Expanding the Gordon-Loeb Model to Cyber-Insurance

Henry R.K. Skeoch[a,*]

[a]*Department of Computer Science, University College London, Gower Street, London, WC1E 6EA, United Kingdom*

Abstract

We present an economic model for decisions on competing cyber-security and cyber-insurance investment based on the Gordon-Loeb model for investment in information security. We consider a one-period scenario in which a firm may invest in information security measures to reduce the probability of a breach, in cyber-insurance or in a combination of both. The optimal combination of investment and insurance under the assumptions of the Gordon-Loeb model is investigated via consideration of the costs and benefits of investment in security alongside purchasing insurance at an independent premium rate. Under both exponential (constant absolute risk aversion) and logarithmic (constant relative risk aversion) utility functions it is found that when the insurance premium is below a certain value, utility is maximised with insurance and security investment. These results suggest that cyber-insurance is a worthwhile undertaking provided it is not overly costly. We believe this model to be the first attempt to integrate the Gordon-Loeb model into a classical microeconomic analysis of insurance, particularly using the Gordon-Loeb security breach functions to determine the probability of an insurance claim. The model follows the tradition of the Gordon-Loeb model in being accessible to practitioners and decision makers in information security.

Keywords: Gordon-Loeb Model, Cybersecurity, Cyber-insurance, Security Investment, Security Economics

1. Introduction

The rapid growth in information technology, especially the internet, during the second half of the twentieth century revolutionised communication between firms and society, particularly the speed with which information can be exchanged. An individual can now undertake banking transactions, shop and correspond with friends and family almost simultaneously on a hand-held device; this set of tasks might have occupied an entire morning just over 20 years ago. However, this speed of interaction has a downside: those with a malicious agenda can undertake nefarious activities as rapidly. Electronic interactions generate a huge amount of data. Any good with value is a potential target for theft[1] and data is no exception. Consumer personal data have value to cyber-criminals looking to perpetrate fraud and such data are now stored or processed by a wide variety of online retailers or service providers. The risks to consumers are recognised by regulations such as the EU General Data Protection Regulation, which can impose significant financially penalties on businesses for data breaches.

The confidentiality-integrity-availability (CIA) triad is a popular framework in information security for understanding risks and identifying potential solutions. It is particularly useful in understanding the risks around data and potential related defences. Certain data needs to be kept confidential - accessible only to the right people at the right time. This might be personal data protected by law as already described, or information of significant commercial or strategic value. Integrity is also important; a crude cyber-attack might aim at corrupting data and thereby cripple either businesses or infrastructure. If appropriate backups are unavailable, or even paper hard copies, corruption or addition of false data could be catastrophic for an organisation. Finally, availability is also critical for maintaining business operations - if an airline cannot access its baggage system, then it would surely be unable to operate properly. The fields of information security and operational research (particularly business continuity) cover these areas; tangible risks arising might be mitigated via either investment in security or insurance.

The field of computer and network security is already addressed by a vast body of literature; Anderson (2001) provides an accessible introduction, tackling the problem from an economic perspective. A wide variety of technical solutions exist covering defence, monitoring and reporting as well as frameworks and policies designed to minimise risks associated with human interaction with technology (such as users' choices of password, access control mechanisms and automated enforcement of security policies). A detailed assessment of the trade-offs associated with these measures is beyond the scope of our work, which aims to address the issue of security investment from a 'top down' rather than 'bottom up' perspective. It is worthwhile considering the nature of the threat, however, albeit one that we will simplify considerably. The view of a 'hacker' in the popular media historically was often of a lone individual operating alone from their own residence, sometimes with significant success[2]. Nowadays, a lone attacker would proba-

*Corresponding author

URL: henry.skeoch.19@ucl.ac.uk (Henry R.K. Skeoch)

[1] For an interesting discussion of the impact of theft on efficiency, see Usher (1987)

[2] A prominent example of this type is Scottish systems administrator, Gary

bly struggle to enumerate all the possible attack surfaces of the systems of a competent organisation without detection and resultant defensive action. However, cyber-attacks have increased in sophistication and complexity with some believed to be state sponsored[3]. With these resources, it is possible to develop custom exploits that have a degree of stealth making them harder to detect and prevent. The idea of cyber-insurance becomes appealing when assessing the risks of compromise via unknown vectors. We make the assumption in this work that a cyber-insurance policy will pay out with certainty in the event of a claim; it should be noted that the alleged involvement of a state entity can complicate matters as an insurer may attempt to argue that such an attack is an 'act of war' and therefore not eligible for a claim. A modern cyber-insurance policy will typically cover not just potential financial losses associated with a cyber-attack but also the costs of forensic computer experts to help assess the extent of any breach. This acts as an inducement to the insurance buyer as forensic expertise would be expensive to retain on a company payroll if it was expected that their services would be only occasionally required.

Gordon and Loeb (2002) proposed a model for decisions on information security investment, in which the probability of a security breach occurring reduces with investment according to a specified function. 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 maximisation 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. 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.

2. Literature Review

The model presented in this work draws on the theory of competitive insurance markets. The formal discipline of insurance economics was arguably established by the work of Borch, Pratt, Arrow and Mossin in the 1960s, following the seminal contributions by von Neumann and Morgenstern on expected utility theory. This was followed by key developments in the 1970s with regard to asymmetric information, particularly the celebrated contributions of Akerlof, Spence and Rothschild & Stiglitz. A literature detailing the theory of insurance supply also subsequently developed. The coverage of insurance economics in this paper aims to inform or remind the reader of some notable contributions to the literature and does not claim to provide a complete survey of a diverse and well-established field. Having provided a summary of relevant literature specifically covering insurance, the Gordon-Loeb model of information security investment is then presented, which forms the basis of the model introduced in this work. The body of derivative literature around the Gordon-Loeb model is also surveyed, including criticisms and extensions to the model. Finally, specific cyber-insurance literature useful for comparison with the results of the model developed in this paper is discussed.

2.1. Insurance Economics

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. In classical economics, expected utility is usually descriptive rather than normative[4]. von Neumann et al. (1944) introduced an axiomatic version of Expected Utility Theory. 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 (1953) and Ellsberg (1961) paradoxes provide noted counterexamples. One of the most famous critiques of expected utility theory known as prospect theory was introduced by Kahneman and Tversky (1979). 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

McKinnon, who was indicted (US Department of Justice (2002)) on charges of compromising almost 100 US military computer systems in the early 2000s

[3] Sanger (2019) and Greenberg (2019) provide interesting and highly readable accounts of alleged activities of this nature. These are journalistic accounts rather than works of scholarship, but make useful contributions given the authors' access to government sources.

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

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.

Borch (1967) proposed a key distinction that the ends or objectives of an economic analysis of insurance ought not to be subservient to the means of analysis available. Borch argued initially that insurance 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 depends on the objectives of the insurance company. Formalising the developed ideas, Borch assumes the insurance company has a complete preference ordering over the set of all profit distributions. Framing the problem in terms of utility, Borch reformulates the original problem reducing the task to maximizing a mathematical expression. Finally, Borch discusses some of the issues involved in applying traditional economic analysis to insurance. He proposes an equilibrium price in which total insurance supply equates to total insurance demand. He goes on to provide a critique of the application of classical market theory in relation to insurance, as there is no natural unit of insurance cover from which to define a price. However, he posits that Pareto-optimality is readily defined for insurance which naturally leads to a Game Theory approach to the problem.

Pratt (1964) introduced $r(x) = -u''(x)/u'(x)$ as a measure of local risk aversion, where $u(x)$ is a utility function for money; this is often known as Arrow-Pratt risk aversion given contemporaneous work by Arrow. Mossin (1968) analysed 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. These are foundational to the theory of insurance demand and are collectively sometimes called the Mossin Theorem. Arrow (1974) considered optimal insurance and generalized deductibles. He demonstrated that a risk averse buyer will prefer a policy offering complete coverage beyond a deductible (an amount of loss below which no claim is paid by an insurer). This form of contract effectively places a cap on the loss of wealth an individual may incur.

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 insurere is not privy; moral hazard is the risk that possessing insurance encourages risky behaviour.

One of the most important contributions in understanding asymmetric information is *The Market for Lemons* by Akerlof (1970). 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 (1973) introduced the idea of signalling within the context of the job market. 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 (1976) consider competitive markets in which the "characteristics of the commodities exchanged are not fully known to at least one of the parties to the transaction." 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 (1977) also found that no stationary equilibrium may exist if all firms have static expectations with regard to the policy offers of other firms. 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 Hölmstrom (1979). 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 (1992) considers how the problem of moral hazard might be solved by provision of a loss-preventative good. For cyber-insurance, an example would be the government providing anti-malware software to the population.

2.1.3. Insurance supply and pricing

A key contribution in explaining the supply of insurance was made by Raviv (1979), who noted that in the earlier model proposed by Arrow(1974), 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. 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. The key result from this analysis is that if the cost of providing insurance is independent of the insurance contract, then the Pareto optimal contract does not have a deductible.

Borch (1981) 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. 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 (1981) 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. 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 (1984) 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.

2.2. Security investment models

2.2.1. The Gordon-Loeb Model and some alternatives

Gordon and Loeb (2002) introduced an economic model that determines that optimal amount to invest to protect a given set of information. The Gordon-Loeb (henceforth GL) model is discussed in full detail in the following section of this paper, 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. 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. Boehme (2010) gives a good summary of security investment models, their terminology and parameters. Huang and Behara (2013) 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. 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 (2006) who provides a counter-example via the use of a logistic function but with quite a few changes to the original Gordon-Loeb assumptions. Willemson (2006) 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. 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 (2012) 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.

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 (2009) proposes a productivity space of information security, specifically considering a productivity regarding threat reduction and a productivity involving vulnerability reduction. In essence, this might be regarded as

an extension of the original Gordon-Loeb model to a two-dimensional case. Tatsumi and Goto (2010) add a timing dimension to the original Gordon-Loeb model using a real options approach. Lelarge (2012) shows that the Gordon and Loeb $1/e$ limit holds under log-convex security functions, and extends the G-L model to a security game where agents consider the implications of their actions on the network with the interesting result that the fulfilled equilibrium is not socially efficient. Gordon et al. (2014) extend the original Gordon-Loeb model to include costs associated with the externalities of security breaches rather than just the firm's private costs. The revised model is sometimes referred to in the literature as the GLLZ model. Farrow and Szanton (2016) propose extensions to the Gordon-Loeb and GLLZ models, based on mathematical equivalency with a generalized homeland security model. Gordon et al. (2016) explain how the Gordon-Loeb model can be used in a practical setting and the intuition underpinning the model's parameters. Young et al. (2016) 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. This work is the closest example we have seen to our approach in the literature and is discussed in detail in Section 3.1.2. Naldi and Flamini (2017) provide a thorough investigation of the productivity parameters in both classes of Gordon-Loeb security breach functions and propose estimators for these parameters. Mazzoccoli and Naldi (2020) expand upon the work of Young et al. (2016) by providing closed form solutions to the same problem; this work is covered in detail in Section 3.1.

2.3. Cyber-insurance

For a topic of apparent significant commercial and intellectual interest, the literature on cyber-insurance appears relatively underdeveloped. The field can be broadly classified into: economic modelling, frameworks and policy, game theory, law and surveys and empirical analysis. However, a search of the ISI Web of Science finds fewer than 100 relevant papers on cyber-insurance. We briefly cover some key papers that in our view make a useful contribution to the literature. Our work contributes primarily to the literature on the economic aspects of cyber-insurance, and therefore give papers in this discipline greatest critical focus.

2.3.1. Economic modelling

An early contribution in this area is by Bojanc and Jerman-Blazic (2008) who 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 but cite the work of Majuca et al. (2006) as cause for concern that some cyber policies may not pay out. Pal et al. (2010) 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. 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 (Shetty et al. (2010), Pal et al. (2014) and Martinelli et al. (2018)). Khalili et al. (2017) suggest that an insurance company can increase profit by insuring both a primary and associated party and that this reduces collective risk. 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 Khalili et al. (2018).

The literature on pure cyber-insurance modelling is fairly limited: Pal et al. (2011) 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. Bodin et al. (2018) 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. Bandyopadhyay and Mookerjee (2019) 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. This appears an area where there is significant opportunity for further work.

The literature on theoretical pricing of cyber-insurance appears particularly sparse and underdeveloped. Saini et al. (2011) attempt to produce a model for deriving utility functions for cyber-insurance, using a university network as an example. Determining the optimal utility function to describe insurance buyer and supplier behaviour is fundamental in developing a sound pricing model, making this a useful contribution. Fahrenwaldt et al. (2018) introduce a polynomial approximation of claims together with a mean-field approach that allows to compute aggregate expected losses and prices of cyber-insurance. However, the limited data publicly available around cyber-insurance would make such a model difficult to validate. Piromsopa et al. (2017) propose a rudimentary cyber-insurance scoring model, which can incorporate existing security standards - this is most applicable to enterprise risk management. Xu and Hua (2019) 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. This is an interesting approach, although the cyberattack model is somewhat simplistic relative to the variety of overall threats.

2.3.2. 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 et al. (2003) 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: Bohme and Kataria (2006) investigate the potential limits of cyber-insurance in the context of the high correlation of potential risks, which Opadhyay et al. (2009) develop arguing that

cyber-insurance tends to be overpriced as insurers cannot estimate the potential secondary losses of customers. Ogut et al. (2011) 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. How the latter would be achieved in practice would depend on regulation.Shackelford (2012) argues that firms should take a proactive stance toward managing cyber attacks implicitly cautioning against over-reliance on cyber-insurance; this perspective is somewhat countered by Laszka and Grossklags (2015) who suggest that insurance providers taking a role in helping improve software security can lead to a more profitable cyber-insurance market.

In terms of papers describing the field, Linton et al. (2014) and Keegan (2014) summarise research in the cyber-security chain while Elnagdy et al. (2016) outline the taxonomy of cyber-risks for cyber-security insurance of the financial industry in cloud computing. There are a few recent papers which propose frameworks for cyber-insurance: Gai et al. (2016), Pavlik and Ieee (2018) and Khalili et al. (2018) who investigate cloud based insurance for big data, organisation insurance and pre-screening and security interdependence respectively. 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.

2.3.3. Game Theory

A reasonable number of papers have attempted to analyse cyber-insurance from a game theoretic perspective. One particularly relevant contribution is from Massaccia et al. (2017) who 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. 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. (2011a) find that equilibria with a joint investment in protection and self-insurance may exist in a one-shot security game. This somewhat contradicts the narrative that cyber-insurance does not improve network security, though the work involves significant simplifying assumptions. 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 and Lui (2012) and Yang and Lui (2014) investigate cyber-insurance as part of a Bayesian network game analysis on security investment. They argue that when insurance is offered at the actuarially fair price (highly unlikely in practice) that the optimal insurance is full coverage. Pal and Hui (2013) propose Bonacich/Eigenvector centralities of network users as an appropriate parameter for differentiating insurance clients. Hayel and Zhu (2015) and Zhang et al. (2017) deploy a game-in-games framework where a zero-sum game is nested within a moral-hazard game problem to model cyber-insurance. Martinelli et al. (2017) investigate how a drop in security investments for non-competitive cyber-insurance markets might be prevented. Rios Insua et al. (2018) model a number of cyber-insurance problems as a network and offer decision making models for cyber-insurance, but the models are not solved or analysed in detail making their merits inconclusive. Woods and Simpson (2018) investigate how aggregated claims data impacts investments in information security using Monte Carlo methods to simulate an extended iterative weakest link model.

2.3.4. 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. Nieuwesteeg et al. (2018) provide the first contemporary legal analysis of cyber-insurance contracts we are aware of focused on the Netherlands. 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 (2018) conducts an analysis contributing to two literatures on organisational compliance: new institutional organisational sociology studies of how organisations respond to legal regulation and sociolegal insurance research on how institutions govern through risk. Talesh concludes that insurers act as de facto compliance managers for organisations dealing with cyber security threats via the provision of risk management services. Heath (2018) 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.

2.3.5. Surveys and empirical analyses

The final category of the cyber-insurance literature we review is concerned with surveys (largely of individual corporate decision makers) and empirical analyses of the state of the cyber-insurance market. Biener et al. (2015) 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. The lack of business and economics literature on cyber-insurance has been identified by Eling and Schnell (2016), who concur with this review in describing the lack of data and modelling approaches in the cyber-insurance literature. Tondel et al. (2016) 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. Marotta

et al. (2017) undertake a highly comprehensive survey of cyber-insurance, albeit analysing only a small number of insurance firms. Their characterisation 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.

One of the most important contributions in the cyber-insurance literature is by Romanosky et al. (2017) who collect and analyse over 100 cyber-insurance policies filed with state insurance commissioners in the United States. This is an important paper, as it exploits insurance regulation in the US that requires the filing of policies and as such is superior to a survey in so far as the content is less likely to be biased. They 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.

There has been a recent trend towards surveys becoming more targeted. Woods et al. (2017) 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. de Smidt and Botzen (2018) 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. 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 (2018) 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. Nurse et al. (2020) investigate the types of data used in pricing cyber-insurance via a qualitative study of professional practitioners including underwriters and actuaries. 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; it is not clear whether the individuals interviewed came just from this individual firm.

3. Model

3.1. Related Work

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:

$$EBIS(z) = [v - S(z, v)]\tau l \tag{1}$$

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):

$$ENBIS(z) = [v - S(z, v)]\tau l - z \tag{2}$$

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:

$$-S_z(z^*, v)\tau l = 1 \tag{3}$$

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[5] 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:

$$z^{I*}(v) = \frac{(v\beta\alpha l\tau)^{1/(\beta+1)} - 1}{\alpha} \tag{4}$$

$$z^{II*}(v) = \frac{\ln(1/-\alpha v l\tau(\ln v))}{\alpha \ln v} \tag{5}$$

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 1 for vulnerability $v = 0.65$ and the corresponding ENBIS for these breach functions.

3.1.2. Critique of the approach of Young et al (2016) to combining the Gordon-Loeb model and Cyber-insurance

Young et al. (2016) adopt a similar conceptual approach to the model introduced in this paper in terms of setting up an optimisation problem incorporating parameters from the Gordon-Loeb model. They propose minimising the expression

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

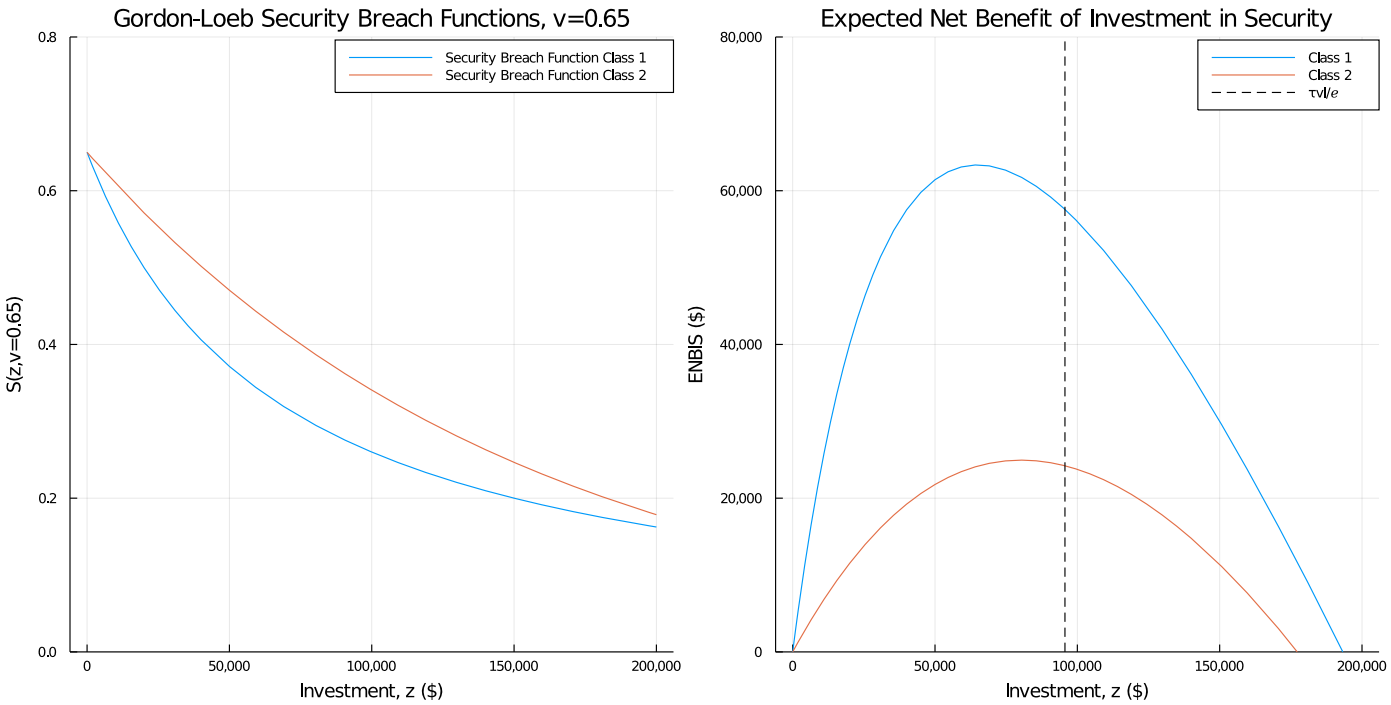

Figure 1: Example Gordon-Loeb security breach functions

$S(z,v)\tau l + z + P$ 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 minimisation 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.

3.1.3. Mazzoccoli and Naldi (2020) on mixed insurance/investment cyber-risk management

Mazzoccoli and Naldi (2020) have recently produced a valuable contribution to the literature on cyber-insurance by pursuing a similar approach to Young et al (2016). 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 paper 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[6]. 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.

3.1.4. Return on Security Investment (ROSI)

Sonnenreich et al. (2006) propose a measure for calculating the value of security expenditure, the return on security investment (ROSI):

$$\text{ROSI} = \frac{(\text{RiskExposure} * \%\text{RiskMitigated}) - \text{SolutionCost}}{\text{SolutionCost}} \tag{6}$$

This measure is broadly similar to ENBIS in the Gordon-Loeb Model (Equation 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 Sonnenreich et al. (2006). This is an area, therefore,

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

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.

3.2. Expanding the Gordon-Loeb model to include cyber-insurance

We introduce a simple model following Rees and Wambach (2008) to describe the microeconomic analysis of a firm aiming to determine its optimal level of cyberinsurance cover. 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 cybersecurity 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[7] of an insurance buyer where there are two states, no-loss and loss:

$$\begin{aligned} E[U] &= (1 - S(v,z))u(B_{sec} - z - P(C)) \\ &+ S(z,v)u(B_{sec} - z - P(C) - \tau l S(z,v) + C) \end{aligned} \tag{7}$$

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 $\tau l S(z,v)$. 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[8], which is increasing and strictly concave implying that the individual is risk averse. The form of the utility function is described in Section 3.3.

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[9] and the decision process around that investment[10] with respect to the system are both excluded. This leaves the model relatively abstract[11] 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 organisations 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 paper 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 maximisation problem then becomes:

$$\begin{aligned} \max_{C \geq 0} \bar{u} &= (1 - S(v,z))u(B_{sec} - z - P(C)) \\ &+ S(z,v)u(B_{sec} - z - \tau l S(z,v) + C - P(C)) \end{aligned} \tag{8}$$

subject to the constraints $P(C) = pC$ where p represents a percentage premium (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 8 yields:

$$\begin{aligned} \bar{u}(C,z) &= (1 - S(v,z))u(B_{sec} - z - pC) \\ &+ S(z,v)u(B_{sec} - z - \tau l S(z,v) + C(1-p)) \end{aligned} \tag{9}$$

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[12] for the problem depicted in equation 8 is:

$$Z = U + \lambda(\frac{v\tau l}{e} - pC - z) \tag{10}$$

where $U = \bar{u}(C,z)$ The Karush-Kuhn-Tucker conditions[13] are (where Z_x denotes the partial derivative of Z with respect to x):

$$\begin{aligned} Z_C = U_C - p\lambda \leq 0 \qquad & C \geq 0 \qquad & C.Z_C = 0 \\ Z_z = U_z - \lambda \leq 0 \qquad & z \geq 0 \qquad & z.Z_z = 0 \\ Z_\lambda = \frac{v\tau l}{e} - pC - z \geq 0 \qquad & \lambda \geq 0 \qquad & \lambda.Z_\lambda = 0 \end{aligned} \tag{11}$$

[7] In economics, the concept of utility aims to describe a set of preferences using mathematical functions

[8] Such a utility function is one that conforms to the four axioms proposed in von Neumann et al. (1944)

[9] 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

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

[11] 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.

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

[13] 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 paper may be found - see, for example, Gass and Fu (2013) for a formal definition. As the form of the utility function is variable, it is important not to lose generality at this stage.

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

$$\lambda = \frac{U_z - U_C}{1 - p} \tag{12}$$

This then implies that the fair premium is given by

$$p = \frac{U_C}{U_z} \tag{13}$$

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 (11) 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 optimisation strategy for model cases rather than producing a closed-form solution *ab initio* that is algebraically intractable and unintuitive to interpret.

3.3. The utility function in relation to insurance

As has been demonstrated, it is possible to make useful judgements regarding the formulation of the optimal security investment/insurance problem without stipulating a precise form for the utility function. As elegantly described by Gollier (2001), however, "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.". 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 analysis, a von Neumann-Morgenstern utility function is needed that allows for an analysis broadly consistent with the two different forms of security breach function in the GL model but is also able to capture the preferences of different types of firms.

It is worth considering the risk tolerance of a firm considering investments in information security. The firm should be risk averse; if it were totally risk tolerant, it would be willing to risk the costs of a breach. It should also be aiming to maximise its wealth as a rational actor. Together, these preferences imply $U' > 0$ and $U'' < 0$. There are three key properties in relation to the utility function that are usually considered, absolute risk aversion:

$$A(z) = -\frac{u''(z)}{u'(z)} \tag{14}$$

prudence:

$$P(z) = -\frac{u'''(z)}{u''(z)} \tag{15}$$

and relative risk aversion:

$$R(z) = -\frac{zu''(z)}{u'(z)} = zA(z) \tag{16}$$

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

$$u(z) = \frac{1 - e^{-az}}{a} \tag{17}$$

or constant relative risk aversion (CRRA):

$$u(z) = \begin{cases} z^{(1-\gamma)}/(1-\gamma) & \text{if } \gamma \neq 1 \\ ln(z) & \text{if } \gamma = 1 \end{cases} \tag{18}$$

As noted by Johnson et al. (2011b), CRRA is an established choice within the cyber-insurance literature though examples of CARA are also found. For completeness, we examine the properties of both CRRA and CARA utility functions in the following simulations of the GLCI model.

4. Method

4.1. Simulation

The simulations of the Gordon-Loeb with Cyber-Insurance (GLCI) model use the following parameters, adapted with slight variations from Gordon et al. (2016) and Naldi and Flamini (2017). 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*. Finally, $\alpha = 1.5 * 10^{-5}$ 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 $\alpha = 1 * 10^{-5}$ 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[14] solutions to the system of equations in (11) 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.

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

5. Discussion

5.1. Model Simulation

5.1.1. 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 9 in Figure 2. The logarithmic utility function is simply $u(.) = ln(.)$ while the exponential function is Equation 17 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 not exceed the imposed constraint. The results broadly suggest that utility is largely maximised at maximum coverage for most reasonable insurance premium rates - the only counterexample in the analysis is for $p > 0.35$ for the Class I SBF with logarithmic utility. The conclusion that maximum coverage is optimal concurs with the game theoretic modelling work in Johnson et al. (2011a) and Yang and Lui (2012).

5.1.2. 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 5.1.1 the maximum cover an insurance buyer would wish to take out is $C_{max} = \tau lS(v, z)$ with maximum cover available respecting the cash constraint is then given by $C = \frac{(1/e)\tau vl - z}{p}$.

Figure 3 shows the variation of maximum available cover subject to the cash 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 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.

5.1.3. 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 optimal amount recommended by the GL model. Figure 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 6 plots the utility functions previously described for buying insurance at the maximum coverage available (as described in section 5.1.2) 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 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 $z^{II*}(v)$.

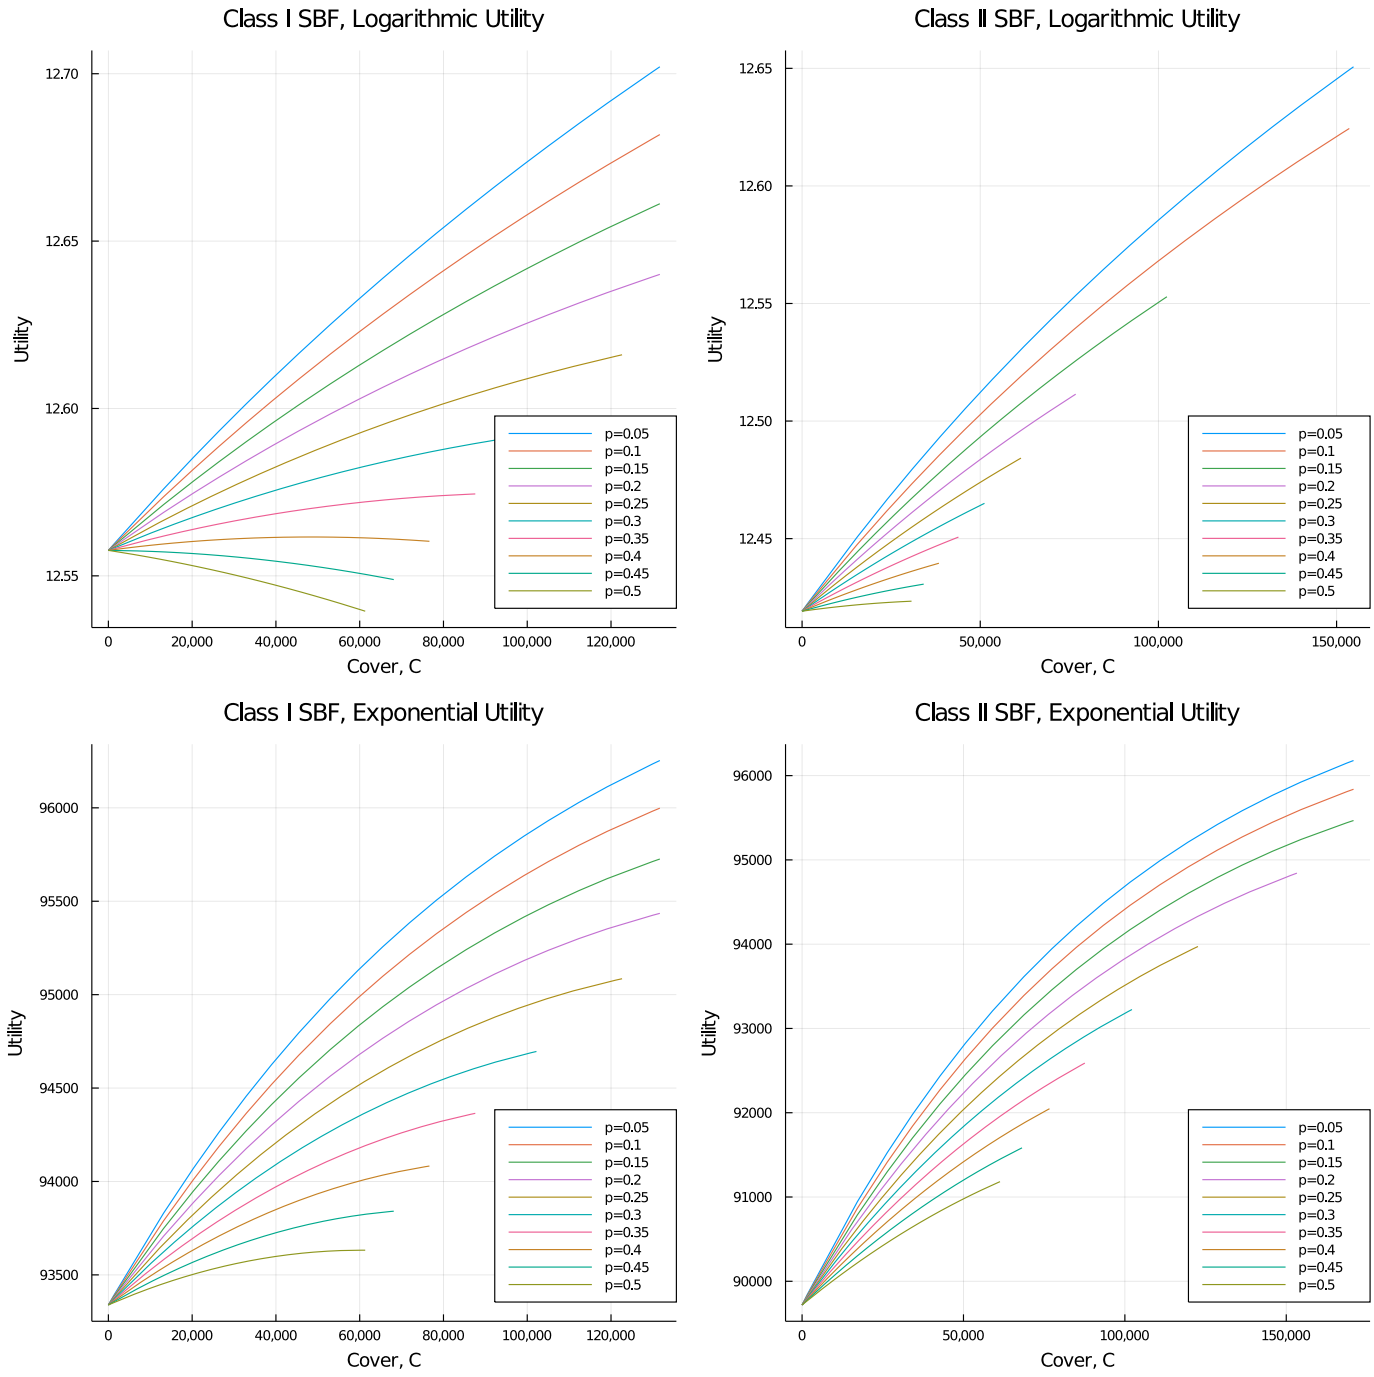

Figure 2: Utility as a function of cover assuming $z = z^*$

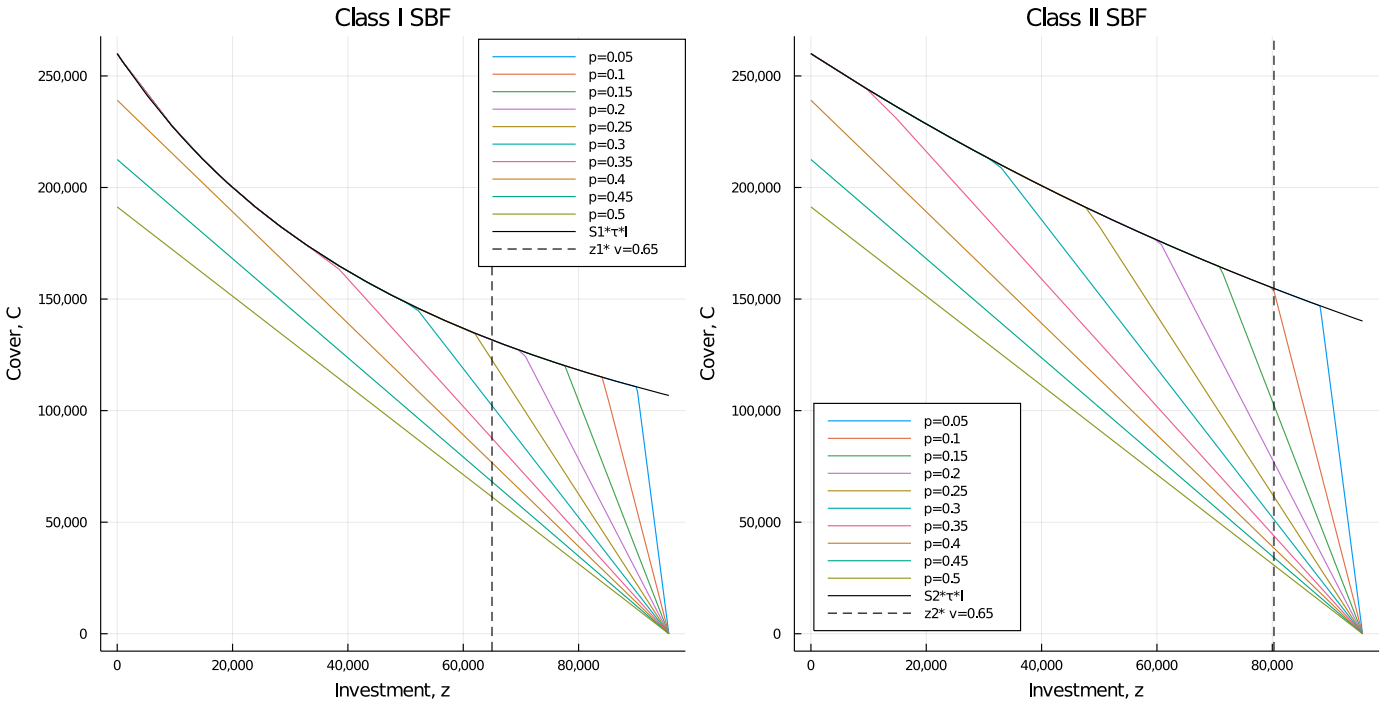

Figure 3: Maximum available cover under the cash constraint at different levels of z

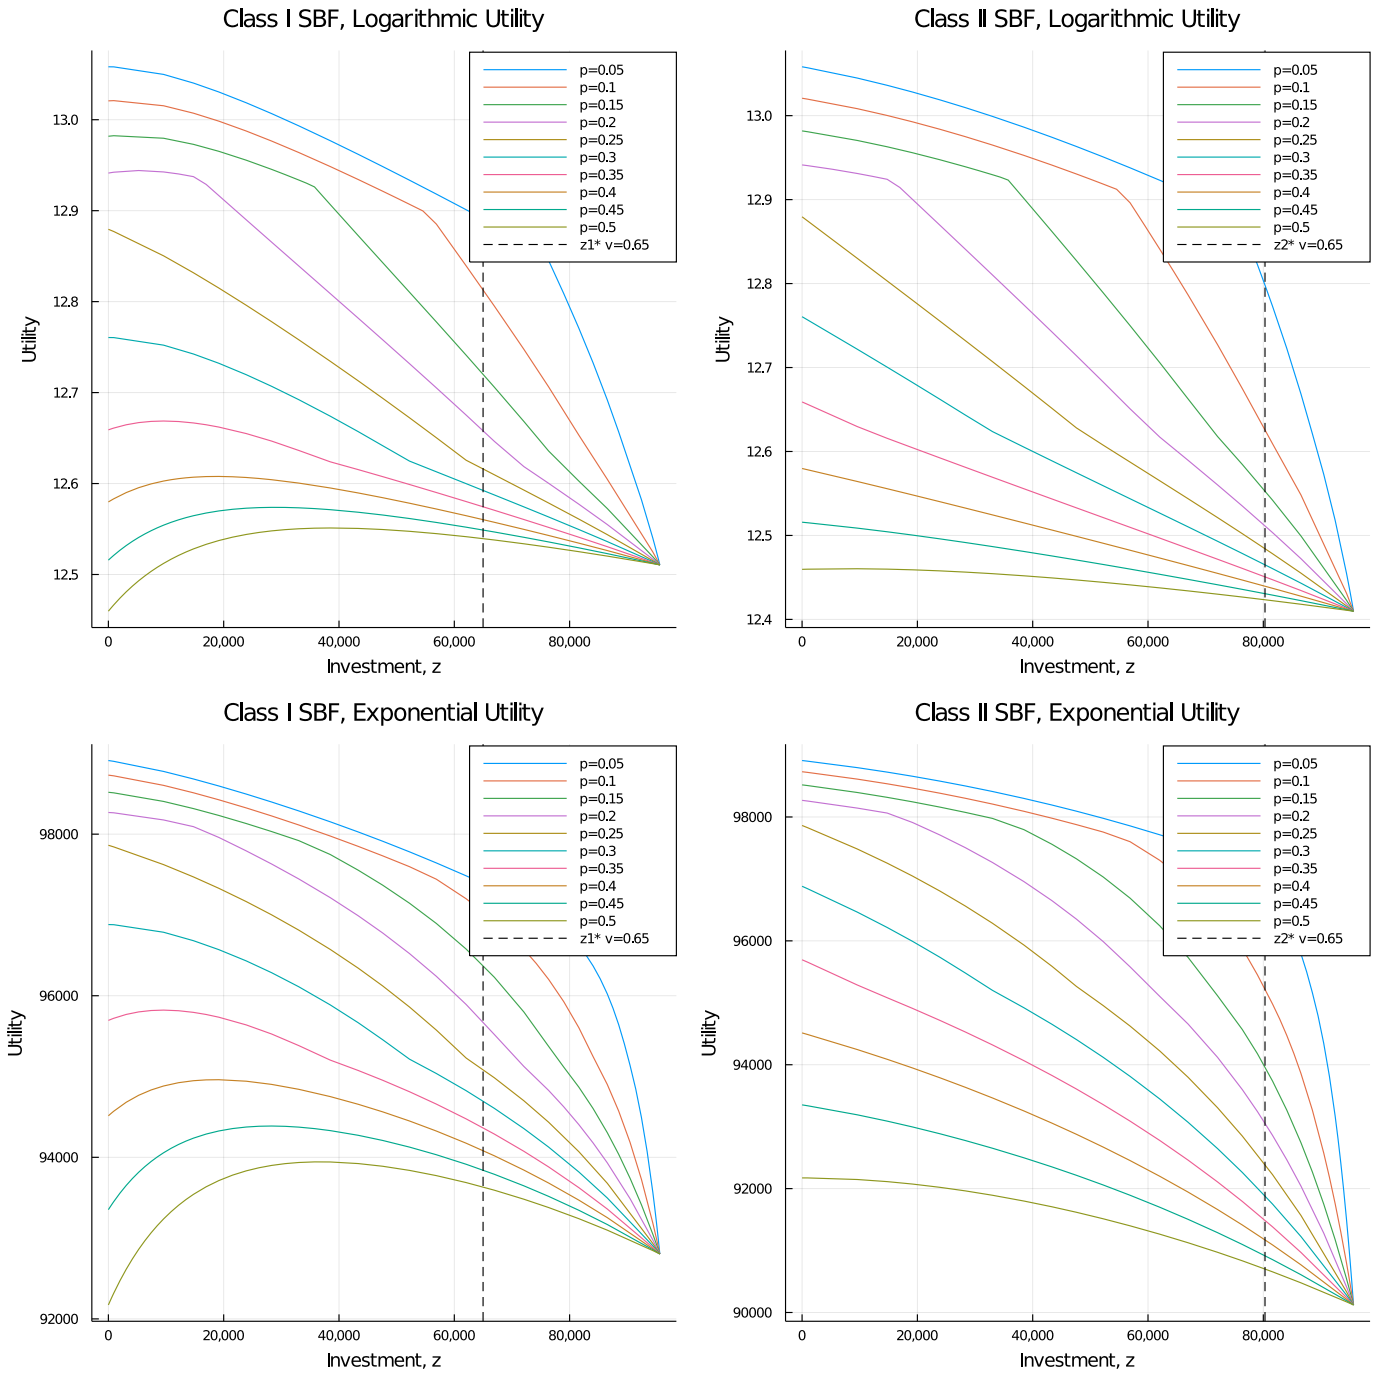

Figure 4: Utility as a function of investment with maximum insurance coverage purchased

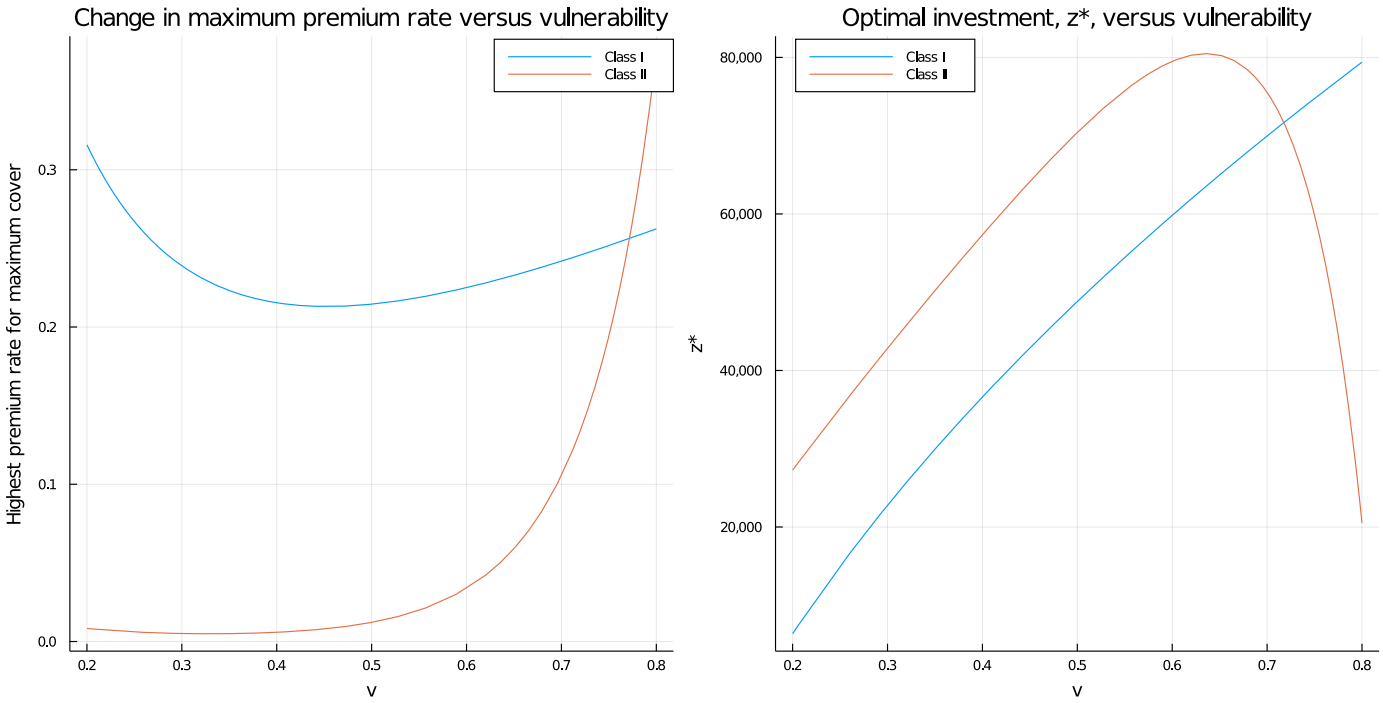

Figure 5: Highest premium rate at which maximum cover may be obtained for vulnerability, v and variation of optimal GL investment with vulnerability, v

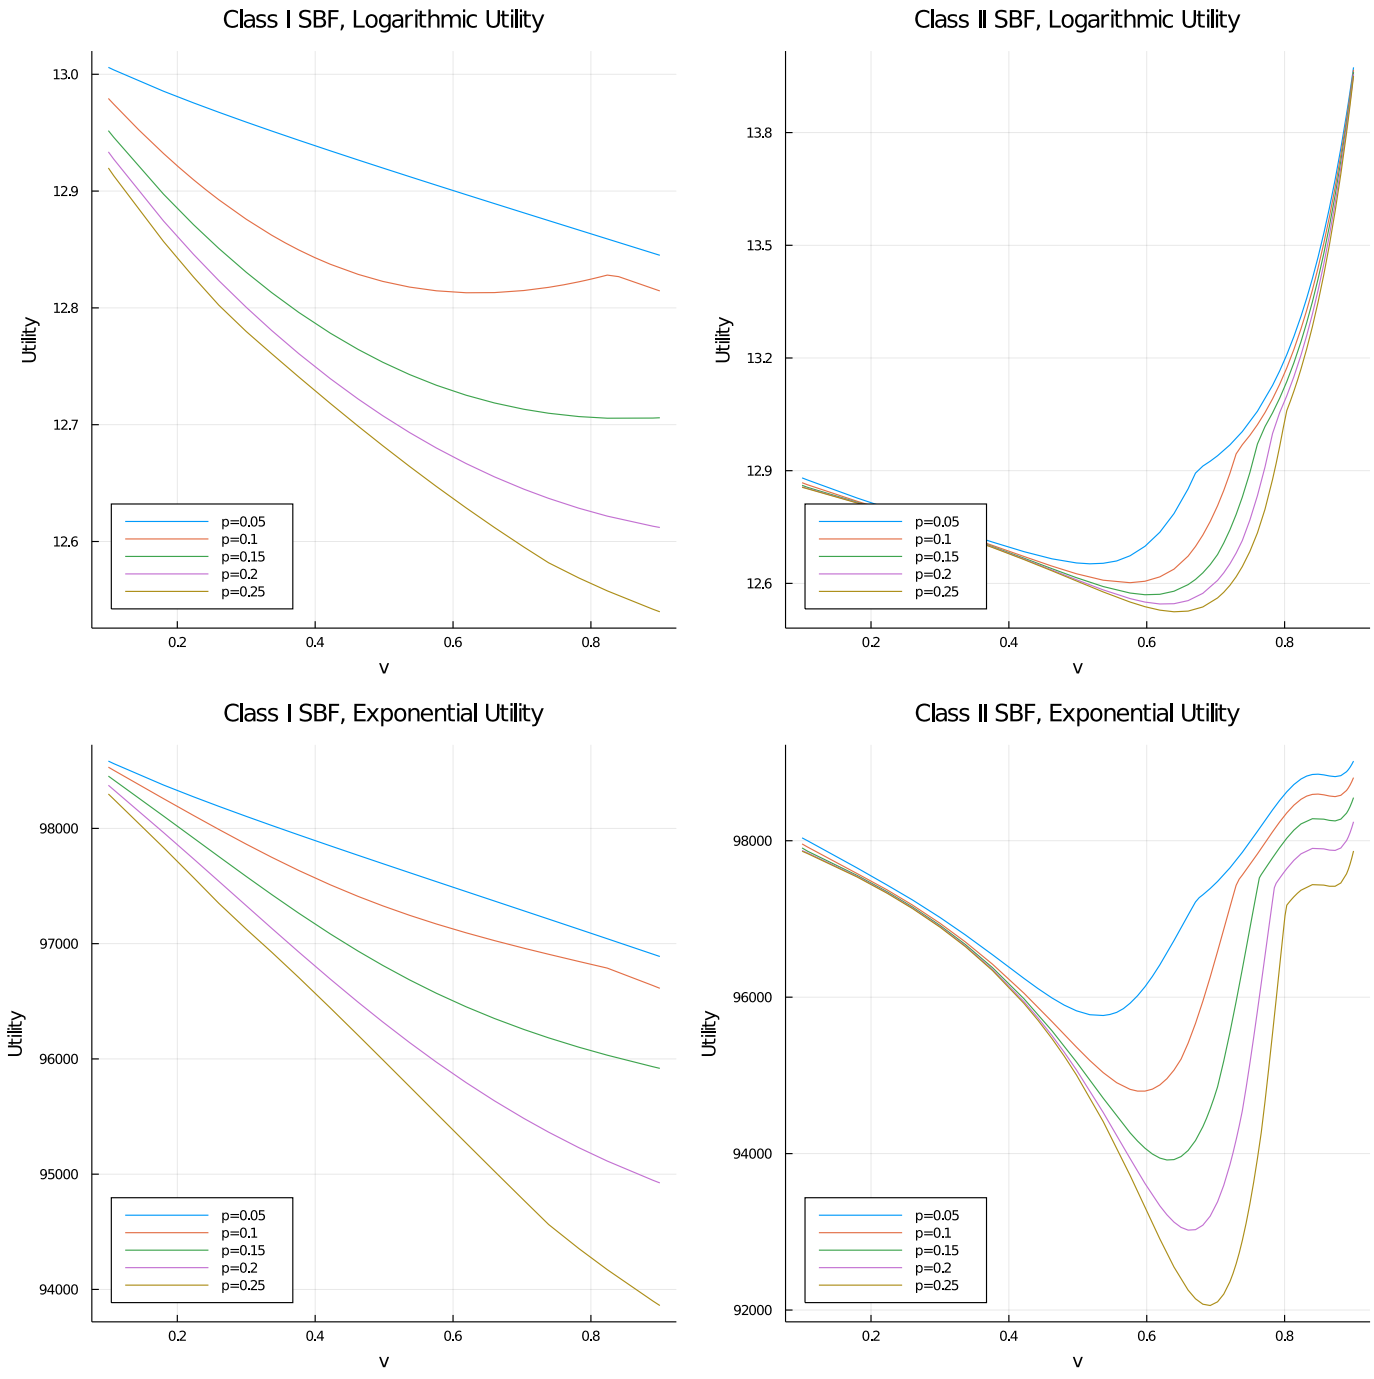

Figure 6: Utility functions for different vulnerabilities assuming investment at the Gordon-Loeb optimum, z^* and maximum coverage respecting the model cash constraint

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 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.

5.2. Model limitations

We believe that the GLCI model meets the initial research objective of assessing whether an expansion of the Gordon-Loeb model can yield useful insights for cyber-insurance, as it demonstrates that the Gordon-Loeb security breach functions can be used within a classical two-state utility maximisation model. 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 optimisation problem, but the simulation approach in this paper 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 organisation 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'[15] 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. (2019) 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). 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 organisation 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

[15] For a survey of early work in this area, see Bringer et al. (2012). Moore (2016) and Tsikerdekis et al. (2018) are also interesting examples of this area of research.

sector to share such information, and indeed to do so might be considered economically irrational (albeit potentially socially responsible). 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 5 and Table 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.

6. Conclusion

This work 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.

7. Acknowledgements

First and foremost to my supervisors, David Pym and Christos Ioannidis at UCL, for their constant support and encouragement. Madeline Carr and Shane Johnson also provided highly valuable feedback, especially on the literature review. Lawrence Gordon and Martin Loeb were extremely generous in response to my early questions related to the research presented in this publication and provided very helpful related work to consider. The anonymous reviewers provided valuable feedback and observations, which greatly improved the paper. This work was supported, in part, by the Engineering and Physical Research Council grant for Doctoral Training EP/R513143/1.

References

Akerlof, G., 1970. The market for "lemons": Quality uncertainty and the market mechanism. Quarterly Journal of Economics 84, 488–500.

Allais, M., 1953. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. Econometrica 21, 503–546. URL: http://www.jstor.org/stable/1907921.

Anderson, R., 2001. Why information security is hard - an economic perspective, in: Seventeenth Annual Computer Security Applications Conference, IEEE. pp. 358–365.

Arrow, K., . Aspects of the theory of risk-bearing.

Arrow, K.J., 1974. Optimal insurance and generalized deductibles. Scandinavian Actuarial Journal 1974, 1–42. URL: http://www.tandfonline.com/doi/abs/10.1080/03461238.1974.10408659.

Bandyopadhyay, T., Mookerjee, V., 2019. A model to analyze the challenge of using cyber insurance. Information Systems Frontiers 21, 301–325. doi:10.1007/s10796-017-9737-3.

Baryshnikov, Y., 2012. It security investment and gordon-loeb's 1/e rule., in: Workshop on the Economics of Information Security 2012.

Biener, C., Eling, M., Wirfs, J.H., 2015. Insurability of cyber risk: An empirical analysis. Geneva Papers on Risk and Insurance-Issues and Practice 40, 131–158. doi:10.1057/gpp.2014.19.

Bodin, L.D., Gordon, L.A., Loeb, M.P., Wang, A., 2018. Cybersecurity insurance and risk-sharing. Journal of Accounting and Public Policy 37, 527–544. doi:10.1016/j.jaccpubpol.2018.10.004.

Boehme, R., 2010. Security Metrics and Security Investment Models. Information Security and Privacy. p. 10–24.

Bohme, R., Kataria, G., 2006. On the limits of cyber-insurance. volume 4083 of *Lecture Notes in Computer Science*. pp. 31–40.

Bojanc, R., Jerman-Blazic, B., 2008. An economic modelling approach to information security risk management. International Journal of Information Management 28, 413–422. doi:10.1016/j.ijinfomgt.2008.02.002.

Borch, K., 1967. The economic theory of insurance - notes for an informal discussion in edinburgh 1 june 1964. ASTIN Bulletin 4, 252–264.

Borch, K., 1981. Is regulation and supervision of insurance companies necessary? Scandinavian Actuarial Journal 1981, 179–190. URL: http://www.tandfonline.com/doi/abs/10.1080/03461238.1981.10432017.

Bringer, M.L., Chelmecki, C.A., Fujinoki, H., 2012. A survey: Recent advances and future trends in honeypot research. International Journal of Computer Network and Information Security 4, 63.

Eling, M., Schnell, W., 2016. What do we know about cyber risk and cyber risk insurance? Journal of Risk Finance 17, 474–491. doi:10.1108/jrf-09-2016-0122.

Eling, M., Zhu, J., 2018. Which insurers write cyber insurance? evidence from the u.s. property and casualty insurance industry. Journal of Insurance Issues 41, 22–56. doi:10.2307/26441191.

Ellsberg, D., 1961. Risk, ambiguity, and the savage axioms. The Quarterly Journal of Economics 75, 643–669. URL: https://doi.org/10.2307/1884324, doi:10.2307/1884324.

Elnagdy, S.A., Qiu, M., Gai, K., 2016. 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. doi:10.1109/CSCloud.2016.46.

Fahrenwaldt, M.A., Weber, S., Weske, K., 2018. Pricing of cyber insurance contracts in a network model. Astin Bulletin 48, 1175–1218. doi:10.1017/asb.2018.23.

Farrow, S., Szanton, J., 2016. Cybersecurity investment guidance: Extensions of the gordon and loeb model. Journal of Information Security 07, 15—-28. doi:10.4236/jis.2016.72002.

Finsinger, J., Pauly, M., 1984. Reserve levels and reserve requirements for profit-maximizing insurance firms, in: Foundations of Insurance Economics. Springer, pp. 685–704.

Gai, K., Qiu, M., Elnagdy, S.A., 2016. A Novel Secure Big Data Cyber Incident Analytics Framework for Cloud-Based Cybersecurity Insurance. 2016 Ieee 2nd International Conference on Big Data Security on Cloud, pp. 171–176. doi:10.1109/BigDataSecurity-HPSC-IDS.2016.65.

Gass, S.I., Fu, M.C. (Eds.), 2013. Karush-Kuhn-Tucker (KKT) Conditions. Springer US, Boston, MA. pp. 833–834. URL: https://doi.org/10.1007/978-1-4419-1153-7_200359, doi:10.1007/978-1-4419-1153-7_200359.

Gollier, C., 2001. The economics of risk and time / Christian Gollier. MIT Press, Cambridge, Mass. ; London.

Gordon, L., Loeb, M., 2002. The economics of information security investment. ACM Transactions on Information and System Security (TISSEC) 5, 438–457.

Gordon, L.A., Loeb, M.P., Lucyshyn, W., Zhou, L., et al., 2014. Externalities and the magnitude of cyber security underinvestment by private sector firms: a modification of the gordon-loeb model. Journal of Information Security 6, 24.

Gordon, L.A., Loeb, M.P., Sohail, T., 2003. A framework for using insurance for cyber-risk management. Communications of the Acm 46, 81–85. doi:10.1145/636772.636774.

Gordon, L.A., Loeb, M.P., Zhou, L., et al., 2016. Investing in cybersecurity: insights from the gordon-loeb model. Journal of Information Security 7, 49.

Greenberg, A., 2019. Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers. Doubleday.

Hausken, K., 2006. Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability. Information Systems Frontiers 8, 338.

Hayel, Y., Zhu, Q., 2015. Attack-Aware Cyber Insurance for Risk Sharing in Computer Networks. volume 9406 of *Lecture Notes in Computer Science*. pp. 22–34.

Heath, B., 2018. Before the breach - the role of cyber insurance in incentivizing data security. George Washington Law Review 86, 1115–1151.

Hölmstrom, B., 1979. Moral hazard and observability. The Bell Journal of Economics 10, 74–91.

Huang, C.D., Behara, R.S., 2013. Economics of information security investment in the case of concurrent heterogeneous attacks with budget constraints. International Journal of Production Economics 141, 255–268. doi:10.1016/j.ijpe.2012.06.022.

Ioannidis, C., Pym, D., Williams, J., Gheyas, I., 2019. Resilience in information stewardship. European Journal of Operational Research 274, 638–653.

Johnson, B., Boehme, R., Grossklags, J., 2011a. Security Games with Market Insurance. volume 7037 of *Lecture Notes in Computer Science*. pp. 117–+.

Johnson, B., Boehme, R., Grossklags, J., 2011b. Security Games with Market Insurance. volume 7037 of *Lecture Notes in Computer Science*. p. 117.

Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica 47, 263–291. URL: http://www.jstor.org/stable/1914185.

Keegan, C., 2014. Cyber security in the supply chain - a perspective from the insurance industry. Technovation 34, 380–381. doi:10.1016/j.technovation.2014.02.002.

Khalili, M.M., Naghizadeh, P., Liu, M., 2018. Designing cyber insurance policies - the role of pre-screening and security interdependence. IEEE Transactions on Information Forensics and Security 13, 2226–2239.

Khalili, M.M., Naghizadeh, P., Liu, M., Ieee, 2017. Embracing Risk Dependency in Designing Cyber-Insurance Contracts. Annual Allerton Conference on Communication Control and Computing, pp. 926–933.

Laszka, A., Grossklags, J., 2015. Should Cyber-Insurance Providers Invest in Software Security?. volume 9326 of *Lecture Notes in Computer Science*. pp. 483–502.

Lee, K., 1992. Moral hazard, insurance and public loss prevention. Journal of Risk and Insurance (1986-1998) 59, 275. URL: http://search.proquest.com/docview/235946196/.

Lelarge, M., 2012. Coordination in network security games: a monotone comparative statics approach. IEEE Journal on Selected Areas in Communications 30, 2210–2219.

Linton, J.D., Boyson, S., Aje, J., 2014. The challenge of cyber supply chain security to research and practice - an introduction. Technovation 34, 339–341. doi:10.1016/j.technovation.2014.05.001.

Majuca, R.P., Yurcik, W., Kesan, J.P., 2006. The evolution of cyberinsurance. arXiv preprint cs/0601020 .

Marotta, A., Martinelli, F., Nanni, S., Orlando, A., Yautsiukhin, A., 2017. Cyber-insurance survey. Computer Science Review 24, 35–61. doi:10.1016/j.cosrev.2017.01.001.

Martinelli, F., Orlando, A., Uuganbayar, G., Yautsiukhin, A., 2017. Preventing the drop in security investments for non-competitive cyber-insurance market, in: International Conference on Risks and Security of Internet and Systems, Springer. pp. 159–174.

Martinelli, F., Orlando, A., Uuganbayar, G., Yautsiukhin, A., 2018. Preventing the Drop in Security Investments for Non-competitive Cyber-Insurance

Market. volume 10694 of *Lecture Notes in Computer Science*. pp. 159–174.
Massaccia, F., Swierzbinski, J., Williams, J., 2017. Cyberinsurance and public policy - self-protection and insurance with endogenous adversaries.
Matsuura, K., 2009. Productivity Space of Information Security in an Extension of the Gordon-Loeb's InvestmentModel. p. 99–119.
Mazzoccoli, A., Naldi, M., 2020. Robustness of optimal investment decisions in mixed insurance/investment cyber risk management. Risk Analysis 40, 550–564.
Moore, C., 2016. Detecting ransomware with honeypot techniques, in: 2016 Cybersecurity and Cyberforensics Conference (CCC), IEEE. pp. 77–81.
Mossin, J., 1968. Aspects of rational insurance purchasing. Journal of Political Economy 76, 553–568.
Munch, P., Smallwood, D., 1981. Theory of Solvency Regulation in the Property and Casualty Insurance Industy. The MIT Press. URL: http://www.nber.org/chapters/c11431.
Naldi, M., Flamini, M., 2017. Calibration of the gordon-loeb models for the probability of security breaches, in: 2017 UKSim-AMSS 19th International Conference on Computer Modelling Simulation (UKSim), pp. 135–140.
von Neumann, J., Morgenstern, O., Rubinstein, A., 1944. Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition). Princeton University Press. URL: http://www.jstor.org/stable/j.ctt1r2gkx.
Nieuwesteeg, B., Visscher, L., de Waard, B., 2018. The law and economics of cyber insurance contracts - a case study. European Review of Private Law 26, 371–420.
Nurse, J.R., Axon, L., Erola, A., Agrafiotis, I., Goldsmith, M., Creese, S., 2020. The data that drives cyber insurance: A study into the underwriting and claims processes .
Ogut, H., Raghunathan, S., Menon, N., 2011. Cyber security risk management - public policy implications of correlated risk, imperfect ability to prove loss, and observability of self-protection. Risk Analysis 31, 497–512. doi:10.1111/j.1539-6924.2010.01478.x.
Opadhyay, T.B., Mookerjee, V.S., Rao, R.C., 2009. Why it managers don't go for cyber-insurance products. Communications of the Acm 52, 68–73. doi:10.1145/1592761.1592780.
Pal, R., Golubchik, L., Ieee, 2010. Analyzing Self-Defense Investments in Internet Security Under Cyber-Insurance Coverage. IEEE International Conference on Distributed Computing Systems. doi:10.1109/icdcs.2010.79.
Pal, R., Golubchik, L., Psounis, K., 2011. Aegis A Novel Cyber-Insurance Model. volume 7037 of *Lecture Notes in Computer Science*. pp. 131–150.
Pal, R., Golubchik, L., Psounis, K., Hui, P., Ieee, 2014. Will Cyber-Insurance Improve Network Security? A Market Analysis. Ieee Infocom, pp. 235–243.
Pal, R., Hui, P., 2013. On Differentiating Cyber-Insurance Contracts A Topological Perspective. 2013 Ifip/Ieee International Symposium on Integrated Network Management.
Pavlik, L., Ieee, 2018. 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. doi:10.1109/mcsi.2018.00036.
Piromsopa, K., Klima, T., Pavlik, L., Ieee, 2017. Designing model for calculating the amount of cyber risk insurance. 2017 Fourth International Conference on Mathematics and Computers in Sciences and in Industry. doi:10.1109/mcsi.2017.41.
Pratt, J.W., 1964. Risk aversion in the small and in the large. Econometrica 32, 122–136.
Raviv, A., 1979. The design of an optimal insurance policy. The American Economic Review 69, 84. URL: http://search.proquest.com/docview/233054428/.
Rees, R., Wambach, A., 2008. The microeconomics of insurance. Foundations and Trends® in Microeconomics 4, 1–163. URL: http://dx.doi.org/10.1561/0700000023, doi:10.1561/0700000023.
Rios Insua, D., Couce-Vieira, A., Musaraj, K., 2018. Some risk analysis problems in cyber insurance economics. Estudios De Economia Aplicada 36, 181–194.
Romanosky, S., Ablon, L., Kuehn, A., Jones, T., 2017. Content analysis of cyber insurance policies - how do carriers write policies and price cyber risk?, in: Workshop on the Economics of Information Security 2017.
Rothschild, M., Stiglitz, J., 1976. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. The Quarterly Journal of Economics 90, 629–649. URL: http://www.jstor.org/stable/1885326.
Saini, D.K., Azad, I., Raut, N.B., Hadimani, L.A., 2011. Utility Implementation for Cyber Risk Insurance Modeling. Lecture Notes in Engineering and Computer Science, pp. 429–432.
Sanger, D.E., 2019. The perfect weapon: War, sabotage, and fear in the cyber age. Broadway Books.
Shackelford, S.J., 2012. Should your firm invest in cyber risk insurance? Business Horizons 55, 349–356. doi:10.1016/j.bushor.2012.02.004.
Shetty, N., Schwartz, G., Felegyhazi, M., Walrand, J., 2010. Competitive Cyber-Insurance and Internet Security. Economics of Information Security and Privacy.
de Smidt, G., Botzen, W., 2018. Perceptions of corporate cyber risks and insurance decision-making. Geneva Papers on Risk and Insurance-Issues and Practice 43, 239–274. doi:10.1057/s41288-018-0082-7.
Sonnenreich, W., Albanese, J., Stout, B., 2006. Return on security investment (rosi)-a practical quantitative model. Journal of Research and practice in Information Technology 38, 45–56.
Spence, M., 1973. Job market signaling. The Quarterly Journal of Economics 87, 355–374.
Talesh, S.A., 2018. Data breach, privacy, and cyber insurance - how insurance companies act as compliance managers for businesses. Law and Social Inquiry-Journal of the American Bar Foundation 43, 417–440. doi:10.1111/lsi.12303.
Tatsumi, K.i., Goto, M., 2010. Optimal timing of information security investment: A real options approach, in: Economics of Information Security and Privacy. Springer, pp. 211–228.
Tondel, I.A., Seehusen, F., Gjaere, E.A., Moe, M.E.G., 2016. Differentiating Cyber Risk of Insurance Customers - The Insurance Company Perspective. volume 9817 of *Lecture Notes in Computer Science*. pp. 175–190.
Tsikerdekis, M., Zeadally, S., Schlesener, A., Sklavos, N., 2018. Approaches for preventing honeypot detection and compromise, in: 2018 Global Information Infrastructure and Networking Symposium (GIIS), IEEE. pp. 1–6.
US Department of Justice, 2002. London, england hacker indicted under computer fraud and abuse act for accessing military computers. last accessed: 21 october 2021. URL: https://www.justice.gov/archive/criminal/cybercrime/press-releases
Usher, D., 1987. Theft as a paradigm for departures from efficiency. Oxford Economic Papers 39, 235–252. URL: http://www.jstor.org/stable/2663263.
Willemson, J., 2006. On the gordon and loeb model for information security investment, in: WEIS Proceedings 2006. URL: https://www.econinfosec.org/archive/weis2006/docs/12.pdf.
Wilson, C., 1977. A model of insurance markets with incomplete information. Journal of Economic Theory 16, 167–207.
Woods, D., Agrafiotis, I., Nurse, J.R.C., Creese, S., 2017. Mapping the coverage of security controls in cyber insurance proposal forms. Journal of Internet Services and Applications 8. doi:10.1186/s13174-017-0059-y.
Woods, D., Simpson, A.C., 2018. Monte carlo methods to investigate how aggregated cyber insurance claims data impacts security investments.
Xu, M., Hua, L., 2019. Cybersecurity insurance - modeling and pricing. North American Actuarial Journal 23, 220–249. doi:10.1080/10920277.2019.1566076.
Yang, Z., Lui, J.C.S., 2012. Security Adoption in Heterogeneous Networks - the Influence of Cyber-Insurance Market. volume 7290 of *Lecture Notes in Computer Science*. pp. 172–183.
Yang, Z., Lui, J.C.S., 2014. Security adoption and influence of cyber-insurance markets in heterogeneous networks. Performance Evaluation 74, 1–17. doi:10.1016/j.peva.2013.10.003.
Young, D., Lopez, J., Rice, M., Ramsey, B., McTasney, R., 2016. A framework for incorporating insurance in critical infrastructure cyber risk strategies. International Journal of Critical Infrastructure Protection 14, 43 – 57. URL: http://www.sciencedirect.com/science/article/pii/S1874548216300439 doi:https://doi.org/10.1016/j.ijcip.2016.04.001.
Zhang, R., Zhu, Q., Hayel, Y., 2017. A bi-level game approach to attack-aware cyber insurance of computer networks. Ieee Journal on Selected Areas in Communications 35, 779–794. doi:10.1109/jsac.2017.2672378.