Annals of Actuarial Science (2026), 1–15
doi:10.1017/S1748499526100311

ORIGINAL RESEARCH PAPER

Bi-level optimization of security investment and insurance pricing

Zixuan Zhang[1], Michail Chronopoulos[1,2] and Ioannis Kyriakou[1]

[1]Faculty of Actuarial Science and Insurance, Bayes Business School, City St George's, University of London, United Kingdom; and [2]Department of Business and Management Science, Norwegian School of Economics, Norway
Corresponding author: Ioannis Kyriakou; Email: ioannis.kyriakou@city.ac.uk

(Received 21 February 2025; revised 21 April 2026; accepted 24 April 2026)

Abstract

We develop a decision-support framework for cyber risk mitigation policies from the perspective of an organization with limited resources for security controls, upgrades, and cyber insurance. To balance the conflicting optimization objectives of the organization 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, such that a more effective upgrade need not attract a commensurately larger budget allocation. We further show that information asymmetry regarding the insurer's risk aversion can lead the defender to a suboptimal budget allocation, resulting in higher realized losses relative to the symmetric-information benchmark.

Keywords: Bi-level optimization; cyber security; insurance

1. Introduction

Advancements in computer information systems have increased the complexity of today's cybersecurity environment, heightening the vulnerability of critical infrastructures to cyber attacks, while threat actors are deploying an increasingly broad range of intelligence-gathering techniques (He *et al.*, 2024). The risk exposure and financial consequences of cyber attacks on organizations are illustrated by 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 of the inadequacies of current policies in addressing cyber risks. The significant economic and societal implications of cyber risk are well-recognized (e.g., see Biener *et al.*, 2015; Cartagena *et al.*, 2020), emphasizing the need for robust risk management solutions (e.g., see Eling & 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, organizations must invest in and maintain up-to-date security controls. These are essential for patching

asset vulnerabilities, helping to minimize the expected present value (PV) 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 organizations 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, organizations resort to cyber insurance (Kesan *et al.,* 2005; Böhme & Schwartz, 2010; Shetty *et al.,* 2010; Pal *et al.,* 2014; Biener *et al.,* 2015). They then face a dual challenge in improving their cybersecurity 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 optimization 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 organization. Both the exploitation time and the impact of an attack are likely to vary randomly, as they depend not only on the skills of the attacker but also on the organization's level of cyber preparedness and response (Fielder *et al.*, 2016). For example, Advanced Persistent Threats (APTs) are origins of considerable cyber risk for organizations (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 organization'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 & Jung, 2018). To this end, in this paper, we develop a decision-support framework for optimal cybersecurity investment. This incorporates the serial nature of a cybersecurity breach, the uncertainty in the time required to exploit a vulnerability, and the strategic interaction between the organization/defender and the insurer, who, due to possible information asymmetry, may exhibit different attitudes toward risk.

The remainder of the paper is organized as follows. Section 2 reviews the related literature on cybersecurity investment, cyber insurance, and information asymmetry, and positions our contribution within these strands. Section 3 introduces the model framework, outlining the assumptions and notation. We then examine the firm's optimization problem in the absence of cyber insurance, extend the analysis to allow for the interaction between the defender and the insurer, and derive the optimal insurance policy design for the insurer. Section 4 presents our numerical analysis, which explores equilibrium budget allocation, insurance coverage, and expected losses under varying attack frequencies, system upgrade effectiveness, and likely information asymmetry. Section 5 concludes the paper by summarizing the main findings and discussing implications for cyber insurance design, along with directions for future research.

2. Related work and advancements

Cyber insurance plays a critical role in an organization's portfolio of mitigation measures, making the interactions between insurers and insureds a key component of a cybersecurity investment strategy. However, this aspect is often overlooked in the cybersecurity economics literature, which primarily focuses on selecting controls to mitigate system vulnerabilities. For example, models for the optimal selection of cybersecurity controls include Smeraldi and Malacaria (2014), who

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

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

explore how to spend a security budget optimally by employing methods that address overlapping controls that exhibit nonlinear relationships, such as optimization algorithms, combinatorial optimization, 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 analyzing defender-attacker interactions. In this framework, the defender chooses a plan to minimize security risk, while the attacker aims to maximize it by exploiting the most effective attack path. This is modeled as a min–max optimization problem, where the attacker maximizes and the defender minimizes 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 optimization 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 organization's assets.

Game-theoretic models that analyze 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 optimization of a firm's cybersecurity 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; however, cyber insurance is not directly integrated into the problem of optimal capital allocation. Similarly, Chong *et al.* (2025) emphasize 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, effectively integrating cybersecurity investment, insurance coverage, and reserving.

Further complicating the strategic interaction between insurers and insureds is information asymmetry, whereby the two parties do not have access to the same information. Within the context of cyber security and cyber insurance, the insurer may lack information regarding the applications and software products installed by network users, as well as regarding the users' network usage habits (Moore, 2010; Böhme, 2010; Pal *et al.*, 2014). There are many obstacles for an insurer in obtaining reliable information about the risk exposure of an insured, and even more obstacles in ensuring that this exposure is maintained at the specified level throughout the policy period. For instance, Pal (2012) addresses information asymmetry in cyber insurance by analyzing three distinct scenarios: mutual ignorance, where neither the insured nor the insurer has information about the insured's cybersecurity investment level; post-contract information acquisition, in which the insurer lacks initial information but the insured gains it after signing the contract; and pre-contract information acquisition, where the insurer remains uninformed while the insured obtains relevant details before entering the contract. In our paper, we explore a critical understudied type of information asymmetry in cyber insurance: the insured has no way of knowing how risk-averse the insurer actually is. This can create a mismatch between what the defender believes about the insurer's risk attitude and the insurer's true level of risk aversion.

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

While the aforementioned literature considers risk mitigation through both cybersecurity measures and cyber insurance, the insurer's decision-making, which in turn influences a company's optimal cybersecurity investment, is often overlooked. This gap is addressed by Zhang and Zhu (2022), who develop a Markov model to capture the 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 minimize losses, favoring contracts with low premiums and high coverage. Conversely, insurers tend to offer contracts with high premiums and low coverage to maximize 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 modeling of the optimal level of resources for securing information (Gordon & 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 framework's applicability not only for investment decision-making but also for risk assessment and management in a cybersecurity context. Second, we develop a bi-level model that captures the strategic interactions between the defender and the insurer. This allows the insureds' decision-making to depend on the insurers' choices, and vice versa, reflecting the interdependent nature of their strategies. Third, by analyzing 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 toward them.

3. Model framework

3.1 Preliminaries

Let the defender's infrastructure consist of $n \in \mathbb{N}$ systems and networks, referred to as assets, that can be compromised by potential hackers (attackers). Each asset $i \in \mathbb{N}$ has $m_i \in \mathbb{N}$ vulnerabilities, that is, software weaknesses (see https://cwe.mitre.org/index.html) that the attacker may exploit. This reflects real-world attacker behavior, where adversaries aim to penetrate as deeply into a network as possible to maximize their expected return from an attack.

These strategic interactions are modeled as a sequence of attack phases, where phase i of an attack refers to the stage in which the attacker aims to compromise asset i by exploiting any of its m_i vulnerabilities, as illustrated in Figure 1. In each phase, the attacker can compromise at most one asset, with successful exploitation potentially leading to privilege escalation or lateral movement within the defender's infrastructure (Niakanlahiji *et al.,* 2020).

The expected impact on the defender from the exploitation of asset i is denoted by L_i. Following standard risk assessment principles Whitman and Mattord (2011), L_i is expressed in terms of

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

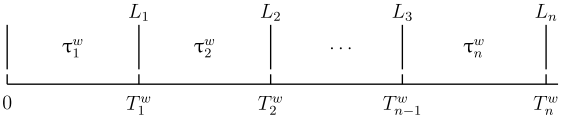

Figure 1. Sequential security breach.

attack likelihood, compromise probability, and loss magnitude. Specifically,

$$L_i = A_i \cdot \langle R_i, S_i \rangle, \tag{1}$$

where A_i denotes the value of asset i, R_i is an m_i-dimensional vector of likelihoods that the attacker attempts to exploit the m_i vulnerabilities of asset i, and S_i is an m_i-dimensional vector of probabilities that each vulnerability is successfully breached. The inner product $\langle R_i, S_i \rangle$ thus represents the overall likelihood of a successful attack against asset i.

The defender has the option of distributing 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, where c is determined endogenously by the insurer (see Section 3.2). The time at which the loss, L_i, of the ith attack is incurred is

$$T_i^w = \sum_{j=1}^{i} \tau_j^w, \tag{2}$$

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

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 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

$$f(w) = \frac{1}{(aw + 1)^b}, \tag{3}$$

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 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

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

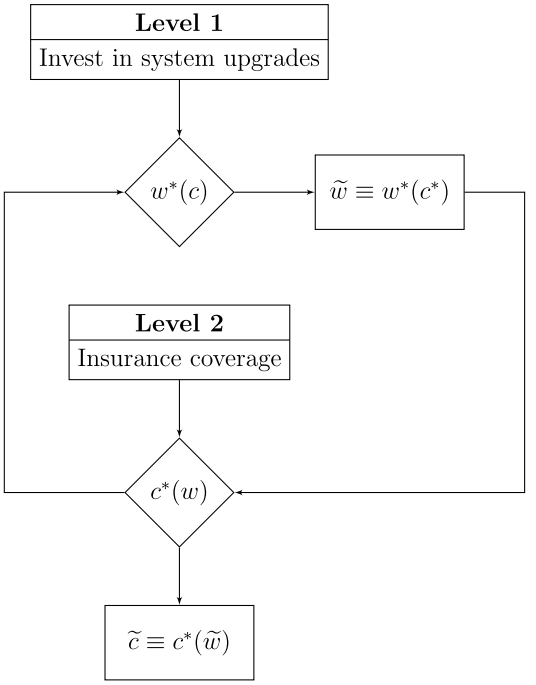

Figure 2. Diagrammatic representation of the bi-level framework capturing the strategic interaction between the defender and the insurer.

the equilibrium investment $\widetilde{w} \equiv w^*(c^*)$ in system upgrades and the equilibrium coverage level $\widetilde{c} \equiv c^*(\widetilde{w})$.

As set out in 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 future losses resulting from cyber attacks. In the event that the defender incurs loss L_i, the insurer reimburses cL_i, where $c \in [0,1]$; $c=0$ corresponds to no coverage, while $c=1$ to full coverage. The defender's PV of 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., $\tau_j^w \sim \text{Exponential}(f(w)\lambda)$, hence $T_i^w \sim \text{Gamma}(i, f(w)\lambda)$ (see equation 2). The distributional properties of the resulting discounted loss $V_i(w)$ are derived in the Appendix. Therefore,

$$\mathbb{E}[V_i(w)] = (1-c)L_i \prod_{j=1}^{i} \mathbb{E}\left[e^{-r\tau_j^w}\right] = (1-c)L_i \left(\frac{f(w)\lambda}{f(w)\lambda + r}\right)^i. \tag{4}$$

The PV over all losses is

$$V(w) = (1-c)\sum_{i=1}^{n} L_i e^{-rT_i^w}, \tag{5}$$

with expectation

$$\mathbb{E}\left[V(w)\right] = (1-c)\sum_{i=1}^{n} L_i \left(\frac{f(w)\lambda}{f(w)\lambda + r}\right)^i. \tag{6}$$

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

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

$$w^*(c) = \underset{w\in[0,1]}{\operatorname{argmin}}\, \mathbb{E}\left[V(w)\right]. \tag{7}$$

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 a cost cL_i when the firm experiences a loss L_i due to a cyber attack. The PV of the insurer's profit is given by

$$S(w) = (1-w)K - c\sum_{i=1}^{n} L_i e^{-rT_i^w}. \tag{8}$$

From (8), 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_{ins}, i.e.,

$$\mathbb{P}(S>0) = \alpha_{ins}.$$

The confidence level $0 \le \alpha_{ins} \le 1$ reflects the insurer's degree of risk aversion, with a larger (smaller) α_{ins} indicating a more (less) conservative insurer. This condition implies that the premium exceeds the cost of insurance coverage with probability α_{ins}. Therefore, the insurer's required level of coverage satisfies

$$(1-w)K = \inf\{z\in\mathbb{R} : \mathbb{P}(Z(w)\le z) \ge \alpha_{ins}\} = \mathrm{VaR}_{\alpha_{ins}}(Z(w)),$$

where $Z(w) = c\sum_{i=1}^{n} L_i e^{-rT_i^w}$ and the Value-at-Risk (VaR) measures the risk by examining the left tail of the PV distribution and is positively homogeneous. Without loss of generality, we adopt VaR as our risk measure, though this choice is not restrictive, and alternative risk measures or utility functions may be employed. By positive homogeneity of VaR, i.e., $\mathrm{VaR}_\alpha(cX) = c\,\mathrm{VaR}_\alpha(X)$ for $c>0$, and rearranging (3.2), we obtain the insurer's required level of coverage as a function of w:

$$c^*(w) = \frac{(1-w)K}{\mathrm{VaR}_{\alpha_{ins}}\left(\sum_{i=1}^{n} L_i e^{-rT_i^w}\right)}. \tag{9}$$

Next, given the required coverage level $c^*(w)$, we determine the PV of the firm's losses, now taking into account the firm's perception of the insurer's level of risk aversion, reflected in α_{def}, which may differ from the insurer's actual level of risk aversion α_{ins} due to information asymmetry. By substituting (9) into (5), the PV of the firm's loss becomes

$$V^*(w) = \left(1 - \frac{(1-w)K}{\mathrm{VaR}_{\alpha_{def}}\left(\sum_{i=1}^{n} L_i e^{-rT_i^w}\right)}\right)\sum_{i=1}^{n} L_i e^{-rT_i^w}. \tag{10}$$

The equilibrium budget allocation strategy then follows as

$$\widetilde{w} = \underset{w\in[0,1]}{\operatorname{argmin}}\, \mathbb{E}\left[V^*(w)\right]. \tag{11}$$

Finally, the equilibrium insurance coverage is obtained as

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

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

Table 1. Impact of the budget allocation ratio w on the optimal insurance coverage level, the expected PV of losses retained by the defender, and the vaR of losses ceded to the insurer, where $Z(w) = c^*(w) \sum_{i=1}^{n} L_i e^{-rT_i^w}$

	$c^*(w)$			$V^*(w)$			$\text{VaR}_{0.95}(Z(w))$		
	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$
$w = 0$	0.6426	0.3609	0.1970	1.7869	6.3909	16.0604	2.7807	4.9512	9.0717
$w = 0.25$	0.5303	0.2994	0.1641	2.0874	6.2280	14.8606	3.3210	8.0059	14.6050
$w = 0.5$	0.3850	0.2183	0.1200	2.4599	6.2534	14.0787	3.9932	5.3781	16.7181
$w = 0.75$	0.2075	0.1180	0.0652	2.8820	6.4142	13.5966	4.7752	7.4194	8.9997
$w = 1$	0	0	0	3.3333	6.6667	13.3333	5.6254	6.0599	12.4463

4. Numerical analysis of equilibrium strategies

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 PV 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 presents the impact of the exogenous budget allocation ratio w on the optimal investment in insurance coverage, the expected PV of losses retained by the defender, and the VaR of losses transferred to the insurer. For illustrative purposes, results are based on the parameter values $n = 150$, $r = 0.1$, $K = 5$, $L_i = 1$, $a = 0.5$, $b = 1$, and $\alpha_{ins} = \alpha_{def} = \alpha = 0.95$, using 10 million simulation runs.[2] 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).

Table 2 reports the equilibrium budget allocation ratio $\widetilde{w}$, insurance coverage level $\widetilde{c}$, and expected PV of losses $\widetilde{V}(\cdot)$, under information asymmetry. Specifically, this asymmetry is reflected in differing percentile levels used in the calculation of the VaR. In the upper panel, we hold constant the defender's perception of the insurer's risk aversion, i.e., $\alpha_{def} = 0.95$, while increasing the insurer's actual level of risk aversion. The results show that the defender's expected losses increase as α_{ins} rises. Intuitively, under information asymmetry, the defender forms decisions based on a fixed belief about the insurer's risk aversion and therefore does not observe the insurer's true level. As a result, the defender maintains a constant budget allocation ratio $\widetilde{w}$. In contrast, the insurer's behavior reflects its true level of risk aversion: the more risk-averse the insurer is relative to the defender's perception, the lower the equilibrium coverage level. For example, when $\lambda = 1$, the defender consistently allocates $\widetilde{w} = 0.35$ of the budget to system upgrades regardless of α_{ins}, yet the coverage level falls from $\widetilde{c} = 0.2897$ to 0.2849 to 0.2532 as α_{ins} rises from 0.90 to 0.95 to 0.975, with expected losses increasing correspondingly from 6.0449 to 6.2226 to 6.3731.

The lower panel of Table 2 isolates the impact of information asymmetry when the insurer's risk aversion level is held fixed, but the defender does not observe its exact value. We find that when the cyberattack frequency is either very low or very high (e.g., $\lambda \leq 0.5$ or $\lambda \geq 2$), the expected losses are identical across all values of α_{def}. This occurs because $\widetilde{w}$ remains constant, indicating that it

[2] By analyzing 27 million devices across 1,803 enterprise networks, Davila and Zou (2025) find that the average enterprise has approximately 35,000 network-connected devices spanning approximately 80 device types, such as endpoints, routers, etc. Trying different values for n yielded (unreported) results that converged quickly to those reported in Table 1 for n as small as 50.

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

Table 2. Equilibrium investment, coverage, and expected PV under asymmetric information

	$(\alpha_{def}, \alpha_{ins}) = (0.95, 0.90)$			$(\alpha_{def}, \alpha_{ins}) = (0.95, 0.95)$			$(\alpha_{def}, \alpha_{ins}) = (0.95, 0.975)$		
	$\widetilde{w}$	$\widetilde{c}$	$\widetilde{V}$	$\widetilde{w}$	$\widetilde{c}$	$\widetilde{V}$	$\widetilde{w}$	$\widetilde{c}$	$\widetilde{V}$
$\lambda = 0.5$	0	0.7037	1.4753	0	0.6431	1.7887	0	0.5970	2.0178
$\lambda = 1$	0.35	0.2897	6.0449	0.35	0.2849	6.2226	0.35	0.2532	6.3731
$\lambda = 1.5$	0.85	0.0555	9.9437	0.85	0.0518	9.9811	0.85	0.0489	10.0111
$\lambda = 2$	1	0	13.3333	1	0	13.3333	1	0	13.3333
	$(\alpha_{def}, \alpha_{ins}) = (0.90, 0.95)$			$(\alpha_{def}, \alpha_{ins}) = (0.95, 0.95)$			$(\alpha_{def}, \alpha_{ins}) = (0.975, 0.95)$		
	$\widetilde{w}$	$\widetilde{c}$	$\widetilde{V}$	$\widetilde{w}$	$\widetilde{c}$	$\widetilde{V}$	$\widetilde{w}$	$\widetilde{c}$	$\widetilde{V}$
$\lambda = 0.5$	0	0.6431	1.7887	0	0.6431	1.7887	0	0.6431	1.7887
$\lambda = 1$	0.25	0.2992	6.2289	0.35	0.2849	6.2226	0.35	0.2695	6.2343
$\lambda = 1.5$	0.75	0.0837	9.9945	0.85	0.0518	9.9811	0.95	0.0178	9.9888
$\lambda = 2$	1	0	13.3333	1	0	13.3333	1	0	13.3333

is optimal for the defender to allocate the entire budget either to system upgrades or to purchasing insurance. When λ takes intermediate values, however, the defender's expected PV of losses is minimized under symmetric information, i.e., when $\alpha_{def} = \alpha_{ins}$, and increases as α_{def} deviates from α_{ins}. This arises because information asymmetry leads the defender to choose a suboptimal budget allocation, resulting in higher realized losses relative to the symmetric-information benchmark. For example, when $\lambda = 1$ and $\alpha_{ins} = 0.95$, the defender optimally sets $\widetilde{w} = 0.35$ under symmetric information, yielding $\widetilde{V} = 6.2226$. If instead the defender underestimates the insurer's risk aversion ($\alpha_{def} = 0.90$), the budget allocation falls to $\widetilde{w} = 0.25$, raising expected losses to $\widetilde{V} = 6.2289$. Conversely, overestimating the insurer's risk aversion ($\alpha_{def} = 0.975$) leaves $\widetilde{w}$ unchanged at 0.35 but results in miscalculated coverage, pushing $\widetilde{V}$ to 6.2343.

Turning to the directional patterns in Table 2, in the upper panel, $\widetilde{w}$ is invariant to α_{ins} across all λ, since the defender does not observe the insurer's true risk aversion and therefore cannot condition on it. The equilibrium coverage $\widetilde{c}$, however, is decreasing in α_{ins}: a more risk-averse insurer demands a higher VaR threshold before committing to coverage, so for any given premium, the coverage offered is lower. Consequently, $\widetilde{V}$ is increasing in α_{ins} for intermediate λ, and remains constant for $\lambda = 2$, where $\widetilde{w} = 1$ and no insurance is purchased regardless. In the lower panel, $\widetilde{w}$ is non-decreasing in α_{def} for intermediate λ: a defender who overestimates the insurer's risk aversion anticipates less coverage being offered and therefore shifts more of the budget toward system upgrades. The coverage $\widetilde{c}$ is correspondingly decreasing in α_{def}, since a higher $\widetilde{w}$ reduces the premium $(1 - \widetilde{w})K$ available to the insurer, mechanically lowering the coverage level. For $\lambda = 0.5$ and $\lambda = 2$, the corner solutions ($\widetilde{w} = 0$ and $\widetilde{w} = 1$, respectively,) are robust to misspecification of α_{def}, so both $\widetilde{c}$ and $\widetilde{V}$ remain constant across columns. The behavior of $\widetilde{V}$ for intermediate λ is non-monotone in α_{def}, with the symmetric-information case ($\alpha_{def} = \alpha_{ins} = 0.95$) yielding the lowest expected losses: any deviation of α_{def} from α_{ins} in either direction leads to a suboptimal allocation and therefore higher realized losses. Intuitively, underestimating the insurer's risk aversion leads to insufficient investment in system upgrades, while overestimating it diverts excessive funds away from insurance; in both cases, the allocation is suboptimal relative to the symmetric-information benchmark.

With these observations in mind, 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 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 decrease 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

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

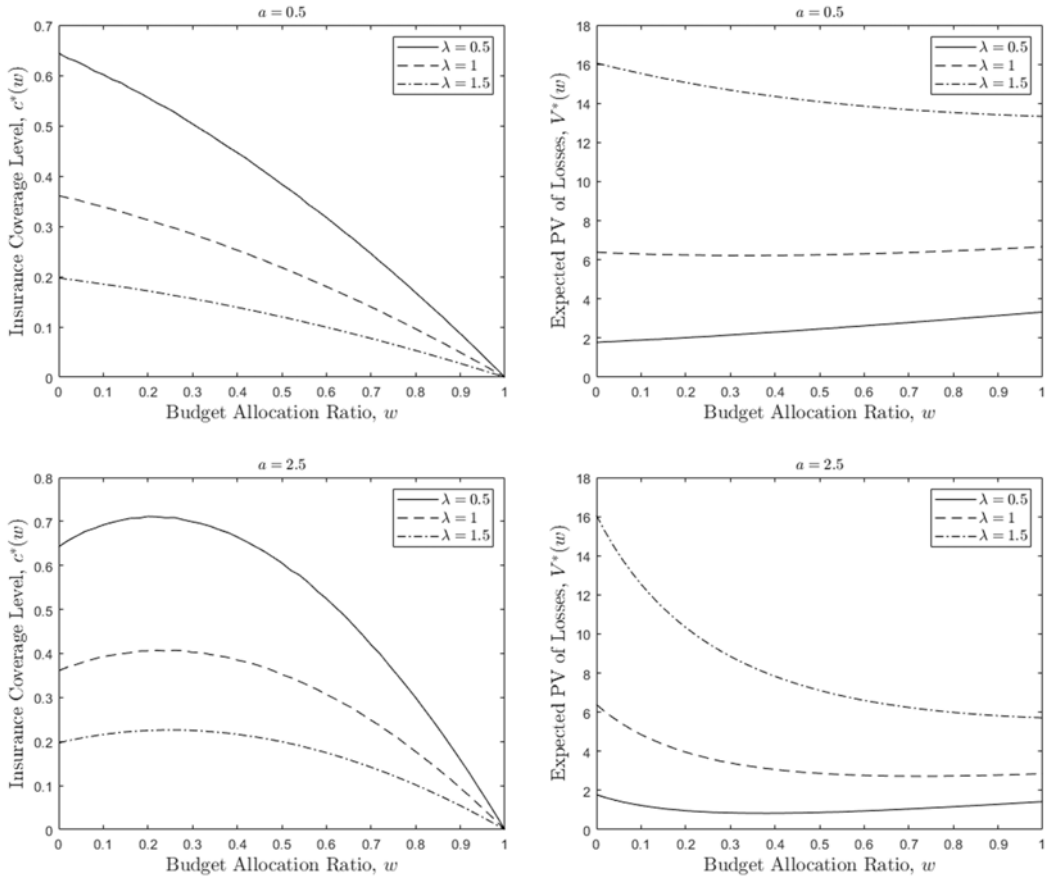

Figure 3. Impact of the exogenous budget allocation ratio w on the insurance coverage level (left) and the expected PV of losses (right) for $a = 0.5$ (top) and $a = 2.5$ (bottom).

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 cost-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, opting for insurance becomes more appealing to the defender, whereas when λ is high, investing in a system upgrade is preferred. This preference for system upgrades when λ is high 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 substantial 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 losses remains minimal even if the company does not invest in self-protection. Under these circumstances, the insurer is willing to provide 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 (3), larger

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

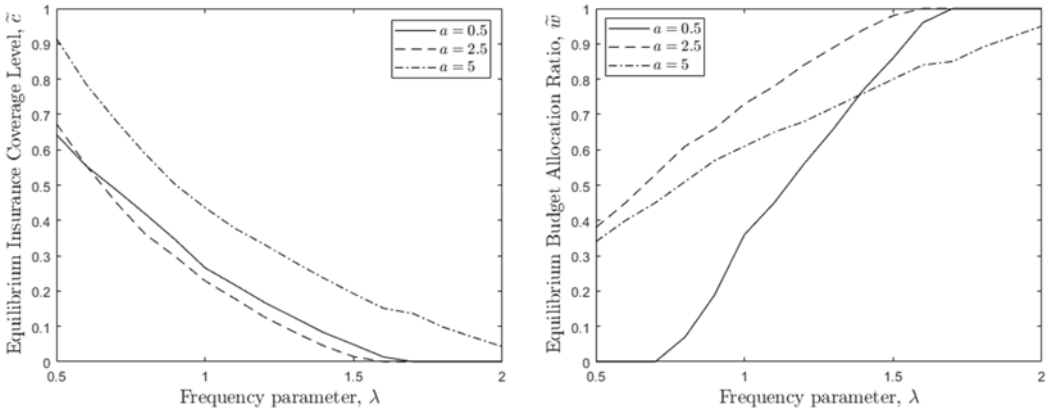

Figure 4. Impact of the frequency parameter λ on the insurance coverage level (left) and the equilibrium budget allocation ratio (right).

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 counterintuitive result can be attributed to the fact that the insurer benefits from either a higher premium or a lower total claim amount, as indicated in (8). 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 larger investment in system upgrades when a is large.

Figure 4 examines 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 or decrease the coverage ratio. However, when λ is high, the premium (see the first term in equation 8) becomes relatively small compared to the claim amount (see the second term in equation 8), 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 from 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.

As shown in 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 insurance when $\lambda < 0.45$, or entirely to system upgrades when $\lambda > 1.7$. 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.

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

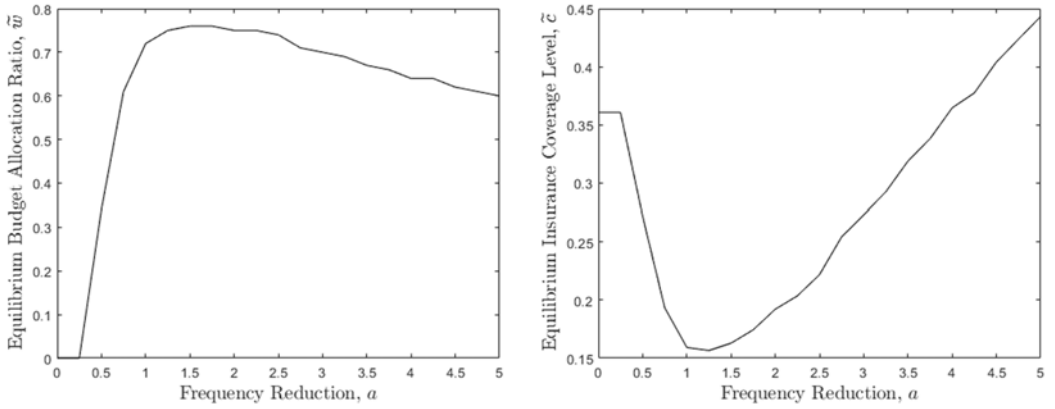

Figure 5. Impact of the attack frequency reduction parameter a on the equilibrium budget allocation ratio (left) and the insurance coverage level (right).

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 initially rises and then decreases with increasing a, while the coverage level does the opposite. 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 $\widetilde{w}$. 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 (3). 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 scope 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 obtained through larger 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 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 cybersecurity 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

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

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 optimization objectives that would enable further analysis of how risk preferences influence the optimal budget allocation problem. A utility-based approach could also be adopted to quantify these preferences and describe the objective functions of different market participants. Lastly, 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. Additional enhancements worthy of consideration in contract design include more advanced underwriting practices, dynamic pricing, and exclusions, as well as explicit treatment of adverse selection and negotiation (Wang, 2019; Awiszus *et al.*, 2023; Arce *et al.,* 2024).

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.

Competing interests. The author(s) declare none.

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**(4), 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**(2), 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**(1), 131–158.

Böhme, **R.** (2010). Security metrics and security investment models. In: *Advances in Information and Computer Security*, pp. 10–24.

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**(3), 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**(2), 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). The advanced persistent threat. *Usenix*, **4**(4), 2013–2016.

Davila, **A.** & **Zou**, **X.** (2025). 2025 Report Exposes Widespread Device Security Risks. URL: https://www.paloaltonetworks.com/blog/network-security/2025-report-exposes-widespread-device-security-risks/?utm_source=chatgpt.com.

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.

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

Gordon, **L. A.** & **Loeb**, **M. P.** (2002). The economics of information security investment. *ACM Transactions on Information and System Security*, **5**(4), 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*. New York, USA: Association for Computing Machinery, pp. 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**(2), 270–309.

Hilton, **S.** (2016). Dyn analysis summary of Friday October 21 attack. Dyn blog 26. URL: http://hub.dyn.com/dyn-blog/dyn-analysis-summary-offriday-october-21-attack.

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**(3), 719–749.

Kelion, **L.** (2013). Cryptolocker ransomware has infected about 250,000 PCs. BBC News techology. URL: www.bbc.com/news/technology-25506020 (accessed 1 May 2018).

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**(3), 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*. Cham, Springer International Publishing, 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**(2), 401–424.

Moore, **T.** (2010). The economics of cybersecurity: Principles and policy options. *International Journal of Critical Infrastructure Protection*, **3**(3), 103–117.

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)*. Washington, D.C., USENIX Association, pp. 559–576,

Nisioti, **A.**, **Loukas**, **G.**, **Rass**, **S.** & **Panaousis**, **E.** (2021). Game-theoretic decision support for cyber forensic investigations. *Sensors*, **21**(16), 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.** (2012). Cyber-Insurance in Internet Security: A Dig into the Information Asymmetry Problem. *The Computing Research Repository*, 1–6. ArXiv:1202.0884.

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*. Cham, Springer International Publishing, 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.

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*. Boston, Springer, 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. ACySE '14*. New York, Association for Computing Machinery, 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.

Whitman, **M. E.** & **Mattord**, **H. J.** (2011). *Principles of information security*. Cengage Learning.

Zhang, **R.** & **Zhu**, **Q.** (2022). Optimal cyber-insurance contract design for dynamic risk management and mitigation. *IEEE Transactions on Computational Social Systems*, **9**(4), 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**(12), 1303–1317.

https://doi.org/10.1017/S1748499526100311 Published online by Cambridge University Press

A. 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 $V_1 = (1-c)L_1 e^{-rT_1}$ that

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

Assuming $\tau_j \sim$ Exponential(λ) for all j, where λ serves as a generic arrival parameter (to connect to the main model, substitute $\lambda \leftarrow f(w)\lambda$; the unmitigated baseline $f(w) = 1$ recovers the expressions as written), we get that

$$\Theta_{V_1}(v) = \left(\frac{v}{(1-c)L_1}\right)^{\frac{\lambda}{r}} \tag{A.1}$$

with associated density function

$$\theta_{V_1}(v) = \frac{\lambda}{r}\left((1-c)L_1\right)^{-\frac{\lambda}{r}} v^{\frac{\lambda}{r}-1} \tag{A.2}$$

and mean

$$\mathbb{E}[V_1] = \int_0^{L_1} v\theta_{V_1}(v)\,dv = \frac{\lambda}{\lambda + r}(1-c)L_1. \tag{A.3}$$

For the general n-phase attack, $V_n = (1-c)L_n e^{-rT_n}$ with

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

Since $T_n \sim$ Gamma(n, λ), we get that

$$\Theta_{V_n}(v) = 1 - \frac{1}{\Gamma(n)}\gamma\left(n, \frac{\lambda}{r}\ln\frac{(1-c)L_n}{v}\right), \tag{A.4}$$

$$\theta_{V_n}(v) = \frac{\lambda^n}{rv\Gamma(n)}\left(\frac{1}{r}\ln\frac{(1-c)L_n}{v}\right)^{n-1}\left(\frac{v}{(1-c)L_n}\right)^{\frac{\lambda}{r}}, \tag{A.5}$$

where $\gamma(\cdot)$ and $\Gamma(\cdot)$ denote the lower incomplete gamma and gamma functions, from which

$$\mathbb{E}[V_n] = (1-c)L_n\prod_{j=1}^{n}\mathbb{E}\left[e^{-r\tau_j^w}\right] = (1-c)L_n\left(\frac{\lambda}{\lambda + r}\right)^n. \tag{A.6}$$

Cite this article: Zhang Z, Chronopoulos M and Kyriakou I (2026). Bi-level optimization of security investment and insurance pricing, *Annals of Actuarial Science*, 1–15. https://doi.org/10.1017/S1748499526100311