Monte Carlo methods to investigate how aggregated cyber insurance claims data impacts security investments

Daniel W. Woods and Andrew C. Simpson
Department of Computer Science, University of Oxford
Wolfson Building, Parks Road, Oxford OX1 3QD, UK
firstname.lastname@cs.ox.ac.uk

Abstract

There is evidence that the availability of cyber insurance is contingent on an applicant's security posture and that premium discounts may apply if the applicant adopts security controls dictated by the insurer. As the cyber insurance market grows in size, questions arise regarding how this situation will affect investment in information security. We investigate how aggregated claims data impacts investments in information security. Monte Carlo methods are used to explore three possible insurer strategies in guiding the policyholder's investments. The results suggest that aggregated claims data can increase the net revenue of all firms, particularly in cases of low security investment or high uncertainty, but these benefits are contingent on the insureds employing diverse defensive configurations.

1 Introduction

Beginning in 2012, policy makers in the US and the EU began to investigate how the cyber insurance industry might drive improvements in cyber security [1]. This attention supported early predictions that the insurance industry would impact information security investment [2]. Insurers could affect security investments decisions by sharing information about cyber attacks, by offering premium discounts to encourage security investments, and even by demanding certain security controls be in place to obtain coverage [1]. Some believe that widespread adoption could mean that "good security [is] rewarded in the marketplace" [2]. Any impact that the insurance industry might have will be compounded by recent forecasts suggesting that global revenues could grow beyond $2.5 billion[1] today (late 2017) to $7.5 billion by 2020[2] and to $14 billion by 2022[3].

Thus far, research investigating the impact of the insurer in this role has been limited. Open questions include: *which strategies will the market structure allow the insurer to pursue?*; *what is the optimal strategy for the insurer to pursue?*; and *are the insured's interests protected in pursuing such a strategy?* Such questions seek to understand how a rational insurer might operate and then consider the impact on the security posture of the insured. This line of thought speaks to Ross Anderson's suggestion that security researchers should ask "what's the source of market power?" [3].

[1] `https://www.pwc.com/gx/en/industries/financial-services/insurance/publications/insurance-2020-cyber.html`

[2] `https://www.pwc.com/gx/en/industries/financial-services/insurance/publications/insurance-2020-cyber.html`

[3] `https://www.alliedmarketresearch.com/press-release/cyber-insurance-market.html`

The ability of the insurer to observe claims across multiple policyholders is one such power. When combined with the risk assessment conducted during the application process, insurers may be able to link defensive posture (the quasi-independent variable) to economic losses (the quasi-dependent variable). Throughout this paper, *claims information* is defined to be the collection of insights regarding the effectiveness of defensive measures provided by these quasi-experiments. Insurers may choose to share claims information with their insureds.

We consider how three market strategies an insurer might adopt can impact the policyholder: the *passive approach* does not share claims information; the *active approach* shares claims information about insureds with different security levels; and the *diverse approach* shares claims information while deliberately maintaining diversity in the security posture of the insureds. Our aim is to explore how each strategy affects both the insurer and the policyholder, paying particular attention to whether their interests align. The contributions potentially have relevance to government policy discussions, which have identified that cyber insurers already share insights from historic claims [1].

In Section 2, we outline how our contribution relates to research into investment models, as well as considering existing work looking at the role of the insurer. Section 3 introduces the Iterated Weakest Link (IWL) [4] investment model, which we extend to the context of cyber insurance in Section 4. Section 5 simulates the results of the insurer adopting three different strategies. Section 6 discusses the context in which our results would be relevant and how they might be taken forward. Section 7 concludes the paper.

2 Related Work

To evaluate research into how the insurance industry impacts information security investments, it is helpful to distinguish between empirical and theoretical work. The former includes collecting contractual documents, interviewing insurers and analysing market-level data, whereas the latter is concerned with models to investigate pricing, simulate market dynamics and estimate correlated risk.

In recent years, the growth of the market has created opportunities for empirical work. In 2015, Biener et al. [5] extracted 994 cases of cyber losses from an operational risk database and asked whether they were insurable losses. A number of authors have directly studied insurance policy wordings to understand what coverage is available [6–9]. These studies ask what might be or what is insurable, without considering how insurance impacts information security.

Franke [10] interviewed 15 insurance professionals based in Sweden. The findings suggest that the insurer plays a role in guiding the policyholder towards adopting security controls, which supports the assumption that purchasing an insurance contract changes how information security is managed. However, a more detailed understanding and evaluation of the impact on security investment was deemed to be out of scope.

Studies [1, 8] of insurance application forms have identified which areas of information security the insurers collect information about. However, Romanosky et al. [8] suggest that only 33% of the US insurers they analysed consider information security when pricing risk, which undermines the conclusions from application form analysis.

Empirical work has identified what could be covered [5], what is covered [7,8,11], the attitudes and practices of the insurance professionals involved [10], what security information is collected [8, 12], and even how a price is calculated from the information collected [8]. Questions remain regarding how each of these aspects affects the market and the corresponding effect on the insured's investment in information security.

Theoretical models may be better suited to answering such questions. This body of work sits

within the field of information security economics, which was founded upon the realisation that misplaced incentives can help explain why many security systems fail [13]. In this vein, several authors [6, 14] have conclude that insurers offering reduced premiums provide incentives for security investment, which corroborates Schneier's early predictions [2].

There have been many attempts to model different aspects of the insurance market. A unifying framework is provided in [15], which draws a distinction between two aspects of the market. The first such aspect is concerned with how security investments accrue benefits to all parties in a system, not just the investor — particularly how these positive externalities can reduce the risk an insurer faces [16–18]. Second, there have been various considerations of systemic risk, in which many firms make insurance claims arising from the same event because of the interdependency of networks.

In [19], Böhme investigates cyber insurance from the perspective of the insurer and identifies the correlated nature of cyber risk as an important consideration. In a subsequent paper [20], Böhme and Kataria suggest that cyber insurance is most suitable for risks with high internal and low global correlation. The authors conclude system managers should emphasise platform diversity — a theme to which we will return in Section 6. However, they do not consider how the insurance industry impacts investments in information security

Rather than extend an insurance model to include security posture, we could instead extend an existing model of cyber risk in the context of the insurance industry. We look to investment models that ask *how much should an organisation invest in information security?* Gordon and Loeb [21] conclude that the answer depends on the shape of the 'security breach function'. Determining which vulnerability to address involves comparing each vulnerability's breach function. Unfortunately, [21] does not contribute a simple method for identifying the security breach function of a given vulnerability.

Böhme [22] suggests that a problem with the Gordon and Loeb model is the direct mapping of security investment to vulnerability. Instead, a direct mapping from investment to 'security level', which then stochastically maps to 'benefits of security', is suggested. The stochastic properties are a result of the indeterminacy of the attacker behaviour, which is assumed to be constant in Gordon and Loeb's model.

Heitzenrater and Simpson [23] link security investments to (prevented) losses using the 'Information Security Breaches Survey' [24], an annual survey to assess breaches in UK organisations. The approach uses the survey's loss data as a baseline and then considers how further investment in various security products affects Annual Loss Expectancy (ALE), contributing a number of decision recommendations that are dependent on the size of the business.

Both [21] and [23] consider a one-time investment that cannot change in response to observed attacks. Böhme and Moore's Iterated weakest link (IWL) model [4] allows a defender to adopt a reactive investment strategy. The attacker's behaviour is modelled stochastically and the defender reduces uncertainty as this behaviour is observed. By introducing a parameter for uncertainty, the model provides a rational explanation for perceived under-investment in security — the defender is adopting a "wait and see" approach.

The IWL model provides three useful features for modelling the cyber insurance market: (i) an appropriate format for security posture; (ii) the uncertainty parameter; and (iii) a temporal dimension. For (i), the binary choice between protecting a given vulnerability or not is in keeping with the cyber insurance application process [12], which predominantly consists of yes–no questions. This can be contrasted with models that assume continuous investment, such as [21]. For (ii), the interviews [10] and pricing decisions [8] have identified that insurers lack information about how security controls relate to cyber risk. Indeed, uncertainty is one of the main challenges insurers face [25]. This can be contrasted with Gordon and Loeb's model, which assumes perfect information about the threat and vulnerability. Finally, (iii) allows one to model the insurer's ability to observe attacks

in claims data and share this information with policyholders.

We now introduce the IWL model in more detail.

3 The Iterated Weakest Link Model

There are three aspects to the Iterated Weakest Link model [4]: the rules, the strategy adopted, and the computation. The rules define whether an attack will take place and the defender's utility conditional on that attack. The strategy determines the defender's choice of defensive configuration across multiple rounds. Finally, the computation involves calculating the expected utility for adopting each strategy. We consider each aspect in turn.

3.1 The Rules

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

$$R_i = ra - p_i za - c_i \tag{1}$$

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

$$x_j = \sup(0, \chi_j) \text{ where } \chi_j \sim \mathcal{N}(\bar{x}_j, \frac{\sigma}{\Delta x}) \text{ for } j = 1, ..., n \tag{2}$$

The defensive configuration $\mathbf{d}_i \in \{0,1\}^n$ describes whether the true cost associated with each vulnerability $x_1, ..., x_n$ is protected or not. We denote the k-th defense by d_k. The cost of employing the defensive configuration $\mathbf{d}_i$ is determined by an $n \times n$ matrix $\mathbf{C}$ that reflects "possible interdependent defenses" [4], so that $c_i = \mathbf{d}_i \mathbf{C} \mathbf{d}_i$. The matrix is set such that

$$c_i = \frac{\rho}{2}k^2 + (1 - \frac{\rho}{2})k \quad \text{where } k = \sum_{i=0}^{n} d_i \tag{3}$$

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

Figure 1 describes a situation in which the defender following the expected cost would result in the first three vulnerabilities being defended. Yet, the fourth and fifth would be unguarded and economically exploitable, contrary to the defender's expectations.

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

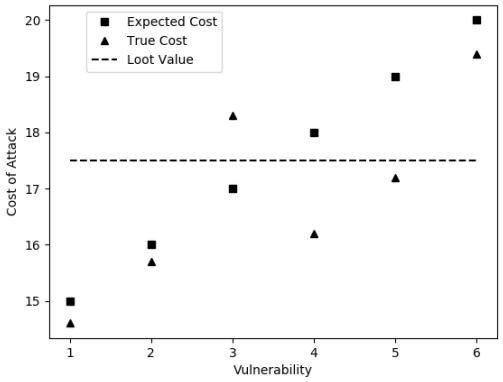

Figure 1: The true costs of exploiting the first, second, fourth and fifth vulnerabilities fall below the loot value, za. As a result, there will be an attack unless these defenses are in place.

3.2 The Strategy

Böhme and Moore [4] assume that the defender is rational and will employ the defensive configuration that maximises revenue. The defender is assumed to be risk-neutral. For a one-round game, this reduces to choosing the defensive configuration with the highest expected utility. For a multiple round game, we must consider how the defensive configuration will be adapted in light of information about attacks.

Böhme and Moore [4] suggest that d_k should not be "tinkered with" once a defense is in place. Otherwise "the direct and indirect ($\rho > 0$) cost of the l-th defense has to be borne for all intermediate rounds" [4]. If a vulnerability x_l is attacked, then the defender gains the information that it is economically viable. The optimal strategy chooses an initial defensive configuration and then places reactive defenses in place only when a vulnerability is exploited.

If the cost of putting additional defenses in place becomes too large, it may be rational to accept future attacks and dis-invest to zero defenses. Additionally, if sunk costs are particularly high, the optimal strategy may involve accepting future attacks without changing the defensive configuration. The next subsection describes how the expected utility for each starting configuration is calculated.

3.3 Calculation

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

$$P(\text{Attack}) = 1 - \prod_{x_k \in B} P(x_k > z) \tag{4}$$

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

In the one-round case, an attack takes place if the true cost of attack x_j of at least one unguarded vulnerability falls below the loot value za, whereas the total number of economically viable unguarded vulnerabilities is important in the multi-round or dynamic case.

Böhme and Moore [4] model each unguarded defense x_j as a Bernoulli random variable. Each has probability of being exploited, given by

$$P(za > x_j) = \Phi(za; x_1 + (j-1)\Delta x, \frac{\sigma}{\Delta x}) \tag{5}$$

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

$$\mu = \sum_{i=k+1}^{n} p_i \text{ and } \sigma = \sum_{i=k+1}^{n} p_i(1-p_i) \tag{6}$$

Böhme and Moore approximate this via $\mathcal{N}(\mu, \sigma)$, which is justified "for suitable parameter choices" [4]. Considering t_{max} rounds, there will be $t_{att} \leq t_{max}$ rounds in which the attacker is successful. The total revenue, for $t_{att} = i$ and initial defensive configuration k, is determined by

$$R(t_{att}, k) = \sum_{t=1}^{t_{att}} (ra - z - c_t) + \sum_{t=t_{att}+1}^{t_{max}} (ra - c_{t_{att}}) \tag{7}$$

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

$$U(k) = \sum_{i=0}^{t_{max}} P(t_{att} = i) R(t_{att} = i, k) \tag{8}$$

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

To summarise, the IWL is a stochastic model that captures a defender interacting under uncertainty with an attacker over multiple rounds. The best strategy for a defender is to only deviate from an initial defensive configuration in response to observed attacks. These provide information related to the true cost of defense of the exploited vulnerability. The true cost of attack is modelled by random realisations of a normal distribution; to simplify calculating expected revenue, an approximation is used to calculate the likelihood of each realisation of these costs [4]. The next section introduces how we extend the IWL to consider multiple policyholders making security decisions.

4 Extending the Model

Our extension, the Iterated Weakest Link – Cyber Security Insurance (IWL-CSI) introduces new rules, new strategies and a new method of computation in extending the original model to consider m policyholders purchasing insurance from a single insurer. In Section 4.1, we introduce those new rules. We outline how the passive, active and diverse approaches translate strategies for the IWL-CSI in Section 4.2. In Section 4.3, we demonstrate how Monte Carlo methods can be used to simulate the IWL-CSI.

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

Table 1: Describing each of the parameters in the motel

4.1 New rules

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

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

To understand (i), consider that network effects empower software monopolies [26], which leads to a lack of so-called "cyber diversity" [27]. The result of policyholders adopting similar information systems is that they share vulnerabilities. Consequently, if the insurer learns that the i-th vulnerability is used to attack one insured, then that same vulnerability will be economically viable in other insured's systems.

We assume (ii) because diverse defensive configurations allow information about the attacker to be shared and collective uncertainty reduced. For example, different organisations protect the same operating system with different security products. The insurer is assumed to have control over which defensive configurations are in place for each policyholder. Insurers have expressed a "desire" to recommend security controls [1] and some insurers even offer "a list of actions to be taken" [10] to improve security. Of course, this ability will be dependent on the wider market; an under-supply of insurance allows the insurer greater freedom to select with whom they enter into contract.

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

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

Table 2: Illustration of the different strategies with $m = 3$.

Defensive costs related to the interdependence of controls employed by an organisation are determined by an $n \times n$ matrix C_{int} (as in the original IWL [4]). Additionally, we introduce an $m \times m$ matrix C_{ext_i} for each of the n possible defenses. These matrices reflect the extent to which defenders can co-operate to take advantage of returns to scale, which may vary depending on the particular defensive investment, d_{i_j}.

The cost to the j-th defender of employing the defensive configuration $\mathbf{d}_{i_j}$ is

$$C_j = \mathbf{d}_{i_j} C_{int} \mathbf{d}_{i_j} + \sum_{k=0}^{n} \mathbf{b}_k C_{ext_k} \mathbf{b}_k \tag{9}$$

where $\mathbf{b}_k \in \{0,1\}^m$ with $b_{k_t} = d_{k_t}$. The j-th element of $\mathbf{b}_k$ represents whether the j-th defender employs controls k. To remain in the scope of this study, the matrices C_{ext_i} are chosen so that defensive costs scale linearly in the number of insureds, calculated using Equation 3. In Section 6.3, we discuss how future work might modify this assumption to explore non-linear scaling.

As in [4], the attacker will exploit each defender's unguarded vulnerability with the lowest cost of attack, unless this is greater than the loot value, za. The true costs of attack are modelled as follows:

$$x_{i_j} = \sup(0, \chi_i) \text{ where } \chi_i \sim \mathcal{N}(\bar{x}_i, \frac{\sigma}{\Delta x}) \text{ for } j = 1, ..., m \tag{10}$$

Consequently, $x_{i_j} = x_{i_k}$ for all $i \in \{1, ..., n\}$ and $j, k \in \{1, ..., m\}$, where n and m are the number of vulnerabilities and insureds respectively. To identify the "weakest link", we must maintain an ordering on the set of vulnerabilities and use the untruncated values χ_i to do so. If more than one vulnerability had a true cost of 0, it is not clear which would be exploited first. For $x_{i_l} = 0 = x_{k_l}$, we say that i is the weakest link if $\chi_i < \chi_k$, where the values of χ_i and χ_k are determined by Equation 10.

4.2 Novel strategies

In this subsection we describe the strategies that characterise the passive, active and diverse approaches. These explore different approaches an insurer might employ to help determine which controls are effective in mitigating losses. At a high level we can say that the passive insurer is hands-off, the active approach ensures the policyholders have different levels of security, and the diverse insurer maintains diverse security configurations among the policyholders. The active and diverse approaches assume different levels of confidence in terms of where the attacks might land. Each strategy is illustrated in Table 2.

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

	x_1	x_2
Policyholder A	x	
Policyholder B	o	x

x'_1	x'_2
	x
o	x

Table 3: An illustration of how the ordering of x_l and x'_l impacts the realisation of attacks.

For the active insurer to share claims information, its policyholders must adopt different defensive postures. Otherwise, they will each gain identical information about the attacker. We assume the m policyholders have different security levels driven by some combination of premium incentives, mandated controls or differing risk aversion levels. This assumptions is supported by evidence [8,10] that insurance coverage is offered to applicants with different security posture (possibly at different prices). This can be modelled by assuming that the i-th policyholder guards one more vulnerability than the $(i-1)$-th policyholder. If a given vulnerability is exploited in any policyholder, it will be protected by every policyholder in the following round. This part of the strategy is facilitated by the active insurer sharing claims information.

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

We now turn to the calculation.

4.3 Calculation using Monte Carlo simulations

First we explain why the approximation used in the original IWL is not suitable for IWL-CSI calculations. We then simulate the original IWL using Monte Carlo methods and validate the results against those found using the approximation described in Section 3.3.

Since the attacker exploits the *weakest link*, the order in which vulnerabilities are exploited is determined by the ordering of the true costs of attack. This was not important in the original IWL [4], in which knowing the number of economically viable vulnerabilities is sufficient to calculate the revenue of that realisation. Whether the vulnerability i is exploited in the $(t-1)$-th or t-th round does not affect the revenue. Consequently, expected revenue is calculated by approximating a series of Bernoulli trials that represent the number of economically viable vulnerabilities.

The order in which vulnerabilities are exploited affects the revenue in the IWL-CSI. For example, consider the two separate realisations of true costs of defense x_l and x'_l in Table 3. In both cases, the expected true cost is $x_1 < x_2$, and $x'_1 < x'_2$, so policyholder B protects the first vulnerability. However, the actual realisation is $x_1 < x_2$ and $x'_1 > x'_2$. Consequently, in the first case, the policyholders gain the information that both x_1 and x_2 are economically viable. In the second case, the policyholders have no additional information as to whether x'_1 is economically exploitable. Consequently, the ordering of the true costs of attack impacts the IWL-CSI. However, the approximation used in the original IWL does not account for this.

The next question is whether simulations provide an appropriate method for calculation. Simulating the IWL using Monte Carlo methods takes into account the ordering of vulnerabilities. To validate this approach, we simulate one defender following the rules and strategy employed in the

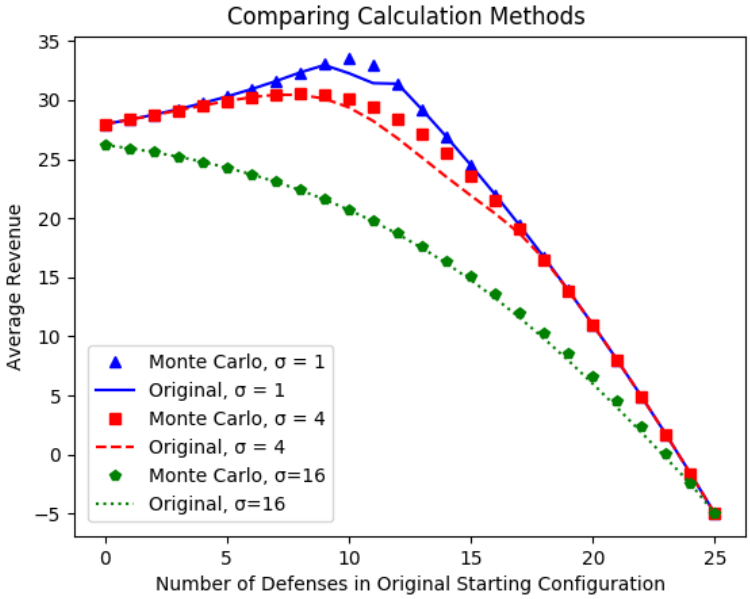

Figure 2: Average revenue of 10^7 Monte Carlo simulations compared to the expected revenue using the method of [4].

original IWL (we consider multiple defenders in Section 5). If we were simulating the natural world, we would have to choose an appropriate distribution. However, we are simulating the IWL, in which the true costs are assumed to be normally distributed.

Figure 2 contains the results of 10^7 simulations of the IWL for the same parameter values as in [4]. The two methods broadly agree; as uncertainty increases, both expected revenue and the optimal number of starting defenses fall. For low values of σ, the methods diverge around the optimal point. Indeed, for $\sigma = 1$ the simulations suggest that the optimal starting configuration is different to the results of [4].

Appendix A presents the standard error of the mean (SEM) for all of the calculations in Figure 2. Even when uncertainty is at its highest, the SEM results do not exceed 0.01. The source of divergence is most likely to be Böhme and Moore's decision to approximate the number of economically viable unguarded vulnerabilities as a Poisson binomial distribution [4].

However, we should not lose the forest for the trees. The IWL is not designed to be a predictive model. Rather, it is designed to gain insights into the strategies different parties might adopt under certain assumptions. Both the original computation and our simulations suggest that the optimal number of starting defenses decreases as uncertainty increases; the original insight of [4] — that sometimes it is rational to "wait and see" with security investments — holds true.

4.4 Summary

The original IWL consists of the rules, strategy and method of computation to consider the strategic interaction between one defender and an attacker. In Section 4.1, we extended the model to consider m defenders. Based on the new rules, we identified three separate strategies an insurer might employ

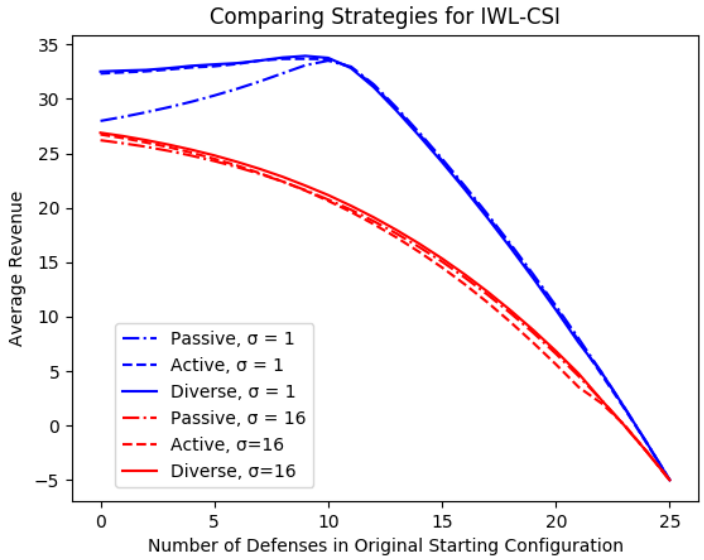

Figure 3: Average revenue for each of the three strategies based on 10^7 simulations with four insureds and no sunk costs.

in Section 4.2. Then, in Section 4.3, we provided some validation for using Monte Carlo methods to compute the original IWL and explained why simulations are better suited to computing the new rules. In Section 5, we will use the new rules and Monte Carlo methods to consider strategies the insurer of m policyholders might employ.

5 Results

We compare the average revenue of the three strategies in Section 5.1 and consider how these might be priced in Section 5.2. In Section 5.3, we look at the variance of claims. We then explore the impact of varying the number of insureds in Section 5.4. Apart from the parameters for uncertainty, sunk costs and numbers of insureds, all of the simulations use the same values as in [4].

5.1 Comparing the passive, active and diverse strategies

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

All of the results in this subsection are based on simulations with four insureds. With this parameter value, the insurer can gain information related to a maximum of 16% of the vulnerabilities per round (often the number will be less). The choice illustrates the differences between the three strategies by choosing a middle road, as evidenced in Section 5.4, which varies the number of insureds.

Figure 3 shows the result of 10^7 simulations without sunk costs. Displaying results like this highlights the optimal initial defensive configuration for a given strategy. For high uncertainty ($\sigma = 16$), it is still rational to under-invest and "wait and see" [4], but sharing claims information does result in a higher revenue. For low uncertainty ($\sigma = 1$), investing less initially but sharing claims information will result in higher average returns.

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

Figures 4 and 5 provide a comparison of active against passive and diverse against active respectively. Figure 4 shows that using claims data is most beneficial when the policyholders begin significantly under-invested and uncertainty is low. A policyholder not sharing claims information suffers many attacks in working out which vulnerabilities are economically viable.

Sharing claims information is least beneficial when the defender invests heavily in defense and uncertainty is low. The cost of maintaining diverse defenses is greater than the value gained by observing attacks because so few attacks are observed. However, Figure 3 shows that high initial investment will result in lower revenues and can only be justified by extreme risk-aversion.

Figure 5 compares the active and diverse strategies, which offer two different approaches to gathering claims information. The diverse strategy is most effective when uncertainty is moderate; the active strategy requires accurate forecasting of the ordering of the true costs of exploitation. When uncertainty is low, the active strategy accurately predicts the ordering and so there is little benefit in adopting the diverse approach. The diverse strategy can tolerate a threshold of inaccuracy. It performs similarly to the active strategy when the threshold is exceeded, which occurs when uncertainty is high and there is a low starting configuration.

Table 4 displays a number of security investment metrics. Broadly speaking, the active and diverse strategies result in adopting a less secure initial defensive configuration with a higher gross return. In every case, the passive strategy results in a higher attack intensity I, which describes the proportion of rounds in which a policyholder suffers an attack.

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

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

When sunk costs are introduced, the active strategy begins to achieve a higher ROSI than the passive strategy even when uncertainty is high. Despite the active strategy's predictions being

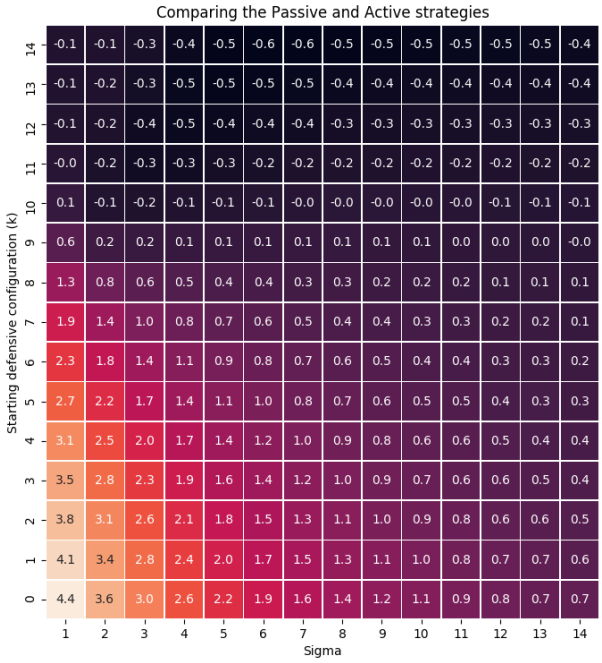

Figure 4: Improvement in revenue gained by adopting the active strategy as opposed to the passive strategy.

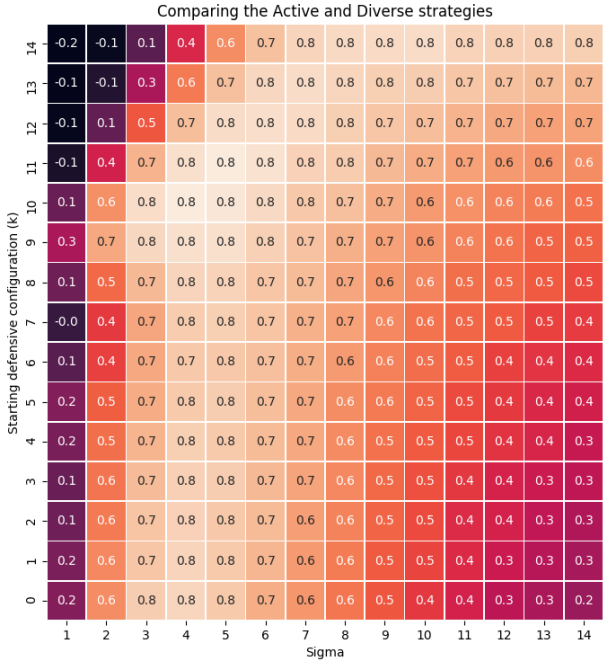

Figure 5: Improvement in revenue gained by adopting the diverse strategy as opposed to the active strategy.

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

Table 4: The results of 10^7 simulations with and without sunk costs.

poor, the few that do gain new information mean the insureds need to make less costly defensive configurations. Indeed, this is even more true for the diverse strategy in the case of moderate uncertainty, which achieves twice the ROSI of the passive strategy. This suggests that sharing claims data is especially valuable in guiding investments when it is costly to change defensive configuration.

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

5.2 Pricing

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

$$ALE = (za)I \tag{11}$$

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

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

$$R = ra - l.ALE - C \tag{12}$$

This is to be contrasted to the case without risk transfer where the policyholder's revenue is dependent on the random realisation of attacks. A risk-averse policyholder might prefer to lose a

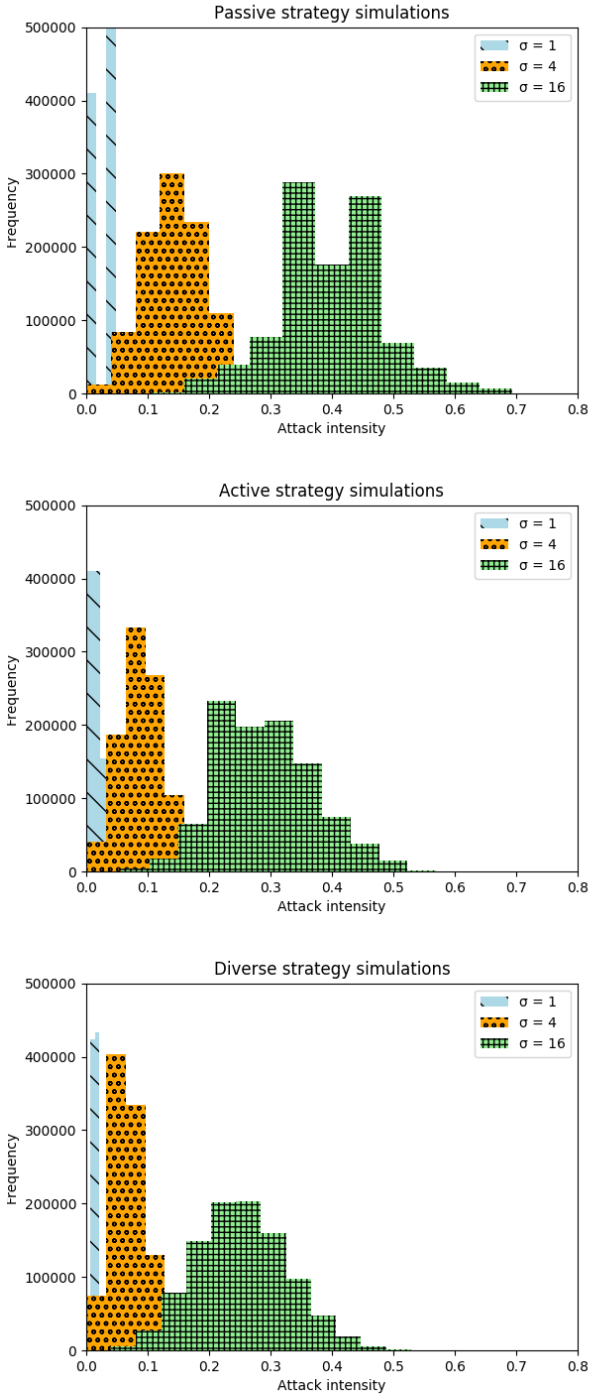

Figure 6: Distribution of claims for the passive, active and diverse strategies.

guaranteed $l.ALE$ rather than risk losing za, even if purchasing insurance results in lower expected revenue than can be expected without purchasing insurance.

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

$$\lambda a - l^* ALE_1 - C_1 = \lambda a - ALE_0 - C_0 \tag{13}$$

Rearranging, the loading factor limit is given by

$$l^* = \frac{ALE_0 + (C_0 - C_1)}{ALE_1} \tag{14}$$

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

$$L^* = (l^* - 1)ALE_1 \tag{15}$$

Table 4 contains the calculations of l^* and L^* based on the simulations. The results suggest that the insurer can charge a higher percentage of the premium to cover operating costs and profits when sunk costs are present. The value to the policyholder in sharing claims information is higher when it is costly to change defensive configuration frequently. However, the size of the premium is far higher in absolute terms when uncertainty is high, which is driven by the higher attack intensity. However, the insurer must hold capital reserves to cover potential claims, which contributes an additional cost. The next subsection investigates this aspect.

5.3 Measures of dispersion

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

Figure 6 highlights how the distribution of claims differs for the optimal starting configuration for each strategy. Comparing the three figures shows that the passive strategy suffers more attacks for each uncertainty value. As uncertainty increases, both the mean and variance of the attack intensity increases. The insurer must hold more capital to account for the higher mean attack intensity and additional capital must be held to account for the variance of the attack intensity. Alternatively, the insurer may choose to purchase reinsurance to cover the tail of these attack intensities.

The passive strategy simulations illustrate how ordering is not important for the original IWL; when $\sigma = 1$ the two bars capture the two cases when 0, 1 or 2 economically viable exploits are left unguarded. Recalling that the passive strategy is equivalent to the original IWL, these results suggest that the "wait and see" approach that was found to be optimal for high uncertainty in [4] leads to a concerning variability in the number of attacks faced. Table 4 shows that both the mean and the standard deviation of the attack intensity are lower for the passive and active strategies for all uncertainty levels. This suggests that less capital would need to be held in reserve for the strategies that involve sharing claims data.

5.4 Varying the number of insureds

The prior results show that sharing claims information is preferable to not sharing it. However, the results have, with a few exceptions, not revealed a significant difference between the active and

diverse strategies. This subsection does not display the results for the passive insurer because the passive insurer does not share information. Consequently, the average revenue is the same for all number of insureds.

Figure 7 varies the number of insureds, displaying the average revenue for the optimal starting configuration for each parameter choice. The active strategy is optimal for between two and four policyholders for all the uncertainty levels we consider. For the diverse strategy, the optimal number of policyholders increases with the level of uncertainty. Figure 7 shows that for $\sigma > 4$, 16 policyholders leads to the highest revenue. In fact, the passive strategy outperforms the active strategy when m and σ are high.

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

Figure 7 allows us to consider the marginal benefit of insuring another policyholder. When uncertainty is low, insuring an additional policyholder will not significantly affect the average revenue for either of the active or diverse approaches. The absolute difference between all numbers of policyholders is less than 0.5 when $\sigma = 1$. As uncertainty increases, the marginal benefit diverges for each strategy. For the active approach, there is a small gain in moving from 2 to 4 policyholders and the marginal benefit of another policyholder is negative. However, for the diverse strategy, this benefit is always positive for $\sigma > 8$ and for $\sigma = 16$ the difference between 2 and 16 policyholders is four times that for $\sigma = 1$. The implications of these results in relation to market composition are discussed in the next section.

6 Discussion

Our motivating concern is the ability of the insurer to aggregate claims information across numerous multiple policyholders. In Section 6.1, we discuss the validity of the assumptions in the IWL-CSI that characterise this ability. We then discuss how these results inform the optimal strategy for the insurer to adopt in Section 6.2. Finally, we discuss future work related to IWL-CSI in Section 6.3.

6.1 Assumptions

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

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

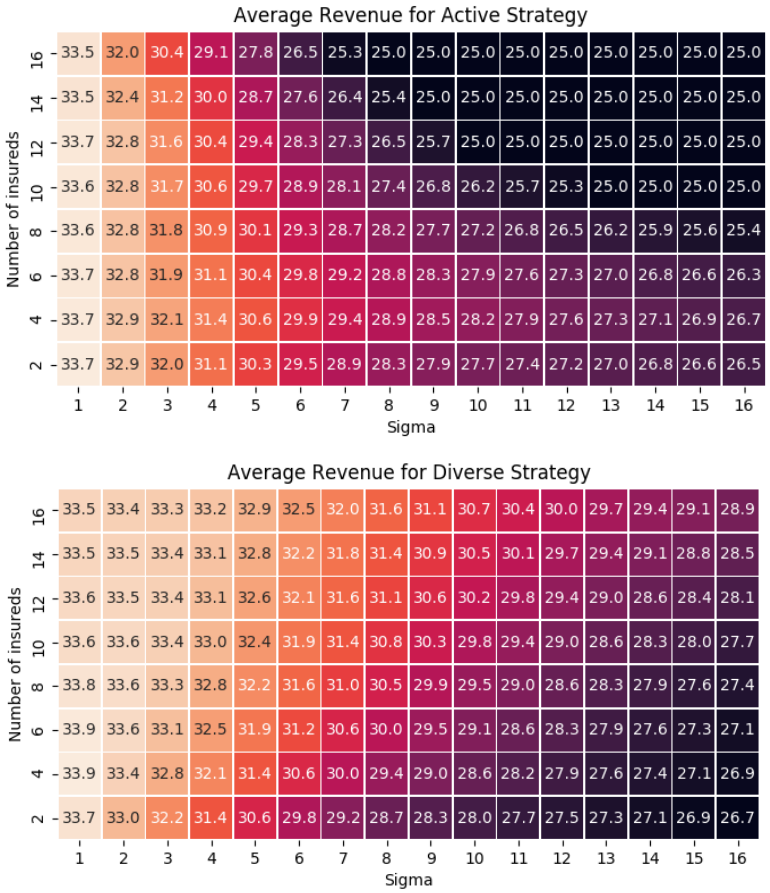

Figure 7: Average revenue for 10^7 simulations with different numbers of policyholders and uncertainty levels for the active and diverse strategies.

Assumption (ii) depends on market dynamics. It is less appropriate in a 'soft market' characterised by an over-supply of insurance. Insurers may not have the freedom to discriminate between applicants based on their security posture, nor be able to mandate security controls, when the insured can easily find coverage elsewhere. The next subsection discusses strategies the insurer might adopt to take advantage of this market power.

6.2 Insurer strategy

The results suggest that an insurer sharing claims information will pay less in claims, as evidenced by the lower attack intensity in Table 4. The ROSI calculations show this increased revenue results from more efficient security investments. Further, sharing claims information reduces the variance of claims meaning that insurers need to hold less capital in reserve.

However, achieving higher revenues relies on security investments being influenced by claims data. Empirical work has shown that insurers do not seem certain about which security controls are important, as two thirds of insurers in one study did not adapt premium prices based on security controls [8]. Our results suggest that this may even be rational — when uncertainty is high, the optimal starting configuration is no defensive investment. Perhaps insurers will begin to require security controls as they observe attacks on their policyholders and discover which defensive measures are effective.

Such a strategy is reliant on dynamically adapting the defensive configuration. In [1], the present authors identified a number of mechanisms by which the insurer can influence security controls of their insureds, which we call *risk selection*, *incentivisation* and *integration*.

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

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

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

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

This question depends on the change to defensive configuration. Unlike in the stylised IWL-CSI model, security controls are not homogeneous and some may be more difficult to implement than others. For example, secure software engineering investments may be more efficient in pre-deployment [28] — long before the insurance policy has been purchased. This raises the concern that the insurance industry's incentives do not align with the insured's long term interest — a concern borne out in the security controls that cyber insurance application forms focus on [12].

Insurers discovering which controls are effective is by no means guaranteed. Collecting information about the security posture of the insured presents problems in terms of the granularity of

the information collected and the reliability of self-reported answers [8]. An application form may report that a firewall is in place without specifying the monitoring processes, let alone whether it is configured correctly. These issues are exacerbated by difficulties establishing the effectiveness of these controls from post-event reporting. To our knowledge, no empirical work has established what information is collected in the claims process, nor how this is used by insurers.

Another concern is whether all the policyholders achieve the same return. As the IWL-CSI strategy relies on defenders adopting different postures, they will suffer different attacks and costs of defense. This raises the question of which defender should adopt the most costly defensive configuration or the one most likely to suffer attack. In theory this can be handled by the market; the insurer should offer premium discounts to defenders adopting more costly defenses. However, it would not be the first market failure in information security if this were not the case [26].

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

Market composition affects the total revenue across all the insureds. Under conditions of high uncertainty, insurers can increase average revenue by taking on an increasing number of policyholders, which raises competition concerns. This is consistent with the accounts stating that large insurers see their claims data as a competitive advantage. Indeed, "insurers with a small amount of claims data were (perhaps understandably) far more enthusiastic about data being shared" [1]. The results suggest that under conditions of high uncertainty, the value of claims data may be such that the marginal benefit of an additional insured drives the market composition towards monopoly. Again, it would not be the first time incentives have tended towards monopoly [26]. However, this analysis does not take into account the risk of aggregated losses which may push the market towards many providers.

In the IWL-CSI, uncertainty and costs of defense are fixed for the entire game. However, in reality these will change over time: organisations will adopt new information systems changing the cost of attack and uncertainty regarding the attacker may rise or fall. The "technical flux of change" undermines the utility of actuarial data over time [29]. There are two competing factors here: (i) changing costs of exploit make old claims data less valuable, and (ii) old claims data may reduce uncertainty about the attacker in future games. The former erodes the competitive advantage of claims data over time while the latter strengthens it. For example, if one insurer holding a larger share of the market leads them to operate with lower uncertainty than another, they can expect to gain higher revenues. There could be a tendency towards monopoly in the market, depending on the balance of (i) and (ii).

6.3 Future work

Investigating the role of aggregating claims information in cyber insurance is a novel research direction and many knowledge gaps remain. Empirical work could probe the validity of the assumptions underpinning our extension of the IWL. The assumption that defenders are homogeneous could be loosened by allowing their expectations regarding the cost of exploiting a vulnerability to differ while still assuming the true cost of exploitation is the same across defenders. Data could be collected to set more insightful parameter values.

Future work could introduce uncertainty into risk assessment and post-breach forensics. Reasoning about the trade-offs between information gained by the insurer and the burden placed on the applicant may provide insights into phenomena like "race-to-the-bottom" cyber risk assessment

standards [1] or information asymmetries [30].

Pricing risks correctly is but one concern of the insurer: they are also concerned with the aggregation of claims [25]. The danger that many policyholders might claim at once provides an incentive for insurers to diversify risk. Future work on the IWL-CSI could consider an insurer with a limited amount of capital depleted each time a policyholder is attacked and build upon the variance of claims analysis seen in Section 5.3.

While the IWL-CSI focuses on the insurer's ability to share claims information, other investigations make considerations for the insurer's shortcomings, such as: imperfect ex-ante assessment of an insured's defensive investment (adverse selection) [31]; inability to observe the insured's engagement in risky behaviour (moral hazard) [30,32]; and how insecurity, when combined with risk averse insureds, increases demand for coverage in the insurer's interest (the demand value of losses) [33]. Each line of study provides isolated insights. However, the "model of models" that can reason about the effect of all of the insurer's abilities and shortcomings in combination will be of most value in analysing the impact of the cyber insurance market.

7 Conclusion

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

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

8 Acknowledgements

The authors thank the anonymous reviewers for their helpful and constructive comments. Daniel Woods' research is funded by the EPSRC via the Centre for Doctoral Training in Cyber Security at the University of Oxford. Dennis Jackson and Martin Dehnel-Wild kindly provided access to computing resources for the simulations.

References

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

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

[3] Ross Anderson. Is it practical to build a truly distributed payment system? *Keynote at the 23rd ACM Conference on Computer and Communications Security*, 2016.

[4] Rainer Böhme and Tyler Moore. The "iterated weakest link" model of adaptive security investment. *Journal of Information Security*, 7(2):81–102, 2016.

[5] Christian Biener, Martin Eling, and Jan Hendrik Wirfs. Insurability of cyber risk: An empirical analysis. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 40(1):131–158, 2015.

[6] Walter Baer. Rewarding IT security in the marketplace. *Contemporary Security Policy*, 24(1):190–208, 2003.

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

[8] Sasha Romanosky, Lillian Ablon, Andreas Kuehn, and Therese Jones. Content analysis of cyber insurance policies: How do carriers write policies and price cyber risk? In *Proceedings of The 16th Workshop on the Economics of Information Security (WEIS 2017)*, 2017.

[9] Angelica Marotta, Fabio Martinelli, Stefano Nanni, Albina Orlando, and Artsiom Yautsiukhin. Cyber-insurance survey. *Computer Science Review*, 24:35–61, 2017.

[10] Ulrik Franke. The cyber insurance market in Sweden. *Computers & Security*, 68:130–144, 2017.

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

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

[13] Ross Anderson and Bruce Schneier. Guest editors' introduction: Economics of information security. *IEEE Security & Privacy*, 3(1):12–13, 2005.

[14] William Yurcik and David Doss. Cyberinsurance: A market solution to the internet security market failure. In *Proceedings of The 1st Workshop on the Economics of Information Security (WEIS 2002)*, 2002.

[15] Rainer Böhme and Galina Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Proceedings of The 9th Workshop on the Economics of Information Security (WEIS 2010)*, 2010.

[16] Hulisi Ogut, Nirup Menon, and Srinivasan Raghunathan. Cyber insurance and IT security investment: Impact of interdependence risk. In *Proceedings of The 4th Workshop on the Economics of Information Security (WEIS 2005)*, 2005.

[17] Jean-Chrysostome Bolot and Marc Lelarge. A new perspective on internet security using insurance. In *Proceedings of The 27th IEEE Conference on Computer Communications (INFOCOM 2008)*, pages 1948–1956. IEEE, 2008.

[18] Xia Zhao, Ling Xue, and Andrew Whinston. Managing interdependent information security risks: A study of cyberinsurance, managed security service and risk pooling. *Proceedings of The 30th International Conference on Information Systems (ICIS 2009)*, pages 49–66, 2009.

[19] Rainer Böhme. Cyber-insurance revisited. In *Proceedings of The 4th Workshop on the Economics of Information Security (WEIS 2005)*, 2005.

[20] Rainer Böhme and Gaurav Kataria. Models and measures for correlation in cyber-insurance. In *Proceedings of The 5th Workshop on the Economics of Information Security (WEIS 2006)*, 2006.

[21] Lawrence Gordon and Martin Loeb. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)*, 5(4):438–457, 2002.

[22] Rainer Böhme. Security metrics and security investment models. In I. Echizen, N. Kunihiro, and R. Sasaki, editors, *Proceedings of The 5th International Workshop on Security (IWSEC 2010)*, volume 6434 of *Lecture Notes in Computer Science*, pages 10–24, 2010.

[23] Chad Heitzenrater and Andrew Simpson. Policy, statistics and questions: Reflections on UK cyber security disclosures. *Journal of Cybersecurity*, 2(1):43–56, 2016.

[24] Department for Business, Innovation & Skills. Information security breaches survey, 2015. Available at `https://www.gov.uk/government/publications/information-security-breaches-survey-2015`. [Online; accessed 27-July-2016].

[25] Rob Thoyts. *Insurance theory and practice*. Routledge, 2010.

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

[27] Mark Stamp. Risks of monoculture. *Communications of the ACM*, 47(3):120, 2004.

[28] Chad Heitzenrater, Rainer Böhme, and Andrew Simpson. The days before zero day: Investment models for secure software engineering. In *Proceedings of The 15th Workshop on the Economics of Information Security (WEIS 2016)*, 2016.

[29] Daniel Geer, Kevin Soo Hoo, and Andrew Jaquith. Information security: Why the future belongs to the quants. *IEEE Security & Privacy*, 99(4):24–32, 2003.

[30] Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. Competitive cyber-insurance and internet security. In T. Moore, D. Pym, and C. Ioannidis, editors, *Economics of Information Security and Privacy: Proceedings of The 9th Workshop on the Economics of Information Security (WEIS 2010)*, pages 229–247. Springer, 2010.

[31] Ranjan Pal, Leana Golubchik, Konstantinos Psounis, and Pan Hui. Will cyber-insurance improve network security? A market analysis. In *Proceedings of The 2014 International Conference on Computer Communications (INFOCOM 2014)*, 2014.

[32] Mohammad Khalili, Parinaz Naghizadeh, and Mingyan Liu. Designing cyber insurance policies in the presence of security interdependence. In *Proceedings of The 12th Workshop on The Economics of Networks, Systems and Computation (NetEcon 2017)*, 2017.

[33] Fabio Massacci, Joe Swierzbinski, and Julian Williams. Cyberinsurance and public policy: Self-protection and insurance with endogenous adversaries. In *Proceedings of The 16th Workshop on the Economics of Information Security (WEIS 2017)*, 2017.

A Error Bars

Standard Error of the Mean (10^{-3})			
k	$\sigma = 1$	$\sigma = 4$	$\sigma = 16$
0	1.59	3.17	4.86
1	1.65	3.28	4.95
2	1.71	3.39	5.03
3	1.77	3.49	5.09
4	1.83	3.59	5.15
5	1.89	3.65	5.20
6	1.95	3.68	5.22
7	2.01	3.67	5.23
8	2.07	3.60	5.24
9	2.09	3.48	5.22
10	1.86	3.28	5.19
11	1.20	3.03	5.13
12	0.48	2.72	5.06
13	0.12	2.38	4.98
14	0.02	2.02	4.86
15	0.00	1.67	4.72
16	0.00	1.34	4.56
17	0.00	1.04	4.37
18	0.00	0.79	4.16
19	0.00	0.58	3.90
20	0.00	0.41	3.61
21	0.00	0.29	3.26
22	0.00	0.19	2.85
23	0.00	0.12	2.35
24	0.00	0.06	1.68
25	0.00	0.00	0.00

Table 5: Standard error of the mean when simulating the original IWL for various values of σ.