Bi-level optimisation of security investment and insurance pricing

Zixuan Zhang[a], Michail Chronopoulos[a,b], Ioannis Kyriakou[a,*]

[a] *Bayes Business School, City St George's, University of London, 106 Bunhill Row, London EC1Y 8TZ, United Kingdom*

[b] *Norwegian School of Economics, Department of Business and Management Science, 5045 Bergen, Norway*

Abstract

We develop a decision-support framework for cyber risk mitigation policies from the perspective of an organisation with limited resources for security controls, upgrades, and cyber insurance. To balance the conflicting optimisation objectives of the organisation and the insurer, we propose a bi-level model that endogenously derives optimal strategies for both parties, accounting for key uncertainties underlying a cyber attack. We find that cyber insurance coverage increases with premium size, though this depends on the effectiveness of system upgrades. Notably, the latter has an ambiguous impact on the equilibrium budget allocation strategy and insurance contract design, meaning that higher effectiveness does not necessarily mandate an analogous capital allocation.

Keywords: Cyber security, bi-level optimisation, insurance

1. Introduction

Advancements in computer information systems have increased the complexity of today's cyber-security environment, heightening the vulnerability of critical infrastructures to cyber attacks, while threat actors are demonstrating a significantly expanding range

[*] Corresponding author. *Competing Interests Statement:* The author(s) declare none. *Data Availability Statement:* The authors provide replication materials openly via https://www.bayes.citystgeorges.ac.uk/faculties-and-research/experts/ioannis-kyriakou. *Funding Statement:* This work was supported by the Society of Actuaries (SOA) Research Institute and the Casualty Actuarial Society (CAS) under the research grant proposal "Bi-level Optimization of Cyber Risk and Insurance Pricing". The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Email addresses: jackzhang_1997@outlook.com (Zixuan Zhang), michalis.chronopoulos@city.ac.uk (Michail Chronopoulos), ioannis.kyriakou@city.ac.uk (Ioannis Kyriakou)

of intelligence-gathering techniques (He et al., 2024). The risk exposure and financial consequences of cyber attacks for an organisation are evident in a wide range of examples. For instance, the SolarWinds hack compromised multiple government systems along with many Fortune 500 companies globally (Oladimeji and Kerner, 2023). The CryptoLocker ransomware attack caused an estimated loss of \$3 million (Kelion, 2013), and the 2016 Dyn cyber attack resulted in the disruption of major internet platforms and services for large swathes of users in Europe and North America (Hilton, 2016). More recently, the Marriott breach exposed personal details of approximately 5.2 million hotel guests (Uberti, 2020), while the Twitter breach led to fraudulent tweets about Bitcoin, generating over $\$100,000$ worth of Bitcoin deposits (Satter, 2023).

Each instance of a data breach or system failure that leads to substantial financial or reputational damage heightens awareness among decision-makers about the inadequacies of current policies in addressing cyber risks. The significant economic and societal implications of cyber risk are well-recognised (e.g., see Biener et al., 2015; Cartagena et al., 2020), emphasising the need for robust risk management solutions (e.g., see Eling and Jung, 2018; Da et al., 2021; Liu et al., 2022; Braun et al., 2023). To address the risk exposure and financial implications of cyber attacks, organisations must invest in and maintain updated security controls. These are essential for patching asset vulnerabilities, which helps minimise the expected present value of an attack's impact by reducing an asset's attack surface or increasing the effort required to breach the asset. However, delivering reliable and robust security for organisations is a capital-intensive process that typically requires a combination of various mitigation measures, and budget constraints often render this strategy economically infeasible. Therefore, to further mitigate cyber risk and improve network resiliency, organisations resort to cyber insurance (Kesan et al., 2005; Böhme and Schwartz, 2010; Shetty et al., 2010; Pal et al., 2014; Biener et al., 2015). They then face a dual challenge when improving their cyber-security posture: gauging the financial impact of cyber breaches and determining the optimal allocation of capital across defence methods and insurance.

Overcoming these challenges requires novel techniques that combine risk assessment and optimisation methods accounting for critical aspects of the attack itself, relevant underlying uncertainties, and strategic interaction between the insurer and the insureds. Key uncertainties associated with an attack include the time required to exploit a vulnerability and the extent of the associated financial impact on the targeted organisation. Both exploitation time and impact due to an attack are likely to vary randomly, as they depend not only on the skills of the attacker but also on the organisation's level of cyber preparedness and response (Fielder et al., 2016). For example, Advanced Persistent Threats (APTs) are origins of considerable cyber risk for organisations (Daly, 2009) that typically breach their targets in phases by exploiting a series of system-, network-, or even user-oriented vulnerabilities (Nisioti et al., 2021; Ahmed et al., 2022). The FireEye M-Trends 2020 Special report found that the mean dwell time for 2019 in the USA was 60 days and in EMEA and APAC 54 days[1].

An in-depth cyber risk assessment enables a more accurate evaluation of an organisation's security posture, helping to prevent potential denial of cyber insurance claims (Panda et al., 2019) and cycles of under- or over-investment that elevate the regulatory risk of corrective policy actions, thus supporting efficient asset-liability management (Kamiya et al., 2021; Eling and Jung, 2018). To this end, in this paper, we develop a decision-support framework for optimal cyber-security investment. This incorporates the serial nature of a cyber-security breach, the uncertainty in the time required to exploit a vulnerability, and the strategic interaction between the organisation/defender and the insurer.

We proceed with Section 2, which reviews related work, provides a detailed discussion of our contributions, and summarises our main findings. In Section 3, we outline our assumptions and notation. We then examine the firm's optimisation problem in the absence of cyber insurance, extending the analysis to allow for the interaction between the defender and the insurer, and derive the optimal insurance policy design for the

[1] `https://www.fireeye.com/current-threats/annual-threat-report/mtrends.html`.

insurer. Section 4 presents policy implications based on numerical results, while Section 5 concludes the paper offering directions for further research.

2. Related work and advancements

Cyber insurance plays a critical role in an organisation's portfolio of mitigation measures, making the interactions between insurers and insureds a key component of a cyber-security investment strategy. However, this aspect is often overlooked in the cyber-security economics literature, which primarily focuses on selecting controls to mitigate system vulnerabilities. For example, models for the optimal selection of cyber-security controls include Smeraldi and Malacaria (2014), who explore how to spend a security budget optimally by employing methods that address overlapping controls exhibiting non-linear relationships, such as optimisation algorithms, combinatorial optimisation, and the classical Knapsack problem. Fielder et al. (2016) propose a methodology for investing in such controls, considering a single value for a vulnerability and several implementation levels for each control. The latter align with the information security levels introduced in the seminal work of Gordon and Loeb (2002).

Building on prior work by Almohri et al. (2016), Khouzani et al. (2019) develop a game-theoretic framework for analysing defender-attacker interactions. In this framework, the defender chooses a plan to minimise security risk, while the attacker aims to maximise it by exploiting the most effective attack path. This is modelled as a min-max optimisation problem, where the attacker maximises and the defender minimises in response to the attacker's action. Additionally, Zheng et al. (2019) cast the problem of optimal control selection as a set covering problem. They first solve a deterministic version to examine incentives for mitigating supply chain vulnerabilities and later introduce constraints and uncertainties in control efficacy. Expanding on Fielder et al. (2016), Panda et al. (2020) propose an optimal control set for protecting healthcare employee groups from social engineering attacks. However, a limitation of these optimisation models is their failure to account for the serial nature of an attack and critical uncertainties, such as the exploitation time of a vulnerability and the associated costs once it is compromised. As a result, these models often overlook the financial implications of such uncertainties on an organisation's assets.

Game-theoretic models that analyse interactions between insurers and insureds include Grossklags et al. (2008), Laszka et al. (2018), and Wang (2019). Specifically, Laszka et al. (2018) employ a two-player signaling game to address information asymmetry between a potential client and an insurer, studying incentives for auditing clients before calculating cyber insurance premiums. In the same line of work, Wang (2019) examines the optimisation of a firm's cyber-security investment decision, whereby a firm must determine how much to invest in both knowledge and expertise, as well as in mitigation measures. The findings indicate that the effectiveness of security spending on specific threats may be diminished if other interdependent security measures are not simultaneously implemented. Insights on how cyber insurance may contribute to risk reduction training are also provided, yet cyber insurance is not directly integrated within the problem of optimal capital allocation. Also, Chong et al. (2025) emphasise the importance of conducting comprehensive cost-benefit analyses for budget-constrained firms that must make informed capital allocation decisions to achieve a balanced cyber risk management strategy integrating effectively cyber-security investment, insurance coverage, and reserving.

While the aforementioned literature considers risk mitigation through both cyber-security measures and cyber insurance, the insurer's decision-making, which, in turn, influences a company's optimal cyber-security investment, is often overlooked. This gap is addressed by Zhang and Zhu (2022), who developed a Markov model to capture cyber risk dynamics and defender decisions regarding mitigation measures, including both controls and cyber insurance. In this framework, defenders receive financial compensation from insurers for losses caused by cyber attacks in exchange for premiums. The defenders' objective is to deploy an optimal combination of controls and cyber insurance to minimise losses, favouring contracts with low premiums and high coverage. Conversely, insurers tend to offer contracts with high premiums and low coverage to maximise profits. Similar to traditional insurance, insurers lack knowledge of local protections implemented by defenders, which can result in inappropriate insurance contracts that significantly harm insurers' profitability.

Our work builds upon three key strands of literature: first, the valuation of serial projects to assess security breach risks progressing in phases (Tsiodra et al., 2023); second, the modelling of the optimal level of resources for securing information (Gordon and Loeb, 2002); and third, the strategic interactions between a defender and an insurer, as explored by Zhang and Zhu (2022). Our contribution is thus threefold. First, we extend the traditional discounted cash flow approach by accounting for key uncertainties and the impact of security upgrades on the likelihood of successful attacks. In doing so, we enhance the applicability of the discounted cash flow approach not only for investment decision-making but also for risk assessment and management in a cyber-security context. Second, we develop a bi-level model that captures the strategic interactions between the defender and the insurer. This allows the insured's decision-making to depend on the insurer's choices, and vice versa, reflecting the interdependent nature of their strategies. Third, by analysing the trade-off involved in allocating a finite budget between controls and cyber insurance, we derive endogenous strategies for both parties.

Our findings indicate that the insurance company tends to offer higher coverage when it receives a larger premium. However, this tendency also depends on the effectiveness of system upgrades. For instance, if a small investment in system upgrades significantly reduces claim frequency, the insurer might be willing to provide high coverage even with a lower premium. Conversely, when the projected frequency of cyber attacks is high, the insurer is inclined to offer lower coverage. In such cases, the defender may find it more advantageous to allocate more capital to system upgrades rather than to insurance. Interestingly, the effectiveness of system upgrades can have a non-monotonic influence on the equilibrium budget allocation strategy and insurance contract design, i.e., greater system upgrade effectiveness does not necessarily imply that the firm should allocate more resources towards them.

3. Model framework

3.1. Preliminaries

Consider an organisation/defender wishing to protect their systems from potential hackers. Assume that the defender's infrastructure consists of $n \in \mathbb{N}$ systems and networks, referred to as assets, that can be compromised by potential hackers (attackers). The adversarial interactions are modelled as a sequence of attack phases, where phase $i = 1, 2, 3, \ldots, n$ corresponds to the stage in which the attacker aims to compromise asset i by exploiting any of its vulnerabilities, as illustrated in Figure 1. We, hence, assume that in each phase the attacker can compromise only one asset, and that successful exploitation can lead to undesirable privilege escalation or lateral movement within the defender's infrastructure (Niakanlahiji et al., 2020; Provos et al., 2003).

The defender has the option to distribute budget K between enhancing the system and purchasing cyber insurance at time 0. More specifically, the defender invests wK, for $w \in [0, 1]$, in a system upgrade and $(1 - w)K$ in cyber insurance. The former aims to decrease the likelihood of cyber attacks, while the latter offers coverage for a fraction $c \in [0, 1]$ of future losses stemming from such attacks. The loss of the ith attack is L_i at time

$$T_i^w = \sum_{j=1}^{i} \tau_j^w, \tag{1}$$

where τ_j^w is the jth random inter-attack duration with probability distribution generally denoted by $G(\cdot)$ (identical for all j).

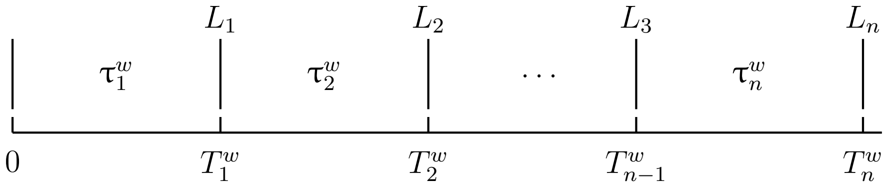

Figure 1: Sequential security breach.

The insurer determines the coverage level, $c^*(w)$, based on the capital $(1 - w)K$

the defender invests in insurance. Given the specifics of the insurance contract, the defender sets optimally the equilibrium budget allocation strategy, $\widetilde{w}$, with corresponding equilibrium coverage level

$$\widetilde{c} = c^* (\widetilde{w}) .$$

Our framework can flexibly accommodate general duration probability distributions. Consistent with Bentley et al. (2020), we adopt the intuitive compound Poisson process with arrival intensity λ to model the impact of mitigations on attack frequency. Following a system upgrade, the likelihood of successful cyber attacks diminishes, and the arrival intensity becomes $f(w)\lambda$, where $0 < f(\cdot) < 1$ depends on the invested funds. Aligning the mitigation models discussed in Gordon and Loeb (2002) with our context yields

$$f(w) = \frac{1}{(aw+1)^b}, \tag{2}$$

where $a > 0$ and $b \geq 1$ are parameters associated with the capital invested in system upgrades. A higher value of a or b represents greater effectiveness of the system upgrade.

3.2. Equilibrium analysis

This section presents the analytical framework within which the objectives of the defender and the insurer are combined to yield equilibrium decisions regarding investment in system upgrades and insurance coverage. A diagrammatic overview of the bi-level framework and the resulting equilibrium is provided in Figure 2. First (Level 1), we formulate the defender's value function, which we will use to derive the capital $w^*(c)$ to be invested in system upgrades. Second (Level 2), the insurer determines the coverage amount $c^*(w)$. This is then passed as input to $w^*(c)$ to produce the equilibrium investment, $\tilde{w} \equiv w^*(c^*)$, in system upgrades and the equilibrium coverage level, $\tilde{c} \equiv c^*(\tilde{w})$.

Following from the previous section, the defender may choose to allocate wK to a system upgrade and $(1-w)K$ to purchasing cyber insurance. This allocation provides coverage for a portion of the future losses resulting from cyber attacks. In the event of the defender incurring loss L_i, the insurer reimburses cL_i, where $c \in [0, 1]$; $c = 0$ corresponds

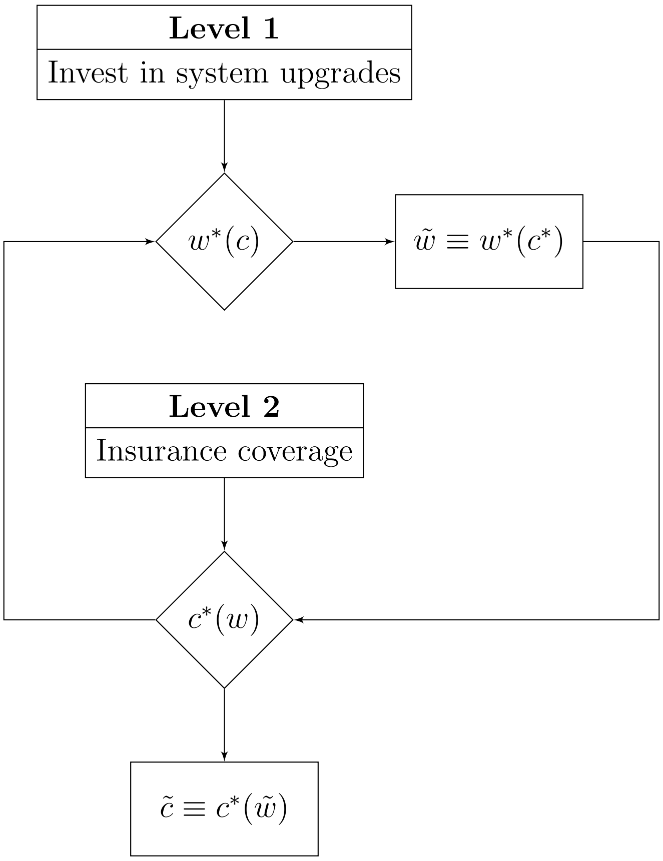
Level 1
Invest in system upgrades
$w^*(c)$
$\tilde{w} \equiv w^*(c^*)$
Level 2
Insurance coverage
$c^*(w)$
$\tilde{c} \equiv c^*(\tilde{w})$

Figure 2: Diagrammatic representation of the bi-level framework capturing the strategic interaction between the defender and the insurer.

to no coverage, while $c = 1$ to full coverage. The defender's PV of the loss in phase $i = 1, 2, 3, \ldots, n$ is

$$V_i(w) = (1-c)L_i e^{-rT_i^w}.$$

Since the arrival of attacks follows a Poisson process, the time intervals between successive attacks are exponentially distributed, i.e., $\tau_j^w \sim \text{Exponential}(f(w)\lambda)$, implying that $T_i^w \sim \text{Gamma}(i, f(w)\lambda)$ (see equation 1). The distribution and density functions of $V_i(w)$ are given, respectively, by

$$\begin{aligned} \Theta_{V_i}(v_i) &= 1 - \frac{1}{\Gamma(i)}\gamma\left(i, \frac{\lambda}{r}\ln\frac{L_i}{v_i}\right), && (3) \\ \theta_{V_i}(v_i) &= \frac{\lambda^i}{rv_i\Gamma(i)}\left(\frac{1}{r}\ln\frac{L_i}{v_i}\right)^{i-1}\left(\frac{v_i}{L_i}\right)^{\frac{\lambda}{r}}. && (4) \end{aligned}$$

(More details are deferred to the appendix.) The resulting mean and variance are

$$\begin{aligned} \mathbb{E}[V_i(w)] &= L_i\prod_{j=1}^{i}\mathbb{E}\left[e^{-r\tau_j^w}\right] = L_i\left(\frac{\lambda}{\lambda+r}\right)^i, && (5) \\ \mathbb{V}\text{ar}[V_i(w)] &= \mathbb{E}\left[V_i^2(w)\right] - \mathbb{E}\left[V_i(w)\right]^2 = L_i^2\left[\left(\frac{\lambda}{\lambda+2r}\right)^i - \left(\frac{\lambda}{\lambda+r}\right)^{2i}\right]. && (6) \end{aligned}$$

The PV over all losses is

$$V(w) = (1-c)\sum_{i=1}^{n} L_i e^{-rT_i^w}, \qquad (7)$$

with expectation

$$\mathbb{E}\left[V(w)\right] = (1-c)\sum_{i=1}^{n} L_i\left(\frac{f(w)\lambda}{f(w)\lambda+r}\right)^i. \qquad (8)$$

The defender's optimisation problem is to derive the value of w that minimises the expected loss for a given c:

$$w^*(c) = \underset{w\in[0,1]}{\text{argmin}}\, E\left[V(w)\right]. \qquad (9)$$

On the other hand, the insurer focuses on designing cyber insurance contracts. The insurer's profit is the premium revenue minus the losses ceded by the firm due to cyber attacks. Specifically, the insurer receives $(1-w)K$ at time 0, but incurs cost cL_i when the firm experiences loss L_i due to a cyber attack. The PV of the insurer's profit is given by

$$S(w) = (1-w)K - c\sum_{i=1}^{n} L_i e^{-rT_i^w}. \tag{10}$$

As suggested by (10), the PV of the insurer's profit depends on the firm's budget allocation plan w. In response, the insurer determines the level of coverage, c, based on the premium received. Here, we assume that the insurer is risk-averse and seeks to achieve a positive profit from the insurance contract with probability α, i.e., $\mathbb{P}(S > 0) = \alpha$. Intuitively, this condition implies that the premium is greater than the cost of the insurance coverage with probability α. Therefore, the insurer's required level of coverage satisfies

$$\begin{aligned}(1-w)K &= \inf\left\{Z \in \mathbb{R} : \mathbb{P}\left(Z \le v\right) \ge \alpha\right\} \\ &= \mathrm{VaR}_\alpha\left(Z\right),\ Z = c\sum_{i=1}^{n} L_i e^{-rT_i^w},\end{aligned} \tag{11}$$

where Value at Risk (VaR) measures riskiness by examining the left tail of the PV distribution and is positively homogenous. A confidence level $0 \le \alpha \le 1$ reflects the insurer's level of risk-aversion, with a larger (smaller) α indicating a more (less) conservative insurer. Without loss of generality, we adopt the VaR as the risk measure; however, this is not restrictive, and other risk measures or utility functions may be used. Rearranging (11), we obtain the insurer's required level of coverage as a function of w:

$$c^*(w) = \frac{(1-w)K}{\mathrm{VaR}_\alpha\left(\sum_{i=1}^{n} L_i e^{-rT_i^w}\right)}. \tag{12}$$

By substituting (12) into equation (7), the PV of the firm's loss becomes

$$V^*(w) = \left(1 - \frac{(1-w)K}{\mathrm{VaR}_\alpha\left(\sum_{i=1}^n L_i e^{-rT_i^w}\right)}\right)\sum_{i=1}^n L_i e^{-rT_i^w}, \tag{13}$$

and the equilibrium budget allocation strategy then follows as

$$\widetilde{w} = \underset{w\in[0,1]}{\mathrm{argmin}}\, \mathbb{E}\left[V^*(w)\right]. \tag{14}$$

Finally, the equilibrium insurance coverage, $\widetilde{c} = c^*(\widetilde{w})$, is obtained.

4. Budget allocation and frequency reduction effects on cyber insurance and system upgrades: a numerical study

This section explores the effects of budget allocation ratios, attack frequency, and the frequency reduction parameter on the equilibrium strategies of a defender and an insurer. We examine how these factors influence the insurance coverage level and the expected present value of losses. We highlight the interplay between system upgrades and insurance, revealing non-monotonic relationships and strategic trade-offs that arise from variations in attack frequency and system upgrade effectiveness.

We begin by exploring how the allocation of resources between system upgrades and cyber insurance influences key outcomes, such as insurance coverage levels and expected losses. Table 1 illustrates the impact of the exogenous budget allocation ratio w on optimal investment in insurance coverage, the expected PV of losses retained by the defender and the VaR of losses transferred to the insurer. For example when $a = 0.5$, the insurer provides higher coverage as the premium $(1-w)K$ increases. However, this increased insurance coverage does not necessarily lead to smaller losses from cyber attacks for the firm. In fact, we observe a non-monotonic relationship with w, particularly for high attack frequencies (see cases $\lambda = 1$ or 2).

Given these first remarks, in Figure 3 we more closely examine how the insurance coverage level (left panel) and the expected PV of losses (right panel) vary with w for

	$c^*(w)$			Expected PV			$\mathrm{VaR}_{0.95}$		
	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$
$w = 0$	0.6442	0.3600	0.1972	1.7789	6.4000	16.0570	2.7898	4.9525	9.1040
$w = 0.25$	0.5305	0.2995	0.1639	2.0867	6.2267	14.8644	3.3035	8.0023	14.5815
$w = 0.5$	0.3839	0.2178	0.1200	2.4642	6.2576	14.0795	4.0214	5.3794	16.7193
$w = 0.75$	0.2068	0.1182	0.0652	2.8843	6.4132	13.5977	4.7814	7.4354	8.9871
$w = 1$	0	0	0	3.3333	6.6667	13.3333	5.6318	6.0900	12.4678

Table 1: Impact of budget allocation ratio, w, on optimal insurance coverage level, the expected PV of losses retained by the defender and the VaR of losses transferred to the insurer. Parameter values used are $r = 0.1$, $K = 5$, $L = 1$, $a = 0.1K/L$, $b = 1$, $\alpha = 0.95$.

different values of a. The upper panel reveals a notable trend: when a is small (i.e., the effectiveness of a system upgrade is low), a decreases in w (that is, an increase in the budget proportion allocated to purchasing insurance) leads to an increase in the level of insurance coverage. This occurs because a low a value implies that investing in a system upgrade yields only marginal reductions in the frequency of cyber attacks and subsequent losses, making insurance a more efficient option. Additionally, the insurer is inclined to offer more extensive coverage when a higher premium is charged. However, as shown in the top-left panel, the coverage level also depends on the frequency of cyber attacks. Intuitively, proportional coverage becomes costly for the insurer when the attack frequency is high; consequently, a lower coverage level is set in such a case.

Interestingly, the top-right panel demonstrates that when a is small and the frequency of cyber attacks (λ) is low (high), opting for insurance (a system upgrade) becomes a more appealing prospect for the defender. This inclination arises because a high λ prompts the insurer to offer minimal coverage in the absence of a system upgrade. In such a case, investing in system upgrades results in a more significant reduction in expected losses from cyber attacks compared to purchasing insurance. Consequently, the equilibrium budget allocation ratio $\widetilde{w}$ is close to 1. Conversely, when λ is low, the expected number of cyber attacks and associated loss remain minimal even if the company does not invest in self-protection. Under these circumstances, the insurer is willing to provide a higher coverage, making insurance a more appealing investment for the company. Thus, $\widetilde{w}$ approaches 0.

When a is large, the impact of w on the insurer's decision and the expected PV of losses becomes more ambiguous, as exhibited in the bottom panel of Figure 3. As shown in (2), larger a implies that investing in a system upgrade can lead to a more significant reduction in the frequency of future cyber attacks. Interestingly, the bottom-left panel indicates that as w increases, the insurer may be willing to offer better insurance coverage (for $w < 0.25$) even if it receives a smaller premium. This counter-intuitive result can be attributed to the fact that the insurer benefits from either increased premium or reduced total claim amount, as indicated in (10). When a is large, the decrease in the expected PV of future claims resulting from the defender's investment in system upgrades surpasses the comparatively smaller premium received. The bottom-right panel also shows that the firm is more likely to benefit from a bigger investment in system upgrades when a is large.

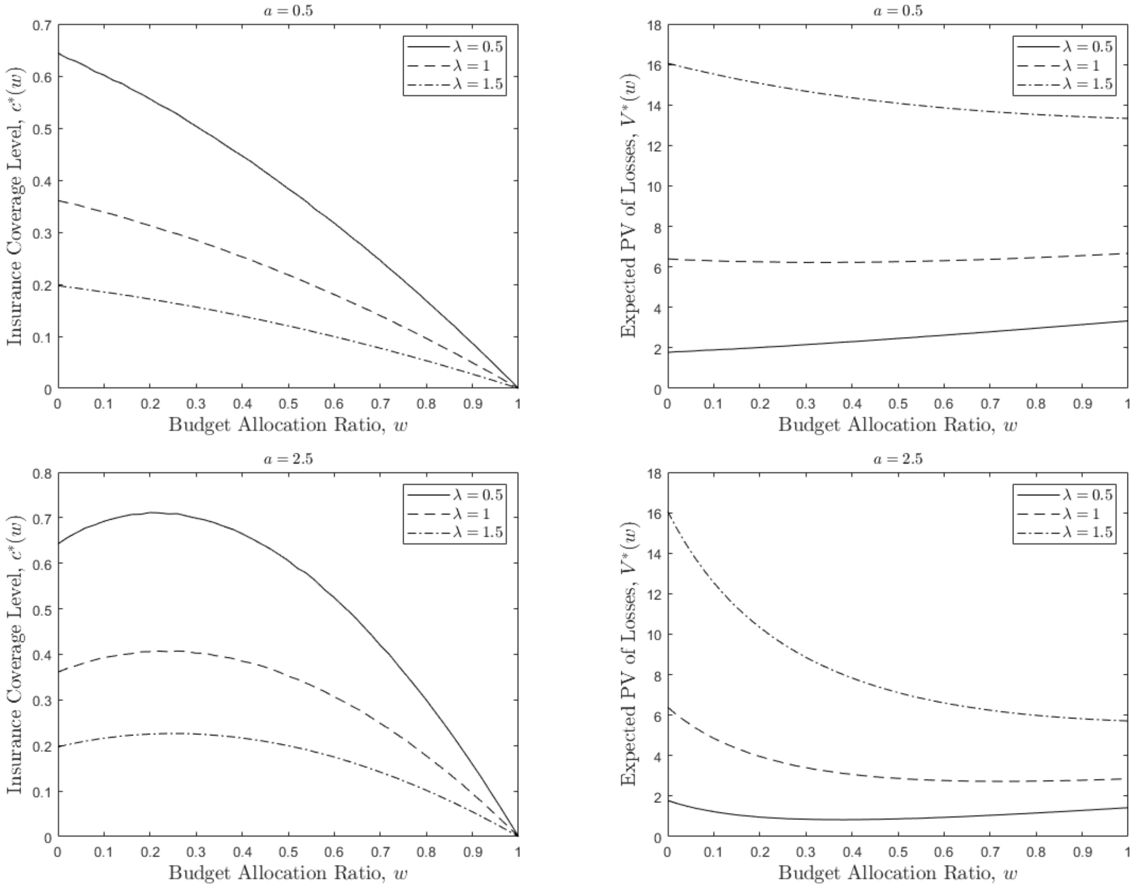

Figure 3: Impact of exogenous budget allocation ratio, w, on the insurance coverage level (left) and the expected PV of losses (right) for $a = 0.5$ (top) and $a = 2.5$ (bottom).

In Figure 4, we examine the influence of cyber attack frequency on the equilibrium strategies of both the defender and the insurer. The left panel shows a decline in the equilibrium insurance coverage level as the frequency of cyber attacks increases. Notably, this coverage level approaches zero when λ becomes exceedingly high. This is because an increase in λ raises both the expected number of cyber attacks experienced by the firm and the claims processed by the insurer. To counterbalance this escalating claim frequency and amount, the insurer may choose to either increase the premium charged or decrease the coverage ratio. However, when λ is high, the premium (see the first term in equation 10) becomes relatively small compared to the claim amount (see the second term in equation 10), making higher premiums less effective. More importantly, this reduces the budget available for system upgrades, leading to weaker frequency reduction. Consequently, the insurer benefits more by offering lower coverage, enabling the firm to allocate more funds for system upgrades. This, in turn, helps curb the frequency of cyber attacks, ultimately benefiting the insurer as well.

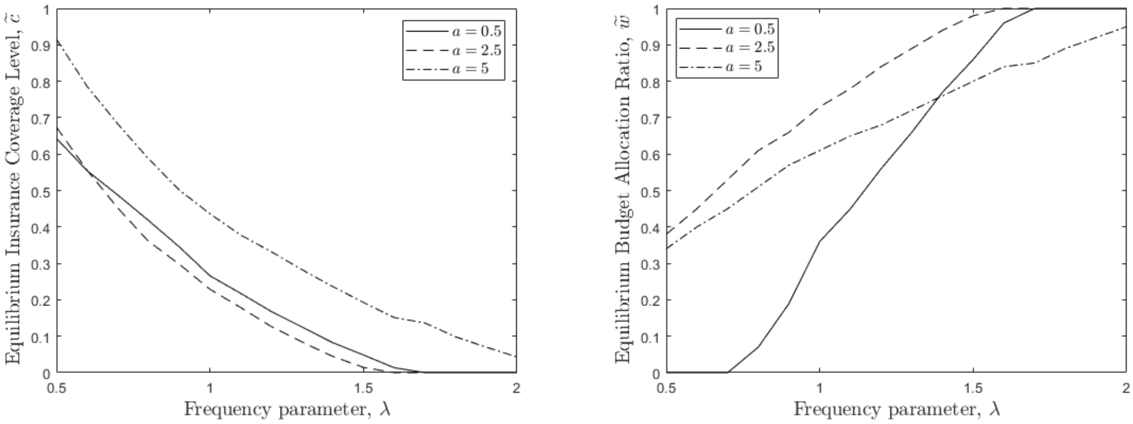

Figure 4: Impact of frequency parameter, λ, on insurance coverage level (left) and equilibrium budget allocation ratio (right).

From the right panel of Figure 4, the equilibrium budget allocation ratio $\widetilde{w}$ increases with λ in all cases, indicating that the firm invests more in system upgrades as the frequency of attacks rises. For example, when $a = 0.5$, the firm tends to allocate its entire budget to purchasing insurance (system upgrades) when $\lambda < 0.45$ ($\lambda > 1.7$). As discussed earlier, when λ is low, the insurer is willing to offer substantial coverage for losses, such as $\widetilde{c} = 65\%$ for $a = 0.5$ and $\widetilde{c} = 90\%$ for $a = 5$. This makes investing in insurance a more attractive option for the firm. However, as the frequency of attacks increases, the insurer has less incentive to provide high coverage levels, even with high premiums. Consequently, the effectiveness of loss reduction through insurance diminishes, making it more advantageous for the firm to allocate a larger portion of its budget to system upgrades. When λ becomes extremely high, the insurance company provides minimal coverage, and $\widetilde{w}$ approaches 1.

Finally, we investigate the impact of the frequency reduction parameter a. As illustrated in Figure 5, there is a non-monotonic relationship between a and the equilibrium strategies of both the firm and the insurer. Specifically, the equilibrium budget allocation ratio (coverage level) initially rises (falls) and then decreases (increases) with increasing a. Intuitively, when a is small, investing in system upgrades does not significantly reduce the frequency of future cyber attacks. Thus, the firm must allocate more resources to purchasing insurance, resulting in a lower value for $\widetilde{w}$. However, as a increases, the system upgrade effectiveness in reducing losses becomes more pronounced; even a small increase in w can substantially decrease the frequency of future attacks, as implied by (2). Consequently, the firm may decide to allocate a larger budget to these upgrades. In response to the marked decrease in premiums, the insurer may reduce the coverage level. When a reaches a high value, the projected frequency of attacks diminishes significantly, potentially approaching zero. This limits the possibility for further loss reduction despite additional investments in system upgrades. Conversely, the insurer faces reduced claim amounts and is inclined to offer higher coverage. Therefore, higher coverage through increased premiums (see right panel for $a > 1.25$) could outweigh the marginal reduction in attack frequency, causing $\widetilde{w}$ to decrease as a increases.

5. Conclusions

In today's digital landscape, cyber insurance has become increasingly essential due to the growing threat of cyber attacks and data breaches. It provides businesses with

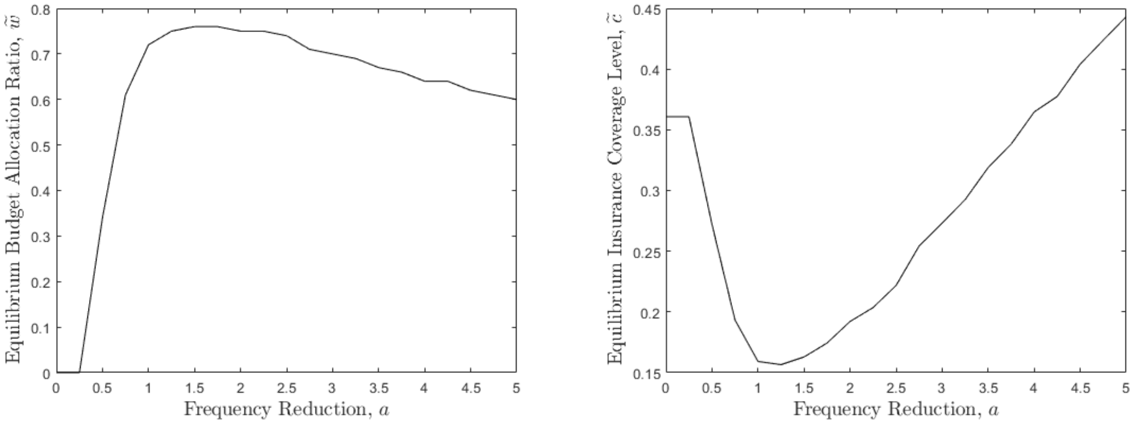

Figure 5: Impact of attack frequency reduction parameter, a, on equilibrium budget allocation ratio (left) and insurance coverage level (right).

financial protection in the event of a cyber incident, helping to mitigate costs, such as forensic investigations, legal fees, customer notifications, and credit monitoring for affected individuals. Without insurance, these expenses can be substantial and potentially devastating for a business. Additionally, cyber insurance incentivises businesses to adopt robust cyber-security measures and protocols. Insurers often require policyholders to meet specific security standards, such as conducting regular assessments and providing employee training, to qualify for coverage. By encouraging proactive risk management practices, cyber insurance reduces the likelihood and severity of cyber incidents.

In this paper, we examine a firm tasked with allocating its limited resources between upgrading its security infrastructure and purchasing cyber insurance. By assessing the risks associated with security breaches and considering the uncertainty in the time required to exploit vulnerabilities in the firm's security infrastructure, as well as the strategic interactions between the firm and an insurer, we derive the optimal strategies for both parties endogenously. Our findings indicate that insurance coverage tends to increase with a higher premium; however, this relationship depends on the system upgrade effectiveness. If a minor investment in system upgrades results in a significant reduction in claim frequency, the insurer may still offer high coverage even if the premium decreases. Conversely, when the frequency of cyber attacks is high, the insurer provides lower coverage, prompting the firm to allocate more capital to system upgrades rather than insurance.

Furthermore, the system upgrade effectiveness can exert a non-monotonic influence on the equilibrium budget allocation strategy and insurance contract design.

Future research directions could involve extending our framework to incorporate alternative optimisation objectives, utilising risk measures, such as Value at Risk and Conditional Value at Risk (CVaR). This would enable an analysis of how risk preferences influence the optimal budget allocation problem, particularly in relation to the decision-maker's level of risk-aversion. Additionally, a utility-based approach could be adopted to quantify these preferences and describe the objective functions of different market participants. Last but not least, the pricing of cyber insurance is inherently complex, as the dynamic and evolving nature of cyber threats undermines the reliability of historical data for forecasting future losses (e.g., expected loss or VaR). Future enhancements in contract design may incorporate more advanced underwriting practices, dynamic pricing, and exclusions, as well as information asymmetry, adverse selection or negotiation (Wang, 2019; Awiszus et al., 2023; Arce et al., 2024).

Appendix

Define, for any $i \geq 1$, $T_i = \tau_1 + \tau_2 + \cdots + \tau_i$ with general distribution function $F_{T_i}(\cdot)$. Consider $i = 1$. We have for $V_1 = L_1 e^{-rT_1}$ that

$$\Theta_{V_1}(v) = \mathbb{P}\left(L_1 e^{-rT_1} \leq v\right) = \mathbb{P}\left(T_1 \geq \frac{1}{r}\ln\frac{L_1}{v}\right) = 1 - F_{T_1}\left(\frac{1}{r}\ln\frac{L_1}{v}\right).$$

Assuming $\tau_j \sim \text{Exponential}(\lambda)$ for all j, we get that

$$\Theta_{V_1}(v) = 1 - \left(1 - e^{-\frac{\lambda}{r}\ln\frac{L_1}{v}}\right) = \left(\frac{v}{L_1}\right)^{\frac{\lambda}{r}}, \tag{15}$$

with associated density function

$$\theta_{V_1}(v) = \frac{\lambda}{r} L_1^{-\frac{\lambda}{r}} v^{\frac{\lambda}{r}-1}, \tag{16}$$

and resulting mean and variance

$$\mathbb{E}\left[V_1\right] = \int_0^{L_1} v\theta_{V_1}(v) = \frac{\lambda}{\lambda+r}L_1, \tag{17}$$

$$\mathbb{V}\mathrm{ar}\left[V_1\right] = \int_0^{L_1} \left(v-\mathbb{E}\left[V_1\right]\right)^2 \theta_{V_1}(v) = \left[\frac{\lambda}{\lambda+2r} - \left(\frac{\lambda}{\lambda+r}\right)^2\right] L_1^2. \tag{18}$$

For the general n-phase attack, $V_n = L_n e^{-rT_n}$ with

$$\Theta_{V_n}(v) = 1 - F_{T_n}\left(\frac{1}{r}\ln\frac{L_n}{v}\right).$$

Since $T_n \sim \mathrm{Gamma}(n,\lambda)$, we get that

$$\Theta_{V_n}(v) = 1 - \frac{1}{\Gamma(n)}\gamma\left(n, \frac{\lambda}{r}\ln\frac{L_n}{v}\right), \tag{19}$$

$$\theta_{V_n}(v) = \frac{\lambda^n}{\Gamma(n)}\left(\frac{1}{r}\ln\frac{L_n}{v}\right)^{n-1} e^{-\frac{\lambda}{r}\ln\frac{L_n}{v}}\frac{1}{rv} = \frac{\lambda^n}{rv\Gamma(n)}\left(\frac{1}{r}\ln\frac{L_n}{v}\right)^{n-1}\left(\frac{v}{L_n}\right)^{\frac{\lambda}{r}},$$

from which

$$\mathbb{E}\left[V_n\right] = L_n\prod_{j=1}^{n} E\left[e^{-r\tau_j^w}\right] = L_n\left(\frac{\lambda}{\lambda+r}\right)^n \text{ and} \tag{20}$$

$$\mathbb{V}\mathrm{ar}\left[V_n\right] = L_n^2\left[\left(\frac{\lambda}{\lambda+2r}\right)^n - \left(\frac{\lambda}{\lambda+r}\right)^{2n}\right]. \tag{21}$$

References

Ahmed, M., Panda, S., Xenakis, C., Panaousis, E., 2022. MITRE ATT&CK-driven cyber risk assessment, in: Proceedings of the 17th International Conference on Availability, Reliability and Security, pp. 1–10.

Almohri, H.M., Watson, L.T., Yao, D., Ou, X., 2016. Security optimization of dynamic networks with probabilistic graph modeling and linear programming. IEEE Transactions on Dependable and Secure Computing 13, 474 – 487.

Arce, D., Woods, D.W., Böhme, R., 2024. Economics of incident response panels in cyber insurance. Computers & Security 140, 103742.

Awiszus, K., Knispel, T., Penner, I., Svindland, G., Voß, A., Weber, S., 2023. Modeling and pricing cyber insurance. European Actuarial Journal 13, 1–53.

Bentley, M., Stephenson, A., Toscas, P., Zhu, Z., 2020. A multivariate model to quantify and mitigate cybersecurity risk. Risks 8, 61.

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

Böhme, R., Schwartz, G., 2010. Modeling cyber-insurance: Towards a unifying framework, in: Workshop on the Economics of Information Security, pp. 1–36.

Braun, A., Eling, M., Jaenicke, C., 2023. Cyber insurance-linked securities. ASTIN Bulletin 53, 684–705.

Cartagena, S., Gosrani, V., Grewal, J., Pikinska, J., 2020. Silent cyber assessment framework. British Actuarial Journal 25, e2.

Chong, W.F., Feng, R., Hu, H., Zhang, L., 2025. Cyber risk assessment for capital management. Journal of Risk and Insurance 92, 424–471.

Da, G., Xu, M., Zhao, P., 2021. Multivariate dependence among cyber risks based on L-hop propagation. Insurance: Mathematics and Economics 101, 525–546.

Daly, M.K., 2009. Advanced persistent threat. Usenix 4, 2013–2016.

Eling, M., Jung, K., 2018. Copula approaches for modeling cross-sectional dependence of data breach losses. Insurance: Mathematics and Economics 82, 167–180.

Fielder, A., Panaousis, E., Malacaria, P., Hankin, C., Smeraldi, F., 2016. Decision support approaches for cyber security investment. Decision Support Systems 86, 13–23.

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

Grossklags, J., Christin, N., Chuang, J., 2008. Secure or insure?: a game-theoretic analysis of information security games, in: Proceedings of the 17th International Conference on World Wide Web, Association for Computing Machinery, New York, USA. p. 209–218.

He, R., Jin, Z., Li, J.S.H., 2024. Modeling and management of cyber risk: a cross-disciplinary review. Annals of Actuarial Science 18, 270–309.

Hilton, S., 2016. Dyn analysis summary of Friday October 21 attack. Dyn blog 26.

Kamiya, S., Kang, J.K., Kim, J., Milidonis, A., Stulz, R.M., 2021. Risk management, firm reputation, and the impact of successful cyberattacks on target firms. Journal of Financial Economics 139, 719–749.

Kelion, L., 2013. Cryptolocker ransomware has 'infected about 250,000 pcs'. BBC News techology.

Kesan, J.P., Majuca, R.P., Yurcik, W., 2005. Cyber-insurance as a market-based solution to the problem of cybersecurity, in: Workshop on the Economics of Information Security, pp. 1–46.

Khouzani, M., Liu, Z., Malacaria, P., 2019. Scalable min-max multi-objective cyber-security optimisation over probabilistic attack graphs. European Journal of Operational Research 278, 894–903.

Laszka, A., Panaousis, E., Grossklags, J., 2018. Cyber-insurance as a signaling game: Self-reporting and external security audits, in: Bushnell, L., Poovendran, R., Başar, T. (Eds.), Decision and Game Theory for Security, Springer International Publishing, Cham. pp. 508–520.

Liu, J., Li, J., Daly, K., 2022. Bayesian vine copulas for modelling dependence in data breach losses. Annals of Actuarial Science 16, 401–424.

Niakanlahiji, A., Wei, J., Alam, M.R., Wang, Q., Chu, B.T., 2020. Shadowmove: A stealthy lateral movement strategy, in: 29th USENIX Security Symposium (USENIX Security 20), USENIX Association, Washington, D.C.. pp. 559–576.

Nisioti, A., Loukas, G., Rass, S., Panaousis, E., 2021. Game-theoretic decision support for cyber forensic investigations. Sensors 21, 5300.

Oladimeji, S., Kerner, S.M., 2023. SolarWinds hack explained: Everything you need to know. URL: https://www.techtarget.com/whatis/feature/SolarWinds-hack-explained-Everything-you-need-to-know.

Pal, R., Golubchik, L., Psounis, K., Hui, P., 2014. Will cyber-insurance improve network security? a market analysis, in: IEEE INFOCOM 2014 - IEEE Conference on Computer Communications, pp. 235–243.

Panda, S., Panaousis, E., Loukas, G., Laoudias, C., 2020. Optimizing investments in cyber hygiene for protecting healthcare users, in: Di Pierro, A., Malacaria, P., Nagarajan, R. (Eds.), From Lambda Calculus to Cybersecurity Through Program Analysis: Essays Dedicated to Chris Hankin on the Occasion of His Retirement. Springer International Publishing, Cham, pp. 268–291.

Panda, S., Woods, D.W., Laszka, A., Fielder, A., Panaousis, E., 2019. Post-incident audits on cyber insurance discounts. Computers & Security 87, 101593.

Provos, N., Friedl, M., Honeyman, P., 2003. Preventing privilege escalation, in: 12th USENIX Security Symposium (USENIX Security 03), USENIX Association, Washington, D.C.. pp. 231–241.

Satter, R., 2023. Twitter hacked, 200 million user email addresses leaked, researcher says. URL: https://www.reuters.com/technology/twitter-hacked-200-million-user-email-addresses-leaked-researcher-says-2023-01-05/.

Shetty, N., Schwartz, G., Felegyhazi, M., Walrand, J., 2010. Competitive cyber-insurance and internet security, in: Moore, T., Pym, D., Ioannidis, C. (Eds.), Economics of Information Security and Privacy, Springer, Boston. pp. 229–247.

Smeraldi, F., Malacaria, P., 2014. How to spend it: optimal investment for cyber security, in: Proceedings of the 1st International Workshop on Agents and CyberSecurity, Association for Computing Machinery, New York. pp. 1–4.

Tsiodra, M., Panda, S., Chronopoulos, M., Panaousis, E., 2023. Cyber risk assessment and optimization: A small business case study. IEEE Access 11, 44467–44481.

Uberti, D., 2020. Marriott reveals breach that exposed data of up to 5.2 million customers. URL: https://www.wsj.com/articles/marriott-reveals-breach-that-exposed-data-of-up-to-5-2-million-customers-11585686590?reflink=desktopwebshare_permalink.

Wang, S.S., 2019. Integrated framework for information security investment and cyber insurance. Pacific-Basin Finance Journal 57, 101173.

Zhang, R., Zhu, Q., 2022. Optimal cyber-insurance contract design for dynamic risk management and mitigation. IEEE Transactions on Computational Social Systems 9, 1087–1100.

Zheng, K., Albert, L.A., Luedtke, J.R., Towle, E., 2019. A budgeted maximum multiple coverage model for cybersecurity planning and management. IISE Transactions 51, 1303–1317.