a linear fashion based on the Gordon-Loeb Security breach function for levels of investment. Their model is solved using a commercial solver add-in to Microsoft Excel. While empirically pragmatic, this lacks mathematical rigour. Furthermore, this minimization is only reliable to specific practical examples where one is assured of the appropriateness of the chosen parameters. The approach taken by Young et al has merit for use in an enterprise situation (for example by a risk department) where a quick calculation is required for analytical approaches, but falls short of the rigour and theoretical consistency provided by a formal economic model such as the GL model.

3.2.2 Mazzoccoli and Naldi on mixed insurance/investment cyber-risk management

Mazzoccoli and Naldi produced a valuable contribution to the literature on cyber-insurance [189], pursuing a similar approach to [305]. A central feature of the Mazzoccoli and Naldi approach is their inclusion of a Gordon-Loeb type security breach function in the premium calculation that might be charged by an insurance company. This differs subtly, but importantly, from the approach in the model presented in this chapter of treating the Gordon and Loeb security breach function as governing the probability of a breach occuring from the perspective of the insurance buyer, rather than the insurance supplier, who is treated as exogenous[3]. Helpfully, Mazzoccoli and Naldi incorporate the possibility of variable coverage and deductibles, which give their model significant real world relevance. Their analysis ultimately focuses on the optimal investment allocation for any given vulnerability level, which provides a useful comparison for our model results. However, our approach aims to provide extensive insight into the implications of variation in the full gamut of relevant parameters on the expected utility of an insurance buyer. We believe the most important contribution of our model is that investment dynamically reduces the breach probability and thus the amount of risk a buyer would wish to insure. We view our work as complementary to the approach of Mazzoccoli and Naldi, though, rather than contradictory.

3.2.3 Return on Security Investment (ROSI)

Sonnenreich et al proposed a measure for calculating the value of security expenditure, the return on security investment (ROSI) [273]:

$$\text{ROSI} = \frac{(\text{Risk Exposure} \times \%\text{Risk Mitigated}) - \text{Solution Cost}}{\text{Solution Cost}} \tag{3.2.1}$$

This measure is broadly similar to ENBIS in the Gordon-Loeb Model (Equation 3.1.2), but is defined in percentage rather than monetary terms. This metric is potentially very useful in a real-world context, the parameters of risk exposure and percentage risk mitigated are extremely difficult to estimated as noted by [273]. This is an area, therefore, where theoretical economic models of security investment may be able to make a useful contribution by providing some initial quantitative inputs that could then be refined based on real world experience. This is a common approach in insurance, where expected loss distributions might be initially simulated but then refined based on losses and claims experienced.

[3] In less formal terms, the insurance supplier is treated as an independent input to the model

3.3 Economic utility

3.3.1 Utility functions

It is possible to formulate theoretical problems in generality without specifying the form of a utility function and solely its inputs. Whilst this allows for determining the conditions required for an optimum, to produce numerical outputs and thus make judgements on real world problems, the utility function needs to be specified. As elegantly described by Gollier, “It is often the case that problems in the economics of uncertainty are intractable if no further assumption is made on the form of the utility function." [124]. The optimal form of utility function is of great importance for solving problems and forms a significant branch of literature in its own right. For the purposes of this research, we work with simple, well established utility functions as the focus of the models in this thesis are to demonstrate how real world problems might be set up as opposed to constructing detailed case studies of real world organizations.

We now explain how to relate preferences and utility functions. There are three key properties in relation to the utility function that are usually considered, absolute risk aversion:

$$A(z) = -\frac{u''(z)}{u'(z)} \tag{3.3.1}$$

prudence:

$$P(z) = -\frac{u'''(z)}{u''(z)} \tag{3.3.2}$$

and relative risk aversion:

$$R(z) = -\frac{zu''(z)}{u'(z)} = zA(z) \tag{3.3.3}$$

The usual procedure for establishing a utility function is to determine the set of risk preference characteristics of the agent in the model, then to choose a utility function that captures these characteristics. The parameters of the utility function can then be set as required for the problem at hand.

Forms

There are two particular classes of utility function that have properties of constant absolute risk aversion (CARA):

$$u(z) = \frac{1 - e^{-az}}{a} \tag{3.3.4}$$

or constant relative risk aversion (CRRA):

$$u(z) = \begin{cases} z^{(1-\gamma)}/(1-\gamma) & \text{if } \gamma \neq 1 \\ ln(z) & \text{if } \gamma = 1 \end{cases} \tag{3.3.5}$$

As noted by [151], CRRA is an established choice within the cyber-insurance literature though examples of CARA are also found. Both forms are used in this thesis, though in cases involving unsophisticated agents, CARA is favoured for ease of manipulation.

State dependence

A key use of utility functions is to optimise decision making across multiple states of the world with different payoffs. For simple insurance pricing examples (see, for example [241]), it is convenient to work with a model involving two states — loss and no loss. Given a loss, L, with probability π of occurrence, insured by a policy costing P providing cover C, the insurance decision problem in terms of expected utility is

$$E[U] = (1-\pi)u(-P) + \pi u(-L - P + C) \tag{3.3.6}$$

where $u(.)$ is the appropriate utility function for the decision maker covered by the problem. It is evident that the utility of the agent is dependent on the different payoffs in different states and the probability of that state of the world occurring. More generally, consider a set of states $s \in S$, with associated losses L_s where each state has an associated probability, π_s. Denoting the overall probability of *any* loss occuring as $\pi = \sum_s \pi_s$, subject to the constraint $0 \leq \pi \leq 1$. Writing the indemnity provided by an insurance policy in state s as C_s, and assuming that there is a component of the premium, P_s attributable to each state Equation 3.3.6 becomes

$$E[U] = (1-\pi)u(-P) + \sum_s \pi_s u(-L_s - P_s + C_s) \tag{3.3.7}$$

This is not simply of theoretical relevance, but is highly relevant for cyber-insurance decisions as a buyer may need to choose between different policies offered by different firms which offer different indemnities against different risks.

3.4 Incorporating cyber-insurance into the Gordon-Loeb Model

We introduce a simple model following Rees and Wambach to describe the microeconomic analysis of a firm aiming to determine its optimal level of cyberinsurance cover [241]. For convenience, this model will be referred to as the Gordon-Loeb with cyber-insurance (GLCI) model. The GLCI model considers the decisions of an individual (for example a Chief Information Security Officer) charged with allocating an annual cyber-security budget, which is treated initially as analogous to the wealth of an individual in a traditional analysis of insurance. A simple model for insurance demand may be formulated in terms of maximising the expected utility (see Section 3.3.1) of an insurance buyer where there are two states, no-loss and loss:

$$\begin{gathered} E[U] = (1-\tau S(z,v))u(B_{sec} - z - P(C)) \\ +\tau S(z,v)u(B_{sec} - z - P(C) - l + C) \end{gathered} \tag{3.4.1}$$

The probability of the loss state is given by a Gordon-Loeb security breach function, $S(z,v)$ - thus, investment in security measures reduces the probability of a loss. The maximum expected loss is denoted by l. The introduction of the GL model SBFs into the utility function of an insurance buyer is, to the best of the author's knowledge, the first example of their use in a classical economic analysis of insurance. C represents the cash coverage of the insurance policy. The case $C = \tau l S(z,v)$ thus implies full cover but depending on the cost of the premium, it may be optimal for the insurance buyer to only take partial cover and accept some residual financial risk. B_{sec} is the security budget (analogous to wealth in the classical insurance model), $P(C)$ is a cash premium, assumed to be a function of cash cover C, which is allowed to vary. $u(.)$ is a von Neumann-Morgenstern utility function[4], which is increasing and strictly concave implying that the individual is risk averse. Constant absolute risk aversion (CARA) and constant relative risk aversion (CRRA) forms are used for comparative purposes in the simulations that follow, which are defined in Section 3.3.1).

The expected utility is a function of two states: one where a breach does not occur and one where a breach occurs. In both states it is assumed that an investment z is made; this investment is allowed to vary but for simplicity, timing effects[5] and the decision process around that investment[6] with respect to the system are both excluded. This leaves the model relatively abstract[7] in relation to a real-world example of security investment and defence, although its one-period nature is arguably comparable to the annual budgeting and investment cycle undertaken by many organizations both in government and industry. Further, the estimation and attribution of economic losses from cybersecurity incidents is a live area of research and there is no reason why the loss parameters in the model could not be expanded as required for a specific use case. The GL model arguably suffers the same limitations as the model in this chapter and these simplifying assumptions still allow for a useful economic analysis of the interaction between security investment and insurance as is evidenced by the enduring popularity of the GL model and the significant body of subsequent literature that has developed. The expected utility maximization problem then becomes:

$$\begin{aligned}\max_{C\geq 0} \bar{u} = (1-\tau S(z,v))u(B_{sec} - z - P(C)) \\ +\tau S(z,v)u(B_{sec} - z - l + C - P(C))\end{aligned} \tag{3.4.2}$$

subject to the constraints $P(C) = pC$ where p represents a percentage premium

[4] Such a utility function is one that conforms to the four axioms proposed in von Neumann and Morgenstern (1944) [211]

[5] Timing effects in an economic model where the unit of measure is money would require treatment of the time value of money and assumptions on interest rates. This would increase the model complexity without yielding significant insights relevant to the research question

[6] This could be a fruitful area of potential further research work

[7] The same observation applies to economic models used in decision making in a range of fields; for example, models of the economy used by Central Banks to inform monetary policy.

(as is conventional in insurance) and $z + pC \leq \frac{v\tau l}{e}$. The value $\frac{v\tau l}{e}$ is the maximum potential value of optimal security investment in the Gordon-Loeb model. The choice of cash constraint is likely in reality to be dictated by a the budgetary preferences of a firm and the Gordon-Loeb maximum potential investment is used as a convenient assumption rather than one that can be rigorously proved as in the GL model. Under the simplifying assumption that p is constant and determined by the insurance supplier, the insurance buyer is faced with the decision as to how much cover to take at that premium. Substituting the first constraint into equation 3.4.2 yields:

$$\begin{aligned} \bar{u}(C,z) = (1-\tau S(z,v))u(B_{sec}-z-pC) \\ +\tau S(z,v)u(B_{sec}-z-l+C(1-p)) \end{aligned} \tag{3.4.3}$$

In this formulation, the level of cover C and defensive security investment z are the only variables in the problem; the vulnerability v is an inherent property of the information set as are τ and l. The Lagrangian[8] for the problem depicted in Equation 3.4.2 is:

$$Z = U + \lambda(\frac{v\tau l}{e} - pC - z) \tag{3.4.4}$$

where $U = \bar{u}(C,z)$ The Karush-Kuhn-Tucker conditions[9] are (where Z_x denotes the partial derivative of Z with respect to x):

$$\begin{aligned} Z_C = U_C - p\lambda \leq 0 \qquad C \geq 0 \qquad C.Z_C = 0 \\ Z_z = U_z - \lambda \leq 0 \qquad z \geq 0 \qquad z.Z_z = 0 \\ Z_\lambda = \frac{v\tau l}{e} - pC - z \geq 0 \qquad \lambda \geq 0 \qquad \lambda.Z_\lambda = 0 \end{aligned} \tag{3.4.5}$$

The third constraint implies that for $\lambda \neq 0$, the solution would imply a commitment of capital up to the Gordon-Loeb maximum. In the case where both cover and

[8] A Langrangian is a function used in mathematical optimization for finding the maxima or minima of a function subject to an equation being satisfied by chosen values of certain variables.

[9] These are the conditions under which an optimal solution to a non-linear programming problem such as the one in the model proposed in this chapter may be found - see, for example, [120] for a formal definition. As the form of the utility function is variable, it is important not to lose generality at this stage.

investment are non-zero, by conditions 1 & 2, we assume $Z_C = Z_z = 0$ then:

$$\lambda = \frac{U_z - U_C}{1 - p} \tag{3.4.6}$$

This then implies that the fair premium is given by

$$p = \frac{U_C}{U_z} \tag{3.4.7}$$

This set of conditions specify the conditions under which a local maximum may exist. However, there is no guarantee that under all sets of conditions that it will. Furthermore, depending on the nature of utility function chosen, solving the system of equations in Equation 3.4.5 has the potential to become a difficult non-linear programming challenge in general terms. Our primary focus is to ascertain whether the model provides useful insights that can guide behaviour towards security investment. This can likely be deduced via the appropriate use of graphical methods to evaluate the model in the first instance to guide an optimization strategy for model cases rather than producing a closed-form solution *ab initio* that is algebraically intractable and unintuitive to interpret.

3.5 Simulation

3.5.1 Method

The simulations of the Gordon-Loeb with Cyber-Insurance (GLCI) model use the following parameters, adapted with slight variations from [130, 209]. We set $l = \$500,000$ with the probability of a threat occuring, $\tau = 0.8$. Both of these parameters are constant in the GL model, which gives an expected loss of \$400,000 before any security investment, z. v is initially set at 0.65, which as previously discussed represents the probability that a threat is successful *once realised.* The budget of the defender, B, is set at \$600,000 — this is equal to $L + z^*$ with extra margin to avoid the risk of negative inputs to the logarithmic utility functions. Finally, $\alpha = 1.5 \times 10^{-5}$ in Class I and II breach functions and $\beta = 1$ for the Class I breach functions. This choice of α was informed partially by the existing literature, where $\alpha = 1 \times 10^{-5}$ is often used; this produced some erratic behaviour within the Class II security breach function whereas the slightly higher α provides well bounded results for both classes of security breach function. The parameter values used in the simulation give well-bounded results and allow for a thorough examination of the model behaviour. Graphical analysis was generated using the *Plots.jl* package within the Julia language.

The GLCI model simulations are presented using both logarithmic and exponential utility functions in the form of plots of the utility functions varying different model parameters. Initially, closed form[10] solutions to the system of equations in Equation 3.4.5 were sought but it became clear that this approach was unlikely to prove fruitful given the large number of variables in the model and small number of constraints. Furthermore, the choice of utility function could be varied depending on the use case and consequently plotting the utility functions imposing the relevant model constraints is sufficient for evaluating the focal research question of this work.

3.5.2 Optimal investment per the Gordon-Loeb Model, variable cover

We first consider the simple case where a firm invests the optimal amount recommended by the Gordon-Loeb model, z^* and then investigates the possibility of cyber-insurance with varying cover and different premium rates observable in the market. To illustrate this case, we plot both logarithmic and exponential utility functions for Equation 3.4.3 in Figure 3.2. The logarithmic utility function is simply $u(.) = ln(.)$ while the exponential function is Equation 3.3.4 setting $a = 10^{-5}$. These utility functions will be used for the remainder of the simulations in this work. A key model assumption is that the total cost of investment and insurance premia should not exceed $(1/e)\tau vl$. Having invested an amount, z, the GL model states that there is a commensurate reduction in the probability of a breach being successful. Utility functions are therefore plotted up to cover $C = min(\tau lS(v, z^*), \frac{(1/e)\tau vl - z^*}{p})$. This ensures that the monetary amount spent on security investment and insurance does

[10] Equations produced using software to resolve the symbols contained within the utility functions

not exceed the imposed constraint. The results broadly suggest that utility is largely maximised at maximum coverage, which concurs with comparable game theoretic modelling work [151, 301].

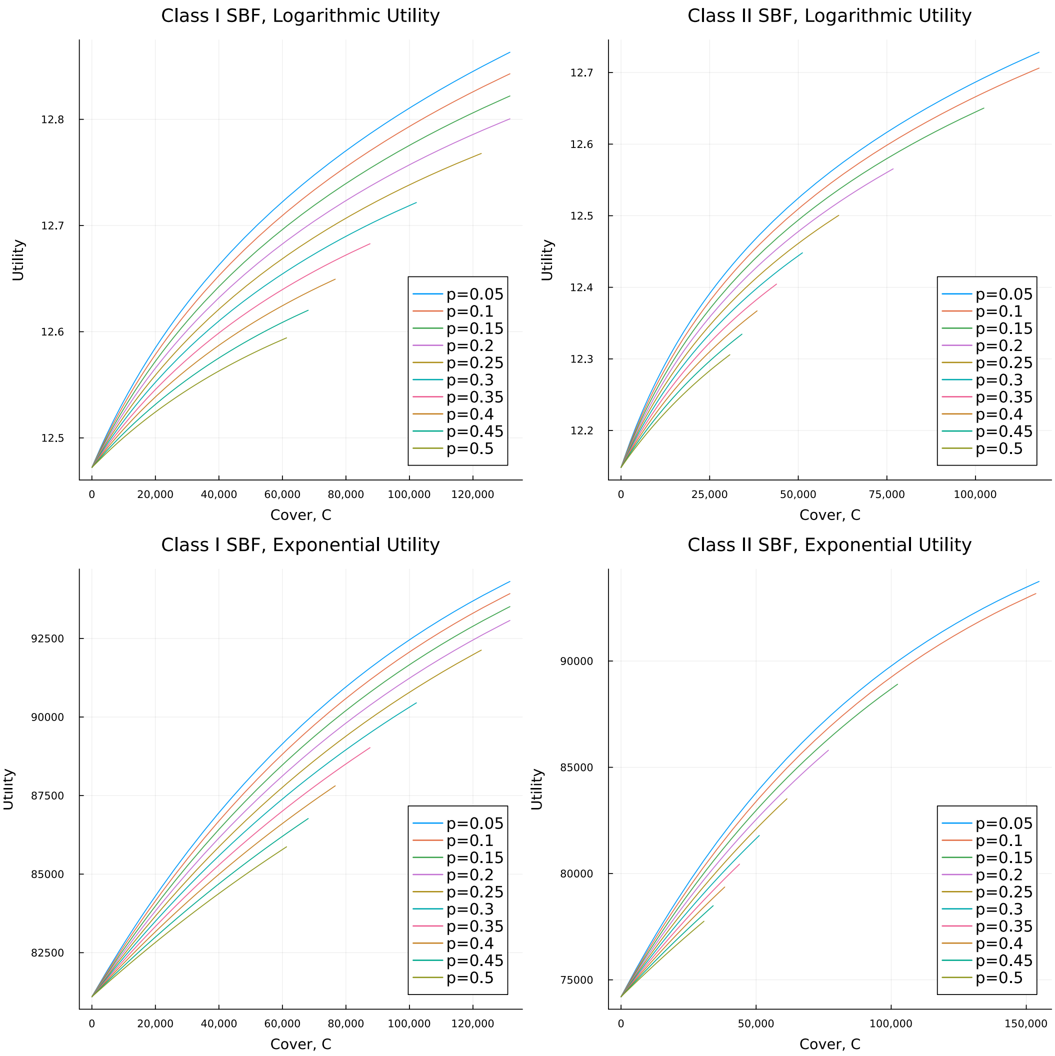

Figure 3.2: Utility as a function of cover assuming $z = z^*$

Variable investment, maximum cover

Relaxing the assumption that the firm first invests the optimal amount into protecting its information allows us to consider the competing interaction between spend on insurance and investment. As in Section 3.5.2 the maximum cover an insurance buyer would wish to take out is $C_{max} = \tau l S(v, z)$ with maximum cover available respecting the cash constraint is then given by $C = \frac{(1/e)\tau v l - z}{p}$.

Figure 3.3 shows the variation of maximum available cover subject to the cash

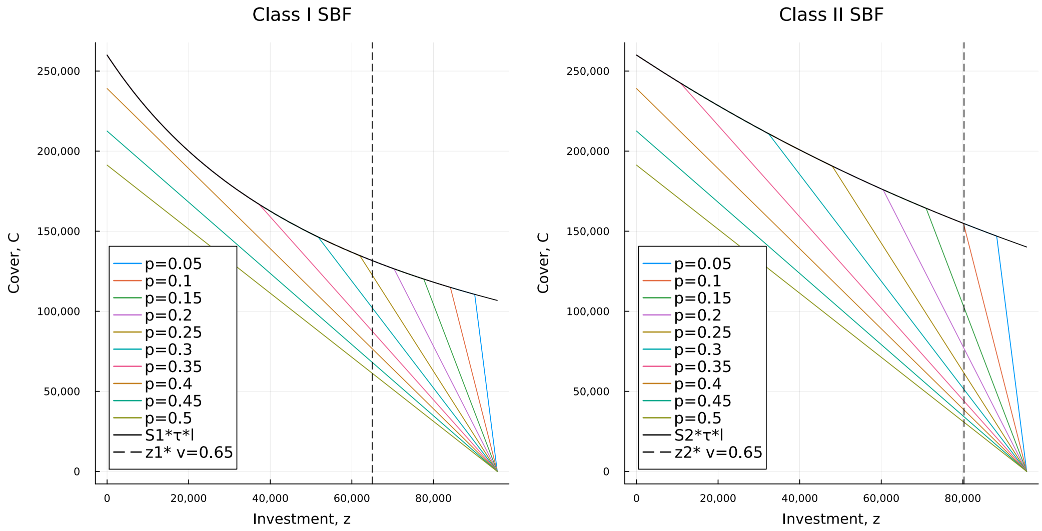

Figure 3.3: Maximum available cover under the cash constraint at different levels of z

cost constraint with premium rates, along with the optimal GL values of investment for reference and the theoretical maximum cover at each value of z. For class I SBFs, it is possible for an insurance buyer to obtain full coverage at z^* for premia less than 25% in our model setup. However, for a corresponding class II SBF, only premia below about 10% offer full cover under the terms of the model. Figure 3.4 illustrates the utility functions in the case of variable investment. The relevant optimum level of investment specified by the GL model is plotted as a dotted vertical line. Under the cover decision we have outlined, it is clear that insurance is usually preferable to investment in our example model set-up at all but very high insurance premium rates. This is an interesting result as the utility functions plotted incorporate the expected benefits of a reduction in breach probability. Economically this makes sense - if the cost of insuring a risk is lower than the cost of reducing it to a certain level then it makes sense to take out the insurance.

3.5.3 Premium versus vulnerability under optimal security investment

Thus far simulations have had fixed $v = 0.65$. It is interesting to consider the effect of varying v, especially for the second class of GL security breach functions, which are exponential in v. To do so, it is assumed that the insurance buyer invests the

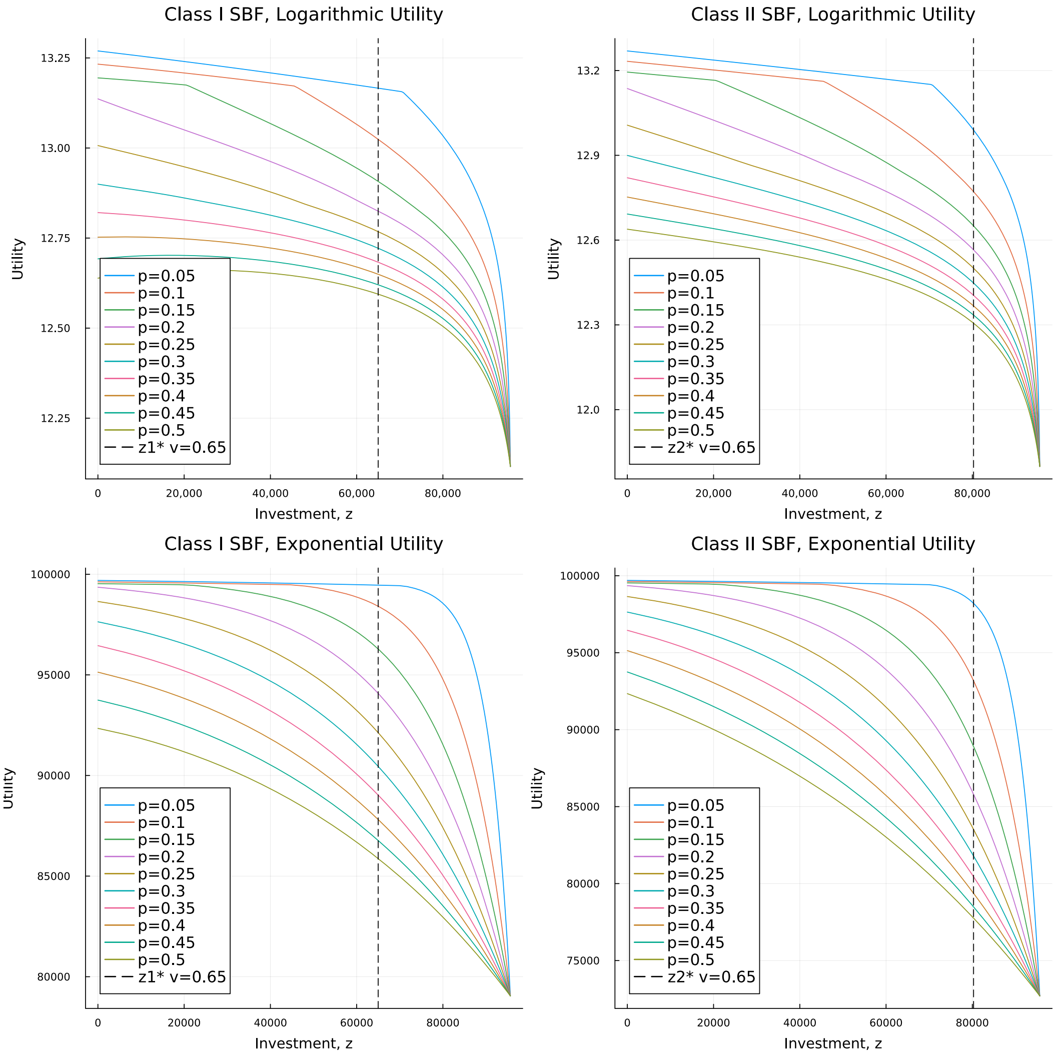

Figure 3.4: Utility as a function of investment with maximum insurance coverage purchased

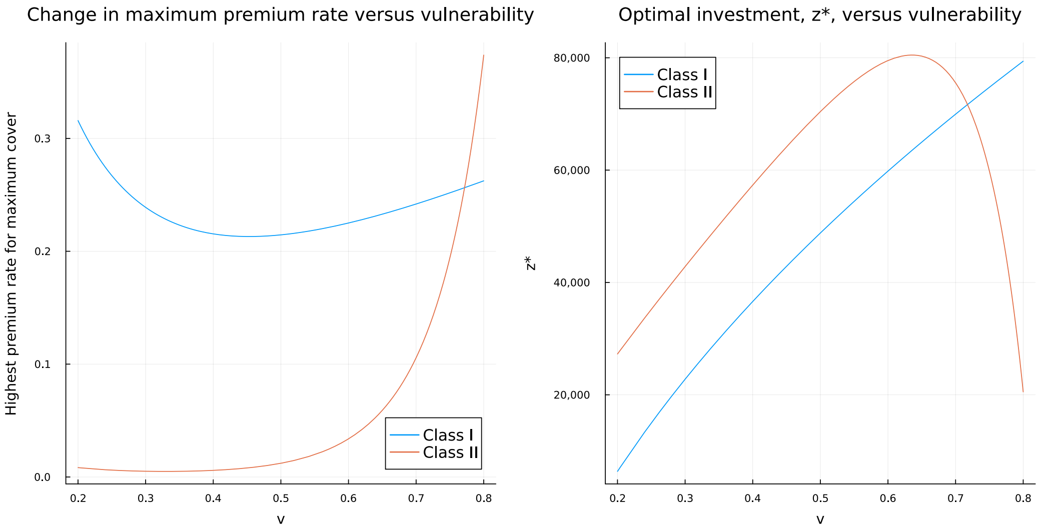

Figure 3.5: Highest premium rate at which maximum cover may be obtained for vulnerability, v and variation of optimal GL investment with vulnerability, v

optimal amount recommended by the GL model. Figure 3.5 plots the variation of the highest premium at which full cover can be achieved with v for both GL SBF classes and also how the GL optimum investment, z^* varies with v with the other model parameters as specified previously. Figure 3.6 plots the utility functions previously described for buying insurance at the maximum coverage available (as described in section 3.5.2) as a function of the vulnerability, v. The main use of this analysis is to demonstrate how the sensitivity of the utility to the premium rate varies at different values of v.

Table 3.1 provides an alternative presentation of this analysis. For each vulnerability, v, the maximum investment under the GL model is calculated followed by the optimum for class I and II SBFs. The expected probability of breach after the investment is then calculated. The maximum cash available to the insurance buyer for insurance purchase is then calculated, from which the maximum premium rate at which full relative cover may be achieved is then calculated. For the class I SBF, this is relatively high; however for Class II SBFs, the relatively higher level of z^{II*} compared with z^{I*} means that it is difficult to achieve full coverage. It should be noted that Class II SBFs start to produce somewhat erratic results as $v \to 1$ given the form of $z^{II*}(v)$.

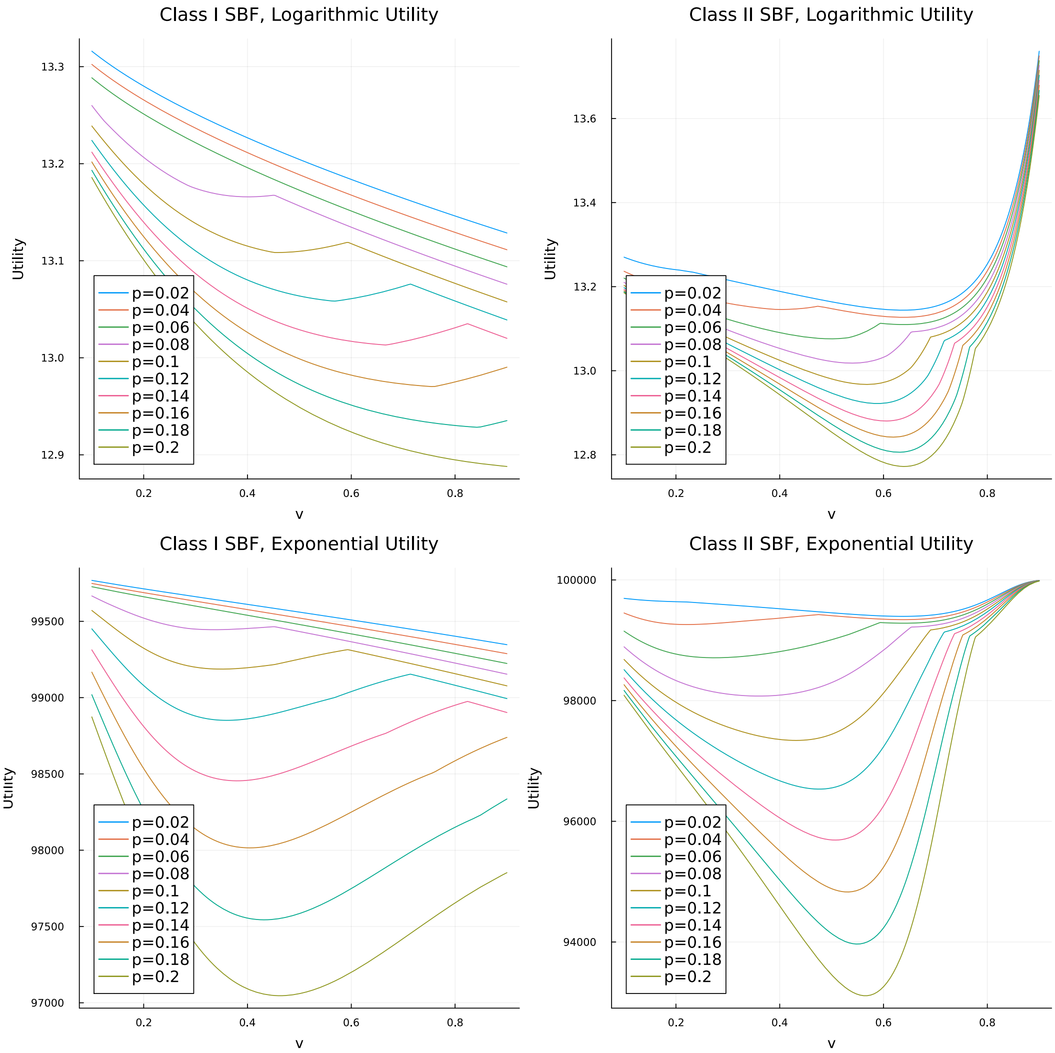

Class I SBF, Logarithmic Utility

Class II SBF, Logarithmic Utility

Class I SBF, Exponential Utility

Class II SBF, Exponential Utility

Figure 3.6: Utility functions for different vulnerabilities assuming investment at the Gordon-Loeb optimum, z^* and maximum coverage respecting the model cash constraint

v	z_{max}(\$)	z^{I*}(\$)	z^{II*}(\$)	$S^I(v, z^{I*})$	$S^{II}(v, z^{II*})$	P^I_{max}(\$)	P^{II}_{max}(\$)	p^I_{max}(%)	p^{II}_{max}(%)
0.20	29,430	6,363	27,264	0.183	0.104	23,067	2,166	31.6	5.2
0.25	36,788	14,983	35,207	0.204	0.120	21,805	1,581	26.7	3.3
0.30	44,146	22,776	42,826	0.224	0.138	21,369	1,320	23.9	2.4
0.35	51,503	29,943	50,203	0.242	0.159	21,561	1,300	22.3	2.0
0.40	58,861	36,613	57,336	0.258	0.182	22,248	1,525	21.5	2.1
0.45	66,218	42,878	64,140	0.274	0.209	23,340	2,079	21.3	2.5
0.50	73,576	48,803	70,413	0.289	0.240	24,773	3,163	21.5	3.3
0.55	80,933	54,439	75,772	0.303	0.279	26,494	5,162	21.9	4.6
0.60	88,291	59,824	79,506	0.316	0.326	28,467	8,785	22.5	6.7
0.65	95,649	64,989	80,292	0.329	0.387	30,659	15,357	23.3	9.9
0.70	103,006	69,959	75,541	0.342	0.467	33,047	27,465	24.2	14.7
0.75	110,364	74,755	59,829	0.354	0.579	35,609	50,534	25.2	21.8

Table 3.1: Sample parameters for different vulnerabilities, v. $\tau l = 400,000$. $P_{max} = z_{max} - z^*$, i.e. the maximum cash available to pay an insurance premium after investing at the optimal level given by the GL model. p is the highest premium rate at which coverage equal to the expected loss after investment, $\tau l S(v, z^*)$, can be achieved.

3.6 Remarks

3.6.1 Model limitations

The GLCI model demonstrates that the Gordon-Loeb security breach functions can be used within a classical two-state utility maximization model yielding useful insights. In particular, the GLCI model offers insight into the competing dynamics of purchasing insurance coverage versus investing in security. However, it is inherently abstract in drawing on the Gordon-Loeb model and classical microeconomic treatment of maximising expected utility. This abstraction brings advantages in terms of ease of use and adaptability but this is at the expense of realism. In a real-life scenario, the trade-offs between security investment and insurance are likely to be more subtle and also not exogenous as our model assumes. A specific example of this is the approach to the insurance premium, which is likely to be unique to each insurance buyer and their specific circumstances. The model treats the premium rate, p as an independent, market observed variable. Further, the insurance premium for each utility curve is static and the buyer has the choice of purchasing varying levels of cover at that rate combined with an investment in security subject to a cash constraint proportional to the 'value' of the dataset. In reality, the baseline market observed premium is likely to be a reducing function of the investment in security, z, as the insurer is likely to account for the reduction in breach probability effected by the client. The insurance problem has been framed from the perspective of the insurance buyer (the decision-maker in the model) as this naturally follows from the Gordon-Loeb model. However, a useful extension would be to include a more sophisticated premium rate term. Unfortunately, cyber-insurance premium data is extremely difficult to obtain in the public domain as the inputs are of high commercial sensitivity to insurance companies. A model with dynamic premia would also further increase in complexity as an optimization problem, but the simulation approach in this chapter would likely yield useful insights.

A further problem is that the nature of loss introduced in the Gordon-Loeb model is hard to reconcile with real world scenarios in the context of insurance. The concept of loss for many lines of insurance is relatively straightforward to understand; if an individual's vehicle is stolen for example, one motivation (beyond the fact that in most countries it is a legal requirement) is that the insurance should cover the cost of replacing the vehicle as well as any damage inflicted by the driver on other vehicles or persons. However, for data, what is the economic notion of loss? One interpretation would be the regulatory costs of a breach an organization might suffer, but these cannot necessarily be covered by insurance. One purpose of regulation such as GDPR could be argued to be protecting consumers by providing a significant financial deterrent to firms from not investing in appropriate security measures. This presents an issue of possible moral hazard around cyber-insurance; the fact that a firm can recover some of its costs if data is stolen is of scant benefit to consumers, for example, if their valuable personal data is stolen. Where cyber-insurance does have a useful role to play is in assisting firms with forensic computing resources to identify the extent of a breach once identified, to help patch any vulnerabilities and to aid with system recovery in the event of a ransomware attack, for example. These dynamics are rather difficult to properly encapsulate in a the simple parameters of loss and coverage. A further issue is that once data is stolen, it can be duplicated, so differs from many conventional economic goods in terms of potential recovery. There are also issues of reputational damage to a firm that must be considered following a data breach; these could provide some motivation for the purchase of an annuity-type structure as part of an insurance package as one would expect the effects of a data breach to gradually fade from public memory over time.

The behaviours of those involved in attacking and defending a set of information are also of interest, though perhaps are better represented via a game theoretic treatment of the problem rather than in a classical economic model. However, the notion of a constant threat probability in the GL model is possibly one of the more problematic assumptions in a real world sense. It is helpful to treat attacks as arising from nature in an initial evaluation of the problem, but it would equally be relatively straightforward to attempt to measure the frequency of general attacks (e.g. via the use of a 'honeypot' [47, 201, 282, 36]) and then including a parameter to account for the risks of a firm being a specific target. There is also the question of the behaviour of the defender (the insurance buyer in our model). Ioannidis et al discuss the notion of a steward, who is able to intervene under certain conditions to either slow the degradation of a system's operating capacity (promoting sustainability) or return a system to its intended state (resilience) [148]. Under the right circumstances, the presence of a steward might help to turn a major data breach into a minor one and thus reduce the tendency for loss. The steward for an organization might be its cyber-security team, who if deemed capable by an insurer, would likely result in it quoting a lower premium.

Thus far, we have negated the supply side of the insurance market, which our model treats as a readily available commodity at uncertain price. In reality, most insurance policies will have a coverage limit and responsible insurers will have clearly defined and enforced risk limits. A particular issue with cyber-insurance is the ability to offset risk. A common strategy among insurers appears to be offering consulting services as part of the insurance package, which generates revenue that helps to form a compensation pool in the event of an insurance claim while also lowering the risk that such a claim will occur. There is an issue of adverse selection inherent in cyber-insurance; a naïve view might be that the insurance buyer poses greater risk as an insurer cannot know all the details of the insurance buyer's activities. However, in reality, the insurer likely has a great information advantage; there are only a limited number of cyber-insurers who are likely to have proprietary pricing models and datasets of breaches and vulnerabilities assembled from a multitude of customers and sources. It is very difficult for firms in a sector to share such information, and indeed to do so might be considered economically irrational (albeit potentially socially responsible) and in some potential instances unlawful given competition law. There is no guarantee also that an insurer will agree to provide coverage at an economically satisfactory level, and the insurance buyer must be assured that the policy is likely to pay out as it expects. The GLCI model helps to quantify what the economically satisfactory level might be (see Figure 3.5 and Table 3.1). However, the model inherently assumes that in the loss state with probability of breach given by the Gordon-Loeb security breach functions, the policy will pay out with certainty. This is difficult to parametrise *a priori*, but a distribution of cyber-insurance payouts might be obtained or modelled to incorporate this uncertainty.

3.7 Summary

This chapter has demonstrated that the Gordon-Loeb model for investment in information security can be used to build a model for cyber-insurance based on maximising the expected utility of an insurance buyer. This model suggests that when the Gordon-Loeb recommended optimum is invested in security measures, then utility is maximised at full coverage for reasonable insurance premium rates subject to a cash constraint that the total spend on security measures and insurance cannot exceed the maximum amount stipulated by the Gordon-Loeb model. We demonstrate that for each of the two classes of Gordon-Loeb security breach function, there is a maximum premium rate at which cover can be purchased equal to the maximum expected loss from a breach after the investment has been made while respecting an imposed cash constraint that the total spent on security investment and insurance cannot exceed $(1/e)$ of the maximum total expected loss. The abstract nature of the model means that it simplifies the intricate trade-offs and decisions of a real-life security investment problem. Nevertheless, it establishes in a rigorous economic sense that cyber-insurance can be a cost effective solution in addition to security investment.

Pricing Cyber-Insurance Based on System Structure 4

4.1 Background

Cyber-insurance policies provide financial mitigation against defined cyber-risk events up to a specified limit, for which the buyer of insurance pays the insurance provider a premium (fee). The definition of events and calculation of both premia and losses are non-trivial and require significant administrative, financial and legal resources. As such, the areas of cyber-risk and cyber-insurance present a multitude of interesting research problems.

For many years, there have been attempts in the academic community to organize and structure the research agenda related to cyber-insurance, an early example of which is a proposal from 2010 by Böhme and Schwarzfor a unifying cyber-insurance modelling framework [38], updated in [37]. Despite considerable subsequent research endeavour, several comprehensive literature reviews [88, 183, 18, 267] suggest that the cyber-insurance research field remains fragmented. This fragmentation motivated a number of eminent researchers and practitioners in the field to reiterate the need for a well-defined research agenda for cyber-risk and insurance [101]. One key identified contribution is to quantify cyber-risk and its relationship with organizational structure.

Insurers will often attempt to assess the security posture of the insurance seeking entity via questionnaire forms [294, 247] and potentially via some limited investigation of their own (for example, perimeter scanning). The use of questionnaires is potentially problematic in cyber-insurance. A company will typically be required to provide answers 'to the best of its knowledge'. This creates the risk that (assuming no dishonesty), if a company does not have a complete grasp of its own security posture or maturity, the price of its insurance policy may not correctly capture the risk dynamics of the organization. For example, asking the question 'Do you have anti-virus software and/or firewalls?' may seem straightforward. However, there are a number of deficiencies in the question. An answer 'yes' does not mean that every machine on the network has anti-virus software installed; it says nothing about the robustness of the firewall configuration or any potential rule exemptions set up. Either of these details could be easily exploited by an attacker to introduce ransomware, for example, into an enterprise network. As such, both the detail of the system architecture and *how* the system is monitored are important.

4.2 Problem statement

The key question for pricing cyber-insurance is: what is the probability of a successful attack on an organization and what are the expected losses from a successful attack? Estimating these requires a threat model to estimate the frequency of attacks and a methodology to estimate the size of losses. Mapping out the detail of an organization's network is one possible starting point, but this only reveals details about the systems an organization uses as opposed to its revenue generating operations, which ultimately determine the scope of possible losses related to a cyber-risk incident. There is therefore a need for an integrated description of the structure of an organization and the IT systems that support delivery of its objectives. Once this relationship is described, the potential scope of losses resulting from cyber-incidents and their probability need to be estimated. Based on these estimates, an insurance company can then quote a premium for reimbursing a customer for these losses should they occur.

This chapter proposes a modelling framework drawing on the fields of operations management, security and economics that can be used to support moving from describing an organization to pricing a cyber-insurance policy. Our motivation is primarily to contribute to the conceptual debate on how to address the significant modelling challenges associated with cyber-insurance pricing and so this chapter does not directly calculate premiums for real-world organizations. However, we aim to provide sufficient illustration that a practitioner equipped with appropriate data could quickly implement it for real-world insurance assessments.

4.3 Theory

4.3.1 Grammar

We shall use the term object to represent a modelling target of interest, which may be (without restriction) a system, phenomenon or some other entity. A language is needed to describe an object. The language may consist of words, diagrams or mathematics. The language has a grammar, which describes the connectivity, properties and components of the language. The components of the language are descriptors of the system and its security objectives. Properties describe the state of the system. Connectivity stipulates how the components interact. In a simple system with a modest set of objectives, there may be just two states in which case economic analysis is not required. However, for more complex systems, the state of the system may only be partially observable, requiring agents to make decisions based on beliefs rather than certainty. These behaviours may be dictated either by policies or by preferences. It may be possible to simplify a representation of a system by making assumptions. The security properties of the system may then be explored by varying parameters within the bounds of the assumptions.

4.3.2 Describing systems

Single systems

The most intuitive way for humans to understand systems is by using diagrams. From the earliest stages of education, children learn to draw pictures that represent their subjective view of the world. At the fundamental level, any computer system is simply a flow of electrons across a large number of transistors, which either permit or block electron flow. This is used to represent truth or falsity (conventionally represented as a binary number, i.e. 0 or 1). The precise details are far too numerous and complex for the human mind to reasonably process, which means that a model or description is needed to analyse the system. One could construct a systems representation along a number of objectives, focused around location, resources, or processes. Fundamentally, a description requires a representation of *what* the components of a system are; *how* they interact; *which* properties relate to each component; and perhaps some measure of their importance depending on the application of the description.

Distributed systems

A distributed system may be defined as "one in which components located at networked computers communicate and coordinate their actions only by passing messages" [81]. Many of the systems of interest for security modelling are distributed systems. The most famous distributed system is of course the internet, which could be viewed in a number of different ways. The worldwide web (WWW) is the most famous component of the internet and can be considered in simple terms as a system for storing and accessing information.

A significant cybersecurity concern is the interaction between internet facing and non-internet facing components of organizational systems. The then raises the question of how the interaction of systems should best be described. This thesis does not directly model distributed systems in the classical sense, but it is important to be aware of this discipline as the models introduced in this thesis could be further expanded to incorporate a rigorous analysis of distributed systems [56].

Interacting systems

Significant research endeavour has been devoted to developing a branch of logic that describes the interactions of different components of a system. One of the seminal contributions is the synchronous calculus of communicating systems (SCCS) by Milner [199]. This work inspired many significant subsequent contributions, which are too numerous to summarise here. The work in this thesis does not directly make use of formal logic, but the principles underpinning calculi such as SCCS are used as the basis of the model developed in Chapter 4.

4.3.3 Describing security

Objectives

Security may be defined in many different ways for a diverse range of applications. In the context of information security, which is the focus of this thesis, a suitable definition is 'the process by which it is ensured that just the right agents have just the right access to just the right (information) resources at just the right time' [236]. The process of ensuring security is specified by *declarative objectives*, which state what is required to achieve the desired state of security. Declarative objectives should not be confused with *operational objectives*, which stipulate how security is expected to be implemented (access control and backups, for example). In a purely binary view of security, objectives might be deemed either true of false. However, if the objectives are measurable in smaller intervals, then the security objectives may be described using parameters which can be used to populate economic models to guide decision making.

4.3.4 Parameters

As discussed in Section 1.2.1, a commonly used set of parameters in security analyses are confidentiality, integrity, and availability. These provide the necessary set of parameters for capturing security objectives in many situations. The extent to which these objectives are met for each of the parameter may be measurable. If the parameters are measurable, then confidentiality, integrity and availability may be defined over the closed interval $[0, 1]$, which is of benefit when working with utility functions and economic models. The parameters may be further simplified into criticality and sensitivity, where criticality is availability and some aspects of integrity and sensitivity is confidentiality and those aspects of integrity not included in criticality. Criticality and sensitivity are arguably more straightforward parameters to work with in instances where independence of the two attributes renders many problems easier to solve, such as probabilistic analysis or multi-attribute utility functions.

4.3.5 Security Maturity Models

Security Maturity Models (SMMs) are a tool used by organizations to describe their security posture in a structured manner. They achieve this by stipulating a framework on how to group processes and resources related to cyber-security and then assess how developed, or mature, is the posture of the organization. The application of security maturity models will be comprehensively demonstrated in Chapter 4. The exact specification and terminology vary from model to model, but most follow the basic structure of:

- *Practices* (p_i) describe single security activities and are *how* maturity is achieved
- *Domains* (d_i) provide a structured set of security practices, grouped by area.
- *Objectives* (o_i) are *what* the organization must do for a practice to be deemed as met.
- *Maturity* (m_i) states the developmental state of objectives, usually assigning a level from a predetermined set.

Any of these elements may be aligned fully or partially with existing cybersecurity standards. The preeminent global standard is ISO27001, which comprises controls organised by category. Within the maturity model, practices are equivalent to controls in the standard and domains are the same as categories.

4.4 Related work

Ruan proposed a framework, *cybernomics*, comprising a combination of methods for estimating the value of digital assets [250]. The work suggests risk units as outputs but has a different focus from the research in our work, as the model of Ruan does not explicitly aim to capture the precise organizational architecture and structure of model organizations. Erola et al present a system that calculates cyber value-at-risk via sequential Monte Carlo simulations [98]. The work includes an example case study based on data provided by an insurance company using risk factors provided by a model vendor. The model proposed in our work might be used to structure a cyber value-at-risk simulation by identifying relevant facets of a target organization to inform simulation parameters and as such the model of Erola et al and ours might usefully be combined. Calvo and Beltran propose a model for adaptive security controls [50] that is only loosely related to the approach we propose, but may have relevance for insurers who wish to take an active interest in their customers' security posture from a remediative perspective.

The specific literature on security maturity models relevant to this work is not especially developed. One possible reason for this is that security maturity models are of significant use to managers and practitioners, but are relatively recent and therefore it may be too early to expect an empirical literature to have developed. For completeness, we review a selection of relevant papers here. Mettler provides insight into the development of maturity models from a design science research perspective [194]. Rea-Guaman et al introduce a taxonomy for assessing cybersecurity capability maturity models [240]. Kour et al discuss a cybersecurity maturity model specifically for the railway sector; however, the chapter provides a useful general overview of the various cybersecurity maturity models in existence [162]. Couretas provides a good general introduction to cyber-modelling [69]. Cebula and Young propose a taxonomy of operational cyber-risks, organized into actions of people, systems and technology failures, failed internal processes and external events [57]. This may prove useful for the purpose of insurance contract construction.

4.5 Chapter organization

The remainder of this chapter is organized as follows. Section 4.6 introduces the structure of our proposed modelling framework — that is, a conceptual tool for constructing models; it is not a model in itself — without providing precise details of the instantiation of its components. We wish to emphasize that the motivation of this chapter is to introduce a *flexible* framework adaptable to a range of modelling applications and that just a small subset of its potential is explored in this introductory work. Section 4.7 provides a theoretical overview of the tools that we recommend be used to deploy the modelling framework, namely entity relationship diagrams, security maturity models, and utility functions. Section 4.8 applies the modelling framework to three prototypical (i.e. fictional) organizations with a diverse range of security requirements. This provides a demonstration of how the modelling framework might be applied to real world organizations. Finally Section 4.9 discusses issues pertinent to real world implementation of the model, its advantages and disadvantages, and potential further work.

4.6 Modelling framework overview

As discussed in Section 4.1, the aim of the modelling framework introduced in this chapter is to connect a description of an organization's structure and systems with the economic parameters required to price a corresponding cyber-insurance policy via an assessment of the security posture of the organization using maturity models. This approach reconciles system-centric and economic treatments of a cyber-insurance decision problem. It is intended to structure and guide the cyber-insurance assessment and not to act as a universal model for cyber-insurance pricing.

At this introductory stage, we focus purely on introducing the components of the modelling framework and the required parameters rather than characterizing them in detail. The overall aim of the modelling framework is to support making insurance assessments that take account of both the overall structure of an organization and the intricacies of its systems. However, it is not our intention that the precise implementation of the modelling framework should be constrained; different target organizations and/or insurers may require different parameters. Specifically, requirements for the detail of an assessment vary across the insurance industry. Firms writing a large number of small limit policies are far less likely to require a detailed assessment of an insurance customer than a firm writing multi-million dollar limits for a select number of multinational organizations.

We aim to contribute a methodology for making insurance judgements and a set of parameters that are suitably flexible to be relevant across a range of possible analyses. In due course, full examples of how the modelling framework might be used will be provided. Section 4.7 will give examples on how our proposed modelling tools might be used to construct insurance policies and Section 4.8 will use the modelling framework to deliver case studies for three hypothetical organizations from separate industry groups to demonstrate its functioning and relevance.

4.6.1 Preliminaries

Representing organizational structure

We propose representing organizations using three layers: a management layer, a services layer, and a systems layer. The management layer describes the structure of the organization — including, but not limited to, what it provides and how it achieves it. The systems layer depicts the infrastructure used in the organization, such as servers, databases or other relevant elements such as manufacturing equipment. The link between the management and systems layers is depicted by service layers. The amount of detail captured by these layers is an important consideration. A description of every communication is likely to be unusable, but a simple overview of management structure too simplistic. The necessary level of detail for an insurance assessment is likely to vary for different types of organizations and insurer risk appetites. While the description we propose might be adequately represented in a table for simple instances, a diagrammatic representation represents a better choice in terms of legibility and ease of interpretation. With an appropriate schematic formulation, a well-constructed diagram might fulfil a diverse range of requirements without needing to be redrawn for different audiences. This work aims to demonstrate that entity relationship diagrams (ERDs) offer just the right level of abstraction to fulfil this purpose as their notational flexibility allows the user to implement their desired level of abstraction rather than prefiguring an organizational description. ERDs will be fully introduced and explained in Section 4.7.3.

Describing organizational security posture

Having established how to represent an organization and its systems, we now require appropriate risk metrics. We propose using criticality and sensitivity as measures for the components of the system that an organizational policymaker may wish to protect. The practice of framing security analyses using the categories of phenomena confidentiality, integrity and availability (often termed the 'CIA triad') is commonplace [9, 236]. As illustrated in Figure 4.1, confidentiality, integrity and availability

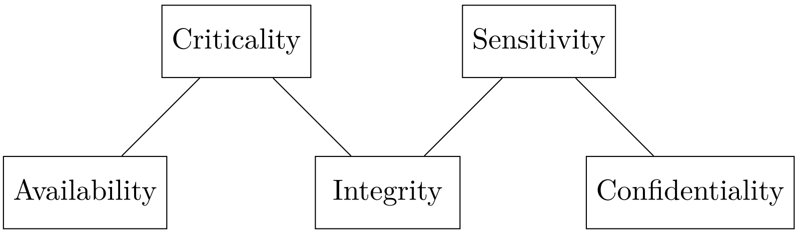

Figure 4.1: Relationship between criticality/sensitivity and confidentiality/integrity/availability

can be simplified to criticality and sensitivity. This assumes that availability maps to criticality; confidentiality to sensitivity; and some aspects of integrity map to criticality but others to sensitivity. For the purposes of the organizational assessments in this work, criticality and sensitivity are the simplest possible metrics that permit description of a wide range of different organizational types and architectures while achieving consistency with recognized industry norms. However, a wider range of categories could easily be deployed according to the requirements of the modeller.

As already outlined, it is important to consider how the structure of an organization and its systems may create risks of some form of cyber-attack and also how that attack might be detected. For a very simple organization, some form of 'fit and forget' controls might suffice, but for more complex organizations, the security posture is likely to evolve over time. A framework is accordingly required that is able to capture the notion of change and describe the communications within an organization and externally in a rigorous and structured manner. No single approach is likely to perfectly suit every instance and, consequently, it is probably optimal to settle on a class of models and choose a specific model for the case under consideration.

The modelling framework introduced in this chapter uses *security maturity models* to parametrize the security posture of an organization. The term maturity in this context means how developed the security controls and monitoring are in relation to an organization's processes, objectives and risk tolerance. An analysis of security maturity needs to produce outputs that can be compared in the context of the type of organizational structure, the details of the systems architecture and the characteristics of internal and external communications. One might construct a measure of the openness of an organization and its risk of being attacked or another measure of the desirability of an organization as a target. If correctly developed, a measure of security can be used to predict and analyze attacks and resultant losses, which in turn might be used to aid the construction of probability distributions, which are a key input to actuarial modelling in insurance.

Economic modelling

For the purposes of our modelling framework, an economic model is a function or framework that maps a number of (ideally parsimonious) parameters to an output. There are two clear categories of function of relevance to the challenge at hand: pricing formulae and utility functions. Pricing formulae are a category of functions that, as the name suggests, output the price for a financial instrument based on a number of inputs. Perhaps the most famous example in the finance literature is the Black-Scholes equation [83], which outlines the price for an option (the right but not the obligation to buy or sell a security). Utility functions stem from seminal work by von Neumann and Morgenstern[1] and are broadly speaking a mathematical description of the preferences of a decision-maker. Utility functions require careful formulation and evaluations when used in economic modelling. As this work is intended to outline a conceptual framework for combining systems and economic modelling, we will demonstrate the use of utility functions for model convenience, but it is important to be aware that there is a rich risk management literature (see, for example, Gollier [123] for a starting point) on their use and, for some modelling applications, there may be alternatives to utility functions that better capture the modelled trade-offs.

Within the field of security economics, Gordon and Loeb proposed a model connecting investment in information security and the probability of a security breach [128]. This spawned a derivative field of literature [209, 131, 267]. Whilst the Gordon-Loeb model is attractive in its tractability and simplicity, the choice of

[1] See Section 2.1.1

parameters is at best subjective and at worst contentious. Consequently, rigour, care and justification are required for an implementation of the model to have practical use. The usual technique in economics for populating model parameters is to use econometrics, which provides mathematical tools for deriving parameters from data in a broad sense. However, this is difficult for cyber-risk problems as there is a paucity of publicly available data and that available is other sparse or has a very short history. This is a significant (and at times valid) criticism of the use of economics in security and it is hoped that the contributions of this research illustrate a potential alternative means of deriving useful insights from economic models.

4.6.2 Modelling framework structure and parameters

Figure 4.2 outlines the structure of the modelling framework. The description of the maturity model draws on the models presented in Table 4.2 and described in Section 4.7.2. We wish to emphasise that the form of the security maturity model shown in the specification is not intended to be prescriptive. At these stage, we simply introduce the parameters that comprise the proposed modelling framework and to retain flexibility do not constraining their form and range of values. The aim of the modelling framework is to illustrate how models from different disciplines might usefully be combined to improve cyber-insurance pricing and not to dictate the best way to describe organizational security posture.

- The parameters C_i and S_i, for $1 \leq i \leq 3$, represent the components of the operations architecture — organizational structure (1), service layer (2), and systems architecture (3) — that deliver respectively those aspects of the operations for which criticality and sensitivity are required by the policy-maker to be protected.
- The parameters p_i, d_i, o_i, m_i denote respectively representations of the effectiveness in protecting C_i and S_i of the practices, domains, objectives, and maturity of the layers of the architecture as assessed by the chosen maturity model.
- The key aspects of this system dependency are the following: the interdepen-

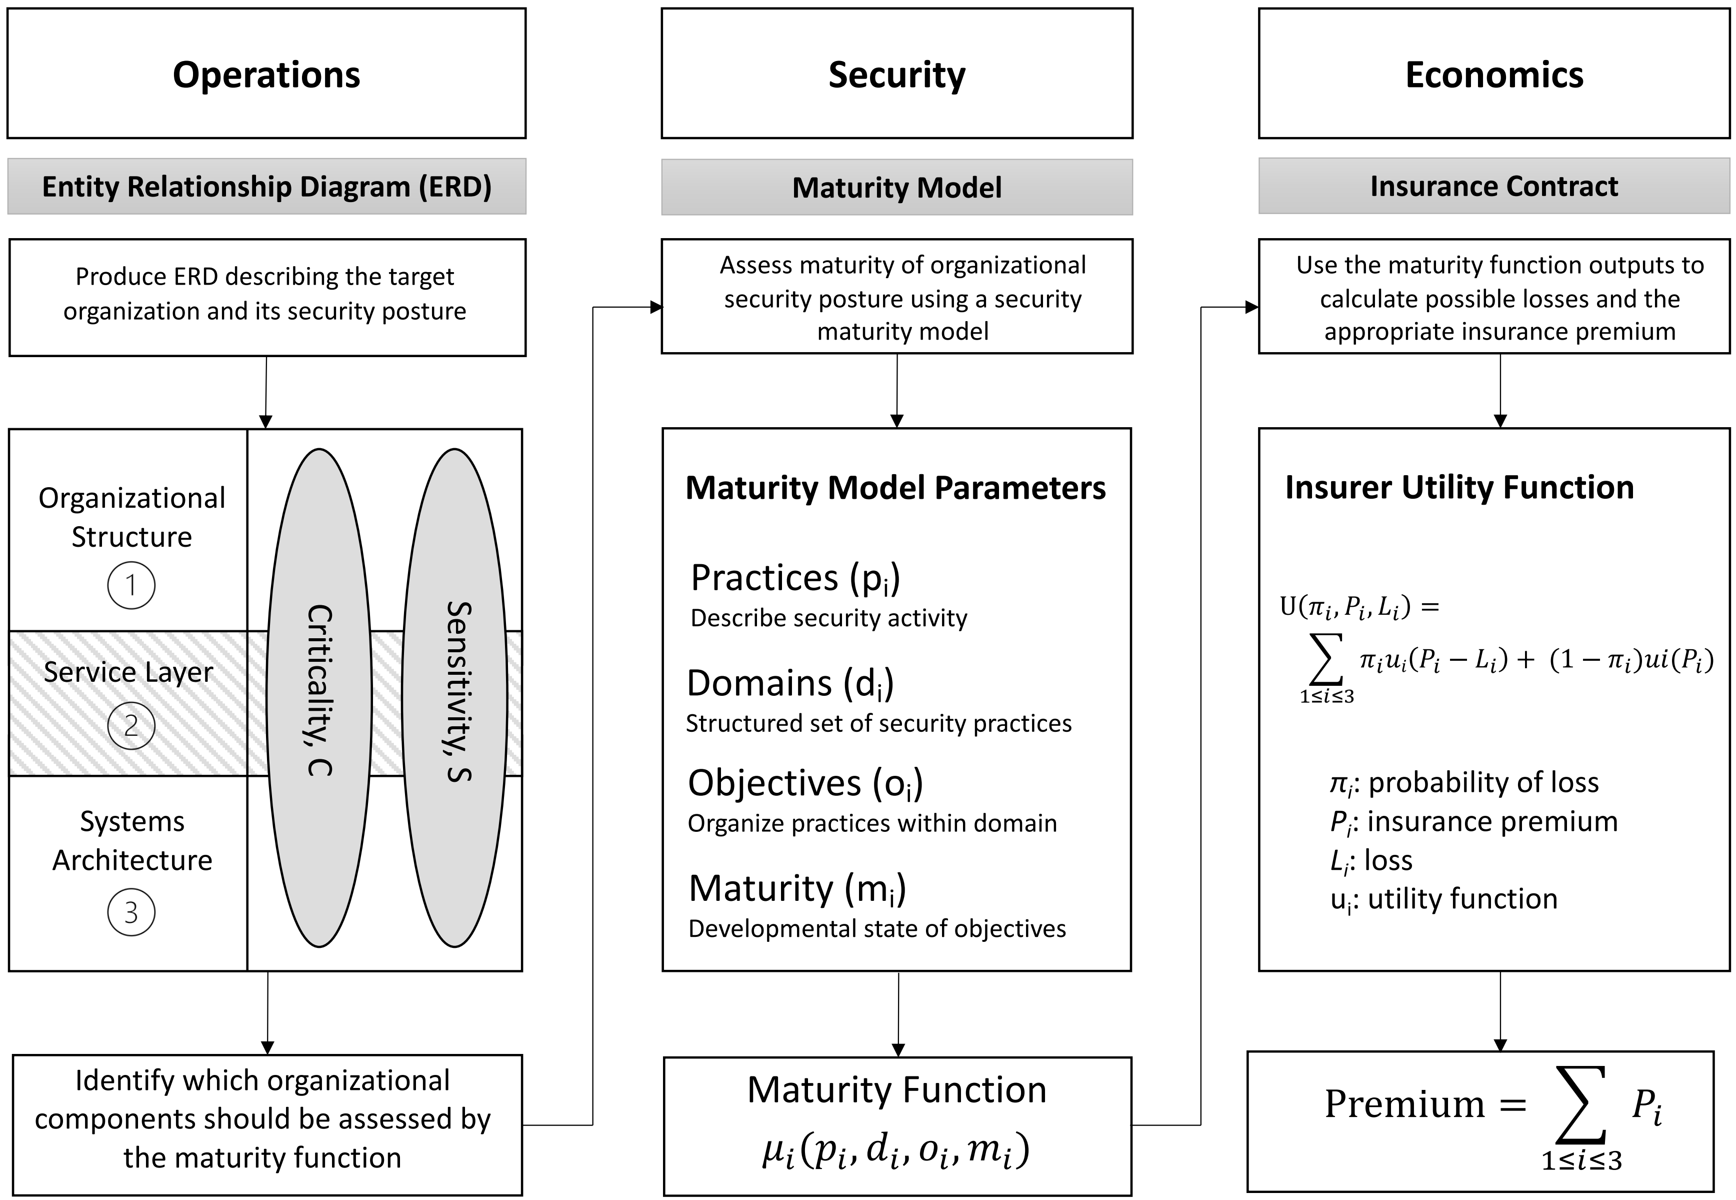
Operations
Entity Relationship Diagram (ERD)
Produce ERD describing the target organization and its security posture
Organizational Structure
1
Service Layer
2
Systems Architecture
3
Criticality, C
Sensitivity, S
Identify which organizational components should be assessed by the maturity function
Security
Maturity Model
Assess maturity of organizational security posture using a security maturity model
Maturity Model Parameters
Practices (p_i)
Describe security activity
Domains (d_i)
Structured set of security practices
Objectives (o_i)
Organize practices within domain
Maturity (m_i)
Developmental state of objectives
Maturity Function
$\mu_i(p_i, d_i, o_i, m_i)$
Economics
Insurance Contract
Use the maturity function outputs to calculate possible losses and the appropriate insurance premium
Insurer Utility Function
$U(\pi_i, P_i, L_i) = \sum_{1 \leq i \leq 3} \pi_i u_i(P_i - L_i) + (1 - \pi_i) ui(P_i)$
π_i: probability of loss
P_i: insurance premium
L_i: loss
u_i: utility function
$\text{Premium} = \sum_{1 \leq i \leq 3} P_i$

Figure 4.2: Conceptual overview of model

dencies between the system's layers and the interdependencies between the individual components of the system. These interdependencies are represented explicitly, through relationships and attributes, in the Entity Relationship Diagrams [59] that we describe in Section 4.7.1. Our proposed use of Entity Relationship Diagrams in this context is illustrated in the case studies that we present in Section 4.8.

- The maturity function $\mu_i(p_i, o_i, d_i, m_i)$ transforms the parameters of the chosen maturity model into parameters to be input into an economic model. The form and outputs of the maturity function depend on both the chosen maturity model and the target economic model. The maturity model may be based on or aligned with accepted security standards such as ISO27001, though this is not a requirement.

- A condition on the formulation of our model is that the analysis of the system's security posture that is provided by the maturity model must support the instantiation of the required utility function for the insurance contract.

- The loss-generating function, $L_i(\mu_i, C_i, S_i)$, specifies the monetary losses that may be incurred by a deviations in C_i and/or S_i from their values in the intended operating state specified by the policymaker. The mapping of these deviations to losses is determined by μ_i.

- The probability of a loss occurring due to a deviation in C_i and/or S_i is denoted by $\pi_i \in [0, 1]$. The calculation of the probability may be a function purely of the properties and/or maturity of the operations architecture component itself (endogenous factors) or take into account external influences or observations (exogenous factors).

- $P_i(\pi_i, L_i)$ are the components of the premium charged by the insurer from its assessment of each component of the operations architecture. The premium is a function of loss, L_i and π_i, the probability of that loss occurring. The final premium charged for the contract is $P = \sum\limits_{1 \leq i \leq 3} P_i$.

- The utility function for the insurer is defined as

$$\pi_i, P_i, L_i) = \sum_{1 \leq i \leq 3} \pi_i u_i(P_i - L_i) + (1 - \pi_i) u_i(P_i) \qquad (4.6.1)$$

The objective of an economically rational insurer is to set premia that maximize this utility function. u_i is a utility function that describes the risk preferences of the insurer. Partial utilities are used to allow specific properties of each architectural component to be captured.

4.7 Setting up the modelling framework

This section provides definitions and examples of the tools deployed in the modelling framework following the informal introduction of the relevant ideas in the previous section. The modelling framework uses three tools drawn from the fields of Operations, Security, and Economics:

Operations *Entity Relationship Diagrams* describe the organization and its systems

Security *Maturity Models* assess the security posture of the organization

Economics *Utility Functions* price the insurance contract

Before considering each of the three tools in detail, we briefly describe how they fit together to deliver the model. Entity relationship diagrams provide a description of three elements of the organization: its organizational structure, its systems architecture and the service layer that connects the two aforementioned elements. The complete entity relationship diagram may then be used to ascertain which of its components are susceptible to degradation in either criticality or sensitivity or both. With this description established, a maturity model is then used to describe the security posture of the organization. The maturity model is used to generate a maturity function, $\mu_i(p_i, d_i, o_i, m_i)$. The parameters practices (p_i) and domains (d_i) are provided by the maturity model. Objectives (o_i) are determined by a policy-maker. In simple cases, the insurer might be the sole policy-maker but for more complex organizations, high-level objectives from the insurer may be more precisely implemented by a policy-maker within the organization (for example, a CISO[2]). The maturity m_i specifies how well the objectives are currently met by the security practices of the organization.

The set of elements susceptible to degradation in criticality or sensitivity are potentially loss-generative in a monetary sense. In order to price an insurance policy, an insurer needs an expected loss for the policy and a probability of loss for that policy. In the model proposed here, the insurer uses a utility function to price the premium. It should be noted that the utility function provided for illustration is a demand-side utility function. Whilst an insurer should be risk-neutral if it has the ability to pool risk, for pricing purposes, it is assumed that the premium calculated by correctly applying the modelling framework is fair. Under the assumption that the fair premium equilibrates supply and demand of insurance in the absence of competition, then each insurer would charge the fair premium calculated using Equation 4.6.1 in the first instance. The maturity function provides the parameters needed to estimate probabilities (π_i) and losses (L_i) for each element of the organizational structure, which then generate a respective premium for each element. The sum of these premia then gives the overall premium for the policy. As an insurer gains experience of losses over time, it will likely gain insight into the relationship between practices and the parameters π_i and L_i and thus refine the objectives required to provide insurance cover or alter its pricing.

The utility function used for illustration of the modelling framework in Equation 4.6.1 uses partial utilities to allow specific parameters to be deployed for each layer of the organizational structure. Certain organizations may conduct operations that pose a high risk of loss from the perspective of the insurer, yet have a high level of maturity. Using the partial utility approach allows for pricing of cyber-insurance to be conducted by modification of existing industry group policy factors — this might be referred to as a modular approach to pricing by a practitioner.

It should be noted that reimbursement of a loss by an insurance company is not automatic. The holder of the insurance policy must first present a claim to the insurance company. The insurer then assesses the validity of the claim, and how much

[2] Chief Information Security Officer

of it to reimburse. Individuals fulfilling this function are known as loss-adjusters.

4.7.1 Entity Relationship Diagrams

Overview

A description of organizational structure and systems architecture requires a language that is able to express systematically the key components of a structure and their relationships. There is a rich literature on formal systems modelling, which is too large and diverse for us to review fully within the scope of this chapter. Examples of this discipline include the strongly compositional approach to systems modelling (e.g., among many, [33, 198, 66, 56]), UML [251, 45, 94], and system dynamics [111, 276].

In the model proposed in this research, we contend that the language of entity relationship diagrams has just the right level of conceptual analysis, abstraction, and descriptive power to sufficiently characterize the structure of an organization and its systems so as to support an analysis of its security posture for the purposes of an insurance assessment. The formal logic associated with the strongly compositional approach to systems modelling (see, for example, [66]), in which compositionality extends to the relationship between logical properties of systems and their execution dynamics, is not necessary for the purposes of insurance assessments. Insurance is always uncertain and, therefore, the modelling task is to *estimate* to the best of the modeller's ability rather than aiming to *prove* the properties of a system in this way. The weaker compositionality of ERDs, which is evidently sufficient to describe the composite structure of the organizations typically envisaged, will suffice.

The distinction between ERDs and UML is somewhat more subtle. UML is a language for creating diagrams whereas ERDs are a type of diagram. The nature of the modelling task for assessing organization structure relationships is a subjective task and consequently, the use of ERDs is more appropriate for our use case. System dynamics provides a methodology for rationalizing the behaviour of complex systems; it has a clear potential application to network modelling but is not particularly helpful for attempting to describe a diverse range of systems using a common set of models and parameters.

The ERD language has a grammar, originally introduced by Chen(1976) [59] for representing data structure, which describes the connectivity and properties of the components of the language.

- *Entities* are 'things', which can be distinctly identified. These might, for example, be organizational departments or assets, network resources, or control devices such as firewalls or intrusion detection systems (IDSs). Entities may be either informational or physical.
- *Relationships* describe associations between entities.
- *Attributes* are associated with both entities and relationships. They are parameters of the entity or relationship they describe.
- A system is described by connecting entities, relationships, and attributes in a graph. The representation of the connection may describe properties of that connection, such as cardinality or modality.

The description of a system via this language is known as an entity relationship diagram (hereafter referred to as an ERD). The components of an ERD map to the grammatical structure of natural language as depicted in Table 4.1.

ERD	Natural Language
Entity	Noun
Relationship	Verb
Attribute (Entity)	Adjective
Attribute (Relationship)	Adverb

Table 4.1: Mapping of ERD and natural language grammar

Properties of the connectivity between the different elements of an ERD, such as *cardinality* (whether the relationship is one-to-one, one-to-many, many-to-one, or many-to-many) and *modality* (whether the relationship is mandatory or optional). Of course, other logical relationships might be established and depicted as required.

ERD notation was originally developed as a model for data. A key motivation for its genesis was to reconcile different models for data (see Chen (1976) [59] for a detailed discussion and references). In this chapter, we aim to demonstrate that there is significant potential in using ERDs to describe complex systems more generally with great flexibility and compatibility with a range of logical requirements. For example, in formal systems modelling, combining systems diagrams is known as *composition*. ERDs deliver this by being readily combinable using appropriate combinators (such as lines or arrows) to link separate diagrams. This is a way of delivering what, in common language, is known as scalability.

ERD notation

Having established the motivation for deplying ERDs in this work, we now introduce their notation. The examples in this work are limited to entities with optional attributes, relationships without attributes, and lines to represent connectivity (Figure 4.3).

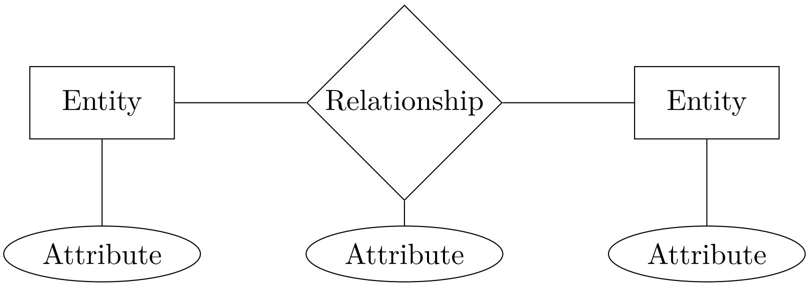

Figure 4.3: A simple entity relationship diagram

This provides the sufficient level of abstraction to apply the model introduced by this work and exemplified by the case studies in Section 4.8. However, we wish to re-emphasize that the level of abstraction is a *choice* and not a *prefiguration*. Depending of the desired intricacy of a particular ERD, more detailed notation might be required. This work makes use of Chen notation [59] to describe the components of a system. The foundations of this notation have already been outlined in Section 4.7.1 and for completeness we now describe the notation in full.

- *Entities* are something distinctly identifiable, such as data, an organizational department, a server, a firewall. They are the fundamental building blocks of the ERD language.
- *Weak Entities* are entities that cannot be identified by their attributes alone;

they require another entity to constitute a unique identifier. A simple example is a room, which cannot exist without a building.

- *Associative Entities*[3] are entities that resolve many-to-many relationships. For example, in a database of student — class relationships, one could construct an enrollment entity with a teach relationship to a teacher and which connects students to classes.
- *Attributes* are properties of entities or relationships. For example, in an organization, the Finance function might have the attribute 'department'.
- *Key Attributes* are attributes which uniquely identify an entity, for example, an employee identification number.
- *Derived Attributes* are calculated from other attributes. For example, an employee's age will increase over time but their date of birth remains constant. Age is therefore a derived attribute.
- *Multi-valued Attributes* are attributes that may have more than one value associated with the key.
- *Relationships* describe how entities interact.
- *Weak Relationships* are non-identifying relationships; formally the primary key of one of the related entities does not contain a primary key component of the other related entities.

The diagrammatic representation of the above components is provided in Figure 4.4. Chen used a solid line to indicate a mandatory relationship between components and a dashed line to indicate optional relationships. A richer set of combinators are provided by Barker [23] notation (Figure 4.5) and Crow's foot notation [99](Figure 4.6). Both notations allow for cardinality (number of relationships) and modality (whether the relationship is optional) to be fully conveyed. Chen notation only allows for a binary description of modality between two components.

[3] See [283] for a careful explanation

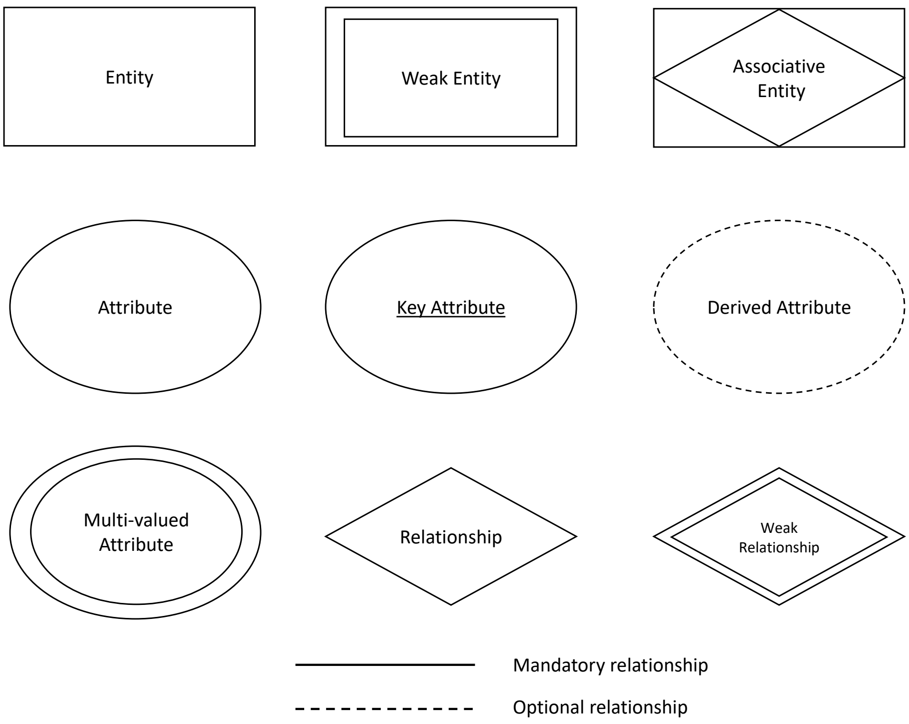

Figure 4.4: Chen notation

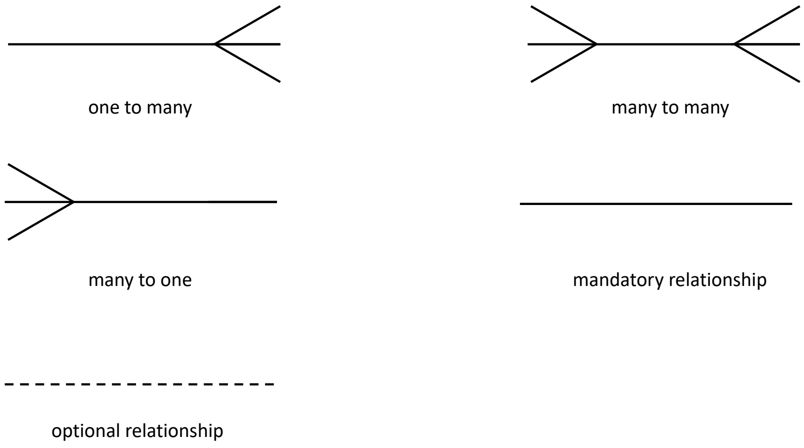

Figure 4.5: Barker notation

Example ERD

Figure 4.7 shows how the payroll function of an organization could be represented using an ERD. This example has been constructed to illustrate the variety of Chen notation that could be used to describe a system. It is intended as a stylized example rather than claiming to be an optimal representation of a payroll function. In the

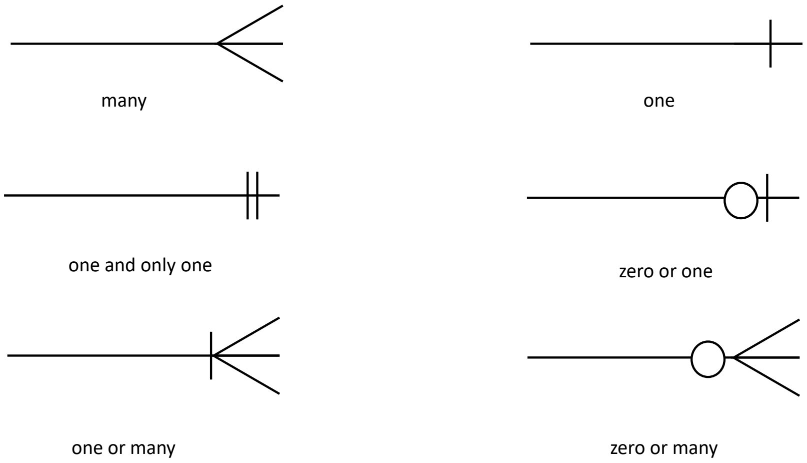

Figure 4.6: Crow's foot notation

example, there are two organizational departments, Finance and Human Resources, which are entities with an attribute denoting that they are departments. The finance department is responsible for paying employees, which are an entity with a multi-valued attribute 'resource' to recognize that employees may have multiple attributes. There is a relationship 'pays' between the finance department and the employees. There is a payroll database, which is a a weak entity with the attribute database; this database contains the entity records with attribute data. It is classified as a weak entity here as it exists to fulfill a purpose and without the payroll function, it would not exist. The contents of the database are delivered via an associative entity, 'records', within the service layer. The service layer in the model is a natural location for associative entities depending on the nature of the organization or entity being depicted as the service layer connects the management and systems layers of the organization in the model. The records have the keyed attribute, 'Data', as each employee record has a unique identifier. The shading in the diagram denotes which elements are critical and which are sensitive; the use of this will be explained in detail in Section 4.7.2.

The role of entity relationship diagrams within the model is to

- describe the regions or zones for which different maturities can be assigned,
- show for which of these regions criticality and sensitivity are required to be protected, and

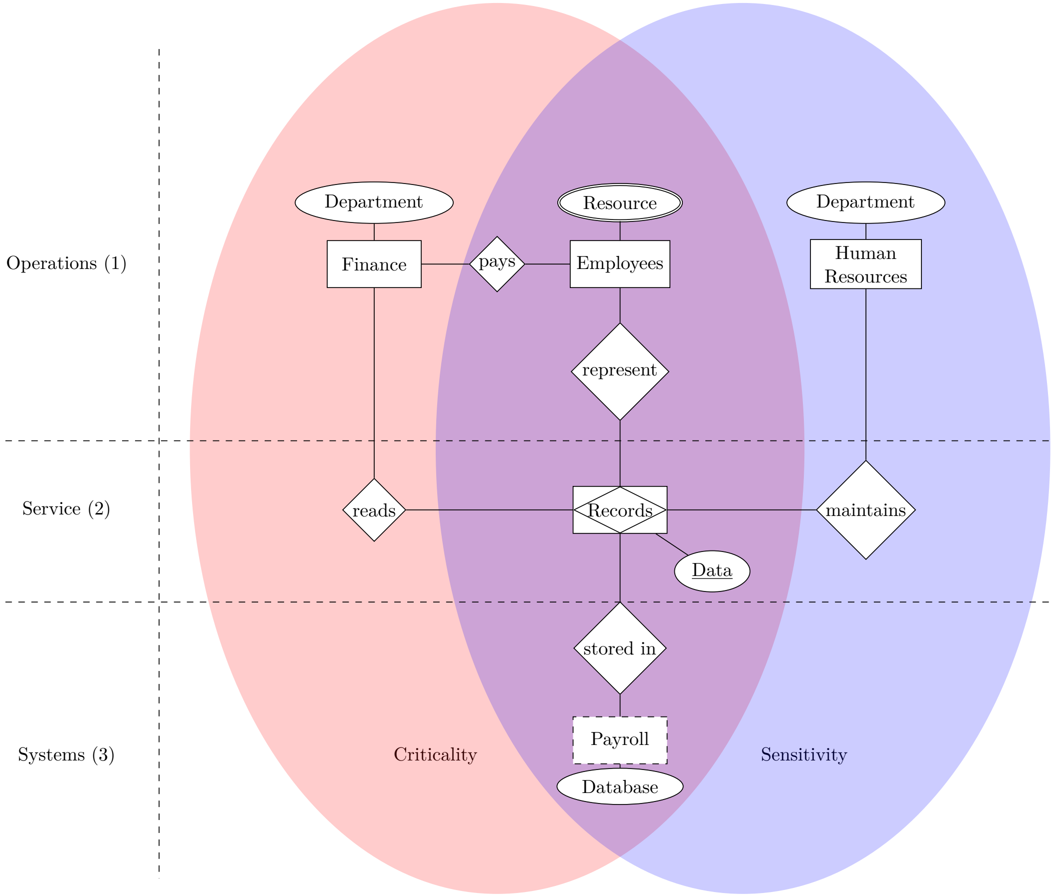

Figure 4.7: Payroll model with criticality and sensitivity

- represent the connectivity between different elements of the organizational structure.

The entity relationship diagrams generated in the organizational assessment could also potentially be used to identify *vulnerabilities* arising from connectivity within the organization and its systems so that an assessment of controls can be undertaken.

We contend that our choice of ERDs to represent the structure of the system is at the right level of abstraction to be able to capture the different security postures in different regions of the architecture without introducing complexity or detail that is not needed for this purpose. Furthermore, we observe that, if required, ERDs can be used to give more detailed descriptions of specific components of the description of an organization. This can be done in a manner that is compositional, with a more complex model of a component replacing a simpler one (or vice versa). Moreover, we would suggest that our choice of such a minimal representation of the architecture of the underlying system supports the scaling of our methodology to larger, more complex systems.

4.7.2 Security Maturity Models

Security Maturity Models (SMMs) are a tool used by organizations to describe their security posture in a structured manner. They achieve this by stipulating a framework on how to group processes and resources related to cyber-security and then assess how developed, or mature, is the posture of the organization. SMMs are not just a tool for self-analysis; they are gaining increasing prominence as a reference for the minimum standards required by external contractors to sensitive organizations, such as government or the military. There is no one standard SMM in current use. Table 4.2 lists the SMMs that appear to be in common use based on web searches. While we believe it to be comprehensive, it is not intended to be exhaustive. The exact specification and terminology vary from model to model, but most follow the basic structure of:

Model	Authoring organization
Personnel Security Maturity Model	UK CPNI (now part of NCSC)
IoT Security Maturity Model	Industrial Internet Consortium
IA Maturity Model (HMG)	NCSC
Cybersecurity Maturity Model Certification	US Department of Defense
Community Cyber Security Maturity Model	UTSA CIAS
ICS-SCADA	ENISA
C2M2	US DOE

Table 4.2: Currently used security maturity models

- *Practices* (p_i) describe single security activities and are *how* maturity is achieved
- *Domains* (d_i) provide a structured set of security practices, grouped by area.
- *Objectives* (o_i) are *what* the organization must do for a practice to be deemed as met.
- *Maturity* (m_i) states the developmental state of objectives, usually assigning a level from a predetermined set.

Any of these elements may be aligned fully or partially with existing cybersecurity standards. The preeminent global standard is ISO27001, which comprises controls organised by category. Within the maturity model, practices are equivalent to controls in the standard and domains are the same as categories. We assume here that the impact on premia of the degree of compliance with standards such as ISO27001 derives from the effect of the level of compliance on the parameters mentioned in the maturity model (Figure 4.2 and above).

Other examples of cybersecurity standards include the NIST Cybersecurity Framework in the USA, Cyber Essentials in the UK, the Essential Eight in Australia and the BSI IT-Grundschutz in Germany. Cyber-security standards generally need to have broad applicability and coverage whereas maturity models may be tailored to a specific use case. The maturity model allows for the same structure and rigour as the cyber-security standards but may place greater emphasis on certain elements of security.

To illustrate how this works in practice, we use an example from the CMMC[4]. A practice might be to disconnect a user after n minutes of inactivity, which is AC.L2-3.1.11 Session Termination in the CMMC. The practice is part of the Access Control (AC), one of 14 domains within the CMMC, which are aligned with the NIST standard SP 800-171. An objective to which this practice might belong could be 'Ensure unintended users do not gain access to the system', which might include other practices from the AC domain relevant to the organizational system. Within the CMMC, to attain AC Maturity Level 1[5], the practices Authorized Access Control, Transaction & Function Control, External Connections, and Control Public Information are required to be implemented. A further 18 incremental practices are required to achieve Level 2 for this domain.

[4] A good resource for the CMMC is CyberAssist (https://ndisac.org/dibscc/cyberassist), provided by the US National Defense Information Security and Analysis Center

[5] The CMMC is, at the time of writing, transitioning from v1 to v2 with the number of maturity levels reducing to 3. This example uses the v2 specification

Security maturity model effectiveness parameters

In Section 4.6.2, we stated that the security maturity model is used to produce a maturity function, μ_i, which takes as parameters p_i, d_i, o_i, and m_i which denote respectively representations of the effectiveness in protecting criticality and sensitivity of the practices, domains, objectives, and maturity of the components of the system architecture as assessed by the chosen maturity model. At that juncture, we refrained from giving a general characterization. We now explain how the parameters *might* be populated for the purposes of assessing payroll function in Figure 4.7 and the case studies that follow in Section 4.8. The precise and rigorous population of the effectiveness parameters for a maturity model is a significant and detailed piece of work and is intended to be undertaken by a specialist team within an organization (or, possibly, external consultants).

- p_i is a vector with elements taking values across $[-1, 1]$. The elements represent whether each practice in the maturity model is not met (-1), not relevant (0) or met (1).

- d_i is a vector with elements taking values across $[0, 1]$ representing the importance of each domain to the insurer or policymaker for their objectives.

- o_i is a vector with elements taking values across $[0, 1]$ representing how well the assessed organization meets the stipulated objectives.

- m_i states the overall maturity level achieved by the organization across $[0, 1]$. Different maturity models have different level structures; thus, $m_i = 0$ denotes complete immaturity and $m_i = 1$ denotes complete maturity.

These parameters would allow an insurer selling a specific insurance policy to quickly evaluate the maturity level of the security of organizations seeking insurance in a robust and systematic manner. An insurer might then compile a table or matrix of premium rates compared with maturity to allow automatic policy pricing.

Example security maturity assessment of payroll function

In the payroll function depicted in Figure 4.7, there are three entities in the Management layer: Finance, Employees and Human Resources. The Finance Department is defined as critical as without money being directed as required, the organization cannot function. Human Resources is defined as sensitive, as it is responsible for confidential data and maintaining its integrity. The service layer compromises records, which are read by Finance and maintained by HR. The records are stored in the payroll database, which resides within the systems layer.

	Criticality	**Sensitivity**
Management	(C1) Transactions must be processed when required and recorded	(S1) Employee data must be kept integral and confidential
Service	(C2) Records must be available	(S2) Records must only be accessed or modified by approved persons
Systems	(C3) The database must be available except for scheduled maintenance and must be integral	(S3) The database must have appropriate information security controls

Table 4.3: Objectives for payroll example

Table 4.3 states objectives for the payroll function that might be determined by an organizational policymaker. We use the notation C and S to number those objectives that relate to criticality and sensitivity, respectively. Insurance companies tend to use some standardized premium rates adjusted for various factors, which are typically multiplicative [247]. In the initial example payroll example, deviations from the objectives outlined in Table 4.3 with respect to criticality and sensitivity might incur losses. However, these might be mitigated by adopting security *practices* stated in the maturity model. First, each of the objectives needs to be mapped to *domains* from the maturity model (in this case, the CMMC). Table 4.4 shows one possible choice of d_i for the payroll example.

The next step in the procedure for maturity assessment is to populate the matrix of practices, p_i, according to which practices are met. This is clearly highly dependent

Domain	C1	C2	C3	S1	S2	S3
Access Control (AC)	0	0	0	0	1	1
Audit and Accountability (AU)	1	0	0	1	0	0
Awareness and Training (AT)	1	0	0	1	0	0
Configuration Management (CM)	0	1	0	0	1	1
Identification and Authentication (IA)	0	0	0	0	1	1
Incident Response (IR)	1	1	0	0	0	0
Maintenance (MA)	0	0	1	0	0	0
Media Protection (MP)	0	0	0	0	0	1
Personnel Security (PS)	0	0	0	0	1	0
Physical Protection (PE)	0	0	1	0	0	1
Risk Assessment (RA)	1	0	0	0	0	0
Security Assessment (CA)	0	0	0	1	0	0
System and Communications Protection (SC)	1	0	0	0	0	1
System and Information Integrity (SI)	1	0	1	0	1	1

Table 4.4: d_i parameters for payroll example

on the precise organizational policies and systems implementation and is consequently difficult to stylize in a meaningful way for the payroll example. However, as an illustration of how this might be performed, Table 4.5 summarizes the number of practices required for each domain within the maturity model to achieve either Level 1 or Level 2 within the maturity model. At time of writing, the required practices to achieve Level 3 within the CMMC had not yet been made publicly available.

For the purposes of the simple payroll example, we limit the scope of the assessment to the 17 practices that comprise Level 1 within the CMMC (Table 4.6). The reader interested in the detailed description of the practices is referred to the CMMC Model Overview[52]. Table 4.6 should be read as the necessary set of practices an organization should meet to achieve an insurance premium pricing of Level 1 in this example. Some practices apply only to certain layers; for example, AC.L1-3.1.22, details external connections, which would be specified at the service layer (2) within the organizational model. The extent to which an organization meets these practices in the assessment of its maturity will determine how much of a discount it would receive relative to the baseline premium charged by an insurance company.

Domain	Level 1	Level 2	Total
Access Control (AC)	4	18	22
Audit and Accountability (AU)	0	9	9
Awareness and Training (AT)	0	3	3
Configuration Management (CM)	0	9	9
Identification and Authentication (IA)	2	9	11
Incident Response (IR)	0	3	3
Maintenance (MA)	0	6	6
Media Protection (MP)	1	8	9
Personnel Security (PS)	0	2	2
Physical Protection (PE)	4	2	6
Risk Assessment (RA)	0	3	3
Security Assessment (CA)	0	4	4
System and Communications Protection (SC)	2	14	16
System and Information Integrity (SI)	4	3	7
Total	17	93	110

Table 4.5: Number of practices for each Domain in the CMMC by Level

CMMC Practice	Description	p_1	p_2	p_3
AC.L1-3.1.1	Authorized Access Control	1	1	1
AC.L1-3.1.2	Transaction and Function Control	0	1	1
AC.L1-3.1.20	External Connections	0	1	0
AC.L1-3.1.22	Control Public Information	0	0	0
IA.L1-3.5.1	Identification	0	1	1
IA.L1-3.5.2	Authentication	1	1	1
MP.L1-3.8.3	Media Disposal	1	0	0
PE.L1-3.10.1	Limit Physical Access	1	0	0
PE.L1-3.10.3	Escort Visitors	1	0	0
PE.L1-3.10.4	Physical Access Logs	1	1	0
PE.L1-3.10.5	Manage Physical Access	1	0	0
SC.L1-3.13.1	Boundary Protection	1	1	1
SC.L1-3.13.5	Public-Access System Separation	0	0	0
SI.L1-3.14.1	Flaw Remediation	1	0	0
SI.L1-3.14.2	Malicious Code Protection	0	0	1
SI.L1-3.14.4	Update Malicious Code Protection	0	1	1
SI.L1-3.14.5	System and File Scanning	0	1	1

Table 4.6: p_i parameters for payroll example

4.7.3 Pricing cyber-insurance using utility functions

For the purposes of this model, a cyber-insurance policy is a contract that provides reimbursement for losses, L at a cost of premium, P, which may be expressed as a percentage premium rate, p. Usually, an insurance buyer will not be able to buy unlimited insurance but will be offered *cover* up to a certain amount. However, the amount of cover an insurer provides represents a limit on its potential cash premium income and the optimal maximum limit offered will accordingly be carefully assessed by the insurer. In this model, we make the simplifying assumption that for the purpose of pricing premia, cover is *ex ante* equal to potential losses. The decision of the insurer is framed in terms of utility maximization in the sense established by von Neumann and Morgenstern [211]. The utility function, u_i, describes the risk preferences of the insurance company (in this case) and may take a number of forms. A commonly used utility function describes constant absolute risk aversion (CARA) preferences and takes the form

$$u(x) = \frac{1 - e^{-\alpha x}}{\alpha}$$

The coefficient of risk aversion

$$-\frac{u''(x)}{u'(x)} = \alpha$$

for this function, where $u'(x)$ and $u''(x)$ are, respectively, the first and second derivatives of the utility function. The choice of utility function would be important for real-world deployment of our model, but is a minor consideration for introducing the model.

A simple means of pricing an insurance contract is to fix the expected utility as a function of payoffs in two states: loss and no-loss [241]. As stated in Section 4.6.2, for the insurer in this model, this is

$$U(\pi_i, P_i, L_i) = \sum_i \pi_i u_i(P_i - L_i) + (1 - \pi_i) u_i(P_i) \tag{4.7.1}$$

where the suffix, i, represents the different components of the organizational architecture. While Equation 4.7.1 may appear simple, the calculation of π_i, the probability of a loss and L_i, the size of the loss are non-trivial. We now proceed to explain how these may be derived from entity relationship diagrams via security maturity models using the payroll example in Figure 4.7.

The security maturity model has two key effects on the insurance assessment:

how *likely* and how *large* might losses be given the maturity assessment of the organizational architecture. Assume that the probability of a loss, π_i may be represented as

$$\pi_i \sim D(\mu_i, C_i, S_i) \tag{4.7.2}$$

where D is a distribution function, relating the expected probability of loss L to criticality, sensitivity and the maturity function, μ_i (where $\sim$ denotes 'is distributed as'). Losses are given by

$$L_i(\mu_i, C_i, S_i) = \lambda_i(\mu_i, \Delta C_i, \Delta S_i) \tag{4.7.3}$$

$$\Delta C_i = C_i - \bar{C}_i, \quad \Delta S_i = S_i - \bar{S}_i \tag{4.7.4}$$

where λ_i is a function specific to the insurer for estimating economic losses from deviations in criticality (C_i) and sensitivity (S_i) informed by the maturity function, μ_i. $\bar{C}_i$ and $\bar{S}_i$ represent, respectively, the intended state of criticality and sensitivity. These states will vary across industries, with varying emphases on specific risks.

The reason for having different λ functions for each insurer is different insurers may have different information, experience or appetite for particular risks. For instance, one insurer might be highly exposed to one industry and thus there is negative utility associated with writing further insurance to that sector; for another, it might diversify the existing portfolio and thus have positive utility.

It is important to note that L_i is not required to be strictly positive for all layers of the model. The operations of the organization, described by the management layer, determine the scope of possible losses to the business that may result from deviations in criticality and sensitivity as previously described. However, it is possible that additional maturity in terms of practices at the level of the systems layer may reduce the scope of potential losses. Therefore, it is possible that the ultimate components of the premium calculated by the insurer related to the service and systems architecture layers, P_2 and P_3 respectively, may be negative depending on the assessment of maturity. This is consistent with the practice of insurers offering discounts on baseline policy quotes for meeting certain conditions. However, it may be the case that certain elements of the architecture create greater risks (for example, external connections to a database), and therefore merit charging more premium in that case. Figure 4.8 illustrates how the organizational representation first introduced in Figure 4.2 outlines a possible procedure for estimating the parameters L_i for the different layers of the organizational structure.

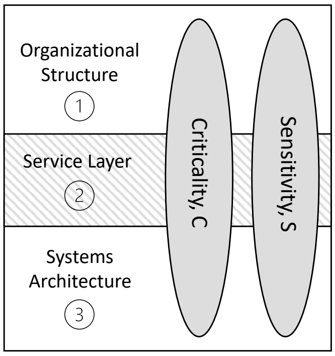

L_1: Estimate potential operating losses due to deviations in criticality or sensitivity. These reflect the business operations of the organization, its revenue streams and intellectual property.

L_2: Identify features of the service architecture supporting interaction between operations and systems that mitigate (-) or amplify potential losses (+).

L_3:Identify controls that mitigate the scope of potential losses (-)

Identify vulnerabilities that may arise from the structure of the system (+)

Figure 4.8: Procedure for estimating potential losses

One advantage of the systematic ERD-maturity-utility modelling framework is that once calibrated by an insurer according to its pricing requirements, the framework could provide consistent comparisons between different insurance risks. This allows for its potential use in accumulation modelling, sometimes known as clash exposures, in the insurance industry [277]. The aim of this modelling is to determine the risk of correlation in losses if certain categories of event occur. A well-diversfied insurance portfolio should not contain policies with high levels of clash risk. This is a potential problem for cyber-insurance in particular as some of the usual diversification strategies such as insuring organizations in a wide range of locations, industry groups and annual revenue may be less effective given the interconnected nature of cyber-space and the concentration of computer operating systems.

The expected-utility-maximization approach adopted here for pricing premia is compatible with other established methodology and models in the literature. The Gordon-Loeb security breach functions [128] could be used as the probability distribution for losses in Equation 4.7.2. This has precedent in the literature [189, 267]. For the premium pricing model here, one could easily replace investment, z in the Gordon and Loeb model with the maturity function μ which could be justified by the notion that increased investment in security should bring increased security maturity; otherwise, the investment is not a rational activity. Beyond the Gordon and Loeb Model, Mazzoccoli and Naldi provide a helpful survey of security breach probability models that might aid a corporate decision-maker or insurer model the relevant threats [190]. This approach over time could be blended with a distribution bootstrapped from experience cyber-losses but any available datasets currently have significant issues with heterogeneity of losses and limited time coverage rendering meaningful statistical inference difficult.

The means of estimating the distribution of expected losses for each layer is non-trivial and there are two possible approaches — construct a distribution *a priori* or fit one from observed data. The latter is what the actuarial department of an insurance company specializes in. For expositional simplicity, we will use the security breach functions (SBF) from the Gordon & Loeb model [128] and specifically, type 1 functions. In the Gordon & Loeb Model, it is assumed that an information set has an inherent probability of breach or vulnerability, v, which is reduced by investment, z. The type 1 SBFs take the form

$$S^I(z, v) = \frac{v}{(\alpha z + 1)^\beta} \tag{4.7.5}$$

α and β are measures of the productivity of information security. For the purposes of populating the utility function, we assume that the probability of a breach follows the Gordon & Loeb type 1 SBF, but replace investment with the met-objectives parameter, o_i, from the maturity model. This is justified by the assumption that to achieve objectives, the insurance-seeking firm must have made an investment in security to implement the practices that allow the objective to be deemed met:

$$\pi_i = \frac{v(C_i, S_i)}{(\alpha_i o_i + 1)_i^\beta} \tag{4.7.6}$$

The higher the met-objectives parameter, the lower the probability of a breach from the insurer baseline and the lower the premium. $v(C_i, S_i)$ represents the insurer's best estimate of loss probability for the insured component. α_i and β_i can be adjusted as required to reflect the insurer's confidence in the efficacy of the objective, or updated over time based on experience. This allows the fundamental form of the analysis and contract to be preserved while allowing for an update on the pricing model based on experience.

The final component required to be defined is the loss function. To compute the loss function, we assume that the meeting of practices stipulated in the maturity model limits the potential for losses to arise via lowering the potential degradation of criticality and sensitivity. Then

$$\lambda_i(\mu_i, \Delta C_i, \Delta S_i) = \frac{\Lambda^{(C_i)}\Delta C_i + \Lambda^{(S_i)}\Delta S_i}{(1+p_i)^{\gamma_i}} \tag{4.7.7}$$

Λ represents the maximum potential estimated loss to the layer from a degradation in confidentiality or sensitivity, while γ is a parameter for the productivity of effectiveness of practices. This might be aligned to the maturity parameter, m_i, for example.

The final step required to populate the utility function is to consider policy limits and coverage. Until now, we have considered the insurance premium purely in monetary terms, which is sufficient for a conceptual overview of the utility function. In reality, the premium is usually decomposed into two components: a premium rate, r (usually expressed in %) and an amount of coverage (also known as a limit), up to which losses will be reimbursed. The implementation of the limit may vary; in this example, we assume that there is a limit for each layer of the organization for consistency with the model. If we assume that coverage is a function of maturity, then

$$P_i = r_i \kappa_i m_i \tag{4.7.8}$$

where κ_i represents the policy limit determined by the insurer for the assessed component.

For simplicity, let us assume that $u_i(x) = x$ and for this example that utility is linear.[6] Taking the above definitions and simplifying Equation 4.7.7, we may now

[6] Optimizing the commonly used economic utility functions is typically algebraically intractable

write the utility function

$$U = \sum_i r_i m_i \kappa_i - \frac{v(C_i, S_i)}{(\alpha_i o_i + 1)^{\beta_i}} . \frac{\Lambda^{(C_i)} \Delta C_i + \Lambda^{(S_i)} \Delta S_i}{(1 + p_i)^{\gamma_i}} \tag{4.7.9}$$

with the constraint

$$\Lambda^{(C_i)} + \Lambda^{(S_i)} \leq m_i \kappa_i \tag{4.7.10}$$

Estimating the loss component of the cyber-insurance premium pricing function is less well developed in the literature. This is likely in part because this task is very dependent on the organization in question and the value of its information assets and thus it is hard to model in generality. Woods et al describe the mapping of questions on insurer proposal forms with existing ISO security standards [294]; this suggests that there is some evidence from industry practice that the idea of using security maturity models for pricing insurance contracts has merit. Likewise, Romanosky et al, describe how insurance companies tend to use factors and modifiers for policy pricing based on observations about the insured entity [247]. The use of maturity models to systematically produce such modifiers would be more objective and if standardized across insurers might lead to more efficient risk transfer.

It is important to note the following: insurance companies have long experience of the magnitude and frequency of losses that arise in organizations based on their size, industry sector, and location. Consequently, their calculations of premia will start from a baseline determined by these considerations. The contribution of the methodology proposed here is to provide a framework for calculating the effects of cyber-based risk on the frequency and magnitude of losses. This is achieved through a security analysis of the relationship between the operational structure of an organization and its information systems.

and requires numerical methods.

4.8 Using the modelling framework

4.8.1 Case studies

This section provides a thought experiment in which we explore how our methodology applies to a range of organizations with differing threat environments and security postures. The analysis is structured to correspond to the modeling approach depicted in Figure 4.2. First, the operations of the prototypical organization are described from a security perspective. The key elements of security maturity in relation to protecting the criticality and sensitivity of the operations, service, and systems layers of the organization are then discussed. Finally, the potential resultant economic losses from degradation in criticality or sensitivity are considered. In real world examples, two prominent threats are data breaches and ransomware, which compromise sensitivity and criticality respectively although some ransomware attacks also involve data leakage in which case both criticality and sensitivity are compromised.

In this thought experiment, we give an analysis of three prototypical (fictional) organizations from three different industries:

- Online retail
- Consumer banking
- Pharmaceuticals manufacture

These industries have a common basic requirement for information security yet the emphasis and implementation will be different in each. In *online retail*, availability of the elements of the organizational structure that process and fulfil sales is paramount. For the *consumer bank*, there is an equal emphasis on confidentiality and integrity; while availability is important, it arguably cannot come at the expense of risking loss of customer funds. Within *pharmaceuticals manufacture*, integrity is of primary importance in the manufacturing process; in research and development, confidentiality is also very important. In addition to the specific details relevant to each, all three industries are customer facing and therefore have external information channels, which might be exploited by an attacker.

The intention with the following case studies is to show how the model introduced in the previous sections might be applied across a range of potential organizations. This flexibility is important for an insurance company that is likely to insure organizations from different industries across a single line of coverage such as cyber-insurance. We make no claim that the analysis that follows represents the optimal security posture for the different industry groups, which would of course require a thorough analysis of their precise systems and operational practices. Rather, the following analyses should be considered as stylized examples. The prototypical organizations are constrained to three departments which prioritize, respectively, confidentiality, integrity and availability. As explained in Section 4.6.1, integrity may be considered as the intersection of criticality and sensitivity. For expositional simplicity, only basic features of ERD notation are used as they provide sufficient detail for demonstrating the use of the model in the case studies. This does not preclude the use of richer notation to describe complex real-world organizations such as those with delicate intellectual property considerations, for example. The organizational departments are *entities*, which sit in the operations layer of the entity relationship diagram. Each department has an associated *entity* in the systems layer, connected by an *entity* in the service layer. The interaction between each *entity* is described by *relationships*. *Attributes* convey descriptions or details of the entities or relationships that are important for the maturity assessment. This allows for potential attack surfaces to be quickly identified with a cursory read of the ERD for each organization.

In each of the case studies, both a maturity assessment and a computation of premia is required. We provide these in some, though not complete, detail for the case of *online retail* and remark that similar analyses can be established for each of *consumer bank* and *pharmaceuticals manufacture*, the details of which are, for brevity, elided.

4.8.2 Online retail

Operations

The online retailer in this example uses a website to market and sell goods to customers. The operations of the retailer (Figure 4.9) are represented by the entities *customer service operations*, *stock control*, and *customer interaction*, which respectively require availability, integrity and confidentiality to be protected. For a retailer, the foremost security priority is availability/criticality as this presents the greatest potential source of revenue losses and therefore would be the risk that the retailer would place greatest emphasis on insuring against. Such insurance is known as business interruption.

Confidentiality is also important for the online retailer given that it handles large quantities of customer data including payment details that may be required to be protected by legislation in the countries in which it operates. This places the online retailer at risk of a regulatory fine for data breaches and also loss of reputation of customer confidence. However, this risk is secondary in terms of immediate potential direct losses relative to availability. The customer interaction department is responsible for dealing with customers. There is a distinction to be made however, between communications such as marketing and automatic notifications.

Integrity is the lowest security priority for the online retailer. The prototypical retailer has a stock control department which is responsible for recording the stock inventory for the retailer. In an optimally run retailer, this would be robust and in real time, but this provides an illustration of the concept of maturity in this context. In a retailer with low maturity of stock control, updates of inventory might be manual and thus there is potential risk of an out-of-stock item being sold to a customer. By contrast, a high volume, established online retailer will likely have sophisticated inventory management technologies that aim to prevent such a scenario from occurring. However, loss of integrity is an isolated business risk in this case and would therefore likely receive relatively little priority in insurance terms.

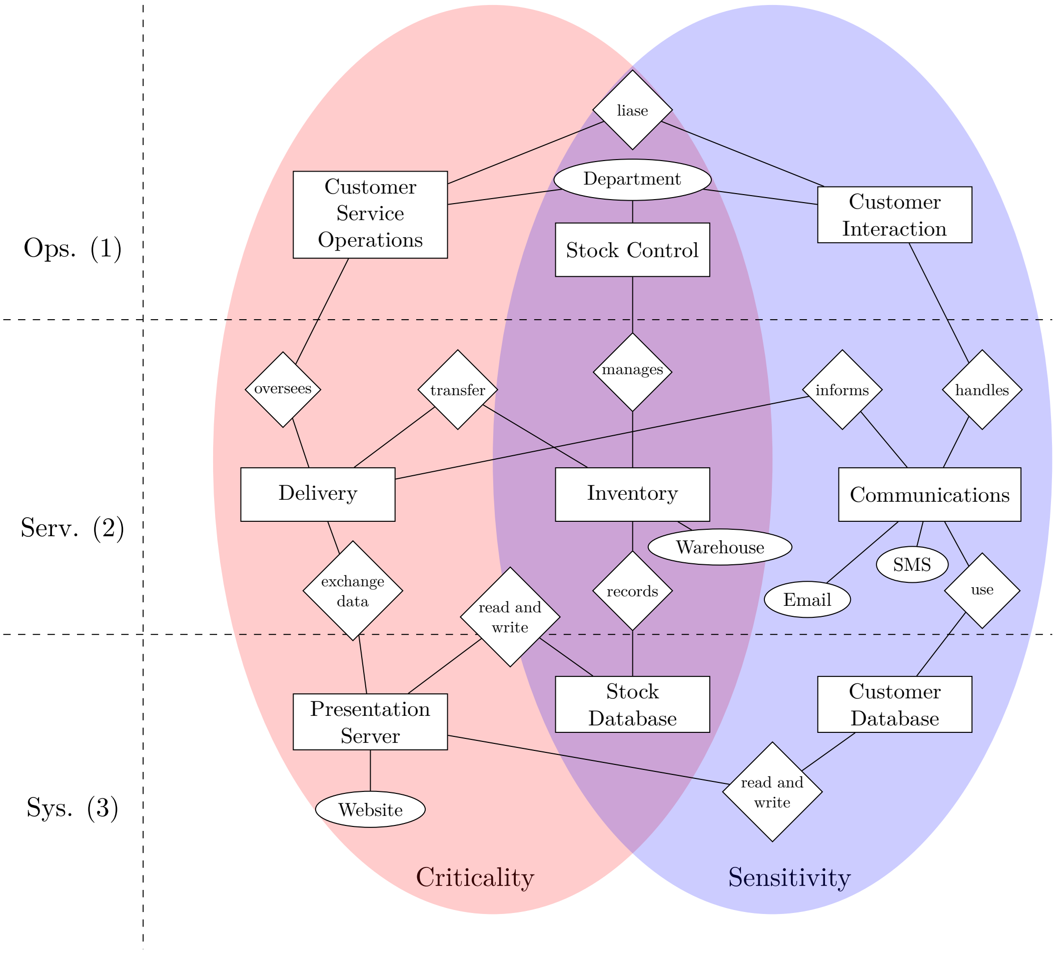

Figure 4.9: Entity relationship diagram for online retailer

Maturity assessment

As stated in Section 4.6.2, a requirement of our model is that the analysis of the system's security posture provided by the maturity model must support the instantiating of the required utility function for the insurance contract. The analysis must deliver parameters representing the maturity of protecting the security parameters criticality and sensitivity, which as previously described are delivered via practices organized into domains and objectives. For the online retailer, the most important security parameter is criticality and specifically the level of safeguards or redundancy in the system that may keep the operations available. The second most important output is an assessment of the protection of confidentiality and the scope for loss of data. The cost of such potential breaches should be treated separately within the utility function rather than the maturity model.

For an online retailer, one might start the maturity assessment by identifying domains as a baseline for calculating the extent to which the security posture is achieved:

1. Cloud usage
2. Incident response
3. Access control
4. Maintenance
5. System and communications protection
6. System and information integrity.

These are intended to provide the necessary set of domains to instantiate the utility function of the insurer rather than provide a complete description of the online retailer's security maturity. For this example, we will not stipulate individual practices within the domains as these are specific to the precise implementation of the system and controls. To move from domains to maturity, we will consider a number of sample objectives, which may also be regarded as insurability criteria. The extent to which these objectives are met by an insurance client in turn determine the maturity and thus the discount or premium relative to the insurer's standard baseline. Table 4.7 states a few sample objectives, which are intended to serve as examplars rather than comprising an exhaustive set of objectives. In a full insurance assessment, a substantive matrix of objectives and assessment criteria would be required.

Within our model, the maturity function is calculated for each layer — Management/Operations, Services, and Systems — of the organization. One resultant consideration for the insurer is how to measure the required parameters of the maturity function for each layer and thus calculate them. For the operations/management layer, this is most likely achieved via asking structured questions of the organization or inspecting existing policies should the insurance buyer be willing to share them.

The service layer would require a combination of the approach for the operations/-management layer and a systematic evaluation of controls. For the systems layer, the controls can likely be verified more rigorously via, for example, perimeter scanning of defences. The assessment of the operations/management layer is the most subjective and therefore most uncertain and that of the systems layer the most objective and therefore most certain. This uncertainty will be reflected in the formulation of the utility function in Section 4.8.2.

	Criticality	**Sensitivity**
Management	(C1) Incident response plan in place and tested	(S1) Information controls deployed (S2) Staff training on handling sensitive information enforced and documented
Service	(C2) Customer facing services must have minimal downtime	
Systems	(C3) Systems should not have vulnerable points of failure	(S3) Databases must be secured against unpermitted modification

Table 4.7: Objectives for online retailer maturity

Table 4.8 provides domain importance (d_i) parameters for the objectives listed in Table 4.7. As described in Section 4.6.2, it is important to understand how the analysis presented above depends upon the specific interdependencies in the system. In this case, we can see this by considering Table 4.7.

Domain	C1	C2	C3	S1	S2	S3
Cloud Usage	1	1	1	1	0	1
Awareness and Training	1	0	0	0	1	0
Incident Response	1	1	0	0	0	1
Access Control	0	0	0	1	0	1
Maintenance	0	1	0	0	0	0
System and Communications Protection	0	1	1	1	0	1
System and Information Integrity	0	1	1	0	0	1

Table 4.8: d_i parameters for online retailer maturity

In each of the components of Table 4.7, the C and S values are determined not only by the entities directly concerned but also by the relationships of these entities to other entities as specified by the ERD for the online retailer's system (see Figure 4.9). For example, the criticality assessment of the 'Presentation Server', in the System Layer, is dependent on the criticality assessment of ‘Delivery’ in the Service Layer and, also in the Service Layer, on the sensitivity assessment of the ‘Customer Database’, and the integrity assessment of the ‘Stock Database’; and so on.

Economic losses

The online retailer will primarily require coverage for losses from business interruption. In an insurance policy this might be specified as coverage up to a certain number of business days with an initial deductible. Coverage may also be required for potential legal and investigative costs associated with a data breach and loss of customer data. In order for the insurer to compute a premium for the policy, as per Equation 4.7.1, there are two variables needed for each layer, i, of the ERD representation for the online retailer (Figure 4.9): π_i, the probability of a loss, and L_i, the expected magnitude of a loss. As emphasized in Section 4.7.3, the maturity model is expected to deliver the potential deviation from the baseline assumptions of the insurance company for the firm in question.

Computing premia

As discussed in Sections 4.6.2 and 4.7.3, the utility equation for the insurer takes the form

$$U = \sum_i r_i m_i \kappa_i - \frac{v(C_i, S_i)}{(\alpha_i o_i + 1)^{\beta_i}} . \frac{\Lambda^{(C_i)} \Delta C_i + \Lambda^{(S_i)} \Delta S_i}{(1 + p_i)^{\gamma_i}} \tag{4.8.1}$$

To populate the utility function parameters, the approach discussed here will require the maturity, m_i of components; the extent to which practices, p_i, in the maturity model are met; and the extent to which objectives, o_i, specified by the insurer in the maturity model are met. For simplicity in this example, we assume that $\bar{C_i} = \bar{S_i} = 1$ and that both C_i and S_i take values across $[0, 1]$. This provides a simple means of combining these parameters with a potential maximum loss or policy limit.

Table 4.9 illustrates the parameters that might be used to populate Equation 4.8.1.

The following assumptions are made:

- Simulations would be run to estimate deviations in criticality and sensitivity under various scenarios. A starting point would be Lloyd's of London realistic disaster scenarios (RDS) for cyber incidents.
- The values for maturity, practices, and objectives are purely indicative in this example. For simplicity, in this example, the parameters representing maturity, practices met, and objectives met are set to the same values for each assessed component. As discussed in Section 4.8.2, the objectivity of assessment and accordingly ability to deem practices met increases from the operations layer, through the services layer, to the systems layer. The illustrative values for $m_1 < m_2 < m_3$ are deliberately chosen to reflect this.
- The maximum potential loss per layer is purely illustrative and the policy limit for each layer in the example is set to this.
- The loss probabilities are purely illustrative; an insurer would either estimate these using actuarial methods or based on real loss experience.
- The productivity parameters are set to 1 for simplicity. The online retailer operates a fairly open system that prioritizes criticality and thus the efficacy of controls may be reduced relative to a more closed system.
- The premium rate would be solved for using standard numerical methods of choice. The focus of this discussion is to illustrate how to produce an assessment, not to estimate specific premia.

The preceding analysis is intended only to provide an indication of how a utility function might be constructed. However, the broad form should be applicable across the case studies. The equation could potentially be solved analytically using readily available software; for more complex formulations, Monte Carlo simulations of losses would be an acceptable strategy.

Parameter	Description	Cust. Serv. Ops. (1)	Inventory (2)	Cust. Data (3)
m_i	Maturity	0.5	0.6	0.7
p_i	Practices met	0.5	0.6	0.7
o_i	Objectives met	0.5	0.6	0.7
Λ_i	Maximum potential loss per layer	\$100k	\$200k	\$300k
κ_i	Policy limit	\$100k	\$200k	\$300k
$v(C_i, S_i)$	Loss probability	0.05	0.02	0.01
$\alpha_i, \beta_i, \gamma_i$	Productivity parameters	1	1	1
$\Delta C_i, \Delta S_i$	Potential degradation in C and S	Derived from simulated scenarios, e.g. Lloyds RDS		
r_i	Premium rate	Outputs of chosen solution method of utility function		

Table 4.9: Summary of utility function parameters for online retail example

4.8.3 Consumer Bank

We represent the prototypical consumer bank via the entities *records*, *payments* and *customer interaction*, for which respectively availability, integrity and confidentiality are required to be protected. The basic operation of the bank is as follows: customers register transactions either withdrawing or depositing funds. These transactions are mandated via interfaces, such as banking cards or online banking facilities. The transactions are processed by the payments department and, if validated, records are then updated.

In terms of security objectives, integrity of the payments system is marginally most important as manipulation of this entity places the bank at greatest financial risk. Confidentiality of customer data is also of high priority, in particular appropriate protocols need to be in place to ensure customer accounts are resistant to misuse. This is typically achieved through the use of password mechanisms and other means of verifications that have become more sophisticated over time. Finally, availability of the records system need to be protected. While this is something a bank would strive to achieve, it is not uncommon of for banks to implement system downtime for maintenance while allowing routine payment processing to continue to take place. Further, in many banking systems, there is a lag between a payment being taken and its being recorded; for credit card transactions, this can be several business days.

Maturity assessment

We proceed to outline the maturity assessment process for the consumer bank in a similar fashion to Section 4.8.2 for the online retailer, though with a lesser level of detail. The issue of banking security is complex and well-studied, and in this relatively brief discussion it is not possible to conclusively cover the issue. Chapter 12 of [11] provides a lively discussion of both the history of security models used in banking and contemporary methodology. The priority of the assessment is to deliver how susceptible the banking system is to integrity attacks. A formal model applicable to the security policy of a bank was proposed by Clark and Wilson [61]. In brief, the model describes how an unconstrained data item (UDI) may be transformed into a

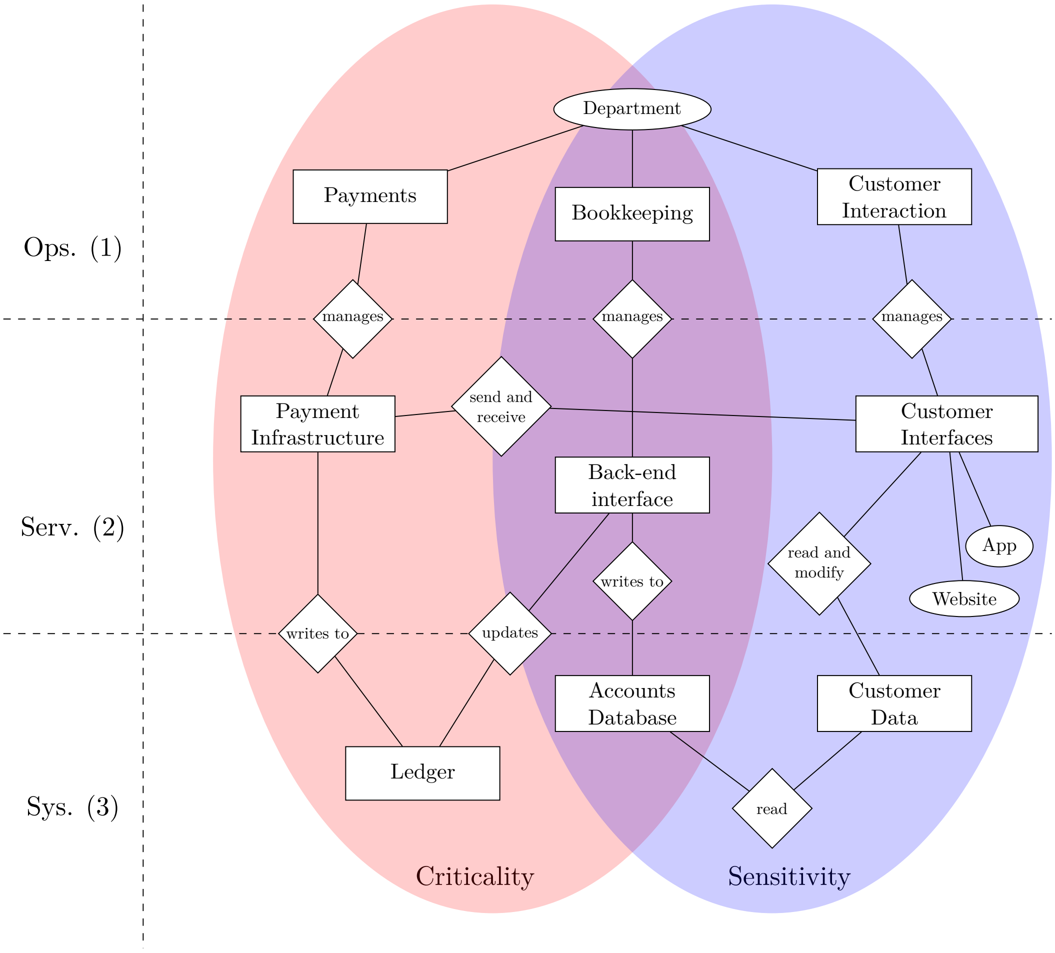

Figure 4.10: Entity relationship diagram for consumer bank

constrained data item (CDI). The integrity of a CDI must be preserved, which is validated through integrity verfication procedures (IVPs). Modification of CDIs is effected via a transformation procedure (TP), which under some circumstances may transform a UDI into a CDI.

Handling and mitigating losses is routine activity for a consumer bank; the role of cyber-insurance is to cover losses outwith the ordinary course of business that are generated either by a failure in technology or circumvention of controls. Effectively, the relevant parameters the maturity model for the consumer bank ought to deliver are:

- How likely is a catastrophic event?
- What might it cost?

As noted in the aforementioned discussion by Anderson, malicious insider risk is one of the foremost potential loss generators for a consumer bank. One area of insurance coverage that is relevant for a bank is technology errors and omissions coverage. This would deliver compensation for damages caused by coding or system configuration errors. Such an assessment would need to be kept highly confidential as it otherwise might provide a 'how-to manual' for potential attackers. Referring to Figure 4.10, the *Customer Interfaces* are a potential point of weakness. If sensitivity of these is not properly protected, a malicious insider might be able to impersonate a customer to misdirect funds. In the event the malicious insider has sufficient access to credentials, it may be hard for the bank to separate fraudulent from genuine activity. This suggests that the maturity model domain 'Personnel Security' would be of high priority for all three *Department* entities within the bank.

Within the systems layer, preservation of the *Accounts Database* is of the utmost importance. A potential avenue of attack would be for an employee in *Bookkeeping* to manipulate the back-end interface to divert funds from accounts in modest quantities ('skimming'). In the service layer, criticality of *Payment Infrastructure* must be preserved. As seen in Figure 4.10, *Payment Infrastructure* has the relationships *writes to* with *Ledger*, *updates* with *Back-end Interface*, and *send and receive* with *Customer Interfaces*. It is clear from these relationships that a failure of *Payment Infrastructure* would cause significant issues across the organization. Therefore, it would need to be given high priority in the maturity assessment of the bank.

Economic losses

The largest risk the consumer bank faces on a regular basis is cyber-enabled fraudulent activity. This may be direct or indirect — using the bank's customers as a conduit. The elements of the bank with significant exposure to criticality are likely to be closed systems; interbank systems are usually not directly exposed to the internet and are will be rigidly firewalled. Given the importance of availability of such systems, banks will usually invest in infrastructure such as uninterrupted power supplies, multiple data centres and disaster recovery sites. The major economic losses a consumer bank might seek coverage against would be disaster recovery beyond that already budgeted for in its operating model; fraud in excess of a certain threshold; recovery from a significant incident causing a material loss of revenue.

Computing Premia

This analysis follows Section 4.8.2, but rather than providing a full utility function, we discuss a few relevant considerations an insurer may wish to consider. Recall that the form of the insurer utility function is

$$U(\pi_i, P_i, L_i) = \sum_i \pi_i u_i(P_i - L_i) + (1 - \pi_i) u_i(P_i) \tag{4.8.2}$$

We now briefly sketch some considerations for populating π_i and L_i.

Layer	π_i	ΔC_i	ΔS_i
Ops/Mgmt	Incidence of management failures	Failure of Payment Infrastructure	Customer redress
Services	Incidence of insured losses due to fraud	Extortion or ransomware	Strength of controls in customer interfaces
Systems	Incidence of system failures	Ledger controls	Customer data controls

Table 4.10: Utility function parameters for the consumer bank

The methodology used for parametrizing the online retailer maturity is arguably too general for the case of the consumer bank. It should be remembered that in many jurisdictions, banks are subject to rigorous stress tests and regulation. In the United Kingdom, for example, there already exists a protocol for testing the maturity of cyber-security for banks, CBEST [22]. Stress test parameters could be used to populate the utility function; these have the advantage of being measured both in the spaces of probability and losses. One potential insurance strategy might be for the insurer to offer coverage against losses of a probability below a specified threshold (for example, 1%). An excess-of-loss policy, covering losses $L_{min} \leq L \leq L_{max}$ might be the best starting point for constructing the utility function. Based on the maturity assessment, the other required parameters such as maximum available coverage could then be populated and the premium rate calculated.

4.8.4 Pharmaceuticals

Operations

We represent the operations departments of the prototypical pharmaceutical manufacturer via the entities *logistics*, *factory* and *research & development*, for which respectively availability, integrity and confidentiality are required to be protected. The foremost priority for the pharmaceutical manufacturer is the integrity of the factory and manufacturing processes as compromise of this potentially places lives at risk. The research and development department will hold sensitive information, both commercial and personal, about potential drug candidates and clinical trials, which must be protected. The availability of the logistics entity is the lowest priority for the pharmaceutical manufacturer compared with the risks of compromise of integrity or confidentiality. The ERD for the pharmaceutical manufacturer is presented in Figure 4.11.

Maturity assessment

A relevant security model for the pharmaceutical manufacturer is due to Biba [31]. The Biba model is essentially a read up, write down model. By means of illustration, consider a drug, which has a specific chemical structure. The goal of pharmaceutical manufacturing is to reliably produce that chemical structure from a range of chemicals according to a specified synthetic procedure. It is evident that the elements of the process responsible for manufacture need to be able to read the synthetic procedure, but there is no reason why such elements should be able to modify the procedure.

Relative to the other two organizations we have analysed, the pharmaceutical manufacturer operates a relatively closed system. Accordingly, an established 'off-the-shelf' model of maturity such as the CMMC would arguably be sufficient for assessing the pharmaceutical manufacturer from an insurance perspective. Particular attention should be directed at the controls in place to protect the sensitivity of the *sensitive data* database. There is also an integrity objective in terms of the interaction between the *manufacturing processes* and the database. Many of the risks faced by the pharmaceutical manufacturer will be covered by traditional insurance

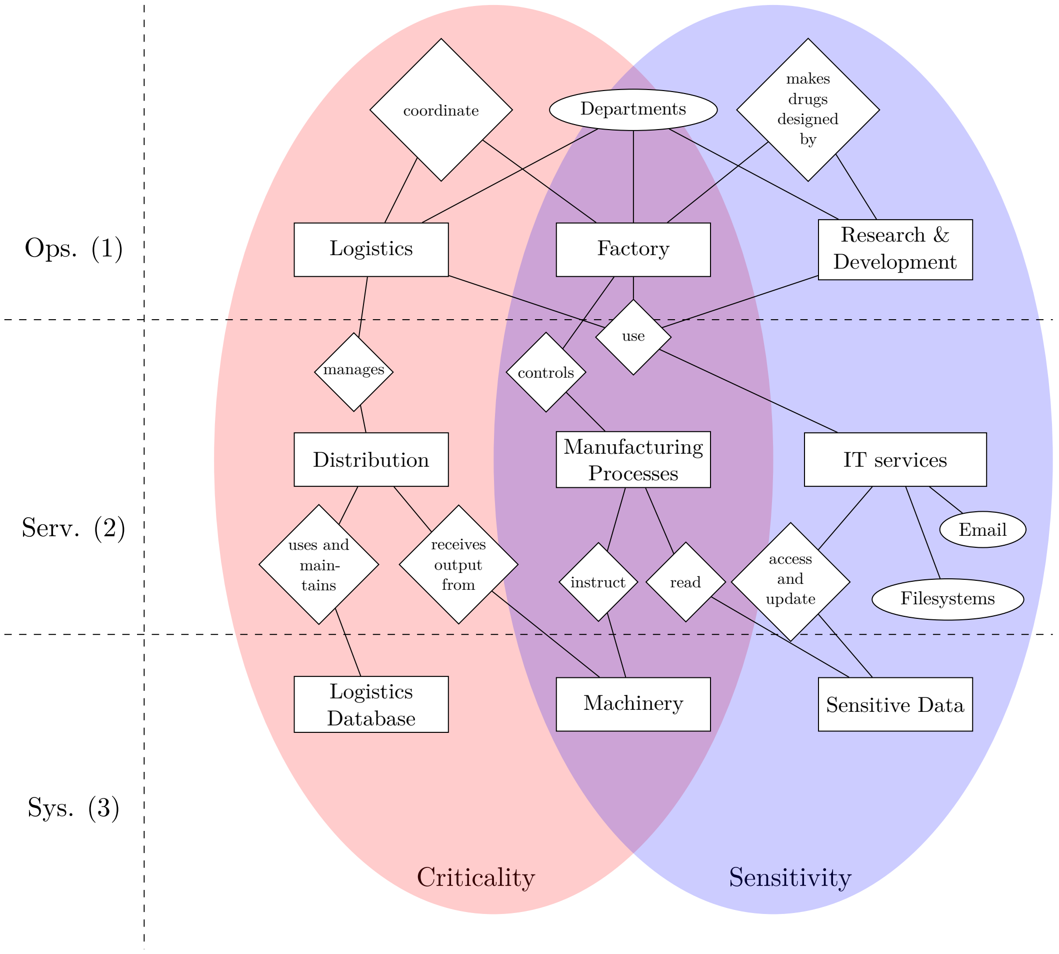

Figure 4.11: Entity relationship diagram for pharmaceutical manufacturing

policies and therefore the scope of the maturity assessment for the cyber-insurance policy will be very specific. An example is the criticality of the systems and service infrastructure supporting the *logistics* function. If there is disruption in this area as a result of a cyber-incident, it may cost the company revenue and inflict reputational damage.

Economic losses

The pharmaceutical manufacturer has the most closed information systems of the three prototypical organizations considered and also has the greatest potential to put in place rigid controls without meaningfully impeding business operations. From an insurance perspective, it is unlikely that the pharmaceutical manufacturer will experience regular losses that it cannot contain itself and the company itself would most likely want to insurer against rare, catastrophic events with potential to seriously damage its business.

Computing Premia

The computation of premia would follow a similar analysis to Section 4.8.2 for the online retailer. We refrain from giving a parametrized utility function for the pharmaceutical manufacturer as the procedure and process have been thoroughly demonstrated in the other case studies. The utility function for the pharmaceutical company would involve instantiating the probability of losses for each layer, π_i, and the expected magnitude of losses, L_i. Pricing of the pharmaceutical policy would likely be most efficient using a deductible — an amount of losses the company itself bears before the insurance company begins to reimburse losses. This would result in a lower upfront premium, which would be rational for a pharmaceutical company which is the most able to protect itself from threats of the three prototypical organizations we have considered.

4.8.5 Reflecting on the Case Studies

- Our chosen case studies illustrate the applicability of our proposed methodology across a variety of industries, with varying risk profiles and security postures.
- They also illustrate how the varying risk profiles and security postures give rise to premia that differ from the baseline of sector profile.
- Finally, the parametrization of maturity illustrates how domains and practices within a sector can give rise to sector-specific accumulation risk. For example, it may be the case that a specific system is widely used in an industry for access control. Should an authentication bypass vulnerability arise, then multiple insured firms in that industry could potentially be affected. This is commonly termed 'supply chain risk'.

4.9 Discussion

4.9.1 Real world considerations for using the modelling framework

In Section 4.8, we explained how to deploy the model for prototypical organizations. Here, we outline some considerations for a real-world insurance assessment. Implementing the model for an insurance assessment of an organization requires the following steps:

- Construct the Entity Relationship Diagram
- Choose the maturity model and parameters
- Identify and assemble relevant data to perform the maturity assessment and compute required inputs (probabilities and losses) for the utility function
- Estimate the utility function.

Briefly, most insurance companies have some form of baseline premium relative to a company's size and risk [247, 298], which is then adjusted for various stated factors such as policy limit, controls in place, prior claims history, and so on. The company may use a questionnaire to gather information that can be used to determine which factors to apply. The actuarial department of an insurance company is responsible for determining the required probability distributions to be used in policy pricing. These will usually be derived from prior claims experienced, but also may reference appropriate external parameters and/or forecasts. Ultimately, the modelling framework proposed in this chapter is intended to help assist this process, not replace it, and to help address some of the difficulties in mapping data of past cyber-insurance losses to potential future losses. These include issues such as technological evolution and the dynamic nature of cyber-threats; for example, the emergence of ransomware over the past few years, which prior loss data could not have predicted.

The first step in the insurance assessment is to construct the Entity Relationship Diagram per Section 4.7.1, containing only the information that is necessary to guide the maturity assessment and insurance pricing. Categorizing risks in terms of criticality and sensitivity allows for the security posture of the organization, its layers, services, and systems to be quickly interpreted. The choice of maturity model should be such that it allows for meaningful comparison between similar organizations for the purposes of policy pricing. It can also be deployed to be aligned with the factors in existing insurance assessments, but with greater depth, structure, and specificity than some questionnaires available in the public domain. Structuring the assessment in this way also allows for any required data to be clearly identified. Finally, the utility function should be framed so that the probability of expected losses and their expected magnitude can be computed based on the maturity model outputs.

It is worth reiterating that the primary determinants of potential losses for an organization are the disruption to its revenue generating activities or costs incurred by a cyber-incident rather than a generic estimation of 'cyber-damages'. Our modelling framework is thoroughly grounded in the structure of an organization and therefore ensures that the treatment of cyber-losses is proportional and appropriate relative to the organization under consideration. At present, there appears to be disagreement between *ex ante* and *ex post* estimates of cyber-losses [279, 298, 91]. One factor worth considering in relation to *ex ante* losses is the extent of employee training and engagement with organizational security policy. This is covered under the 'awareness and training' domain of the CMMC maturity model, and belongs to a set of security considerations commonly termed 'human factors'. As malicious emails are a common means of malware delivery by threat actors, the extent to which human factors form a risk of insurance losses arising should be given due consideration when using the modelling framework.

We would stress that there is no one correct way to deploy the modelling framework and that it has been developed with the aim of being flexible and adaptable to different use cases while providing some consistency of overall approach.

Compatibility with existing cyber-security standards

Our model is compatible with existing cyber-security standards, which as discussed in Section 4.7.2 form the basis of many maturity models. However, these standards are not perfect, as discussed in Chapter 28 of Anderson [11]. Anderson contends that standards such as ISO27001 for security management are deficient in engendering improved security outcomes. The reason for this is that firms are audited and supply the auditors with information regarding the controls. This usually relies on the principle of ‘good faith’, but a misunderstanding or lack of knowledge about the efficacy of controls can have catastrophic consequences. As Anderson points out, many firms with large data breaches have been ISO27001 certified, yet the breach still occurs. There are third party firms, such as Bitsight, who offer scanning services, which are popular in the cyber-insurance industry. Yet the security score they provide is only applicable to the systems layer of the model we have proposed, and standards if self-certified by management cover only the operations layer. The model we have suggested provides an integrated approach across the three layers of an organization: operations/management, service, and systems. The model allows for an insurance company to identify potential risks and price the premium appropriately. A company might be motivated to consider the benefits of investment in security to reduce its insurance premium following a maturity assessment.

4.9.2 Weaknesses

Developing the Entity Relationship Diagrams

The Entity Relationship Diagrams for an organization could quickly become complex and careful consideration must be given in a real-world deployment as to how much information to include. A requirement for the model to be realistic is that sufficient data be available to populate the maturity assessment and utility function. We have also negated the risks of asymmetric information and/or moral hazard in our model. These are important considerations for pricing of insurance.

Subjectivity in Choice of Maturity Model

The choice of maturity model determines the pricing outputs of our model. It is possible that a particular maturity model could be chosen to attempt to move pricing in a particular desired direction by an insurer, or a company might use a maturity model that deliberately covers suspected weaknesses in its own controls.

4.9.3 Extensions and refinements

Cloud usage

In the systems assessment, one potential avenue of investigation is the interaction between cloud systems and insurance. The extent of cloud usage is potentially important as some of the potential criticality issues that insurance might be sought to mitigate may be covered by the service level agreement (SLA) with the cloud provider. In the event that the online retailer is fully reliant on cloud-based infrastructure for its systems layer, then a specialized form of insurance cover may also be available. (Table 4.11). A cloud provider is uniquely positioned to assess the maturity of its customer base, subject to the relevant consents and establishes a synergy with the cyber-insurer.

Cloud Provider	Insurer	Reference
Google	Allianz/Munich Re	[206]
Microsoft	At-Bay	[196]
AWS	Cowbell/Swiss Re	[70]

Table 4.11: Cloud provider and cyber-insurer collaborations

Extending the systems modelling

A potential extension of the work would be to integrate the model based on ERDs, maturity models, and utility with a strongly compositional approach to systems modelling (again, e.g., among many, [33, 198, 66, 56]). Such an extension would support the scaling of our approach to larger, more complex systems.

ERDs as an insurance template

There is potential for an insurance company to use ERDs as a template for an insurance assessment along the lines of what has been described in this chapter. The nuances of ERD notation could be used to depict what perils an insurance company is willing to cover and which they are not. The approach to constructing such a template is not trivial and is likely to be a considerable research and practical endeavour in its own right.

5 Modelling the Cyber-Insurance Market with Risk Transfer via Reinsurance

Cyber-insurance has attracted considerable attention in the literature as a research topic and is now a significant insurance market in its own right, with $7.2bn of direct written premium in 2022 in the US domestic market alone[13] while reinsurance brokers estimate the global market may total $14bn[134]. Commercial estimates suggest that up to 45% of premium is ceded to reinsurers in the cyber-insurance market [3][118]. Yet the interaction between insurers and reinsurers in the cyber-insurance market has received surprisingly little attention in theacademic cyber-insurance literature in comparison to industry publications [119, 134, 46, 221]. This chapter aims to help partially address this gap by considering the asymmetry of information exchange and the uncertain time profile of damage revelation in relation to the cyber-insurance market and its interaction with reinsurers. It is then questioned whether reinsurers are sufficiently incentivised to participate in the cyber-insurance market on a long-term basis given the significant difficulties in achieving *ex-post* efficient information exchange. Cyber risks are a relatively new multi-faceted phenomena and the type of attacks and their impact may change in an unanticipated manner. It is therefore important to understand the resultant issues that may arise and the ability of the market to absorb unexpected losses as otherwise the sustainability of the market is threatened.

5.1 Background

5.1.1 Insurance market structure

We now briefly review the structure of the insurance market and the interaction of its various associated entities and parties. A thorough analysis of the cyber-insurance market requires the role and function of the different participants in the market to be defined.

We assume here that the insurance buyer is a firm who buys insurance coverage via an insurance broker. The broker obtains quotes from different insurance firms provided by their underwriters. An underwriter is responsible for managing a book of insurance policies to deliver specified performance targets. These may vary according to the experience and skill of the underwriter (underwriters with a proven track record may be permitted to write either more premium or cover riskier entities than less experienced colleagues), the markets they cover and the risk tolerance of the provider of the insurance capital. Contrary to what might be expected, underwriting is not purely a statistical exercise. The dynamics of the exchanges between underwriters and brokers are complex, in particular with respect to information exchange which may be highly asymmetric. The job of the underwriter is to make a subjective judgement on the likelihood of the risks (prospective policyholders) they are presented with experiencing a loss and whether these can be underwritten at a premium rate which the underwriter believes is likely to be profitable. This judgement requires a certain amount of skill as while a high insurance rate is more profitable, it will attract less demand than a more attractive rate. The key objective is to price the policy such that the desired mix of risk characteristics is obtained by the insurance firm. Underwriters are assisted by actuaries, who provide mathematical to assist the underwriting .

While the underwriter is the key decision maker at each insurance company in our model, insurance companies usually have multiple underwriters with different areas of expertise in terms of peril and geography - by writing policies covering different perils, insurance companies can reduce their average expected loss by diversification. Insurance brokers act as the intermediary between the insurance company and its underwriters and the end-user of the insurance. For corporate insurance, companies will typically ask their broker to prepare an insurance proposal covering a range of potential losses; these are known as lines in the insurance industry. Property, Casualty & Professional (Liability), Aerospace, and Maritime are well-known examples. The role of the broker is to obtain the best possible terms for its clients - both in terms of premium and depth of coverage. This requires the broker to have an excellent knowledge of the different insurance firms in the market and their reputation. Underwriters will aim to build a strong business relationship with leading brokers in the hope that they will receive a strong allocation of available premium.

Reinsurance companies provide insurance to insurance companies. The main reason for their existence, informally, is to smooth the potential loss profile of insurance companies who otherwise might only be able to write more modest quantities of premium or hold greater capital reserves to cover potential rare outsize losses. Reinsurers also act as a potential clearinghouse for information within the market as the reinsurer will have visbility over the portfolio contents of a range of insurers (known as cedents, which rival insurance companies in the market cannot directly observe.

Cyber-insurance presents a particularly interesting case of insurance market dynamics. The nature of the insured is particularly important as a large firm with high turnover is likely to present a more interesting and economically lucrative target for attackers, but may have better defences than a smaller firm. However, barring a systemic vulnerability the risks of significant losses in a well diversified portfolio of numerous low-limit small-medium enterprise policy may be a more profitable undertaking for a firm. An insurance company will usually obtain reinsurance to manage either tail risks associated with its portfolio (excess of loss) or to reduce its overall exposure (quota share).

5.1.2 Technological advancement, information deficiency and cyber-insurance

One particular issue for understanding loss risks stemming from cyber-incidents is the difficulty in framing the potential future scope of losses. Kurz (2023) [167] provides a very readable account of the attendant challenges for economic reasoning related to advances in technology. Estimates of significant loss are usually calculated by the exposure management department of an insurance company and may be either probabilistic or deterministic (based on stated realistic disaster scenarios). Exposure management traditionally is used to ascertain the risks from a natural catastrophe. In this scenario, the attacker is nature and the vectors are either wind (hurricane) or water (flooding). The questions the model for premium calculation must address are the geographical scope of the damage which determines the expected frequency of claims and the ferocity of the natural disaster which determines the expected severity. While nature is inherently unpredictable, nevertheless past experience of weather patterns gives some basis for modelling expected future losses. The relatively brief (at least in the history of insurance) history of cyber-risks and the constant evolution of technology, its integration in an ever increasing number of processes and the sophistication and capability of attackers makes such comparative predictions regarding potential losses extremely difficult. When designing cyber-insurance policies, it is important for the insurer to be highly specific in terms of the coverage and for the reinsurer to have a clear understanding of the risk dynamics it assumes if these policies are ceded. Figure 5.1 outlines a range of possible cyber losses organised by frequency and severity. This relatively simple graphic encapsulates the potential modelling challenges associated with cyber-insurance and reinsurance. The scope for cyber-losses is determined by the evolution of technology; at the time of writing, generative artificial intelligence and quantum computing are examples of two emerging technologies that have significant security implications.

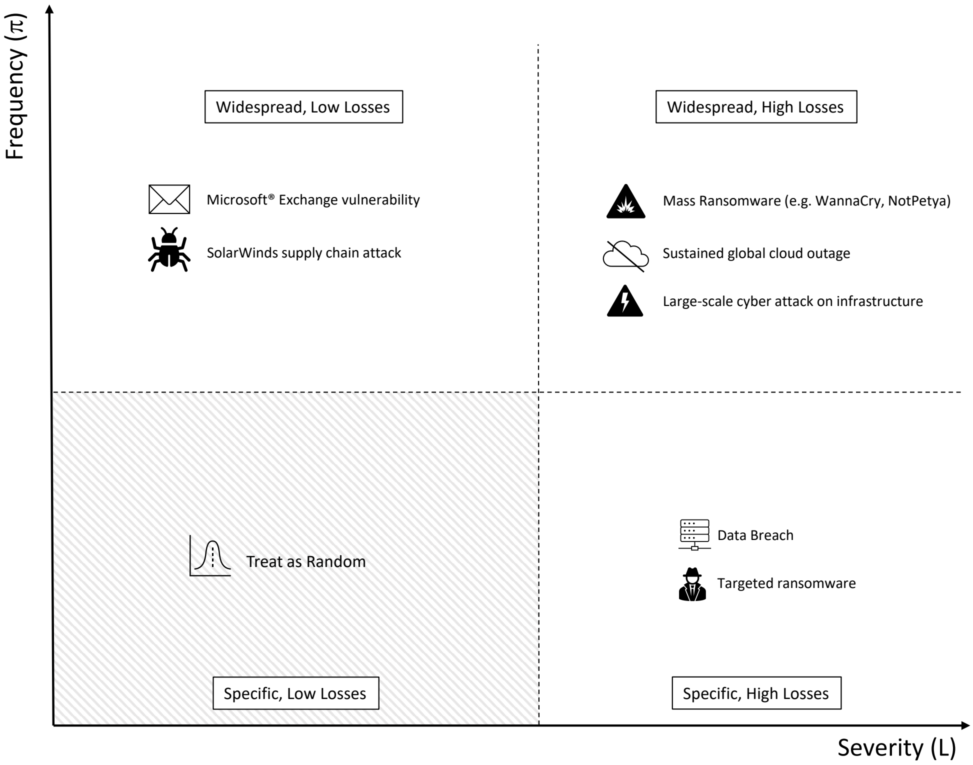
Frequency (π)
Widespread, Low Losses
Widespread, High Losses
Microsoft® Exchange vulnerability
SolarWinds supply chain attack
Mass Ransomware (e.g. WannaCry, NotPetya)
Sustained global cloud outage
Large-scale cyber attack on infrastructure
Treat as Random
Data Breach
Targeted ransomware
Specific, Low Losses
Specific, High Losses
Severity (L)

Figure 5.1: Categorisation of cyber losses by frequency and severity

5.1.3 What claims might arise in relation to cyber-insurance?

Barely a day passes without news of an emerging cyber-attack or other risk. It is important to realise that while these may be extremely disruptive for individuals, companies or societies, not every cyber-attack results in an insurance loss. An insurance loss may be defined as a loss resulting in a claim being paid by an insurer, whereas an economic loss is the total loss to an insured from the peril.

, a ransomware attack (without data exfiltration) would largely result in first party claims for network interruption and recovery costs. In contrast, a large data breach can incur significant third party costs. It is worth noting that such third party claims might arise several years after the policy is written. This is an important feature in cyber-insurance; a prominent relatively recent example was large hotel chain Marriott suffered a cyber-attack commencing in 2014 that was undetected until September 2018[1], which has been one of the largest cyber-insurance claims seen thus

[1] This data breach is widely documented on the WWW from a variety of sources; for an insurance perspective see, for example, [4]

far. A reinsurer might have expected to retain the bulk of cedent premium income, only to find large claims emerging later.

As cyber-risk is such a nascent class of business, the insurance industry is still adapting to understanding how to price the risks, which sectors are most vulnerable and how best to assess underwriting standards. This creates a risky environment to the reinsurer, particularly as cyber is likely to be a relatively small line in their overall business and they may therefore lack the requisite technical expertise to truly evaluate the risks. An interesting example is Solarwinds vulnerability[2], which proved very widespread. However, it appears that the main motivation of the attackers was espionage rather than financial gain and consequently, bar investigative costs, there is little likelihood of significant cyber-claims as a result.

5.1.4 What is the motivation for reinsurance involvement in the cyber-insurance market?

A possible motivation for early entrants to the cyber-reinsurance market is to build market share and hope to capture premium rate increases as the market becomes more popular. Gallagher Re (2022) [119], a leading reinsurance broker with a specialist focus on cyber, have argued that reinsurers, technological solutions and cyber-security practice may converge to create a 'virtuous cycle of capital protection'. As insurers gain more knowledge about the likely distribution of losses, underwriting standards may be tightened. There is nevertheless a clear information advantage possessed by the ceding insurer about the 'quality' of their insurance portfolio relative to the reinsurer, which raises the issue of adverse selection. A rational reinsurer will pay extremely close attention to the information they are given by the cedent with the past loss history of the portfolio often a key feature. The fact that so much premium is ceded suggest also that insurance carriers are themselves nervous of the quantity of risks insured relative to the likelihood of losses. This begs the question as to why reinsurers would rationally increase capital allocations to the cyber-insurance market if the originating insurer is not comfortable with the risks. One possibility is that the reinsurer may have extended scope to absorb losses from cyber-risks more readily

[2] See, for example, Devanny et al (2021)[78]

in a diversified portfolio and may further be able to charge elevated premia if the cedent insurer is desperate to offload the risk.

.. developing a model of a cyber-insurance market, which is stylized yet aims to be representative of the existing cyber-insurance market. . A key argument of this chapter is that there may be a diverse range of beliefs among market participants about the dynamics of cyber-risk and resultant losses. We demonstrate via simulations that under this assumption, reinsurance is only sometimes optimal for insurers. Economic theory on efficiency is consistent with this conclusion. This implies that insurers may need to rely on external sources of risk-tolerant capital (such as insurance-linked securities (ILS, which are sometimes called catastrophe bonds). Further, there is societal benefit in better information sharing on cyber-losses, which should see convergence in beliefs.

5.1.5 Relation to existing literature

This chapter applies well-established economic theory to the cyber-insurance market in a novel manner. Its contribution is that it is the first attempt, to the best of our knowledge, to consider the specific interaction of reinsurance capital and cyber-insurance via simulations that are representative of the existing market.

5.2 Economic Theory on Efficiency

We now consider how to relate well-established economic arguments on efficiency to insurance of cyber-risks. A rational buyer of insurance will likely aim to purchase a policy via a market in order to achieve a price they deem acceptable (ideally optimal). The aim of a well functioning market is to match buyers and sellers of a particular good and to establish a fair price for that good. Efficiency is often used as a measure for the efficacy of risk transfer in a market and can be defined in two ways: *ex-ante* (before the transaction) or *ex-post* (after the transaction). Ex-ante efficiency requires conditions, that we shall demonstrate are extremely hard to satisfy. Ex-post efficiency can be realised but requires an exchange of information. It should be noted that an efficient ex-post premium if realised would be considered a (i.e.

an actuarially fair premium); this is distinct and different from a premium that meets customer expectations and is subjectively viewed as acceptable based on risk tolerance or beliefs. A lack of efficiency does not mean that transactions will not take place, but creates a comparative advantage for the party with greater access to, or possession of less noisy, information.

5.2.1 Ex-ante efficiency

According to the Arrow-Debreu model[14][77], a complete market has:

1. Negligible transaction costs and therefore perfect information
2. Every asset has a price in every possible state of the world

Both of these assumptions are highly unlikely to be valid for cyber-insurance markets. For an asset such as a stock or bond, which may be continually traded, price is a legitimate marker of information. However, commercial insurance contracts are struck at discrete time periods and are valid for a specific length of time only. These typically operate on a yearly basis with key renewal points throughout the year dictated by market convention. Further, there is a significant cost of operating for the insurance company that is typically passed onto the customer via the premium. Most insurance contracts are non-fungible and non-transferable, unlike many publicly traded financial instruments such as stocks or bonds. . This is a fundamental feature of insurance markets and implies that the first condition of completeness within the Arrow-Debreu model is unlikely to be satisfied. The second assumption that every asset has a price in every possible state of the world is equally not realistic as in reality insurance companies may decline to quote for a particular policy if the insuring party considers the risks outside of their tolerance.

5.2.2 Ex-post efficiency

Starr suggests that a set of valuation decisions is ex-post efficient if that there be no redistribution that will increase some traders realized utility while decreasing no traders realized utility [275]. Alternatively, as interpreted by Feiger, there exists no alternative feasible set which is sure to be Pareto improving, looking back from the state which actually occurs. [105]. The Arrow interpretation of states of the world is convenient for an insurance analysis as certain states of the world are loss-triggering. There are a diverse range of possibilities for attempting to frame these states one possible utility driven approach is to model the utility of the protector of an information set using confidentiality, integrity and availability and constructing potential attacks degrading these properties in terms of deviations from their preferred state. A cyber-insurance policy can cover a wide range of potential losses, an interesting case being costs of specialist IT consultants to help diagnosis and recovery after a data breach, for example. A data breach confidentiality but if the system from which the data is taken is somehow modified by a malicious actor to facilitate the theft, then it is also an attack on integrity. Recently, ransomware attacks have become a prominent cyber-threat adding a further risk of loss of availability.

A particular issue for cyber-insurance is the risk of a catastrophic cyber event. A problem with establishing distributions for catastrophic events is that the sample space is often sparse as these events tend not to occur too often. Despite computer systems and networks being societally ubiquitous in most developed countries, public data about cyber-incidents and computer mishaps of the standard required to properly price cyber-insurance contracts remains lacking. Returning to the definitions of Starr and Feiger, these require careful interpretation in the context of cyber-insurance. Consider the scenario in which an entity suffers a loss as a result of a cyber-attack, which is deemed 'with high confidence' by relevant National Cybersecurity Agencies to have been state sponsored. In an efficient market, it ought to be the case that a loss is experienced and thus constitutes a valid claim. However, . . In the event of a significant cyber-attack, the world reaches a state whereby losses are generated. These claims ought to be paid, yet there is a clear potential . Wolff (2023) [292] has produced an extensive survey that relates existing literature on the role of insurance in forming *de facto* regulation to the development of war exclusions in cyber-insurance, concluding that industry leading this development may have far-reaching consequences.

5.2.3 Rational belief equilibria

Kurz compares rational expectations equilibria, in which all agents know the true probability distribution of prices, with rational belief equilibria, in which no one knows the true distribution of prices and each agent must form their own belief about it [166]. Even at first sight, it appears intuitive that the latter category of equilibrium is likely to better characterise cyber-insurance decisions given that a claim to know the path of future technological development with even a degree of confidence is almost certainly fallacious. Kurz's theory of rational belief equilibria relies on the system being stationary for the purposes of agents generating forecasts. The theory identifies a set $B(Q)$ of beliefs compatible with the data generated under Q, which cannot be rejected by the data. At first sight, this may appear a significant issue for analysis of cyber-risk. However, one possibility is that there exists a brittle equilibrium for a finite period of time, subject to shocks. Eventually a shock, or paradigm shift in the sense of Kuhn [163], may perturb the market from its state of equilibrium. This causes market participants to abandon their beliefs but then upon stabilisation a new set of beliefs may be formed. For example, the ransomware epidemic post-WannaCry makes for an interesting case study. This introduced a hitherto less well considered generator of potential losses, which insurers had to adjust for in their policies and subsequently triggered a marked increase in premiums charged to the market .

5.3 Model

We now introduce a model for describing the dynamics of a reinsurance market. We use standard results in the microeconomic theory of insurance without derivation for brevity. The motivation for this is to outline in formal economic terms the structure of an insurance market with reinsurance, from which theoretical simulations may be developed.

5.3.1 Insurance buyer

Before formulating the model for a market, we establish the baseline decision of a buyer of insurance facing two states — loss and no loss — with probability π and $1-\pi$, respectively. The corresponding utility function is

$$E[U] = (1-\pi)u(W - P(C)) + \pi u(W - P(C) - L + C - D) \tag{5.3.1}$$

u is the constant absolute risk aversion (CARA) utility function,

$$u(w) = \frac{1 - e^{-\alpha w}}{\alpha} \tag{5.3.2}$$

where α is a constant. For the purposes of this research, we chose CARA as it is a commonly used utility function and sufficiently captures the trade-offs we wish to model. Other forms of the utility function might be deployed to represent more complex buyer preferences. The parameters in Equation 5.3.1 are W, representing the initial wealth of the insurance buyer; $P(C)$, the premium paid for an amount of insurance coverage, C; and D, the deductible[3] set by the insurer. We shall assume that

$$P(C) = pC \tag{5.3.3}$$

where p represents a premium *rate*. We emphasise that the customer chooses the coverage amount C, up to a limit permitted by the insurer and observes the premium rate, p, from different insurance companies. L is the loss experienced in the loss state. In the event that there are multiple loss states, denoted by s, we assume that these belong to a finite and countable set of states, S, such that $s \in S$, with a corresponding loss for that state, L_s. Specifying an initial endowment, W_0, and representing the total cash premium paid as P, Equation 5.3.1 may be restated

$$E[U] = (1 - \sum_s \pi_s)u(W_0 - P) + \sum_s \pi_s u(W_0 - P - L_s + C_s - D_s) \tag{5.3.4}$$

[3] The amount of losses which must be borne by the insurance buyer

Both Equations 5.3.1 and 5.3.4 are equivalent and for the unsophisticated cyber-insurance buyer, Equation 5.3.1 is a sufficient formulation of the problem. However, when considering the supply dynamics of the cyber-insurance and reinsurance markets, it would be expected that the insurance company consider the different states that may be loss generating. We assume that the objective of the insurance buyer is to maximise their utility.

Assumption 1 *The insurance buyer aims to maximize their utility*

5.3.2 Supply of insurance

Having established the theoretical decision framework for the insurance buyer, we now establish a formal model determining the supply of cyber-insurance. Following Hammond (1981)[136], we consider the actions of consumers in the economy:

$$x^i(s) = [x^i_t, x^i_{t+1}(s)] \tag{5.3.5}$$

i represents an individual consumer of a total I consumers in the marketplace. As before, s represents a contingent state of the world, and it is assumed that the set of possible states, S is finite. The vector of total insurance demand, $\mathbf{x}_t = [C^1_t, C^2_t, \ldots, C^i_t]$.

We assume that there are J insurers in the market, each with a supply of insurance

$$\mathbf{y}^j(s, \mathbf{x}) = [\mathbf{y}^j_t(\mathbf{x}_t), \mathbf{y}^j_{t+1}(s, \mathbf{x}_{t+1})] \tag{5.3.6}$$

$\mathbf{y}^j_t$ is an i-length vector of the units of insurance sold by insurer j to customer i at time t and consequently, which, expressed in monetary terms is identical to cover, C. It is assumed that each customer i has an exclusive policy with its chosen insurer j. Each insurer has a premium vector,

$$\boldsymbol{P}^j = [P_1, P_2, \ldots, P_i] \tag{5.3.7}$$

representing the premium it charges to each customer. This vector may be time dependent. For conciseness of presentation, we will henceforth drop time subscripts as the analysis in this chapter is restricted to a single period.

Insurer objectives

We assume the insurer formulates its decisions on insurance supply, $\mathbf{y}^j(s, \mathbf{x})$ via the following parameters (see Chapter 3.5 of[242]):

- K: the reserve capital held by each insurer.
- P: the total premium income for each insurer.
- X: the claim costs (losses) experienced, described by probability function $F(X)$ with differentiable density $f(X)$ defined over the interval $[0, X_{\max}]$.
- D: the total deductible enforced by the insurer.
- W_0: the initial wealth of the insurer - this may be thought of as shareholder equity, for example, or syndicate (non-regulatory) capital.
- W: the residual wealth the insurer has after paying claims. If the amount of claims is greater than $A \equiv P + K + W_0$, the insurer faces ruin.
- r: the risk-free interest rate for the relevant period.

We assume that each insurer has zero utility condition and its objective is to maximise its wealth

$$W_j = W_0 + P_j - \int_0^{A_j} \frac{X_j - D_j}{1 + r} dF_j(X) \tag{5.3.8}$$

subject to the constraint

$$W_j + K_j > 0 \tag{5.3.9}$$

The intuition underpinning Equation 5.3.8 is that the insurer has a trade-off between the amount of premium it collects and the risk of claim associated with that premium. It may also set a deductible to mitigate moral hazard. Accordingly, the insurer should assess the probable maximum loss of claims according to its distribution and ensure that it has sufficient capital to pay the claims. The claims are discounted by the risk-free rate, r, as it is assumed that the insurer will earn interest on its earned premium over the period. The optimal set of allocations for the insurer would be to policies that maximise the wealth/capital ratio W_j/K.

Assumption 2 *Insurers aim to maximise their wealth*

Assumption 3 *The probability distribution, $F_j(X)$ is subjective to each insurer in the sense of de Finetti (1974) [76]. This will be justified in Section 5.3.4.*

5.3.3 Introducing reinsurance

In order to reduce risk exposure, the insurer may also seek to purchase reinsurance. There are two categories of reinsurance considered in this work: quota share and excess-of-loss. Reinsurers are consequently concerned with determining the probability of two types of extreme events: those resulting in single large losses from a particular client (concentrated losses and those resulting in widespread repeated claims across cedents (contagion/). In the event that this distribution is objective, then this would lead to a universal fair price for insurance. Reinsurers in turn will have their own subjective distributions and charge the expected value of their own distributions to clients. While this may be commercially reasonable, such prices are not fair in a strict economic sense. The existence of reinsurance serves to allow insurers to smooth their subjective expected loss distributions, which clearly implies risk aversion as opposed to neutrality. In short, intermediation implies imperfection[4]. Including reinsurance, Equation 5.3.8 becomes:

$$W_j = W_0 + (P_j - R_j) - \int_0^{A_j} \frac{X_j - D_j - I_j}{1+r} dF_j(X) \tag{5.3.10}$$

where the parameters are as above, with the addition of R, which represents the cost of reinsurance to the insurer and I_j, which is the amount of losses indemnified by the reinsurance policy purchased. The constraint $W_j + K > 0$ continues to apply. Notation-wise, in similar fashion to Section 5.3.2, we use vectors to describe

[4] Skiadas (2013) [270] presents an interesting analysis on this topic

reinsurance supply. We assume that there are k reinsurers, who charge $\mathbf{r}^k$ rates to insurer j and denote the supply vector of reinsurance as $\mathbf{z}^k$.

For a simple quota share policy,

$$R = \rho P$$

where ρ is the proportion of the portfolio ceded and then

$$I(L) = \rho L$$

However, in cases involving excess of loss or other reinsurance treaties, the calculation is more involved. Miccolis (1977) [195] provides an exposition of some standard mathematical techniques for describing excess of loss calculations. In the case of excess of loss, the indemnification equation becomes:

$$I(L) = (L - N)^+ - (L - N - M)^+ \tag{5.3.11}$$

M and N are parameters for an excess of loss policy covering \$M(mn) of losses in excess of \$N(mn). For simplicity, it is assumed that each insurer can purchase only a single excess-of-loss policy from each reinsurer. It would seem rational for the purposes of our discussion that the insurers seek to buy reinsurance above the aforementioned value A, losses above which the firm becomes insolvent.

The reinsurance market

We assume that there are K reinsurers in the market who provide reinsurance capacity. The reinsurer aims to maximise wealth in similar fashion to the insurer (Equation 5.3.8), but does not include a deductible:

$$W_k = W_0 + R_k - \int_0^{A_k} \frac{I_k(X)}{1+r} dF_k(X) \tag{5.3.12}$$

R_k is the total reinsurance premia received and $I_k(X)$ denotes expected reimbursements paid out to cedents. The reinsurer is subject to the capital constraint

$W_k + K_k > 0$. Note that we allow for the reinsurer and insurer to have different beliefs about the expected distribution of losses. As with the insurers, A_k, represents the amount of reinsurance payouts above which the reinsurer would be insolvent.

Assumption 4 *The reinsurer may have a different belief from the insurer regarding the distribution of risks.*

5.3.4 Modelling cyber-risks

We have thus far considered losses related to cyber-risk in an abstract sense as setting up the theoretical framework for evaluating the interaction between buyers, insurers, and reinsurers does not require the functions dictating these losses to be instantiated. However, simulating the decision making to analyse the potential for efficiency in the market does require some sample distributions. We use standard results in probability theory without derivation (the reader wishing to understand the background more thoroughly is referred to any standard statistical text on probability theory; Williams (1991) [289] is a particularly accessible and carefully explained introduction). While using formal probability theory is not essential for simulating the results in this chapter, it is beneficial to apply theoretical rigour as this helps to highlight some features specific to cyber-risk that are potentially problematic for formulating traditional actuarial insurance assessments.

We start by defining a probability triple $(\Omega, \mathcal{F}, \mathbf{P})$. Ω is a set representing the sample space of *all events.* ω represents a sample point of the sample space. The σ-algebra[5], $\mathcal{F}$, on Ω, is known as the family of events[6]. Denoting an event by A, we may write

$$\mathcal{F} = \{A | A \subseteq \Omega, A \in \mathcal{F}\} \tag{5.3.13}$$

The intuitive explanation in relation to cyber-insurance is that $\mathcal{F}$ is the collection of

[5] The definition of a σ-algebra is a collection of subsets of a set that is closed (stable) under any countable number of set operations. This is important for working with probabilities, where the probabilities of all possible outcomes must sum to 1.

[6] See Chapter 2 of [289]

events covered by a policy that may trigger a claim and then, possibly, a loss to the insurer. If $\mathcal{F}$ is the Borel[7] σ algebra on the set of real numbers, then there exists a unique probability measure on $\mathcal{F}$ for any cumulative distribution function. Letting X be a random variable on $(\Omega, \mathcal{F}, \mathbf{P})$,

$$\begin{tikzcd}
 & \Omega \arrow[r, "X"] & \mathbb{R} \\
{[0,1]} & \mathcal{F} \arrow[l, "\mathbf{P}"'] & \mathcal{B} \arrow[l, "X^{-1}"']
\end{tikzcd} \tag{5.3.14}$$

Informally, this means that so long as there is a collection of events that obeys certain mathematical properties, it is possible to assign a probability to an event using a probability distribution function. One interesting outcome is that a key assumption of probability theory is that the system is stable. This is a potentially problematic assumption for cyber-risk as there have been clear examples of previously unconsidered threats developing. However, insurance policies comprise a set of event definitions as part of the policy, which are contractually binding (albeit open to legal dispute). The importance of careful policy wording is consequently readily apparent. As will shortly be explained, underwriting cyber-insurance policies requires an assumption of subjective, temporary stationarity in distributions. This is a realistic assumption in the context of industry practice, where (re)insurance policies last for a year and then are re-priced based on updated distributions resulting from supply-demand dynamics and claims experienced.

Why use subjective probabilities to model cyber-risks

Assumption 3 in Section 5.3.2 stated that the probability distributions that govern insurance supply are subjective in the sense of de Finetti [76]. We now provide the intuition behind and justification for this assumption before moving to consider the form of distribution that might be used to model cyber-insurance decisions.

In Section 5.2.2, we outlined the conditions required for ex-post efficiency. Considering these in the context of cyber-insurance, we conclude that ex-post efficiency is unlikely to hold and almost certainly cannot be implemented at the time when the underwriting decision is made. Unless of course, the true probability distribution

[7] The Borel σ-algebra, $\mathcal{B}(\mathbb{R})$, is the smallest σ-algebra containing all open intervals in $\mathbb{R}$

attached to the known and finite states of nature is known and shared by all participants. Such condition is the foundation of the theory of rational expectations. This is synonymous with the existence of a stationary distribution. One way of defining a stationary process is to say that its moments are time-independent, which means that the average value of the measurements is a constant. Such distributions are foundational for the existence of efficient equilibria under risk.

It is usual in macroeconomics to depict technological progress as a Markov chain [173, 48]. If the depicted process has started far from its invariant distribution, then it is also non-stationary, but easy to predict as it will approach the limiting distribution that is ultimately stationary. However, in a short epoch, it will appear as non-stationary. Whether technological progress has such a limiting distribution is an unresolved question. Over the long-run it appears to have exhibited a definite trend, with some downwards transitions attributable to natural disasters, wars epidemics etc. In the short run, local approximations can be derived, and expectations can be formed, however agents will splice different segments depending upon their horizons and discount rates. The imposition of rational expectations restrictions upon this structure can only be justified if all agents have identical preferences and endowments, a condition that by construction does not hold. For Markov chains with non-stationary transition probabilities, no steady-state typically exists and almost nothing in the non-stationary setting is computable in closed-form.

It is hard to imagine that there is any way to truly predict an arbitrary non-stationary process. This is because as soon as one postulates a future path another can always reverse it, without creating any problems of consistency with earlier data. In a more general case one might lower expectations, not to actually predicting well, but to predicting with low regret. To this effect agents can choose their most suitable approximations selecting the time span and use their best computational algorithms.

In the absence of a universally accepted probability distribution, *ex-post* efficiency is almost impossible to attain. Of course, there are opportunities which can best utilised only with *ex-ante* knowledge of the state of nature. In its simplest form, it is the choice of technique in production/product that depends upon the expected state. However, a more interesting situation arises when the expected state conditions the preferred level of production. Expecting the cyber-insurance market to quote premia at all levels that are consistent with ex-post efficiency is rather unrealistic. The very nature of the underlying processes does favour the existence of a generally accepted stationary distribution. Rational agents will behave as if they are *ex-ante* efficient using their own expectations of losses based on their subjective probability distributions taken over their own sample spaces.

The evolution of cyber-threats will be conditioned of the path of technological improvements in both elements of information and communications technology, software and hardware. The future path of such advances may be partly predictable based on well-established empirical regularities, such as Moores law that famously predicted that the number of transistors on integrated circuits would double every two years, i.e. at an annual rate of about 40% [202]. Others[8], looking at related data came to the conclusion that predictions of particular technological IT innovations, such as hard drives may be approximated using exponential functions. A very useful exposition of this attempt, using smooth functions the predict technological progress is Farmer and Lafond (2016) [103].

Yet technological advances undergo structural breaks, where both the level of technology in terms of some of its main characteristics and its future direction change. A prominent example at present is the introduction of quantum computing, which will alter radically reduce computational time and thus has implications for the robustness of cryptographic protocols that are currently infeasible to attack on a realistic timeframe.

Technological progress is achieved by the complex interactions of two main human pursuits. The organised knowledge as it appears in scientific papers, submitted patents, recipes, protocols, routines and probably informal know-how, acquired through ‘learning by doing’ in a long process of imitation and repetition. The development of science, technology, innovation and production require both codification and knowledge.

It seems unlikely that such dual processes can be tamed into a smooth parametric function with time invariant parameters, shared by all participants. If anything, in

[8] For example, [27, 114, 208]

the absence of such shared beliefs, it is expected that for market participants whose welfare depends upon such developments, their decision making will be based on arbitrarily diverse anticipations. These are individually efficient decision makers because they act on the basis of all the information available at the time. It is clear therefore that by and large insurance contracts on expected losses based of future technological developments, that are subject to structural changes, cannot be written on generally accepted parameters, to deliver Arrow-Debreu type ex-ante efficient premia. All the participants are efficient in terms of fully exploiting their private anticipations of losses, but the quoted premia at the two levels will not result in fully efficient in the Pareto sense economic outcomes.

Probability Distributions

For the simulations in Section 5.4, we separate the expected distribution of losses into the number of expected claims (frequency) and the average expected loss per claim (severity). This is a very common method for actuarial modeling and is described in most standard texts, for example Panjer (2006) [227]. Its appropriateness to categorising cyber-risks was described in Section 5.1.3 and summarised in Figure 5.1. We assume that frequencies follow a Poisson distribution and severities a log-normal distribution. The Poisson distribution is a standard starting point for frequency modelling (see, among many, [197]). There is no clear consensus in the empirical literature on which distribution is most appropriate for describing the severity of cyber-losses (see in particular [91, 298]). We use the log-normal distribution as a starting point as it is well-understood and straightforward to configure. We use simulated rather than empirical distributions as the aim of the simulations is to examine whether efficiency is theoretically possible, whereas markets in practice are very unlikely to be efficient. The probability of k events occurring in a unit of time represented by the Poisson distribution is

$$f(k, \lambda) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{5.3.15}$$

where λ is the expected number of events. The log-normal distribution assumes

$$\ln(X) \sim \mathcal{N}(\mu, \sigma) \tag{5.3.16}$$

that is the natural logarithm of variable X is normally distributed with mean, μ, and standard deviation, σ, which are defined as

$$\mu = \frac{\mu_X^2}{\sqrt{\mu_X^2 + \sigma_X^2}} \quad \text{and} \quad \sigma^2 = \ln\left(1 + \frac{\sigma_X^2}{\mu_X^2}\right) \tag{5.3.17}$$

μ_X and σ_X^2 are the mean and variance, respectively, of the variable X. The probability density functions and cumulative distributions functions for the log-normal distribution are readily available in any standard resource on statistics and are omitted for brevity.

Combining probability distributions

Cyber-insurance policies cover a diverse range of first- and third-party risks and consequently, there is probably no one distribution that actually covers all relevant risks. Accordingly, it is desirable to consider a combination of possible risks. Unfortunately, probability distribution functions are rather difficult to combine with a closed-form solution (see, for example, Nadarajah et al (2018) [207]) and require analytical solutions. A common strategy is to use a package such as Mathematica [293]. However, there is an alternative approach which is to use Monte Carlo-type simulations. Section 5.4 will show how these can be deployed to yield useful insights on insurance decisions, the results of which do not require sophisticated mathematics to formulate or interpret.

5.4 Simulations

We consider simulations of a cyber-insurance market with reinsurance over a single period. We assume that losses arise in the period of the insurance policy and are recorded at the time they arise. Policy data is confidential to insurance companies and consequently, the simulations are established for model convenience but are constructed to replicate real-world insurance market dynamics. We use Poisson distributions for the frequency of losses and log-normal distributions for the severity of losses (details of these distributions and their associated functions may be found in any standard statistical text). The Poisson distribution is a common choice for modelling claim frequencies in insurance (see, among many excellent references, [228][197]). There is no clear consensus in the literature on the optimal distribution for modelling the severity of cyber-related claims, but the log-normal distribution has been shown to be a reasonable approximation in the limited empirical studies to date (e.g. Eling et al (2019) [91], Woods et al (2021)[298]. The use of the joint frequency-severity distribution approach follows Panjer (2006) [227]. We assume a common set of contracts across insurers varying in limit size.

The analysis considers only variation in coverage and premium. We assume arbitrarily a market size of $500mn total coverage. The simulations were computed using the Julia programming language. We found the *Distributions.jl*[175], *QuadGK*[152], and *Plots.jl*[60] packages particularly useful in facilitating the presentation analysis. Unless otherwise specified, *Monte Carlo* type simulations were run 100,000 times.

The goal of the simulations is to illustrate how capital supply from the reinsurance market to the insurance market and then to buyers is inherently inefficient as pricing is influenced by the diversity of opinions regarding the frequency and severity of losses even with relatively simple standard distributions. The simulations might be applied to a variety of insurance markets, but they have been constructed to be representative of the existing cyber-insurance market based on the authors' interaction with insurance market professionals.

5.4.1 Preliminaries

Familiarity with the insurance market is not a prerequisite for understanding and interpreting the simulations that follow. We have taken care to explain the terms used and ensure parameters are fully defined and explained. However, the reader unfamiliar with corporate insurance may find the following definitions helpful as a reference. These may be safely skipped for those experienced in either the practice or study of insurance.

- μ_L: The average expected loss in monetary (cash) terms.
- σ_L: The standard deviation of losses in monetary (cash terms).
- $F^{-1}(p)$: The loss value that occurs with probability p according to the cumulative distribution function F. If $p = 0.95$, then in 95% of cases, the loss is expected to be less than or equal to the output of this function.
- Loss ratio: the percentage of cash premiums collected by an insurance company for a specified period (usually a year) paid out as losses.
- Frequency: the number of claims in a period.
- Severity: the average loss per claim.
- Cover/Exposure: the total maximum losses that could result from a policy/-portfolio respectively.
- Expected loss: the mean loss from a policy/portfolio.
- Technical premium: the cash premium or premium rate (calculated as the ratio Expected Loss/Cover) corresponding to the expected loss. This is the premium income at which the insurer can be expected to break even.
- Simulated loss: the average loss from running N simulations based on random sampling of the expected loss distribution. This can only be computed once the portfolio is formed, so we assume that premiums are calculated based on expected loss values.
- Ceding commission: the percentage premium paid back to a cedent by a reinsurer to cover underwriting expenses and other costs.

It is important to note the sequencing of the insurance transactions in the market. The insurance buyers observe a premium rate and based on this decide how much cover to buy. The insurance provider then has obtained a portfolio. Based on the risk characteristics of that portfolio, the insurer may look to enter into a reinsurance contract to eliminate some potential risk. The simulations assume that insurers and reinsurers target a specific loss ratio *ex ante* to determine pricing.

5.4.2 Simulation Strategy

We consider three simulations:

1. A benchmark simulation.
2. One reinsurer, five insurers with different portfolios comprised of different weights of five common contracts, buyers not considered.
3. One insurer, one reinsurer, different buyer price sensitivities.

These simulations are distinct from each other, though have broadly consistent parameters where possible. The aim of the benchmark simulation is to demonstrate the approach used to generate loss distributions and also to instantiate buyer utility functions to show that if the buyer has a different expectation of loss severity from the insurer, then full insurance coverage may not be utility maximising.

The second simulation starts with a reinsurer who has a range of distributions its actuaries consider acceptable. The reinsurer attempts to offer reinsurance to achieve a target loss ratio and so quotes a reinsurance rate to the market. The market consists of five insurers who have portfolios that range from a large number of small loss risks (called Insurer Alpha) to a small number of large loss risks (called Insurer Echo) with Insurers Beta, Charlie and Delta having portfolios that move progressively between the two extremes. This aims to replicate the structure of the cyber-insurance market in a stylised form and contrast the appropriate reinsurance strategy for the different types of insurer.

It should be noted that the premia in the simulations may vary from those witnessed in the market and in some cases appear very large. The simulations are intended to guide the reader through an application of the economic theory and market model from a theoretical perspective and demonstrate the difficulty of establishing efficiency rather than aiming to be a simulation of the real-world cyber-insurance market.

5.4.3 Simulation 1: Benchmark simulation

We first consider a simple simulation before starting to examine the effects of varying market structure and pricing variables. This simulation assumes the following:

- There is only one insurance policy offered in the market, with a limit of \$1mn.
- for each policy is \$500k, with standard deviation \$250k.
- The frequency of losses is simulated under two scenarios where 10% and 50%of policies are expected to experience a loss, respectively.
- There are 100 buyers, five insurers and one reinsurer in the market. For simplicity, we model total losses for the market and assume they are evenly distributed.
- Losses are simulated with 100,000 runs and random sampling of the severity and frequency distributions.
- Distributions are shared by all market participants.

Figure 5.2 plots the probability distribution functions of the severity distribution and the two frequency distributions. The severity distribution is log-normal with parameters $\mu = 13.0$ and $\sigma = 0.22$; the two frequency distributions are Poisson with λ of 10 and 50, respectively. The PDF values for the severity distribution are very small because of the units of the loss; the integral of the PDF across the function domain must sum to 1. Running a simulation, the expected loss distribution for the two frequency distributions can be obtained. This is presented in Figure 5.3. The values on the y-axis of Figure 5.3 simply represent the number of times each loss value range in the histogram appears in the simulation. Each bar in the histogram has a width of \$0.5mn. This is simply chosen for aesthetic reasons. The main emphasis is on the shape of the distributions rather than the precise frequency count in the histogram.

Having examined the distributions, we now consider the pricing of the policies. Table 5.1 shows the expected and simulated losses for the distributions in Figure 5.3.

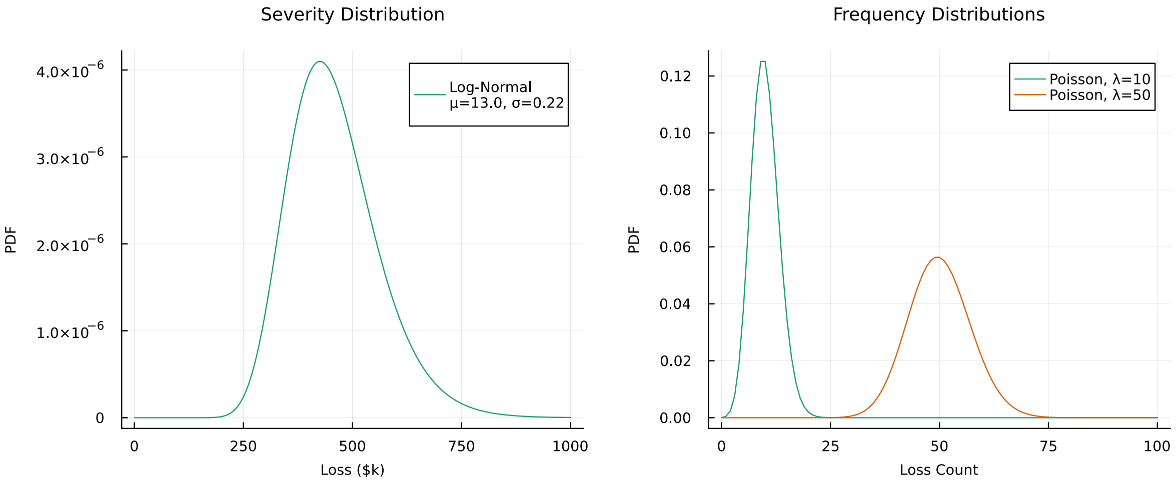

Figure 5.2: Benchmark severity (LHS) and frequency (RHS) distributions

Note that

$$\text{Expected Loss} = \text{Expected Frequency} \times \text{Expected Severity} \times \text{Number of polices}$$

The ratio of the Expected Loss and the Exposure ($100mn in this example) gives what is known in insurance as the technical premium rate. Accordingly, the technical premium would be 5% for the 10% frequency scenario and 25% for the 50% frequency scenario. The simulated losses are lower than the expected (mean) losses because of the skew of the log-normal distribution.

Frequency	Expected Loss	Simulated Loss
10%	$5mn	$4.6mn
50%	$25mn	$22.9mn

Table 5.1: Expected versus simulated benchmark distribution losses

To simulate reinsurance pricing, we first fit a log-normal distribution to the joint distribution with 50% loss frequency as previously described. . We consider reinsurance only for the 50% loss frequency distribution as guided by the reported loss ratios in Table 5.17, which suggest relatively high frequencies of losses have been experienced by the actual market. Using the *fit* functions in *Distributions.jl*, we obtain a log-normal distribution with $\mu = 16.9$ and $\sigma = 0.27$. Under this distribution, the cumulative probability of a loss exceeding $50mn is extremely small, therefore we price the reinsurance policies using excess-of-loss of $50mn. Using the cumulative

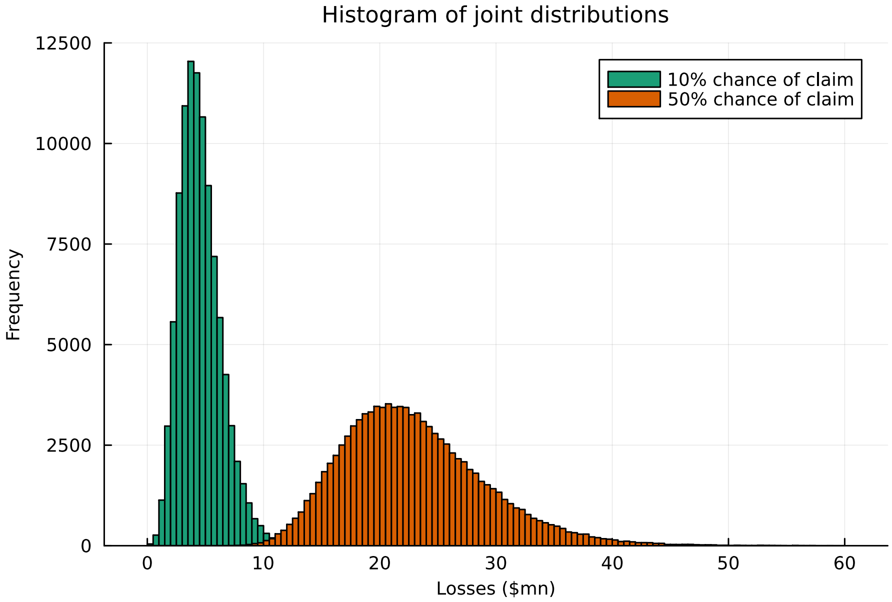

Figure 5.3: Simulated loss distributions

probability functions for the estimated distribution, we can then obtain premium rates for the reinsurance, which, multiplied by the amount of reinsurance required, gives the cost of reinsurance. We then re-run the simulations of losses for the insurer assuming no losses are incurred above the threshold at which reinsurance cover binds. We can then obtain the simulated loss with reinsurance. The results are presented in Table 5.2.

Reinsurance	Reinsurance Premium Rate	Reinsurance Cover	Technical Reinsurance Premium	Simulated Net Loss
\$25mn xs \$25mn	32.2%	\$25mn	\$8.1mn	\$13.1mn
\$20mn xs \$30mn	12.5%	\$20mn	\$3.8mn	\$18.7mn
\$15mn xs \$35mn	4.2%	\$15mn	\$0.6mn	\$21.3mn
\$10mn xs \$40mn	1.3%	\$10mn	\$0.1mn	\$22.4mn

Table 5.2: Reinsurance premia for excess-of-loss policies on the 50% chance of claim distribution.

This reinsurance pricing may be considered efficient because both the reinsurer providing coverage and the insurer seeking reinsurance have the same expected loss distribution.

5.4.4 Simulation 2: Reinsurance supply and price

Having considered the case where all parties agree on the same distribution, we relax this assumption and start to consider divergence in distributions of expected losses. We begin by considering the objective of the reinsurer. We assume a log-normal distribution of total losses. This is the distribution the reinsurance company believes represents the losses experienced from a pool of cedents. The reinsurance company needs to model different potential loss ratios. Initially, we assume cover is fixed at a maximum of \$500mn. Table 5.3 presents a number of log-normal distributions. These are purely for illustrative purposes; in a real world situation, the reinsurer would model the distribution based on experience and data. However, it is helpful to consider a range of distributions to understand how the shape of the distribution may affect pricing.

		μ_L	σ_L	μ	σ	$F^{-1}(0.995)$
Distributions	A	\$10mn	\$10mn	15.8	0.69	\$42mn
	B	\$20mn	\$20mn	16.5	0.69	\$84mn
	C	\$30mn	\$30mn	16.9	0.69	\$126mn
	D	\$40mn	\$40mn	17.2	0.69	\$169mn
	E	\$50mn	\$50mn	17.4	0.69	\$211mn
	F	\$60mn	\$60mn	17.6	0.69	\$253mn

Table 5.3: Table of reinsurance distributions

Within this table, $F^{-1}(0.995)$ represents the maximum loss with 99.5% certainty within the distribution. This is the probability value used under the Solvency II insurance regulation to determine the required capital a firm must hold. The probability density function and cumulative distribution functions for the distributions in Table 5.3 are plotted in Figure 5.4. Note that the scale of the loss axis is shortened to \$100mn as the probability density function returns extremely low values beyond this point.

To estimate the premium rate, we consider the following. The reinsurer targets a

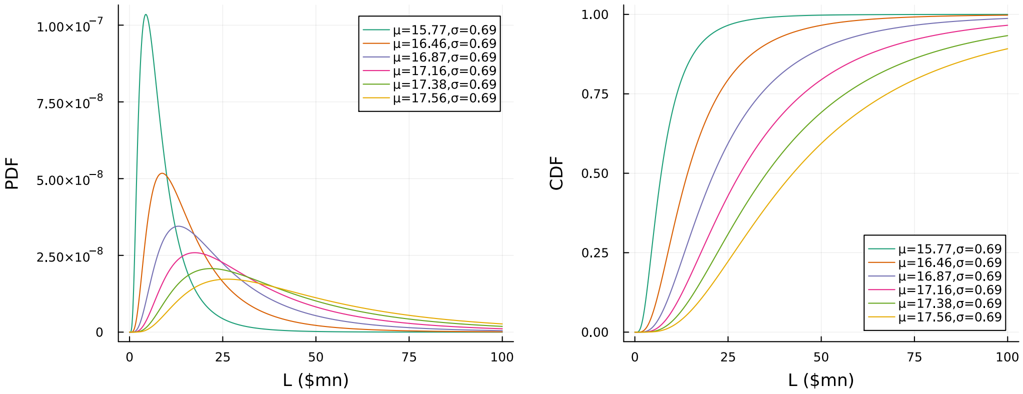

Figure 5.4: Reinsurer loss distributions

loss ratio (a common performance metric in the insurance industry). Total losses from the portfolio are then written

$$L = \text{L.R.} \times \sum_J r_J C_J \tag{5.4.1}$$

Losses experienced are also given by

$$L = \sum_J E[I_J] \tag{5.4.2}$$

We assume there is a single rate for reinsurance such that $r_J = r \forall J$. Then,

$$r = \frac{\sum_J E[I_J]}{\text{L.R.} \times \sum_J C_J} \tag{5.4.3}$$

Denoting $\bar{C} = \sum_j C_J$ and noting that $\sum_J E[I_J] = \int_0^{\bar{C}} If(I)dI$ where $f(I)$ is the probability density function of an appropriate distribution, we obtain

$$r = \frac{\int_0^{\bar{C}} If(I)dI}{\text{L.R.} \times \bar{C}} \tag{5.4.4}$$

This integral can be evaluated numerically, for example using *QuadGK* in Julia.

Suppose the reinsurer believes that Distribution C best describes expected losses to the portfolio and targets a loss ratio of 50%. The rate of reinsurance charged is then 11% (Table 5.4). Premium income for the reinsurer will be $55mn. Note that

		Premium rates								
Loss ratios→		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Distributions	A	0.18	0.09	0.06	0.04	0.04	0.03	0.03	0.02	0.02
	B	0.36	0.18	0.12	0.09	0.07	0.06	0.05	0.04	0.04
	C	0.54	0.27	0.18	0.13	0.11	0.09	0.08	0.07	0.06
	D	0.72	0.36	0.24	0.18	0.14	0.12	0.1	0.09	0.08
	E	0.9	0.45	0.3	0.22	0.18	0.15	0.13	0.11	0.1
	F	1.08	0.54	0.36	0.27	0.22	0.18	0.15	0.13	0.12

Table 5.4: Illustrative premium rates for target loss ratios under different distributions at cover fixed at \$500mn

per Table 5.3, in Distribution C, the 99.5% upper bound for losses is \$126mn.

Insurance supply

We assume for simplicity that there are five insurance contracts in the market with different limits: \$500k, \$1mn, \$2mn, \$5mn, \$10mn. We assume that there is a uniform individual and independently distributed probability of loss for each contract: The expected severity in the above contracts is assumed to be log-normally distributed per Table 5.5 and the frequency $\sim$ Poisson($\pi_l k$) where k is the number of contracts. Table 5.6 contains a sample portfolio for a panel of 5 insurers for illustrative purposes to run a loss simulation. The technical premium is the premium income that equates to the expected loss. Equivalently, this is the premium written at which the insurer would expect to break even.

Limit	μ_L	σ_L	Frequency (π_L)	Expected Loss ($\pi_L.\mu_L$)	Premium (Exp. Loss/Limit)
\$500k	\$200k	\$125k	0.1	\$20k	4%
\$1mn	\$400k	\$350k	0.15	\$60k	6%
\$2mn	\$1mn	\$1mn	0.16	\$160k	8%
\$5mn	\$2.5mn	\$1.25mn	0.2	\$500k	10%
\$10mn	\$4mn	\$4mn	0.3	\$1.2mn	12%

Table 5.5: Insurance contracts in the market

In reality, insurers do not attempt to break even but rather aim to produce a profit to provide a return on investment to the source of their capital. One simple

	Policy count grouped by policy limit						
Insurer	\$500k	\$1mn	\$2mn	\$5mn	\$10mn	Total Exposure	Technical Premium
Alpha	200	0	0	0	0	\$100mn	\$4.0mn
Beta	100	50	0	0	0	\$100mn	\$5.0mn
Charlie	50	20	15	5	0	\$100mn	\$7.1mn
Delta	30	0	5	5	5	\$100mn	\$9.9mn
Echo	0	0	0	0	10	\$100mn	\$12.0mn
Total	380	70	20	10	15	\$500mn	\$38.0mn

Table 5.6: Insurance policies written by insurance panel

objective might to not exceed a target loss ratio. This is achieved via an additional charge to the insurance buyer over the technical premium known as a loading[9]. The loading is calculated:

$$\text{Loading} = \frac{1}{\text{Total Exposure}} \times \left(\frac{\text{Technical Premium}}{\text{Target Loss Ratio}} - \text{Technical Premium} \right) \tag{5.4.5}$$

The variation between loading and loss ratio for the insurance portfolios in Table 5.6 is plotted in Figure 5.5. The variation in target loss ratios may occur for a number of reasons, such as rate of return on capital demanded by the capital source (as discussed in Section 5.1.1, prior loss experience, or other variable expenses. The loading also may aim to capture any skew in the actuarial distribution.

Table 5.7 shows the calculated loadings for each insurer in the simulation assuming a target loss ratio of 50%. For ease of comparison, we keep the target loss ratio constant across the insurer panel and also the overall exposure.

Insurer	Technical Premium	Target Loss Ratio	Exposure	Technical Premium Rate	Loading	Weighted Average Charged Premium Rate
Alpha	\$4.0mn	50%	\$100mn	4.0%	4.0pp	8.0%
Beta	\$5.0mn	50%	\$100mn	5.0%	5.0pp	10.0%
Charlie	\$7.1mn	50%	\$100mn	7.1%	7.1pp	14.2%
Delta	\$9.9mn	50%	\$100mn	9.9%	9.9pp	19.8%
Echo	\$12.0mn	50%	\$100mn	12.0%	12.0pp	24.0%

Table 5.7: Calculating premium loading rates for insurance companies based on simulated losses. The loading rate is expressed in percentage points.

[9] See, for example, Benjamin (1986) [26] for a discussion.

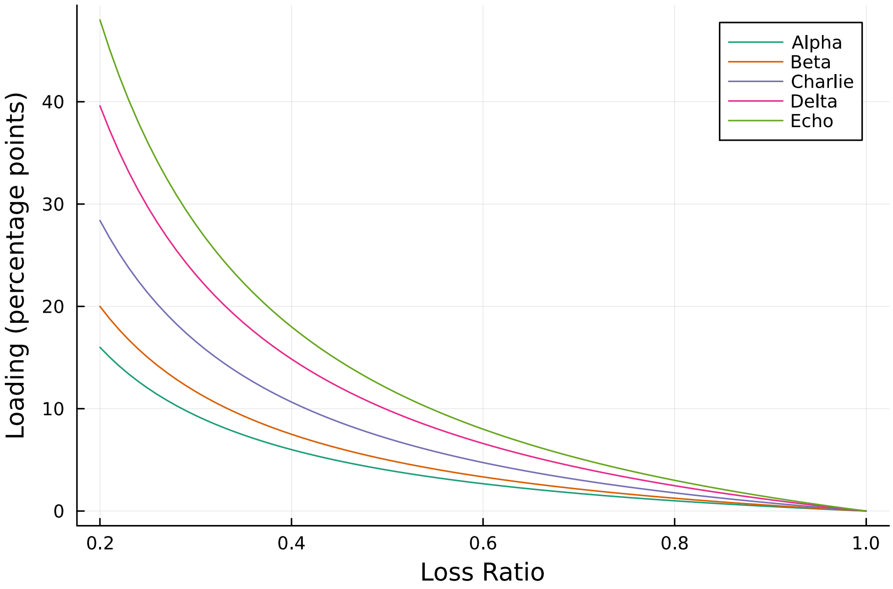

Figure 5.5: Loading versus target loss ratio for different insurance portfolios

Interaction between insurance and reinsurance

The total expected losses for the cyber-insurance market depicted in Table 5.7 are \$38.0. The technical premium is equal to the expected losses in monetary terms. In Section 5.4.4 we stated that for a log-normal distribution with mean and standard deviation of \$40mn and target loss ratio of 0.5, the premium charge would be 14% for the reinsurer (distribution D). The usual process of reinsurance in quota share is that the reinsurer assumes a stated percentage of portfolio losses. The reinsurance contract (or treaty) is priced[10] via a ceding commission and reinsurance margin. In this case, the reinsurance margin is already accounted for in the 14% premium rate as this was calculated to give the required reinsurer loss ratio. The ceding commission is paid back to the ceding insurer to compensate them for underwriting expenses. The ceding commission is defined as the average premium rate (Table 5.7) less the cost of reinsurance (14%). Inspecting Table 5.7 once more, we can see for insurers Charlie, Delta, and Echo, the average premium rate of the portfolio exceeds the reinsurance cost. Therefore, the ceding commission for these insurers would be positive. However, for insurers Alpha and Beta, their weighted average premium rate is below that charge for reinsurance, implying a negative ceding commission. If

[10] Clark (2014) [62] is a highly approachable introducing to reinsurance pricing

Alpha or Beta believe that their assumed distributions are correct, this would not be rational behaviour. For the other insurers, purchasing reinsurance would reduce profits for *expected losses*. However, the value of reinsurance will become apparent once we consider the effect of capital.

Having established the target pricing for each insurer *ex ante*, we now consider simulating *ex post* losses. The profit equation for the insurer, may be written:

$$\begin{aligned}\$\text{Profit}(L) = \,&\$\text{Premium Written} \times (1-\rho)\\ &+ (\$\text{Exposure} \times \rho \times \%\text{Ceding Commission})\\ &- L\end{aligned} \tag{5.4.6}$$

$$L = \begin{cases} \text{Loss}(1-\rho) & \text{if } D = 0\\ \text{Loss} & \text{if } D > 0, \text{Loss} \leq D\\ D + (1-\rho)(\text{Loss} - D) & \text{if } D > 0, \text{Loss} > D \end{cases} \tag{5.4.7}$$

ρ is the fraction of the portfolio ceded to the reinsurer and D is a deductible. We restrict our analysis in this simulation solely to policies without deductibles, but provide for their inclusion for completeness.

Simulation Procedure

For each insurance portfolio in Table 5.5 we simulate losses via the following procedure.

1. Set severity distribution for each contract as in Table 5.5.

2. Set frequency distribution as per Table 5.5 — Poisson $\sim \pi_L.k$ where k is the number of each contract contained in the portfolio.

3. Randomly sample the frequency of expected losses for each contract in the portfolio, to generate a number of losses for *each contract*, N_{loss}.

4. Randomly sample from the severity distribution for each contract N_{loss} times, sum and record the losses.

5. Run the above process 100,000 times.

The results of the simulations are presented in Table 5.8 (histograms of the generated loss distributions are provided in Figure 5.6). The table contains the premium income for each insurer as previously determined, a capital level assumed to be held by the insurer equal to the average baseline loss in the simulation and reserves defined,

$$\text{Reserves} = \text{Premium Written} + \text{Capital} \tag{5.4.8}$$

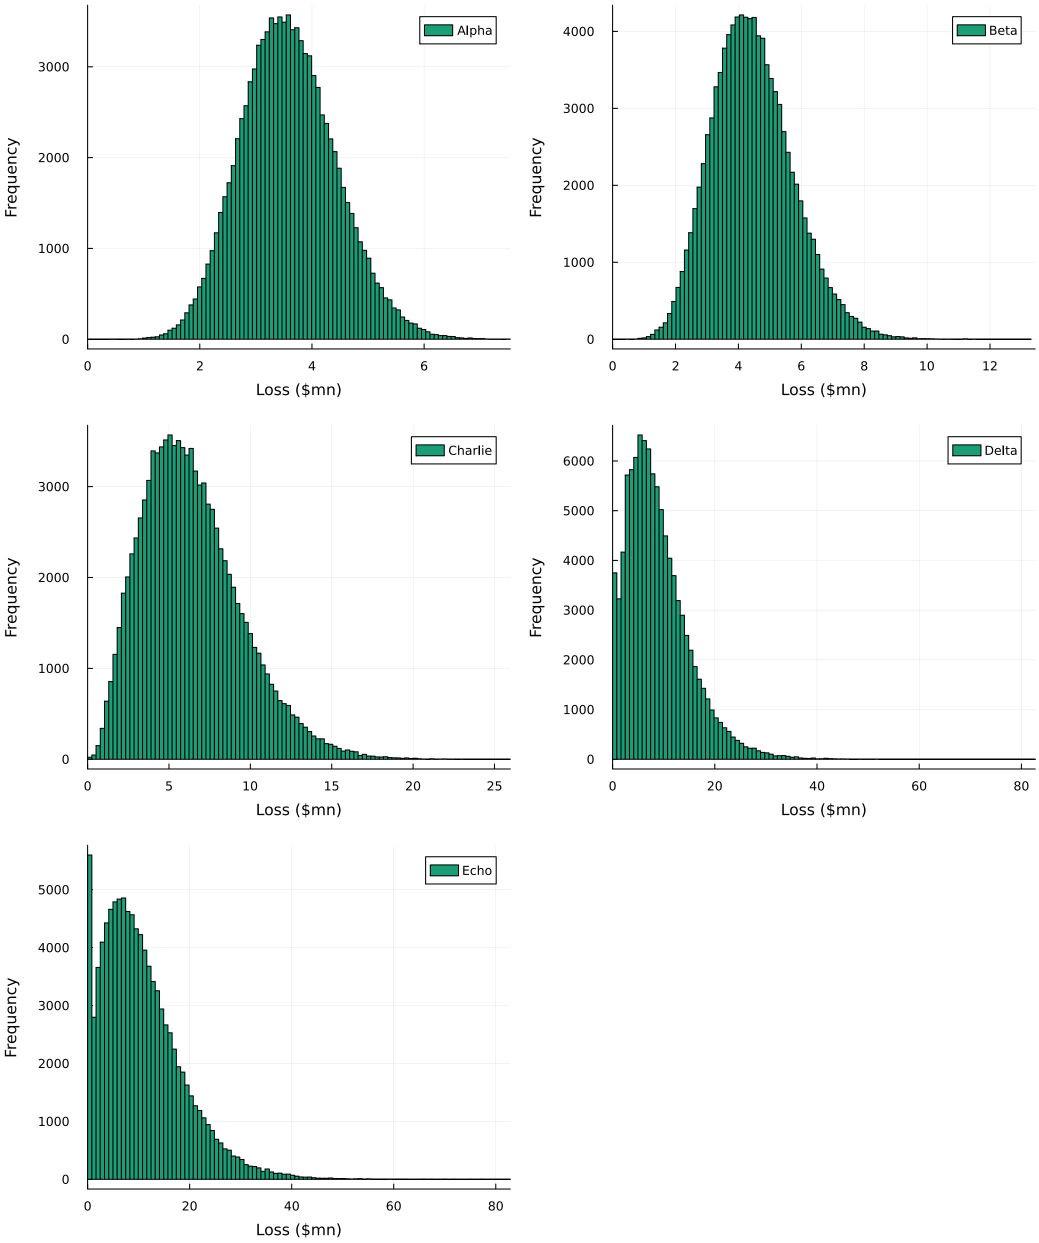

Figure 5.6: Insurer Simulated Loss Distributions (Section 5.4.4)

Along with the simulated loss values, we also calculate loss values for a ‘stress test’ type scenario, calculating the maximum loss in 95% and 97.5% of cases[11]. This is done via using the *quantile* function of *Distributions.jl* to calculate the respective frequency and severity at $F^{-1}(0.95)$ and $F^{-1}(0.975)$. The required values are then readily obtained. With these values obtained, we may now proceed to consider the interaction between reinsurance and the insurance portfolios.

	Assets			Losses			
Insurer	Premium Income	Capital	Reserves	Simulation Baseline Average	Simulation Baseline SD	95% Stress Test	97.5% Stress Test
Alpha	\$8.0mn	\$3.6mn	\$11.6mn	\$3.6mn	\$0.8mn	\$8.2mn	\$9.4mn
Beta	\$10.0mn	\$4.4mn	\$14.4mn	\$4.4mn	\$1.3mn	\$13.6mn	\$17.4mn
Charlie	\$14.2mn	\$6.4mn	\$20.6mn	\$6.4mn	\$3.0mn	\$28.0mn	\$36.6mn
Delta	\$19.8mn	\$8.9mn	\$28.7mn	\$8.9mn	\$6.2mn	\$51.2mn	\$64.9mn
Echo	\$24.0mn	\$10.8mn	\$34.8mn	\$10.8mn	\$7.9mn	\$53.1mn	\$77.0mn

Table 5.8: Simulated Losses

Considering the effect of capital

Suppose, as per Table 5.8 that the insurer has a capital buffer, which initially, is equal to the simulated average losses on its portfolio. We now examine the optimal reinsurance fraction which means the insurer would remain solvent in the event of losses of a specified magnitude. We consider ρ values for both the 95% and 97.5% stress tests. This means calculating the value of ρ which would set $\$\text{Profit}(L) = -K$ (Equation 5.4.6). The required expression is

$$\bar{\rho}(L_{\text{stress}}) = \frac{(L_{\text{stress}} - \$\text{Premium Written} - K)}{L_{\text{stress}} - \$\text{Premium Written} + (\$\text{Exposure} \times \%\text{CC})} \tag{5.4.9}$$

where $\%CC$ is the percentage ceding commission.

The solvency threshold for the insurer is Reserves $= L_{stress}$. If Reserves $> L_{stress}$ then we set $\bar{\rho} = 0$ as the insurer does not need reinsurance at this stress test loss level as it would remain solvent without it. For insurer Alpha, reserves exceed the stress test losses at both thresholds, while for Beta, reserves exceed only the 95% stress test loss. Figure 5.7 and Table 5.9 show the complete results of the

[11]

analysis. Starting with Table 5.9, it appears that neither Alpha nor Beta should buy reinsurance. In the 97.5% stress-test, Beta is insolvent even with reinsurance. This suggests that Beta would need to implement a higher loading than that initially calculated to pass the stress test. For Charlie, Delta, and Echo, there is benefit in purchasing reinsurance as a quota share policy, though the optimal fractions appear fairly high. Consequently, the insurers might decide to buy less than the optimum but set capital higher. However, this then means that the market is not efficient. Table 5.9 also shows the profit each insurer would receive if *ex post* losses equal the simulated baseline with no reinsurance; with reinsurance at the $\rho^{0.95}$ fraction; and with reinsurance at the $\rho^{0.975}$ fraction. For Charlie, given that its weighted average premium rate is close to the reinsurer objective, it receives a scant ceding commission. Consequently, there is an opportunity cost of \$4.0-5.2mn of purchasing quota share at the optimum relative to baseline simulated profit of \$7.8mn. In a market where information is shared, there should not be an opportunity cost. For Delta and Echo, the purchase of quota share appears more attractive because of the more generous ceding commission. These are deliberately extreme examples, but in practice suggest that bargaining may occur between different insurers and reinsurers over the ceding commission, which introduces inefficiency into the market.

				Profit if Losses=Simulation Baseline ($mn)		
Insurer	Ceding Commission	$\bar{\rho}^{0.95}$	$\bar{\rho}^{0.975}$	$\rho = 0$	$\rho = \bar{\rho}^{0.95}$	$\rho = \bar{\rho}^{0.975}$
Alpha	-6.0%	0.00	0.00	4.4	4.4	4.4
Beta	-4.0%	0.00	0.87	5.6	5.6	-2.7
Charlie	0.2%	0.53	0.71	7.8	3.8	2.4
Delta	5.8%	0.60	0.71	10.9	7.8	7.3
Echo	10.0%	0.47	0.67	13.3	11.7	11.1

Table 5.9: Reinsurance ceding fractions that maintain insurer solvency at different stress-test values

Figure 5.7 presents a more detailed picture of the simulations that yield the optimal ρ. For each insurer portfolio, we plot the insurer profit (Equation 5.4.6) as a function of losses for values of ρ between 0 and 1. The capital held (i.e. the average simulated loss as already discussed) is represented as a horizontal line and the average simulated losses for the 95% and 97.5% stress tests are represented as vertical lines. The intersection of the average simulated loss and the stress test allows for the optimal $\$rho$ to be read from the graphs. In the case of Alpha, it can be seen that on the $\rho = 0$ profit line, at the two stress test loss values (Points A and B), the profit exceeds the capital held. For reinsurance to be worth purchasing, the $\rho = 0$ profit line must be less than the capital horizontal lines at the stress test losses. Taking Echo as an example, with stress test loss of 95%, we can see that the horizontal capital and vertical loss lines intersect between the profit lines for $\rho = 0.4$ and $\rho = 0.6$ (Point C) As may be verified from Table 5.9, the reinsurance fraction for this case is 0.47. The comparable intersection for the 97.5% stress test (Point D) is at $\rho = 0.67$.

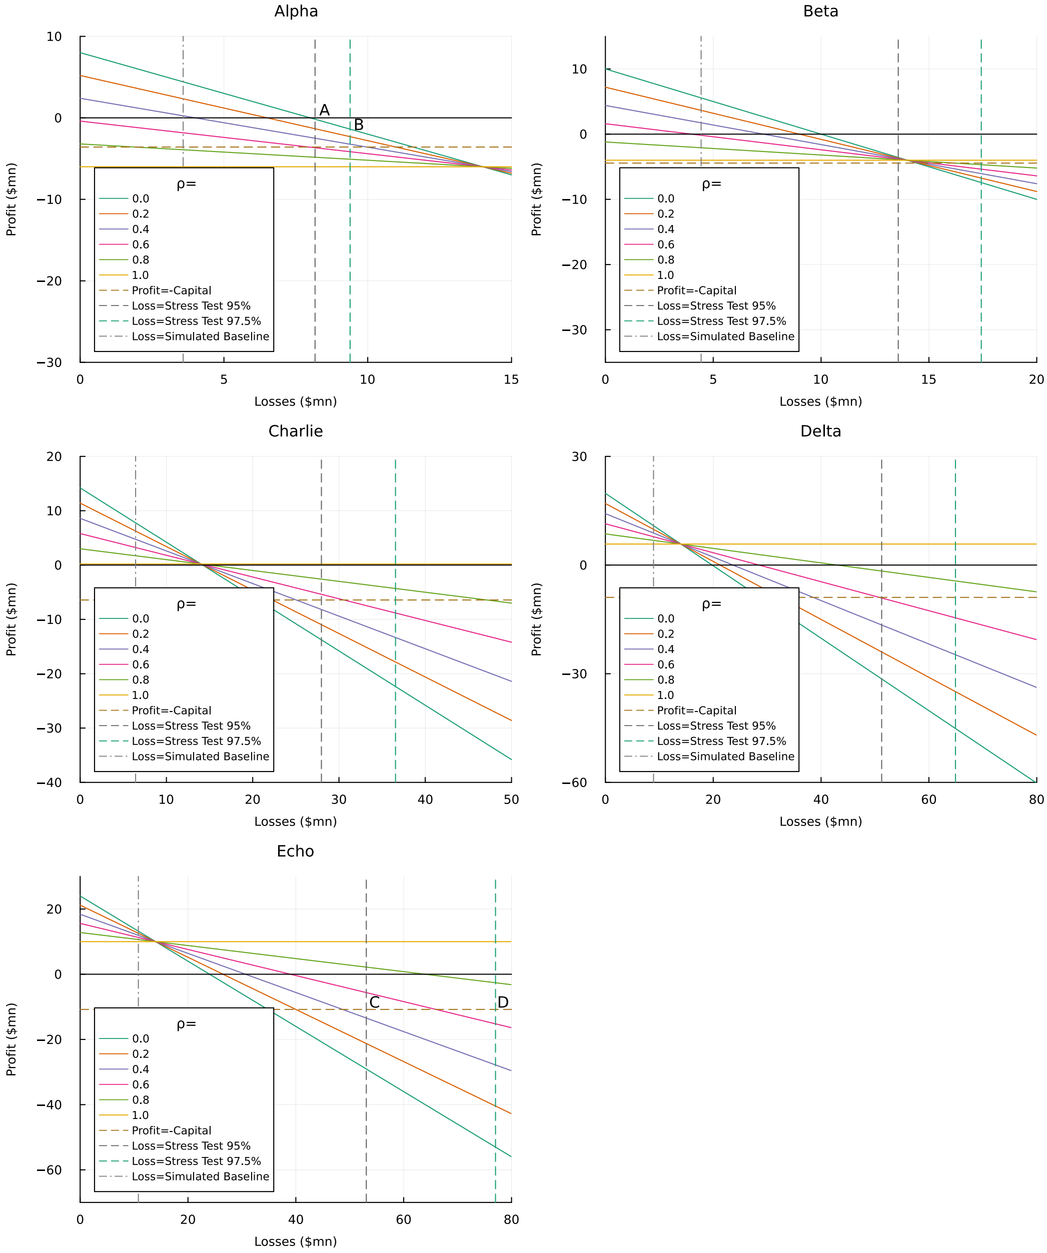

Figure 5.7: Insurer profit with varying quota share proportions (ρ).

Insurer	XL Reinsurance Coverage	Probability of Loss > Cash Premium	Technical XL Reinsurance Premium	QS Reinsurance Premium at $\bar{\rho}^{0.975}$
Alpha	\$1.4mn xs \$8.0mn	0.0%	\$0.0mn	\$0.0mn
Beta	\$7.4mn xs \$10.0mn	0.0%	\$0.0mn	\$8.3mn
Charlie	\$22.4mn xs \$14.2mn	1.6%	\$0.4mn	\$5.4mn
Delta	\$45.1mn xs \$19.8mn	5.9%	\$2.7mn	\$3.6mn
Echo	\$53.0mn xs \$24.0mn	6.6%	\$3.5mn	\$2.2mn

Table 5.10: Excess of Loss Pricing Example

Excess of Loss

Having considered the quota share case, it is worth considering the case of excess-of-loss insurance as an alternative to quota share for the insurers. Rather than using the capital buffer approach, we consider a simpler objective: that the insurer rather than holding a capital buffer buys insurance from a reinsurer to cover losses in excess of its cash premium income up to the limit of the 97.5% stress test loss value. To calculate the required parameters, we can use the simulated baseline losses already calculated in Section 5.4.4. From these, we compute the number of instances of losses in the vector of generated losses that exceed the cash premium income but are less than the 97.5% stress test loss value.

The results are contained in Table 5.10. The portfolios of Alpha and Beta generate expected losses well below the level of cash premium income (see Figure 5.6) and accordingly there is little benefit in excess-of-loss insurance. For Charlie, Delta, and Echo, it is interesting to note that the combined technical premium is \$7mn. Recall that in Section 5.4.4, Distribution A in Table 5.4 gives the loss ratios for the reinsurer versus quoted premium for an expected \$10mn of losses. If we assume that the reinsurer requires a loss ratio of 0.5 or better, then the minimum premium it will charge is 4%. For the insurance buyers, only for Echo is buying excess-of-loss reinsurance cheaper than buying quota share. Consequently, for each insurance portfolio, there is a different optimal reinsurance contract from the perspective of the insurance company seeking reinsurance.

The excess-of-loss premia calculated in Table 5.10 are computed using the individual joint distributions of frequency and severity for each of the five insurance companies. These are known only to each of those insurance companies alone and are not visible to the reinsurer. Consequently, there are information asymmetries between the insurers seeking reinsurance and the reinsurer. Insurers Delta and Echo know that the fair insurance premium rate for the excess-of-loss contracts specified in Table 5.10 are 5.9% and 6.6%, respectively. However, the reinsurer would offer these contracts at 4% premium rate based on its own distribution. Consequently, the insurers can, under these assumptions, buy reinsurance cheaper than its fair cost based on their avantageous knowledge of the 'true' distribution rather than the reinsurer's distribution which assumes simple log-normal distribution of a set of risks at a particular expected loss value. This illustrates how inefficiency and therefore financial imbalances between insurance and reinsurance may emerge as a consequence of different expected loss distributions, unlike in Table 5.2 where the reinsurer and insurer(s) had the same distribution of expected losses.

5.4.5 Simulation 3: Insurance buyers of variable risk

We now consider a simulation in which buyers have heterogeneous preferences and risk tolerance. The interactions of the real insurance market are hard to model as insurance customers interact with insurance companies via insurance brokers who act as an intermediary. The flow of business is directed therefore partly by relationships (and so is not efficient in a traditional economic sense). However, it is possible to construct some simulations of insurance demand based on different characteristics and illustrate the utility demand model and how this may affect reinsurance pricing.

The insurance buyer faces a single utility maximization decision: for a given premium rate, how much cover does the agent with to purchase. This could be formalised in terms of expected utility (Equation 5.3.4) via variation of the risk aversion parameter, α, but this is not necessary for the example presented here. The insurance company must choose premium rates that it believes will not excessively deplete its capital for a certain level of risks, or plan to cede premium to reinsurance to cover that risk as demonstrated in the previous section. We will retain the contract limit structure from Table 5.5 for this analysis, meaning that insurance buyers choose one of the five contracts.

We will now assume that the more coverage the buyer takes, the more sophisticated its assessment of the risks are. This places a constraint on the amount of loading the insurer can apply to the higher limit contracts. We will, as previously, fix the total *potential* cover available in the market at \$500mn and consider how this may be allocated among buyers. However, as will be illustrated, the risks associated with some contracts make them commercially unviable even if theoretically priceable. Table 5.11 sets out some arbitrary premia based on the subjective beliefs of the respective buyers, and the maximum number of contracts available in the market based on the overall capacity of \$500mn. We wish to stress that these numbers are established purely for model convenience and to illustrate the further difficulties to establishing efficiency under heterogeneous buyer beliefs. The assumption of market size is required to price potential reinsurance on insurer policies.

For this analysis, we set the expected severity loss mean equal to a quarter of the policy limit and the standard deviation to half the mean. Unlike in the previous section, we will allow the distribution of expected losses to vary with different clients and have a mixture of buyers considered low-, medium-, and high-risk with different distributions accordingly. We assume that the variation in risk characteristics of the three buyer groups is expressed through variation in frequency.

	Highest premium rate at which a buyer takes full coverage			Maximum number of customers		
Limit	Low risk	Medium risk	High risk	Low risk	Medium risk	High risk
\$500k	14%	20%	26%	46	46	46
\$1mn	13%	18%	23%	32	32	32
\$2mn	12%	16%	20%	16	16	16
\$5mn	11%	14%	17%	8	8	8
\$10mn	10%	12%	14%	4	4	4

Table 5.11: Insurance buyer premium ceilings

We assume that reinsurers consider the risks involved for the three different risk categories and apply distributions A, C, and E (Table 5.3) to low, medium and high risks effectively, and target loss ratios of 0.3, 0.5, and 0.7 respectively. This means that the reinsurance charges for the portfolios are 6%, 11% and 13%.

We now consider the distributions associated with the different contracts. Table 5.12 shows the severity and frequency distributions for each policy. We have fixed the severity on each contract and assumed that riskier clients have a higher expected frequency of claims. This assumption could, of course, be varied further, but this approach suffices for the purposes of this example. From this, we simulate the losses with 100,000 runs and derive the expected loss for the entire set of possible contracts. This is shown in Table 5.13 along with the expected average loss per contract derived

.

	Severity			Frequency, Poisson(λ)		
Limit	μ_L	σ_L	Distribution	Low risk	Medium risk	High risk
\$500k	\$125k	\$62.5k	LogNormal(11.6,0.22)	4.6	11.5	23
\$1mn	\$250k	\$125k	LogNormal(12.3,0.22)	6.4	12.8	19.2
\$2mn	\$500k	\$250k	LogNormal(13.0,0.22)	4	8	12
\$5mn	\$1.25mn	\$625k	LogNormal(13.9,0.22	2	4	6
\$10mn	\$2.5mn	\$1.25mn	LogNormal(14.6,0.22)	1	2	3

Table 5.12: Distribution specification for insurance contracts offered to buyers

With this calculated, we can then derive the technical premium for each contract, which is shown in Table 5.14. Comparing with Table 5.11, we can see that for

	Expected Loss (Total, $mn)			Expected Loss per Contract ($k)		
Limit	Low risk	Medium risk	High risk	Low risk	Medium risk	High risk
$500k	0.53	1.32	2.63	11	29	57
$1mn	1.47	2.93	4.40	46	92	138
$2mn	1.84	3.67	5.50	115	229	344
$5mn	2.28	4.59	6.89	285	574	861
$10mn	2.30	4.59	6.86	574	1,147	1,716

Table 5.13: Expected losses for policies

the $5mn and $10mn limits, the high risk technical premium is higher than what customers are willing to pay. It may be possible in this case for the insurer to instigate a deductible and reduce the premium. Otherwise, margin is very limited for medium-risk $5mn and $10mn limits, which might also motivate introducing a deductible.

	Technical Premium (%)		
Limit	Low risk	Medium risk	High risk
$500k	2.3	5.7	11.5
$1mn	4.6	9.2	13.8
$2mn	5.7	11.5	17.2
$5mn	5.7	11.5	17.2
$10mn	5.7	11.5	17.2

Table 5.14: Technical premium for insurance contracts

We now consider the capital requirements associated with the insurance policies. Table 5.15 shows the expected losses for $F^{-1}(0.995)$ and $F^{-1}(0.5)$ for frequency and severity respectively for both the whole set of contracts and also per contract. Each insurer must decide how to allocate its available capital and how much reinsurance to purchase. Rather than calculating sample portfolios, we will simply calculate the reinsurance fraction that is optimal based on Equation 5.4.6.

	Stress Test Loss (Total, $mn)			Stress Test Loss per Contract ($k)		
Limit	Low risk	Medium risk	High risk	Low risk	Medium risk	High risk
$500k	1.2	2.3	4.0	27	51	87
$1mn	3.1	5.1	6.9	98	161	217
$2mn	4.5	7.2	9.8	280	447	615
$5mn	6.7	11.2	14.5	839	1,398	1,817
$10mn	8.9	13.4	17.9	2,236	3,354	4,472

Table 5.15: Stress Test losses for policies, with Frequency set at $F^{-1}(0.995)$, Severity at $F^{-1}(0.5)$

Based on the Stress Test loss values, and assuming that the insurer writing each contract holds capital equal to the expected value of losses for the contract (Table 5.13), we can then derive the optimal reinsurance fraction for each contract. As in the prior section (Table 5.9), this is calculated by calculating the reinsurance fraction that sets the profit to the insurer equal to $-K$, i.e. at the level of loss given in the Stress Test, the insurer breaks even if it holds this proportion of reinsurance. As the buyers of the smaller contracts are less knowledgeable and will accept a higher premium, the reinsurance fraction is lower as the insurer writes more premium. However, the reinsurance fraction increases from an average of 20% for the \$500k limit contract to as high as 64% for the medium-risk \$10mn limit contract. It is clear from this analysis that while it is possible to achieve risk transfer between insurance buyer, insurance company and reinsurer, for a simulated market, achieving convergence of distributions is extremely unlikely as each party is incentivized to maximize their profit rather than target efficiency.

	Optimal reinsurance fraction for each contract		
Limit	Low risk	Medium risk	High risk
\$500k	0.23	0.23	0.20
\$1mn	0.31	0.30	0.25
\$2mn	0.41	0.40	0.36
\$5mn	0.51	0.53	0.48
\$10mn	0.63	0.64	0.60

Table 5.16: Optimal reinsurance purchase fraction for each contract implied by stress test values

We have stopped short of simulating the allocation of policies to individual insurers as to model competitive market dynamics under uncertainty with heterogeneous beliefs is a complex problem that in itself might fill multiple papers. However, it is hoped that the simulation presented illustrates the additional dynamics that heterogeneous buyer beliefs brings to the challenges of modelling cyber-insurance and re-insurance. To place the simulation results in context with the US cyber-insurance market, in 2020, according to the NAIC [210], there were approximately 4 million cyber-insurance policies written in the US market, with the top 20 insurers taking 68% market share. The report for 2021 does not provide a policy number, but notes that almost 50% of cyber-insurance premia were ceded.

5.5 Discussion

The simulations show the difficulty of achieving economic efficiency in an artificial cyber-insurance market even using relatively standard distributions and contract structures. However, as has been stressed, just because a market is not efficient does not mean that transactions cannot take place. We now consider some of the further informational barriers to facilitating smooth transfer of cyber-risk. Issues of data transparency, attack measurement, and reporting — making relevant data publicly available — are particularly crucial in enabling agents to make informed pricing decisions.

5.5.1 Information asymmetry

By and large insurance and reinsurance companies operate in environments where high quality precision signals about loss risks exist. For example, in the case of natural catastrophes, their frequencies are well known and established over many periods. Further, there are enough tail events to help construct reasonable approximations of extremes. When it comes to events regarding human interactions, such as crime, illness, death or accidents, these are reported by statute to the relevant central authorities. This data is publicly available. In both these cases agents at all levels share the public signals and can condition their private expectations on good quality evidence. Of course, there may be variability in the accuracy of private expectations based on individual interpretation of the data or circumstances. This set-up allows the buyers of insurance the calculate their expected loss in a well informed manner and the insurance companies, based on the public information, can quote a premium. In turn the reinsurers share the same beliefs as no further information is available to them regarding the likelihood of the different states of nature.

When it comes to cyber-risk and cyber-insurance, the state of data curation and sharing is far more nascent than for other insurance perils and it is reasonable to argue that there is no high quality public signal to inform all agents' priors. In the regulation of the aviation industry, it is standard to require reporting of 'near misses' so that lessons can be learnt and procedures updated to lessen the risk of future accidents. It is possible that this might be addressed by vendor telemetry — an insurer might have a series of recommended cyber-security solution providers that their clients could sign up for as part of their insurance package who would share data with the insurer. This raises potential issues of confidentiality.

5.5.2 Cyber-insurer loss experience

The United States National Association of Insurance Commissioners publishes an annual report on the cyber-insurance market derived from its Property/Casuality Annual Statement [210]. Table 5.17 presents this information for the four years currently available. In 2018 and 2019, the data was presented separately for standalone and package policies but in 2020 and 2021 was presented for combined policies. We have adjusted for this to present the data on a consistent basis.

It is notable that the ransomware epidemic from 2020 to 2021 had a marked effect on experienced loss ratios for some insurers[12]. However, there are pockets of differentiation. For example, the Hartford Insurance Company specialises in insurance for smaller companies, creating a fairly well diversified portfolio of insurance contracts where the holders are unlikely to fall victim to sophisticated, targeted ransomware attacks given the potential revenue available. For these companies, basic defences and security software should help mitigate against losses.

Figure 5.8 plots the losses experienced in the underwriting year versus the premium written and a linear trend line with intercept fixed at 0. The slope of the fitted trend line is then the loss ratio. The average loss ratio remained fairly stable across the two years, but it is striking that less than 30% of premia received was, on average, retained by the underwriting insurer. The aforementioned NAIC report states that some 50% of premia for cyber-insurance was ceded to the reinsurance market.

[12] This has been widely reported in the trade press – see, for example, [65]

	Direct Written Premium ($mn)				Loss Ratio			
Firm	2018	2019	2020	2021	2018†	2019†	2020	2021
CHUBB LTD GRP	320.73	355.28	404.14	473.07	28.6%	27.7%	61.0%	76.9%
FAIRFAX FIN GRP	38.15	65.01	108.69	436.45	23.4%	51.6%	55.7%	51.9%
AXA INS GRP	255.87	229.68	293.03	421.01	57.2%	65.7%	98.2%	86.5%
TOKIO MARINE HOLDINGS	44.59	46.91	78.16	249.79	30.6%	17.1%	51.1%	43.8%
AMERICAN INTL GRP	232.31	226.20	228.42	240.61	36.1%	55.4%	100.6%	130.6%
TRAVELERS GRP	146.23	178.53	206.82	232.28	22.4%	32.1%	85.5%	72.7%
BEAZLEY INS CO INC	110.95	150.94	177.75	200.88	7.8%	22.0%	47.9%	38.7%
CNA INS GRP	83.36	94.72	119.61	181.38	26.9%	33.2%	105.7%	87.5%
ARCH INS GRP	—	—	—	171.94	—	—	—	9.2%
AXIS CAPITAL GRP	76.00	97.31	133.55	159.06	7.2%	18.5%	46.2%	105.2%
ZURICH INS GRP	43.32	43.67	64.43	151.87	18.2%	86.9%	40.4%	76.9%
LIBERTY MUT GRP	66.50	68.38	41.86	138.22	38.9%	23.3%	30.0%	95.2%
SOMPO GRP	34.05	49.71	72.59	133.52	56.7%	29.3%	114.1%	54.3%
BCS INS GRP	69.50	76.06	86.58	132.04	10.4%	32.9%	59.1%	80.1%
HARTFORD FIRE 7 CAS GRP	39.70	49.74	102.86	123.16	16.4%	31.6%	25.4%	16.3%

Table 5.17: Cyber-insurer loss experience in the US market
(†denotes weighted average by DWP)
Source: NAIC, Researcher calculations

There is some evidence to support the premise of a disconnect between expected and experienced losses in cyber-insurance pricing. Woods et al (2021) [298] develop a distribution of cyber-losses based on insurance company filings in the United States. They note that their model significantly under-predicts losses in relation to *ex-post* losses reported in other literature. The under-pricing of premia implies that either

- Insurers believe they can diversify loss risk.
- Customers were not willing to pay the technical premium and insurers are pursuing a 'loss-leader' strategy.

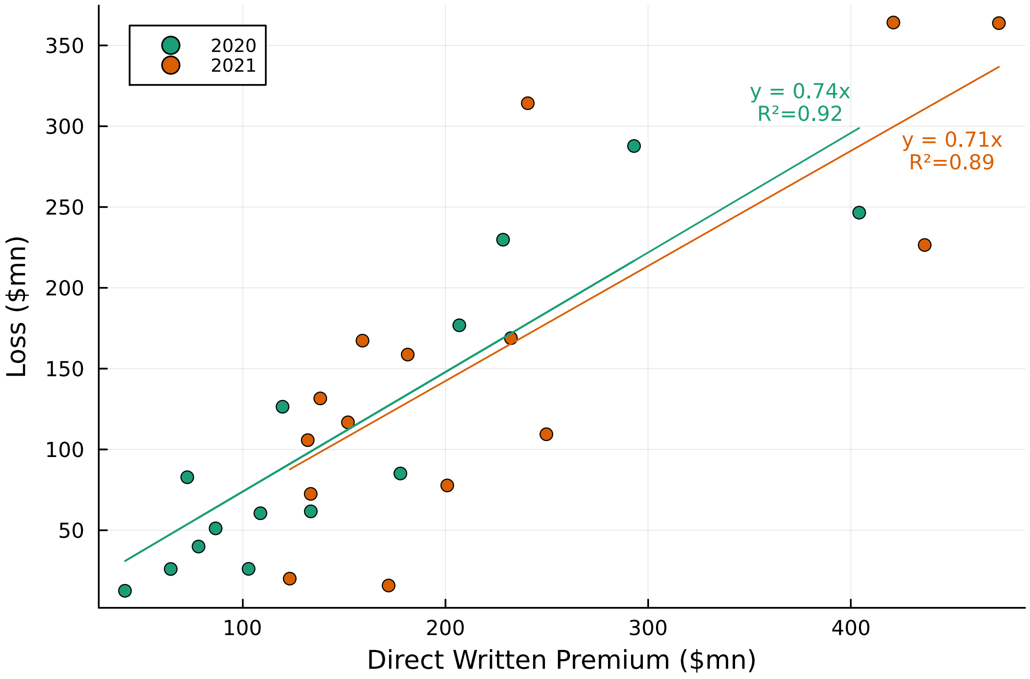

Figure 5.8: US cyber-insurer losses vs premium written

The entry of Arch Insurance also merits comment. Arch insurance provides capacity[13] to a relatively new managing general agent (MGA), Coalition Inc., providing 'active cyber-insurance'. Active cyber-insurance is a relatively new product, which merges the roles of an outsourced security provider and a traditional cyber-insurer. This reduces some of the risks of asymmetric information transfer associated with cyber-insurance from the perspective of the insurer. The trade-off between

[13] https://www.coalitioninc.com/en-ca/announcements/Arch-Insurance-Backs-Coalition-With-Long-term-Capacity-Across-Cyber-Insurance-Programs

cyber-insurance and security investment has been modelled by Mazzoccoli and Naldi (2020) [189] and Skeoch (2022) [268].

Comparison to simulated results

Comparing the experienced losses by insurance companies, our assumption regarding the adoption by re-insurance firms of their own private distributions for both severity and frequency of successful cyber-attacks and subsequently losses at this stage of development of this nascent market seems well-grounded on the available evidence.

The evolution of proportional losses across 15 major insurance companies over the period 2018-2021 presented in Table 5.17 reveals a somewhat unstable path. Both the average loss and its distribution exhibits both wide variability and an increasing trend. Specifically in 2018, average losses were 25.3% of the premia collected and this measure has monotonically increased to 68.3% by 2021. At the same time the maximum losses have more than doubled from 57% to 130% by 2021. The cross-sectional standard deviations exhibit the same monotonic trended pattern.

Attempting to fit a log-normal distribution over the whole period for the companies in the sample using the same methodology for fitting such distributions in the simulations shows that the kernel[14] of the empirical distribution deviates significantly from the normal and reveals slight bi-modality (Figure 5.9). It is also notable that the fitted distributions underestimates the tail of large losses, which is arguably a significant consideration for reinsurance companies.

Faced with such movements of the cross sectional distributions, meaningful aggregation of the losses experienced by individual insurance companies does not seem effective. In the light of this (admittedly cursory) review of the statistical evidence presented in this paper, Assumption 4 in Section 5.3.3 seems justified.

5.5.3 Loss transparency

We consider what happens if agents only selectively claim on losses from an insurer. In an insurance analysis, it is usually assumed that every agent is aware of the attacks they experience. This is a reasonable assumption for some categories of

[14] See Epanechnikov (1969) [96]

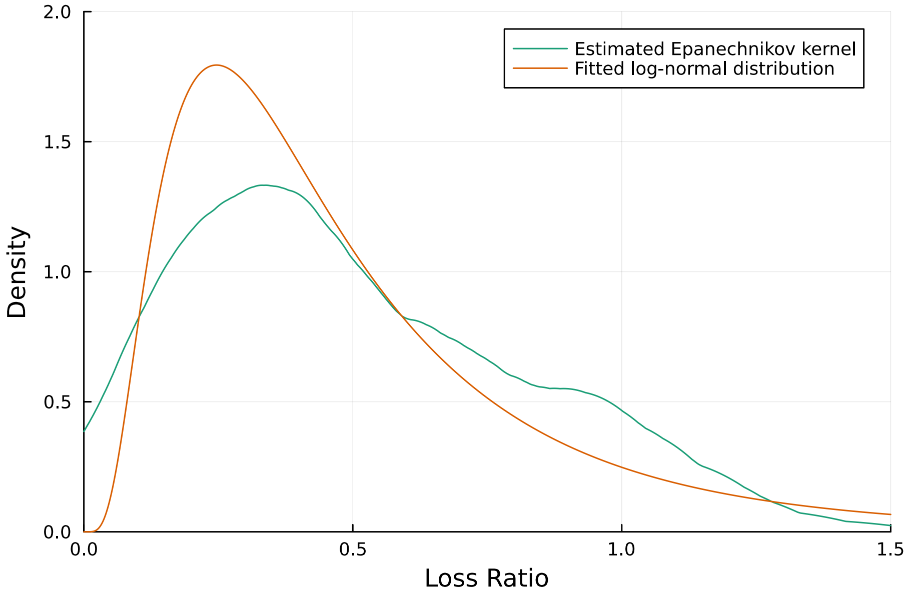

Figure 5.9: Epanechnikov kernel versus fitted log-normal distribution for NAIC reported cyber-insurance loss ratios, 2018-2021

cyber-incidents, such as ransomware, although other cyber-incidents such as data breaches might not be detected until some time after the event. Agents report some attacks to an insurer and thus a claim is made; some attacks go undisclosed (in insurance, this is known as IBNR — incurred but not reported). More formally, at time t, the agent may be aware of the attack and its damage so the state of the world in which the attack occurs, s is known to them. The agent might inform the insurer about the state so the insurance knowledge of the state s is conditional on the revelation of the agent. Now, the insurer knows that their distribution is not the objective one but only a partial revelation due to the agents selectively choosing to report losses. The insurer then tries to approximate the objective distribution but it will be with error. In the event that reinsurers know that different insurers have different approximations of the true distribution, they will use some kind of averaging across these approximations to quote reinsurance premiums. The results are:

- No insurer is offered a fair premium given their approximation of the true distribution.
- No agent is offered a fair premium as the insurance offer is based on a distribu-

tion different to their own.

- Objectively measured data is absent at all levels because reporting is a choice.

5.5.4 Consistency of reference

There is a significant problem with the standard actuarial modelling cycle approach to cyber-insurance: the evolution of systems over time, which is quite unique in its complexity in relation to other perils. Calibration of models using events such as WannaCry have poor future predictive power as the security vulnerabilities it exploited have been patched, Windows XP is less widespread than it was and the operating systems that replaced it have better, though of course not perfect, security by design. In economics, this can be couched in clients' Bayesian updating of their distributions; they do not and cannot observe attacks on other clients (other than indirectly via media reports) so there is no need to converge to a stationary distribution at the client level. The consequence of this is that the insurers and reinsurers may have a better understanding of the fair price of risk, but buyers do not share the same concern and thus are not willing to pay the demanded premium for the insurance.

5.5.5 Supply and demand

In the insurance industry, it is common to describe the state of the market as 'hard' or 'soft'. In a soft market, supply exceeds demand placing downward pressure on premium, whereas in a hard market the converse is true. Often the experience of losses in a particular class of business will result in a market hardening. This has important implications for the pricing of cyber-insurance by a vendor. In a soft market, the insurer must charge the lowest premium it can actuarially justify to build market share. In a hard market, the insurer should charge the highest realistic premium possible. If the market were efficient, it would converge to some form of equilibrium but if not it may swing between financial imbalances. There is evidence that in the early stages of the cyber-insurance industry, some insurers operated a very experimental approach to pricing. Woods (2023) [297] provides an account of one large US insurer, , whose Chief Operating Officer admitted that their early cyber-insurance models were a “complete guess”. The same insurer then suffered loss ratios of 100% and 130% in 2020 and 2021, respectively (Table 5.17), suggesting that even if refined and updated, the pricing models may have underestimated the claim frequency or severity.

5.5.6 Further Work

We have considered simulations in which losses are uncorrelated. An interesting next step would be to consider the correlation of losses and implement the modeling strategy presented in this chapter using more complex loss-generating functions, such as those reviewed in Section 2.3.2, than the simple joint distributions of severity and frequency used in this chapter. It would also be instructive to compare the results of simulations of distributions proposed by Eling et al (2019) [91] and Woods et al (2021) [298], with insurer loss data. Claims data is deeply confidential to insurance companies, however, so the results of such analysis would unlikely be able to be widely disseminated unless extensively anonymised.

In the simulations, we focused on the supply dynamics of insurance and in particular the interaction between insurers and reinsurers. The model provides for consideration of buyer preferences, which at this stage we have explored only briefly in the first simulation to illustrate how buyer utility can affect coverage. A further piece of work would be to explore the price sensitivity of buyers of insurance coverage and how these preferences propagate through the information chain to reinsurers.

5.5.7 Conclusions

This chapter has developed an artificial yet realistically structured model of the cyber-insurance market considering all three levels of agent interactions. The model incorporates the demand choices of the consumers/buyers of cyber-insurance, their suppliers — insurance companies offering contracts — and reinsurance companies providing additional underwriting capacity.

The extent to which an insurance market facilitates smooth risk transfer is linked to the sharing of information by participants regarding the distribution of losses.

We argue that this condition is very unlikely to hold in the cyber-insurance market. Disagreements on loss expectations means that cyber-insurance contact pricing will be considered inefficient at both the retail and wholesale levels, leading to lower societal benefit. The purpose of this chapter was to quantify such inefficiency within the confines of a three-tier market under miscellaneous types of disagreements in loss expectations among the participants at each tier.

To establish a benchmark to gauge the extent of inefficiency, we have simulated a simple market where all agents share a distribution of losses based on two loss frequencies. From this simulation, we obtained the efficient measures of reinsurance premium and the proportional participation of reinsurers. We found that simulated loss reduction to the insurers is almost identical to the cost of reinsurance (bar small statistical errors), as expected. This case represents the economically efficient market outcome.

Maintaining all the behavioural parameters from the first simulation, we then proceeded to compute expected losses and reinsurance premiums based on diverse distributions held by insurance companies and reinsurers. Both insurers and reinsurers independently price premiums to meet target loss ratios based on distinct and subjective distributions. Under conditions where losses are close to the modal simulated value, insurers are typically not incentivised to buy reinsurance. However, when considering relatively extreme losses under a ‘stress test’ type scenario, the value of reinsurance emerges to some insurers whose distributions are relatively heavy tailed in comparison to others. For such insurers, the upfront cost of such reinsurance is justified by the avoidance of ruin under high loss scenarios.

Even within the confines of this simple example, the divergence in distributions, expectations and objectives demonstrates that efficient pricing is hard to achieve. It should be noted that whilst there are specialists in cyber-insurance operating within the reinsurance market, cyber-insurance itself competes with other lines of insurance for allocation of specialty reinsurance capital. Based on this, we used a uniform cost of reinsurance in the second of our two simulations. This is the outcome of the reinsurer holding a private loss distribution. This condition may reduce the reinsurance capital allocated to cyber-insurance.

Our findings suggest that the cyber-insurance market will continue to face potential financial imbalances. That is, it will be highly profitable for some participants and costly for others. This is already evident in data on cyber-insurer loss data (Table 5.17). There has been considerable progress in the academic literature on theoretical modelling of cyber-losses and on empirical analysis. However, access to reliable and transparent data remains a problem for researchers as insurance claims data is confidential and highly guarded. Braun et al have noted that an insurance-linked securities market to support cyber-insurance may struggle to develop without better cyber-modelling [43]. . Without a means of accessing reliable data on cyber-losses, insurance buyers will have to continue to form highly subjective probability distributions. In a recent paper, Bajoori et al argue for the creation of an official registry of cyber-security experts with a duty to report [19], which has also been proposed by the UK Government[15].

The cyber-insurance market is still at as stage of relative infancy. The current institutional setup does not appear fully conducive to the delivery of efficient market outcomes at this juncture. Achieving efficiency requires commonly held beliefs and stationary loss distributions. Whether such conditions can be achieved and maintained is questionable given the dynamic nature of cyber-threats. Our provisional conclusions are that the most likely market structure will involve firms specialising in particular insurance contracts covering different ranges of loss limits, with varying access to reinsurance based on these contracts. The overall outcome will be that the capital capacity of this market will be below its optimal size under shared informational conditions.

[15] https://www.ncsc.gov.uk/information/ncsc-assured-cyber-security-consultancy

6 Threat Modelling Ransomware Attacks on Enterprise Networks

Threat modelling is a useful exercise for risk analysis on cyber-insurance policies, for as discussed in Chapter 5, the lack of public reporting of cyber-security incidents and evolving nature of technology complicates traditional actuarial modelling of risks. NIST SP800-53 defines threat modelling as, “a form of risk assessment that models aspects of the attack and defense sides of a logical entity, such as a piece of data, an application, a host, a system, or an environment” [214].

In insurance, there are typically two commonly used forms of risk assessment: deterministic and probabilistic assessments [221]. These may be complementary; usually the deterministic assessment is used specifically to determine capital limits for insurers via ‘realistic disaster scenarios’ while the probabilistic assessment is used more broadly for both pricing individual policies and risk analysis.

As demonstrated in Table 5.17 in Chapter 5, 2021 saw significant losses to cyber-insurance companies, much of which is likely due to the emergence of ransomware. This chapter introduces a model for evaluating the potential impact of different ransomware attacks on a model network, which could be of use to an insurer in formulating loss estimates for insured risks based on the structure of the insured’s network and defences.

6.1 Background

Ransomware is a commonly used term for computer code that uses encryption to compromise the availability of files and/or a system with the aim of extracting a ransom from the victim. The concept of ransomware has been discussed in the literature since 1996 when Young and Yung introduced the concept of 'cryptovirology' where they envisaged encryption being used offensively to extort money from a system owner [304]. The literature on ransomware spans the fields of computer science, crime science and economics. Ransomware was largely of theoretical interest until relatively recently owing to the difficulty for criminal enterprises to extract payments from victims. Legitimate financial institutions are in general prohibited from engaging in or facilitating criminal transactions and this rendered cross-border extraction of ransom payments problematic. However, the development of crypto-currencies has facilitated pseudonymous monetary transactions outside conventional financial channels. This has proved to be an effective enabler of cyber-dependent [285] crime such as ransomware.

A Trend Micro white paper in conjunction with Osterman Research in 2021 reported that 50% of surveyed firms lacked the capability to prevent ransomware attacks [244]. The increased use of cyber-insurance by firms risks transferring the burden of costs from the attacked business to the insurer depending on the exact terms of insurance coverage. The ability to model the extent to which an individual risk[1] may be affected by ransomware is of critical importance to a responsible insurer's underwriting strategy in determining the expected frequency of claims and also the severity[2]. Use of economic models such as the one presented in this research may help insurers better understand how to price the risks of ransomware and thus provide better coverage for firms.

[1] In insurance, it is common to refer to policyholders as risks

[2] Average loss per claim

6.1.1 Modelling ransomware requires a clear, disciplined approach

The potential costs of a ransomware attack may be mapped using the confidentiality-integrity-availability (CIA) framework. Assuming the encryption process is perfectly reversible with the appropriate key, the integrity of the information may be unaffected, though there is always the risk of corruption. If the ransomware allows the attacker access to or exfiltration of information, then there is a potential for breach of confidentiality. The encryption methodology can vary in sophistication, but even a relatively rudimentary encryption methodology would prove hard to crack within a limited timeframe given to pay the ransom (unless the ransomware is a common variant using a key that has already been cracked). The interaction between the attacker and defender is potentially nuanced and complex in a ransomware attack relative to other cybercrimes [252]. In ransomware, there may be direct interaction and bargaining between an attacker and defender, whereas other malware such as spyware, keyloggers and so on may compromise the confidentiality of information or its integrity, but the interaction between attacker and defender is usually indirect.

The motivation for the research depicted in this chapter is to develop a model for ransomware infections and defence that has broad accessibility and applicability. It is important, however, to ground this against an established and well-studied class of models to allow for a range of existing techniques to be used to study the model output. This approach is inherently vulnerable to the criticism of abstraction from real world cases. However, a model that perfectly replicates every detail of a system is likely to be difficult to efficiently solve. Thus, there is a balance to be struck between choosing a set of sufficiently sparse parameters to depict the problem and losing practical significance.

A particular challenge in modelling ransomware is finding a framework to capture the intricate architecture of different networks and different attackers. Partially Observable Markov Decision Process models appear to be particularly useful for describing a ransomware infection as they can capture uncertainty about the state of a system from the perspective of the observer. Further the transition structure allows for either deterministic or stochastic outcomes or a mixture of both. Whilst the representation of a network system within such models is relatively simple compared with the complex structure of protocols, privileges and interaction that comprise a network, a POMDP model at least captures the core features of the system architecture in a way that most conventional game theory based models may not.

6.1.2 Distinguishing different ransomware attacks

For the purposes of this research, ransomware is considered to take two forms: worm-like malware without attacker interaction and malware launched by a strategic attacker [137, 220]. The former automated threat can be countered by antivirus companies updating detection signatures and software vendors patching known exploited vulnerabilities [230]. This type of ransomware is economically similar to a mass marketing effort, where a criminal enterprise hopes to gain large numbers of small ransoms. This has been well researched and documented, with backup of data often used as the key defensive strategy. The wide availability of secure (in so far as anything can be) cloud storage mitigates to some extent against risks of the loss of information availability but does not solve the risk of a breach of confidentiality. For individuals, ransomware insurance may be of value in covering the costs of a replacement device should an expensive piece of equipment be rendered inoperable (especially if a backup of data is available).

Backup is however only partially effective for the latter type of ransomware, which targets enterprise networks. Typically, this type of ransomware is introduced via either a malicious email attachment; via direct unauthorised network access (Remote Desktop Protocol, for example); or by exploiting a vulnerability in a system. Under the assumption that the main objective of such an attack is to render key network nodes unavailable in the hope of extracting a ransom, if a prior backup exists it will be of a de facto vulnerable configuration that if restored may be immediately compromised again. There may be a mitigating patch or configuration alteration available, but this is not guaranteed. In the event that the attacked organization places most weight on pure information, a backup is useful. However, this may not solve the potentially significant financial risk of a loss of business operations. The most significant public example of this is the shipping conglomerate, Maersk, who suffered a global logistics outage as a result of the NotPetya malware[3].

6.1.3 Ransomware may incur reparative costs

Once an organization is aware it has been compromised, it may enlist the help of a specialist company providing 'post-breach services' [68]. This may be paid for either by the victim itself, or increasingly commonly by an insurer as part of a cyber-insurance policy. The trade-offs between the cost and benefits of these services is an emerging but potentially fruitful area of research. The decision process following a ransomware infection ranges from attempting only the minimum remedial actions needed to clear the immediate ransomware infection (including paying the demanded ransom) to a complete replacement of all information technology infrastructure including a 'clean install' of all operating systems and software. A rational firm paying for the clean-up itself would choose the minimum cost needed to contain its eventual potential loss. Losses include potential third party claims in case of data leakage, loss of turnover due to business interruption and direct expenses related to combatting the ransomware infection. An interesting question emerges when an insurer is paying for the cost of the clean-up. In this case, subject to the limits of policy and the risks of affecting future premia, there is a potential incentive to spend more than the minimum amount if the firm is not paying for the post-breach costs itself. It is clearly in the interests of the post-breach specialist to maximise its income from such an operation.

6.2 Theory

The model introduced in this chapter is based on partially observable Markov decision processes (POMDPs) (see Kälbling et al [153]), which are a class of stochastic models for decisions based on partial observation of the state of a system. POMDPs are

[3] See, for example, [133] for an interesting account.

generalisations of Markov Decision Processes[4] and are characterised as a 7-tuple

$$(S, A, T, R, \Omega, O, \gamma)$$

where

- S : Set of states
- A : Actions
- T : Conditional transition probabilities between states
- $R : S \times A \to \mathbb{R}$ (reward function)
- Ω : Observations
- O : Conditional observation probabilities
- γ : Discount function

POMDPs may be implemented in the Julia language using the *POMDPs.jl* package [85]. Julia has several advantages for this type of work: legibility of code and outputs via using symbols to represent key parameters, speed of computation and finally the ability to define custom types. A POMDP where observations are known with certainty is a Markov Decision Process (MDP). Section 6.4.2 carefully walks through the construction of the POMDP used in this chapter in detail.

6.2.1 Solutions

Once a problem has been characterised using a POMDP, simulations can be run to evaluate the range of reward outputs based on different actions. POMDPs have a wide range of applications, most notable in fully or partially autonomous decision-making in fields such as robotics or air traffic control [165, 180]. In these instances, the future potential states and actions need to be evaluated to determine the best possible outcome, or reward. The algorithms that are used to calculate (where an exact solution is possible) or best estimate the maximum reward are known as solvers [260, 286, 164]. Without perfect foresight, a decision-maker needs to react to

[4] See, for example, [235] for an introduction to MDPs

the state which they observe, which is known as a belief [249, 172]. The set of actions corresponding to a particular belief is termed a policy [55]; a policy represents the decisions to be taken under different scenarios, or states of the system. The aim of a POMDP solver is to derive or estimate the optimal policy.

6.3 Related work

6.3.1 Game theory models of ransomware infection

A small but high quality body of literature around the economics of ransomware has developed, chiefly organised around a game theoretic treatment of bulk ransomware attacks. Laszka et al model ransomware as a multistage, multidefender game with mitigation via backup [169]. In the game, the first stage is organizations and attackers choosing their backup and attack efforts respectively. In stage two, each organization becomes compromised; those falling victim decide whether to pay the ransom. August et al provide an extremely thorough economic treatment of the problem of software with vulnerabilities potentially exploitable to deliver ransomware [16]. They examine a downstream endogenous recovery decision that influences an upstream security decision. They note that a limitation of prior literature is that the possibility of negative security externalities is not captured. The work is particularly focused around the trade-off between software pricing and potential for ransom, which while of theoretical interest is practically less intuitive as the monetary cost of software is just one factor governing its adoption or utilization.

Cartwright et al develop two prior game theoretic models[5] of kidnapping to ransomware [54]. Their set of payoffs comprises: criminal does not infect computer; release of files for ransom & not caught; files destroyed & not caught; criminal caught after release of files; criminal caught after destroying files. Li and Liao consider a multi-stage game [174]. In stage 1, the attacker launches ransomware attacks on N victims. In stage 2, after observing random, R, victims decide whether or not to pay it. In stages 3 & 4, the attacker follows up with decision making. An interesting innovation by Li and Liao is the introduction of a reputation score for the ransomware

[5] Selten's game [258] and Lapan & Sandler [168]

originator, which guides the decision of the defender on whether to pay the ransom.

Ryan et al consider how the development of targeted ransomware has affected the dynamics of ransomware negotiations, concluding that imperfect information results in a non-trvial optimal strategy for the attacker [252]. Galinkin frames the ransomware defence problem as a lottery and considers how best to remove the incentives based on data from actual ransomware attacks, concluding that off-site backups incentivised by governments are the strongest deterrents [117]. Yin et al conduct a game-theoretic analysis of ransomware via attacker-defender and defender-insurer games [303]. They find that backup strategies are abandoned when recovery becomes too expensive and that the introduction of insurance leads to moral hazard.

6.3.2 POMDP models of penetration testing

There is a reasonably developed, though arguably fairly concentrated, body of literature on the use of partially observable Markov decision process models (POMDPs) for penetration testing. POMDPs will be fully introduced in Section 6.4.2. In brief, they are a class of models that guide structured decision making under uncertain observations. The decision-maker receives a belief regarding the state of a system, from which they may take certain actions. The decision-maker then receives a reward that is a function of the next state reached, which is determined via a set of transitions. The work applying POMDPs to penetration testing is organised around the identification of potential attack paths within a system from a defensive perspective. However, this methodology is equally applicable to the decisions of an attacker albeit the attacker may be more risk averse with regard to potential detection. One possible reason why this study is not more popular is that vulnerabilities in systems can be esoteric and the POMDP model therefore both too general and abstract to usefully model the cases. However, for a broad economic analysis of systems vulnerability, these models may yield useful insights.

Some of the literature is inconclusive regarding the efficacy of POMDPs to model attack chains. Sarraute et al represents an early use of POMDPs for modelling penetration testing by considering the planning of attacks under uncertainty [254], which was later refined to conclude that in general penetration testing is not POMDP solving, for the reason that the specificity of the models is inherently limiting set against the continually evolving information security landscape [255]. Hoffman provides a taxonomy of models in respect of the previous research, but again highlights the limitations of decision models in fully capturing human behaviour [143].

Mehta et al discuss how POMDPs can be used to inform resilient systems design, which is clearly of relevance to understanding how to defend against ransomware attacks [192]. Ghanem and Chen highlight the value in using automated reinforcement learning to replicate and analyzing complex penetration tests far faster than even an expert human might be able to [121]. Schwartz et al present two different POMDP-based penetration testing models [19]. The work appears relatively abstract but the introduction of a discount factor is interesting.

6.3.3 Incident response and recovery

An interesting consideration in cyber-insurance policies is coverage of incident response services, providers of which may be called upon to assist a company respond to a ransomware attack. Woods and Böhme conduct (to the best of our knowledge) the first survey of how insurers address this particular problem [295]. They find that insurers tends to nominate a panel of firms to provide services to insured parties, split between legal, forensics and communications experts. The panel sizes range from just 5 firms (Allianz) to 50 (AIG) within the top 20 US cyber-insurance carriers who make such information public. Woods and Boḧme highlight that the question as to whether insurers have resulted in a worsening of the 2021 ransomware epidemic is an empirical one to which they are not aware of any answers. Further, they fail to distil any stylised facts about ransom procedures finding “considerable variation across insurers and providers”. Without such information, it is arguably difficult for firms to plan a strategy *ex ante* and it is this decision making process that our model aims to assist with. Filiz et al conduct an interesting study into the effectiveness of ransomware decryption tools [108]; the malware in this study is largely of that encountered in the wild rather than the targeted strains covered by the research in this paper.

6.3.4 Business continuity insurance

Business continuity insurance is a long-standing line of insurance, which traditionally covered computer systems and data records under the 'all other contents' definition under 'property damage' [122]. Glynn et al also note that "commercial combined policies have generally sought to exclude hacking attacks and losses flowing from viruses, corruption of data, etc" [122] This clearly therefore excludes ransomware attacks, which are arguably better covered under a cyber-insurance policy.

6.3.5 Network malware models

Jacob et al present an interesting treatment of the issues that might need to be addressed in an automata model of malware, in particular interaction and concurrency [149]. Dalla Preda and Di Giusto offer a formalization of this thinking via the κ-calculus [73]. Cam develops a combined POMDP/logistic regression model for minimising the impact of a malware infection [51]. Liu presents a thorough theoretic analysis of ransomware spreading across a network incorporating its specific topology using an adjacency matrix [179]. The model assumes that the dynamic state of each network node is statistically dependent on the states of its neighbouring nodes. Hu et al use Bayesian attack graphs to model the interactions between a multi-stage attacker and a network, formulating the defence problem as a POMDP [145].

6.4 Model

6.4.1 Problem statement and economic considerations

We define a ransomware attack as the introduction of a malicious process that uses encryption to compromise the availability of a system by attacking the integrity of the system, manipulating existing processes and resources, potentially with loss of confidentiality as well. Confidentiality, integrity and availability are harder to represent mathematically than monetary costs. The economic concept of utility is helpful in this situation as it provides a way to describe the preferences of a decision-maker (often called an agent in the economics literature). A so-called multi-attribute utility function [2] can be defined to represent the preferences of the decision-maker in this problem:

$$U_{defender} = U(\kappa, \iota, \alpha) \tag{6.4.1}$$

where κ represents confidentiality, ι integrity and α availability. $U(\kappa, \iota, \alpha)$ is a multi-attribute utility function, which may vary according to the preferences of the defender. We assume for simplicity that this takes the value of 1 for a system operating according to its specified parameters. This framework accounts for the expected benefits of security investment including insurance coverage in a rigorous manner. It also allows for a decision-maker placing greater weight on the security parameters in formulating choices over the immediate monetary costs of an attack.

The concept of integrity is important in the attack as for a large-scale ransomware attack to be effective, a process needs to be introduced into the target system with sufficient privileges to effect encryption of key files beyond the privileges of the initially compromised user. Targeted resources may include credentials (passwords, keys etc), configuration files (for access control or firewalls). Manipulation of firewalls is particularly important if the attacker seeks to exfiltrate data from the attacked system, though this is not likely to be the primary motivation of a ransomware attack but rather a strategy by the attacker to increase the likelihood of ransom payment. Table 6.1 summarises the utility impact of various different types of attack. It is important to note that the parameters governing behavioural response to a ransomware attack may not just be limited to the direct financial costs of a ransom or business interruption, but also to the impact of an attack on confidentiality of data. If an attacker can encrypt a file, then they can likely also exfiltrate it and companies may need to verify data either decrypted or restored from backups. Accordingly, the model developed in the subsequent sections takes account not only of monetary

Attack	Confidentiality	Integrity	Availability
Data Breach	x	x	
Locker Ransomware		x	x
Double-extortion	x	x	x

Table 6.1: Attack impact on utility

costs, but also confidentiality, integrity, and availability from a utility perspective.

6.4.2 POMDP Model Structure

The set of states[6] for the POMDP model are

$$\text{S} = \{: \text{clean}, : \text{infected}, : \text{locked}, : \text{offline}\}$$

The rationale for choosing these states is that when a machine is compromised, it is not certain that the files contained on it will immediately be encrypted or that it will be locked[7]. In a targeted ransomware attack, the attacker may wish to compromise multiple systems within the network before attempting to extort a ransom. The model states may apply to the system as a whole in the case of a single machine, or to individual machines in the multi-machine cases contained within a vector. The 'offline' state is a terminal state for the single machine case and if a systemically critical machine such as a domain controller is offline in the multi-machine case.

The overall set of available actions is defined as

$$\text{A} = \{: \text{observe}, : \text{repair}, : \text{shutdown}, : \text{pay}\}$$

Within each state, only certain actions are available (Figure 6.2). The actions apply to the system as a whole rather than individual machines. This sacrifices some potential realism but has the benefit of significantly reducing potential dimensional complexity in the model transition structure. For the purposes of this work, the actions $\{: \text{shutdown}, : \text{pay}\}$ are assumed to be terminal. Thus, paying the ransom restores the system to its original clean state without possible reinfection. The

State	Observe	Repair	Shutdown	Pay
Clean	x	(x)		
Infected	x	x		
Locked			x	x
Offline				

Table 6.2: Model actions

[6] In Julia, the : prefix denotes a symbol

[7] This could be thought of as analogous to an incubation period in viruses targeting living organisms

monetary reward structure for the model is depicted in Table 6.3. In addition to monetary rewards, the reward function can also update the utility function, U_{defender} based on the action taken and the resultant state, s'. The set of observations, $\Omega \in S$, are equivalent to the model states. These are accompanied by an observation accuracy parameter, $p_{obs} = [0 \rightarrow 1]$. A key assumption is that there is ambiguity only as to whether a machine is infected with the attacking malware. This means that for $o \in$: clean, : locked, : offline $p_{obs} = 1$ (i.e. the observer sees the current state) but for $s =$: infected the observer receives observation :infected with probability p_{obs} or :clean with probability $1 - p_{obs}$. The intuition behind this is that some strains of ransomware may initially be stealthy and therefore hard to observe before the ransomware starts to encrypt files. p_{obs} could equivalently be interpreted as the level of competence of malware detection defences.

Parameter	Description
$r_{observe}$	Cost of observation
r^{+}_{repair}	Cost of successful repair
r^{-}_{repair}	Cost of unsuccessful repair
$r_{shutdown}$	Cost of shutting down system
r_{ransom}	Cost of ransom payment

Table 6.3: Model reward structure

Figure 6.1 depicts the transition probability structure of the model for a single machine. It should be noted that the actions $\{\text{Shutdown}, \text{Pay Ransom}\}$ are deterministic whereas other actions cause the Julia program to return a probability distribution of potential states, which can then be sampled. While the transition structure represents a simplification of the progress of an attack, the aim of the model is to capture the broad dynamics of an attack rather than to model each individual stage intricately.

An important feature of the transition model structure is the two-stage process of ransomware infecting and then encrypting a machine. The justification for this is that in sophisticated ransomware attacks, the attackers may spend time implementing command and control infrastructure and attempting to gain privileges before attempting to launch the ransomware and making demands. It is also not a given that ransomware will prove effective at encrypting data on a given machine. It

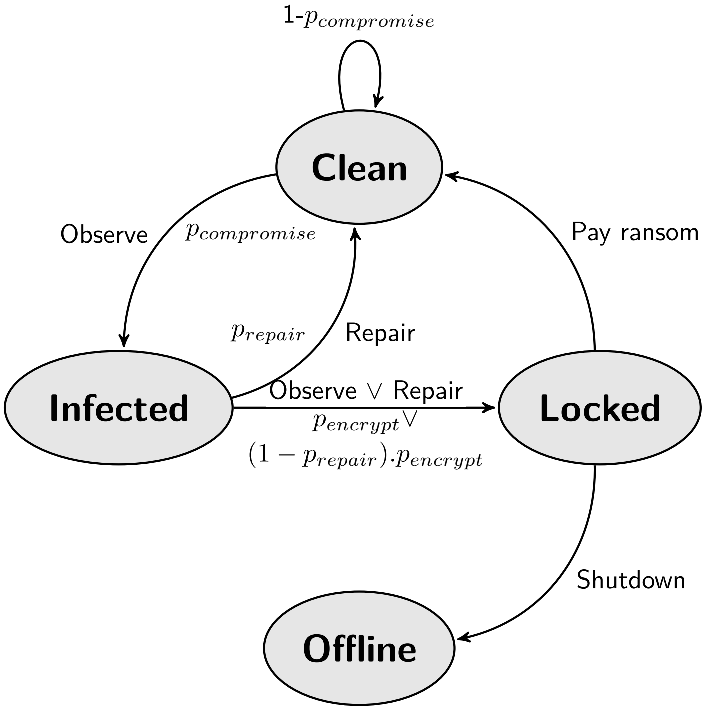

Figure 6.1: Stylised transition probability structure for a single machine

may be possible for an attacker to launch malicious code on a machine or network, but there is a risk that the code fails to execute as expected due to insufficient privileges, active defences or a combination of measures. This failed attempt would likely be spotted by monitoring personnel, who would then trigger the 'repair' action and attempt to remove the attackers and/or malware from the system.

Expanding the model to a network of machines

The ability to define custom types in *Julia* allows for a ready extension of the single-machine model to a network of machines via a pseudo-objected oriented approach. Each machine within a network is represented by the constructor *NetworkedMachine* with fields:

name, category $\ni \{Workstation, Fileserver, DomainController\}$, boolean initial_vector, boolean isCritical and importance $\ni \{: low, : medium, : high\}$.

The properties of each different type of machine in this specific model are outlined in Table 6.4[8]: A function, *ModelNetwork*, takes as arguments the number of each type of machine, and then constructs a vector of *NetworkedMachines*. Each machine

[8] These could of course be altered as needed for modelling of a specific use case

Machine Category	Name	Initial Vector	Critical	Importance
:workstation	ws	true	false	:low
:fileserver	fs	false	false	:medium
:domaincontroller	dc	false	true	:high

Table 6.4: Machine Specification

is automatically assigned a name corresponding to the abbreviations in Table 6.4 and an integer number. By setting up the problem in this fashion, the POMDP simulations can interact with the network in a manner that is realistic. The inclusion of the initial vector property is particularly important as this allows for fine control of infection modelling with respect to the network topology and privilege structure. For example, the typical initial vectors for ransomware infections are either spear phishing of malware or credential theft for remote access [238]. In a well-managed network, file servers and domain controllers should not arguably be readily internet facing. This distinction allows the POMDP model to simulate the lateral movement phase of a ransomware attack, which is important for realism.

The states and observations are contained within N-length vectors, $\mathbf{s}$ and $\mathbf{o}$, where N represents the total number of machines in the network. The vectors are ordered in strictly ascending numerical order $dc \rightarrow fs \rightarrow ws$[9]. The set of actions is applied to the system as a whole. There is an argument for having the repair action target individual machines, but it is assumed that a repair action could be scripted and deployed rapidly across the whole network to affected machines (via Powershell or other administrative tools). It is possible that an attacker could attempt to disable this type of administrative control. This is a motivation behind including the probability of repairing an infection failing within the model. However, in an enterprise network, the early stage privileges granted will likely be limited solely to those of the user of that machine who ordinarily should not have such privileges. The available actions are similar to that of the single-machine model: if any machine in the network is observed infected, the repair action becomes available. If a machine is locked, then the shutdown or pay ransom actions become available. In the special case where a domain controller becomes locked, rendering the network unusable,

[9] The vector of states for a single domain controller, single fileserver, three workstation network would thus be $[\mathrm{dc1}, \mathrm{fs1}, \mathrm{ws1}, \mathrm{ws2}, \mathrm{ws3}]$

the defender faces an ultimatum of either paying the ransom or shutting down the network.

Within the transitions, it is assumed that once a low importance machine is infected, then the attackers move to infect machines within the network. Separate probabilities are included for low, medium and high importance machines (Table 6.5) to allow for different ransomware strategies to be considered. These probabilities are assumed initially to be independent, but this assumption could be relatively easily refined if required for a particular case of interest. There is an argument for considering a network infection model rather than using simple probabilities. However, this would be most justified for a case in which the aim of the attacking malware is to indiscriminately infect as many possible machines and the dynamics of a ransomware attack may be more nuanced. In terms of simulations, the transition probabilities could be parameterised based on the number of infections, but this would add significant complexity to the model.

Probability	
$p_{compromise}$	Probability ransomware initially infects low importance machines
p_{spread_low}	Probability ransomware spreads to other low importance machines
p_{spread_medium}	Probability the ransomware spreads across the network to a medium-importance machine
p_{spread_high}	Probability the ransomware spreads across the network to a high-importance machine
p_{repair}	Probability network cleansed of ransomware before it is locked/files are encrypted
$p_{encrypt}$	Probability that once a machine is infected with ransomware, it becomes locked
p_{obs}	Probability of observations being correct

Table 6.5: Network transition probabilities

6.4.3 Pricing ransomware insurance

As discussed in Section 2.3.6, insurance carriers collect summary data regarding the networks of those looking to purchase cyber-insurance. Realistically, an individual underwriter is likely to have a time constraint in terms of fully evaluating this data. This is especially the case for relatively small policy limits or small/medium enterprise (SME) firms, where a firm may have written thousands of policies or the potential premium intake is modest. The POMDP model presented in this research allows for a representation of a network based on summary data about the number of the machines and is complementary to an underwriting strategy based on mapping specific firm characteristics to past claims. It may also help cyber-insurance firms evaluate policy restrictions - what a firm must do for a claim on an insurance policy to be valid.

Some rudimentary mathematical details of a simple insurance pricing model follow. An insurer writes a policy, $P(p, t, C(\epsilon))$ where p is the premium rate, t is the period of coverage (usually a year), C is the amount of coverage (in monetary units) and ϵ represents the terms of coverage (exclusions, details, sublimits etc.). A policy holder may make claims on losses, l experienced during that year. The insurer will determine whether the claim is valid or not; the policyholder may contest the findings at which point the matter enters the legal rather than purely economic domain. The aim of the insurance company is to ensure that $\sum p_i C_i > \sum l_i$. A rational policyholder will only buy the policy if $pC \leq \sum E[l_i \lambda_i]$ where λ_i is the expected probability of that loss occurring.

The insurance company and buyer compete on information with respect to the decision. The insurance company will have knowledge of the market and risks but the insured may have greater understanding of its own risks. The time dynamic of losses is particularly important for cyber-insurance. In a data breach, costs may be claimed for multiple years after the event, which is problematic for the insurance company who may have by that stage considered the premium intake from the year in question as profit.

In respect of ransomware, for the purposes of this research

$$\epsilon \ni \text{BusinessInterruption}, \text{RansomCosts}, \text{BreachInvestigativeCosts}$$

Within the model, each of these heads of cover has its own separate sub-limit, which will be agreed by the carrier and insured. The POMDP model actions can be mapped to insurance claim states:

$$: \text{shutdown} \rightarrow \text{BusinessInterruption}$$

$$: \text{pay} \rightarrow \text{RansomCosts}$$

$$: \text{repair} \rightarrow \text{BreachInvestigativeCosts}$$

One can then run simulations of the POMDP with different reward (cost) values and probabilities with different confidence weightings to aim to derive the optimal premium.

6.5 Simulations

An initial sensitivity analysis is presented varying different parameters within the model. Two simulations are then introduced: a simple stepwise simulation of the POMDP using three different policies to familiarise the reader with the model structure, and a simulation demonstrating the insurance pricing strategy described in Section 6.4.3. Within these simulations, it is assumed that payment of the ransom restores the system to its original uncompromised state with no risk of reinfection. In reality, this outcome is not guaranteed.

6.5.1 Sensitivity analysis

Varying transition probabilities and network size

The simplest sensitivity analysis is to vary each of the different probabilities within the transaction structure separately, while holding the others constant at $p = 0.5$. The size of the network is initially set at 10 machines, comprising 1 domain controller, 1 fileserver and 8 workstations. This is arbitrary, but seems a reasonable starting point. Separate POMDPs are constructed in *Julia* varying $p = 0.1 \rightarrow 0.9$ in 0.1 step intervals for each probability depicted in Table 6.6. The output variable is average number of simulation steps taken until all domain controllers in the network are locked, which effectively represents the problem absolute terminal state. The simulations were run 10,000 times; this value was chosen as it yielded a good balance of convergence and relatively modest computation time ($< 10s$). Unsurprisingly, only varying the probability that the infection spreads to a high importance machine or the probability that once infected a machine is locked have significant bearing on the number of steps for which the simulation runs before reaching a terminal state. This simply verifies that the transition probability structure is operating as designed.

Next, the effect of the size of the network on the number of steps before all high

p	$p_{compromise}$	p_{spread_low}	p_{spread_medium}	p_{spread_high}	$p_{encrypt}$
0.1	6	5	5	13	13
0.2	5	5	5	8	8
0.3	5	5	5	6	6
0.4	5	5	5	5	6
0.5	5	5	5	5	5
0.6	5	5	5	4	5
0.7	5	5	5	4	4
0.8	5	5	5	4	4
0.9	5	5	5	4	4

Table 6.6: Sensitivity analysis: number of simulation steps until terminal state reached varying single probability variable, holding others constant at 0.5

importance machines are locked are investigated. Table 6.7 shows the results of this simulation, again with 10,0000 runs. In this simulation, the probabilities are all fixed at a specific value. The transition probabilities determine the ultimate speed with which ransomware can lock a network, so one would expect the number of steps before a terminal state is reached to be inversely proportional to the probabilities. Increasing the network size modestly increases the number of average steps, which a simulation runs.

	#(high, medium, low) importance machines				
p	(1,1,8)	(2,2,16)	(3,3,24)	(4,4,32)	(5,5,40)
0.1	22	28	33	36	38
0.2	11	14	16	18	19
0.3	8	10	11	12	13
0.4	6	7	8	9	10
0.5	5	6	7	7	8
0.6	4	5	6	6	6
0.7	4	4	5	5	5
0.8	4	4	4	4	4
0.9	3	3	3	3	4

Table 6.7: Sensitivity analysis: number of simulation steps until terminal state reached varying all probabilities to p with different network sizes.

Finally, the effect on varying the number of domain controllers (i.e. high importance machines) in the network is tested (Table 6.8). As in the prior analysis, all transition probabilities are set at value p and the simulations are run 10,000 times. It is found that increasing the number of domain controllers in the network generally increases the amount of steps before a terminal state is reached except for at extremely high probabilities of infection/spread/encryption. This suggests that for a network with reasonable defences, there is economic benefit to having multiple domain controllers, perhaps in a failsafe-type configuration.

	# Domain Controllers				
p	1	2	3	4	5
0.1	21	28	33	36	38
0.2	11	14	16	18	19
0.3	8	10	11	12	13
0.4	6	7	8	9	9
0.5	5	6	7	7	8
0.6	4	5	6	6	6
0.7	4	4	5	5	5
0.8	4	4	4	4	5
0.9	3	3	3	4	4

Table 6.8: Sensitivity analysis: number of simulation steps until terminal state reached varying number of domain controllers in network, 5 fileservers, 100 workstations.

Effect of transition probability variation on utility

The next simulation run is to check how the utility parameters evolve as an infection spreads without intervention (i.e. the POMDP action is held at :observe). A network of 5 domain controllers, 5 fileservers and 40 workstations is used. This is an arbitrary choice but using a large network allows for variation to be more readily observed as demonstrated by the results in Table 6.7

The utility components are defined as follows:

- C: 1 - (%medium and high importance machines infected or locked)
- I: 1 - (%machines infected)
- A: 1 - (%machines locked)

First $p_{encrypt}$ is varied, holding all other probabilities constant at 0.5 (Figure 6.2). As expected, because availability is solely a function of encryption, there is divergence only in this parameter. This provides a useful test that the simulations are running as expected. Next, all probabilities are held constant except for $p_{compromise}$ and p_{spread_low}, which should have an effect particularly on integrity and availability (in this simulation, there is a 50% chance that an infected machine become locked in the following step). As shown in Figure 6.3, there is no variation in confidentiality but both integrity and availability decrease rapidly as a function of $p_{compromise}$ and p_{spread_low}.

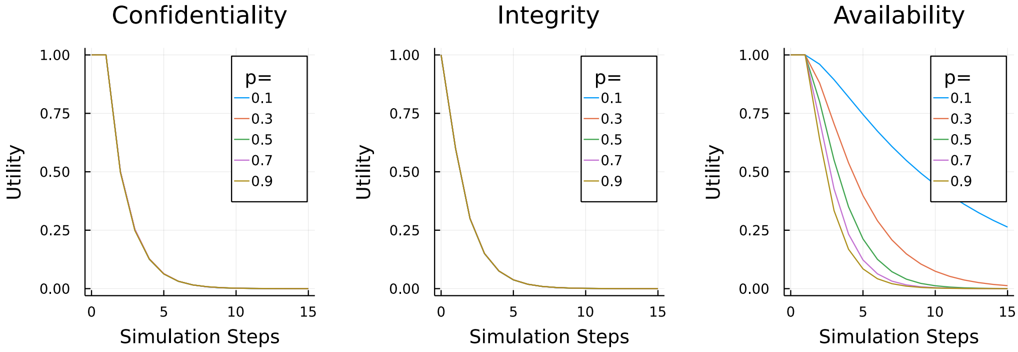

Figure 6.2: Utility versus number of simulation steps. $p_{encrypt} = p$; all other probabilities set to 0.5

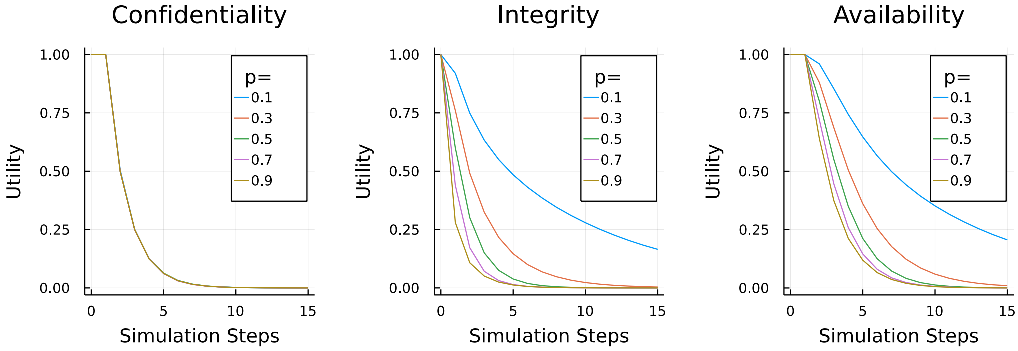

Figure 6.3: Utility versus number of simulation steps. $p_{comrpomise}, p_{spread\_low} = p$; all other probabilities set to 0.5

Finally, p_{spread_medium} and p_{spread_high} are varied (Figure 6.4). This has the largest impact on confidentiality as expected given its definition, but also some impact on integrity and availability though to a lesser extent given that there are 40 low importance machines in the sample network but only 5 medium and 5 high importance ones.

These results are designed to illustrate how simple utility metrics can be used to gain a picture of the evolution of a moderately complex and uncertain simulation and the variation of key parameters. Such plots could be used to simulate the impact of complex technical defences and present the results to non-technical key decision-makers. The subsequent simulations in this chapter will demonstrate some

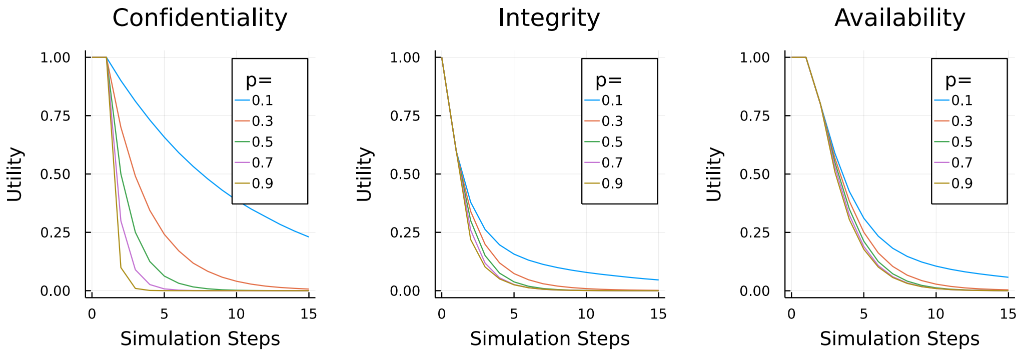

Figure 6.4: Utility versus number of simulation steps. $p_{spread\_medium}, p_{spread\_high} = p$; all other probabilities set to 0.5

applied cases of the effect of different defence strategies against various strains of ransomware with divergent characteristics.

6.5.2 Ransomware infection

Specification

This simulation considers defence against three different ransomware strain attack scenarios:

- Rare/Sophisticated
- Common/Unsophisticated
- 50-50 Baseline

Table 6.9 provides a full specification for each scenario. The first of these is designed to replicate a targeted strain of highly effective ransomware, which is not commonly observed in the wild but once inside a system proves very effective at facilitating lateral movement and ultimately conferral of domain administrator privileges. The probability of initial infection is set relatively low at 0.1, but if a low importance machine is compromised then it spreads quickly to other machines. The probability of successfully infecting medium and high importance machines are set lower at 0.6 and 0.5 respectively to reflect the fact that these servers may be actively monitored and likely have some more sophisticated defences and/or policies aimed to prevent them being susceptible to malicious activity. $p_{encrypt} = 0.7$ for this strain. The second strain studied is commonly observed, self-propagating ransomware such as

WannaCry, which is readily eliminated by appropriate tools. Here, the probability of initial infection is set at a very high 0.9, but the probability of repair is set at 0.8; thus there is a decent, but not certain, chance that this strain might be cleared from any workstation it infects. It is assumed that its locking/encrypting methodology is not that sophisticated, expressed by $p_{encrypt} = 0.3$. Finally, as the name suggests, the 50-50 baseline scenario sets all probabilities in the model to 0.5. Observation accuracy is set at 70% initially to create the possibility of inaccurate observations and consequent policy errors. For completeness, a simple discount factor of 0.95 is set, though this is not required for simulations but would be used if applying a solver to the system.

	Rare/Sophisticated	Common/Unsophisticated	50-50
Rewards			
$r_{observe}$	-1	-1	-1
r^{+}_{repair}	-2	-2	-2
r^{-}_{repair}	-10	-10	-10
$r_{shutdown}$	-150	-150	-150
r_{ransom}	-50	-50	-50
Probabilities			
$p_{infection}$	0.1	0.9	0.5
p_{spread_low}	0.8	0.3	0.5
p_{spread_medium}	0.6	0.1	0.5
p_{spread_high}	0.5	0.1	0.5
p_{repair}	0.2	0.8	0.5
$p_{encrypt}$	0.7	0.3	0.5
Other			
$Obs.Acc.$	0.7	0.7	0.7
$Disc.Fac.$	0.95	0.95	0.95

Table 6.9: Simulation Specifications

The reward parameters selected are intended to be largely illustrative and are arguably the most transparent component of the model. There is a small penalty for observation, which is designed to represent the cost of monitoring a network. Separate rewards are included for successful and failed repairs (r^{+}_{repair} and r^{-}_{repair} respectively). Intuitively, an unsuccessful repair means likely further investigative costs or expense to attempt to remove the ransomware for the network, such as hiring specialist help. The costs of shutting down the network are deliberately set as higher (i.e. more negative in reward terms) than paying the ransom. The aim of this simulation is to investigate how defensive actions affect the resultant outcomes and consequently, the rewards are simply a means of 'keeping score'. While abstract in relation to real world costs, this approach is consistent with conventions within the game theory and decision model literature.

The three strains are tested on a sample network containing 1 domain controller, 2 fileservers and 10 workstations. This network size was chosen to provide a reasonably sized attack surface but to be of a manageable size for debugging purposes. When aiming to solve, or at least simulate a POMDP, it is conventional to evaluate the effect of different policies. A policy in this context is a specification of actions corresponding to a belief (in this model, the belief is simply the observations). Three policies are evaluated: the 'cautious policy', the 'gambler policy' and the 'random policy'.

- **Cautious policy**: attempt repair if infected; shut down if domain controller encrypted/locked; never pay ransom.
- **Gambler policy**: observe until a system becomes encrypted/locked at which point pay ransom.
- **Random policy**: take random action from set of available actions corresponding to received observation.

The simulations are run in step-wise fashion:

1. Initial vector of states s and observations o set fully clean
2. Receive optimal action a from policy p based on o
3. Determine next state s' from transition $t(s, a)$
4. Compute reward $r(s, a, s')$
5. Record s, a, s', r
6. If action is terminal, terminate simulation
7. Set $s = s'$ and compute observations $o(s)$

8. Repeat from (2) until maximum number of steps reached or terminal action taken

Results

For each POMDP and policy, the simulations were run 10,000 times and the history recorded. A maximum of 15 steps was permitted in each individual simulation - per Table 6.7, this is likely to be sufficient to fully capture the simulation steps in most outcomes. As expected, the average reward (Table 6.10) for the cautious policy is much lower than the gambler policy, given that the cautious policy prohibits ransom payment and the cost of shutting down the system is higher than the ransom. This is particularly apparent for the rare but dangerous strain of ransomware. However, for the common but benign strain, as it is far less likely that the key network infrastructure is locked, the difference between average rewards is much smaller. The conclusion of these simulations is that paying the ransom is the best economic strategy.

	Rare/Sophisticated	Common/Unsophisticated	50-50
Cautious Policy	-146	-68	-136
Gambler Policy	-53	-52	-52
Random Policy	-111	-112	-111

Table 6.10: Simulation Results - average wealth

Figures 6.5, 6.6, and 6.7 show the distribution of the number of simulation steps before the simulation terminates across the 10,000 runs. This provides an insight into the variability of the length of the simulation and the impact of the policy chosen. For the rare/sophisticated ransomware strain, the cautious policy shows the greatest variability. This is likely because the probability of repair is low, and thus the chances of the domain controller becoming locked are relatively high, as suggested by the average reward returned being close to the cost of shutting down the system. The gambler policy in contrast results in much shorter run times. For the common/unsophisticated strain, there is almost a deterministic distribution of outcomes, which makes sense given that the probability of a repair is much higher than the probability of the infection spreading.

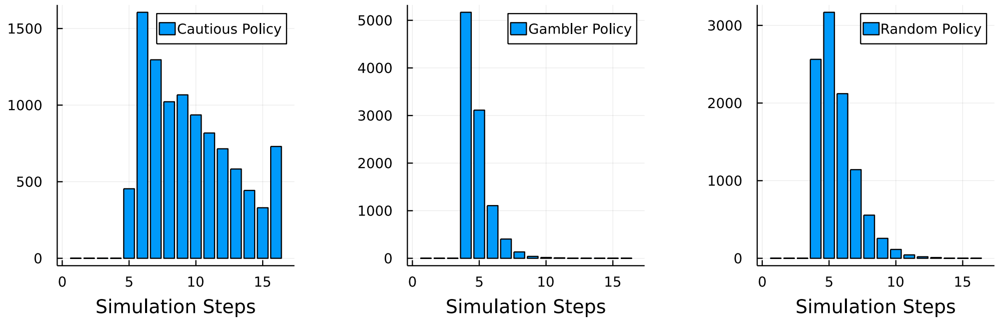

Figure 6.5: Simulation step distribution for rare/sophisticated strain

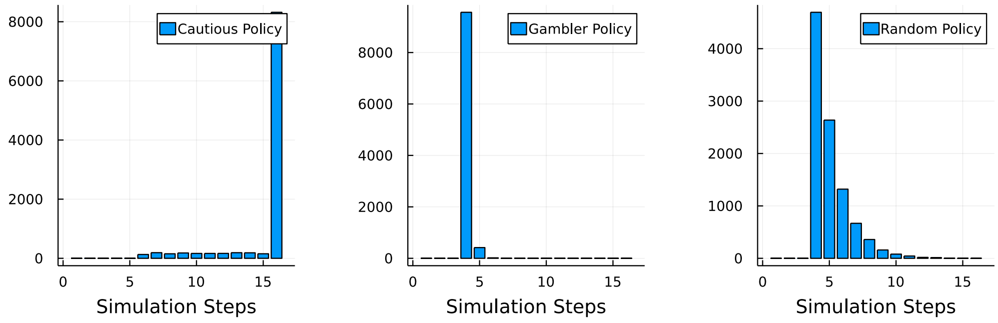

Figure 6.6: Simulation step distribution for common/unsophisticated strain

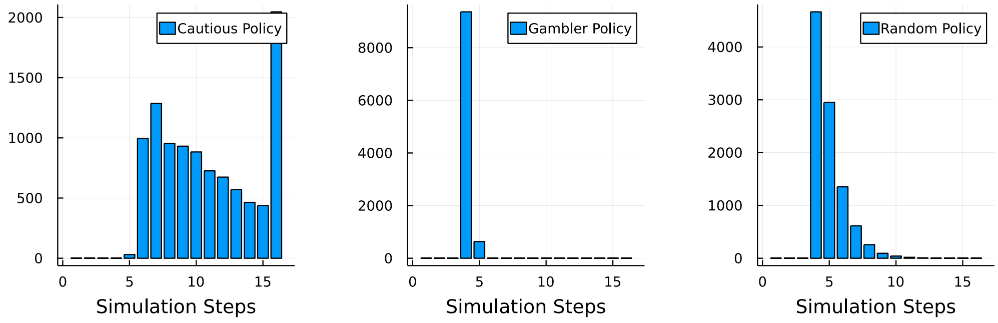

Figure 6.7: Simulation step distribution for 50/50 strain

The utility for each state in the simulations was calculated (as in Section 6.5.1). Figures 6.8, 6.9 and 6.10 show the average components of the utility for each ransomware strain and each policy. The x-axis of each subplot represents the number of simulation steps and the y-axis the numerical utility, ranging from 0 to 1. For each of the 10,000 runs, the number of simulation steps taken was recorded and transformed into a vector so that the average is correctly calculated. For the rare/dangerous strain, the gambler policy maximises utility whereas the cautious policy drastically underperforms the benchmark random policy. For the common/unsophisticated strain, however, the cautious policy performs notably better, with the domain controller locked in only 25% of simulations and on average fewer than 50% of network machines either infected with ransomware or locked. The 50/50 strain is intended as a control; the gambler policy has notably less potential for randomness in the outcomes whereas in the cautious policy, the repair action proves ineffective at stemming the spread and progress of the ransomware.

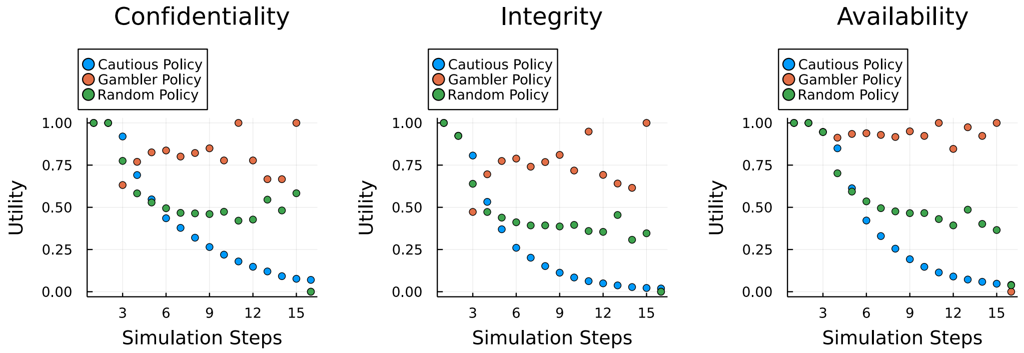

Figure 6.8: Utility evolution for rare/sophisticated strain

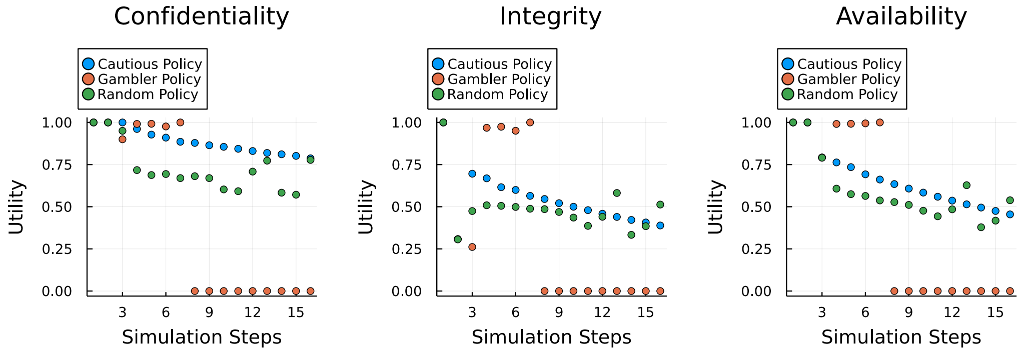

Figure 6.9: Utility evolution for common/unsophisticated strain

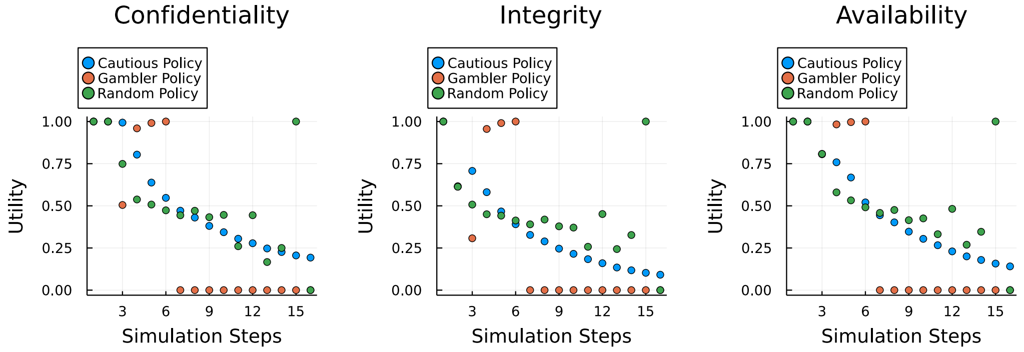

Figure 6.10: Utility evolution for 50/50 strain

It is relatively straightforward to assess the effect of varying the parameters within the different POMDPs on the average rewards received for the different policies. Figure 6.11 shows the average reward for the cautious/unsophisticated strain POMDP varying the probability that an infection spreads to the domain controller once in the system (p_{spread_high}). This illustrates that the crossover point between the gambler policy being the optimal strategy that the cautious policy occurs at a fairly low probability of overall domain controller locking. This is largely because the gambler policy immediately pays the ransom thus preventing the infection from spreading to the domain controller.

A useful experiment is to investigate the effects of varying the probability of a repair being successful on the reward (Figure 6.12) for the common/unsophisticated ransomware strain where this probability should have greatest impact. This provides a useful check as to the robustness of the policies as only the cautious policy reward should vary with p_{repair}. While in a real world decision, the evaluation of defences would not be undertaken purely on the basis of probabilities, this nevertheless illustrates the sort of cost-benefit analysis that might be undertaken when planning security investments. For a real world problem, a POMDP could be constructed and solved for the optimal policy to help determine an incident response plan, for example.

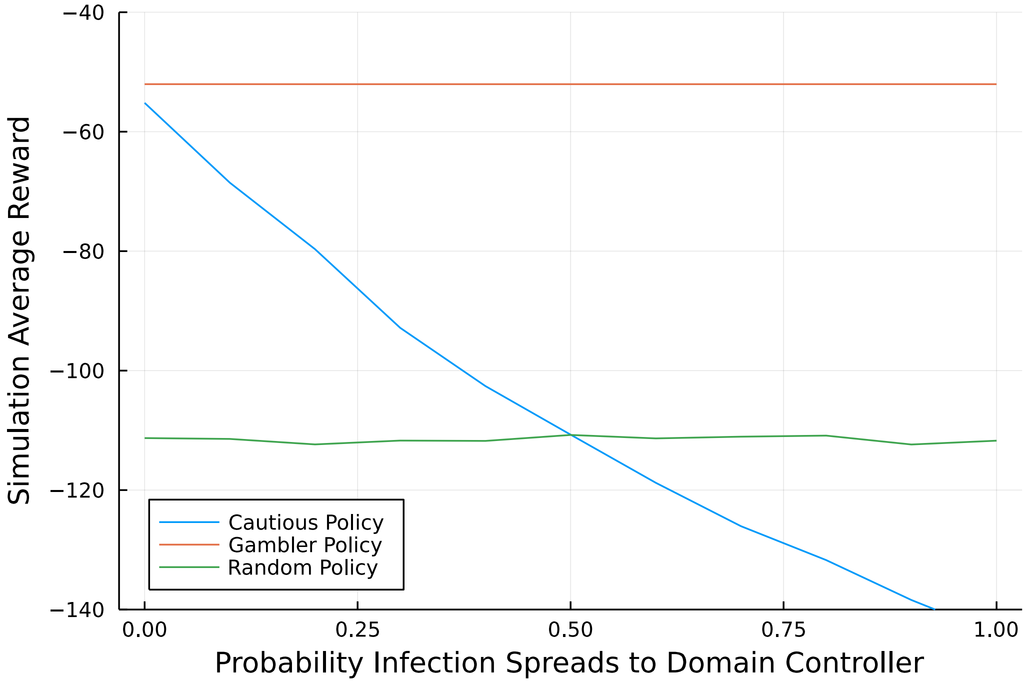

Figure 6.11: Varying probabilities of domain controller compromise, Common/Unsophisticated Strain

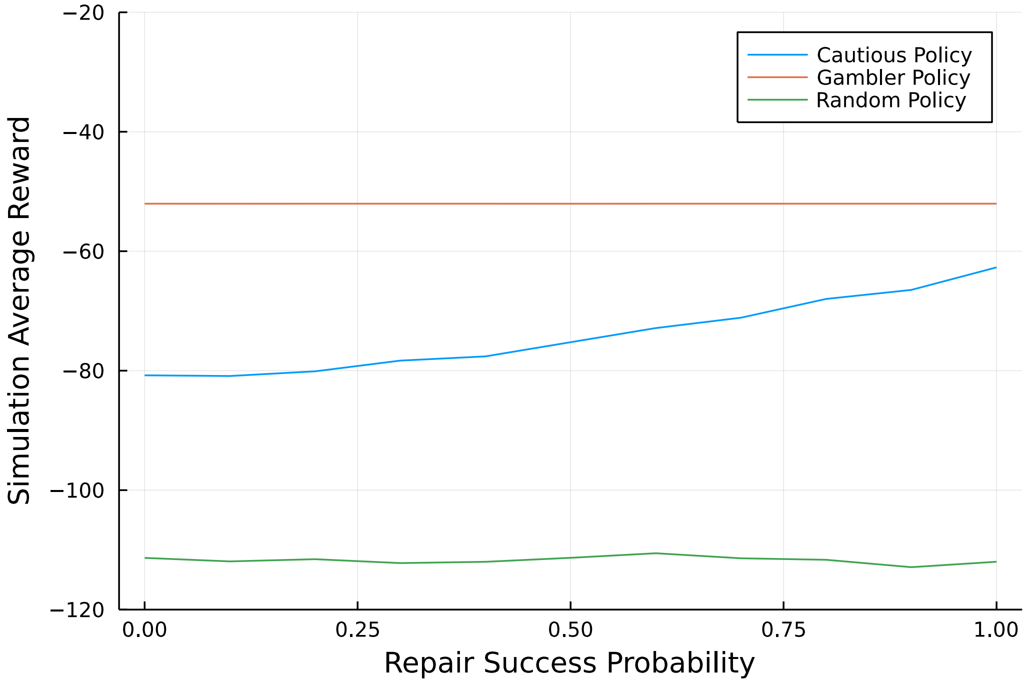

Figure 6.12: Varying probabilities of successful repair, Common/Unsophisticated Strain

6.5.3 Insurance pricing

This section considers a simple insurance pricing example against ransomware based on the POMDP model introduced. An insurer is considering pricing ransomware insurance for an organization. The organization has a network of 1,000 workstations, 10 fileservers and 5 domain controllers. The policy has the following features:

- Ransom paid up to $1mn only if all domain controllers locked.
- Repair costs paid. A successful repair costs $10,000 and an unsuccessful one $50,000 each time.
- Business interruption costs of $10mn in the event of more than 50% machines in network locked by ransomware.

For this example, it is assumed that the insurer is aware of six different strains of ransomware with differing characteristics. For consistency with the simulations discussed in Section 6.5.2, three variants of the rare/sophisticated and common/unsophisticated strains are considered. The parameters for POMDPs representing these different strains are given in Table 6.11.

	R/S(i)	R/S(ii)	R/S(iii)	C/U(i)	C/U(ii)	C/U(iii)
Rewards						
$r_{observe}$	0					
r^{+}_{repair}	-10,000					
r^{-}_{repair}	-50,000					
$r_{shutdown}$	-10,000,000					
r_{ransom}	-1,000,000					
Probabilities						
$p_{infection}$	0.1	0.1	0.1	0.9	0.9	0.9
p_{spread_low}	0.8	0.7	0.6	0.3	0.2	0.1
p_{spread_medium}	0.8	0.7	0.6	0.3	0.2	0.1
p_{spread_high}	0.8	0.7	0.6	0.3	0.2	0.1
p_{repair}	0.1	0.2	0.3	0.7	0.8	0.9
$p_{encrypt}$	0.8	0.7	0.6	0.3	0.2	0.1
Other						
$Obs.Acc.$	0.7					
$Disc.Fac.$	0.95					

Table 6.11: POMDP specifications for insurer-specified ransomware strains

Suppose that an insurance company has appetite for writing 1000 policies insuring against ransomware attacks. Then, assume that the insurance company has an expected distribution of claims resulting from each strain, the arrival of which is Poisson distributed with λ parameters given in Table 6.12. The Poisson distribution gives the frequency of expected claims and the POMDP simulation the severity of the claim. This approach is similar to that used in the simulations developed in Chapter 5.

Strain	R/S(i)	R/S(ii)	R/S(iii)	C/U(i)	C/U(ii)	C/U(iii)
λ	10	15	20	25	30	35

Table 6.12: Poisson parameters for ransomware strains

A simulation may then be run according to the following protocol:

1. For each ransomware strain, randomly sample the appropriate Poisson distribution to obtain number of ransomware attacks, n to simulate.

2. Simulate n attacks per the protocol outlined in Section 6.5.2 and record rewards.

3. Repeat 1000 times (equal to the maximum number of policies the insurer is willing to write), and take average.

The results of this simulation are presented in Table 6.13. The total expected losses to the portfolio are \$722.2mn, or \$722,200 per contract. Accordingly this would be the actuarially fair premium the insurance company would quote for insuring per the policy terms and conditions. This poses an interesting decision for the insurance buyer; the cost of insurance is only half that of a ransom demanded under stringent conditions but is only 0.7% of the total possible business interruption cover available under the policy.

6.6 Further work

The simulation demonstrates a proof-of-concept of a POMDP approach to modelling ransomware. Ideally, the next steps in the work would be to use the framework to evaluate decision making in specific scenarios and systems architecture and feedback

Ransomware Strain	Average Loss per attack ($mn)	Total Simulated Loss ($mn)
R/S(i)	9.2	92.0
R/S(ii)	10.0	146.8
R/S(iii)	10.0	201.0
C/U(i)	9.4	235.0
C/U(ii)	0.8	22.8
C/U(iii)	0.7	24.5
Sum		722.2

Table 6.13: Results of insurance claim simulation

is welcomed as how this might be most usefully achieved. The representation of the network was constructed with the aim of replicating sample networks such as an Active Directory network and the demonstration of its usage in this work is fairly simple. A potentially interesting expansion of the simple approach would be to introduce labelled transitions to formally describe the privilege structure between machines. This might allow for incorporating user accounts and privilege structures within the communications and may be of use in threat modelling to describe various different potential attack vectors from unintended use of privileges (for example from compromise of service account credentials). It should be noted that the construction of the model potentially allows for separate model networks to be constructed, representing an Active Directory Forest, for example.

There is the potential to introduce significant complexity into models such as the one presented in this work. For simplicity, it is assumed that a ransom payment results in full decryption and restoration of the system to its original state. This may not be the case in reality and there would be potential scope to incorporate this into future model simulations. Equally, it is assumed that once a machine is infected, if not repaired, it is encrypted or locked with fixed probability. If an attacker is able to gain introduce command and control (C2C) functionality, then this might not be the case.

The model presented within this research focused on a single POMDP and assumes no costs to the criminal actor. An interesting expansion of the model may be to simulate such costs on the criminal actor (for example, resource constraints, risk of discovery within a network etc.). The criminal actor might also be simulated as a reinforcement learning (RL) agent; one could also potentially introduce a defender

(RL) agent as well.

6.7 Summary

This chapter has introduced a POMDP model for simulating ransomware attacks either on a single machine or a network of machines of varying importance. The results of a simple simulation of different types of ransomware attack highlight that economically the least costly financial outcome is usually to pay the ransom at the first chance, although this is a scenario that is unlikely to be encouraged by governmental authorities or insurers. It is hoped that this model may be useful for helping frame simulations of complex attacks and in developing optimal defence strategies. The applicability of such model results to a simple insurance pricing example has also been demonstrated, highlighting how cover could be adapted based on risk perception.

Conclusion 7

7.1 Closing remarks

This thesis deploys economics to tackle several different problems in modelling information security, resulting in three separate models and one modelling framework. Chapter 3 expands the well-established Gordon-Loeb model to include cyber-insurance as part of the security investment decision problem. This introduces an important dimension to the problem: the decision-maker can invest to reduce the probability of a breach and/or insure the risk. This is arguably a more complete model for a corporate decision-maker compared with the standalone Gordon-Loeb Model. Chapter 4 provides a modelling framework that covers all parts of the insurance assessment procedure for a cyber-insurance policy: describing the company seeking insurance, assessing its security posture, and assisting guide the policy pricing. This framework is arguably a much richer vehicle for facilitating discussions between an insurance buyer and seller than the simple, standard questionnaire that is commonplace in the insurance industry. It is hoped that the framework makes a contribution towards helping companies better describe their own security posture should they lack the resources or expertise so to do.

In the language of Economists, Chapter 4 introduces a ‘bottom-up’ or microscopic analysis of the cyber-insurance problem, focusing on the specific detail of individual components of an ecosystem and describing how these might be aggregated and scaled. Chapter 5 in contrast takes a ‘top down’ view of cyber-insurance, considering the necessary conditions for a sustainable market at a macroscopic level. Even on a cursory analysis of economic efficiency arguments, it is clear that at present it is almost impossible for the cyber-insurance market to be efficient. However, it is not beyond reason that one day it might be. The key question to address is whether technological growth might reach or even at least trend to a terminal state. Additionally, the threat landscape may evolve to a point where greater international coordination and enforcement diminishes the frequency of cyber-incidents. At the time of writing, this appears to be wishful thinking. Should the *status quo* persist, it appears that better coordination and data sharing is the first step towards smoothing the frictions in the cyber-insurance market. Alongside this, with reinsurance capacity heavily used already, cyber-insurers are likely to need to find other sources of external, risk-tolerant capital to grow underwriting capacity.

Chapter 6 considers how to model ransomware attacks using a model based on Markov Decision Processes. While highly abstract as introduced in the thesis, the model could very easily be adapted for use in calculating the effect of different ransomware strains on real-world organizations. The challenge is to find an organization that is willing to share such details and spend the time working on producing models. Initially, it was hoped that this thesis would be far more empirical in nature than has transpired. There are two reasons for this: first, publicly available useful quantitative data on cyber-incidents is hugely lacking and most studies to date have relied upon proprietary or external information. Second, companies remain apparently reluctant to share data on cyber-security incidents. There are tentative signs that regulatory pressure may change this, for example, the Securities and Exchange Commission in the United States has introduced rules on mandatory disclosure of cyber-incidents for public companies [284]. This may allow the academic community to start to garner and curate information appropriately and thereby advance the empirical security economics field.

The ultimate conclusion from the work presented in the thesis is: modelling problems in information security is indeed hard. It often requires novel thinking and there is occasionally relatively little in terms of literature to guide endeavours, especially when approaching a problem from a multi-disciplinary angle. Approaches will often be criticised as too abstract by practitioners or as too generic by those in the research community more used to granular modelling as opposed to the language of economics. Yet, it appears that using economics for information security problems is worthwhile. In particular, the insurance industry is becoming increasingly keen to refine its approach to exposure management for cyber-risk and in modelling catastrophe risk. Simply being able to ascribe a number to the worst possible outcomes of cyber-events is a starting point in creating boundaries to the problem. These might then spur discussions that clarify thinking on the modelling of cyber-risk and deliver better outcomes and resilience for society at large, which is now heavily reliant on technology and has delegated great responsibility to those who control it in terms of security and privacy.

7.2 Further work

The ERD-SMM-utility modelling framework introduced in Chapter 4 has significant potential to be applied to real-world case studies. A particularly interesting exercise would be to compare an assessment delivered via the framework to insurance questionnaires and examine whether it might price an insurance policy differently. The simulations in Chapter 5 could be expanded to include dynamic behavioural adjustments based on the information endowed to the various agents. This is a complex piece of work, but one that might yield fruitful insights on the behaviour of participants in the cyber-insurance market.

Finally, the ransomware model in Chapter 6 might be combined with the descriptions of system structure in Chapter 4 to provide a richer model than the simple network structure used for ransomware modelling in the thesis. All these outputs might then be combined to populate the Gordon-Loeb Model expanded with cyber-insurance introduced in Chapter 3 to produce an organizational decision model for specific security investments versus cyber-insurance in contrast to the relatively abstract setup used to introduce the model.

All of these suggestions would require significant collaboration with an insurance carrier and/or broker and the data required to deliver these is, in many cases, proprietary and confidential. In time, better public databases and incident reporting might help spur more empirical research endeavours along the lines discussed.

Bibliography

[1] Knut K Aase. “The Nash bargaining solution vs. equilibrium in a reinsurance syndicate”. In: *Scandinavian Actuarial Journal* 2009.3 (2009), pp. 219–238.

[2] Ali E Abbas. *Foundations of multiattribute utility.* Cambridge University Press, 2018.

[3] Manuel Adam, Maren Josefs, Simon Ashworth, Johannes Bender, and Taoufik Gharib. *Cyber Risks In A New Era: Reinsurers Could Unlock The Cyber Insurance Market.* Sept. 2021. URL: `https://www.spglobal.com/ratings/en/research/articles/210929-cyber-risks-in-a-new-era-reinsurers-could-unlock-the-cyber-insurance-market-12118547`.

[4] AIR worldwide. *AIR Estimates Losses for the Marriott Breach Will Be Between USD 200 Million and USD 600 Million.* Dec. 18, 2018. URL: `https://www.air-worldwide.com/news-and-events/press-releases/AIR-Estimates-Losses-for-the-Marriott-Breach-Will-Be-Between-USD-200-Million-and-USD-600-Million/`.

[5] G.A. Akerlof. “The market for lemons: Quality uncertainty and the market mechanism”. In: *Quarterly Journal of Economics* 84.3 (1970), pp. 488–500. ISSN: 00335533.

[6] Armen A Alchian. “The basis of some recent advances in the theory of management of the firm”. In: *The Journal of Industrial Economics* (1965), pp. 30–41.

[7] M Allais. "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine". In: *Econometrica* 21.4 (1953), pp. 503–546. URL: http://www.jstor.org/stable/1907921.

[8] Liz Allen, Alison OConnell, and Veronique Kiermer. "How can we ensure visibility and diversity in research contributions? How the Contributor Role Taxonomy (CRediT) is helping the shift from authorship to contributorship". In: *Learned Publishing* 32.1 (2019), pp. 71–74. DOI: https://doi.org/10.1002/leap.1210. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/leap.1210. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/leap.1210.

[9] James P. Anderson. *Computer security technology planning study.* Tech. rep. Electronic Systems Division, National Technical Information Service, US Department of Commerce, 1972.

[10] R Anderson. "Why information security is hard - an economic perspective". In: *Seventeenth Annual Computer Security Applications Conference.* Vol. 2001-. IEEE, 2001, pp. 358–365. ISBN: 0769514057.

[11] Ross Anderson. *Security engineering: a guide to building dependable distributed systems.* John Wiley & Sons, 2020.

[12] Erik Angner and George Loewenstein. "Behavioral economics". In: *Handbook of the philosophy of science: Philosophy of economic* (2007), pp. 641–690.

[13] Aon PLC. *U.S. Cyber Market Update: 2022 U.S. Cyber Insurance Profits and Performance.* 2022. URL: https://www.aon.com/getmedia/438dfae5-3004-4f60-9698-d85fb6770868/20230920-2022-us-cyber-market-update.pdf.

[14] Kenneth Arrow. *Aspects of the theory of risk-bearing.* 1965.

[15] Kenneth J Arrow. "Optimal insurance and generalized deductibles". In: *Scandinavian Actuarial Journal* 1974.1 (1974), pp. 1–42. ISSN: 0346-1238. URL: http://www.tandfonline.com/doi/abs/10.1080/03461238.1974.10408659.

[16] Terrence August, Duy Dao, and Marius Florin Niculescu. "Economics of ransomware: Risk interdependence and large-scale attacks". In: *Management Science* 68.12 (2022), pp. 8979–9002.

[17] AXA Research Fund. *Cyber insurance risks: evaluating the cyber costs of cyber risks.* URL: `https://axa-research.org/funded-projects/socio-economy-new-tech/cyber-insurance-risks-evaluating-the-cyber-costs-of-cyber-risks`.

[18] Baharuddin Aziz, Suhardi, and Kurnia. "A systematic literature review of cyber insurance challenges". In: *2020 International Conference on Information Technology Systems and Innovation (ICITSI)*. IEEE. 2020, pp. 357–363.

[19] Elnaz Bajoori, Tristan Caulfield, and Christos Ioannidis. "Cyber Security Service Providers- Should we leave them alone?" In: *Workshop on Approaches to Modelling Heterogeneous Interacting Systems, Grenada.* In Association with Financial Cryptography. 2022.

[20] Adrian Baldwin, Iffat Gheyas, Christos Ioannidis, David Pym, and Julian Williams. "Contagion in cyber security attacks". In: *Journal of the Operational Research Society* 68.7 (2017), pp. 780–791.

[21] Tridib Bandyopadhyay and Vijay Mookerjee. "A model to analyze the challenge of using cyber insurance". In: *Information Systems Frontiers* 21.2 (2019), pp. 301–325. ISSN: 1387-3326. DOI: `10.1007/s10796-017-9737-3`.

[22] Bank of England. *CBEST Threat Intelligence-Led Assessments.* June 13, 2022. URL: `https://www.bankofengland.co.uk/financial-stability/operational-resilience-of-the-financial-sector/cbest-threat-intelligence-led-assessments-implementation-guide`.

[23] Richard Barker. *CASE Method: entity relationship modelling.* Addison-Wesley Longman Publishing Co., Inc., 1990.

[24] Carlos Barreto, Alvaro A. Cardenas, and Galina Schwartz. *Cyber-Insurance for Cyber-Physical Systems.* 2018 IEEE Conference on Control Technology and Applications. 2018, pp. 1704–1711. ISBN: 978-1-5386-7698-1.

[25] Yuliy Baryshnikov. "IT Security Investment and Gordon-Loeb's 1/e Rule." In: *Workshop on the Economics of Information Security 2012*. 2012.

[26] Sidney Benjamin. "Loadings for insurance premiums". In: *Geneva Papers on Risk and Insurance* (1986), pp. 110–125.

[27] Christopher L Benson and Christopher L Magee. "Quantitative determination of technological improvement from patent data". In: *PloS one* 10.4 (2015), e0121635.

[28] Lawrence A Berger, J David Cummins, and Sharon Tennyson. "Reinsurance and the liability insurance crisis". In: *Journal of risk and Uncertainty* 5 (1992), pp. 253–272.

[29] Yannick Bessy-Roland, Alexandre Boumezoued, and Caroline Hillairet. "Multivariate Hawkes process for cyber insurance". In: *Annals of Actuarial Science* 15.1 (2021), pp. 14–39.

[30] Yannick Bessy-Roland, Alexandre Boumezoued, and Caroline Hillairet. "Multivariate Hawkes process for cyber insurance". In: *Annals of Actuarial Science* 15.1 (2021), pp. 14–39. DOI: 10.1017/S1748499520000093.

[31] Kenneth J Biba. *Integrity considerations for secure computer systems*. Tech. rep. Mitre Corporation MTR-3153, 1975.

[32] Christian Biener, Martin Eling, and Jan Hendrik Wirfs. "Insurability of Cyber Risk: An Empirical Analysis". In: *Geneva Papers on Risk and Insurance-Issues and Practice* 40.1 (2015), pp. 131–158. ISSN: 1018-5895. DOI: 10.1057/gpp.2014.19.

[33] G. Birtwistle. *DEMOS: Discrete Event Modelling on Simula*. Springer, 1979.

[34] Lawrence D. Bodin, Lawrence A. Gordon, Martin P. Loeb, and Aluna Wang. "Cybersecurity insurance and risk-sharing". In: *Journal of Accounting and Public Policy* 37.6 (2018), pp. 527–544. ISSN: 0278-4254. DOI: 10.1016/j.jaccpubpol.2018.10.004.

[35] Rainer Boehme. "Security Metrics and Security Investment Models". In: *Information Security and Privacy*. Information Security and Privacy, 2010, pp. 10–24.

[36] Rainer Böhme and Gaurav Kataria. "On the limits of cyber-insurance". In: *Trust, Privacy, and Security in Digital Business, Proceedings*. Ed. by S. FischerHubner, S. Furnell, and C. Lambrinoudakis. Vol. 4083. Lecture Notes in Computer Science. 2006, pp. 31–40. ISBN: 3-540-37750-6.

[37] Rainer Böhme, Stefan Laube, and Markus Riek. "A fundamental approach to cyber risk analysis". In: *Variance* 12.2 (2019), pp. 161–185.

[38] Rainer Böhme and Galina Schwartz. "Modeling cyber-insurance: towards a unifying framework." In: *Workshop on the Economics of Infomation Security*. 2010.

[39] Rok Bojanc and Borka Jerman-Blazic. "An economic modelling approach to information security risk management". In: *International Journal of Information Management* 28.5 (2008), pp. 413–422. ISSN: 0268-4012. DOI: 10.1016/j.ijinfomgt.2008.02.002.

[40] Karl Borch. "The economic theory of insurance - notes for an informal discussion in Edinburgh 1 June 1964". In: *ASTIN Bulletin* 4.3 (1967), pp. 252–264. ISSN: 05150361.

[41] Karl Borch. "Is regulation and supervision of insurance companies necessary?" In: *Scandinavian Actuarial Journal* 1981.3 (1981), pp. 179–190. ISSN: 0346-1238. URL: http://www.tandfonline.com/doi/abs/10.1080/03461238.1981.10432017.

[42] Alexander Braun. "Pricing in the Primary Market for Cat Bonds: New Empirical Evidence". In: *Journal of Risk and Insurance* 83.4 (2016), pp. 811–847. DOI: https://doi.org/10.1111/jori.12067. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jori.12067. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/jori.12067.

[43] Alexander Braun, Martin Eling, and Christoph Jaenicke. “Cyber insurance-linked securities”. In: *ASTIN Bulletin: The Journal of the IAA* (2023), pp. 1–22.

[44] Alexander Braun and Carolyn Kousky. *Wharton Risk Center Primer: Catastrophe Bonds*. July 2021. URL: `https://riskcenter.wharton.upenn.edu/wp-content/uploads/2021/07/Cat-Bond-Primer-July-2021.pdf`.

[45] Ruth Breu, Ursula Hinkel, Christoph Hofmann, Cornel Klein, Barbara Paech, Bernhard Rumpe, and Veronika Thurner. “Towards a formalization of the unified modeling language”. In: *ECOOP’97Object-Oriented Programming: 11th European Conference Jyväskylä, Finland, June 9–13, 1997 Proceedings 11*. Springer. 1997, pp. 344–366.

[46] Oliver Brew. *The all risk cyber challenge*. 2023. URL: `https://global.lockton.com/re/en/news-insights/lockton-re-cyber-report-says-market-needs-cyber-product-clarity`.

[47] Matthew L Bringer, Christopher A Chelmecki, and Hiroshi Fujinoki. “A survey: Recent advances and future trends in honeypot research”. In: *International Journal of Computer Network and Information Security* 4.10 (2012), p. 63.

[48] Max T Brozynski and Benjamin D Leibowicz. “Markov models of policy support for technology transitions”. In: *European Journal of Operational Research* 286.3 (2020), pp. 1052–1069.

[49] Alan Calder and Steve G Watkins. *Information security risk management for ISO27001/ISO27002*. It Governance Ltd, 2010.

[50] Miguel Calvo and Marta Beltrán. “A Model For risk-Based adaptive security controls”. In: *Computers and Security* 115 (2022), p. 102612. ISSN: 0167-4048. DOI: `https://doi.org/10.1016/j.cose.2022.102612`. URL: `https://www.sciencedirect.com/science/article/pii/S0167404822000116`.

[51] Hasan Cam. “Online detection and control of malware infected assets”. In: *MILCOM 2017 - 2017 IEEE Military Communications Conference (MILCOM)*. 2017, pp. 701–706. DOI: `10.1109/MILCOM.2017.8170869`.

[52] Carnegie Mellon University and The Johns Hopkins University Applied Physics Laboratory LLC. *Cybersecurity Maturity Model Certification (CMMC) Model Overview*. Tech. rep. 2021. URL: `https://dodcio.defense.gov/Portals/0/Documents/CMMC/ModelOverview_V2.0_FINAL2_20211202_508.pdf`.

[53] Robert L Carter. *Reinsurance*. Springer Science & Business Media, 2013.

[54] Edward Cartwright, Julio Hernandez Castro, and Anna Cartwright. "To pay or not: game theoretic models of ransomware". In: *Journal of Cybersecurity* 5.1 (2019).

[55] Anthony R Cassandra. "A survey of POMDP applications". In: *Working notes of AAAI 1998 fall symposium on planning with partially observable Markov decision processes*. Vol. 1724. 1998.

[56] Tristan Caulfield and David Pym. "Improving security policy decisions with models". In: *IEEE Security & Privacy* 13.5 (2015), pp. 34–41.

[57] James L Cebula and Lisa R Young. *A taxonomy of operational cyber security risks*. Tech. rep. Carnegie-Mellon Univ Pittsburgh Pa Software Engineering Inst, 2010.

[58] David Chaum. "Security without Identification: Transaction Systems to Make Big Brother Obsolete". In: *Commun. ACM* 28.10 (Oct. 1985), pp. 1030–1044. ISSN: 0001-0782. DOI: `10.1145/4372.4373`. URL: `https://doi.org/10.1145/4372.4373`.

[59] Peter Pin-Shan Chen. "The entity-relationship model — toward a unified view of data". In: *ACM transactions on database systems (TODS)* 1.1 (1976), pp. 9–36.

[60] Simon Christ, Daniel Schwabeneder, Christopher Rackauckas, Michael Krabbe Borregaard, and Thomas Breloff. "Plots.jl – a user extendable plotting API for the julia programming language". In: (2023). DOI: `https://doi.org/10.5334/jors.431`. URL: `https://openresearchsoftware.metajnl.com/articles/10.5334/jors.431/`.

[61] David D Clark and David R Wilson. “A comparison of commercial and military computer security policies”. In: *1987 IEEE Symposium on Security and Privacy*. IEEE. 1987, pp. 184–184.

[62] David R. Clark. “Basics of Reinsurance Pricing”. In: *CAS Actuarial Study Note* (2014). URL: `%7Bhttps://www.casact.org/sites/default/files/old/studynotes_clark_2014.pdf%7D`.

[63] Ronald Harry Coase. *The nature of the firm*. Springer, 1995.

[64] Alma Cohen and Liran Einav. “Estimating risk preferences from deductible choice”. In: *American economic review* 97.3 (2007), pp. 745–788.

[65] Carolyn Cohn. *Insurers run from ransomware cover as losses mount*. Nov. 2021. URL: `https://www.reuters.com/markets/europe/insurers-run-ransomware-cover-losses-mount-2021-11-19/`.

[66] M. Collinson, B. Monahan, and D. Pym. *A Discipline of Mathematical Systems Modelling*. College Publications, 2012.

[67] Matthew Collinson, Brian Monahan, and David Pym. *A Discipline of Mathematical Systems Modelling*. College Publications, Aug. 2012. ISBN: 978-1-904987-50-5.

[68] Alena Yuryna Connolly and Hervé Borrion. “Reducing ransomware crime: analysis of victims payment decisions”. In: *Computers & Security* 119 (2022), p. 102760.

[69] Jerry M Couretas. *An introduction to cyber modeling and simulation*. John Wiley & Sons, 2018.

[70] Cowbell. *Cowbell and Swiss Re Partner to Offer First Ever Cyber Insurance Program Dedicated to Cloud Workloads*. Aug. 11, 2022. URL: `https://cowbell.insure/news-events/pr/cowbell-and-swiss-re-partner-on-cyber-insurance-for-cloud/`.

[71] David M Cutler, Amy Finkelstein, and Kathleen McGarry. “Preference heterogeneity and insurance markets: Explaining a puzzle of insurance”. In: *American Economic Review* 98.2 (2008), pp. 157–162.

[72] Richard M Cyert and Charles L Hedrick. “Theory of the firm: Past, present, and future; an interpretation”. In: *Journal of Economic Literature* 10.2 (1972), pp. 398–412.

[73] Mila Dalla Preda and Cinzia Di Giusto. “Hunting Distributed Malware with the κ-Calculus”. In: vol. 6914. Aug. 2011, pp. 102–113. ISBN: 978-3-642-22952-7. DOI: 10.1007/978-3-642-22953-4_9.

[74] Savino Dambra, Leyla Bilge, and Davide Balzarotti. “SoK: Cyber insurance–technical challenges and a system security roadmap”. In: *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE. 2020, pp. 1367–1383.

[75] Angelos Dassios and Ji-Wook Jang. “Pricing of catastrophe reinsurance and derivatives using the Cox process with shot noise intensity”. In: *Finance and Stochastics* 7.1 (2003), pp. 73–95.

[76] Bruno. De Finetti. *Theory of probability : a critical introductory treatment / Bruno De Finetti ; translated by Antonio Machí and Adrian Smith.* eng. Wiley series in probability and mathematical statistics. London: Wiley, 1974.

[77] Gerard Debreu. *Theory of value: An axiomatic analysis of economic equilibrium.* Vol. 17. Yale University Press, 1959.

[78] Joe Devanny, Ciaran Martin, and Tim Stevens. “On the strategic consequences of digital espionage”. In: *Journal of Cyber Policy* 6.3 (2021), pp. 429–450.

[79] Hoang Dinh Thai, Dusit Niyato, and Ping Wang. *Optimal Cost-Based Cyber Insurance Policy Management for Mobile Services.* 2017 IEEE 86th Vehicular Technology Conference. 2017. ISBN: 978-1-5090-5935-5.

[80] Georges Dionne and Scott E Harrington. *Foundations of insurance economics: readings in economics and finance.* Vol. 14. Springer Science & Business Media, 2013.

[81] Jean Dollimore, Tim Kindberg, George Coulouris, and Gordon Blair. *Distributed systems: concepts and design.* eng. International computer science series. Pearson Education, 2011. ISBN: 0273760599.

[82] Jacques H Drèze and Erik Schokkaert. “Arrows theorem of the deductible: moral hazard and stop-loss in health insurance”. In: *Journal of Risk and Uncertainty* 47 (2013), pp. 147–163.

[83] Darrell Duffie. “Black, Merton and Scholes: Their central contributions to economics”. In: *The Scandinavian Journal of Economics* 100.2 (1998), pp. 411–423.

[84] Louis Eeckhoudt and Christian Gollier. “The impact of prudence on optimal prevention”. In: *Economic Theory* 26 (2005), pp. 989–994.

[85] Maxim Egorov, Zachary N. Sunberg, Edward Balaban, Tim A. Wheeler, Jayesh K. Gupta, and Mykel J. Kochenderfer. “POMDPs.jl: A Framework for Sequential Decision Making under Uncertainty”. In: *Journal of Machine Learning Research* 18.26 (2017), pp. 1–5. URL: `http://jmlr.org/papers/v18/16-300.html`.

[86] Liran Einav, Amy Finkelstein, Iuliana Pascu, and Mark R Cullen. “How general are risk preferences? Choices under uncertainty in different domains”. In: *American Economic Review* 102.6 (2012), pp. 2606–2638.

[87] Martin Eling. “Cyber Risk and Cyber Risk Insurance - Status Quo and Future Research”. In: *Geneva Papers on Risk and Insurance-Issues and Practice* 43.2 (2018), pp. 175–179. ISSN: 1018-5895. DOI: `10.1057/s41288-018-0083-6`.

[88] Martin Eling and Werner Schnell. “What do we know about cyber risk and cyber risk insurance?” In: *The Journal of Risk Finance* (2016).

[89] Martin Eling and Werner Schnell. “What do we know about cyber risk and cyber risk insurance?” In: *Journal of Risk Finance* 17.5 (2016), pp. 474–491. ISSN: 1526-5943. DOI: `10.1108/jrf-09-2016-0122`.

[90] Martin Eling and Werner Schnell. “Capital requirements for cyber risk and cyber risk insurance: an analysis of solvency II, the US Risk-based capital standards, and the swiss solvency test”. In: *North American Actuarial Journal* 24.3 (2020), pp. 370–392.

[91] Martin Eling and Jan Wirfs. "What are the actual costs of cyber risk events?" In: *European Journal of Operational Research* 272.3 (2019), pp. 1109–1119.

[92] Martin Eling and Jingjing Zhu. "Which Insurers Write Cyber Insurance? Evidence from the U.S. Property and Casualty Insurance Industry". In: *Journal of Insurance Issues* 41.1 (2018), pp. 22–56. ISSN: 1531-6076. DOI: 10.2307/26441191.

[93] Daniel Ellsberg. "Risk, Ambiguity, and the Savage Axioms". In: *The Quarterly Journal of Economics* 75.4 (Nov. 1961), pp. 643–669. ISSN: 0033-5533. DOI: 10.2307/1884324. URL: https://doi.org/10.2307/1884324.

[94] R Elmasri, Shamkant B Navathe, R Elmasri, and SB Navathe. *Fundamentals of Database Systems</Title*. Springer, 2000.

[95] Sam Adam Elnagdy, Meikang Qiu, and Keke Gai. *Understanding Taxonomy of Cyber Risks for Cybersecurity Insurance of Financial Industry in Cloud Computing*. 2016 IEEE 3rd International Conference on Cyber Security and Cloud Computing. 2016, pp. 295–300. ISBN: 978-1-5090-0946-6. DOI: 10.1109/CSCloud.2016.46.

[96] Vassiliy A Epanechnikov. "Non-parametric estimation of a multivariate probability density". In: *Theory of Probability & Its Applications* 14.1 (1969), pp. 153–158.

[97] "Equilibrium in a Reinsurance Market". In: *Econometrica* 30.3 (1962), pp. 424–444. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1909887 (visited on 08/03/2022).

[98] Arnau Erola, Ioannis Agrafiotis, Jason R.C. Nurse, Louise Axon, Michael Goldsmith, and Sadie Creese. "A system to calculate Cyber Value-at-Risk". In: *Computers and Security* 113 (2022), p. 102545. ISSN: 0167-4048. DOI: https://doi.org/10.1016/j.cose.2021.102545. URL: https://www.sciencedirect.com/science/article/pii/S0167404821003692.

[99] Gordon C Everest. "Basic data structure models explained with a common example". In: *Proc. Fifth Texas Conference on Computing Systems*. 1976, pp. 18–19.

[100] Matthias A. Fahrenwaldt, Stefan Weber, and Kerstin Weske. "Pricing of cyber-insurance contracts in a network model". In: *Astin Bulletin* 48.3 (2018), pp. 1175–1218. ISSN: 0515-0361. DOI: 10.1017/asb.2018.23.

[101] Gregory Falco, Martin Eling, Danielle Jablanski, Virginia Miller, Lawrence A Gordon, Shaun Shuxun Wang, Joan Schmit, Russell Thomas, Mauro Elvedi, Thomas Maillart, et al. "A research agenda for cyber risk and cyber insurance". In: *Workshop on the Economics of Information Security (WEIS)*. 2019.

[102] Hanming Fang and Giuseppe Moscarini. "Morale hazard". In: *Journal of Monetary Economics* 52.4 (2005), pp. 749–777.

[103] J Doyne Farmer and François Lafond. "How predictable is technological progress?" In: *Research Policy* 45.3 (2016), pp. 647–665.

[104] Scott Farrow and Jules Szanton. "Cybersecurity Investment Guidance: Extensions of the Gordon and Loeb Model". In: *Journal of Information Security* 07.02 (2016), pp. 15–28. ISSN: 2153-1234. DOI: 10.4236/jis.2016.72002.

[105] George Feiger. *Diverse Anticipations, Rational Anticipations, Ex Ante Efficiency and Ex Post Efficiency*. Graduate School of Business, Stanford University, 1976.

[106] Shaohan Feng, Zehui Xiong, Dusit Niyato, and Ping Wang. "Competitive Security Pricing in Cyber-Insurance Market - A Game-Theoretic Analysis". In: *2018 IEEE 88th Vehicular Technology Conference*. IEEE Vehicular Technology Conference Proceedings. 2018. ISBN: 978-1-5386-6358-5.

[107] Shaohan Feng, Zehui Xiong, Dusit Niyato, Ping Wang, Shaun Shuxun Wang, and Yang Zhang. "Cyber Risk Management with Risk Aware Cyber-insurance in Blockchain Networks". In: *2018 IEEE Global Communications Conference*. IEEE Global Communications Conference. 2018. ISBN: 978-1-5386-4727-1.

[108] Burak Filiz, Budi Arief, Orcun Cetin, and Julio Hernandez-Castro. “On the Effectiveness of Ransomware Decryption Tools”. In: *Computers and Security* 111 (2021), p. 102469. ISSN: 0167-4048. DOI: https://doi.org/10.1016/j.cose.2021.102469. URL: https://www.sciencedirect.com/science/article/pii/S0167404821002935.

[109] Jörg Finsinger and Mark Pauly. “Reserve levels and reserve requirements for profit-maximizing insurance firms”. In: *Foundations of Insurance Economics*. Springer, 1984, pp. 685–704.

[110] Peter C Fishburn. “Expected utility: An anniversary and a new era”. In: *Journal of Risk and Uncertainty* 1 (1988), pp. 267–283.

[111] Jay Wright Forrester. “Industrial dynamics”. In: *Journal of the Operational Research Society* 48.10 (1997), pp. 1037–1041.

[112] Ulrik Franke. “The cyber insurance market in Sweden”. In: *Computers & Security* 68 (2017), pp. 130–144. ISSN: 0167-4048. DOI: 10.1016/j.cose.2017.04.010.

[113] Kenneth A Froot and Paul GJ O’Connell. “The pricing of US catastrophe reinsurance”. In: *The Financing of Catastrophe Risk*. University of Chicago Press, 1999, pp. 195–232.

[114] Jeffrey L Funk and Christopher L Magee. “Rapid improvements with no commercial production: How do the improvements occur?” In: *Research Policy* 44.3 (2015), pp. 777–788.

[115] Keke Gai, Meikang Qiu, and Sam Adam Elnagdy. “A Novel Secure Big Data Cyber Incident Analytics Framework for Cloud-Based Cybersecurity Insurance”. In: 2016 IEEE 2nd International Conference on Big Data Security on Cloud. 2016, pp. 171–176. ISBN: 978-1-5090-2403-2. DOI: 10.1109/BigDataSecurity-HPSC-IDS.2016.65.

[116] Keke Gai, Meikang Qiu, and Houcine Hassan. “Secure cyber incident analytics framework using Monte Carlo simulations for financial cybersecurity

insurance in cloud computing". In: *Concurrency and Computation — Practice & Experience* 29.7 (2017). ISSN: 1532-0626. DOI: 10.1002/cpe.3856.

[117] Erick Galinkin. "Winning the Ransomware Lottery". In: *International Conference on Decision and Game Theory for Security*. Springer. 2021, pp. 195–207.

[118] Gallagher Re. *Cyber in the 2020s: A question of capacity*. White Paper. Gallagher Re, Apr. 21, 2021. URL: https://www.ajg.com/gallagherre/-/media/files/gallagher/gallagherre/cyber-capacity-whitepaper.pdf (visited on 08/30/2021).

[119] Gallagher Re. *The Future of Cyber (Re)insurance*. White Paper. Gallagher Re, 2022. URL: https://www.ajg.com/gallagherre/-/media/files/gallagher/gallagherre/future-of-cyber-reinsurance.pdf.

[120] "Karush-Kuhn-Tucker (KKT) Conditions". In: *Encyclopedia of Operations Research and Management Science*. Ed. by Saul I. Gass and Michael C. Fu. Boston, MA: Springer US, 2013, pp. 833–834. ISBN: 978-1-4419-1153-7. DOI: 10.1007/978-1-4419-1153-7_200359. URL: https://doi.org/10.1007/978-1-4419-1153-7%5C_200359.

[121] Mohamed C Ghanem and Thomas M Chen. "Reinforcement learning for efficient network penetration testing". In: *Information* 11.1 (2020), p. 6.

[122] Damian Glynn, Sue Taylor, and Steven Nock. *The Basic Business Interruption Book*. 2020. URL: https://www.cila.co.uk/cila/download-link/sig-downloads/business-interruptions/371-cila-the-basic-business-interruption-book-2020/file.

[123] Christian Gollier. *The economics of risk and time*. eng. Cambridge, Mass.: MIT Press, 2001. ISBN: 0262072157.

[124] Christian Gollier. *The economics of risk and time / Christian Gollier*. Cambridge, Mass. ; London: MIT Press, 2001. ISBN: 0262072157.

[125] Christian Gollier. "Optimal insurance design of ambiguous risks". In: *Economic Theory* 57.3 (2014), pp. 555–576.

[126] Christian Gollier, James K Hammitt, and Nicolas Treich. “Risk and choice: A research saga”. In: *Journal of risk and uncertainty* 47 (2013), pp. 129–145.

[127] L. A. Gordon, M. P. Loeb, and T. Sohail. “A framework for using insurance for cyber-risk management”. In: *Communications of the Acm* 46.3 (2003), pp. 81–85. ISSN: 0001-0782. DOI: 10.1145/636772.636774.

[128] Lawrence A Gordon and Martin P Loeb. “The economics of information security investment”. In: *ACM Transactions on Information and System Security (TISSEC)* 5.4 (2002), pp. 438–457.

[129] Lawrence A Gordon, Martin P Loeb, William Lucyshyn, Lei Zhou, et al. “Externalities and the magnitude of cyber security underinvestment by private sector firms: a modification of the Gordon-Loeb model”. In: *Journal of Information Security* 6.01 (2014), p. 24.

[130] Lawrence A Gordon, Martin P Loeb, Lei Zhou, et al. “Investing in cybersecurity: insights from the Gordon-Loeb model”. In: *Journal of Information Security* 7.02 (2016), p. 49.

[131] Lawrence A Gordon, Martin P Loeb, and Lei Zhou. “Integrating cost–benefit analysis into the NIST Cybersecurity Framework via the Gordon–Loeb Model”. In: *Journal of Cybersecurity* 6.1 (2020), tyaa005.

[132] Andrew Granato, Andy Polacek, et al. “The growth and challenges of cyber insurance”. In: *Chicago Fed Letter* 426 (2019), pp. 1–6.

[133] Andy Greenberg. *The Untold Story of NotPetya, the Most Devastating Cyberattack in History.* Aug. 22, 2018. URL: https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/.

[134] Guy Carpenter. *THROUGH THE LOOKING GLASS: Interrogating the key numbers behind todays cyber market.* 2023. URL: https://www.guycarp.com/content/dam/guycarp-rebrand/pdf/Insights/2023/Guy_Carpenter_Cyber_(Re)insurance_Market_Report_Publish_rev%20.pdf.

[135] Jeroen Van Der Ham. “Toward a better understanding of cybersecurity”. In: *Digital Threats: Research and Practice* 2.3 (2021), pp. 1–3.

[136] Peter J Hammond. "Ex-ante and ex-post welfare optimality under uncertainty". In: *Economica* 48.191 (1981), pp. 235–250.

[137] Nihad A Hassan and Nihad A Hassan. "Ransomware Families: The Most Prominent Ransomware Strains". In: *Ransomware Revealed: A Beginners Guide to Protecting and Recovering from Ransomware Attacks* (2019), pp. 47–68.

[138] Kjell Hausken. "Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability". In: *Information Systems Frontiers* 8.5 (Dec. 2006), p. 338.

[139] Yezekael Hayel and Quanyan Zhu. "Attack-Aware Cyber Insurance for Risk Sharing in Computer Networks". In: *Decision and Game Theory for Security, Gamesec 2015*. Ed. by M. H. R. Khouzani, E. Panaousis, and G. Theodorakopoulos. Vol. 9406. Lecture Notes in Computer Science. 2015, pp. 22–34. ISBN: 978-3-319-25594-1.

[140] Brendan Heath. "Before the Breach - The Role of Cyber Insurance in Incentivizing Data Security". In: *George Washington Law Review* 86.4 (2018), pp. 1115–1151. ISSN: 0016-8076.

[141] Caroline Hillairet and Olivier Lopez. "Propagation of cyber incidents in an insurance portfolio: counting processes combined with compartmental epidemiological models". In: *Scandinavian Actuarial Journal* 2021.8 (2021), pp. 671–694.

[142] Caroline Hillairet, Anthony Réveillac, and Mathieu Rosenbaum. "An expansion formula for Hawkes processes and application to cyber-insurance derivatives". In: *arXiv preprint arXiv:2104.01579* (2021).

[143] Jörg Hoffmann. "Simulated Penetration Testing: From "Dijkstra" to "Turing Test++"". In: *Proceedings of the International Conference on Automated Planning and Scheduling*. Vol. 25. 1. 2015.

[144] Bengt Hölmstrom. “Moral Hazard and Observability”. In: *The Bell Journal of Economics* 10.1 (1979), pp. 74–91. ISSN: 0361915X.

[145] Zhisheng Hu, Minghui Zhu, and Peng Liu. “Adaptive Cyber Defense Against Multi-Stage Attacks Using Learning-Based POMDP”. In: *ACM Trans. Priv. Secur.* 24.1 (Nov. 2020). ISSN: 2471-2566. DOI: `10.1145/3418897`. URL: `https://doi.org/10.1145/3418897`.

[146] C. Derrick Huang and Ravi S. Behara. “Economics of information security investment in the case of concurrent heterogeneous attacks with budget constraints”. In: *International Journal of Production Economics* 141.1 (2013), pp. 255–268. ISSN: 0925-5273. DOI: `10.1016/j.ijpe.2012.06.022`.

[147] Steven R Hursh. “Behavioral economics”. In: *Journal of the experimental analysis of behavior* 42.3 (1984), pp. 435–452.

[148] Christos Ioannidis, David Pym, Julian Williams, and Iffat Gheyas. “Resilience in information stewardship”. In: *European Journal of Operational Research* 274.2 (2019), pp. 638–653.

[149] Grégoire Jacob, Eric Filiol, and Hervé Debar. “Malware as interaction machines: a new framework for behavior modelling”. In: *Journal in Computer Virology* 4.3 (2008), pp. 235–250.

[150] Jongkil Jeong, Joanne Mihelcic, Gillian Oliver, and Carsten Rudolph. “Towards an improved understanding of human factors in cybersecurity”. In: *2019 IEEE 5th International Conference on Collaboration and Internet Computing (CIC)*. IEEE. 2019, pp. 338–345.

[151] Benjamin Johnson, Rainer Boehme, and Jens Grossklags. “Security Games with Market Insurance”. In: *Decision and Game Theory for Security: Gamesec 2011*. Ed. by J. S. Baras, J. Katz, and E. Altman. Vol. 7037. Lecture Notes in Computer Science. 2011, p. 117. ISBN: 978-3-642-25279-2.

[152] Steven G. Johnson. *QuadGK.jl: Gauss–Kronrod integration in Julia*. `https://github.com/JuliaMath/QuadGK.jl`. 2013.

[153] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. "Planning and acting in partially observable stochastic domains". In: *Artificial intelligence* 101.1-2 (1998), pp. 99–134.

[154] Daniel Kahneman and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk". In: *Econometrica* 47.2 (1979), pp. 263–291. ISSN: 00129682, 14680262. URL: `http://www.jstor.org/stable/1914185`.

[155] Marek Kaluszka and Andrzej Okolewski. "An extension of Arrow's result on optimal reinsurance contract". In: *Journal of Risk and Insurance* 75.2 (2008), pp. 275–288.

[156] Daniel Kasper. "Analyzing the Feasibility of Cyber Bonds by Stochastically Solving a Copula-based Model with Differential Evolution". PhD thesis. Jan. 2019. DOI: `10.13140/RG.2.2.30180.40325`.

[157] Christopher Keegan. "Cyber security in the supply chain - A perspective from the insurance industry". In: *Technovation* 34.7 (2014), pp. 380–381. ISSN: 0166-4972. DOI: `10.1016/j.technovation.2014.02.002`.

[158] Mohammad Mahdi Khalili, Parinaz Naghizadeh, and Mingyan Liu. "Embracing Risk Dependency in Designing Cyber-Insurance Contracts". In: *2017 55th Annual Allerton Conference on Communication, Control, and Computing*. Annual Allerton Conference on Communication Control and Computing. 2017, pp. 926–933. ISBN: 978-1-5386-3266-6.

[159] Mohammad Mahdi Khalili, Parinaz Naghizadeh, and Mingyan Liu. "Designing Cyber Insurance Policies - The Role of Pre-Screening and Security Interdependence". In: *IEEE Transactions on Information Forensics and Security* 13.9 (2018), pp. 2226–2239. ISSN: 1556-6013. DOI: `10.1109/tifs.2018.2812205`.

[160] Mohammad Mahdi Khalili, Parinaz Naghizadeh, and Mingyan Liu. "Designing cyber insurance policies - The role of pre-screening and security interdependence". In: *IEEE Transactions on Information Forensics and Security* 13.9 (2018), pp. 2226–2239.

[161] Oleg Kolesnikov, Alexander Markov, Daulet Smagulov, and Sergejs Solovjovs. *Cyber bonds and their pricing models*. 2019. arXiv: `1911.06698 [q-fin.RM]`.

[162] Ravdeep Kour, Ramin Karim, and Adithya Thaduri. "Cybersecurity for railways — A maturity model". In: *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit* 234.10 (2020), pp. 1129–1148.

[163] Thomas S. Kuhn. *The structure of scientific revolutions / Thomas S. Kuhn.* eng. 2nd enl. ed. International encyclopedia of unified science ; v.2 ; no.2. Chicago ; University of Chicago Press, 1970. ISBN: 0226458032.

[164] Hanna Kurniawati. "Partially observable markov decision processes and robotics". In: *Annual Review of Control, Robotics, and Autonomous Systems* 5 (2022), pp. 253–277.

[165] Hanna Kurniawati, David Hsu, and Wee Sun Lee. "Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces." In: *Robotics: Science and systems*. Vol. 2008. Citeseer. 2008.

[166] Mordecai Kurz. "On rational belief equilibria". In: *Economic Theory* 4.6 (1994), pp. 859–876.

[167] Mordecai Kurz. *The market power of technology: Understanding the second gilded age*. Columbia University Press, 2023.

[168] Harvey E Lapan and Todd Sandler. "To bargain or not to bargain: That is the question". In: *The American Economic Review* 78.2 (1988), pp. 16–21.

[169] Aron Laszka, Sadegh Farhang, and Jens Grossklags. "On the economics of ransomware". In: *International Conference on Decision and Game Theory for Security*. Springer. 2017, pp. 397–417.

[170] Aron Laszka and Jens Grossklags. "Should Cyber-Insurance Providers Invest in Software Security?" In: *Computer Security - Esorics 2015, Pt I*. Ed. by G. Pernul, P. Y. A. Ryan, and E. Weippl. Vol. 9326. Lecture Notes in Computer Science. 2015, pp. 483–502. ISBN: 978-3-319-24174-6.

[171] Kangoh Lee. “Moral Hazard, Insurance and Public Loss Prevention”. In: *Journal of Risk and Insurance (1986-1998)* 59.2 (1992), p. 275. ISSN: 00224367. URL: `http://search.proquest.com/docview/235946196/`.

[172] Wee Lee, Nan Rong, and David Hsu. “What makes some POMDP problems easy to approximate?” In: *Advances in neural information processing systems* 20 (2007).

[173] Shanling Li, Richard Loulou, and Atiqur Rahman. “Technological progress and technology acquisition: Strategic decision under uncertainty”. In: *Production and Operations Management* 12.1 (2003), pp. 102–119.

[174] Zhen Li and Qi Liao. “Ransomware 2.0: to sell, or not to sell a game-theoretical model of data-selling Ransomware”. In: *Proceedings of the 15th International Conference on Availability, Reliability and Security.* 2020, pp. 1–9.

[175] Dahua Lin, John Myles White, Simon Byrne, Douglas Bates, Andreas Noack, John Pearson, Alex Arslan, Kevin Squire, David Anthoff, Theodore Papamarkou, Mathieu Besançon, Jan Drugowitsch, Moritz Schauer, and other contributors. *JuliaStats/Distributions.jl: a Julia package for probability distributions and associated functions.* July 2019. DOI: `10.5281/zenodo.2647458`. URL: `https://doi.org/10.5281/zenodo.2647458`.

[176] Jonathan D. Linton, Sandor Boyson, and John Aje. “The challenge of cyber supply chain security to research and practice - An introduction”. In: *Technovation* 34.7 (2014), pp. 339–341. ISSN: 0166-4972. DOI: `10.1016/j.technovation.2014.05.001`.

[177] Haibo Liu, Qihe Tang, and Zhongyi Yuan. “Indifference pricing of insurance-linked securities in a multi-period model”. In: *European Journal of Operational Research* 289.2 (2021), pp. 793–805. ISSN: 0377-2217. DOI: `https://doi.org/10.1016/j.ejor.2020.07.028`. URL: `https://www.sciencedirect.com/science/article/pii/S0377221720306391`.

[178] Liu Liu, Olivier De Vel, Qing-Long Han, Jun Zhang, and Yang Xiang. “Detecting and preventing cyber insider threats: A survey”. In: *IEEE Communications Surveys & Tutorials* 20.2 (2018), pp. 1397–1417.

[179] Wanping Liu. “Modeling Ransomware Spreading by a Dynamic Node-Level Method”. In: *IEEE Access* 7 (2019), pp. 142224–142232. DOI: 10.1109/ACCESS.2019.2941021.

[180] Zouhair Mahboubi and Mykel J Kochenderfer. “Autonomous air traffic control for non-towered airports”. In: *Proc. USA/Eur. Air Traffic Manage. Res. Develop. Seminar*. 2015, pp. 1–6.

[181] Ruperto P Majuca, William Yurcik, and Jay P Kesan. “The evolution of cyberinsurance”. In: *arXiv preprint cs/0601020* (2006).

[182] Angelica Marotta, Fabio Martinelli, Stefano Nanni, Albina Orlando, and Artsiom Yautsiukhin. “Cyber-insurance survey”. In: *Computer Science Review* 24 (2017), pp. 35–61. ISSN: 1574-0137. DOI: 10.1016/j.cosrev.2017.01.001.

[183] Angelica Marotta, Fabio Martinelli, Stefano Nanni, Albina Orlando, and Artsiom Yautsiukhin. “Cyber-insurance survey”. In: *Computer Science Review* 24 (2017), pp. 35–61.

[184] Fabio Martinelli, Albina Orlando, Ganbayar Uuganbayar, and Artsiom Yautsiukhin. “Preventing the drop in security investments for non-competitive cyber-insurance market”. In: *International Conference on Risks and Security of Internet and Systems*. Springer. 2017, pp. 159–174.

[185] Fabio Martinelli, Albina Orlando, Ganbayar Uuganbayar, and Artsiom Yautsiukhin. “Preventing the Drop in Security Investments for Non-competitive Cyber-Insurance Market”. In: *Risks and Security of Internet and Systems, Crisis 2017*. Ed. by N. Cuppens, F. Cuppens, J. L. Lanet, A. Legay, and J. GarciaAlfaro. Vol. 10694. Lecture Notes in Computer Science. 2018, pp. 159–174. ISBN: 978-3-319-76687-4.

[186] Fabio Massacci, Joe Swierzbinski, and Julian Williams. *Cyberinsurance and Public Policy - Self-Protection and Insurance with Endogenous Adversaries*. Conference Paper. 2017.

[187] Ana J. Mata. “Pricing Excess of Loss Reinsurance with Reinstatements”. In: *ASTIN Bulletin* 30.2 (2000), pp. 349–368. DOI: 10.2143/AST.30.2.504640.

[188] Kanta Matsuura. “Productivity Space of Information Security in an Extension of the Gordon-Loeb Investment Model”. In: 2009, pp. 99–119.

[189] Alessandro Mazzoccoli and Maurizio Naldi. “Robustness of Optimal Investment Decisions in Mixed Insurance/Investment Cyber Risk Management”. In: *Risk Analysis* 40.3 (2020), pp. 550–564.

[190] Alessandro Mazzoccoli and Maurizio Naldi. “An Overview of Security Breach Probability Models”. In: *Risks* 10.11 (2022), p. 220.

[191] Julie McNally. *Improving Public-Private Sector Cooperation on Cyber Event Reporting.* Proceedings of the 8th International Conference on Information Warfare and Security. 2013, pp. 147–153. ISBN: 978-1-909507-11-1.

[192] Vineet Mehta, Paul D Rowe, Gene Lewis, Ashe Magalhaes, and Mykel Kochenderfer. “Decision-theoretic approach to designing cyber resilient systems”. In: *2016 IEEE 15th International Symposium on Network Computing and Applications (NCA).* IEEE. 2016, pp. 302–309.

[193] Per Hakon Meland, Inger Anne Tondel, Marie Moe, and Fredrik Seehusen. “Facing Uncertainty in Cyber Insurance Policies”. In: *Security and Trust Management.* Ed. by G. Livraga and C. Mitchell. Vol. 10547. Lecture Notes in Computer Science. 2017, pp. 89–100. ISBN: 978-3-319-68063-7.

[194] Tobias Mettler. “Maturity assessment models: a design science research approach”. In: *International Journal of Society Systems Science (IJSSS)* 3.1/2 (2011), pp. 81–98.

[195] Robert S Miccolis. “On the theory of increased limits and excess of loss pricing”. In: *PCAS LXIV* 27 (1977).

[196] Microsoft News Center. *Microsoft and At-Bay partner to offer data-driven cyber insurance coverage.* Sept. 29, 2021. URL: `https://news.microsoft.com/2021/09/29/microsoft-and-at-bay-partner-to-offer-data-driven-cyber-insurance-coverage/`.

[197] Thomas Mikosch. *Non-life insurance mathematics : an introduction with stochastic processes / Thomas Mikosch.* eng. Universitext. Berlin: Springer, 2004. ISBN: 3540406506.

[198] R. Milner. *The Space and Motion of Communicating Agents.* Cambridge University Press, 2009.

[199] R. (Robin) Milner. *Communication and concurrency / Robin Milner.* eng. Prentice-Hall international series in computer science. New York ; Prentice Hall, 1989. ISBN: 0131149849.

[200] Kevin D Mitnick and William L Simon. *The art of deception: Controlling the human element of security.* John Wiley & Sons, 2003.

[201] Chris Moore. "Detecting ransomware with honeypot techniques". In: *2016 Cybersecurity and Cyberforensics Conference (CCC).* IEEE. 2016, pp. 77–81.

[202] Gordon E Moore et al. *Moores law at 40.* 2006.

[203] Jan Mossin. "Aspects of Rational Insurance Purchasing". In: *Journal of Political Economy* 76.4 (1968), pp. 553–568. ISSN: 00223808.

[204] Sendhil Mullainathan and Richard H Thaler. *Behavioral economics.* 2000.

[205] Patricia Munch and Dennis Smallwood. "Theory of Solvency Regulation in the Property and Casualty Insurance Industy". In: *Studies in Public Regulation.* The MIT Press, 1981. URL: `http://www.nber.org/chapters/c11431`.

[206] Munich Re. *Pioneering cyber insurance: Munich Re partners with Google Cloud and Allianz.* Mar. 2, 2021. URL: `https://www.munichre.com/en/company/media-relations/media-information-and-corporate-news/media-information/2021/pioneering-cyber-insurance.html`.

[207] Saralees Nadarajah, Yuanyuan Zhang, and Tibor K Pogány. "On sums of independent generalized Pareto random variables with applications to insurance and CAT bonds". In: *Probability in the Engineering and Informational Sciences* 32.2 (2018), pp. 296–305.

[208] Béla Nagy, J Doyne Farmer, Quan M Bui, and Jessika E Trancik. "Statistical basis for predicting technological progress". In: *PloS one* 8.2 (2013), e52669.

[209] M. Naldi and M. Flamini. "Calibration of the Gordon-Loeb Models for the Probability of Security Breaches". In: *2017 UKSim-AMSS 19th International Conference on Computer Modelling Simulation (UKSim)*. 2017, pp. 135–140.

[210] National Association of Insurance Commissioners. *Report on the cyber insurance market*. Oct. 2022. URL: `https://content.naic.org/sites/default/files/cmte-c-cyber-supplement-report-2022-for-data-year-2021.pdf`.

[211] John von Neumann, Oskar Morgenstern, and Ariel Rubinstein. *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press, 1944. ISBN: 9780691130613. URL: `http://www.jstor.org/stable/j.ctt1r2gkx`.

[212] Bernold Nieuwesteeg, Louis Visscher, and Bob de Waard. "The Law and Economics of Cyber Insurance Contracts - A Case Study". In: *European Review of Private Law* 26.3 (2018), pp. 371–420. ISSN: 0928-9801.

[213] Sokratis Nifakos, Krishna Chandramouli, Charoula Konstantina Nikolaou, Panagiotis Papachristou, Sabine Koch, Emmanouil Panaousis, and Stefano Bonacina. "Influence of human factors on cyber security within healthcare organisations: A systematic review". In: *Sensors* 21.15 (2021), p. 5119.

[214] NIST Joint Task Force. *Security and Privacy Controls for Information Systems and Organizations*. Sept. 1, 2020. DOI: `https://doi.org/10.6028/NIST.SP.800-53r5`.

[215] Dusit Niyato, Hoang Dinh Thai, Ping Wang, and Zhu Han. "Cyber Insurance for Plug-In Electric Vehicle Charging in Vehicle-to-Grid Systems". In: *IEEE Network* 31.2 (2017), pp. 38–46. ISSN: 0890-8044. DOI: `10.1109/mnet.2017.1600321nm`.

[216] Jason RC Nurse, Louise Axon, Arnau Erola, Ioannis Agrafiotis, Michael Goldsmith, and Sadie Creese. "The Data that Drives Cyber Insurance: A Study into the Underwriting and Claims Processes". In: (2020).

[217] Tim O'reilly. *What is web 2.0*. O'Reilly Media, Inc., 2009.

[218] Hulisi Ogut, Srinivasan Raghunathan, and Nirup Menon. "Cyber Security Risk Management - Public Policy Implications of Correlated Risk, Imperfect Ability to Prove Loss, and Observability of Self-Protection". In: *Risk Analysis* 31.3 (2011), pp. 497–512. ISSN: 0272-4332. DOI: 10.1111/j.1539-6924.2010.01478.x.

[219] Tridib Bandy Opadhyay, Vijay S. Mookerjee, and Ram C. Rao. "Why IT Managers Don't Go for Cyber-Insurance Products". In: *Communications of the Acm* 52.11 (2009), pp. 68–73. ISSN: 0001-0782. DOI: 10.1145/1592761.1592780.

[220] Harun Oz, Ahmet Aris, Albert Levi, and A Selcuk Uluagac. "A survey on ransomware: Evolution, taxonomy, and defense solutions". In: *ACM Computing Surveys (CSUR)* 54.11s (2022), pp. 1–37.

[221] Darren Pain. *CYBER RISK ACCUMULATION: Fully tackling the insurability challenge.* 2023. URL: https://www.genevaassociation.org/sites/default/files/2023-11/cyber_accumulation_report_91123.pdf.

[222] Ranjan Pal and Leana Golubchik. "Analyzing Self-Defense Investments in Internet Security Under Cyber-Insurance Coverage". In: *2010 International Conference on Distributed Computing Systems.* IEEE International Conference on Distributed Computing Systems. 2010. ISBN: 978-0-7695-4059-7. DOI: 10.1109/icdcs.2010.79.

[223] Ranjan Pal, Leana Golubchik, and Konstantinos Psounis. "Aegis A Novel Cyber-Insurance Model". In: *Decision and Game Theory for Security - Gamesec 2011.* Ed. by J. S. Baras, J. Katz, and E. Altman. Vol. 7037. Lecture Notes in Computer Science. 2011, pp. 131–150. ISBN: 978-3-642-25279-2.

[224] Ranjan Pal, Leana Golubchik, Konstantinos Psounis, and Pan Hui. "Will Cyber-Insurance Improve Network Security? A Market Analysis". In: *2014 Proceedings IEEE Infocom.* IEEE Infocom. 2014, pp. 235–243. ISBN: 978-1-4799-3360-0.

[225] Ranjan Pal and Pan Hui. *On Differentiating Cyber-Insurance Contracts A Topological Perspective.* 2013 IFIP/IEEE International Symposium on Integrated Network Management. 2013, pp. 836–839. ISBN: 978-3-901882-51-7.

[226] Sakshyam Panda, Daniel W. Woods, Aron Laszka, Andrew Fielder, and Emmanouil Panaousis. "Post-incident audits on cyber insurance discounts". In: *Computers & Security* 87 (2019). ISSN: 0167-4048. DOI: `10.1016/j.cose.2019.101593`.

[227] Harry H Panjer. *Operational risk: modeling analytics.* John Wiley & Sons, 2006.

[228] Harry H. Panjer and Gordon E. Willmot. *Insurance risk models.* eng. Schaumburg, Ill: Society of Actuaries, 1992. ISBN: 0938959255.

[229] Pietro Parodi. *Pricing in General Insurance / Parodi, Pietro.* eng. 1st edition. 2014. ISBN: 9781466581487.

[230] Manveer Patyal, Srinivas Sampalli, Qiang Ye, and Musfiq Rahman. "Multi-layered defense architecture against ransomware". In: *International Journal of Business and Cyber Security* 1.2 (2017).

[231] Lukas Pavlik. *Identifying and Modeling the Impact of Cyber Threats in the Field of Cyber Risk Insurance.* 2018 5th International Conference on Mathematics and Computers in Sciences and Industry. 2018, pp. 118–121. ISBN: 978-1-5386-7500-7. DOI: `10.1109/mcsi.2018.00036`.

[232] Lukas Pavlik and Roman Jasek. *Possibilities pricing of the information system by providing insurance against cyber risk.* Knowledge for Market Use 2016 - Our Interconnected and Divided World. 2016, pp. 345–351. ISBN: 978-80-87533-14-7.

[233] Krerk Piromsopa, Tomas Klima, and Lukas Pavlik. *Designing model for calculating the amount of cyber risk insurance.* 2017 Fourth International Conference on Mathematics and Computers in Sciences and in Industry. 2017, pp. 196–200. ISBN: 978-1-5386-2820-1. DOI: `10.1109/mcsi.2017.41`.

[234] John W. Pratt. "Risk Aversion in the Small and in the Large". In: *Econometrica* 32.1/2 (1964), pp. 122–136. ISSN: 00129682.

[235] Martin L Puterman. "Markov decision processes". In: *Handbooks in operations research and management science* 2 (1990), pp. 331–434.

[236] David Pym. *Category errors in (information) security: how logic can help.* Nov. 3, 2015. URL: https://www.benthamsgaze.org/2015/11/03/category-errors-in-information-security-how-logic-can-help/.

[237] Tashfiq Rahman, Rohani Rohan, Debajyoti Pal, and Prasert Kanthamanon. "Human factors in cybersecurity: a scoping review". In: *The 12th International Conference on Advances in Information Technology.* 2021, pp. 1–11.

[238] Kellyn A Wagner Ramsdell and Kristin E Esbeck. "Evolution of ransomware". In: (2021). URL: https://healthcyber.mitre.org/wp-content/uploads/2021/08/Ransomware-Paper-V2.pdf.

[239] Artur Raviv. "The Design of an Optimal Insurance Policy". In: *The American Economic Review* 69.1 (1979), p. 84. ISSN: 00028282. URL: http://search.proquest.com/docview/233054428/.

[240] Angel Marcelo Rea-Guaman, Tomás San Feliu, Jose A Calvo-Manzano, and Isaac Daniel Sanchez-Garcia. "Comparative study of cybersecurity capability maturity models". In: *International conference on software process improvement and capability determination.* Springer. 2017, pp. 100–113.

[241] Ray Rees and Achim Wambach. "The Microeconomics of Insurance". In: *Foundations and Trends in Microeconomics* 4.12 (2008), pp. 1–163. ISSN: 1547-9846. DOI: 10.1561/0700000023. URL: http://dx.doi.org/10.1561/0700000023.

[242] Ray Rees, Achim Wambach, et al. "The microeconomics of insurance". In: *Foundations and Trendső in Microeconomics* 4.1–2 (2008), pp. 1–163.

[243] Jürgen Reinhart. "Discussion on A comprehensive model for cyber risk based on marked point processes and its applications to insurance(Zeller, Scherer)". In: *European Actuarial Journal* 12.1 (2022), pp. 87–88.

[244] Osterman Research. *How to Reduce the Risk of Phishing and Ransomware.* White Paper. July 1, 2021.

[245] David Rios Insua, Aitor Couce-Vieira, and Kreshnik Musaraj. “Some Risk Analysis Problems in Cyber Insurance Economics”. In: *Estudios De Economia Aplicada* 36.1 (2018), pp. 181–194. ISSN: 1133-3197.

[246] 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: *Workshop on the Economics of Information Security 2017.* 2017.

[247] Sasha Romanosky, Lillian Ablon, Andreas Kuehn, and Therese Jones. “Content analysis of cyber insurance policies: How do carriers price cyber risk?” In: *Journal of Cybersecurity* 5.1 (2019), tyz002.

[248] Michael Rothschild and Joseph Stiglitz. “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information”. In: *The Quarterly Journal of Economics* 90.4 (1976), pp. 629–649. ISSN: 00335533, 15314650. URL: `http://www.jstor.org/stable/1885326`.

[249] Nicholas Roy, Geoffrey Gordon, and Sebastian Thrun. “Finding approximate POMDP solutions through belief compression”. In: *Journal of artificial intelligence research* 23 (2005), pp. 1–40.

[250] Keyun Ruan. “Introducing cybernomics: A unifying economic framework for measuring cyber risk”. In: *Computers and Security* 65 (2017), pp. 77–89.

[251] James Rumbaugh, Ivar Jacobson, and Grady Booch. *The Unified Modeling Language Reference Manual.* eng. 2nd ed. Addison-Wesley, 2005. ISBN: 0321245628.

[252] Pierce Ryan, John Fokker, Sorcha Healy, and Andreas Amann. *Dynamics of targeted ransomware negotiations.* 2021. arXiv: `2110.00362 [math.DS]`.

[253] Dinesh Kumar Saini, Imran Azad, Nitin B. Raut, and Lingaraj A. Hadimani. “Utility Implementation for Cyber Risk Insurance Modeling”. In: *World Congress on Engineering, Wce 2011, Vol I.* Ed. by S. I. Ao, L. Gelman, D. W. L.

Hukins, A. Hunter, and A. M. Korsunsky. Lecture Notes in Engineering and Computer Science. 2011, pp. 429–432. ISBN: 978-988-18210-6-5.

[254] Carlos Sarraute, Olivier Buffet, and Jörg Hoffmann. “POMDPs make better hackers: Accounting for uncertainty in penetration testing”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 26. 1. 2012.

[255] Carlos Sarraute, Olivier Buffet, and Jörg Hoffmann. “Penetration testing== POMDP solving?” In: *arXiv preprint arXiv:1306.4714* (2013).

[256] Harris Schlesinger and Neil A Doherty. “Incomplete markets for insurance: An overview”. In: *Journal of Risk and Insurance* (1985), pp. 402–423.

[257] Bruce Schneier. *Beyond fear: Thinking sensibly about security in an uncertain world*. Vol. 10. Springer, 2003.

[258] Reinhard Selten. *Models of strategic rationality*. Vol. 2. Springer Science & Business Media, 2013.

[259] Scott J. Shackelford. “Should your firm invest in cyber risk insurance?” In: *Business Horizons* 55.4 (2012), pp. 349–356. ISSN: 0007-6813. DOI: 10.1016/j.bushor.2012.02.004.

[260] Guy Shani, Joelle Pineau, and Robert Kaplow. “A survey of point-based POMDP solvers”. In: *Autonomous Agents and Multi-Agent Systems* 27 (2013), pp. 1–51.

[261] Carl Shapiro, Hal R Varian, Shapiro Carl, et al. *Information rules: A strategic guide to the network economy*. Harvard Business Press, 1999.

[262] Eric D Shaw. “The role of behavioral research and profiling in malicious cyber insider investigations”. In: *Digital investigation* 3.1 (2006), pp. 20–31.

[263] Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. *Competitive Cyber-Insurance and Internet Security*. Economics of Information Security and Privacy. 2010, p. 229. ISBN: 978-1-4419-6966-8.

[264] Sudha Singh, S. C. Dutta, and D. K. Singh. "Information Security and Its Insurance in the World of High Rise of Cybercrime Through a Model". In: *Proceedings of Fifth International Conference on Soft Computing for Problem Solving.* Ed. by M. Pant, K. Deep, J. C. Bansal, A. Nagar, and K. N. Das. Vol. 437. Advances in Intelligent Systems and Computing. 2016, pp. 93–98. ISBN: 978-981-10-0451-3.

[265] Henry Skeoch and Christos Ioannidis. "The barriers to sustainable risk transfer in the cyber-insurance market". In: *arXiv preprint arXiv:2303.02061* (2023).

[266] Henry Skeoch and David Pym. "Pricing cyber-insurance for systems via maturity models". In: *arXiv preprint arXiv:2302.04734* (2023).

[267] Henry R.K. Skeoch. "Expanding the Gordon-Loeb model to cyber-insurance". In: *Computers and Security* 112 (2022), p. 102533. ISSN: 0167-4048. DOI: `https://doi.org/10.1016/j.cose.2021.102533`. URL: `https://www.sciencedirect.com/science/article/pii/S0167404821003576`.

[268] Henry R.K. Skeoch. "Expanding the Gordon-Loeb model to cyber-insurance". In: *Computers & Security* 112 (2022), p. 102533.

[269] Henry RK Skeoch. "Modelling Ransomware Attacks using POMDPs". In: *Workshop on the Economics of Information Security.* 2022. URL: `https://weis2022.econinfosec.org/wp-content/uploads/sites/10/2022/06/weis22-skeoch.pdf`.

[270] Costis Skiadas. "Smooth ambiguity aversion toward small risks and continuous-time recursive utility". In: *Journal of Political Economy* 121.4 (2013), pp. 775–792.

[271] Guido de Smidt and Wouter Botzen. "Perceptions of Corporate Cyber Risks and Insurance Decision-Making". In: *Geneva Papers on Risk and Insurance-Issues and Practice* 43.2 (2018), pp. 239–274. ISSN: 1018-5895. DOI: `10.1057/s41288-018-0082-7`.

[272] Adam Smith. *An inquiry into the nature and causes of the wealth of nations. By Adam Smith, ... In three volumes.* eng. Dublin: printed for Messrs. Whitestone, Chamberlaine, W. Watson, Potts, S. Watson [and 15 others in Dublin], 1776.

[273] Wes Sonnenreich, Jason Albanese, and Bruce Stout. "Return on security investment (ROSI)-a practical quantitative model". In: *Journal of Research and practice in Information Technology* 38.1 (2006), pp. 45–56.

[274] Michael Spence. "Job Market Signaling". In: *The Quarterly Journal of Economics* 87.3 (1973), pp. 355–374. ISSN: 00335533.

[275] Ross M Starr. "Optimal production and allocation under uncertainty". In: *The Quarterly Journal of Economics* 87.1 (1973), pp. 81–95.

[276] John Sterman. "System Dynamics: systems thinking and modeling for a complex world". In: (2002).

[277] Kayla Strong. "Multi-line insurance clash management". In: Cambridge Centre for Risk Studies 2018 Summit. 2018.

[278] Shauhin A. Talesh. "Data Breach, Privacy, and Cyber Insurance - How Insurance Companies Act as Compliance Managers for Businesses". In: *Law and Social Inquiry-Journal of the American Bar Foundation* 43.2 (2018), pp. 417–440. ISSN: 0897-6546. DOI: 10.1111/lsi.12303.

[279] Unal Tatar, Omer Keskin, Hayretdin Bahsi, and C. Ariel Pinto. *Quantification of Cyber Risk for Actuaries: An Economic-Functional Approach.* 2020. URL: https://www.soa.org/49c222/globalassets/assets/files/resources/research-report/2020/quantification-cyber-risk.pdf.

[280] Ken-ichi Tatsumi and Makoto Goto. "Optimal timing of information security investment: A real options approach". In: *Economics of Information Security and Privacy.* Springer, 2010, pp. 211–228.

[281] Inger Anne Tondel, Fredrik Seehusen, Erlend Andreas Gjaere, and Marie Elisabeth Gaup Moe. "Differentiating Cyber Risk of Insurance Customers - The Insurance Company Perspective". In: *Availability, Reliability, and Security*

in Information Systems, Cd-Ares 2016, Paml 2016. Vol. 9817. Lecture Notes in Computer Science. 2016, pp. 175–190. ISBN: 978-3-319-45507-5.

[282] Michail Tsikerdekis, Sherali Zeadally, Amy Schlesener, and Nicolas Sklavos. "Approaches for preventing honeypot detection and compromise". In: *2018 Global Information Infrastructure and Networking Symposium (GIIS)*. IEEE. 2018, pp. 1–6.

[283] Charles Tupper. *Data architecture: from zen to reality*. Elsevier, 2011.

[284] United States Securities and Exchange Commission. *SEC Adopts Rules on Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure by Public Companies*. URL: `https://www.sec.gov/news/press-release/2023-139`.

[285] David S. Wall. "The Internet as a Conduit for Criminal Activity". In: *Information Technology and the Criminal Justice System*. Ed. by A Pattavina. 2015th ed. Thousand Oaks, California: Sage Publications, 2005, pp. 77–98. URL: `https://papers.ssrn.com/abstract=740626` (visited on 04/29/2020).

[286] Erwin Walraven and Matthijs Spaan. "Accelerated vector pruning for optimal POMDP solvers". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 31. 1. 2017.

[287] Shaun S. Wang. "Integrated framework for information security investment and cyber insurance". In: *Pacific-Basin Finance Journal* 57 (2019). ISSN: 0927-538X. DOI: `10.1016/j.pacfin.2019.101173`.

[288] J Willemson. "On the Gordon and Loeb Model for Information Security Investment". In: *WEIS Proceedings 2006*. 2006. URL: `https://www.econinfosec.org/archive/weis2006/docs/12.pdf`.

[289] David Williams. *Probability with martingales*. Cambridge university press, 1991.

[290] Charles Wilson. "A model of insurance markets with incomplete information". In: *Journal of Economic Theory* 16.2 (1977), pp. 167–207. ISSN: 0022-0531.

[291] Ralph A Winter. “The liability crisis and the dynamics of competitive insurance markets”. In: *Yale J. on Reg.* 5 (1988), p. 455.

[292] Josephine Wolff. “The role of insurers in shaping international cyber-security norms about cyber-war”. In: *Contemporary Security Policy* (2023), pp. 1–30.

[293] Stephen Wolfram. *Mathematica: a system for doing mathematics by computer.* Addison Wesley Longman Publishing Co., Inc., 1991.

[294] Daniel Woods, Ioannis Agrafiotis, Jason RC Nurse, and Sadie Creese. “Mapping the coverage of security controls in cyber insurance proposal forms”. In: *Journal of Internet Services and Applications* 8.1 (2017), pp. 1–13.

[295] Daniel Woods and Rainer Böhme. “How Cyber Insurance Shapes Incident Response: A Mixed Methods Study”. In: Workshop on the Economics of Information Security. June 1, 2021. URL: https://weis2021.econinfosec.org/wp-content/uploads/sites/9/2021/06/weis21-woods.pdf.

[296] Daniel Woods and Andrew C. Simpson. *Monte Carlo methods to investigate how aggregated cyber insurance claims data impacts security investments.* Conference Paper. 2018.

[297] Daniel W Woods. “A Turning Point for Cyber Insurance”. In: *Communications of the ACM* 66.3 (2023), pp. 41–44.

[298] Daniel W Woods, Tyler Moore, and Andrew C Simpson. “The county fair cyber loss distribution: Drawing inferences from insurance prices”. In: *Digital Threats: Research and Practice* 2.2 (2021), pp. 1–21.

[299] Xiaoying Xie, Charles Lee, and Martin Eling. “Cyber insurance offering and performance: An analysis of the US cyber insurance market”. In: *The Geneva Papers on Risk and Insurance-Issues and Practice* 45 (2020), pp. 690–736.

[300] Maochao Xu and Lei Hua. “Cybersecurity Insurance - Modeling and Pricing”. In: *North American Actuarial Journal* 23.2 (2019), pp. 220–249. ISSN: 1092-0277. DOI: 10.1080/10920277.2019.1566076.

[301] Zichao Yang and John C. S. Lui. "Security Adoption in Heterogeneous Networks - the Influence of Cyber-Insurance Market". In: *Networking 2012, Pt Ii*. Vol. 7290. Lecture Notes in Computer Science. 2012, pp. 172–183. ISBN: 978-3-642-30054-7.

[302] Zichao Yang and John C. S. Lui. "Security adoption and influence of cyber-insurance markets in heterogeneous networks". In: *Performance Evaluation* 74 (2014), pp. 1–17. ISSN: 0166-5316. DOI: 10.1016/j.peva.2013.10.003.

[303] Tongxin Yin, Armin Sarabi, and Mingyan Liu. "Deterrence, Backup, or Insurance: A Game-Theoretic Analysis of Ransomware". In: Workshop on the Economics of Information Security. June 1, 2021. URL: https://weis2021.econinfosec.org/wp-content/uploads/sites/9/2021/06/weis21-yin.pdf.

[304] A. Young and Moti Yung. "Cryptovirology: extortion-based security threats and countermeasures". In: *Proceedings 1996 IEEE Symposium on Security and Privacy*. 1996, pp. 129–140. DOI: 10.1109/SECPRI.1996.502676.

[305] Derek Young, Juan Lopez, Mason Rice, Benjamin Ramsey, and Robert McTasney. "A framework for incorporating insurance in critical infrastructure cyber risk strategies". In: *International Journal of Critical Infrastructure Protection* 14 (2016), pp. 43–57. ISSN: 1874-5482. DOI: https://doi.org/10.1016/j.ijcip.2016.04.001. URL: http://www.sciencedirect.com/science/article/pii/S1874548216300439.

[306] Gabriela Zeller and Matthias Scherer. "A comprehensive model for cyber risk based on marked point processes and its application to insurance". In: *European Actuarial Journal* 12.1 (2022), pp. 33–85.

[307] Rui Zhang, Quanyan Zhu, and Yezekael Hayel. "A Bi-Level Game Approach to Attack-Aware Cyber Insurance of Computer Networks". In: *IEEE Journal on Selected Areas in Communications* 35.3 (2017), pp. 779–794. ISSN: 0733-8716. DOI: 10.1109/jsac.2017.2672378.

[308] Philip R Zimmerman and Peter Ludlow. *How PGP works/why do you need PGP*. The MIT Press Cambridge, 1996.

Appendix

The following replicates the CRediT statements [8] for those papers from which the below chapters were derived. All other work in the thesis is solely the work of the author except where otherwise indicated.

Chapter 4: Pricing Cyber-Insurance Based on System Structure

Henry Skeoch: Conceptualization, Methodology, Writing — Original Draft, Visualization

David Pym: Conceptualization, Methodology, Writing — Original Draft, Visualization, Supervision, Funding acquisition.

Chapter 5: Modelling the Cyber-Insurance Market with Risk Transfer via Reinsurance

Henry Skeoch: Conceptualization, Methodology, Software, Formal analysis, Writing — Original Draft, Writing — Review and Editing, Visualization, Project administration.

Christos Ioannidis: Methodology, Formal analysis, Writing — Original Draft, Supervision.