Annals of Actuarial Science (2026), 1–29
doi:10.1017/S1748499526100359

ORIGINAL RESEARCH PAPER

A stochastic Gordon–Loeb model for optimal cybersecurity investment under clustered attacks

Giorgia Callegaro[1], Claudio Fontana[1], Caroline Hillairet[2] and Beatrice Ongarato[3]

[1]Dept. of Mathematics, "Tullio Levi-Civita", Università degli Studi di Padova, Italy; [2]ENSAE Paris, France; and [3]Institute of Mathematical Stochastics, Technische Universität Dresden, Germany
Corresponding author: Claudio Fontana; Email: fontana@math.unipd.it

(Received 05 September 2025 UTC; revised 10 March 2026 UTC; accepted 01 June 2026 UTC)

Abstract

We develop a continuous-time stochastic model for optimal cybersecurity investment under the threat of cyberattacks. The arrival of attacks is modeled using a Hawkes process, capturing the empirically relevant feature of clustering in cyberattacks. Extending the Gordon–Loeb model, each attack may result in a breach, with breach probability depending on the system's vulnerability. We aim at determining the optimal cybersecurity investment to reduce vulnerability. The problem is cast as a two-dimensional Markovian stochastic optimal control problem and solved using dynamic programming methods. Numerical results illustrate how accounting for attack clustering leads to more responsive and effective investment policies, offering significant improvements over static and Poisson-based benchmark strategies. Our findings underscore the value of incorporating realistic threat dynamics into cybersecurity risk management.

Keywords: benefit-cost analysis; cyber-risk; Hawkes process; security breach; stochastic optimal control

1. Introduction

Cyber-risk is nowadays widely acknowledged as one of the major sources of operational risk for organizations worldwide. The 2024 ENISA Threat Landscape Report (ENISA, 2024) documents "a notable escalation in cybersecurity attacks, setting new benchmarks in both the variety and number of incidents, as well as their consequences." According to the AON 9th Global Risk Management Survey,[1] cyberattacks and data breaches represent the foremost source of global risk faced by organizations, with the second biggest risk being business interruption, which is itself often a consequence of cyber-incidents. A recent poll on Risk.net confirms information security and IT disruption as the top two sources of operational risk for 2025.[2] According to IBM, the global average cost of a data breach has reached nearly 5 M USD in 2024, an increase of more than 10% over the previous year.[3]

The rapid and widespread emergence of cyber-risk as a key source of operational risk has led to a significant increase in cybersecurity spending. In the 2025 ICS/OT cybersecurity budget survey of the SANS Institute (2025), 55% of the respondents reported a substantial rise in cybersecurity budgets over the previous two years. This trend underscores the importance of adopting effective cybersecurity investment policies that balance risk mitigation with cost efficiency.

[1]Source: https://www.aon.com/en/insights/reports/global-risk-management-survey.
[2]Source: https://www.risk.net/risk-management/7961268/top-10-operational-risks-for-2025.
[3]Source: https://www.ibm.com/reports/data-breach.

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Institute and Faculty of Actuaries. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

The problem of optimal cybersecurity investment has been first addressed in the seminal work of Gordon and Loeb (2002). In their model, reviewed in Section 2.1 below, the decision maker can reduce the vulnerability to cyberattacks by investing in cybersecurity. The optimal expenditure in cybersecurity is determined by maximizing the expected net benefit of reducing the breach probability. The Gordon–Loeb model laid the foundations for a rigorous quantitative analysis of cybersecurity investments and has been the subject of numerous extensions and generalizations: we mention here only some studies that are closely related to our context, referring to Fedele and Roner (2022) for a comprehensive overview. The key ingredient of the Gordon–Loeb model is represented by the security breach probability function (see Section 2.1), which has been further analyzed in Huang and Behara (2013) and Mazzoccoli and Naldi (2022). The risk-neutral assumption of the original model (Gordon & Loeb, 2002) has been relaxed to accommodate risk-averse preferences in Miaoui and Boudriga (2019). The practical applicability of the Gordon–Loeb model for guiding cost-effective cybersecurity investment decisions has been examined in Gordon *et al.* (2020) within the National Institute of Standards and Technology (NIST) cybersecurity framework.

The original Gordon–Loeb model is a static model and, therefore, does not allow addressing the crucial issue of the optimal timing of investment decisions. Adopting a real-options approach, Gordon *et al.* (2003) and Tatsumi and Goto (2010) have proposed dynamic versions of the model that allow analyzing the optimal timing and level of cybersecurity investment. Closer to our setup, a dynamic extension of the Gordon–Loeb model has been developed in Krutilla *et al.* (2021), considering the problem of optimal cybersecurity investment over an infinite time horizon and assuming that cybersecurity assets are subject to depreciation over time, while future net benefits of cybersecurity investment are discounted.

An effective cybersecurity investment policy must be adaptive and evolve in response to changing threat environments. As noted by Zeller and Scherer (2022), a key feature of cyber-risk is its dynamic nature due to the rapid technological transformation and the evolution of threat actors. Similarly, Balzano and Marzi (2025) emphasize the need for adaptable and responsive cybersecurity policies in order to face the challenge of dynamic cyberattacks. The framework of Krutilla *et al.* (2021) is based on a deterministic model and, therefore, cannot capture the dynamic behavior of cyber-risk. Addressing this need, the main contribution of this work consists in proposing a modeling framework for optimal cybersecurity investment in a dynamic stochastic setup, allowing for investment policies that respond in real time to randomly occurring cyberattacks. Our work contributes both to cyber-risk modeling and to cyber-risk management, as categorized in the recent survey by He *et al.* (2024). Moreover, our stochastic modeling framework takes into account the empirically relevant feature of temporally clustered cyberattacks, as discussed in the next subsection.

1.1 Modeling cyberattacks with Hawkes processes

A distinctive feature of our modeling framework, which will be described in Section 2.2, is the use of a Hawkes process to model the arrival of cyberattacks. First introduced by Alan G. Hawkes in Hawkes (1971), these stochastic processes are used to model event arrivals over time and are particularly suited to situations where the occurrence of one event increases the likelihood of subsequent events (self-excitation), thereby generating temporally clustered events. In our context, we denote by N_t the cumulative number of cyberattacks up to time t, modeled as a Hawkes process, and by λ_t the associated stochastic intensity (hazard rate), representing the instantaneous likelihood of an attack occurring.

The self-exciting property of the Hawkes process $(N_t)_{t\geq 0}$ is captured by the specification of its stochastic intensity:

$$\lambda_t = \alpha + (\lambda_0 - \alpha)e^{-\xi t} + \beta \sum_{i=1}^{N_t} e^{-\xi(t-\tau_i)}, \tag{1}$$

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

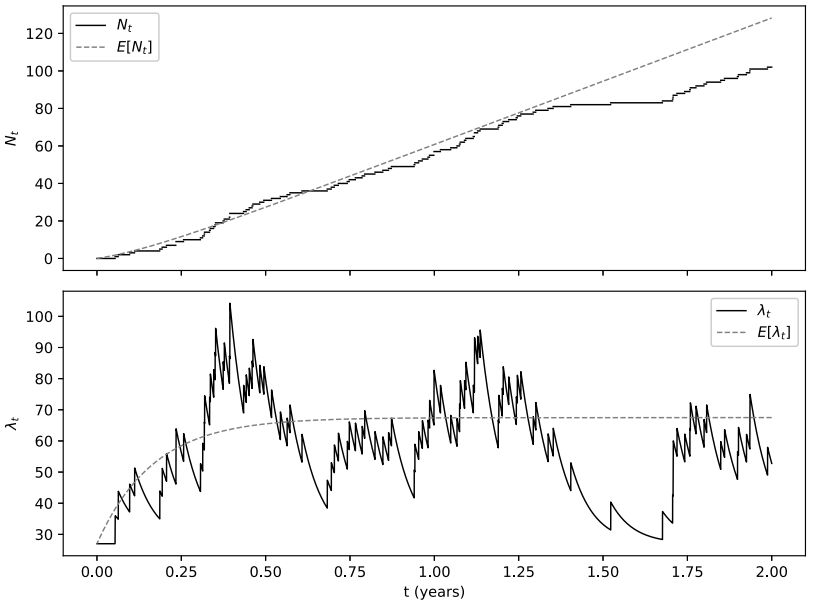


Figure 1. One simulated trajectory of N (top) and λ (bottom), for $\alpha = 27$, $\lambda_0 = 27$, $\xi = 15$, and $\beta = 9$.

for all $t \geq 0$, where

- $\alpha > 0$ is the long-term mean intensity;
- $\lambda_0 > 0$ is the intensity at the initial time $t = 0$;
- $\xi > 0$ is the exponential decay rate;
- $\beta > 0$ determines the magnitude of self-excitation;
- $(\tau_i)_{i \in \mathbb{N}^*}$ are the random times at which attacks occur.

Figure 1 shows a simulated trajectory of N and λ, showing the clustering behavior induced by the self-exciting mechanism described above. General presentations of the theory and the applications of Hawkes processes can be found in Laub *et al.* (2021) and Lima (2023).

This modeling choice is particularly relevant in the context of cyber-risk. Cyberattacks frequently occur in bursts, for instance, following the discovery of a vulnerability or due to the propagation of malware across interconnected systems (see Nguyen *et al.*, 2024). Such clustered patterns are not adequately captured by memoryless models, such as those based on Poisson processes.

Empirical evidence supports the appropriateness of the Hawkes framework for modeling cyber-risk. A contagious behavior in cyberattacks has been documented in Baldwin *et al.* (2017), analyzing the threats to key internet services using data from the SANS Institute. Using data from the Privacy Rights Clearinghouse, it has been empirically demonstrated in Bessy-Roland *et al.* (2021) that Hawkes-based models provide a more realistic representation of the interdependence of data breaches compared to Poisson-based models. The recent work (Boumezoued *et al.*, 2023) reinforces this perspective by calibrating a two-phase Hawkes model to cyberattack data, taking into account the publication of cyber-vulnerabilities. These studies provide strong support for modeling cyberattacks via Hawkes processes, as described in more detail in Section 2.2.

1.2 Goals and main contributions

In this work, we address the challenge of optimal cybersecurity investment under temporally clustered cyberattacks, in line with the empirical evidence reported above. In particular, we aim at

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

studying the adaptive investment policy that best responds in real time to the random arrival of cyberattacks, within a framework that balances realism with analytical tractability. To this end, we develop a continuous-time stochastic extension of the classical Gordon–Loeb model, describing attack arrivals with a Hawkes process. The model incorporates key operational features such as technological obsolescence and the decreasing marginal effectiveness of large investments. The resulting optimization problem is cast as a stochastic optimal control problem and solved via dynamic programming methods. We develop efficient numerical schemes to compute the optimal policy, and we quantify the benefits of dynamic investment strategies under clustered attacks. By integrating risk dynamics into the cybersecurity investment problem, our framework provides new insights into how organizations can better allocate resources to mitigate cyber-risk.

1.3 Structure of the paper

In Section 2, we recall the original Gordon–Loeb model and introduce our continuous-time stochastic extension. In Section 3, we formulate the cybersecurity investment problem and characterize the optimal policy. Section 4 details the model parameters and the numerical methods used in our analysis. Section 5 presents the results of our numerical analysis and discusses their practical implications for cyber-risk management. Section 6 discusses the relevance of our study from an insurance perspective, also in relation to the calculation of premia for cyber-insurance contracts. Section 7 concludes.

2. The model

We study the decision problem faced by an entity (a public administration or a large corporation) that is threatened by a massive number of randomly occurring cyberattacks with a temporally clustered pattern. As in the Gordon–Loeb model (reviewed in Section 2.1), not all cyberattacks result in successful breaches of the entity's system. The success rate of each attack depends on the system's vulnerability, which the entity can mitigate by investing in cybersecurity.

2.1 The Gordon–Loeb model

The Gordon–Loeb model, introduced in 2002 in the seminal work (Gordon & Loeb, 2002), provides a static framework for determining the optimal investment in cybersecurity to protect a given information set under the threat of cyberattacks. In our context, the information set corresponds to the entity's IT infrastructure. In the Gordon–Loeb model, the information set is characterized by three key parameters, all assumed to be constant:

- $p \in [0, 1]$: the probability that a cyberattack occurs;
- $v \in [0, 1]$: the baseline probability that a cyberattack successfully breaches the information set (vulnerability);
- $\ell \geq 0$: the loss incurred when a breach occurs.

The vulnerability v refers to a baseline level of security, reflecting the cybersecurity measures already in place before any new decision is taken. Without any further investment in cybersecurity, the expected loss is $vp\ell$. To mitigate its vulnerability, the entity may invest an additional amount $z \geq 0$ in cybersecurity. The variable z therefore represents incremental cybersecurity investment beyond the existing baseline protection. In particular, $z = 0$ means that no further investment is made, but the system still retains the prior protection embodied in v. Throughout the paper, absence of cybersecurity investment will refer to the case $z = 0$. The effectiveness of the investment z is measured by a security breach probability function $S(z, v)$, which represents the probability that an attack successfully breaches the information set, given investment level z

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

and the vulnerability v. The resulting expected loss is thus $S(z, v)p\ell$. Gordon and Loeb require the function S to satisfy the properties listed in the following assumption.

Assumption A.

(A1) $S(z, 0) = 0$, *for all* $z \geq 0$, *i.e., an invulnerable information set always remains invulnerable;*
(A2) $S(0, v) = v$, *i.e., in the absence of investment, the information set retains its baseline vulnerability;*
(A3) S *is decreasing and convex in* z, *so that* $S_z(z, v) < 0$ *and* $S_{zz}(z, v) > 0$, *for all* $z \geq 0$, *i.e., cybersecurity investment reduces breach probability with diminishing marginal effectiveness.*

Gordon and Loeb consider two classes of security breach probability functions, which satisfy Assumption A:

$$S_I(z, v) = \frac{v}{(az + 1)^b} \quad \text{and} \quad S_{II}(z, v) = v^{az+1}, \tag{2}$$

for some parameters $a, b > 0$.

In Gordon and Loeb (2002), the optimal investment in cybersecurity is determined by maximizing the Expected Net Benefit of Investment in Information Security (ENBIS), defined as follows:

$$\text{ENBIS}(z) := (v - S(z, v))p\ell - z. \tag{3}$$

The ENBIS function quantifies the net trade-off between the benefit (captured by the reduction in the expected loss due to the investment in cybersecurity) and the direct cost of investing. Under Assumption A, the optimal investment level z^* is determined by the following first-order condition:

$$-S_z(z^*, v)p\ell - 1 = 0.$$

Remark 1. For both classes of security breach functions in (2), Gordon and Loeb show that the optimal cybersecurity investment never exceeds $1/e \approx 37\%$ of the expected loss:

$$z^* < \frac{vp\ell}{e}.$$

2.2 A continuous-time model driven by a Hawkes process

We now introduce a continuous-time model for randomly occurring cyberattacks. As discussed in Section 1.1, we want to capture the empirically relevant feature of clustering of cyberattacks while retaining the key elements of the original Gordon–Loeb model reviewed in Section 2.1.

The arrival of cyberattacks is described by a Hawkes process $N = (N_t)_{t\geq 0}$, defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with N_t representing the number of cyberattacks up to time t, for all $t \geq 0$. As explained in Section 1.1, the process N is characterized by a self-exciting stochastic intensity $(\lambda_t)_{t\geq 0}$ solving the following stochastic differential equation:

$$\mathrm{d}\lambda_t = \xi(\alpha - \lambda_t)\,\mathrm{d}t + \beta \mathrm{d}N_t, \quad \lambda_0 > 0, \tag{4}$$

whose explicit solution is given by (1). In the next proposition (adapted from Dassios & Zhao, 2013), we compute the expectation of some basic quantities, which will be used later.[4]

[4]The condition $\beta < \xi$ in Proposition 1 ensures that the L^1-norm of the self-excitation kernel of the Hawkes process is strictly less than one. This guarantees that the process is non-explosive, meaning that it generates almost surely a finite number of events over any finite time interval. The same condition also corresponds to the stationarity condition, widely adopted in the theory of Hawkes processes since the seminal work of Hawkes and Oakes (1974).

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Proposition 1. *Let $(N_t)_{t\geq 0}$ be a Hawkes process with intensity $(\lambda_t)_{t\geq 0}$ given in (4), with $\beta < \xi$. Then, for all $t \geq 0$,*

$$\mathbb{E}[\lambda_t] = \frac{\alpha\xi}{\xi-\beta} + e^{-(\xi-\beta)t}\left(\lambda_0 - \frac{\alpha\xi}{\xi-\beta}\right),$$

$$\mathbb{E}[N_t] = \int_0^t \mathbb{E}[\lambda_s]\mathrm{d}s$$

$$= \frac{\alpha\xi}{\xi-\beta}t - \frac{1}{\xi-\beta}\left(\lambda_0 - \frac{\alpha\xi}{\xi-\beta}\right)\left(e^{-(\xi-\beta)t} - 1\right).$$

We denote by $(\tau_i)_{i\in\mathbb{N}^*}$ the random jump times of the process N, representing the arrival times of cyberattacks. For each $t \geq 0$, we denote by $\mathcal{F}_t := \sigma(N_s; s \leq t)$ the sigma-field generated by the Hawkes process N up to time t, representing the information generated by the history of the attack timings up to time t. In the following, this is assumed to be the information set available to the decision maker (see also Remark 3 below).

In the absence of cybersecurity investment, each attack is assumed to breach the entity's IT system with fixed probability ν (vulnerability). For each $i \in \mathbb{N}^*$, we introduce a Bernoulli random variable B_i^ν of parameter ν, where the event $\{B_i^\nu = 1\}$ corresponds to a successful breach caused by the i-th attack. In the event of a breach, the entity incurs a random monetary loss η_i, realized at the attack time τ_i. Otherwise, if $B_i^\nu = 0$, the attack is blocked and no loss occurs at time τ_i.

The families of random variables $(B_i^\nu)_{i\in\mathbb{N}^*}$ and $(\eta_i)_{i\in\mathbb{N}^*}$ are assumed to satisfy the following standing assumption.

Assumption B. *The family $(\eta_i)_{i\in\mathbb{N}^*}$ is composed of i.i.d. positive random variables in $L^1(\mathbb{P})$. The families $(\eta_i)_{i\in\mathbb{N}^*}$ and $(B_i^\nu)_{i\in\mathbb{N}^*}$ are mutually independent and independent of N.*

Remark 2. The assumption that the losses $(\eta_i)_{i\in\mathbb{N}^*}$ are independent of the attack arrival process N can be restrictive from an empirical viewpoint. In practice, both the frequency and the severity of cyber losses can increase during periods of clustered attacks. From a modeling standpoint, this assumption can be relaxed, for instance, by allowing the loss distribution to depend on the current intensity, making use of the techniques recently developed in Hillairet et al. (2025). However, introducing such dependency would significantly complicate the analysis of the optimization problem in Section 3 and require precise information on the conditional distribution of losses given the attack dynamics. In the absence of sufficient publicly available data on cyber losses, we retain the independence assumption as a tractable and transparent modeling choice.

The cumulative loss incurred over a planning horizon $[0, T]$, in the absence of cybersecurity investment, is given by:

$$L_T^0 := \sum_{i=1}^{N_T} \eta_i B_i^\nu. \tag{5}$$

In our dynamic model, the entity can react to the evolving threat environment by investing in cybersecurity in order to mitigate its vulnerability to cyberattacks. Investment occurs continuously throughout the planning horizon $[0, T]$ and is described by a non-negative investment rate process $(z_t)_{t\in[0,T]}$. For each $t < T$, the quantity z_t represents the increase in the level of cybersecurity over the infinitesimal time interval $[t, t + \mathrm{d}t]$ and is assumed to be chosen based on the information set $\mathcal{F}_{t-}$ available to the decision maker just before time t.

Remark 3. In the terminology of stochastic processes, the requirement that the investment rate z_t must be chosen on the basis of $\mathcal{F}_{t-}$ corresponds to requiring the process $(z_t)_{t\in[0,T]}$ to be predictable with respect to the filtration generated by N. In particular, at any attack time τ_i, the value z_{τ_i} is determined by the information $\mathcal{F}_{\tau_i-}$: the decision maker may adjust investment after

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

observing an attack, but not in a way that anticipates the attack at the same instant. We point out that, in our setup, the outcomes of previous attacks (i.e., whether an attack has successfully breached the system or not) do not carry any relevant informational content for decision making, as they do not affect the dynamics of future attack arrivals.

Investment in cybersecurity is subject to rapid technological obsolescence (see, e.g., Hayes & Bodhani, 2013). In line with Krutilla *et al.* (2021), we take into account this significant aspect in our model by introducing a depreciation rate $\rho > 0$. The cybersecurity level reached at time t is then defined as follows, for all $t \in [0, T]$:

$$H_t^z = H_0 e^{-\rho t} + \int_0^t e^{-\rho(t-s)} z_s \mathrm{d}s, \tag{6}$$

which equivalently, in differential form, reads as follows:

$$\mathrm{d}H_t^z = \left(z_t - \rho H_t^z\right) \mathrm{d}t, \qquad H_0^z = H_0 \geq 0.$$

As in Krutilla *et al.* (2021), we interpret the cybersecurity level as an aggregated asset, which can be thought of as a combination of technological infrastructures, software, and human expertise.

In our continuous-time framework, we let the breach probability evolve dynamically with the current cybersecurity level. More specifically, suppose that the decision maker adopts an investment policy $z = (z_t)_{t\in[0,T]}$. At each attack time τ_i, a breach is assumed to occur with probability

$$S(H_{\tau_i}^z, v), \tag{7}$$

where $H_{\tau_i}^z$ is given by (6) evaluated at $t = \tau_i$ and S is a security breach probability function satisfying Assumption A, as in the original Gordon–Loeb model. Hence, the probability that the i-th attack successfully breaches the IT system depends on the cybersecurity level $H_{\tau_i}^z$ reached at the attack's time τ_i. In turn, $H_{\tau_i}^z$ is determined by the investment realized over the time period $[0, \tau_i]$, taking into account technological obsolescence. If the i-th attack breaches the IT system, then the entity incurs a loss of η_i; otherwise, the attack is blocked and the entity does not suffer any loss at time τ_i.

Remark 4. The proposed model allows for adaptive real-time cybersecurity investment. More specifically, the arrival of an attack triggers an increased likelihood of further attacks within a short timeframe due to the form (4) of the intensity. The decision maker can respond in real-time by increasing cybersecurity investment, which in turn reduces future breach probabilities through the function S in (7). The optimal investment policy will be determined in Section 3, while the practical importance of allowing for an adaptive real-time investment strategy – rather than a static policy as in the original Gordon–Loeb model – will be empirically analyzed in Section 5.3.

Analogously to the case without investment in cybersecurity, we can write as follows the cumulative losses L_T^z incurred on the time interval $[0, T]$ when investing in cybersecurity according to a generic rate $z = (z_t)_{t\in[0,T]}$:

$$L_T^z := \sum_{i=1}^{N_T} \eta_i B_i^{S(H_{\tau_i}^z, v)}, \tag{8}$$

where $(B_i^{S(H_{\tau_i}^z, v)})_{i\in\mathbb{N}^*}$ is a family of random variables taking values in $\{0, 1\}$ and satisfying the following assumption.

Assumption C. *For any process $(z_t)_{t\in[0,T]}$, it holds that*

$$\mathbb{P}\left(B_i^{S(H_{\tau_i}^z, v)} = 1 \,\middle|\, \mathcal{F}_T\right) = S(H_{\tau_i}^z, v), \qquad \textit{for all } i \in \mathbb{N}^*,$$

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

where $(H^z_t)_{t\in[0,T]}$ *is determined by* $(z_t)_{t\in[0,T]}$ *as in (6). Moreover, for each* $i\in\mathbb{N}^*$, *the random variables* $B_i^{S(H^z_{\tau_i},\nu)}$ *and* η_i *are conditionally independent given* $\mathcal{F}_T$.

Remark 5. The cumulative loss process $(L^z_t)_{t\in[0,T]}$ defined as in (8) constitutes a marked Hawkes process in the terminology of point processes (see Brémaud, 1981). In our modeling framework, the marks (losses) are endogenous and depend on the dynamically evolving cybersecurity level $(H^z_t)_{t\in[0,T]}$.

For strategic decision making, a key quantity is represented by the expected losses due to cyber-attacks over the time interval $[0, T]$ when adopting a suitable cybersecurity policy. This is the content of the following proposition, which will be fundamental for addressing the optimal investment problem in Section 3. We denote by $\bar\eta := \mathbb{E}[\eta_i]$ the expected loss resulting from a successful breach, for all $i\in\mathbb{N}^*$.

Proposition 2. *Under Assumptions B and C, it holds that*

$$\mathbb{E}\left[L^0_T\right] = \bar\eta\,\nu\,\mathbb{E}\left[\int_0^T \lambda_t \mathrm{d}t\right],$$

$$\mathbb{E}\left[L^z_T\right] = \bar\eta\,\mathbb{E}\left[\int_0^T S(H^z_t,\nu)\lambda_t \mathrm{d}t\right].$$

Therefore, the expected net benefit of investment is

$$\mathbb{E}\left[L^0_T - L^z_T\right] = \bar\eta\,\mathbb{E}\left[\int_0^T \big(\nu - S(H^z_t,\nu)\big)\,\lambda_t \mathrm{d}t\right]. \tag{9}$$

Proof. Let $z=(z_t)_{t\in[0,T]}$ be an arbitrary cybersecurity investment rate process. Applying the tower property of conditional expectation and making use of Assumptions B and C, we can compute

$$\begin{aligned}
\mathbb{E}\left[L^z_T\right] &= \mathbb{E}\left[\sum_{i=1}^{N_T} \eta_i B_i^{S(H^z_{\tau_i},\nu)}\right] \\
&= \mathbb{E}\left[\sum_{i=1}^{N_T} \mathbb{E}\left[\eta_i B_i^{S(H^z_{\tau_i},\nu)}\,\middle|\,\mathcal{F}_T\right]\right] \\
&= \bar\eta\,\mathbb{E}\left[\sum_{i=1}^{N_T} S(H^z_{\tau_i},\nu)\right] \\
&= \bar\eta\,\mathbb{E}\left[\int_0^T S(H^z_t,\nu)\mathrm{d}N_t\right] \\
&= \bar\eta\,\mathbb{E}\left[\int_0^T S(H^z_t,\nu)\lambda_t \mathrm{d}t\right],
\end{aligned}$$

where the last step follows by definition of intensity (see, e.g., Brémaud, 1981, Definition II.D7), together with the continuity of the process H^z, which implies that $H^z_{\tau_i}$ coincides with the left-limit $H^z_{\tau_i-}$, for all $i\in\mathbb{N}^*$. The first equation in the statement of the proposition follows as a special case by taking $z\equiv 0$. □

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

3. Optimal cybersecurity investment

In this section, we determine the optimal cybersecurity investment policy in the model setup introduced in Section 2.2. In the spirit of the original Gordon–Loeb model, we aim at characterizing the investment rate process $z^* = (z_t^*)_{t\in[0,T]}$ which maximizes the net trade-off between the benefits and the costs of cybersecurity over a planning horizon $[0, T]$.

To ensure the well-posedness of the optimization problem, we constrain the admissible investment policies to a suitably defined admissible set $\mathcal{Z}$.[5] In line with Remark 3, in the following definition we restrict our attention to processes that are predictable with respect to the filtration generated by N.

Definition 1. *The admissible set $\mathcal{Z}$ is defined as the set of all non-negative predictable processes $(z_t)_{t\in[0,T]}$ such that $\mathbb{E}[\int_0^T z_t^2 \mathrm{d}t] < \infty$.*

We now formulate the central optimization problem, which generalizes the benefit-cost trade-off function in (3) to a dynamic stochastic setting. The objective is to maximize the expected net benefit of cybersecurity investments:

$$\sup_{z\in\mathcal{Z}} \mathbb{E}\left[L_T^0 - L_T^z - \int_0^T \left(\delta z_t + \frac{\gamma}{2} z_t^2\right) \mathrm{d}t + U(H_T^z)\right], \tag{10}$$

where L_T^0 and L_T^z are defined in (5) and (8), respectively, and the state variables λ and H^z satisfy the dynamics

$$\mathrm{d}\lambda_t = \xi(\alpha - \lambda_t)\mathrm{d}t + \beta \mathrm{d}N_t, \tag{11}$$

$$\mathrm{d}H_t^z = \left(z_t - \rho H_t^z\right) \mathrm{d}t. \tag{12}$$

In the objective functional (10), the term $\mathbb{E}[L_T^0 - L_T^z]$ represents the reduction in the expected losses due to the investment in cybersecurity. Differently from (3), we consider in (10) a non-linear cost function $z \mapsto \delta z + \gamma z^2/2$, for $\delta, \gamma > 0$. The non-linearity penalizes irregular or highly concentrated investment strategies, reflecting real-world constraints and incentivizing smoother, more sustained cybersecurity efforts (e.g., continuous IT updates versus abrupt large-scale interventions). The term $U(H_T^z)$ accounts for the residual utility of the cybersecurity level reached at the end of the planning horizon. This accounts for the fact that cybersecurity investment carries a long-term benefit, since the entity does not cease to exist after the planning horizon. As usual, the function $U : \mathbb{R}_+ \to \mathbb{R}$ is assumed to be an increasing and concave utility function.

Up to a rescaling of the model parameters, there is no loss of generality in taking $\delta = 1$. Hence, making use of Proposition 2, we can equivalently rewrite problem (10) as follows:

$$\sup_{z\in\mathcal{Z}} \mathbb{E}\left[\int_0^T \left(\bar{\eta}\big(\nu - S(H_t^z, \nu)\big)\lambda_t - z_t - \frac{\gamma}{2} z_t^2\right) \mathrm{d}t + U(H_T^z)\right]. \tag{13}$$

Remark 6. The objective functional (10) is linear in the losses $(\eta_i)_{i\in\mathbb{N}^*}$. More specifically, formula (13) shows that the loss distribution enters into the optimization problem only through its first moment. In turn, this implies that the optimal policy derived in Theorem 1 below depends only on $\bar{\eta}$ and not on higher-order moments of the loss distribution. Under non-linear objective functionals, higher moments and the tail behavior of the loss distribution would affect the optimal policy.

Problem (13) is a bi-dimensional stochastic optimal control problem, where the stochastic intensity process $(\lambda_t)_{t\in[0,T]}$ acts as an additional state variable beyond the controlled process

[5] We point out that, as a direct consequence of the Cauchy–Schwarz inequality, the integral in (6) is always well-defined for every process $z \in \mathcal{Z}$.

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

$(H_t^z)_{t\in[0,T]}$. Due to the Markovian structure of the system, dynamic programming techniques can be applied for the solution of (13). To this end, we first introduce the following notation, for any $(t,\lambda,h)\in[0,T]\times(0,\infty)\times\mathbb{R}_+$:

- for all $s\in[t,T]$,
$$H_s^{t,h;z} := he^{-\rho(s-t)} + \int_t^s e^{-\rho(s-v)} z_v \mathrm{d}v \tag{14}$$
represents the cybersecurity level reached at time s when starting from level $H_t = h$ at time t and investing according to $z\in\mathcal{Z}$;
- for all $s\in[t,T]$,
$$\lambda_s^{t,\lambda} := \alpha + (\lambda-\alpha)e^{-\xi(s-t)} + \beta \sum_{i=N_t+1}^{N_s} e^{-\xi(s-\tau_i)}$$
represents the stochastic intensity at time s when starting from value $\lambda_t = \lambda$ at time t.

For any stopping time τ taking values in $[0,T]$, we denote by $\mathcal{Z}_\tau$ the admissible set $\mathcal{Z}$ restricted to the stochastic time interval $[\tau,T]$.

We define as follows the benefit-cost trade-off functional J associated with a given investment rate process z:

$$J(t,\lambda,h;z) := \mathbb{E}\left[\int_t^T \bar{\eta}\big(v - S\big(H_s^{t,h;z},v\big)\big)\lambda_s^{t,\lambda}\mathrm{d}s - \int_t^T \left(z_s + \frac{\gamma}{2}z_s^2\right)\mathrm{d}s + U\big(H_T^{t,h;z}\big)\right].$$

As a consequence of Definition 1 together with the concavity of the function U, the functional J is always well-defined and finite. Consequently, the value function associated with the stochastic optimal control problem (13) is given by

$$V(t,\lambda,h) := \sup_{z\in\mathcal{Z}_t} J(t,\lambda,h;z). \tag{15}$$

In our dynamic model, the value function $V(t,\lambda,h)$ encodes the benefit-cost trade-off of cybersecurity investment over the residual planning horizon $[t,T]$, when considered at time t with current cybersecurity level h and intensity λ.

Remark 7. Standard arguments allow proving that the function V has the following behavior with respect to λ and h:

- for every $(t,h)\in[0,T)\times\mathbb{R}_+$, the map $\lambda\mapsto V(t,\lambda,h)$ is strictly increasing: the benefit of cybersecurity investment is greater in the presence of a greater risk of cyberattacks;
- for every $(t,\lambda)\in[0,T]\times(0,\infty)$, the map $h\mapsto V(t,\lambda,h)$ is strictly increasing and concave: increasing the current cybersecurity level h always improves the expected future benefit, but the marginal value of additional protection decreases as the cybersecurity level h grows.

The optimization problem formulated in (10) falls within the framework of the optimal control of a piecewise deterministic Markov process, in the sense of Davis (1984). Under the admissibility conditions imposed in Definition 1, the performance functional J is well-posed for every control $z\in\mathcal{Z}_t$. As a consequence, the same techniques used in the proof of Pham (2009, Theorem 3.3.1) yield that the value function satisfies the dynamic programming principle. More precisely, for all (t,λ,h) in $[0,T)\times(0,\infty)\times\mathbb{R}_+$ and for every stopping time τ taking values in $[t,T]$, it holds that

$$V(t,\lambda,h) = \sup_{z\in\mathcal{Z}_t} \mathbb{E}\left[\int_t^\tau \bar{\eta}\big(v - S\big(H_s^{t,h;z},v\big)\big)\lambda_s^{t,\lambda}\mathrm{d}s - \int_t^\tau \left(z_s + \frac{\gamma}{2}z_s^2\right)\mathrm{d}s + V\big(\tau,\lambda_\tau^{t,\lambda},H_\tau^{t,h;z}\big)\right].$$

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

We now proceed to characterize the value function V as the solution to a Hamilton-Jacobi-Bellman (HJB) partial integro-differential equation (PIDE). Further theoretical results on the value function, together with a verification theorem under suitable regularity assumptions, are proved in Ongarato (2026, Section 2.3).

Theorem 1. *Suppose that the value function V defined in (15) is of class $C^{1,1,1}$ (i.e., continuously differentiable in all its arguments). Then, the function V solves the following HJB-PIDE:*[6]

$$\begin{aligned} &\frac{\partial V}{\partial t} + \xi(\alpha - \lambda)\frac{\partial V}{\partial \lambda} - \rho h \frac{\partial V}{\partial h} \\ &+ \lambda(V(t, \lambda + \beta, h) - V(t, \lambda, h)) + \bar{\eta}(v - S(h, v))\lambda \\ &+ \frac{\left(\left(\frac{\partial V}{\partial h} - 1\right)^+\right)^2}{2\gamma} = 0, \\ &V(T, \lambda, h) = U(h). \end{aligned} \tag{16}$$

In addition, the optimal investment rate process z^* is given by

$$z^* = \frac{\left(\frac{\partial V}{\partial h} - 1\right)^+}{\gamma}. \tag{17}$$

Proof. In view of the assumption that V is of class $C^{1,1,1}$, standard arguments based on Itô's formula together with (11) and (12) imply that V satisfies the following HJB equation (see, e.g., Bensoussan & Chevalier-Roignant, 2024, Section 5.2):

$$\begin{aligned} 0 &= \sup_{z \geq 0} \left(\frac{\partial V}{\partial t} + \xi(\alpha - \lambda)\frac{\partial V}{\partial \lambda} - \rho h \frac{\partial V}{\partial h} + z\frac{\partial V}{\partial h} \right. \\ &\qquad + \lambda\big(V(t, \lambda + \beta, h) - V(t, \lambda, h)\big) \\ &\qquad \left. + \bar{\eta}(v - S(h, v))\lambda - z - \frac{\gamma}{2}z^2 \right) \\ &= \frac{\partial V}{\partial t} + \xi(\alpha - \lambda)\frac{\partial V}{\partial \lambda} - \rho h \frac{\partial V}{\partial h} \\ &\quad + \lambda(V(t, \lambda + \beta, h) - V(t, \lambda, h)) + \bar{\eta}(v - S(h, v))\lambda \\ &\quad + \sup_{z \geq 0} \left(z\frac{\partial V}{\partial h} - z - \frac{\gamma}{2}z^2 \right). \end{aligned}$$

The supremum in the last line is given by

$$\sup_{z \geq 0} \left(z\frac{\partial V}{\partial h} - z - \frac{\gamma}{2}z^2 \right) = \begin{cases} 0, & \text{if } \frac{\partial V}{\partial h} \leq 1, \\ \frac{1}{2\gamma}\left(\frac{\partial V}{\partial h} - 1\right)^2, & \text{otherwise,} \end{cases}$$

and is reached by the optimal control given in equation (17). □

A numerical method for the solution of the PIDE (16) will be presented in Section 4.2 and then applied in Section 5.

The optimal investment rate z_t^* in equation (17) depends on the current time t, on the current level λ_t of the stochastic intensity, and on the current cybersecurity level H_t. In particular, the dependence on λ_t makes z_t^* adaptive, meaning that it reacts to the random arrival of cyberattacks. Since the arrival of an attack increases the likelihood of further attacks due to the self-exciting

[6] For brevity of notation, in the statement and in the proof of this theorem, we omit to denote explicitly the dependence of V on its arguments (t, λ, h).

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

behavior of the Hawkes process, this enables the decision maker to strategically increase the cybersecurity investment in order to raise the cybersecurity level as a defense for the incoming attacks. This important feature will be numerically illustrated in Section 5.4.3.

Remark 8. The optimal policy described in equation (17) admits a clear economic interpretation: it is worth investing in cybersecurity whenever the marginal benefit of the investment is greater than its marginal cost. This insight aligns with the earlier results of Krutilla et al. (2021) in a dynamic but deterministic setup.

We close this section with the following result, which provides an explicit lower bound for the value function V.

Proposition 3. *For every* $(t, \lambda, h) \in [0, T] \times (0, \infty) \times \mathbb{R}_+$, *it holds that*

$$V(t, \lambda, h) \geq J(t, \lambda, h; \rho h), \tag{18}$$

where

$$J(t, \lambda, h; \rho h) = U(h) - \rho h \left(1 + \frac{\gamma}{2}\rho h\right)(T - t) + \bar{\eta}(\nu - S(h, \nu)) \left(\frac{\alpha\xi}{\xi - \beta}(T - t) - \frac{1}{\xi - \beta}\left(\lambda - \frac{\alpha\xi}{\xi - \beta}\right)\left(e^{-(\xi-\beta)(T-t)} - 1\right)\right).$$

Proof. By definition of the value function (15), it holds that $V(t, \lambda, h) \geq J(t, \lambda, h, z)$, for any given $z \in \mathcal{Z}_t$. In particular, the constant process $\bar{z} \equiv \rho h$ belongs to $\mathcal{Z}_t$, and, therefore, inequality (18) holds. In view of equation (14), we have that

$$H_s^{t,h;\rho h} = he^{-\rho(s-t)} + \int_t^s e^{-\rho(s-\nu)}\rho h \mathrm{d}\nu = h,$$

for all $s \in [t, T]$. Therefore, we obtain that

$$J(t, \lambda, h; \rho h) = \bar{\eta}\big(\nu - S(h, \nu)\big)\mathbb{E}\left[\int_t^T \lambda_s^{t,\lambda}\mathrm{d}s\right] - \rho h \left(1 + \frac{\gamma}{2}\rho h\right)(T - t) + U(h).$$

The expectation $\mathbb{E}[\int_t^T \lambda_s^{t,\lambda}\mathrm{d}s]$ can be computed by a straightforward adaptation of Proposition 1 (compare also with Dassios & Zhao, 2011, Theorem 3.6), thus completing the proof. □

Remark 9. The lower bound obtained in Proposition 3 is associated with a fixed investment rate that offsets technological obsolescence by maintaining the cybersecurity level constant over time (this follows directly from equation (12)). In Section 5.3, we numerically show that the optimal dynamic investment policy characterized in Theorem 1 consistently outperforms any constant investment strategy, highlighting the value of real-time adaptability in cybersecurity investment.

4. Numerical methods

In this section, we describe the parameters' choice and the numerical methods adopted for the solution of the optimization problem introduced in Section 3.

4.1 Specification of the model parameters

We report in Tables 1–3 the *standard set* of the model parameters. Unless mentioned otherwise, the numerical analysis will be performed using the standard set of parameters.

We employ a security breach probability function S of class I, as defined in (2). The parameters' values in Table 1 are consistent with those determined in Naldi and Flamini (2017) through a calibration based on synthetic data and also used as benchmark inputs in subsequent works (see,

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Table 1. Specification of the security breach function

Function type	v	a	b
S_I	0.65	10^{-1}	1

Table 2. Parameters of the stochastic intensity

α	ξ	β	λ_0
27	15	9	27

Table 3. Parameters of the optimization problem

δ	γ	$\tilde{\eta}$(k$)	$U(h)$	ρ	T
1	0.05	10	$\sqrt{h}$	0.2	1

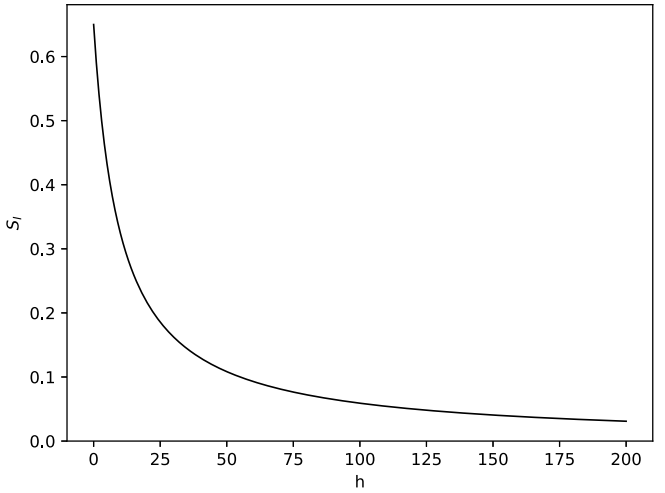


Figure 2. Security breach function (parameters as in Table 1).

e.g., Skeoch, 2022; Mazzoccoli & Naldi, 2020). Taking v, a, b as in Table 1, the function $h \mapsto S_I(h, v)$ is plotted in Figure 2.

The parameters of the stochastic intensity of the Hawkes process (see Table 2) are chosen to generate on average approximately 60 cyberattacks per year. We believe that this is a reasonable figure, in the absence of reliable estimates of the number of cyberattacks targeting a single entity.[7]

Remark 10. For the standard set of parameters, we have $\lambda_0 = \alpha$, and so the stochastic intensity λ_t can be expressed as follows:

$$\lambda_t = \lambda_0 + \beta \sum_{i=1}^{N_t} e^{-\xi(t-\tau_i)}. \tag{19}$$

[7]Empirical estimates of the intensity of cyberattacks can be found in the recent works (Bessy-Roland *et al.*, 2021; Boumezoued *et al.*, 2023; Li & Mamon, 2023). However, these estimates are not suitable for our purposes, since they are based on the number of attacks at a worldwide scale, while our model takes the viewpoint of a single entity. An average of 60 cyberattacks per year is in line with the number of attacks per organization reported in the Cyber Readiness Report 2024 by Hiscox (see Hiscox, 2024).

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Table 4. Meta-parameters for Algorithm 1

$\lambda_{\min}$	$\lambda_{\max}$	$\Delta\lambda$	$h_{\min}$	$h_{\max}$	Δh
27	216	1	0	50	0.5

We consider a one-year planning horizon ($T = 1$) and set an average loss of 10k for each successful breach, resulting in a total expected annual loss of approximately 390k without cybersecurity investments, which is in the same order of magnitude of Skeoch (2022). The depreciation rate is set at $\rho = 0.2$, consistently with the technological depreciation rates considered in Krutilla *et al.* (2021). The parameter γ is set at a rather low value in order to avoid an excessive penalization of large investment rates. Finally, we choose $U(h) = \sqrt{h}$, representing a strictly increasing and concave CRRA utility function.

4.2 Numerical solution of the HJB-PIDE

As shown in Section 3, determining the optimal cybersecurity investment requires the solution of the non-linear PIDE (16). Due to the complexity of the problem, one cannot expect to find explicitly an analytical solution, and, hence, numerical methods are required. We opt for the *method of lines*, as described in Yuan (1999). This technique consists of discretizing the PIDE in the spatial domain $(\lambda, h) \in (0, \infty) \times \mathbb{R}_+$ but not in time and then in integrating the semi-discrete problem as a system of ODEs. In our setting, we discretize the (λ, h) dimensions with a central difference and then numerically solve the resulting ODE system. Similarly to the case of PIDEs arising in Lévy models (see, e.g., Cont & Voltchkova, 2005), the unbounded space domain $(0, \infty) \times \mathbb{R}_+$ is localized into a bounded domain $[\lambda_{\min}, \lambda_{\max}] \times [h_{\min}, h_{\max}]$. We refer to Algorithm 1 for a precise description of the implementation of this method.

In our implementation, the algorithm's meta-parameters are specified as in Table 4.

The value $h_{\min} = 0$ corresponds to the absence of cybersecurity investment, while $h_{\max} = 50$ represents an upper bound that is rarely achieved in our setup under the standard parameter set. We choose $\lambda_{\min} = \lambda_0$, which coincides with the lower bound of the stochastic intensity λ_t; see equation (19). We set $\lambda_{\max} = \mathbb{E}[\lambda_T] + 7\sqrt{\mathrm{Var}(\lambda_T)} \approx 216$ in order to ensure that the truncation of the intensity domain does not have any material impact on our numerical results. The value function V is extrapolated beyond $[\lambda_{\min}, \lambda_{\max}]$ by setting

$$V(t, \lambda, h) = V(t, \lambda_{\max}, h), \quad \text{for all } \lambda > \lambda_{\max},$$

analogously to the scheme implemented in Ly Vath *et al.* (2025, Section 5.1). When plotting the function V in Section 5, we shall consider a subinterval of $[\lambda_{\min}, \lambda_{\max}]$: intensity values close to $\lambda_{\max}$ are rarely achieved and might lead to numerical instabilities.

We have implemented Algorithm 1 in Python, using the built-in ODE solver scipy.integrate .solve_ivp. We make use of an implicit Runge–Kutta method of the Radau IIA family of order 5 (see Hairer *et al.*, 1993 for further details). The computations were performed on an Intel Xeon E5520 CPU equipped with 32 GB of RAM, requiring approximately 23 hours to compute the value function and the corresponding optimal control. To empirically assess the convergence of Algorithm 1, we verified that standard discretization error metrics decrease consistently as the spatial grid in (λ, h) is refined.

4.3 Optimal cybersecurity investment rate

Besides determining the optimal net benefit of cybersecurity investments, we aim at computing the real-time adaptive strategy that best responds to the arrival of cyberattacks. To this end, after solving the PIDE (16) via Algorithm 1, we compute numerically the optimal investment rate given in equation (17) along a simulated sequence of cyberattacks. This entails simulating a trajectory of the stochastic intensity $(\lambda_t(\omega))_{t \in [t_{\text{init}}, T]}$, starting from an initial cybersecurity level H_{init} at time

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Algorithm 1 Numerical solution of the PIDE (16)

1: Set $\lambda_{\min}, \lambda_{\max}, h_{\min}, h_{\max}$.
2: Discretize $[\lambda_{\min}, \lambda_{\max}]$, with $\lambda_0 = \lambda_{\min}$, $\lambda_N = \lambda_{\max}$ and $\lambda_n - \lambda_{n-1} = \Delta\lambda$, for $n = 1, \ldots, N$.
3: Discretize $[h_{\min}, h_{\max}]$, with $h_0 = h_{\min}$, $h_M = h_{\max}$ and $h_m - h_{m-1} = \Delta h$, for $m = 1, \ldots, M$.
4: Set $V_{n,m}(t) := V(t, \lambda_n, h_m)$, for all n and m.
5: Approximate the partial derivatives w.r.t. λ:

$$\frac{\partial V}{\partial \lambda}(t, \lambda_n, h_m) \approx \frac{V_{n+1,m}(t) - V_{n-1,m}(t)}{2\Delta\lambda},$$
$$\frac{\partial V}{\partial \lambda}(t, \lambda_0, h_m) \approx \frac{V_{1,m}(t) - V_{0,m}(t)}{\Delta\lambda},$$
$$\frac{\partial V}{\partial \lambda}(t, \lambda_N, h_m) \approx \frac{V_{N,m}(t) - V_{N-1,m}(t)}{\Delta\lambda}.$$

6: Approximate the partial derivatives w.r.t. h:

$$\frac{\partial V}{\partial h}(t, \lambda_n, h_m) \approx \frac{V_{n,m+1}(t) - V_{n,m-1}(t)}{2\Delta h},$$
$$\frac{\partial V}{\partial h}(t, \lambda_n, h_0) \approx \frac{V_{n,1}(t) - V_{n,0}(t)}{\Delta h},$$
$$\frac{\partial V}{\partial h}(t, \lambda_n, h_M) \approx \frac{V_{n,M}(t) - V_{n,M-1}(t)}{\Delta h}.$$

7: Let $\tilde{n} = \frac{\lfloor \beta \rfloor}{\Delta\lambda}$ and set

$$V(t, \lambda_n + \beta, h_m) \approx V_{(n+\tilde{n})\wedge N, m}(t).$$

8: Solve the ODE system given for all n, m by

$$V'_{n,m}(t) = \xi(\lambda_n - \alpha)\frac{V_{n+1,m}(t) - V_{n-1,m}(t)}{2\Delta\lambda} + \rho h \frac{V_{n,m+1}(t) - V_{n,m-1}(t)}{2\Delta h}$$
$$-\lambda_n(V_{n+\tilde{n}\wedge N,m}(t) - V_{n,m}(t)) - \bar{\eta}(\nu - S(h_m, \nu))\lambda_n - \frac{\left(\left(\frac{V_{n,m+1}(t) - V_{n,m-1}(t)}{2\Delta h} - 1\right)^+\right)^2}{2\gamma},$$
$$V_{n,m}(T) = U(h_m).$$

t_{init}. Our numerical method for the computation of the optimal investment rate is described in Algorithm 2 and will be numerically implemented in Sections 5.4.2 and 5.4.3.

5. Results and discussion

In this section, we report some numerical results that illustrate the key properties and implications of the model. In particular, we are interested in assessing the benefit of adopting the optimal dynamic cybersecurity investment policy.

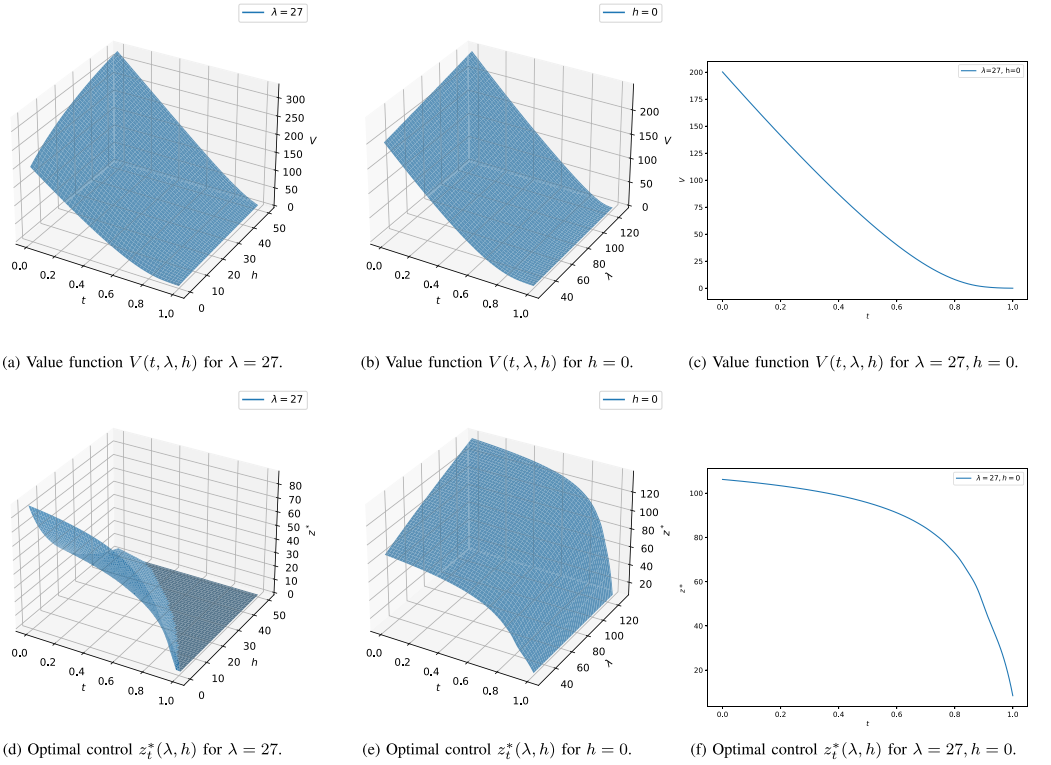
5.1 Value function and optimal cybersecurity policy

Figure 3 displays the value function V and the optimal cybersecurity investment rate z^*. In panels Figure 3(a) and (d), we plot, respectively, V and z^* for fixed intensity $\lambda = 27$, varying t and h. Coherently with Remark 7, we observe that the value function is increasing in h, while the optimal investment rate is decreasing. This behavior reflects the fact that higher cybersecurity levels yield greater benefits and reduce the need for further cybersecurity investments. In panels Figure 3(b) and (e) we plot, respectively, V and z^* for fixed $h = 0$, varying t and λ. Coherently with Remark 7,

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Algorithm 2 Numerical computation of the optimal control

Set $t_{\min}, t_{\max}, \lambda_{\min}, \lambda_{\max}, h_{\min}, h_{\max}$.
Discretize $[t_{\min}, t_{\max}]$, with $t_0 = t_{\min}, t_I = t_{\max}$ and $t_i - t_{i-1} = \Delta t$, for $i = 1, \dots, I$.
Discretize $[\lambda_{\min}, \lambda_{\max}]$, with $\lambda_0 = \lambda_{\min}, \lambda_N = \lambda_{\max}$ and $\lambda_n - \lambda_{n-1} = \Delta\lambda$, for $n = 1, \dots, N$.
Discretize $[h_{\min}, h_{\max}]$, with $h_0 = h_{\min}, h_M = h_{\max}$ and $h_m - h_{m-1} = \Delta h$, for $m = 1, \dots, M$.
Compute $V(t_i, \lambda_n, h_m)$ and $z^*(t_i, \lambda_n, h_m)$, for $i = 0, \dots, I$, $n = 0, \dots N$, $m = 0, \dots, M$.
Simulate a trajectory $\lambda_{t_i}(\omega)$, $i = 0, \dots, I$.
For the initial time $t_{\text{init}} \geq t_{\min}$, set $\bar{i} := \text{argmin}_i\{|t_i - t_{\text{init}}|\}$.
Consider the initial state $H_{t_{\bar{i}}} = H_{\text{init}}$:
for i in $\bar{i}, \dots, I$, **do**
    set $k := \text{argmin}_\ell\{|\lambda_\ell - \lambda_{t_i}(\omega)|\}$;
    set $j := \text{argmin}_m\{|h_m - H_{t_i}|\}$;
    let $z^*_{t_i} = z^*(t_i, \lambda_k, h_j)$;
    $H^{z^*}_{t_{i+1}} := H^{z^*}_{t_i} - \rho H^{z^*}_{t_i}\Delta t + z^*_{t_i}\Delta t$.
end for



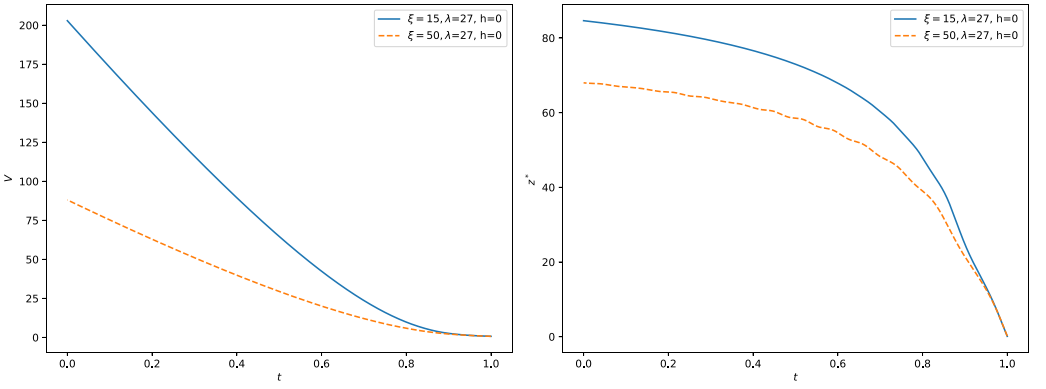
(a) Value function $V(t, \lambda, h)$ for $\lambda = 27$. (b) Value function $V(t, \lambda, h)$ for $h = 0$. (c) Value function $V(t, \lambda, h)$ for $\lambda = 27, h = 0$.

(d) Optimal control $z^*_t(\lambda, h)$ for $\lambda = 27$. (e) Optimal control $z^*_t(\lambda, h)$ for $h = 0$. (f) Optimal control $z^*_t(\lambda, h)$ for $\lambda = 27, h = 0$.

Figure 3. Value function and optimal investment rate computed under the standard parameters set.

we observe that both the value function and the optimal investment rate are increasing in λ. This is explained by the fact that, in the presence of a higher risk of cyberattacks, investing in cybersecurity becomes more valuable due to the larger potential of mitigating expected losses. As can be seen from panels Figure 3(c) and (f), both the value function and the optimal investment rate decrease over time. This is due to the fact that, under the standard parameter configuration (see

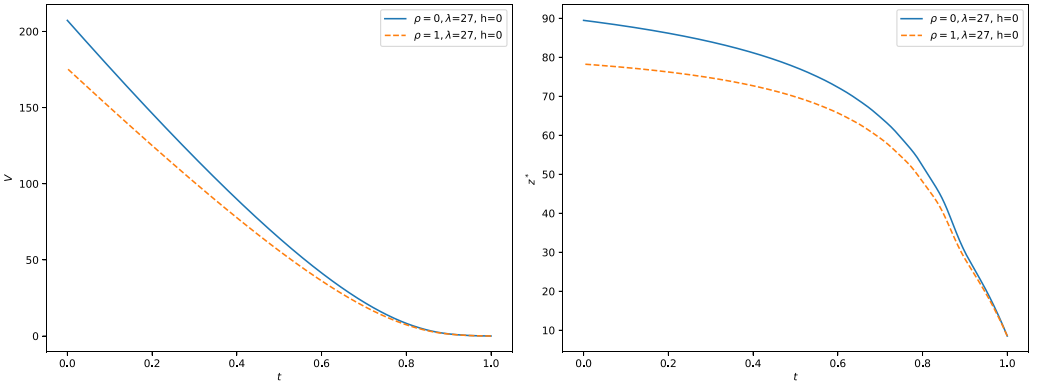
https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press



(a) Value function for different ξ, $\lambda = 27, h = 0$.

(b) Optimal control for different ξ, $\lambda = 27, h = 0$.

Figure 4. Value function and optimal investment rate for $\xi = 15$ and $\xi = 50$, for fixed h and λ.



(a) Value function $V(t, \lambda, h)$ for $\lambda = 27, h = 0$.

(b) Optimal control $z_t^*(\lambda, h)$ for $\lambda = 27, h = 0$.

Figure 5. Value function and optimal investment rate for $\rho = 0$ and $\rho = 1$, for fixed h and λ.

Section 4.1), the residual utility $U(H_T)$ of cybersecurity plays a relatively minor role and, therefore, the value of additional cybersecurity investment declines as the end of the planning horizon $[0, T]$ approaches.

5.2 Parameter sensitivity

In the proposed model, the parameter ξ determines the clustering behavior of cyberattacks. Specifically, higher values of ξ correspond to a more rapid decay of the intensity following each attack, thereby reducing the likelihood of temporally clustered attacks. To evaluate the impact of clustered cyberattacks, we compare in Figure 4 the value function and the optimal investment rate under two scenarios: $\xi = 15$ (more clustered attacks) and $\xi = 50$ (less clustered attacks). We observe that both the value function and the optimal investment rate are substantially greater in the case $\xi = 15$: if cyberattacks occur in clustered patterns, it is optimal to invest more in cybersecurity in order to mitigate the risk of large cumulative losses arising from rapid attack sequences. This finding underscores the critical importance of accounting for clustering dynamics in the optimal management of cyber-risk.

We also analyze the role of obsolescence in cybersecurity investment decisions, motivated by the analysis in Krutilla *et al.* (2021), which highlights its significance in a dynamic setup. Figure 5 displays the value function and the optimal investment rate under two contrasting depreciation

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

scenarios: $\rho = 0$ (no obsolescence) and $\rho = 1$ (high obsolescence).[8] We observe that $\rho = 1$ leads to smaller values for V and z^*, in line with the findings of Krutilla *et al.* (2021) in a deterministic setup. Our results confirm that a rapid depreciation of cybersecurity effectiveness reduces both the expected net benefit and the incentive to invest.

5.3 Comparison with a static investment strategy

The optimal investment rate z^* characterized in Theorem 1 represents the real-time adaptive cybersecurity policy that best responds to the arrival of cyberattacks.

In order to assess whether the adoption of an adaptive dynamic strategy provides a tangible benefit, we compare it against the best constant investment strategy, i.e., the strategy $z_t = \bar{z}$, for all $t \in [0, T]$, that maximizes the benefit-cost trade-off functional J. When investing according to a constant rate $\bar{z}$, the benefit-cost trade-off functional J takes the following form:

$$\begin{aligned} J(t, \lambda, h; \bar{z}) = &\int_t^T \bar{\eta}\big(\nu - S\big(H_s^{t,h;\bar{z}}, \nu\big)\big)\mathbb{E}\big[\lambda_s^{t,\lambda}\big]\mathrm{d}s \\ &- (T-t)\left(\bar{z} + \frac{\gamma}{2}\bar{z}^2\right) + U\big(H_T^{t,h;\bar{z}}\big), \end{aligned} \tag{20}$$

where the cybersecurity level $H^{t,h,\bar{z}}$ is given by

$$H_s^{t,h;\bar{z}} = he^{-\rho(s-t)} + \frac{\bar{z}}{\rho}\big(1 - e^{-\rho(s-t)}\big).$$

The optimal constant investment rate $\bar{z}^*$ solves the problem

$$J(t, \lambda, h; \bar{z}^*) := \sup_{\bar{z} \in \mathbb{R}_+} J(t, \lambda, h; \bar{z}). \tag{21}$$

In view of Proposition 1, the expectation $\mathbb{E}[\lambda_s^{t,\lambda}]$ in (20) can be computed in closed form. Therefore, the optimization problem (21) reduces to a deterministic maximization with respect to a scalar variable, which can be easily solved numerically. To this effect, we adopt the built-in global scalar optimizer scipy.optimize.differential_evolution in Python.

We quantify the relative gain obtained by investing according to the optimal dynamic policy z^* versus the constant policy $\bar{z}^*$ by computing the following quantity:

$$\%\text{gain}(t, \lambda, h) := 100 \times \frac{V(t, \lambda, h) - J(t, \lambda, h; \bar{z}^*)}{J(t, \lambda, h; \bar{z}^*)}. \tag{22}$$

Figure 6 displays the relative gain over time for varying cybersecurity levels h, for $\lambda = 27$ fixed. At the initial time $t = 0$, the gain reaches 15% for $h = 0.5$, 14% for $h = 1$, and 12% for $h = 2$, while it is 9.04% for $h = 5$, 5.7% for $h = 10$, and 2.6% for $h = 20$. These results show that the optimal dynamic investment strategy z^* consistently outperforms the best constant strategy $\bar{z}^*$, underscoring the importance of adaptive and responsive cybersecurity investments. The fact that the gain is rather small for large initial cybersecurity levels is coherent with the findings in Section 5.1: when the initial cybersecurity level is already high, the benefit of further investments diminishes, thereby reducing the relative advantage of the optimal policy. Moreover, a further analysis shows that the gain increases monotonically with respect to λ, indicating that the advantage of adopting the dynamic optimal policy (17) becomes more pronounced in high-risk scenarios.

[8] In practical terms, a depreciation rate of $\rho = 1$ implies that a given initial cybersecurity level H_0 depreciates by over 73% over a one-year period.

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

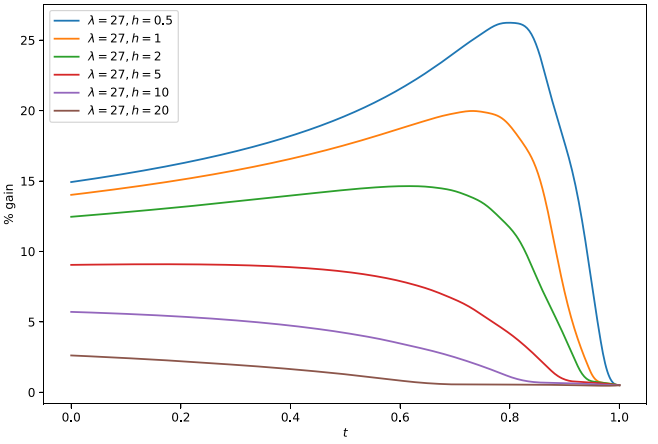


Figure 6. Relative gain with respect to the optimal constant investment rate.

5.4 Comparison with a standard Poisson model

To further assess the impact of clustered cyberattacks, we compare our model, which features a self-exciting Hawkes process, with a simplified version based on a standard Poisson process. A Poisson process $P = (P_t)_{t\in[0,T]}$ is characterized by a constant intensity λ^P and does not capture any temporal dependence in the arrival of attacks. Indeed, conditionally on $P_T = n$, the attack times are distributed as the order statistics of n i.i.d. random variables uniformly distributed on $[0, T]$, for every $n \in \mathbb{N}$. This setup can be recovered as a special case of the model introduced in Section 2.2 by setting $\beta = 0$ and $\alpha = \lambda_0$ in the intensity dynamics (4).

We consider the same optimization problem as in Section 3, and we replace the Hawkes process N with a Poisson process P of constant intensity λ^P and denote the resulting optimal investment rate by z^{P*}. A key observation is that, in this case, problem (13) reduces to a deterministic optimal control problem. The associated value function, $V^P(t, h)$, solves the following PDE:

$$\begin{aligned}&\frac{\partial V^P}{\partial t} - \rho h \frac{\partial V^P}{\partial h} + \lambda^P \bar{\eta}\,(\nu - S(h,\nu)) + \frac{\left(\left(\frac{\partial V^P}{\partial h} - 1\right)^+\right)^2}{2\gamma} = 0,\\ &V^P(T,h) = U(h).\end{aligned} \tag{23}$$

This PDE can be numerically solved using a scheme similar to Algorithm 1. Analogously to Theorem 1, the optimal investment rate z^{P*} in the Poisson model is given by

$$z^{P*} = \frac{\left(\frac{\partial V^P}{\partial h} - 1\right)^+}{\gamma}. \tag{24}$$

Remark 11. The optimal policy z^{P*} is deterministic. This is due to the fact that, in the Poisson model, the occurrence of a cyberattack does not carry any informational content.

We compare the Hawkes-based model with two Poisson-based benchmarks:

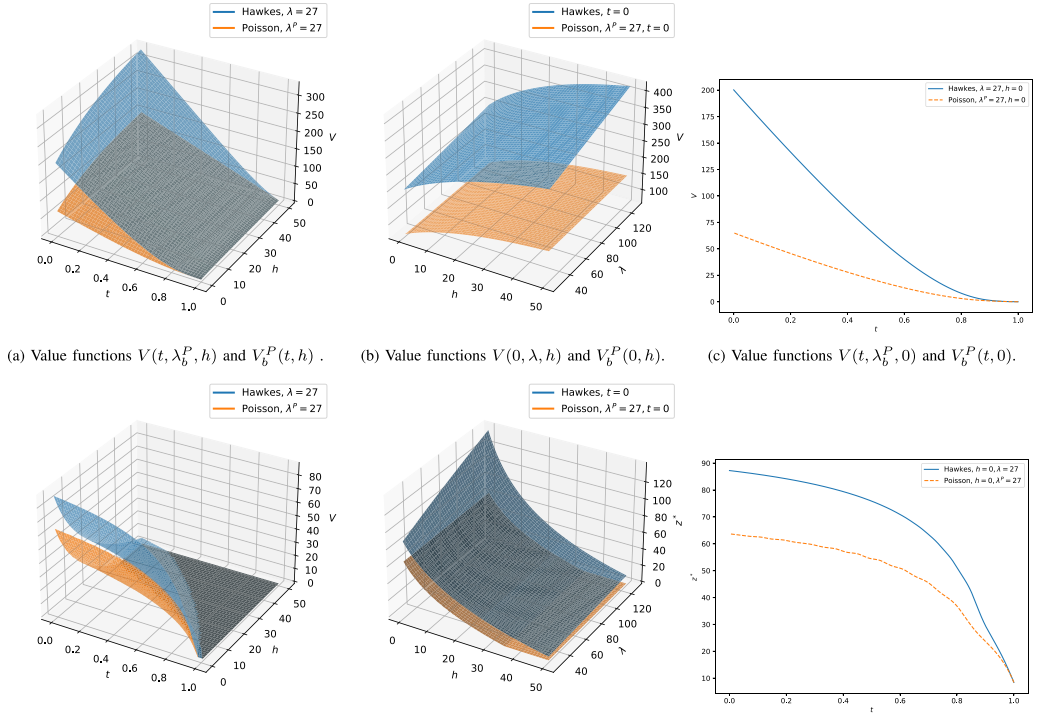
(i) a Poisson model with intensity λ_b^P chosen as

$$\lambda_b^P = \lambda_0 = 27; \tag{25}$$

(ii) a Poisson model with intensity λ_e^P chosen as

$$\lambda_e^P = \frac{\lambda_0 \xi}{\xi - \beta} + \frac{1 - e^{-\xi T}}{T(\xi - \beta)}\left(\lambda_0 - \frac{\lambda_0 \xi}{\xi - \beta}\right) \approx 61. \tag{26}$$

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press



(a) Value functions $V(t, \lambda_b^P, h)$ and $V_b^P(t, h)$.

(b) Value functions $V(0, \lambda, h)$ and $V_b^P(0, h)$.

(c) Value functions $V(t, \lambda_b^P, 0)$ and $V_b^P(t, 0)$.

(d) Optimal controls $z_t^*(\lambda_b^P, h)$ and $z_t^{P^*,b}(h)$.

(e) Optimal controls $z_0^*(\lambda_b^P, h)$ and $z_0^{P^*,b}(h)$.

(f) Optimal controls $z_t^*(\lambda_b^P, 0)$ and $z_t^{P^*,b}(0)$.

Figure 7. Comparison with a Poisson model with constant intensity $\lambda_b^P = 27$.

The first case corresponds to a Poisson process with the same baseline intensity of the Hawkes process. This scenario can be thought of as the situation where the entity underestimates the likelihood of cyberattacks (possibly due to relying on a limited or unrepresentative dataset) and considers it to be constant over time. In the second case, in view of Proposition 2, the value λ_e^P is chosen so that $\mathbb{E}[P_T] = \mathbb{E}[N_T]$, ensuring that the Hawkes-based model and the Poisson model with intensity λ_e^P generate the same expected number of cyberattacks over the planning horizon $[0, T]$. This reflects a case where the average attack frequency is estimated correctly, but the clustering dynamics are ignored.

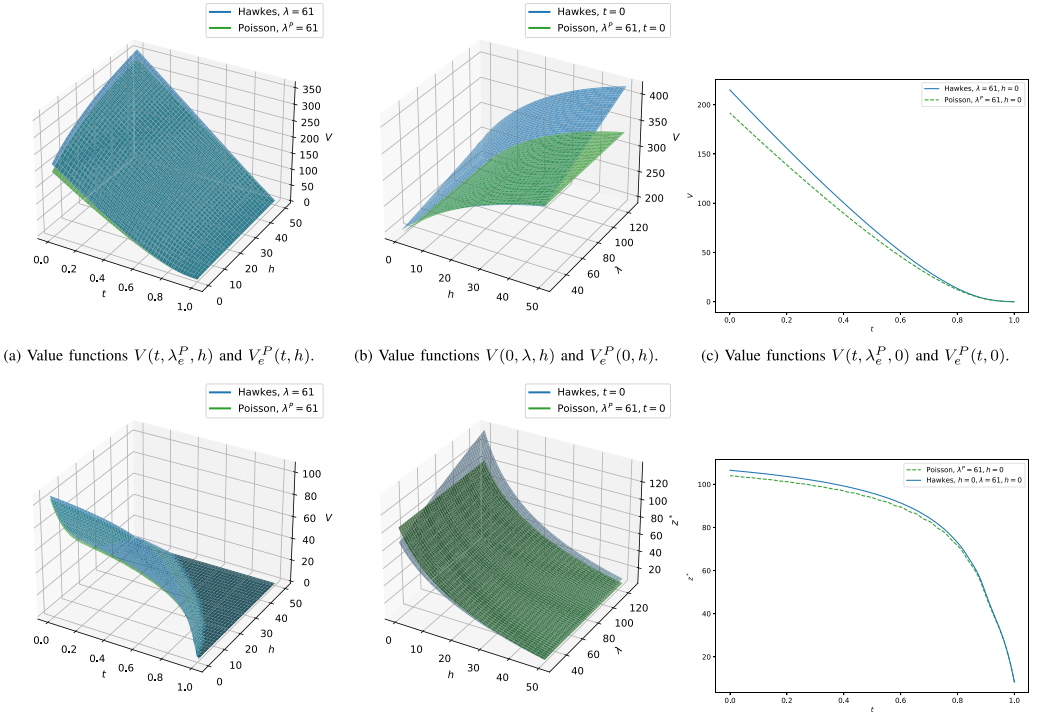
We shall make use of the following notation:

- $V_b^P(t, h)$ is the value function associated with the PDE (23) for the Poisson model with intensity λ_b^P specified in (25), and $z_t^{P*,b}(h)$ is the associated optimal control;
- $V_e^P(t, h)$ is the value function associated with the PDE (23) for the Poisson model with intensity λ_e^P specified in (26), and $z_t^{P*,e}(h)$ is the associated optimal control;
- $V(t, \lambda, h)$ is the value function associated with the PIDE (16), and $z_t^*(\lambda, h)$ is the associated optimal control.

5.4.1 Value functions and optimal cybersecurity policies

Figure 7 displays the results of the comparison with the Poisson model (i) with intensity λ_b^P. We observe that both the value function and the optimal cybersecurity investment rate under the Hawkes-based model consistently dominate their counterparts in the Poisson model (i) across the entire planning horizon. This is a direct consequence of the fact that $\lambda_t \geq \lambda_b^P$, for all $t \in [0, T]$.

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press



(a) Value functions $V(t, \lambda_e^P, h)$ and $V_e^P(t, h)$.

(b) Value functions $V(0, \lambda, h)$ and $V_e^P(0, h)$.

(c) Value functions $V(t, \lambda_e^P, 0)$ and $V_e^P(t, 0)$.

(d) Optimal controls $z_t^*(\lambda_e^P, h)$ and $z_t^{P^*,e}(h)$.

(e) Optimal controls $z_0^*(\lambda, h)$ and $z_0^{P^*,e}(h)$.

(f) Optimal controls $z_t^*(\lambda_e^P, 0)$ and and $z_t^{P^*,e}(0)$.

Figure 8. Comparison with a Poisson model with constant intensity $\lambda_e^P = 61$.

In other words, the Poisson model (i) not only disregards the temporal clustering of cyberattacks but also systematically underestimates their frequency. As a result, the perceived benefit of cybersecurity investment is lower, leading in turn to a suboptimal investment strategy.

Figure 8 reports the comparison with the Poisson model (ii) with intensity λ_e^P. Panels Figure 8(a) and (d) show that the benefit of cybersecurity investment and the optimal investment rate are slightly greater in the presence of clustered attacks (Hawkes-based model). This finding is confirmed in Panels Figure 8(c) and (f), which compare the value functions and optimal investment rates for fixed values $\lambda = \lambda_e^P$ and $h = 0$. Further insight is provided by panels Figure 8(b) and (e), which display, respectively, the value functions and the optimal investment rates at the initial time $t = 0$, across varying intensity levels. We can observe that the difference between the Hawkes and the Poisson models is negligible for small values of λ, while it becomes increasingly pronounced at higher values of λ. Interestingly, panel Figure 8(e) shows that the optimal investment rate under the Hawkes model may be either higher or lower than that in the Poisson model, depending on whether the current intensity λ exceeds λ_e^P or not. This feature will be analyzed in more detail in Section 5.4.3 below. Overall, these findings indicate that even when the average attack intensity is correctly estimated, neglecting the temporal clustering of cyberattacks can lead to suboptimal cybersecurity investment decisions.

5.4.2 Relative gain

Proceeding similarly to Section 5.3, we now evaluate the additional benefit derived from implementing the optimal adaptive policy z^*, as defined in (17), relative to the dynamic but deterministic policy z^{P*} derived in a Poisson-based model. We assume that the underlying model is the one introduced in Section 2.2 and compute the value function V given in (15) via Algorithm 1, using

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Table 5. Meta-parameters for Algorithm 2

$h_{\min}$	$h_{\max}$	Δh	$\lambda_{\min}$	$\lambda_{\max}$	$\Delta\lambda$	t_{init}	H_{init}	$\lambda_t(\omega)$
0	50	0.5	27	216	1	t	h	λ^P

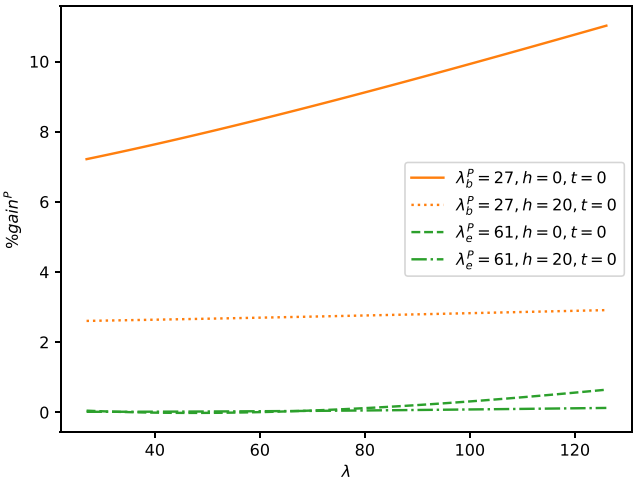


Figure 9. Relative gain with respect to the Poisson deterministic strategy, as defined in (27).

the standard parameter set described in Section 4.1. When employing the deterministic strategy z^{P*}, the expected net benefit from cybersecurity investment is quantified as follows:

$$J(t,\lambda,h;z^{P*}) = \int_t^T \left(\bar{\eta}\big(\nu - S\big(H_s^{t,h;\,z^{P*}},\nu\big)\big)\mathbb{E}\big[\lambda_s^{t,\lambda}\big]\right) \mathrm{d}s - \int_t^T \left(z_s^{P*} + \frac{\gamma}{2}\big(z_s^{P*}\big)^2\right) \mathrm{d}s + U\big(H_T^{t,h;\,z^{P*}}\big),$$

where $H_s^{t,h;\,z^{P*}}$ is defined as in (14) with $z = z^{P*}$ and $\mathbb{E}[\lambda_s^{t,\lambda}]$ can be computed explicitly by Proposition 1. Once the PDE (23) is numerically solved, z^{P*} can be computed via Algorithm 2 taking λ constant. Our numerical implementation of the Poisson-based model adopts the parameters specified in Table 5.

The gain of the optimal investment policy z^* with respect to z^{P*} is computed as follows, in analogy to (22):

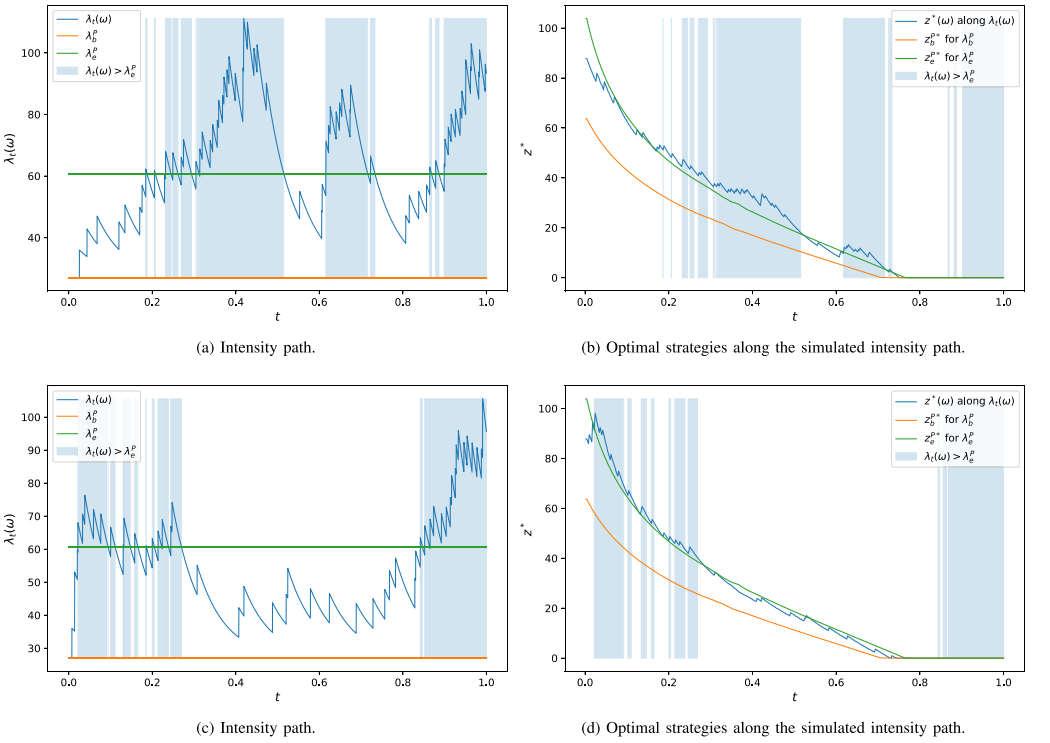
$$\%\text{gain}^P(t,\lambda,h) := 100 \times \frac{V(t,\lambda,h) - J(t,\lambda,h;z^{P*})}{J(t,\lambda,h;z^{P*})}. \tag{27}$$

Figure 9 reports the quantity $\%\text{gain}^P(t,\lambda,h)$ comparing the Hawkes-based model against two Poisson-based benchmarks with intensities λ_b^P and λ_e^P, as considered above. For $h = 0$, the gain increases with λ, ranging between 7.6% and 11.4% for λ_b^P and between 0.04% and 0.6% for λ_e^P. For $h = 20$, the gain becomes nearly constant in λ. The fact that the gain for λ_e^P is limited can be explained by the fact that the objective functional (10) is linear with respect to the losses and λ_e^P is chosen in such a way that the Poisson-based model generates the same expected losses of our Hawkes-based model.

5.4.3 Adaptive dynamics of the optimal investment policy

Finally, we illustrate the adaptive behavior of the optimal investment policy given in equation (17). While the overall improvement over a Poisson-based strategy may appear limited in terms

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press



(a) Intensity path.

(b) Optimal strategies along the simulated intensity path.

(c) Intensity path.

(d) Optimal strategies along the simulated intensity path.

Figure 10. Optimal strategies along simulated intensity paths.

of overall gain (see Figure 9), the key strength of our approach lies in its capacity to dynamically adjust the cybersecurity investment in response to the arrival of cyberattacks. To this effect, panels Figure 10(a) and (c) display two simulated paths of the Hawkes intensity $(\lambda_t)_{t\in[0,T]}$, alongside the constant intensities λ_b^P and λ_e^P defined in Section 5.4. The corresponding optimal investment policies are shown in panels Figure 10(b) and (d). Consistent with the analysis in Section 5.4.1, the optimal investment rate z_t^* is always larger than the Poisson-based benchmark z_b^{P*}, due to the fact that $\lambda_t \geq \lambda_b^P$, for all $t \in [0, T]$. In contrast, the comparison with the benchmark strategy z_e^{P*} is more nuanced. We highlight in cyan the time intervals during which $\lambda_t > \lambda_e^P$. Our simulations reveal that when $\lambda_t \leq \lambda_e^P$, the adaptive strategy z_t^* closely aligns with z_e^{P*}. However, when $\lambda_t > \lambda_e^P$, especially during extended periods resulting from clusters of cyberattacks, the investment rate z_t^* increases markedly, exceeding the corresponding deterministic strategy. This shows that the optimal cybersecurity investment policy z_t^* can react in real-time to rapid sequences of cyberattacks. Finally, under the standard parameter set, the investment rate naturally declines toward the end of the planning horizon $[0, T]$, as the accumulated cybersecurity level suffices to mitigate future risk.

6. Implications for cyber-insurance

In this section, we connect our dynamic model for optimal cybersecurity investment to actuarial aspects of cyber-insurance. From an insurance viewpoint, the cybersecurity investment policy acts as a form of self-insurance (see Böhme & Schwartz, 2010): by investing in cybersecurity, the decision maker changes the distribution of aggregate losses over the planning horizon. Cybersecurity investment is, therefore, a form of prevention, which represents a crucial aspect of

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

cyber-insurance, as discussed in Section 6.1 below. In turn, prevention measures such as cybersecurity investment have an impact on the determination of cyber-insurance premia, as shown in Section 6.2 below.

6.1 The role of prevention in cyber-insurance

Cyber-insurance contracts increasingly go beyond the pure compensation of financial losses and often include prevention requirements and bundled cyber-risk mitigation services. In this sense, cyber-insurance combines risk transfer with risk mitigation services and incentives, promoting best practices in cyber-risk management (see Lopez *et al.,* 2025, for a recent analysis).

In our model, prevention is represented by a cybersecurity investment policy $z \in \mathcal{Z}$ and has a precise actuarial meaning: it changes the aggregate losses L_T^z over the planning horizon by reducing the breach probability at attack times from the baseline ν to $S(H_{\tau_i}^z, \nu)$. Note that, by lowering the breach probability at attack times, in our model prevention does not only reduce the expected loss but also reduces loss variability. As a result, any underwriting or pricing rule that depends on the level of risk should, in principle, depend on the prevention policy (or on observable proxies of it, as discussed below). Our framework allows quantifying the effect of optimal prevention, which is an essential input for the insurer in order to determine cybersecurity requirements, incentive schemes, and premium differentiation rules across policyholders (see Section 6.2).

A key practical issue is that prevention is only partially observable, creating asymmetric information and moral hazard. Two implementation channels discussed in the literature are consistent with our framework. First, prevention may be delegated to the insurer or to an IT security provider in partnership with the insurer, which is particularly relevant for small firms or policyholders with limited internal expertise. This scheme makes prevention partly contractible and allows joint pricing of risk transfer together with mitigation services (see Zeller & Scherer, 2023). Second, when prevention is handled internally by the entity (as a self-protection measure), insurers may rely on incentive-based contracting. One approach is to condition the premium on minimum protection measures, combined with audits and penalties if requirements are not met (see Hofmann, 2007; Lelarge & Bolot, 2009). Another approach is to rely on principal–agent incentive schemes, as considered in Brachetta *et al.* (2025), Mastrolia and Yan (2025), linking part of the compensation to measurable reduction in risk or losses. In both cases, the insurer needs quantitative tools to assess the impact of prevention measures, as considered in our study.

6.2 Implications for premium determination

To quantify the impact of prevention on cyber-insurance premia, we apply the standard deviation principle (see, e.g., Schmidli, 2017), for an arbitrary cybersecurity investment policy $z \in \mathcal{Z}$:

$$\pi\left(L_T^z\right) := \mathbb{E}\left[L_T^z\right] + \theta\, \sigma\left(L_T^z\right), \tag{28}$$

where $\mathbb{E}[L_T^z]$ is the pure premium and $\theta\sigma(L_T^z)$ is a loading proportional to the standard deviation $\sigma(L_T^z)$ of L_T^z, with θ typically in $[0.1, 0.5]$. This principle is stable with respect to the monetary unit and captures the impact of prevention on both the expected loss and the dispersion of losses.

We first compare the expected aggregate loss without additional investment, $\mathbb{E}[L_T^0]$, with the expected aggregate loss under the optimal prevention policy z^*, denoted $\mathbb{E}[L_T^{z^*}]$, corresponding to (28) with $z = z^*$ and $\theta = 0$. We compute $\mathbb{E}[L_T^0]$ by Proposition 2, while $\mathbb{E}[L_T^{z^*}]$ is estimated by Monte Carlo over 100,000 samples, where z^* is computed via Algorithm 2 starting from $h = 0$. Under the standard parameter set, we obtain (in k$)

$$\mathbb{E}\left[L_T^0\right] = 394.98 \quad \text{and} \quad \mathbb{E}\left[L_T^{z^*}\right] = 141.77,$$

corresponding to a reduction of approximately 65% in the pure premium under the optimal dynamic prevention strategy.

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

The variance of the losses L_T^z associated to a given cybersecurity investment policy $z \in \mathcal{Z}$ is given in the following lemma. We denote $\bar{\eta} := \mathbb{E}[\eta_i]$ and $\sigma_\eta^2 := \mathrm{Var}(\eta_i)$, for $i \in \mathbb{N}^*$.

Lemma 1. *Suppose that Assumptions B and C hold. Assume furthermore that $\eta_i \in L^2(\mathbb{P})$, for all $i \in \mathbb{N}^*$. Then, for every $z \in \mathcal{Z}$, it holds that*

$$\mathrm{Var}(L_T^z) = \mathbb{E}\left[\int_0^T \Big(\sigma_\eta^2\, S(H_t^z, \nu) + \bar{\eta}^2 S(H_t^z, \nu)\big(1 - S(H_t^z, \nu)\big)\Big)\lambda_t \mathrm{d}t\right] + \bar{\eta}^2\, \mathrm{Var}\left(\int_0^T S(H_t^z, \nu)\, \mathrm{d}N_t\right).$$

Proof. Let L_T^z be defined as in (8), where $z = (z_t)_{t\in[0,T]}$ is an admissible cybersecurity investment process. By the law of total variance (see, e.g., Billingsley, 1986, Problem 34.10), we can compute

$$\mathrm{Var}\big(L_T^z\big) = \mathbb{E}\left[\mathrm{Var}\left(\sum_{i=1}^{N_T} \eta_i B_i^{S(H_{\tau_i}^z, \nu)} \middle| \mathcal{F}_T\right)\right] + \mathrm{Var}\left(\mathbb{E}\left[\sum_{i=1}^{N_T} \eta_i B_i^{S(H_{\tau_i}^z, \nu)} \middle| \mathcal{F}_T\right]\right) \tag{29}$$

For the first term on the right-hand side in (29), it holds that

$$\mathrm{Var}\left(\sum_{i=1}^{N_T} \eta_i B_i^{S(H_{\tau_i}^z, \nu)} \middle| \mathcal{F}_T\right) = \sum_{i=1}^{N_T} \mathrm{Var}\left(\eta_i B_i^{S(H_{\tau_i}^z, \nu)} \middle| \mathcal{F}_T\right)$$
$$= \sum_{i=1}^{N_T} \Big(\sigma_\eta^2\, S(H_{\tau_i}^z, \nu) + \bar{\eta}^2 S(H_{\tau_i}^z, \nu)\big(1 - S(H_{\tau_i}^z, \nu)\big)\Big),$$

while for the second term we have that

$$\mathbb{E}\left[\sum_{i=1}^{N_T} \eta_i B_i^{S(H_{\tau_i}^z, \nu)} \middle| \mathcal{F}_T\right] = \bar{\eta} \sum_{i=1}^{N_T} S(H_{\tau_i}^z, \nu).$$

Therefore, we can write that

$$\mathrm{Var}\big(L_T^z\big) = \sigma_\eta^2\, \mathbb{E}\left[\sum_{i=1}^{N_T} S(H_{\tau_i}^z, \nu)\right] + \bar{\eta}^2\, \mathbb{E}\left[\sum_{i=1}^{N_T} S(H_{\tau_i}^z, \nu)\big(1 - S(H_{\tau_i}^z, \nu)\big)\right] + \bar{\eta}^2\, \mathrm{Var}\left(\sum_{i=1}^{N_T} S(H_{\tau_i}^z, \nu)\right).$$

Arguing as in the proof of Proposition 2 yields the formula for $\mathrm{Var}(L_T^z)$ in the statement of the lemma. □

In the case $z \equiv 0$ (hence $S(0, \nu) = \nu$), the formula of Lemma 1 simplifies to

$$\mathrm{Var}\big(L_T^0\big) = \mathbb{E}[N_T]\big(\sigma_\eta^2 \nu + \bar{\eta}^2 \nu(1 - \nu)\big) + \bar{\eta}^2 \nu^2 \mathrm{Var}(N_T).$$

This expression can be made fully explicit by using the closed-form expressions for $\mathbb{E}[N_T]$ and $\mathrm{Var}(N_T)$ available for Hawkes processes (see Hillairet & Réveillac, 2025, Section 2.1). For the optimal policy z^*, the quantities appearing in Lemma 1 are computed by Monte Carlo over 100,000 samples.

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Table 6. Standard deviations

$\mathbb{E}[\eta]$	$\mathrm{Var}(\eta)$	$\sigma(L_T^0)$	$\sigma(L_T^{z^*})$	$1-[\sigma(L_T^{z^*})/\sigma(L_T^0)]$
10	10	118.56	51.21	56.80%
10	50	125.04	56.58	54.75%
10	100	132.70	62.65	52.79%

Table 7. Premia (std deviation principle)

$\mathbb{E}[\eta]$	$\mathrm{Var}(\eta)$	$\pi(L_T^0)$	$\pi(L_T^{z^*})$	$1-[\pi(L_T^{z^*})/\pi(L_T^0)]$
10	10	430.55	157.13	63.50%
10	50	432.5	158.74	63.32%
10	100	434.79	160.57	63.07%

Table 6 reports the standard deviations of L_T^0 and $L_T^{z^*}$ for different values of σ_η^2 (with $\bar{\eta}$ fixed), showing a reduction of roughly 53%–57% under optimal prevention. Table 7 reports the corresponding premia under (28) with $\theta = 0.3$. Across the scenarios considered, the premium under optimal prevention is about 63% lower than in the no-investment benchmark. These results quantify how dynamic optimal cybersecurity investment translates into lower premia.

The values $\pi(L_T^{z^*})$ and $\pi(L_T^0)$ provide natural benchmark lower and upper premia corresponding to optimal prevention and to the absence of additional cybersecurity investment, respectively. More generally, when policyholders are classified into prevention classes using observable proxies (e.g., cybersecurity scores or audits), the model supports premium differentiation across classes: better prevention (corresponding to higher cybersecurity levels) implies lower risk and, therefore, lower actuarial premia under standard premium principles. In this sense, the gap between $\pi(L_T^0)$ and $\pi(L_T^{z^*})$ provides a concrete estimate of the premium range induced by prevention under the standard deviation principle.

7. Conclusions

In this work, we introduce a dynamic and stochastic extension of the Gordon–Loeb model (Gordon & Loeb, 2002) for optimal cybersecurity investment, incorporating temporally clustered cyberattacks via a Hawkes process. Our modeling framework captures the empirically observed phenomenon of attack bursts, thus offering a more realistic representation of the current cyber-risk environment. We formulate the cybersecurity investment decision problem as a two-dimensional stochastic optimal control problem, maximizing the expected net benefit of cybersecurity investments. We allow for adaptive investment policies that respond in real-time to the arrival of cyberattacks.

Our numerical results demonstrate that the optimal cybersecurity investment policy consistently outperforms both static benchmarks and Poisson-based models that ignore clustering. In particular, even when Poisson models are calibrated to match the expected attack frequency, they fail to capture the implications of attack clustering on investment timing and magnitude, thus leading to suboptimal investment decisions. Our findings indicate that the optimal dynamic strategy is able to react promptly to attack clusters, offering substantial improvements in expected net benefit in high-risk scenarios. From an actuarial perspective, we interpret cybersecurity investment as a form of self-insurance, showing that improved prevention lowers both expected losses and loss variability, which in turn reduces cyber-insurance premia under standard premium principles. Overall, our results underscore the importance of accounting for dynamic and stochastic threat patterns in cybersecurity planning. The proposed framework supports risk

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

managers, insurers, and policymakers in designing responsive cybersecurity investment strategies and cyber-insurance pricing schemes that are consistent with the evolving cyber-risk landscape.

Future research directions include the empirical calibration to sector-specific cyber incident data, the consideration of risk aversion with respect to losses resulting from cyberattacks, and the integration of cyber-insurance as a complementary tool for risk mitigation (see Awiszus *et al.*, 2023; Dou *et al.*, 2020; Mazzoccoli & Naldi, 2020; Miaoui & Boudriga, 2019; Öğüt *et al.*, 2011; Skeoch, 2022; Skeoch & Ioannidis, 2024) for some recent studies in this direction). Our framework can also be applied from the viewpoint of an insurance firm that provides insurance against losses due to cyberattacks, thus laying the foundations for the development of Cramér–Lundberg-type models (see, e.g., Mikosch, 2009) for cyber-insurance. Finally, our modeling setup can also be extended to multivariate Hawkes processes (as considered in Embrechts *et al.*, 2011), or in the more general versions of Bielecki *et al.* (2022, 2023) to differentiate among multiple types of cyberattacks (see Bentley *et al.*, 2020) for a multivariate generalization of the static Gordon–Loeb model).

Acknowledgements. The authors are thankful to Alessandro Calvia, Salvatore Federico, Luca Grosset, Eleonora Losiouk, Daniele Marazzina, Andrea Pallavicini, Andrea Perchiazzo, Mario Putti, and Edit Rroji for valuable discussions on the topic of this work. Constructive comments by three anonymous reviewers and an associate editor are gratefully acknowledged.

Data availability statement. Data availability is not applicable to this articles as no new data were created or analyzed in this study. If needed, the authors agree to make available the code used for the numerical results illustrated in Section 5.

Funding statement. This work was supported by the Europlace Institute of Finance; Fondazione Gini; Unione Matematica Italiana and European Union-Next Generation EU [PRIN-2022BE-MMLZ]; and the Chair Stress Test, RISK Management and Financial Steering, led by the French Ecole Polytechnique and its Foundation and sponsored by BNP Paribas. The funders had no role in study design, decision to publish, or preparation of the manuscript.

Competing interests. All authors declare they do not have competing interests.

References

Awiszus, **K.**, **Knispel**, **T.**, **Penner**, **I.**, **Svindland**, **G.**, **Voß**, **A.**, & **Weber**, **S.** (2023). Modeling and pricing cyber insurance: Idiosyncratic, systematic, and systemic risks. *European Actuarial Journal*, **13**(1), 1–53.

Baldwin, **A.**, **Gheyas**, **I.**, **Ioannidis**, **C.**, **Pym**, **D.**, & **Williams**, **J.** (2017). Contagion in cyber security attacks. *Journal of the Operational Research Society*, **68**(7), 780–791.

Balzano, **M.**, & **Marzi**, **G.** (2025). At the cybersecurity frontier: Key strategies and persistent challenges for lusiness leaders. *Strategic Change*, **34**(2), 181–192.

Bensoussan, **A.**, & **Chevalier-Roignant**, **B.** (2024). Stochastic control for diffusions with self-exciting jumps: An overview. *Mathematical Control and Related Fields*, **14**(4), 1452–1476.

Bentley, **M.**, **Stephenson**, **A.**, **Toscas**, **P.**, & **Zhu**, **Z.** (2020). A multivariate model to quantify and mitigate cybersecurity risk. *Risks*, **8**(2), 61.

Bessy-Roland, **Y.**, **Boumezoued**, **A.**, & **Hillairet**, **C.** (2021). Multivariate Hawkes process for cyber insurance. *Annals of Actuarial Science*, **15**(1), 14–39.

Bielecki, **T. R.**, **Jakubowski**, **J.**, & Niewęgłowski (2022). Construction and simulation of generalized multivariate Hawkes processes. *Methodology and Computing in Applied Probability*, **24**, 2865–2896.

Bielecki, **T. R.**, **Jakubowski**, **J.**, & Niewęgłowski (2023). Multivariate Hawkes processes with simultaneous occurrence of excitation events coming from different sources. *Stochastic Models*, **39**(3), 537–565.

Billingsley, **P.** (1986). *Probability and measure* (2nd ed.). John Wiley & Sons.

Boumezoued, **A.**, **Cherkaoui**, **Y.**, & **Hillairet**, **C.** (2023). Cyber risk modeling using a two-phase Hawkes process with external excitation. arXiv preprint arXiv: 2311.15701.

Brachetta, **M.**, **Callegaro**, **G.**, **Ceci**, **C.**, & **Sgarra**, **C.** (2025). Optimal self-protection via BSDEs for risk models with jump clusters. *SIAM Journal on Control and Optimization*, **63**(6), 3990–4017.

Brémaud, **P.** (1981). *Point processes and queues: Martingale dynamics*. Springer.

Böhme, **R.**, & **Schwartz**, **G.** (2010). Modeling cyber-insurance: Towards a unifying framework. In *Workshop on the Economics of Information Security (WEIS), Harvard, June 2010*.

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Cont, **R.**, & **Voltchkova**, **E.** (2005). A finite difference scheme for option pricing in jump diffusion and exponential Lévy models. *SIAM Journal on Numerical Analysis*, **43**(4), 1596–1626.

Dassios, **A.**, & **Zhao**, **H.** (2011). A dynamic contagion process. *Advances in Applied Probability*, **43**(3), 814–846.

Dassios, **A.**, & **Zhao**, **H.** (2013). Exact simulation of Hawkes process with exponentially decaying intensity. *Electronic Communications in Probability*, **18**(62), 1–13.

Davis, **M.** (1984). Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society: Series B*, **46**(3), 353–376.

Dou, **W.**, **Tang**, **W.**, **Wu**, **X.**, **Qi**, **L.**, & **Xu**, **X.** (2020). An insurance theory based optimal cyber-insurance contract against moral hazard. *Information Sciences*, **527**, 576–589.

Embrechts, **P.**, **Liniger**, **T.**, & **Lin**, **L.** (2011). Multivariate Hawkes processes: An application to financial data. *Journal of Applied Probability*, **48(A)**(3), 367–378.

ENISA. (2024). ENISA threat landscape 2024 report. https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024.

Fedele, **A.**, & **Roner**, **C.** (2022). Dangerous games: A literature review on cybersecurity investments. *Journal of Economic Surveys*, **36**(1), 157–187.

Gordon, **L. A.**, & **Loeb**, **M. P.** (2002). The economics of information security investment. *ACM Transactions on Information and System Security*, **5**(4), 438–457.

Gordon, **L. A.**, **Loeb**, **M. P.**, & **Lucyshyn**, **W.** (2003). Information security expenditures and real options: A wait-and-see approach. *Computer Security Journal*, **19**(2), 1–7.

Gordon, **L. A.**, **Loeb**, **M. P.**, & **Zhou**, **L.** (2020). Integrating cost–benefit analysis into the NIST cybersecurity framework via the Gordon–Loeb model. *Journal of Cybersecurity*, **6**(1), tyaa005.

Hairer, **E.**, **Nørsett**, **S. P.**, & **Wanner**, **G.** (1993). *Solving ordinary differential equations II: Stiff and differential-algebraic problems*. Springer.

Hawkes, **A. G.** (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, **58**(1), 83–90.

Hawkes, **A. G.**, & **Oakes**, **D.** (1974). A cluster process representation of a self-exciting process. *Journal of Applied Probability*, **11**, 493–503.

Hayes, **J.**, & **Bodhani**, **A.** (2013). Cyber security: Small firms under fire. *Engineering & Technology*, **8**(6), 80–83.

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.

Hillairet, **C.**, **Peyrat**, **T.**, & **Réveillac**, **A.** (2025). Multivariate self-exciting processes with dependencies. arXiv preprint arXiv: 2503.15958.

Hillairet, **C.**, & **Réveillac**, **A.** (2025). Explicit correlations for the Hawkes processes. *Stochastics-an International Journal of Probability and Stochastic Processes*, 1–29.

Hiscox. (2024). Cyber readiness report 2024. https://www.hiscoxgroup.com/cyber-readiness.

Hofmann, **A.** (2007). Internalizing externalities of loss prevention through insurance monopoly: An analysis of interdependent risks. *The Geneva Risk and Insurance Review*, **32**(1), 91–111.

Huang, **D. C.**, & **Behara**, **R. S.** (2013). Economics of information security investment in the case of concurrent heterogeneous attacks with budget constraints. *International Journal of Production Economics*, **141**(1), 255–268.

Krutilla, **K.**, **Alexeev**, **A.**, **Jardine**, **E.**, & **Good**, **D.** (2021). The benefits and costs of cybersecurity risk reduction: A dynamic extension of the Gordon and Loeb model. *Risk Analysis*, **41**(10), 1795–1808.

Laub, **P. J.**, **Lee**, **Y.**, & **Taimre**, **T.** (2021). *The elements of Hawkes processes*. Springer.

Lelarge, **M.**, & **Bolot**, **J.** (2009). Economic incentives to increase security in the internet: The case for insurance. In *IEEE INFOCOM 2009, Rio de Janeiro, Brazil* (pp. 1494–1502).

Li, **Y.**, & **Mamon**, **R.** (2023). Modelling health-data breaches with application to cyber insurance. *Computers and Security*, **124**, 102963.

Lima, **R.** (2023). Hawkes processes modeling, inference, and control: An overview. *SIAM Review*, **65**(2), 331–374.

Lopez, **O.**, **Denuit**, **M.**, **Ghossoub**, **M.**, **Trufin**, **J.**, **Kher**, **J.**, **Maillart**, **A.**, **Raes**, **E.**, **Rapior**, **H.**, **Skoubani**, **M.-A.**, & **Spoorenberg**, **B.** (2025). Cyber risk: Quantification, stress scenarios, mitigation, and insurance, CAS research paper.

Ly Vath, **V.**, **Scotti**, **S.**, & **Gaigi**, **M.** (2025). Optimal harvesting under uncertain environment with clusters of catastrophes. *Journal of Economic Dynamics and Control*, **179**, 105165.

Mastrolia, **T.**, & **Yan**, **H.** (2025). Agency problems and adversarial bilevel optimization under uncertainty and cyber threats, working paper. https://arxiv.org/abs/2505.08989.

Mazzoccoli, **A.**, & **Naldi**, **M.** (2020). Robustness of optimal investment decisions in mixed insurance/investment cyber risk management. *Risk Analysis*, **40**(3), 550–564.

Mazzoccoli, **A.**, & **Naldi**, **M.** (2022). An overview of security breach probability models. *Risks*, **10**(11), 220.

Miaoui, **Y.**, & **Boudriga**, **N.** (2019). Enterprise security economics: A self-defense versus cyber-insurance dilemma. *Applied Stochastic Models in Business and Industry*, **35**(3), 448–478.

Mikosch, **T.** (2009). *Non-life insurance mathematics: An introduction with the Poisson process* (2nd ed.). Springer.

Naldi, **M.**, & **Flamini**, **M.** (2017). Calibration of the Gordon–Loeb models for the probability of security breaches. In *19th International Conference on Computer Modelling & Simulation (UKSim)*.

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press

Nguyen, **D. D. A.**, **Alain**, **P.**, **Autrel**, **F.**, **Bouabdallah**, **A.**, **François**, **J.**, & **Doyen**, **G.** (2024). How fast does malware leveraging Eternal Blue propagate? The case of Wanna Cry and NotPetya. In *2024 IEEE 10th international conference on network softwarization (NetSoft)* (pp. 399–404).

Öğüt, **H.**, **Raghunathan**, **S.**, & **Menon**, **N.** (2011). Cyber security risk management: Public policy implications of correlated risk, imperfect ability to prove loss, and observability of self-protection. *Risk Analysis*, **31**(3), 497–512.

Ongarato, **B.** (2026). Hawkes and affine processes in risk modeling: Theory and applications in finance and cybersecurity. *Ph.D. dissertation*. University of Padova.

Pham, **H.** (2009). *Continuous-time stochastic control and optimization with financial applications*. Springer.

SANS Institute (2025). 2025 ICS/OT cybersecurity budget: Spending trends, challenges, and the future. https://www.sans.org/white-papers/2025-ics-ot-cybersecurity-budget-spending-trends-challenges-future/.

Schmidli, **H.** (2017). *Risk theory*. Springer.

Skeoch, **H. R.** (2022). Expanding the Gordon-Loeb model to cyber-insurance. *Computers and Security*, **112**, 102533.

Skeoch, **H. R.**, & **Ioannidis**, **C.** (2024). The barriers to sustainable risk transfer in the cyber-insurance market. *Journal of Cybersecurity*, **10**(1), tyae003.

Tatsumi, **K.-i.**, & **Goto**, **M.**2010). Optimal timing of information security investment: a real options approach. In **T. Moore**, **D. Pym** & **C. Ioannidis** (Eds.), *Economics of information security and privacy* (pp. 211–228). Springer.

Yuan, **S.** (1999). ODE-oriented semi-analytical methods. *Computational Mechanics in Structural Engineering*, 375–388.

Zeller, **G.**, & **Scherer**, **M.** (2022). A comprehensive model for cyber risk based on marked point processes and its application to insurance. *European Actuarial Journal*, **12**, 33–85.

Zeller, **G.**, & **Scherer**, **M.** (2023). Risk mitigation services in cyber insurance: Optimal contract design and price structure. *The Geneva Papers on Risk and Insurance - Issues and Practice*, **48**, 502–547.

Cite this article: Callegaro G, Fontana C, Hillairet C and Ongarato B (2026). A stochastic Gordon–Loeb model for optimal cybersecurity investment under clustered attacks, *Annals of Actuarial Science*, 1–29. https://doi.org/10.1017/S1748499526100359

https://doi.org/10.1017/S1748499526100359 Published online by Cambridge University Press