Cyber-Insurance as a Signaling Game: Self-Reporting and External Security Audits

Aron Laszka[1], Emmanouil Panaousis[2], and Jens Grossklags[3]

[1] Department of Computer Science, University of Houston
[2] Surrey Centre for Cyber Security, University of Surrey
[3] Department of Informatics, Technical University of Munich

Published in the proceedings of the 9th Conference on
Decision and Game Theory for Security (GameSec 2018).

Abstract. An insurer has to know the risks faced by a potential client to accurately determine an insurance premium offer. However, while the potential client might have a good understanding of its own security practices, it may also have an incentive not to disclose them honestly since the resulting information asymmetry could work in its favor. This information asymmetry engenders adverse selection, which can result in *unfair premiums* and *reduced adoption of cyber-insurance*. To overcome information asymmetry, insurers often require potential clients to self-report their risks. Still, clients do not have any incentive to perform thorough self-audits or to provide comprehensive reports. As a result, insurers have to complement self-reporting with external security audits to verify the clients' reports. Since these audits can be very expensive, a key problem faced by insurers is to devise an auditing strategy that deters clients from dishonest reporting using a minimal number of audits. To solve this problem, we model the interactions between a potential client and an insurer as a two-player signaling game. One player represents the client, who knows its actual security-investment level, but may report any level to the insurer. The other player represents the insurer, who knows only the random distribution from which the security level was drawn, but may discover the actual level using an expensive audit. We study the players' equilibrium strategies and provide numerical illustrations.

1 Introduction

Technological innovations, such as artificial intelligence and ubiquitous computing, are becoming integral parts of our lives, and providing us with many benefits. But these developments also bring new threats, and the insurance industry is playing catch-up to keep pace with the rapid rise of cyber-risks. Cyber-threat remains one of the most significant—and growing—risks facing businesses. For example, a UK government survey estimated that in 2014, 81% of large corporations and 60% of small businesses have suffered a security breach. The average cost of breaches, in the UK, was between £600,000 and £1.15 million for large businesses and between £65,000 and £115,000 for small businesses [1]. Further, in 2016, more than 1.1 billion identities were stolen in data breaches, almost double the number stolen in 2015 [23]. In aggregate, Forbes reports that cyber-crime losses will be more than US$2.1 trillion by 2019 [9].

Unfortunately, even with the strongest cyber-security controls purchased and implemented, an organization is at risk of being compromised. As such, apart from security measures, responses to cyber-security risk include outsourcing it by purchasing cyber-insurance coverage. However, 60% of Fortune 500 companies still lack any insurance against cyber-incidents, primarily due to a lack of cover currently available for many types of cyber-risk [14].

Cyber-insurance, as any other field of insurance, faces a number of challenges [6,15,16]. In particular, *asymmetry of information* and the resulting *adverse selection* caused by organizations being reluctant to share their actual levels of cyber-risk may present significant premium pricing obstacles to insurers. It is, therefore, perhaps unsurprising that insurers tend to offer a pricing structure that charges companies similar rates regardless of the underlying actual risks [18]. However, if a cyber-insurer cannot differentiate between clients based on their security level and therefore cannot offer differentiated premiums, the insurer will not be able to sustain a profitable business [3].

Typically, insurers require organizations to self-report on their security level in order to determine premiums. Prior to setting the premium, the insurer must then decide whether the security level reported by the client must be confirmed by undertaking some audit (e.g., penetration testing). Although it is beneficial for the insurer to know the exact security level of its potential client so that it can ask for a *fair* premium, there is a cost associated with conducting an audit.[4] The insurer has two options: (i) to trust that the security level the client reported is true and compute the premium based on this level, thereby saving audit costs; or (ii) not to trust the reported security level and perform an audit to reveal the real security level despite having to pay for an audit to take place.

After the insurer offers a premium, the client must decide whether it will accept the offer and be underwritten, or whether it will not use cyber-insurance at all. This is an important decision to be made, and it has been noticed that many organizations, especially small-to-medium enterprises, decide not to purchase cyber-insurance due to the incurred financial costs [11].

Contributions: The aim of this research is to introduce a new model to study optimal strategies for self-reporting security levels (for organizations) and undertaking audits (for insurers). The insurers' strategy aims to ensure that the actual security levels of their clients have been elicited and therefore "fair" contracts (coverage, premium) are put in place.

More concretely, we model the interactions between a potential client and an insurer as a two-player signaling game, where the organization plays the role of the sender, while the insurer plays the role of the receiver. We assess our game model using numeric simulations to derive the probability of reporting each type,

[4] In fact, the cost of penetration testing, cyber-security risk assessment and related services is non-trivial and quickly increases with the size of an organization. See, for example, the pricing examples at: `https://www.trustnetinc.com/pricing/`.

Table 1. List of Symbols

Symbol	Description
$\mathcal{S}$	Set of organization types (i.e., security levels)
P_t	Probability of the organization's type being t
t, T	Organization's real type (realization, random variable)
r, R	Organization's reported type (realization, random variable)
p	Cyber-insurance premium
$\boldsymbol{\rho}^t$	Reporting strategy of organization with real type t
$\boldsymbol{a}$	Insurer's strategy for auditing the organization
$\boldsymbol{p}^A$	Insurer's strategy for premium selection after auditing
$\boldsymbol{p}^N$	Insurer's strategy for premium selection without auditing
W	Organization's initial wealth
L	Organization's loss in case of a cyber-incident
U	Organization's utility function
C	Insurer's cost for auditing the organization

the audit probabilities, and the insurance premiums for various audit cost values. The proposed game-theoretic model can form the basis of a framework that can further accelerate the adoption of cyber-insurance.

2 Model

We model the interactions between an *insurer* and a potential client, whom we will call the *organization*, as a two-player single-shot game. For a list of symbols that are used in our model, see Table 1. We assume that the organization has a type $t \in \mathcal{S}$, where $\mathcal{S}$ is a finite set of types. Type t models the level of security investments and the combination of security measures that the organization implements. For simplicity, we let type t be equal to the estimated probability of the organization not suffering a cyber-incident.

The organization applies for insurance coverage and the insurer determines the premium as follows:

- First, the organization's type t is drawn randomly from the set of types $\mathcal{S}$ according to a known distribution[5]. We let T denote the random variable taking the value of the organization's type, and we let P_t denote the probability that the organization's type is t (i.e., $P_t = \Pr[T = t]$).
- Second, the organization chooses a type $r \in \mathcal{S}$ that it reports to the insurer. The organization's choice may be randomized based on a mixed strategy. We let R denote the random variable taking the value of the reported type.
- Based on the reported type r, the insurer decides whether to audit the organization or not. If the insurer chooses to audit the organization, then the true type t is revealed, but the insurer incurs a constant auditing cost C.
- Finally, based on the type t or r (depending on whether the organization has been audited), the insurer chooses a premium p that is asked from the organization in exchange for insurance coverage. The organization rejects the

[5] Randomness models the insurer's a priori uncertainty regarding what type of organization it faces.

coverage if doing so increases its utility; otherwise, it accepts the coverage and pays the premium.

2.1 Strategies

An organization's strategic choice is to select what type to report to the insurer. We let $\boldsymbol{\rho}^t$ denote the mixed strategy of an organization with real type t, where ρ_r^t is the probability that the organization reports type r (i.e., $\rho_r^t = \Pr[R = r \,|\, T = t]$). Note that we assume that the organization's strategic choice does not include coverage acceptance or rejection (i.e., we assume that coverage is rejected if and only if it is not worth purchasing). This is similar to assuming that the organization makes coverage decisions but restricting the solutions to subgame perfect equilibria (i.e., prohibiting non-credible threats of not purchasing insurance).

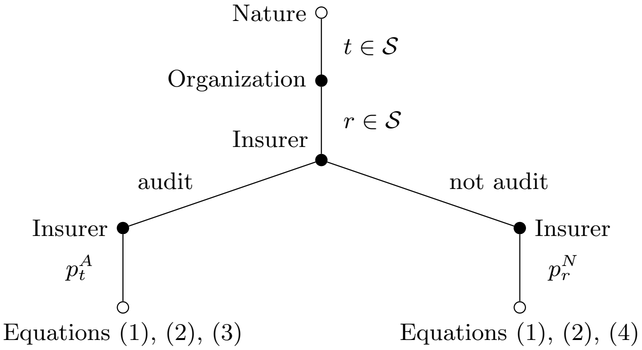

Fig. 1. Tree representation of the game. The players' payoffs are given by Equations (1), (2), (3), and (4).

The insurer's first strategic choice is to decide whether to audit the organization or not. Before auditing, the insurer does not know the organization's real type t, but it does know the exogenous parameter values of the model[6], which include the distribution from which the type was drawn (i.e., it knows the probabilities P_t), as well as the organization's reporting strategies $\boldsymbol{\rho}$. We let $\boldsymbol{a}$ denote the insurer's strategy, where a_r is the probability that the insurer audits an organization with reported type r. The insurer's second strategic choice is to choose a premium p. First, we let $\boldsymbol{p}^N$ denote the insurer's strategy given that it has not performed an audit, where p_r^N is the premium asked from an organization with reported type r. Second, we let $\boldsymbol{p}^A$ denote the insurer's strategy given that it has performed an audit, where p_t^A is the premium asked from an organization with real type t.

2.2 Payoffs

Now, we define the players' payoffs in the various outcomes of our game. As it is standard in the cyber-insurance literature, we capture the risk aversion of clients

[6] These may be learned from statistics that are available to the insurer.

using a concave utility function, initial wealth, and potential losses. First, the organization's payoff (i.e., utility), if it accepts coverage is

$$\mathcal{U}_t^{org,acc}(p) = U(W - p), \tag{1}$$

where W is the organization's initial wealth, and U is its utility function, which we assume to be continuous, monotonically increasing, and concave.

Second, the organization's payoff if it rejects coverage is

$$\mathcal{U}_t^{org,rej} = (1 - t) \cdot U(W - L) + t \cdot U(W), \tag{2}$$

where L is its loss in case of a cyber-incident. The two terms correspond to the cases of suffering a cyber-incident and not suffering one, respectively.

If the insurer audits the organization, its payoff (i.e., profit) is

$$\mathcal{U}^{ins,aud}(t, p) = (p - (1 - t) \cdot L) \cdot 1_{\{\text{insurance accepted}\}} - C, \tag{3}$$

where $1_{\{\text{insurance accepted}\}}$ is equal to 1 if the organization purchases insurance, and 0 otherwise. If the insurer does not audit, then its payoff is

$$\mathcal{U}^{ins,noaud}(t, p) = (p - (1 - t) \cdot L) \cdot 1_{\{\text{insurance accepted}\}}. \tag{4}$$

Note that the insurer does not learn the true value of t if it does not audit the organization; however, its payoff still depends on t.

Given mixed-strategy profile $(\boldsymbol{\rho}, (\boldsymbol{a}, \boldsymbol{p}^N, \boldsymbol{p}^A))$, the expected utility of an organization with type t is

$$\mathbb{E}\left[\mathcal{U}_t^{org}\right]\left(\boldsymbol{\rho}, \boldsymbol{a}, \boldsymbol{p}^N, \boldsymbol{p}^A\right) = \sum_{r \in \mathcal{S}} \rho_r^t \Bigg[a_r \cdot \max\left\{\mathcal{U}_t^{org,acc}(p_t^A),\ \mathcal{U}_t^{org,rej}\right\} + (1 - a_r) \cdot \max\left\{\mathcal{U}_t^{org,acc}(p_r^N),\ \mathcal{U}_t^{org,rej}\right\}\Bigg],$$

while the insurer's expected utility is

$$\mathbb{E}\left[\mathcal{U}^{ins}\right]\left(\boldsymbol{\rho}, \boldsymbol{a}, \boldsymbol{p}^N, \boldsymbol{p}^A\right) = \sum_{t \in \mathcal{S}} P_t \sum_{r \in \mathcal{S}} \rho_r^t \Bigg[a_r \cdot \mathcal{U}^{ins,aud}(t, p_t^A) + (1 - a_r) \cdot \mathcal{U}^{ins,noaud}(t, p_r^N) \Bigg].$$

2.3 Solution Concept

We are interested in finding an equilibrium of our game, which can capture the long-term insurance market equilibrium. Since our model is essentially a signalling game, we use *perfect Bayesian Nash equilibrium* as the solution concept.

After receiving the reported level r, the insurer's belief regarding the potential client's real type can be expressed using Bayes' rule as

$$\Pr\left[T = t | R = r\right] = \frac{\Pr\left[T = t, R = r\right]}{\Pr\left[R = r\right]} = \frac{P_t \cdot \rho_r^t}{\sum_{t' \in \mathcal{S}} P_{t'} \cdot \rho_r^{t'}}.$$

A mixed-strategy profile $(\boldsymbol{\rho}^*, (\boldsymbol{a}^*, \boldsymbol{p}^{N^*}, \boldsymbol{p}^{A^*}))$ is an equilibrium if

- for each security level $t \in \mathcal{S}$, the strategy $\boldsymbol{\rho}^{t^*}$ maximizes the expected utility of an organization with level t given the insurer's strategy $(\boldsymbol{a}^*, \boldsymbol{p}^{N^*}, \boldsymbol{p}^{A^*})$:

$$\boldsymbol{\rho}^{t^*} \in \operatorname{argmax}_{\boldsymbol{\rho}^t} \mathbb{E}\left[\mathcal{U}_t^{org}\right]\left(\boldsymbol{\rho}^t, \boldsymbol{a}, \boldsymbol{p}^N, \boldsymbol{p}^A\right);$$

- for each reported security level $r \in S$, the strategy $(\boldsymbol{a}^*, \boldsymbol{p}^{N^*}, \boldsymbol{p}^{A^*})$ maximizes the expected utility of the insurer given its belief regarding the potential client's real type t:

$$(\boldsymbol{a}^*, \boldsymbol{p}^{N^*}, \boldsymbol{p}^{A^*}) \in \operatorname{argmax}_{(\boldsymbol{a}, \boldsymbol{p}^N, \boldsymbol{p}^A)} \sum_{t \in \mathcal{S}} \Pr[T = t \,|\, R = r]\Big[a_r \cdot \mathcal{U}^{ins,aud}(t, p_t^A) \\ + (1 - a_r) \cdot \mathcal{U}^{ins,noaud}(t, p_r^N)\Big].$$

3 Preliminary Analysis

Next, we provide some necessary conditions on the players' best responses.

Lemma 1. *An organization of type t accepts insurance coverage for premium p if and only if $p \leq \hat{p}_t$, where*

$$\hat{p}_t = W - U^{-1}\left((1-t) \cdot U(W-L) + t \cdot U(W)\right). \tag{5}$$

Proof. By definition, an organization with type t accepts coverage for premium p if and only if

$$\begin{aligned} \mathcal{U}_t^{org,acc}(p) &\geq \mathcal{U}_t^{org,rej} \\ U(W-p) &\geq (1-t) \cdot U(W-L) + t \cdot U(W) \\ W - p &\geq U^{-1}\left((1-t) \cdot U(W-L) + t \cdot U(W)\right) \\ p &\leq W - U^{-1}\left((1-t) \cdot U(W-L) + t \cdot U(W)\right) := \hat{p}_t. \end{aligned}$$

□

Lemma 2. *In an equilibrium, the premium $p_t^{A^*}$ that an insurer requests from an organization with type t after an audit is $p_t^{A^*} = \hat{p}_t$ if $\hat{p}_t \geq (1-t) \cdot L$. Otherwise, the insurer asks for some premium $p_t^{A^*} > p_t^*$, which will always be rejected by the organization.*

Proof. If the insurer has audited an organization and found its type to be t, then its payoff for premium p will be

$$\begin{aligned} \mathcal{U}^{ins,aud}(t,p) &= (p - (1-t) \cdot L) \cdot 1_{\{\text{insurance accepted}\}} - C &(6) \\ &= (p - (1-t) \cdot L) \cdot 1_{\{p \leq \hat{p}_t\}} - C. &(7) \end{aligned}$$

When $p \leq \hat{p}_t$, the first derivative of the payoff $\mathcal{U}^{ins,aud}(t,p)$ is

$$\frac{\partial \mathcal{U}^{ins,aud}(t,p)}{\partial p} = 1; \tag{8}$$

hence, the maximum on interval $(-\infty, \hat{p}_t]$ is attained at $\hat{p}_t$, and the maximum payoff is $\hat{p}_t - (1-t) \cdot L - C$. When $p > \hat{p}_t$, the payoff is always $-C$ since the organization rejects coverage. Hence, $\hat{p}_t$ is an optimal premium $p_t^{A^*}$ if and only if

$$\hat{p}_t - (1-t) \cdot L - C \geq -C \tag{9}$$

$$\hat{p}_t \geq (1-t) \cdot L; \tag{10}$$

otherwise, premiums greater than $\hat{p}_t$ are optimal, which will be rejected. □

Lemma 3. *In an equilibrium, the premium $p_r^{N^*}$ that an insurer requests without an audit from an organization with reported type r is either $p_r^{N^*} \in \{\hat{p}_t \,|\, t \in \mathcal{S}\}$, which may be accepted by some organizations, or $p_r^{N^*} > \max\{\hat{p}_t \,|\, t \in \mathcal{S}\}$, which will be rejected by any organization.*

Proof. If the insurer has not audited an organization that reported type r, then its payoff for premium p will be

$$\sum_{t \in \mathcal{S}} \Pr[T = t | R = r] \mathcal{U}^{ins,noaud}(t, p) \tag{11}$$

$$= \sum_{t \in \mathcal{S}} \Pr[T = t | R = r] \left(p - (1-t) \cdot L\right) \cdot 1_{\{\text{insurance accepted}\}} \tag{12}$$

$$= \sum_{t \in \mathcal{S}} \Pr[T = t | R = r] \left(p - (1-t) \cdot L\right) \cdot 1_{\{p \leq \hat{p}_t\}}. \tag{13}$$

The values $\{\hat{p}_t \,|\, t \in \mathcal{S}\}$ divide the set of possible premium values $[0, \infty)$ into $|\mathcal{S}| + 1$ contiguous intervals, the last one being $(\max\{\hat{p}_t \,|\, t \in \mathcal{S}\}, \infty)$. The payoff is strictly increasing on each interval, except for the last one. On the last interval, $(\max\{\hat{p}_t \,|\, t \in \mathcal{S}\}, \infty)$, the payoff is always zero. Therefore, the optimal premium $p_r^{N^*}$ is either one of the values $\{\hat{p}_t \,|\, t \in \mathcal{S}\}$ or any value $p_r^{N^*} > \max\{\hat{p}_t \,|\, t \in \mathcal{S}\}$, which will be rejected by any organization. □

4 Numerical Illustrations

In this section, we present numerical illustrations of our model. We let $\mathcal{S} = \{0.5, 0.65, 0.8, 0.95\}$, $W = 10$, $L = 5$, the utility function U be the natural logarithm function, and the organization's type t be drawn according to $P_{0.5} = 0.125, P_{0.65} = 0.375, P_{0.8} = 0.375$, and $P_{0.95} = 0.125$; we used numerical search to find equilibria for various audit costs C.

Figure 2 shows the organization's and the insurer's expected payoffs in equilibrium as functions of the audit cost C. The organization's expected payoff remains steady if it is secure and has little incentive to misreport. But in the case of an organization with security level 0.5, the payoff increases when the auditing cost reaches the value of 0.22. On the other hand, the insurer's expected payoff decreases with the auditing cost, but the rate of reduction is fairly small.

Figure 3 shows the probabilities $\Pr[T = t]$ of reporting various types as functions of the audit cost C. We observe rampant misreporting since the organization's security level is either $t = 0.5$ or $t = 0.65$ with probability $0.5 =$

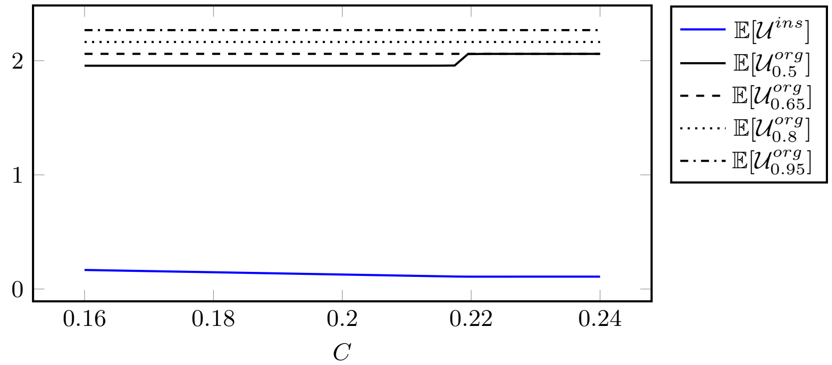

Fig. 2. Players' payoffs in equilibrium with various audit cost values.

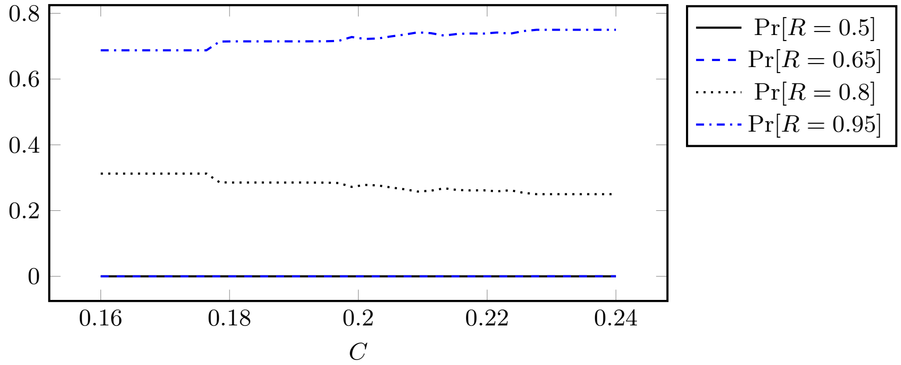

Fig. 3. Probability of reporting each type in equilibrium with various audit cost values.

$P_{0.5} + P_{0.65}$, but it *never* reports these low levels (i.e., $\Pr[R = 0.5] = 0$ and $\Pr[R = 0.65] = 0$). We also see that the probability $\Pr[R = 0.95]$ of misreporting a higher, "more suspicious" level increases as audits become more expensive.

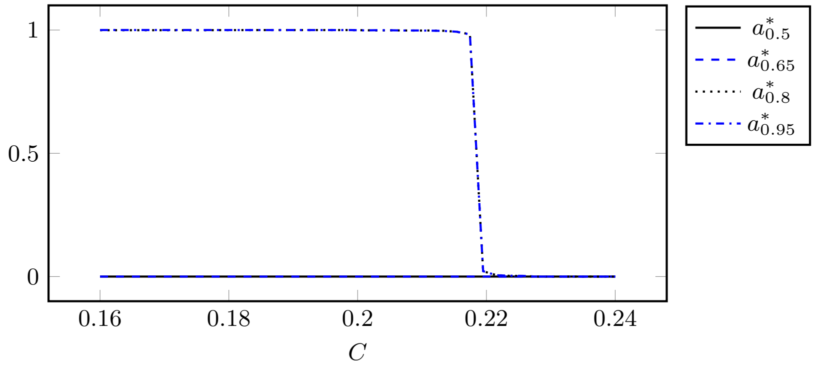

Fig. 4. Audit probabilities in equilibrium with various audit cost values.

Figure 4 shows the equilibrium auditing probabilities $\boldsymbol{a}^*$ as functions of the audit cost C. Interestingly, the results show that in an equilibrium, the insurer does not conduct audits for reported security levels equal to or less than 0.8.

For reported level 0.95, we observe a sharp threshold: the insurer always audits if the cost of the audit is less than 0.22, but never audits if it is greater.

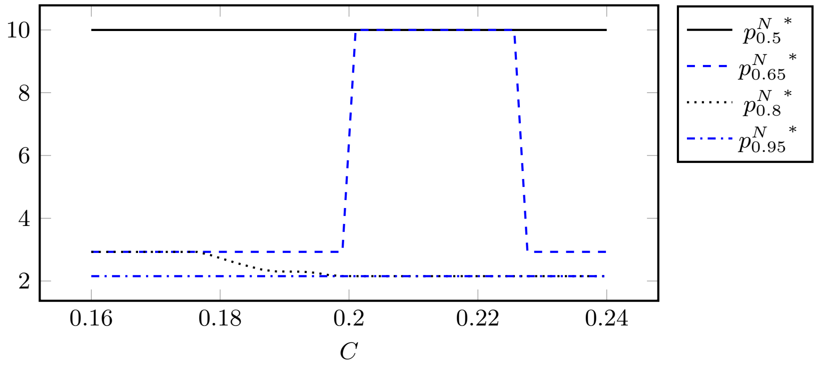

Fig. 5. Premiums (without audit) in equilibrium with various audit cost values.

Figure 5 shows the equilibrium premiums without audit $\boldsymbol{p}^{N^*}$ as functions of the audit cost C. We notice that for the lowest security level assessed, 0.5, the premium is the highest one and remains steady for the entire range of audit costs. Such a security level means having a 50% chance of getting compromised and therefore the insurer must ask for a sizable premium. For security levels 0.8 and 0.95, when audit costs are higher than 0.2, the equilibrium premiums are identical and lower than for the other two security levels, 0.5 and 0.65. In future work, we will conduct further experiments to understand the behaviour of the insurer's strategy in terms of premiums in the equilibrium, considering both the audited and non-audited cases.

5 Related Work

We discuss two classes of related work: literature on cyber-insurance—more specifically on adverse selection, moral hazard, and information asymmetries—and literature on security audits.

Cyber-insurance and Information Asymmetry: Some of the main factors that hinder cyber-risk management via cyber-insurance are *risk correlations*, *interdependence*, and *information asymmetries* [8]. Among these, we focus on information asymmetries, which are taken into account by many articles in the field of security economics and cyber-insurance [4,12,17,20].

Shetty et al. [21,22] prove a proposition providing a condition, which when satisfied, states that any insurance contract with security levels unobservable by the insurers strictly decreases the utility of the users, leading to a missing insurance market. Yet, with insurers present, and security levels contractible, in any equilibrium, full client coverage is offered.

Schwartz et al. [19] investigate the occurrence of a "lemon market" [2] when insurers cannot differentiate between different risk behaviors of clients. Lack of rich information about user choices and activities, leads to information asymmetry, which worsens the usual insurers' problems of moral hazard and adverse selection. They prove that no matter how small the fraction of malicious users is, an equilibrium does not exist, and therefore the cyber-insurance market is missing. In addition, they claim that due to adverse selection, cyber-insurers would not underwrite contracts conditioning user premiums on their security.

Bandyopadhyay et al. [5] build an economic model describing an optimal cyber-insurance contract and the optimal claim strategy for the insured firm. They show that insured firms optimally transfer more risk through insurance contracts under information symmetry than otherwise. They also present a circle of steps going from information symmetry for the cyber-insurance market to information asymmetry when the client first realizes the effect of IT risks, to end up back in information symmetry when the insurer finds out about the altered claiming and buying behavior as well as the underlying reasons.

Security Audits for Cyber-insurance Underwriting: In [3], Baer and Parkinson suggest that both insurers and clients are sophisticated in dealing with security assessments regarding cyber-insurance coverage decisions. Cyber-insurers demand audits by independent consultants on a case-by-case basis, depending on the risks to be insured and the client requirements with regards to policy limits [10]. For example, the largest cyber-insurance underwriter, called AIG, asks prospective clients to complete an "Information Security Self Assessment" online. The results of such self-assessments determine whether the insurer will undertake a security audit on client's premises to bind coverage.

Böhme [7] argues that security audits can generate positive utility overcoming information asymmetries in a scenario focused on "solving coordination problems." According to this, the players themselves decide about their own security investment and whether or not to give away information about the resulting security level. Security audits are a tangible way to derive such security levels. The author states that it is difficult to measure the security level of products due to: (i) the difficulty of specifying all security requirements; and (ii) attacks neither occur deterministically nor is their occurrence observable in real time. The hardness of measurement implies significant effort to undertake a meaningful audit and requires special knowledge and experience. The difficulty of the audit increases disproportionately to the complexity of the system due to the non-linear growth of interdependencies among different assets.

Khalili et al. [13] suggest that recent advances in Internet measurement combined with machine learning techniques enable accurate quantitative assessments of security posture at a firm level. They claim that this can be used as a tool to perform an initial security audit, or pre-screening, of a prospective client to better enable premium discrimination and the design of customized policies.

6 Conclusions

Cyber-insurers face the challenge of devising a policy that is "reasonable" for the client to purchase but profitable for the insurer as well. To elicit risk levels for premium calculations, the insurer either asks the organization to conduct some self-assessment and report it, or it undertakes an audit to identify the real security level with certainty. Further, the possibility of being audited by the insurer may incentivize the organization to report truthfully. However, such an audit introduces costs for the insurer, which may be relatively high.

We introduced a new model to study optimal strategies for self-reporting security levels (for organizations) and undertaking expensive audits (for insurers). The insurers' strategy aims to ensure that the actual security levels of their clients have been elicited and therefore "fair" contracts (coverage, premium) are put in place. More concretely, we modeled the interactions between a potential client and an insurer as a two-player signaling game, where the organization plays the role of the sender, while the insurer plays the role of the receiver. To the best of our knowledge, this paper is the first to attempt studying incentives for auditing potential clients before cyber-insurance premium calculations. The proposed model may form the basis of a framework that can further accelerate the adoption of cyber-insurance.

Our model and analyses do have certain limitations, which we intend to improve upon in future work. First, future work may allow the insurer to offer multiple levels of coverage (for different premiums) to an organization; i.e., when the insurer computes premiums, it also associates these premiums with certain degrees of loss recovery. Secondly, future work may allow the insurer to perform cyber-forensics when a claim is filed in order to reveal whether the organization was honest when reporting its security level. In this case, the insurer avoids auditing costs, but may still be able to deter clients from misreporting. We shall investigate the trade-offs between forensics and auditing costs. Furthermore, we will consider penalties for untruthful organizations. Punishment of such behavior may be realized in the form of increased premiums or reduction of some reputation metric that affects any future cyber-insurance contract of the organization. More ambitiously, our plan is to work with cyber-insurers to acquire realistic data as part of a recently funded research project.

Acknowledgement: We thank the anonymous reviewers for their comments. The research activities of Jens Grossklags are supported by DIVSI. Emmanouil Panaousis' work is supported by the H2020-MSCA-RISE-2018 SECONDO project.

References

1. ABI. Cyber risk insurance. `https://www.abi.org.uk/products-and-issues/products/business-insurance/cyber-risk-insurance/`. [Online; accessed 19-June-2017].
2. George Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, pages 488–500, 1970.
3. Walter Baer and Andrew Parkinson. Cyberinsurance in IT security management. *IEEE Security & Privacy*, 5(3):50–56, 2007.
4. Tridib Bandyopadhyay, Vijay Mookerjee, and Ram Rao. Why IT managers don't go for cyber-insurance products. *Communications of the ACM*, 52(11):68–73, 2009.
5. Tridib Bandyopadhyay, Vijay Mookerjee, and Ram Rao. A model to analyze the unfulfilled promise of cyber insurance: The impact of secondary loss. Technical report, University of Texas at Dallas, 2010.

6. Rainer Böhme. Cyber-insurance revisited. In *Workshop on the Economics of Information Security (WEIS)*, 2005.
7. Rainer Böhme. Security audits revisited. In *Proceedings of the 16th International Conference on Financial Cryptography and Data Security (FC)*, pages 129–147. Springer, 2012.
8. Rainer Böhme and Galina Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Workshop on the Economics of Information Security (WEIS)*, 2010.
9. Forbes. Worldwide cybersecurity spending increasing to $170 billion by 2020. `https://www.forbes.com/sites/stevemorgan/2016/03/09/worldwide-cybersecurity-spending-increasing-to-170-billion-by-2020/#5804298e6832`, March 2016.
10. Lawrence Gordon, Martin Loeb, and Tashfeen Sohail. A framework for using insurance for cyber-risk management. *Communications of the ACM*, 46(3):81–85, 2003.
11. HM-Government. UK cyber security: The role of insurance in managing and mitigating the risk. `https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/415354/UK_Cyber_Security_Report_Final.pdf`, June 2015.
12. Annette Hofmann. Internalizing externalities of loss prevention through insurance monopoly: An analysis of interdependent risks. *Geneva Risk and Insurance Review*, 32(1):91–111, 2007.
13. Mohammad Mahdi Khalili, Parinaz Naghizadeh, and Mingyan Liu. Designing cyber insurance policies: The role of pre-screening and security interdependence. *IEEE Transactions on Information Forensics and Security*, 13(9):2226–2239, 2018.
14. KPMG. Seizing the cyber insurance opportunity. `https://assets.kpmg.com/content/dam/kpmg/xx/pdf/2017/07/cyber-insurance-report.pdf`, July 2017.
15. Aron Laszka and Jens Grossklags. Should cyber-insurance providers invest in software security? In *Proceedings of the 20th European Symposium on Research in Computer Security (ESORICS)*, pages 483–502, 2015.
16. Aron Laszka, Benjamin Johnson, Jens Grossklags, and Mark Felegyhazi. Estimating systematic risk in real-world networks. In *Proceedings of the 18th International Conf. on Financial Cryptography and Data Security (FC)*, pages 417–435, 2014.
17. Marc Lelarge and Jean Bolot. Economic incentives to increase security in the internet: The case for insurance. In *Proceedings of the 28th IEEE International Conf. on Computer Communications (INFOCOM)*, pages 1494–1502. IEEE, 2009.
18. Philip Low. Insuring against cyber-attacks. *Computer Fraud & Security*, 2017(4):18–20, 2017.
19. Galina Schwartz, Nikhil Shetty, and Jean Walrand. Cyber-insurance: Missing market driven by user heterogeneity. Technical report, UC Berkeley, 2010.
20. Galina Schwartz, Nikhil Shetty, and Jean Walrand. Why cyber-insurance contracts fail to reflect cyber-risks. In *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 781–787. IEEE, 2013.
21. Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. Competitive cyber-insurance and internet security. In *Economics of Information Security and Privacy*, pages 229–247. Springer, 2010.
22. Nikhil Shetty, Galina Schwartz, and Jean Walrand. Can competitive insurers improve network security? In *Proceedings of the 2010 International Conference on Trust and Trustworthy Computing (TRUST)*, pages 308–322. Springer, 2010.
23. Symantec. ISTR: Internet security threat report. `https://www.symantec.com/security-center/threat-report`, April 2017.