

Received 28 March 2023, accepted 25 April 2023, date of publication 3 May 2023, date of current version 10 May 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3272670

RESEARCH ARTICLE

Cyber Risk Assessment and Optimization: A Small Business Case Study

MARIA TSIODRA¹, SAKSHYAM PANDA², (Member, IEEE), MICHAEL CHRONOPOULOS^{3,4},
AND EMMANOUIL PANAOUSIS²

¹Business Intelligence, The Open University, MK7 6BJ Milton Keynes, U.K.

²Faculty of Engineering and Science, University of Greenwich, SE10 9LS London, U.K.

³Bayes Business School, City, University of London, EC1Y 8TZ London, U.K.

⁴Department of Business and Management Science, Norwegian School of Economics, 5045 Bergen, Norway

Corresponding author: Michael Chronopoulos (michalis.chronopoulos@city.ac.uk)

The work of Maria Tsiotra was supported by the European Commission under the H2020 cyberSecurity Platform for vRtualised 5G cybEr Range services (SPIDER) Project under Grant 833685. The work of Sakshyam Panda and Emmanouil Panaousis was supported in part by the U.K. National Cyber-Security Centre (NCSC), and in part the Research Institute for Sociotechnical Cyber Security (RISCS) under the fraMEwoRk to Model and IncenTivise Cyber Security Investment Decisions (MERIT) Project. The work of Michail Chronopoulos was supported in part by the European Commission under the H2020 SPIDER Project under Grant 833685, in part by the U.K. National Cyber-Security Centre (NCSC), and in part the Research Institute for Sociotechnical Cyber Security (RISCS) under the MERIT Project.

ABSTRACT Assessing and controlling cyber risk is the cornerstone of information security management, but also a formidable challenge for organisations due to the uncertainties associated with attacks, the resulting risk exposure, and the availability of scarce resources for investment in mitigation measures. In this paper, we propose a cybersecurity decision-support framework, called CENSOR, for optimal cyber security investment. CENSOR accounts for the serial nature of a cyber attack, the uncertainty in the time required to exploit a vulnerability, and the optimisation of mitigation measures in the presence of a limited budget. First, we evaluate the cost that an organisation incurs due to a cyber security breach that progresses in stages and derive an analytical expression for the distribution of the present value of the cost. Second, we adopt a Set Covering and a Knapsack formulation to derive and compare optimal strategies for investment in mitigation measures. Third, we validate CENSOR via a case study of a small business (SB) based on: (i) the 2020 Common Weakness Enumeration (CWE) top 25 most dangerous software weaknesses; and (ii) the Center for Internet Security (CIS) Controls. Specifically, we demonstrate how the Knapsack formulation provides solutions that are both more affordable and entail lower risk compared to those of the Set Covering formulation. Interestingly, our results confirm that investing more in cybersecurity does not necessarily lead to an analogous cyber risk reduction, which indicates that the latter decelerates beyond a certain point of security investment intensity.

INDEX TERMS Cybersecurity, operational research, set covering, knapsack, software weaknesses, control optimisation.

I. INTRODUCTION

Breakthroughs and advancements in the area of computer information systems have improved the operational efficiency of critical infrastructures but have also rendered these substantially more vulnerable to cyber attacks. The cyber risk exposure and financial consequences these attacks entail for an organisation can be demonstrated through a range of examples. Among the most recent breaches is that at Marriott that

revealed personal details of approximately 5.2 million hotel guests. Also, the breach at Twitter allowed fraudulent tweets about Bitcoin generating more than \$100, 000 Bitcoin worth, while the Solarwinds hack managed to compromise multiple government systems along with many fortune 500 companies, globally. The latter, resulted in an 8% fall in the share price of FireEye after it disclosed information about the attack,¹

The associate editor coordinating the review of this manuscript and approving it for publication was Moussa Ayyash¹.

¹<https://www.cnn.com/2020/12/08/fireeye-shares-fall-after-security-company-discloses-cyberattack.html>

and is expected to cost cyber insurers \$90 million for incident response and forensic services.²

This situation becomes more challenging for small businesses. According to the World Bank,³ small businesses represent about 90% of businesses and more than 50% of employment worldwide. With access to sensitive and personally identifiable information, UK small businesses are both a soft and lucrative target for exploitation by threat actors. While large organisations are highly likely to have resources to invest in cyber security activities, small businesses, in most cases, not only have scarce budgets and time to invest in protecting themselves but also lack basic cyber security know-how. Evidently, the Cyber Security Breaches Survey 2022,⁴ published by the Department for Digital, Culture, Media & Sport (DCMS), states that almost four in ten businesses (39%) identified a cyber-attack in the last 12 months, while most of them do not have the ability to prevent these attacks nor to undertake any incident response. The same survey identified that one-third of small businesses in the UK were attacked at least once a week over the past year and one in five have suffered losses due to cyber incidents.

These examples demonstrate how cyber security is a critical defensive manoeuvre as well as a strategic decision that may increase an organisation's competitive advantage. Furthermore, they emphasise the increasing need for developing economic models to assess cyber risk and derive insights on how to invest in measures to mitigate it. Indeed, although cyber defence is a standard part of enterprises' agenda, hackers improve upon their techniques, thereby increasing the cyber risk levels around the world. Consequently, assessing cyber risk is not just a necessary process that enterprises must conduct, but a natural way to realise the exact weaknesses, threats, and current security level of an organisation towards mitigating this risk. The existence of several tools that assess cyber risk and the vast number of papers and textbooks in this field demonstrate the importance of the cyber risk assessment domain for both industry and academia, as well as the wide variety of challenges to be addressed within this domain [1], [2], [3]. Among the challenges that organisations must tackle to improve their cyber security posture, are those of gauging the financial impact of cyber breaches and selecting the optimal set of mitigation measures.

Addressing these challenges is vital, as, for example, the General Data Protection Regulation (GDPR) poses fines up to €20 million, or, in the case of an undertaking, up to 4% of the total turnover of the preceding financial year, whichever is higher [4]. Overcoming these challenges requires the development of novel techniques that combine risk assessment and control optimisation in a way that accounts for critical aspects of the attack itself, the relevant underlying uncertainties, and constraints associated with the selection of mitigation

measures [5], [6], [7], [8]. Examples of key uncertainties associated with an attack are the *time required to exploit* a vulnerability before moving to the next one and the *extent of the associated impact* (i.e. financial cost) to the targeted organisation. Indeed, both exploitation time (also called "exploit window") and impact due to an attack are likely to vary randomly, as they depend not only on the skills of the attacker but also on the level of cyber preparedness and response of the organisation [9]. Therefore, prior to investing in mitigation measures, it is essential to carry out an in-depth cyber risk assessment to achieve an accurate evaluation of the organisation's security posture. Failing to carry out this task may result in cycles of under- or over-investment, which raises regulatory risk when corrective policy actions are required, and lead to denial of cyber insurance claims [10].

Additionally, organisations race to catch up with the ever-growing cyber threat landscape, as a wide range of threat actors demonstrate a significantly increasing range of intelligence-gathering techniques. For example, Advanced Persistent Threats (APTs) are motivated by political or economic reasons, and are origins of considerable cyber risk for the organisations they are targeting [11]. Typically, an APT breaches its targets in *phases* reflecting the process in which an adversary gradually exploits a series of system-, network-, or even user-oriented vulnerabilities [12], [13]. APTs are considered a major threat as they not only attack in multiple phases but also over a variable time period, also called dwell time. This denotes the time an adversary remains in a network undetected or the time required for their attribution [14], depending on whether we study pre- or post-incident events. The FireEye M-Trends 2020 Special report found that the mean dwell time for 2019 in the USA is 60 days and in EMEA and APAC is 54 days.⁵

In this paper, we develop a decision-support framework, called CENSOR (Cyber risk assessment and control optimisation framework), for optimal cyber security investment that accounts for the serial nature of a cyber security breach, the uncertainty in the time required to exploit a vulnerability (also called "exploit window") and the optimisation of mitigation measures in light of scarce financial resources [15]. Consequently, the contribution of our work is threefold: First, we incorporate key uncertainties into the traditional *discounted cash flow* (DCF) approach, thereby making it more suitable not only for investment decision-making but also for risk assessment and management within a cyber security context. Second, we combine the cyber risk assessment framework with two different models, based on Set Covering and Knapsack, for deriving optimal sets of measures for mitigating cyber risk. Third, we develop a use case of a small business (SB) to validate and demonstrate the application potential of our framework in cyber security investment decision-making. Specifically, we have used the 2020 Common Weakness Enumeration (CWE)

⁵<https://www.fireeye.com/current-threats/annual-threat-report/mtrends.html>

²<https://www.isaca.org/resources/news-and-trends/industry-news/2020-top-cyberattacks-of-2020-and-how-to-build-cyberresiliency>

³<https://www.worldbank.org/en/topic/smefinance>

⁴<https://www.gov.uk/government/statistics/cyber-security-breaches-survey-2022>

top 25 most dangerous software weaknesses and the Center for Internet Security (CIS) controls version 7.1. This case study is used to compare our findings to the set of guidelines for small businesses published by the UK government.⁶

CENSOR can support high-level governance and management frameworks and guidelines by providing insights and recommendations through a scientific risk assessment methodology. This paper demonstrates how mathematical models and simulations can be used to estimate the potential impact of cyber attacks and to identify risks that need attention, which further leads to effective cyber risk control strategies. CENSOR, in particular, will support decision makers to better understand their risks and prioritise the implementation of security controls based on their risk profile, budget constraint and security objectives. This approach will lead to informed risk management decisions for organisations seeking to improve their cybersecurity posture using frameworks like NIST Cyber Security Framework and ISO 27000 series.

The remainder of the paper is organised as follows. In Section II, we discuss work most relevant to this paper, and, in Section III, we introduce assumptions and notation. Next, in Section IV, we begin the analysis by deriving the analytical expression of the expected present value (PV) of the first, second and n phases of a cyber attack. This section further evaluates the selection of the best Security Packages (set of security controls) using Set Covering formulation and Knapsack Optimisation. Section V presents the use case and discusses the results obtained from our framework, while Section VI concludes this paper by offering future research directions.

II. RELATED WORK

A strand of the cyber security economics literature draws on the theory of investment and project valuation under uncertainty [16], with the main objective to derive the expected value of investment in cyber security controls along with the investment threshold price and the probability of investment within a given time horizon [17], [18]. This methodology, also known as *real options*, addresses the problem of investment under uncertainty while reflecting the value from embedded managerial discretion. For example, [19] extend the framework of [20] by showing that information-sharing regarding vulnerabilities can decrease uncertainty about risks, and, in turn, the value of deferment options. More recently, [17] develop a real options model to cast the cyber security investment problem as one of selecting a subset of uncertainty-reducing mitigation measures, whose availability is controlled by decision-makers and their size is log-normally distributed. The contribution of this work is to improve the efficiency of cyber security investments, by balancing the costs of mitigation against their incremental uncertainty-reduction impact on cyber security loss expectancy. In the same line of work, [21] analyse how uncertainty over both the cost of an attack and the arrival

of a control impacts the optimal time of investment in cyber security.

Although the optimal time of investment in cyber security controls is an important problem, especially considering the intensity and irreversibility of this capital expenditure as well as the various underlying uncertainties, the main limitation of the aforementioned literature is two-fold: First, decisions for mitigation of threats and protection of a network must be taken promptly, and, therefore, the value of waiting, which real options theory emphasises, may not be as pronounced as in other industries, e.g. research and development, and energy. Second, real options models can be used to derive the expected value of an investment opportunity along with the investment threshold price, but they do not quantify the degree to which risk is hedged. The latter problem often fits within a security planning process, in terms of optimal selection of countermeasures, yet such models are typically deterministic, as they ignore key uncertainties of cyber attacks.

Examples of models for optimal selection of cyber security controls include [6], who investigate the challenge of how to spend a security budget optimally. They propose methods, such as optimisation algorithms, combinatorial optimisation and the classical Knapsack problem, that can deal with overlapping safeguards that exhibit non-linear relationships. Similarly, [9] propose a methodology for investing in CIS controls, considering a single value for a vulnerability and a number of implementation levels for each control. The latter, represent the information security levels proposed in the seminal work on the economics of information security [22]. Also, [7] extend the methodology proposed in [9] to obtain an optimal set of controls to protect various employee groups of a healthcare organisation from social engineering attacks. Additionally, [23] cast the problem of optimal selection of controls as a Set Covering problem. They first solve a deterministic version to analyse the incentive to implement complementary mitigations to reduce supply chain vulnerabilities, and then extend the deterministic version to allow for limitations on the choice as well as uncertainty over the efficacy of the different controls. Building upon [24], [25] develop a game-theoretic framework to analyse defender-attacker interactions. The defender chooses a security plan to minimise its security risk, while the attacker aims to maximise it via the most effective attack path. This is modelled as a min-max optimisation problem, where the attacker maximises and the defender minimises in response to the reaction of the attacker.

A limitation of the aforementioned optimisation models is that they overlook the serial nature of an attack as well as the uncertainty over the time it takes to exploit a vulnerability and the cost that the system incurs once a vulnerability is compromised. Consequently, the financial implications of these uncertainties on an organisation's assets remain an important open research question. To assist in the anticipation and control of the financial impact of attacks [26], [27], our work builds upon the literature on the valuation of serial projects to develop quantification tools for assessing the risk associated

⁶<https://www.ncsc.gov.uk/collection/small-business-guide>

with a security breach that progresses in phases. Specifically, we first derive the distribution of the PV of the impact of an attack and then develop methods for minimising the expected PV of the impact by optimising the selection of controls. Via a case study, where we cast the optimal selection of controls as a Knapsack and a Set Covering problem, we find that the Knapsack formulation provides solutions that not only are more affordable but also entail lower risk. Also, we find that greater investment intensity does not necessarily result in an analogous reduction of risk, which, in turn, implies that the rate of risk reduction decreases beyond a certain level of investment intensity.

III. CYBER RISK ASSESSMENT AND CONTROL OPTIMISATION FRAMEWORK

In this section, we model cyber risk by adopting a techno-economic approach that couples capital budgeting for valuation of the financial impact of a serial cyber security breach with optimisation of mitigation measures. In summary, this section discusses: i. the underlying system model with an organisation (e.g. Small and Medium-sized Enterprise) that wishes to protect its systems (Defender) and hackers who target the organisation (Attacker); ii. how the different system assets that host vulnerabilities (e.g. CWE) are linked to each other, and, as a result, how a multi-phase attack can sequentially compromise these assets causing damage to the Defender; and iii. how a limited budget poses the challenge of optimally allocating cyber security controls that may overlap in terms of the vulnerabilities they patch.

A. SYSTEM MODEL

We assume that the Defender's infrastructure consists of a number of systems and networks, referred to as assets, that the Defender aims to protect from the Attacker. Each asset $i \in \mathbb{N}$ has a set of $m_i \in \mathbb{N}$ vulnerabilities, i.e. $\mathcal{V}_i = \{v_{i1}, v_{i2}, \dots, v_{im_i}\}$, that the Attacker may exploit. These vulnerabilities are part of software weaknesses, as presented in the CWE.⁷ This assumption is aligned with the real-world behaviour of attackers, who aim to penetrate as deep in a network as they can to increase their expected return from the attack. These adversarial interactions are modelled as a sequence of attack phases, where phase i of an attack refers to the stage in which the Attacker aims to compromise asset i by exploiting any of its vulnerabilities $v_{ij} \in \mathcal{V}_i$, where $j = 1, 2, \dots, m_i$. We assume that in each phase the Attacker can compromise only one asset and that successful exploitation can lead to undesirable privilege escalation or lateral movement within the Defender's infrastructure [28], [29], which presents a new set of vulnerabilities that the Attacker may exploit in order to compromise the subsequent asset.

B. EXPECTED IMPACT

The expected impact from the exploitation of asset i is denoted by K_i . Utilising the broadly accepted risk

assessment formula [30], expected impact = (likelihood of being attacked) \times (probability of being compromised) \times (probable loss), we compute K_i as in (1). This expresses the impact posed to the Defender with regards to the exploitation of asset i . In (1): A_i is the value of asset i , also known as Single Loss Expectancy (SLE) [31]; $R_i = (r_{i1}, r_{i2}, \dots, r_{im_i})$, where r_{ij} is the likelihood of the Attacker attempting to exploit vulnerability v_{ij} , which expresses the degree of attractiveness of a vulnerability to the Attacker and is also referred to as the Annual Rate of Occurrence (ARO) [31]; and $S_i = (s_{i1}, s_{i2}, \dots, s_{im_i})$, where s_{ij} is the probability of the same vulnerability being successfully breached, thereby reflecting the current security level associated with the vulnerability. We denote the likelihood of occurrence of an attack against asset i as, $\langle R_i, S_i \rangle$, i.e. the inner product between R_i and S_i .

$$K_i = A_i \cdot \langle R_i, S_i \rangle. \quad (1)$$

T_i is the time required to exploit asset i and follows a general distribution function denoted by $\Psi_{T_i}(\cdot)$, as shown in Figure 1. Assuming that a successful attack consists of a number of phases, each of them compromising an asset, we compute the aggregate duration of the attack, W_k , as the sum of the exploitation times required to compromise an asset in each phase, i.e. $W_k = \sum_{i=1}^k T_i$, $1 \leq k \leq n$.

To determine the distribution of the PV of the impact associated with the attack, we first express the aggregate impact, Z_n , over n attack phases in (2)

$$\begin{aligned} Z_n &= \underbrace{K_1 e^{-\rho W_1}}_{U_1} + \underbrace{K_2 e^{-\rho W_2}}_{U_2} + \dots + \underbrace{K_n e^{-\rho W_n}}_{U_n} \\ &= \sum_{i=1}^n U_i, \end{aligned} \quad (2)$$

where U_i is the PV of K_i and ρ denotes the subjective discount rate. The PV integrates the concept of discounting into the calculation of the current value of the impact of an attack that may require a substantial amount of time to be carried out. In turn, this supports effective decision-making and facilitates the development of risk measures to assess the financial risk exposure of the Defender [32].

C. CYBER RISK CONTROL

After gauging the impact associated with each vulnerability, CENSOR subsequently focuses on optimising the coverage of vulnerabilities in each asset by determining an optimal Security Package. The latter refers to the set of controls that minimise the expected PV of the impact from an attack. This is done by patching asset vulnerabilities, thereby reducing an asset's attack surface or by increasing the effort required in successfully breaching the asset. CENSOR specifically considers that the implementation of a control will mitigate the expected impact of an attack by reducing the probability of the latter being successful. We denote by $\mathcal{C} = \{C_1, C_2, \dots, C_g\}$ the set of available controls and by E_{ijl} the efficacy of control C_l , $l = 1, 2, \dots, g$, against vulnerability v_{ij} ,

⁷<https://cwe.mitre.org/index.html>




FIGURE 1. Sequential security breach.

implemented at level L_ℓ , $\ell = 1, 2, \dots, h$. Intuitively, $E_{ij\ell}$ reflects the degrees of protection offered by control C_l for a vulnerability v_{ij} .

As the implementation of cyber security controls is not cost free, the associated *direct* and *indirect* costs must be considered by the Defender. According to [33], the former refers to the acquisition, deployment, and maintenance costs of a control, while the latter can be anything else that inflicts loss to the Defender, such as slowing down essential processes due to incompatibility of controls and training employees to get acquainted with the new controls. For the sake of brevity, we assume that each control C_l comes with a set of costs Ξ_l associated with each level of implementation, inclusive of the direct and indirect costs. Table 1 presents a summary of the notation used in this paper.

IV. ANALYSIS

A. IMPACT ASSESSMENT

The first functionality of CENSOR is to assess the PV of the impact from an attack. To achieve this, we perform a phase-wise analysis of an attack to compute: (i) the distribution and expectation of the PV of K_i at each phase $i = 1, 2, \dots, n$; and (ii) the expected impact following the multi-staged attack.

- 1) *First Phase*: We begin with the first of an n -phase attack that starts at time 0 and stops at time $w_1 \equiv \tau_1$, which is a realisation of the random variable $W_1 \equiv T_1$. The PV of K_1 , denoted by u_1 , is described in (3), and the effect of discounting can be interpreted as the reduction in the impact for the Defender due to the exploit window.

$$u_1 = K_1 e^{-\rho \tau_1} \quad (3)$$

Consequently, the cumulative distribution function (CDF), $\Theta_1(\cdot)$, and probability density function (PDF), $\theta_1(\cdot)$, of U_1 are given in (4) and (5), respectively (all proofs can be found in the Appendix).

$$\Theta_{U_1}(u_1) = 1 - \Phi_1 \left(\ln \left(\frac{K_1}{u_1} \right) \rho^{-1} \right) \quad (4)$$

$$\theta_{U_1}(u_1) = \frac{\phi_1 \left(\ln \left(\frac{K_1}{u_1} \right) \rho^{-1} \right)}{\rho u_1} \quad (5)$$

While (4) and (5) assume a generic distribution for T_1 , the CDF and PDF of U_1 when $T_1 \sim \exp(\lambda_1)$ is

described in (6) and (7), respectively.

$$\Theta_{U_1}(u_1) = \left(\frac{u_1}{K_1} \right)^{\frac{\lambda_1}{\rho}} \quad (6)$$

$$\theta_{U_1}(u_1) = \frac{\lambda}{\rho} \left(\frac{u_1}{K_1} \right)^{\frac{\lambda_1}{\rho} - 1} \frac{1}{K_1} \quad (7)$$

Having derived the analytical expression of the CDF and PDF of U_1 , we can derive the main moments of the distribution. Specifically, the expression of the expectation and variance is $\mu_1 = \frac{\lambda_1}{\lambda_1 + \rho} K_1$ and $\sigma_1^2 = \left[\frac{\lambda_1}{\lambda_1 + 2\rho} - \left(\frac{\lambda_1}{\lambda_1 + \rho} \right)^2 \right] K_1^2$, respectively.

- 2) *Second Phase*: We next derive the distribution of the PV of K_2 . Therefore, we begin with the distribution of $W_2 = T_1 + T_2$ and consider the case in which $T_1 \sim \exp(\lambda_1)$ and $T_2 \sim \exp(\lambda_2)$,⁸ so that W_2 follows a hypo-exponential distribution, i.e. $W_2 \sim \text{Hypo}(\lambda_1, \lambda_2)$. Then, the distribution of the PV of K_2 is shown in Proposition 1.

Proposition 1: If $W_2 \sim \text{Hypo}(\lambda_1, \lambda_2)$, then the CDF and PDF of U_2 is described in (8) and (9)

$$\Theta_{U_2}(u_2) = \frac{\lambda_2}{\lambda_2 - \lambda_1} \left(\frac{u_2}{K_2} \right)^{\frac{\lambda_1}{\rho}} - \frac{\lambda_1}{\lambda_2 - \lambda_1} \left(\frac{u_2}{K_2} \right)^{\frac{\lambda_2}{\rho}} \quad (8)$$

$$\theta_{U_2}(u_2) = \frac{\lambda_1 \lambda_2}{\lambda_2 - \lambda_1} \frac{1}{\rho u_2} \left[\left(\frac{u_2}{K_2} \right)^{\frac{\lambda_1}{\rho}} - \left(\frac{u_2}{K_2} \right)^{\frac{\lambda_2}{\rho}} \right]. \quad (9)$$

We can also derive the main moment of the distribution of U_2 , e.g., the mean and the variance are $\mu_2 = \frac{\lambda_1 \lambda_2}{(\lambda_1 + \rho)(\lambda_2 + \rho)} K_2$ and $\sigma_2^2 = \frac{\lambda_1 \lambda_2}{(\lambda_1 + 2\rho)(\lambda_2 + 2\rho)} K_2^2 - \mu_2^2 K_2^2$, respectively.

- 3) *n-th Phase*: Here, we will determine the distribution of the PV of the impact of the attack in the arbitrary n phase. Following the same approach as in the case of $i = 1, 2$, we assume that $W_n \sim \text{Hypo}(\lambda_1, \lambda_2, \dots, \lambda_n)$. The CDF and PDF of U_n is described in Proposition 2.

⁸To emphasise the implications of differences in asset exploitation, we may consider the special case where $\lambda_1 = \lambda_2 = \lambda$. Under this assumption, T_1 and T_2 are i.i.d. random variables, such that $W_2 = T_1 + T_2 \sim \text{Erlang}(2, \lambda)$ ([34]).

TABLE 1. List of symbols.

Symbol	Description
i	Phase of attack or an asset ($i = 1, 2, \dots, n$)
\mathcal{V}_i	Set of vulnerabilities in asset i , $\mathcal{V}_i = \{v_{i1}, v_{i2}, \dots, v_{im_i}\}$
\mathcal{C}	Set of cyber security controls, $\mathcal{C} = \{C_1, C_2, \dots, C_g\}$
\mathcal{L}_l	Set of levels of control C_l , $\mathcal{L}_l = \{L_1, L_2, \dots, L_h\}$, where $l = 1, 2, \dots, g$
Ξ_l	set of cost of each level of control C_l , $\Xi_l = \{\xi_{l1}, \xi_{l2}, \dots, \xi_{lh}\}$
\mathcal{E}_{ijl}	Set of efficacy of each level of control C_l on vulnerability v_{ij} , $\mathcal{E}_{ijl} = \{E_{ijl1}, E_{ijl2}, \dots, E_{ijlh}\}$
v_{ij}	Vulnerability within asset i ($j = 1, 2, \dots, m_i$)
r_{ij}	Probability of vulnerability v_{ij} being targeted (attack occurrence)
s_{ij}	Probability of vulnerability v_{ij} being compromised when attacked (success rate)
$x_{l\ell}$	Indicates whether level L_ℓ , $\ell = 1, 2, \dots, h$, of control C_l is selected
y_{jl}	Indicates whether vulnerability v_{ij} is covered by C_l
A_i	Value of asset i
T_i	Time required to exploit a weakness in asset i
λ_i	Rate parameter for attack phase i
W_k	Total duration of the attack until phase k , $\sum_{i=1}^k T_i$ where $1 \leq k \leq n$
ρ	Discount factor
K_i	Impact from exploiting asset i
U_i	PV of the impact for attack phase i
Z_i	Expected impact of attack phase i
Z_n	Aggregated expected impact for the first k attack phases, where $1 \leq k \leq n$

Proposition 2: If $W_n \sim \text{Hypo}(\lambda_1, \lambda_2, \dots, \lambda_n)$, then the CDF and PDF of U_n is described in (10) and (11)

$$\Theta_{U_n}(u_n) = \sum_{i=1}^n \left(\frac{u_n}{K_n}\right)^{\frac{\lambda_i}{\rho}} \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \quad (10)$$

$$\theta_{U_n}(u_n) = \sum_{i=1}^n \frac{1}{\rho K_n} \left(\frac{u_n}{K_n}\right)^{\frac{\lambda_i}{\rho} - 1} \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \quad (11)$$

The mean and variance of U_n is described in (12) and (13), respectively,

$$\mu_n = K_n \sum_{i=1}^n \frac{\lambda_i}{\lambda_i + \rho} \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \quad (12)$$

$$\sigma_n^2 = K_n^2 \sum_{i=1}^n \frac{\lambda_i}{\lambda_i + 2\rho} \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} - \left[K_n \sum_{i=1}^n \frac{\lambda_i}{\lambda_i + \rho} \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \right]^2, \quad (13)$$

and, since we have already derived the distribution of the PV of the impact for each phase, the expected value of $Z_n = \sum_{i=1}^n U_i$ is given in (14).

$$\begin{aligned} \mathbb{E}[Z_n] &= \sum_{i=1}^n \mathbb{E}[U_i] \\ &= \sum_{b=1}^n K_b \sum_{i=1}^b \frac{\lambda_i}{\lambda_i + \rho} \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \end{aligned} \quad (14)$$

B. SECURITY CONTROLS

The optimisation functionality of CENSOR aims to provide decision support regarding the budget constrained selection of controls. This includes: (i) the minimum number of controls required to patch all vulnerabilities (Approach 1); or (ii) the controls that minimise the expected PV of the impact (Approach 2).

Regarding Approach 1, the Security Package (i.e. combination of security controls) is derived as a solution to a Set Covering problem in order to account for the potential interaction among Security Packages that may overlap in terms of the vulnerabilities they cover.

Regarding Approach 2, we use the risk assessment component of CENSOR along with the improvement of the expected impact when deploying Security Packages. We then derive the Security Package that maximises the improvement of the expected PV of the impact subject to a financial budget. The same function takes into account the cost of each Security Package, which is subtracted by the improvement. Thus, CENSOR optimises the return on security investment, which is defined in [33] as the ratio (benefit of security - the cost of security) / cost of security.

1) SET COVERING FORMULATION

Here, we cast the problem of optimal selection of mitigation measures as a Set Covering problem [35]. The objective is to determine the *minimum number of controls* that offer a baseline coverage of the network's vulnerabilities based on a budget constraint and the desired level of efficacy resulting from each patch. This optimisation is described in (15), where

$x_{l\ell}$ is a binary variable indicating whether a specific level ℓ of a security control C_l is applied. Constraint (16) ensures that each vulnerability is covered by at least one control, while (17) is the budget constraint. Finally, constraint (18) ensures that the choice of controls offers a minimum coverage of all vulnerabilities at a desired level of efficacy, \hat{e} .

$$\min \sum_{l=1}^g x_{l\ell} \quad (15)$$

$$s.t. \sum_{l:E_{ijl\ell}>0} x_{l\ell} \geq 1, \quad x_{l\ell} \in \{0, 1\}, \quad \forall i, j \in \mathbb{N} \quad (16)$$

$$\sum_{l=1}^g x_{l\ell} \xi_{l\ell} \leq B \quad (17)$$

$$E_{ijl\ell} > \hat{e}, \quad \forall i, j \in \mathbb{N}, \quad \hat{e} \in (0, 1) \quad (18)$$

Algorithm 1 Set Covering Problem With Cost and Control Efficacy Constraints

Input: $\mathcal{V}_i, \mathcal{C}, \Xi_l, \mathcal{E}_{ijl}, \hat{e}$

Output: Minimum set of controls with budget and efficacy bound

Function SetCoverEfficacy ($\mathcal{V}_i, \mathcal{C}, \Xi_l, \mathcal{E}_{ijl}, \hat{e}$):

```

for  $C_l$  in  $\mathcal{C}$  do
  price  $\leftarrow \xi_{l\ell} / \text{len}(C_l \cap \mathcal{V})$ 
  if price < cost and  $E_{ijl\ell} > \hat{e}$  then
    cost  $\leftarrow$  price
    cover  $\leftarrow C_l$ 
return (cover, cost)

```

```

while  $\text{len}(\mathcal{V}_i) \neq 0$  and Budget  $\neq 0$  do
  (cover, cost)  $\leftarrow$ 
  SetCoverEfficacy( $\mathcal{V}_i, \mathcal{C}, \Xi_l$ )
   $\mathcal{V}_i \leftarrow \mathcal{V}_i - \text{cover}$ 
  Budget  $\leftarrow$  Budget - cost

```

In its basic implementation, the Set Covering problem is appropriate when the underlying controls are related to patching vulnerabilities, as the degree of their effectiveness is 1, meaning that either the vulnerability is patched or not. However, this implementation does not account for the expected impact of the security breach. Indeed, the solution obtained via (15)-(18) ensures the minimisation of the number of controls, but does not consider whether the proposed controls minimise the expected impact of the attack. Furthermore, when we study preventative controls (e.g. firewalls), we must take into account the degree of a control's efficacy against a vulnerability. In a quantitative approach, the control efficacy falls within the interval (0, 1). Algorithm 1 presents the Set Covering implementation with cost and control efficacy constraints.

2) KNAPSACK FORMULATION

The challenge of optimal budget allocation in cyber security can be addressed through combinatorial optimisation [9].

Unlike Section IV-B1, the objective here is to select the controls that minimise the expected PV of the impact, as expressed in (14). Based on the selection of a control at a particular level, indicated through $x_{l\ell} \in \{0, 1\}$, the probability of exposure can be expressed as $\varepsilon_j = \prod_{l,\ell} (1 - x_{l\ell} E_{ijl\ell} \mathbb{1}_{\{v_{ij} \in C_l\}})$. Notice how ε_j is a strictly decreasing function of $x_{l\ell}$, so that the inclusion of a control will reduce the likelihood of exposure. Thus, CENSOR solves the following Knapsack problem:

$$\min_{\substack{x_{l\ell} \\ l=1,2,\dots,g \\ \ell=1,2,\dots,h}} \sum_{i=1}^n \left\{ \left\{ \prod_{l=1}^g \prod_{\ell=1}^h (1 - x_{l\ell} E_{ijl\ell}) \right\} \cdot \mathbb{E}[Z_i] \right\}, \quad (19)$$

$$\forall i, j \in \mathbb{N}$$

$$s.t. \sum_{l=1}^g \sum_{\ell=1}^h x_{l\ell} \xi_{l\ell} \leq B \quad (20)$$

$$\sum_{\ell=1}^h x_{l\ell} = 1, \quad x_{l\ell} \in \{0, 1\}, \quad \forall l = 1, \dots, g, \quad (21)$$

where $\mathbb{E}[Z_i]$ is a dependent variable obtained from eq (14). The optimal efficacy matrix O is constructed iteratively for all the cost values within the Budget, and for each value the problem is solved considering all the available levels of a control within that cost. The optimal aggregated efficacy value $O[l, cost]$ depends on the control level selected for the l -th cost. For a detailed analysis of 0-1 Knapsack Optimisation using dynamic programming refer to [36].

Algorithm 2 Dynamic Programming Based 0-1 Knapsack Optimisation

Input: $\mathcal{V}_i, \mathcal{C}, \mathcal{L}_l, \Xi_l, \mathcal{E}_{ijl}$

Output: Optimal set of controls and total cost

Function

KnapsackOptimisation ($\mathcal{V}_i, \mathcal{C}, \mathcal{L}_l, \Xi_l, \mathcal{E}_{ijl}$):

```

for  $C_l$  in  $\mathcal{C}$  do
  for cost in Budget do
     $O[C_l, cost] \leftarrow O[C_l - 1, cost]$ 
    for  $\ell$  in  $\mathcal{L}_l$  do
      if cost  $\geq \xi_{l\ell}$  then
         $O[C_l, cost] \leftarrow$ 
        max{ $O[C_l, cost], (O[C_l - 1, cost - \xi_{l\ell}] + (1 - E_{ijl\ell}))$ }

```

V. USE CASE

This section presents our case study to implement our methodology using a sample network topology of an SB. Let the network architecture consists of three layers. Each layer of the architecture signifies the importance of the asset to the organisation. A weakness can lead to a number of vulnerabilities that the Attacker could exploit to compromise an asset in a layer and move to the next layer.

As each organisation is likely to have different assets and value each one differently, defining the direct financial value of an asset to the organisation is a challenge. To tackle this, we assume that all assets in a layer have the same value. Thus, based on the location of an asset in the architecture we define the value of an asset, A_i , as a categorical value, where $A_1 = 500$, $A_2 = 1000$ and $A_3 = 1500$ in the three layers, respectively. Additionally, we have used $\rho = 0.3$, while the cost and efficacy for a control is determined using random uniform distribution. For the purpose of this case study, we consider the 2020 top 25 CWEs⁹ and rank them based on their easiness to find and exploit. We obtain the CWEs from the Common Vulnerabilities and Exposures (CVE) data found within the National Institute of Standards and Technology (NIST) National Vulnerability Database (NVD)¹⁰ for the years 2018 and 2019. The CVE data includes a description, Common Vulnerability Scoring System (CVSS) base scores, vulnerable product configuration, and weakness categorisation information on each vulnerability identified during a specific year.

We primarily utilise the CVSS metrics to acquire parametric values required for CENSOR. CVSS is a publicly available framework that details the characteristics and severity of software vulnerabilities and is built upon three core metric groups: Base, Temporal, and Environment. The *Base metric* represents the intrinsic qualities of a vulnerability that remain unchanged over time and across user environments. The *Temporal metric* reflects the characteristics of a vulnerability that can change over time, while the *Environmental metric* reflects qualities of a vulnerability that are unique to a user's environment. This case study uses the Base metrics to extract the CENSOR parameters, but the rest of the metrics could be potentially used by experts to personalise the values to fuse specific characteristics of CVEs. Table 6 presents the top 25 CWEs for the year 2020 with their number of associated CVEs, CVSS v3 metrics,¹¹ and the Time required to exploit a CWE used in this case study, where:

Frequency: It projects the number of times a CWE is the root cause of a vulnerability. The NVD data for the years 2018 and 2019 consists of 31501 CVEs (excluding ones without a CVSS score) that are associated with one or more CWEs. To determine the frequency of a CWE, we calculate the number of times a CWE is mapped to a CVE in the NVD. We then filter the records based on the top 25 CWEs of the year 2020.

Average CVSS: It details the severity of a CWE, and is represented by the average CVSS base score of all the CVEs that map to a particular CWE.

Average Exploitability Score: It details the ease and the technical means by which the weakness can be exploited by an attacker, and is the normalised average of the CVSS Exploitability metric of all CVEs that maps to a particular

CWE. For example, the average exploitability score of all the 3848 vulnerabilities associated with CWE-79 (Rank 1, Tabl. 6) is 2.520, which when normalised gives a score of 0.433. For normalisation, we have used $Normalised_x = \frac{x - x_{min}}{x_{max} - x_{min}}$ where x is the entity being normalised. We interpret the Exploitability metric with the probability of a vulnerability being targeted by the attacker (r_{ij}).

Average Complexity Score: It details the conditions beyond the Attacker's control that must exist to exploit a vulnerability. This may include conditions that require the collection of additional information about the target and/or vulnerability, or additional resources, such as computational power. We assign a categorical value of 1 for "LOW" and 2 for "HIGH" attack complexity. The average complexity value is the normalised average of the CVSS complexity metric of all the CVEs that maps to a particular CWE. Similar to the average exploitability score, the average complexity score of all the 3848 vulnerabilities associated with CWE-79 is 1.004, which, when normalised, gives a score of 0.004. We interpret the Complexity metric with the probability of a vulnerability being compromised on an attack, s_{ij} .

Average Privileges Required: It details the level of privileges an attack requires to exploit the vulnerability, successfully. We assign a categorical value of {1, 2, 3} unit of time for {NONE, LOW, HIGH} privileges required to exploit a CVE, respectively. It is then calculated as the average of the CVSS Privileges Required metric of all the CVEs that map to a particular CWE.

Average User Interaction: It details whether an additional human user's participation, apart from the attacker, is required to successfully exploit the vulnerability. We associate a categorical value of 0 for "NONE" and 1 unit of time for "REQUIRED" user interaction. It is then calculated as the average of the CVSS User Interaction metric of all the CVEs that maps to a particular CWE.

Time: It refers to the time required to exploit asset i . In our simulations, each layer is associated with a set of unique CWEs and each of them has a number of CVEs. We have computed the average time required to exploit each CWE by using CVSS metrics for all the CVEs under a CWE. A CVE requiring higher privileges and higher user interaction will demand more time from the attacker. Thus, we define, for each CWE the sum of Average Privileges Required and Average User Interaction metrics required to exploit a CVE in asset i , and assume that $\mathbb{E}[T_i] = \frac{1}{\lambda_i}$, where λ_i is the rate parameter which is considered as inversely proportional to the average exploitation time for attack phase i .

Next, we develop a knowledge base between the 2020 top CWEs and the CIS critical security controls version 7.1.¹² The CIS controls are a set of prioritised, globally recognised, and supported security actions that organisations can take to assess and improve their cyber security. For example, control 9.4 specifies the use of host-based firewalls or port filtering tools on end systems to manage communication on networked

⁹https://cwe.mitre.org/top25/archive/2020/2020_cwe_top25.html

¹⁰<https://nvd.nist.gov/vuln/data-feeds>

¹¹<https://www.first.org/cvss/v3.1/specification-document>

¹²<https://www.cisecurity.org/controls/>

TABLE 2. Case study control family and controls distribution.

No.	Control family	Controls	No.	Control family	Controls
1	Inventory and control of hardware assets	1	2	Inventory and control of software assets	3
4	Control use of administrative privileges	2	5	Secure configuration	1
6	Maintenance, monitoring and analysis of audit logs	1	9	Limitation and controls of network ports, protocols and services	1
8	Malware defences	2	7	Email and web browser protection	2
10	Data recovery capabilities	3	11	Secure configuration for network devices	1
12	Boundary defences	2	13	Data protection	2
14	Controlled access based on the need to know	1	17	Implement security awareness and training program	3
16	Account monitoring and control	1	15	Wireless access control	2

devices, while control 17.6 mentions that organisations need to train the workforce on how to identify different forms of social engineering attacks, such as phishing attack, phone scams, and impersonation calls/emails. Table 2 presents the control family and number of controls per family from the CIS controls list that are used in this case study.

The knowledge base defines a binary relation between these basic cyber hygiene controls that every enterprise should apply to protect against the most common attacks and the CWEs. The mapping is shaped considering the scope of impact and possible mitigation detailed against each CWE. The scope of impact reflects the application security areas that is violated if an attacker successfully exploits this weakness. For example, an attacker can impact the Confidentiality, Integrity, Availability and Access Control mechanisms by exploiting CWE-79.¹³ Environment hardening controls such as Firewalls, implementing secure configurations, and installing the latest stable version of security updates can be used to mitigate the impact and reduce the likelihood of an attack. A control can provide a variety of types of mitigation, and the distribution used in this use case can be seen in Table 3. The definition of control efficacy against a vulnerability detailed in CVSS does not support CENSOR for optimal cyber security spending. Thus, we have redesigned the efficacy of a control against each CWE and spread the efficacy between the levels.

As the same weakness can appear at different assets of a layer, we assume that the implementation of a control mitigates all occurrences of that weakness. Otherwise, the security of a network would not increase as it is as strong as the weakest link in the network. For a control, we have assumed two different implementation levels: Level L and Level H. Level L refers to the minimal configuration/integration of a safeguard with lower cost and lower efficacy against a weakness. On the contrary, Level H refers to a better configuration/integration of a safeguard, which, generally, is expensive but provides higher protection against attacks. The cost of a control is considered to be the most important factor guiding the cyber security strategy of an organisation. Unlike [7] and [9], where authors have separately considered the direct cost and indirect cost of a control, this case study considers

the aggregated cost of a control. Our consideration acknowledges the fact that the cost of implementation of a control can vary across organisations. For example, integration of security software for a skilled employee might require less effort compared to a trainee. Furthermore, the indirect costs such as training and maintenance associated with a control could also be different across organisational teams. The CIS controls mapping, the CENSOR code, along with the rest of the data and analysis results can be accessed from the Github repository.¹⁴

A. RESULTS ANALYSIS

In this section, we present the results obtained from applying CENSOR to determine the optimal Security Package and compare it against the set of cyber security guidelines published by the UK's National Cyber security Centre (NCSC) for SBs to improve their cyber security and protect themselves from common attacks.¹⁵ The controls considered in the guidelines focus on malware protection, secure mobile devices, data backup, passwords to protect data and to avoid phishing attacks. The results are obtained using a number of testing conditions that represent different cyber security investment strategies. We begin with a demonstration of the risk assessment functionality of CENSOR as this is formulated in Propositions 1 and 2. Specifically, Fig. 2 illustrates the distribution of the PV of the impact for each phase, and each panel illustrates the analytical and simulated PDF of the PV as well as the 95-th percentile, indicating the worst-case scenario.

The first set of tests presents scenarios that account for the requirement of having controls that mitigate all 25 CWEs. The methodology proposed in Section IV-B1 is used to identify the minimum set of controls required to address all 25 CWEs under different budget and control efficacy bounds. The second set of tests utilises the Knapsack optimisation methodology presented in Section IV-B2 to identify a set of controls that minimises the PV of the impact of an attack given a budget limit.

Case A shown in Table 4 presents the scenario, where any of the controls selected must be implemented at the

¹³<https://cwe.mitre.org/data/definitions/79.html>

¹⁴<https://github.com/SakshyamPanda/CENSOR>

¹⁵<https://www.ncsc.gov.uk/collection/small-business-guide>

TABLE 3. Case study control and mitigation distribution.

Mitigation	Control	Mitigation	Control
Access control	[4.3, 14.6, 16.8]	Firewall	[9.4]
Security software	[2.1, 2.2, 7.1, 11.4]	Input validation	[6.2, 17.5]
Encryption	[10.4, 13.6, 15.7]	Boundary defence	[7.7, 12.1, 12.4, 15.10]
Data protection	[4.2, 10.1, 10.2, 13.2, 17.7, 17.8]	Authorised device	[1.6, 2.6]
Secure configuration	[5.1, 8.4, 8.5]	-	-

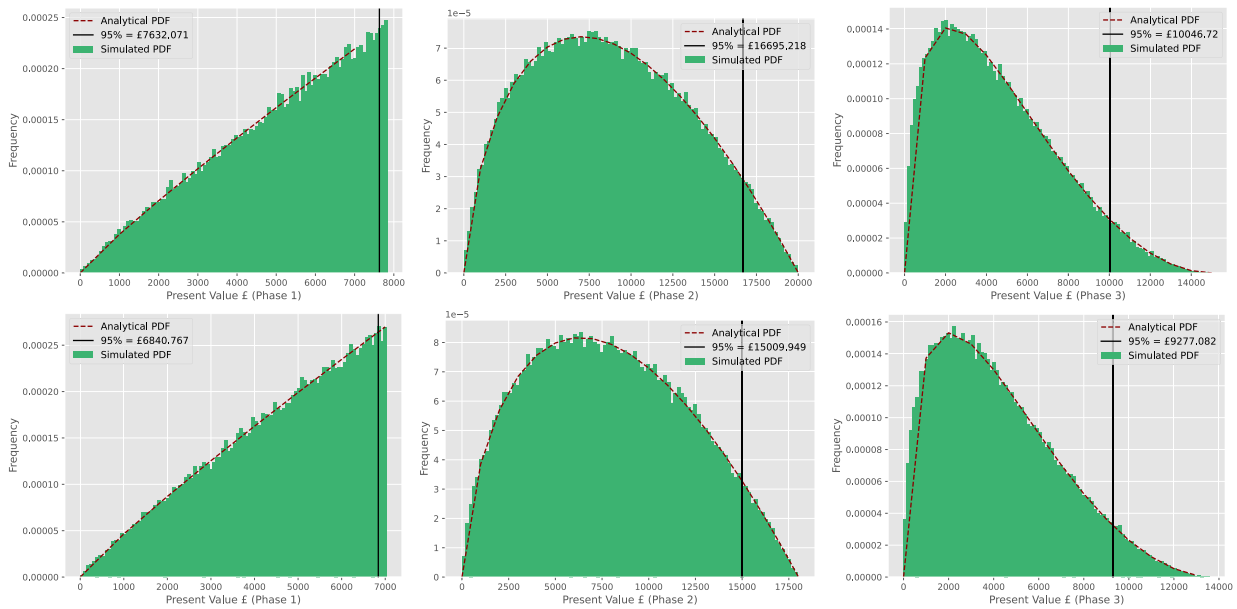


FIGURE 2. Distribution of present value of the expected impact for attack phase 1 (left panel), 2 (middle panel) and 3 (right panel) in case A (top panels) and in case F (bottom panels).

TABLE 4. Case study results with set covering formulation.

Case	Cost	Efficacy	Solution	Reduced Expected Impact	95 th Percentile		
					Attack Phase 1	Attack Phase 2	Attack Phase 3
A	£7,096.46	-	[1.6, 2.1, 4.2, 4.3, 5.1, 6.2, 7.7, 9.4, 10.4]	£42,858.84	£7,632.07	£16,695.22	£10,046.72
B	£4,225.88	-	[2.6, 8.5, 10.4, 11.4, 16.8, 17.5, 17.8]	£44,009.73	£7,792.03	£17,161.33	£10,379.47
C	£5,002.9	-	[2.2, 2.6, 8.5, 10.4, 16.8, 17.5, 17.7]	£40,607.89	£7,178.48	£15,606.24	£9,696.81
D	£4,225.88	10%	[2.6, 8.5, 10.4, 11.4, 16.8, 17.5, 17.8]	£44,009.73	£7,790.96	£17,132.12	£10,361.09
E	£5,002.9	30%	[2.2, 2.6, 8.5, 10.4, 16.8, 17.5, 17.7]	£40,607.89	£7,175.01	£15,645.01	£9,681.88

highest level, while the budget considers the highest level of investment. The Set Covering solution suggests implementing controls associated with all nine mitigations listed in Tabl. 3. In addition to the mitigation controls, the solution suggests four out of five focus areas outlined in the SB cyber security guide. Even though the solution is the most expensive, it does not suggest the implementation of controls for Security Awareness and Training, which predominantly mitigates social engineering attacks.

When a budget constraint is added to support cost-effective decisions, we see that the solution in cases B and C for controls with levels L and H, respectively, cover four out of five focus areas. However, these solutions include controls for

Security Awareness and Training to cover phishing attacks instead of controls for Administrative Privileges, which means that awareness and training of staff must be given a higher priority. Also, when the budget is insufficient for case C, the advanced software asset inventory management tool could potentially be replaced with a standard software asset management tool together with training staff on causes of unintentional data exposure. Similarly, cases D and E present alternative solutions with an option to select controls that are effective against all CWEs by at least a given threshold. Here, case D presents controls with level L that are at least 10% effective against the CWEs, while case E presents controls with level H with the efficacy of at least 30%.

TABLE 5. Case study results with Knapsack formulation.

Case	Budget	Cost	Solution	Reduced Expected Impact	95 th Percentile		
					Attack Phase 1	Attack Phase 2	Attack Phase 3
F	£5,100	£4,848.31	[6.2, 7.1, 8.5, 11.4, 15.10, 17.8]	£38,857.47	£6,840.76	£15,009.95	£9,277.08
F1	£6,200	£5,938.6	[1.6, 2.6, 7.1, 8.5, 11.4, 15.10, 17.8]	£40,078.36	£7,210.97	£15,432.78	£9,500.55
F2	£6,900	£6,807.71	[7.7(L), 8.5, 9.4, 10.1, 10.2, 10.4, 15.10, 17.8]	£43,360.01	£7,713.73	£16,831.24	£10,208.53
F3	£7,400	£7,371.64	[4.2(L), 8.5, 9.4, 10.1, 10.2, 10.4, 11.4, 15.10, 17.5, 17.8]	£43,567.46	£7,699.04	£16,978.47	£10,222.35



(A) Set cover with no constraint. (B) Set cover with budget constraint for subcontrols level L. (C) Set cover with budget constraint for subcontrols level H. (D) Set cover with budget and efficacy bound for subcontrols level L. (E) Set cover with budget and efficacy bound for subcontrols level H. (F) Knapsack Optimisation with budget

FIGURE 3. Control section using Knapsack optimisation with budget £5100.

As investing in cyber security is not always straightforward, decision-makers seek ways to invest such that the overall risk is maximally reduced. The solution in Table 5 highlights that controls for Email and Web Browser Protection, Malware Defences, Secure Configuration for Network Devices, and Security Awareness and Training are necessary to reduce the risk exposure of an organisation. In terms of the SB cyber security guidelines, results indicate that controls to cover phishing attacks and malware defences should be prioritised to reduce the number of social engineering-based attacks. Case F3 presents the solution which covers all five areas of the SB cyber security guide for a cost of £7,371.64. Besides controls to cover SB cyber security guide, the solution recommends implementing Wireless Access Control, which operates to cover additional vulnerabilities and strengthen boundary defences.

The first result that can be observed by comparing Tables 4 and 5 is that the Knapsack optimisation methodology provides solutions that entail a greater reduction of expected impact and lower cost when compared to the solutions of the Set Covering approach. This means that, by implementing the suggested solution in Tabl. 5, we achieve a better reduction of

expected impact and better chances to gain a positive return on the security investment.

Second, results indicate that investing more in security does not necessarily lead to an analogous reduction of risk, as illustrated in Figure 3. Finally, notice that the 95-th percentile corresponding to the solutions obtained via the Knapsack optimisation, in most cases, is lower than that corresponding to the solutions obtained via the Set Covering formulation. This emphasises how the former approach is more suitable for providing solutions for better control of risk within a specified budget. The results further elicit that practitioners must consider methods that optimise their security investment decisions, with respect to return and risk reduction, rather than having a cyber security strategy just as a compliance exercise. Apart from the expected PV of the impact of an attack, other objectives could also be implemented within eq (20), e.g. the value at risk (VaR) and the conditional VaR (CVaR), which renders this approach suitable for both risk-neutral and risk-averse decision-makers. However, the latter is not within the scope of the current paper and is left for future work.

TABLE 6. 2020 Top 25 software weaknesses from CWE labelled with relevant characteristic metrics used in the case study.

Rank	CWE	Name	Freq	Avg CVSS	Avg Exploitability Score (r_{ij})	Avg Complexity Score (s_{ij})	Avg Privileges Required	Avg User Interaction	Time
1	CWE-79	Improper Neutralisation of Input During Web Page Generation ('Cross-site Scripting')	3848	5.797	0.433	0.004	1.538	0.985	2.523
2	CWE-787	Out-of-bounds Write	2013	8.277	0.453	0.148	1.214	0.54	1.754
3	CWE-20	Improper Input Validation	1887	7.332	0.541	0.052	1.437	0.208	1.645
4	CWE-125	Out-of-bounds Read	1602	7.134	0.569	0.019	1.13	0.531	1.661
5	CWE-119	Improper Restriction of Operations within the Bounds of a Memory Buffer	1230	8.057	0.486	0.031	1.28	0.395	1.675
6	CWE-89	Improper Neutralisation of Special Elements used in an SQL Command ('SQL Injection')	911	8.976	0.759	0.013	1.447	0.015	1.462
7	CWE-200	Exposure of Sensitive Information to an Unauthorised Actor	1474	6.004	0.573	0.082	1.4	0.134	1.534
8	CWE-416	Use After Free	935	8.244	0.488	0.052	1.201	0.554	1.755
9	CWE-352	Cross-Site Request Forgery (CSRF)	877	8.074	0.534	0.009	1.048	0.985	2.033
10	CWE-78	Improper Neutralisation of Special Elements used in an OS Command ('OS Command Injection')	778	8.533	0.465	0.027	1.824	0.049	1.873
11	CWE-190	Integer Overflow or Wraparound	860	7.699	0.798	0.007	1.095	0.195	1.29
12	CWE-22	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal')	801	7.268	0.714	0.027	1.406	0.089	1.495
13	CWE-476	NULL Pointer Dereference	537	6.845	0.542	0.041	1.242	0.389	1.631
14	CWE-287	Improper Authentication	473	7.998	0.643	0.112	1.243	0.04	1.283
15	CWE-434	Unrestricted Upload of File with Dangerous Type	349	8.491	0.59	0.032	1.725	0.054	1.779
16	CWE-732	Incorrect Permission Assignment for Critical Resource	431	7.013	0.369	0.07	1.726	0.114	1.84
17	CWE-94	Improper Control of Generation of Code ('Code Injection')	309	8.741	0.604	0.049	1.55	0.11	1.66
18	CWE-522	Insufficiently Protected Credentials	306	7.878	0.534	0.052	1.621	0.02	1.641
19	CWE-611	Improper Restriction of XML External Entity Reference	285	7.871	0.694	0.056	1.344	0.154	1.498
20	CWE-798	Use of Hard-coded Credentials	245	8.738	0.755	0.078	1.176	0.012	1.188
21	CWE-502	Deserialisation of Untrusted Data	225	8.962	0.748	0.062	1.249	0.138	1.387
22	CWE-269	Improper Privilege Management	301	7.337	0.337	0.07	1.844	0.08	1.924
23	CWE-400	Uncontrolled Resource Consumption	268	7.106	0.792	0.034	1.146	0.127	1.273
24	CWE-306	Missing Authentication for Critical Function	206	8.088	0.88	0.024	1.063	0.015	1.078
25	CWE-862	Missing Authorisation	254	6.83	0.573	0.028	1.665	0.087	1.752

VI. CONCLUSION

Efficient cyber security risk management relies on managerial strategies responsive to various uncertainties associated with attacks. The need for such strategies becomes particularly pronounced considering the critical impact attacks may have on organisations and the limited time to make executive decisions. Hence, risk management within the area of cyber security is a considerably delicate task. The presence of uncertainties raises the incentive to postpone decisions and, in turn, the value of waiting, which is often a luxury that cannot be afforded. In this paper, we take into account the serial nature of the attack and the uncertainty in the time

required to exploit a vulnerability and develop a decision support framework to evaluate the risk exposure of an organisation and propose an optimal set of mitigation measures.

We develop a real-world use case using the 2020 CWE top 25 most dangerous software weaknesses and the CIS Controls to evaluate our framework. Our methodology provides security effective and cost-efficient solutions to counteract the most common attacks. We believe our work can assist cyber-security managers to select controls to optimally mitigate risk within a budget, given underlined uncertainties on a cyber-security breach. To the best of our knowledge, our framework is the first that adopts a discounted cash flow

approach to model and assess cyber risk for multi-phase attacks taking into account the uncertainty in the duration of each phase and proposes risk mitigation measures in the form of cyber-security controls to patch the identified weaknesses in the assets.

We aim to use the work presented in this paper to support attack models, such as [37] and [38], considering the strategic aspect of cyber-security interactions. This model would capture the interaction between the defender and the attacker as a non-cooperative game to determine the best responses against strategic attackers. From an application point of view, future work could extend the analysis to obtain recommendations against known threat actors by including adversarial information from CAPEC and MITRE ATT&CK framework.

Additionally, governments advise organisations to get cyber-security compliance certifications to demonstrate compliance with prescribed guidelines. For example, the UK government demands organisations to get certified for Cyber Essentials,¹⁶ which is a government-backed scheme aimed to protect organisations from a range of most common attacks. To be certified, organisations have to satisfy a list of requirements that cover five technical control themes: firewall, secure configuration, user access control, malware protection, and security update management. The Set Covering approach could be a comparable method used to identify controls to meet these requirements. However, as highlighted from the results, this is a basic method that provides ideas on how to invest in cyber-security and does not deal with cyber risk minimisation. On the other hand, the Knapsack Optimisation method could be used to overcome the inefficacy of the Set Covering method to identify controls that minimise cyber risk.

Besides its novel contributions to the field of cyber risk management, the proposed framework exhibits some limitations, which can admittedly become the basis of future work. First, the framework can be utterly improved by taking a more multifaceted approach where three types of cyber risk assessment are employed; threat-based, vulnerability-based, and control-based risk assessment. Such a direction will ultimately introduce new research challenges. These will entail the lack of online numeric data, besides the CVSS one, that will adequately facilitate the establishment of a use case to validate this novel framework in the most realistically way possible. Arguably, we believe that the lack of such datasets can act as a catalyst for researchers to pursue more practical paths to encounter organisations, interview them, and gather data about previous cyber incidents suffered. This collective information can be turned into an apt use case for attaining the requisite validation of the new framework.

To innovate towards this direction is to consider the MITRE ATT&CK framework as the source of data about APT groups, attack tactics, techniques, procedures and mitigation, in the context of cyber forensics [39], [40] and to determine optimal security countermeasures against the APT

groups [13]. Future work may also consider the strategic interactions between the Defender and the Attacker, as it has been previously studied within the literature of game theory as applied to cyber-security.

To effectively apply CENSOR in the real world, it is necessary to integrate it within existing cyber risk assessment tools, which typically require a high level of cybersecurity maturity. These tools identify Common Vulnerabilities and Exposures (CVEs), which can be used as input to CENSOR and lead to the instantiation of its model. However, such tools may not be readily available to small businesses due to the associated costs of implementation and maintenance. This represents a limitation of CENSOR when it comes to creating an impact in real-world scenarios.

It is worth noting that the choice of a small business use case was intended to facilitate the assessment of the framework, and was a proof-of-concept implementation for the purpose of the paper. However, it should be noted that the application of CENSOR is not limited to small businesses, and organisations of any size could potentially use it to optimise their cybersecurity risk level.

With regards to the effectiveness of CENSOR in real-world scenarios, a key challenge is the validation of the efficacy of the different cybersecurity controls [9], [41]. This is an ongoing industrial challenge,¹⁷ which can be addressed by assessing the degree of improvement in the cybersecurity posture of an organisation before and after the deployment of controls. Overall, the efficacy of cybersecurity controls can be assessed through measuring and tracking security-related data over time, such as the number of security incidents or the time taken to detect and respond to an incident. This most likely involves using a team of internal or external security experts to simulate an attack on a system and evaluate the effectiveness of existing controls.

An important direction for further research would be to combine the risk assessment and optimisation framework presented in this paper with the game-theoretic frameworks such as [25], [38], [42], and [43], to explore how strategic interactions impact the optimal choice of controls and support cyber security decisions. Additionally, the optimisation functionality of CENSOR could be extended in terms of its scope to allow for different objective functions, reflected in risk measures, such as the VaR and the CVaR. This will facilitate the analysis of the implication of risk preferences for the optimal selection of controls, and, specifically, how the latter may depend of the level of the decision-maker's risk aversion.

APPENDIX

The general expression of the CDF of the PV is described in (A-1).

$$\begin{aligned}\Theta_{U_1}(u_1) &= \mathbb{P}\left(K_1 e^{-\rho T_1} \leq u_1\right) = \mathbb{P}\left(T_1 \geq \frac{1}{\rho} \ln\left(\frac{K_1}{u_1}\right)\right) \\ &= 1 - \Phi_{W_1}\left(\frac{1}{\rho} \ln\left(\frac{K_1}{u_1}\right)\right)\end{aligned}\quad (\text{A-1})$$

¹⁷<https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/the-risk-based-approach-to-cybersecurity>

¹⁶<https://www.ncsc.gov.uk/cyberessentials/overview>

Note that we can insert the expression for any CDF of W_1 into (A-1) in order to obtain the CDF of U_1 . If $W_1 \sim \exp(\lambda_1)$, then the CDF and PDF of the PV is indicated in (A-2) and (A-3), respectively.

$$\Theta_{U_1}(u_1) = 1 - \left[1 - e^{-\frac{\lambda_1}{\rho} \ln\left(\frac{K_1}{u_1}\right)} \right] = \left(\frac{u_1}{K_1}\right)^{\frac{\lambda_1}{\rho}} \quad (\text{A-2})$$

$$\theta_{U_1}(u_1) = \frac{\lambda_1}{\rho} K_1^{-\frac{\lambda_1}{\rho}} \cdot u_1^{\frac{\lambda_1}{\rho}-1} \quad (\text{A-3})$$

Next, having derived the analytical expression for the distribution of U_1 , we can proceed with the derivation of its key characteristics, e.g. the mean and variance.

$$\mu_1 = \int_0^{K_1} u_1 \frac{\lambda_1}{\rho} K_1^{-\frac{\lambda_1}{\rho}} \cdot u_1^{\frac{\lambda_1}{\rho}-1} du_1 = \frac{\lambda_1}{\lambda_1 + \rho} K_1 \quad (\text{A-4})$$

$$\sigma_1^2 = \mathbb{E}[U_1^2] - \mathbb{E}[U_1]^2 = \left[\frac{\lambda_1}{\lambda_1 + 2\rho} - \left(\frac{\lambda_1}{\lambda_1 + \rho}\right)^2 \right] K_1^2 \quad (\text{A-5})$$

Proof of Proposition 1: Here, we want to determine the distribution of the PV of the phase 2 attack, which is expressed as $U_2 = K_2 e^{-\rho W_2}$. We can obtain the expression of the CDF of U_2 for a generic distribution function of W_2 , $\Phi_{W_2}(\cdot)$.

$$\begin{aligned} \mathbb{P}\left(K_2 e^{-\rho W_2} \leq u_2\right) &= 1 - \mathbb{P}\left(W_2 \leq \frac{1}{\rho} \ln\left(\frac{K_2}{u_2}\right)\right) \\ &= 1 - \Phi_{W_2}\left(\frac{1}{\rho} \ln\left(\frac{K_2}{u_2}\right)\right) \end{aligned} \quad (\text{A-6})$$

In the special case where $W_2 \sim \text{Hypoexp}(\lambda_1, \lambda_2)$ the expression of the CDF of W_2 is:

$$\Phi_{W_2}(w_2) = 1 - \frac{\lambda_2}{\lambda_2 - \lambda_1} e^{-\lambda_1 w_2} + \frac{\lambda_1}{\lambda_2 - \lambda_1} e^{-\lambda_2 w_2} \quad (\text{A-7})$$

Consequently, the CDF and PDF of U_2 is indicated in (A-8) and (A-9), respective.

$$\Theta_{U_2}(u_2) = \frac{\lambda_2}{\lambda_2 - \lambda_1} \left(\frac{u_2}{K_2}\right)^{\frac{\lambda_1}{\rho}} - \frac{\lambda_1}{\lambda_2 - \lambda_1} \left(\frac{u_2}{K_2}\right)^{\frac{\lambda_2}{\rho}} \quad (\text{A-8})$$

$$\theta_{U_2}(u_2) = \frac{\lambda_1 \lambda_2}{\lambda_2 - \lambda_1} \frac{1}{\rho u_2} \left[\left(\frac{u_2}{K_2}\right)^{\frac{\lambda_1}{\rho}} - \left(\frac{u_2}{K_2}\right)^{\frac{\lambda_2}{\rho}} \right] \quad (\text{A-9})$$

Proof of Proposition 2: Next, we determine the CDF of $U_n = K_n e^{-\rho W_n}$, where $W_n = T_1 + T_2 + \dots + T_n$. The expression of the CDF of U_2 for a general distribution function $\Phi_{W_n}(w_n)$ is:

$$\Theta_{U_n}(u_n) = 1 - \Phi_{W_n}\left(\frac{1}{\rho} \ln\left(\frac{K_n}{u_n}\right)\right) \quad (\text{A-10})$$

Depending on the distribution of W_n , we can derive a specific expression for the distribution of the U_n . Here, we assume that $W_n \sim \text{Hypo}(\lambda_1, \lambda_2, \dots, \lambda_n)$ and the CDF and PDF of W_n are described in (A-11) and (A-12), respectively.

$$\Phi_{W_n}(w_n) = \sum_{i=1}^n \left[1 - e^{-\lambda_i w_n} \right] \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \quad (\text{A-11})$$

$$\phi_{W_n}(w_n) = \sum_{i=1}^n \lambda_i e^{-\lambda_i w_n} \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \quad (\text{A-12})$$

Hence, the CDF U_n is obtained by setting $w_n = \frac{1}{\rho} \ln\left(\frac{K_n}{u_n}\right)$ in (A-11) and then substituting (A-11) into (A-10).

$$\Theta_{U_n}(u_n) = 1 - \sum_{i=1}^n \left[1 - \left(\frac{u_n}{K_n}\right)^{\frac{\lambda_i}{\rho}} \right] \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \quad (\text{A-13})$$

$$\theta_{U_n}(u_n) = \sum_{i=1}^n \left[\frac{\lambda_i}{\rho} \left(\frac{u_n}{K_n}\right)^{\frac{\lambda_i}{\rho}-1} \frac{1}{K_n} \right] \prod_{j \neq i} \frac{\lambda_j}{\lambda_j - \lambda_i} \quad (\text{A-14})$$

□

REFERENCES

- [1] R. Flage, T. Aven, E. Zio, and P. Baraldi, "Concerns, challenges, and directions of development for the issue of representing uncertainty in risk assessment," *Risk Anal.*, vol. 34, no. 7, pp. 1196–1207, Jul. 2014.
- [2] Y. Cherdantseva, P. Burnap, A. Blyth, P. Eden, K. Jones, H. Soulsby, and K. Stoddart, "A review of cyber security risk assessment methods for SCADA systems," *Comput. Secur.*, vol. 56, pp. 1–27, Feb. 2016.
- [3] G. Gonzalez-Granadillo, S. A. Menesidou, D. Papamartzivanos, R. Romeu, D. Navarro-Llobet, C. Okoh, S. Nifakos, C. Xenakis, and E. Panaousis, "Automated cyber and privacy risk management toolkit," *Sensors*, vol. 21, no. 16, p. 5493, Aug. 2021.
- [4] P. Voigt and A. Von dem Bussche, "The EU general data protection regulation (GDPR)," in *A Practical Guide*, vol. 10, no. 3152676, 1st ed. Cham, Switzerland: Springer, 2017, p. 5555.
- [5] A. Fielder, S. König, E. Panaousis, S. Schauer, and S. Rass, "Risk assessment uncertainties in cybersecurity investments," *Games*, vol. 9, no. 2, p. 34, Jun. 2018.
- [6] F. Smeraldi and P. Malacaria, "How to spend it: Optimal investment for cyber security," in *Proc. 1st Int. Workshop Agents CyberSecurity*, May 2014, pp. 1–4.
- [7] S. Panda, E. Panaousis, G. Loukas, and C. Laoudias, "Optimizing investments in cyber hygiene for protecting healthcare users," in *From Lambda Calculus to Cybersecurity Through Program Analysis*. Cham, Switzerland: Springer, 2020, pp. 268–291.
- [8] E. Scott, S. Panda, G. Loukas, and E. Panaousis, "Optimising user security recommendations for AI-powered smart-homes," in *Proc. IEEE Conf. Dependable Secure Comput. (DSC)*, Jun. 2022, pp. 1–8.
- [9] A. Fielder, E. Panaousis, P. Malacaria, C. Hankin, and F. Smeraldi, "Decision support approaches for cyber security investment," *Decis. Support Syst.*, vol. 86, pp. 13–23, Jun. 2016.
- [10] S. Panda, D. W. Woods, A. Laszka, A. Fielder, and E. Panaousis, "Post-incident audits on cyber insurance discounts," *Comput. Secur.*, vol. 87, Nov. 2019, Art. no. 101593.
- [11] M. K. Daly, "Advanced persistent threat," *Usenix*, vol. 4, no. 4, pp. 2013–2016, 2009.
- [12] A. Nisioti, G. Loukas, S. Rass, and E. Panaousis, "Game-theoretic decision support for cyber forensic investigations," *Sensors*, vol. 21, no. 16, p. 5300, Aug. 2021.
- [13] M. Ahmed, S. Panda, C. Xenakis, and E. Panaousis, "MITRE ATT&CK-driven cyber risk assessment," in *Proc. 17th Int. Conf. Availability, Rel. Secur.*, Aug. 2022, pp. 1–10.
- [14] N. Pitropakis, E. Panaousis, A. Giannakoulis, G. Kalpakis, R. D. Rodriguez, and P. Sarigiannidis, "An enhanced cyber attack attribution framework," in *Proc. Int. Conf. Trust Privacy Digit. Bus.*, 2018, pp. 213–228.
- [15] S. Panda, "Optimal strategies for cyber security decision-making," Ph.D. dissertation, Dept. Comput. Sci., Univ. Surrey, Guildford, U.K., 2022.
- [16] A. K. Dixit, R. K. Dixit, and R. S. Pindyck, *Investment under uncertainty*. Princeton, NJ, USA: Princeton Univ. Press, 1994.
- [17] M. Benaroch, "Real options models for proactive uncertainty-reducing mitigations and applications in cybersecurity investment decision making," *Inf. Syst. Res.*, vol. 29, no. 2, pp. 315–340, Jun. 2018.
- [18] A. Etheridge and M. Baxter, *A Course Financial Calculus*. Cambridge, U.K.: Cambridge Univ. Press, 2002.

- [19] L. A. Gordon, M. P. Loeb, W. Lucyshyn, and L. Zhou, "The impact of information sharing on cybersecurity underinvestment: A real options perspective," *J. Accounting Public Policy*, vol. 34, no. 5, pp. 509–519, Sep. 2015.
- [20] L. A. Gordon, M. P. Loeb, and W. Lucyshyn, "Information security expenditures and real options: A wait-and-see approach," *Comput. Secur. J.*, vol. 19, no. 2, pp. 1–16, 2003.
- [21] M. Chronopoulos, E. Panaousis, and J. Grossklags, "An options approach to cybersecurity investment," *IEEE Access*, vol. 6, pp. 12175–12186, 2018.
- [22] L. A. Gordon and M. P. Loeb, "The economics of information security investment," *ACM Trans. Inf. Syst. Secur.*, vol. 5, no. 4, pp. 438–457, Nov. 2002.
- [23] K. Zheng, L. A. Albert, J. R. Luedtke, and E. Towle, "A budgeted maximum multiple coverage model for cybersecurity planning and management," *IISE Trans.*, vol. 51, no. 12, pp. 1303–1317, Dec. 2019.
- [24] H. M. J. Almohri, L. T. Watson, D. Yao, and X. Ou, "Security optimization of dynamic networks with probabilistic graph modeling and linear programming," *IEEE Trans. Depend. Secure Comput.*, vol. 13, no. 4, pp. 474–487, Jul. 2016.
- [25] M. Khouzani, Z. Liu, and P. Malacaria, "Scalable min-max multi-objective cyber-security optimisation over probabilistic attack graphs," *Eur. J. Oper. Res.*, vol. 278, no. 3, pp. 894–903, Nov. 2019.
- [26] R. T. Mercuri, "Analyzing security costs," *Commun. ACM*, vol. 46, no. 6, pp. 15–18, Jun. 2003.
- [27] D. Geer, K. Soo Hoo, and J. Andrew, "Information security: Why the future belongs to the quants," *IEEE Secur. Privacy*, vol. 1, no. 2, pp. 32–40, Jul./Aug. 2003.
- [28] A. Niakanlahiji, J. Wei, M. R. Alam, Q. Wang, and B.-T. Chu, "Shadow-Move: A stealthy lateral movement strategy," in *Proc. 29th USENIX Secur. Symp. (USENIX Secur.)*, 2020, pp. 559–576.
- [29] N. Provos, M. Friedl, and P. Honeyman, "Preventing privilege escalation," in *Proc. USENIX Secur. Symp.*, 2003, pp. 1–11.
- [30] M. E. Whitman and H. J. Mattord, *Principles of Information Security*. Boston, MA, USA: Cengage Learning, 2011.
- [31] W. Sonnenreich, J. Albanese, and B. Stout, "Return on security investment (ROSI)—A practical quantitative model," *J. Res. Pract. Inf. Technol.*, vol. 38, no. 1, pp. 45–56, 2006.
- [32] L. A. Gordon, M. P. Loeb, W. Lucyshyn, and L. Zhou, "Increasing cybersecurity investments in private sector firms," *J. Cybersecurity*, vol. 1, no. 1, pp. 3–17, 2015.
- [33] R. Böhme, "Security metrics and security investment models," in *Proc. Int. Workshop Secur. Cham, Switzerland: Springer*, 2010, pp. 10–24.
- [34] M. S. Ross, *Introduction to Probability Models*. New York, NY, USA: Academic, 2010.
- [35] J. E. Beasley and P. C. Chu, "A genetic algorithm for the set covering problem," *Eur. J. Oper. Res.*, vol. 94, no. 2, pp. 392–404, Oct. 1996.
- [36] S. Kumar, A. Sarkar, and A. Sur, "A resource allocation framework for adaptive video streaming over LTE," *J. Netw. Comput. Appl.*, vol. 97, pp. 126–139, Nov. 2017.
- [37] S. Amin, X. Litrico, S. Sastry, and A. M. Bayen, "Cyber security of water SCADA systems—Part I: Analysis and experimentation of stealthy deception attacks," *IEEE Trans. Control Syst. Technol.*, vol. 21, no. 5, pp. 1963–1970, Sep. 2013.
- [38] N. Boumkheld, S. Panda, S. Rass, and E. Panaousis, "Honeypot type selection games for smart grid networks," in *Proc. Int. Conf. Decis. Game Theory Secur. Cham, Switzerland: Springer*, 2019, pp. 85–96.
- [39] A. Nisioti, G. Loukas, A. Laszka, and E. Panaousis, "Data-driven decision support for optimizing cyber forensic investigations," *IEEE Trans. Inf. Forensics Security*, vol. 16, pp. 2397–2412, 2021.
- [40] S. Atefi, S. Panda, E. Panaousis, and A. Laszka, "Principled data-driven decision support for cyber-forensic investigations," 2022, *arXiv:2211.13345*.
- [41] Y. Zhang and P. Malacaria, "Bayesian Stackelberg games for cybersecurity decision support," *Decis. Support Syst.*, vol. 148, Sep. 2021, Art. no. 113599.
- [42] S. Panda, S. Rass, S. Moschogiannis, K. Liang, G. Loukas, and E. Panaousis, "HoneyCar: A framework to configure honeypot vulnerabilities on the internet of Vehicles," *IEEE Access*, vol. 10, pp. 104671–104685, 2022.
- [43] I. Kalderemidis, A. Farao, P. Bountakas, S. Panda, and C. Xenakis, "GTM: Game theoretic methodology for optimal cybersecurity defending strategies and investments," in *Proc. 17th Int. Conf. Availability, Rel. Secur.*, Aug. 2022, pp. 1–9.



MARIA TSIOTRA received the bachelor's degree in mathematics from the University of Brighton, Brighton, U.K. She was an Analyst with the National Audit Office and a Research Assistant with the City, University of London. She is currently a Senior Analyst with The Open University. Her research interests include applied mathematics, energy economics, operational research, software development, and big data analytics.



SAKSHYAM PANDA (Member, IEEE) received the dual M.Sc. degrees in human-computer interaction and design from the KTH Royal Institute of Technology, Sweden, and Aalto University, Finland, in 2018, and the Ph.D. degree in computer science from the University of Surrey, U.K., in 2022. He is currently a Research Associate of cyber security and privacy with the Faculty of Engineering and Science, University of Greenwich, London, U.K. His research interests include

optimization of cyber security, and privacy decisions and robustness against strategic adversaries.



MICHAEL CHRONOPOULOS received the bachelor's degree in mathematics from the University of Ioannina, Ioannina, Greece, the M.Sc. degree in statistics from Western Michigan University, Kalamazoo, and the Ph.D. degree in statistics from University College London, London, U.K. He was a Research Fellow with the London Business School, London, and a Research Associate with the Department of Management Science and Innovation, University College London. He is currently

a Senior Lecturer of actuarial science with the Bayes Business School, City, University of London. His research interests include operations research, real options, game theory, mathematical finance, and energy economics. Specifically, he is interested in the application of financial and operational research methods for pricing financial and energy instruments under risk aversion and competition.



EMMANOUIL PANAOUSIS received the B.Sc. degree in informatics and telecommunications from the National and Kapodistrian University of Athens, Greece, in 2006, the M.Sc. degree in computer science from the Athens University of Economics and Business Greece, in 2008, and the Ph.D. degree in computer science from Kingston University, London, U.K., in 2012. He was a Post-doctoral Researcher with the Queen Mary University of London, U.K., from 2013 to 2014; a Visiting Researcher with Imperial College London, U.K., from 2013 to 2016; a Lecturer and a Senior Lecturer with the University of Brighton, from 2014 to 2017; a Lecturer and a Senior Lecturer with the University of Surrey, from 2017 to 2019; and an Associate Professor with the University of Greenwich, London, U.K., from 2019 to 2022, where he is currently a Professor of Cyber Security. His research has been funded by the European Commission, the U.K. National Cyber Security Centre, and the Engineering and Physical Sciences Research Council. His research and teaching activities focus on socio-technical challenges of cyber security and data privacy.

• • •