

Optimizing Investments in Cyber Hygiene for Protecting Healthcare Users

Sakshyam Panda¹, Emmanouil Panaousis²,
George Loukas², and Christos Laoudias³

¹ University of Surrey, UK

`s.panda@surrey.ac.uk`

² University of Greenwich, UK

`{e.panaousis,g.loukas}@greenwich.ac.uk`

³ University of Cyprus, Cyprus

`laoudias.christos@ucy.ac.cy`

Abstract. Cyber hygiene measures are often recommended for strengthening an organization's security posture, especially for protecting against social engineering attacks that target the human element. However, the related recommendations are typically the same for all organizations and their employees, regardless of the nature and the level of risk for different groups of users. Building upon an existing cybersecurity investment model, this paper presents a tool for optimal selection of cyber hygiene safeguards, which we refer as the Optimal Safeguards Tool (OST). The model combines game theory and combinatorial optimization (0-1 Knapsack) taking into account the probability of each user group to being attacked, the value of assets accessible by each group, and the efficacy of each control for a particular group. The model considers indirect cost as the time employees could require for learning and training against an implemented control. Utilizing a game-theoretic framework to support the Knapsack optimization problem permits us to optimally select safeguards' application levels minimizing the aggregated expected damage within a security investment budget.

We evaluate OST in a healthcare domain use case. In particular, on the Critical Internet Security (CIS) Control group 17 for implementing security awareness and training programs for employees belonging to the ICT, clinical and administration personnel of a hospital. We compare the strategies implemented by OST against alternative common-sense defending approaches for three different types of attackers: Nash, Weighted and Opportunistic. Our results show that Nash defending strategies are consistently better than the competing strategies for all attacker types with a minor exception where the Nash defending strategy, for a specific game, performs at least as good as other common-sense approaches. Finally, we illustrate the alternative investment strategies on different Nash equilibria (called plans) and discuss the optimal choice using the framework of 0-1 Knapsack optimization.

Keywords: Cybersecurity, Cyber hygiene, Healthcare, Optimization, Training and awareness, CIS control, Game theory⁴

⁴ The original publication is available at <https://link.springer.com>

1 Introduction

In the last few years, several cybersecurity incidents have taken place in the healthcare sector, including the WannaCry ransomware, which influenced globally the cybersecurity landscape⁵. The 2018 Ponemon Cost of a Data Breach study⁶ shows that the healthcare industry has the highest cost per record breached in a cyber incident, at \$408. This is almost twice the equivalent cost per record breached in the financial sector. This calls for the effective preparation of healthcare organizations in an ever-evolving cyber attack landscape. An example project that is addressing this from the perspective of training the users in the sector is H2020 CUREX project⁷, which allows a healthcare provider to assess the realistic cybersecurity and privacy risks they are exposed to [1].

Yet, a recent report from Mckinsey⁸ states that almost all companies systematically over-invest in the protection of assets that have no risk while at the same time they under-fund the protection of high-risk assets. Furthermore, regarding bearing costs of cybersecurity controls, in a survey from KPMG⁹, 43% of correspondents stated that they did not increase their cybersecurity budget even though high profile security breaches have been widely known. So, effective risk management is not only about assessing the risk correctly but also about selecting the controls that are optimal given the cost constraints of adopting them. To address the challenge of optimal control selection, in this paper, we formulate a model and tool for suggesting mathematically optimal *cyber hygiene* strategies minimising the cyber risk.

Regarding cyber hygiene, we adopt the recent definition proposed by [2], which relates it to “*the cyber security practices that online consumers should engage in to protect the safety and integrity of their personal information on their Internet enabled devices from being compromised in a cyber-attack.*”

Towards the goal of optimizing cyber hygiene, we extend the model presented in [3] so that:

- the Attacker’s target is a user group (focusing on social engineering attacks) instead of (asset, vulnerability) pair of the system;
- the Indirect cost of a safeguards’ application depends not only on the safeguard itself but also on the size of the user group (i.e., number of users) and more specifically it increases with the group size;
- we adopt an aggregated risk model, as the objective function of Knapsack optimization problem, rather than the weakest-link defending against a variety of attacks that can cause, in total, highest aggregated damage and;
- we use a “small” healthcare case study as a preliminary example to evaluate the OST against other common-sense approaches for a number of attacking strategies that have not been simulated in [3].

⁵ <https://www.telegraph.co.uk/technology/2018/10/11/wannacry-cyber-attack-cost-nhs-92m-19000-appointments-cancelled>.

⁶ <https://securityintelligence.com/series/ponemon-institute-cost-of-a-data-breach-2018>.

⁷ <https://cordis.europa.eu/project/rcn/220350/factsheet/en>.

⁸ <https://www.mckinsey.com/business-functions/risk/our-insights/cyber-risk-measurement-and-the-holistic-cybersecurity-approach>.

⁹ <https://advisory.kpmg.us/content/dam/advisory/en/pdfs/cyber-report-healthcare.pdf>.

Our analysis results show that the game-theoretic approach increases risk control efficacy, by selecting an optimal combination of safeguard application levels, compared with alternative common-sense approaches. In addition, our use case designed for the healthcare domain exhibits a number of interchangeably optimal investment strategies subject to a budget constraint under the framework of 0-1 Knapsack optimization.

The remainder of this paper is organized as follows. Section 2 presents the related work in both the fields of (i) user-oriented cybersecurity safeguards and (ii) optimization of cybersecurity countermeasures including security investments. Section 3 presents both the game-theoretic model used to determine optimal cybersecurity safeguard plans as well as the optimization problem modeled and solved to derive the best ways to invest in these safeguards given a limited available budget. In Section 4, we undertake comparisons of the game-theoretic defending strategies against alternative common-sense approaches as well as we plot the results of the Knapsack optimization to illustrate the optimal investment solutions. Finally, Section 5 concludes this paper by summarizing its main contributions and highlighting future work to be undertaken to further improve the performance and the usability of our model.

2 Related Work

This work has been inspired by a previous work of Fielder et al. [3] where the authors have proposed decision support methodologies for the optimal choice of cybersecurity controls within an investment budget. They have addressed cybersecurity investment decisions by proposing different approaches; a game-theoretic approach, a combinatorial optimization approach and a mix of both called *hybrid*. This paper utilizes the latter method to recommend the optimal choice of safeguards for healthcare organizations. In this section, we discuss two classes of work relevant to this paper: literature on *cyber hygiene in healthcare* - more specifically on the *user-oriented cybersecurity safeguards*, and literature on *optimal selection of cybersecurity safeguards*. Note that the literature covered on optimal selection of cybersecurity safeguards mainly highlight work beyond the literature covered in [3].

2.1 Cyber Hygiene in Healthcare

There have been growing concerns that the existing cybersecurity posture of healthcare organizations are insufficient and this has already impacted the confidentiality [4] and integrity of medical data [5]. Further, many healthcare organizations are still using legacy systems such as Windows XP and Windows NT 3.1 which Microsoft has long stopped supporting¹⁰, allowing adversaries to easily breach the defenses (e.g., WannaCry attacks on NHS¹¹). In general, healthcare

¹⁰ <https://www.itpro.co.uk/public-sector/27740/nine-in-10-nhs-trusts-still-use-windows-xp>.

¹¹ <https://www.nao.org.uk/wp-content/uploads/2017/10/Investigation-WannaCry-cyber-attack-and-the-NHS-Summary.pdf>.

organizations being rich sources of valuable data and relatively weaker security postures have become attractive targets for cybercrime [6]. The weaker security posture that they exhibit is primarily due to lack of adequate cybersecurity budget resulting in limited access to technology and expertise [7].

Besides, investment in cybersecurity has not been traditionally considered essential for healthcare systems as emphasis has predominantly been upon providing patient care and people believed that there would be no motivation to attack them. On the other hand, the increasing use of Internet of Things (IoT) technologies in healthcare has widened the attack surface beyond electronic health record databases and privacy issues to physical safety [8]. Alongside technical aspects, the role of the user in cybersecurity is paramount, as a significant proportion of attacks target the users directly through deceptive means such as application masquerading and spear-phishing. This is particularly the case in healthcare as deceiving a nurse, doctor, healthcare IT professional or administrator can impact the privacy and physical safety of patients [9].

With the increasing usage of technology, the role that humans play in underlying security processes will continually expand. Heartfield and Loukas [10] have developed a framework involving humans to effectively detect and report semantic social engineering attacks against them. Their results illustrate that involving users significantly improves the cyber threat detection rate affirming the importance of the human in cybersecurity. This further depicts that humans can no longer be seen as a threat and/or vulnerability in cybersecurity.

Acknowledging the importance of human in cybersecurity along with the increase in the severity of breaches, security experts, policymakers and governments are urging to improve cyber hygiene. Such et al. [11] have demonstrated that Cyber Essentials¹² have worked well for SMEs in mitigating threats exploiting vulnerabilities remotely using commodity-level exploitation tools. From a human-cyber interaction perspective, Vishwanath et al. [2] have demonstrated that cyber hygiene practices positively impact individuals' cyber attitude which is pivotal to cyber safety. These studies have actively exhibited that even general concepts of basic cyber hygiene work in different organizational contexts and can convincingly reduce cyber risk.

Security training in healthcare has been studied for over 20 years. It ranges from an exploratory analysis of the factors that healthcare professionals need to focus on, up to highly targeted digital applications (e.g., [12]) and platforms for raising awareness of healthcare data privacy and security risks. Furnell et al. [13] discussed the necessity to promote information security issues and the need for appropriate training and awareness initiatives in healthcare institutions. They have highlighted factors to consider while designing training and awareness programmes to familiarize healthcare personnel with basic security concepts and procedures.

The effect to which security training and awareness programmes work for different users has been studied from multiple angles. The authors have shown that specifically for deception-based attacks, such as semantic social engineering

¹² <https://www.gov.uk/government/publications/cyber-essentials-scheme-overview>.

[14], where self-study and work-based training are considerably more effective than formal education in cybersecurity [15]. Besides, the perceived origin of training materials i.e., from security experts, third party agencies, or peers can have large impacts on security outcomes [16].

2.2 Optimal Selection of Cybersecurity Controls

Cybersecurity has become a key factor in determining the growth of organizations relying on information systems as it is not only a defensive measure but also has become a strategic decision providing a competitive advantage over rivalry firms. Further, the potential loss due to cyber incidents has encouraged organizations to imperatively consider cybersecurity investment decisions, especially in deriving the optimum level of investments between risk treatment options. The objective of cybersecurity investment methodologies is to compute an optimal distribution of cybersecurity budget and one of the initial work studying this was performed by Gordon and Loeb [17].

Beyond previous works such as [3,18,19] and the related work investigated there, Nagurney et al. [20] have proposed a game-theoretic supply chain network model with retailers competing to maximize their expected profits. This maximization is based on determining optimal product transactions and cybersecurity investments under budget constraints. Along the direction of optimal cybersecurity investments, Wang [21] investigated the cybersecurity investment balance between acquiring knowledge and expertise, and deploying mitigation techniques. On the other hand, Chronopoulos et al. [22] have opted a real options approach to analyze the performance of optimal cybersecurity controls on organizations. In particular, the authors have analyzed the effects of the cost of cyber attacks and the time of arrival of cybersecurity controls on the organization's optimal strategy. Similar to these papers, our work also considers the choice of the optimal strategy based on the efficacy of the control towards mitigating cyber risks.

Most closely, in terms of methodology, related recent work on optimal cybersecurity investment is [23] where the authors have investigated the balance between investing in self-protection and cyber insurance. The key difference is that their optimization minimizes expected risk and cyber insurance premium, while our model optimizes considering the efficacy of control in mitigating the aggregated residual risk and the security investment budget. Besides this, our work uses a unique combination of game theory and combinatorial optimization inspired by [3].

3 Optimal Cyber Hygiene Safeguards Model

3.1 System Model

Our model assists in acquiring an optimal selection of safeguards using game theory and combinatorial optimization. We assume \mathcal{U} be the set of potential

user groups consisting of employees of a healthcare organization. Any employee of a user group being susceptible to malicious activities can use any of the safeguards from the set of available safeguards \mathcal{S} to improve their defense posture. However, each safeguard has a set of implementation levels \mathcal{L} with each level having different efficacies in improving the security posture of user groups.

Each user group i is associated with an impact value which expresses the level of expected damage to the healthcare organization, given a successful attack against a user of a group i . This impact is equivalent to the overall asset value in association with user group i and may relate to *confidentiality*, *integrity*, and *availability*. We further consider A_i to be a random variable that expresses the overall value of the assets that the user group i has access to. For simplicity, we let the users of a group have the same *access privileges*, thus having access to assets of the same value. Users of different groups have different *access privileges* due to their different roles (e.g., IT personnel, healthcare practitioners, and administration) and access to different assets. The vulnerability of a user group i , i.e., the probability of being compromised by an attack, is captured by the security level S_i exhibited by the user group i . We assume that S_i increases with the number of safeguards applied as well as their application level.

Furthermore, we denote R_i as the threat occurrence, i.e., the probability of a threat to attack the i user group, and L_i as the expected loss associated with a user group i . Using the well-known risk assessment formula, risk = (likelihood of being attacked) x (probability of success of this attack) x probable loss [24], we compute the risk as

$$L_i = R_i S_i A_i. \quad (1)$$

An attack against a user group i is partially mitigated by the efficacy value of the implemented cybersecurity safeguard. The efficacy parameter, modeled as a random variable, depends on the selected application level and can be represented as $E(j, i): \mathcal{L} \times \mathcal{U} \rightarrow [0, 1)$. It is evident from real-world practices that different implementation levels work differently on different users and this has motivated us in considering $E(j, i)$ rather than a single efficacy value for the level j against all user groups i . Note that $E(j, i)$ is determined by the application level j and the user group i . Due to the existence of 0-day vulnerabilities, we assume that $E(j, i) \neq 1$.

Remark 1. Different users have different likelihood of adopting a measure. A cyber hygiene measure works only when it is adopted, and this adaption rate distinguishes human users from systems. For example, a user may decide not to implement a cyber hygiene measure due to unfitting usability (e.g, hard to remember complex passwords) even if the optimization framework recommends otherwise.

Let $S(j, i)$ be the security level of a user group i when level j is implemented and can be expressed as $S(j, i) = 1 - E(j, i)$. Replacing L_i and S_i as $L(j, i)$ and $S(j, i)$, respectively, in formula 1, we compute the *cybersecurity loss* for a

safeguard application level j and target i as

$$L(j, i) = R_i A_i [1 - E(j, i)]. \quad (2)$$

Equation 2 implies the expected damage of the Defender when a user group i is successfully compromised given the investigated safeguard has been applied at level j .

While the application of a cybersecurity safeguard strengthens the defense of the healthcare organization, it is associated with two types of cost namely; *indirect* and *direct*. Examples of indirect cost are System Performance and Usability. We express the indirect cost of an application level j by the random variable $C: \mathcal{C} \times \mathcal{L} \times \mathcal{U} \rightarrow \mathbb{Z}^+$. Note that $C(j, i)$ adheres to the defined property for any safeguard against a user group i . Further, the indirect cost increases with an increase in the level of application of the safeguard i.e.,

$$j > j' \Leftrightarrow C(j, i) \geq C(j', i), \quad \forall j \neq j'. \quad (3)$$

From the above, we derive the *overall expected loss* of the organization when application level j is applied on user group i as

$$\sum_{i=1}^{|\mathcal{U}|} L(j, i) + C(j, i). \quad (4)$$

Each level has also a direct cost expressed by the random variable $F: \mathcal{L} \rightarrow \mathbb{Z}^+$ that maps the safeguards and application levels to the monetary cost of the plan. In this paper, we refer the direct cost to be the available investment budget of the organization. For reference purposes, the symbols used throughout this paper are described in Table 1.

3.2 Game-Theoretic Model for Selection of Safeguards Levels

This section presents a formal model for the selection of safeguard implementation levels for each of the available safeguards. The Defender chooses to implement (or apply as in this paper we use these two terms interchangeably) a cyber hygiene safeguard from \mathcal{S} , while the Attacker chooses to attack a user group from \mathcal{U} . The Defender must decide to apply this safeguard at a specific level (pure strategy) or combination of different levels (mixed strategy) both from \mathcal{L} . The higher the level, the greater is the applied degree of a cyber hygiene safeguard. We refer to the application of a safeguard s at a certain level j as *cybersecurity safeguard plan*. This strategic interaction is modeled as a game where the Defender chooses the level of a safeguard to implement rather than the safeguards from \mathcal{S} .

We define the Cyber Safeguard Game (CSG) between Defender and Attacker, as an *one-shot, bimatrix* game of *complete information* played for any of the safeguards leading to a total number of \mathcal{S} independent games. For simplicity, we have assumed no inter-dependencies between the safeguards, i.e., each safeguard mitigates a portion of the overall risk inflicted by the Attacker [25].

Symbol	Description
\mathcal{S}	Set of safeguards
\mathcal{U}	Set of users
\mathcal{L}	Set of safeguard implementation levels
R_i	Probability of group i to be attacked
S_i	Security level of group i
A_i	Asset value that group i has access to
λ	Maximum application level
U_d	Utility of the Defender
U_a	Utility of the Attacker
$\delta_{\sigma,j}$	Randomized Safeguard Strategy for safeguard σ at application level j
α	Randomized Attacking Strategy
$\alpha(i)$	Probability of attacking group i
L_i	Expected loss from group i
$L(j, i)$	Expected loss from group i when choosing application level j
$L(\delta_{\sigma,j}, i)$	Expected loss from group i when choosing Safeguards Plan $\delta_{\sigma,j}$
$C(j, i)$	Indirect cost of level j when applied to group i
$E(j, i)$	Efficacy of application level j on group i
$E(\delta_{\sigma,j}, i)$	Efficacy of safeguards plan $\delta_{\sigma,j}$ on group i
$\Gamma_{\sigma,\lambda}$	Cyber Safeguard Game for safeguard σ and maximum application level λ
$\delta_{\sigma,\lambda}^{NE}$	Nash Safeguards Plan
$F(\delta_{\sigma,j})$	Financial cost of Safeguards Plan $\delta_{\sigma,j}$
$F(\sigma, j)$	Financial cost of safeguard σ when applied at level j
B	Available financial budget to invest in Nash Safeguards Plans

Table 1: List of Symbols

The set of pure strategies of the Defender consists of all possible application levels, $j \in \mathcal{L}$, while the Attacker's pure strategies are the different user groups $i \in \mathcal{U}$ which could be targeted using attacks such as social engineering. Thus, in CSG a pure strategy profile is a pair of Defender and Attacker actions, $(j, i) \in \mathcal{L} \times \mathcal{U}$ giving a pure strategy space of size $|\mathcal{L}| \times |\mathcal{U}|$. For the rest of the paper, we adopt the convention where the Defender is the row player and the Attacker is the column player.

Each player's preferences are specified by a *payoff function* defined as $U_d : (j, i) \rightarrow \mathbb{R}_-$ and $U_a : (j, i) \rightarrow \mathbb{R}_+$ for the Defender and Attacker, respectively, for the pure strategy profile (j, i) . According to [?], we define a *preference relation* \succsim , when i is chosen by the Attacker, defined by $j \succsim j'$, if and only if $U_d(j, i) \geq U_d(j', i)$. In general, given the set \mathcal{L} of all available application levels of a safeguard, a rational Defender can choose a level (i.e., pure strategy) j^* that is *feasible*, that is $j^* \in \mathcal{L}$, and *optimal* in the sense that $j^* \succsim j, \forall j \in \mathcal{L}, j \neq j^*$; alternatively she solves the problem $\max_{j \in \mathcal{L}} U_d(j, i)$, for a user group $i \in \mathcal{U}$. Likewise, we define the preference relation for the Attacker, where $i \succsim i' \iff U_a(j, i) \geq U_a(j, i')$, for an application level $j \in \mathcal{L}$. CSG is a game defined for each cyber hygiene safeguard and it is realistic to assume that all levels may be available for selection by the Defender. Their availability

depends on the investment budget of the Defender and the overall financial cost of the game solution.

To derive optimal strategies for the Defender, we deploy the notion of *mixed strategies*. Since players act independently, we can enlarge their strategy spaces to allow them to base their decisions on the outcome of random events that create uncertainty to the opponent about individual strategic choices maximizing their payoffs. Hence, both Defender and Attacker deploy randomized (i.e., mixed) strategies. The mixed strategy δ of the Defender is a probability distribution over the different application levels (i.e. pure strategies) where $\delta(j)$ is the probability of applying level j under mixed strategy δ . We refer to a mixed strategy of the Defender as a *Randomized Safeguard Strategy* (RSS). For the finite nonempty set \mathcal{L} , let $\Pi_{\mathcal{L}}$ represent the set of all probability distributions over it, i.e.,

$$\Pi_{\mathcal{L}} := \{\delta \in \mathbb{R}^{+R} \mid \sum_{j \in \mathcal{L}} \delta(j) = 1\}. \quad (5)$$

Therefore a member of $\Pi_{\mathcal{L}}$ is a mixed strategy of the Defender. Likewise, the Attacker's mixed strategy is a probability distribution over the different available user groups. This is denoted by α , where $\alpha(i)$ is the probability of attacking the i -th user group under mixed strategy α . We refer to a mixed strategy of the Attacker as the *Randomized Attacking Strategy* (RAS). Alike (5), we express $\Pi_{\mathcal{U}}$ as the set of all probability distributions over the set of all Attacker's pure strategies (i.e., given by \mathcal{U}). Therefore, a member of $\Pi_{\mathcal{U}}$ is as a mixed strategy of the Attacker. From the above, the set of mixed strategy profiles of CSG is the Cartesian product of the individual mixed strategy sets, $\Pi_{\mathcal{L}} \times \Pi_{\mathcal{U}}$.

Definition 1. (*Support of RSS*) The support of δ is the set of application levels $\{j \mid \delta(j) > 0\}$, and it is denoted by $\text{supp}(\delta)$.

Definition 2. (*Support of RAS*) The support of α is the set of healthcare user groups $\{i \mid \alpha(i) > 0\}$, and it is denoted by $\text{supp}(\alpha)$.

The above definitions state that the subset of applications levels (resp. user groups) that are assigned positive probability by the mixed strategy δ (resp. α) is called the *support* of δ (resp. α). Note that a pure strategy is a special case of a mixed strategy, in which the support is a single action.

Now that we have defined the mixed strategies of the players, we define CSG as the finite strategic game

$$\Gamma := \langle (\text{Defender, Attacker}), \Pi_{\mathcal{L}} \times \Pi_{\mathcal{U}}, (U_d, U_a) \rangle. \quad (6)$$

For a given mixed strategy profile $(\delta, \alpha) \in \Pi_{\mathcal{L}} \times \Pi_{\mathcal{U}}$, we denote by $U_d(\delta, \alpha)$, and $U_a(\delta, \alpha)$ the expected payoff values of the Defender and Attacker, where the expectation is due to the independent randomization according to mixed strategies δ , and α . This can be formally represented as

$$U_d(\delta, \alpha) := \sum_{j \in \mathcal{L}} \sum_{i \in \mathcal{U}} U_d(j, i) \delta(j) \alpha(i), \quad (7)$$

and similarly

$$U_a(\boldsymbol{\delta}, \boldsymbol{\alpha}) := \sum_{j \in \mathcal{L}} \sum_{i \in \mathcal{U}} U_a(j, i) \boldsymbol{\delta}(j) \boldsymbol{\alpha}(i). \quad (8)$$

By using the preference relation we can say that, for an Attacker's mixed strategy $\boldsymbol{\alpha}$, the Defender prefers to follow the RSS $\boldsymbol{\delta}$ as opposed to $\boldsymbol{\delta}'$ (i.e., $\boldsymbol{\delta} \succsim \boldsymbol{\delta}'$), if and only if $U_d(\boldsymbol{\delta}, \boldsymbol{\alpha}) \geq U_d(\boldsymbol{\delta}', \boldsymbol{\alpha})$.

Definition 3. *The Defender's (resp. Attacker's) best response to the mixed strategy $\boldsymbol{\alpha}$ (resp. $\boldsymbol{\delta}$) of the Attacker (resp. Defender) is an RSS (resp. RAS) $\boldsymbol{\delta}^{\text{BR}} \in \Pi_{\mathcal{L}}$ (resp. $\boldsymbol{\alpha}^{\text{BR}} \in \Pi_{\mathcal{U}}$) such that $U_d(\boldsymbol{\delta}^{\text{BR}}, \boldsymbol{\alpha}) \geq U_d(\boldsymbol{\delta}, \boldsymbol{\alpha}), \forall \boldsymbol{\delta} \in \Pi_{\mathcal{L}}$ (resp. $U_a(\boldsymbol{\delta}, \boldsymbol{\alpha}^{\text{BR}}) \geq U_a(\boldsymbol{\delta}, \boldsymbol{\alpha}), \forall \boldsymbol{\alpha} \in \Pi_{\mathcal{U}}$).*

Remark 2. The game-theoretic solutions that we propose in the next section involve *randomization*. For instance, in a mixed equilibrium, each player's randomization leaves the other *indifferent* across her randomization support. These choices can be deliberately randomized, however these are not the only equilibria interpretations. For instance, the probabilities over the pure actions (i.e., application level or user group pure selections) can represent (i) time averages of an "adaptive" player, (ii) a vector of fractions of a "population", where each player type adopts pure strategies and, (iii) a "belief" vector that each player has about the other regarding their behavior.

3.3 CSG solutions

Given the definition of CSG and its components, we derive optimal strategies for the Defender. First, we investigate the problem of determining best RSSs and RASs (i.e., mixed strategies), for the Defender and the Attacker respectively, when both players are strategic and play simultaneously.

As we have not explicitly defined the *strategic type* of Attacker, we consider different types of solutions based on various Attacker behaviors. This analysis will allow us to draw robust conclusions regarding the *overall optimal* Defender strategy, which will minimize expected damages *regardless of the Attacker type*.

The most commonly used solution concept in game theory is that of *Nash Equilibrium* (NE) [?]. This concept captures a steady state of the play of the CSG in which both Defender and Attacker hold the correct expectation about the other players' behavior and they act rationally. A NE dictates optimal responses to each other's actions, keeping the others' strategies fixed, i.e., strategy profiles that are resistant against unilateral deviations of players.

Definition 4. *In any Cyber Safeguard Game, a mixed strategy profile $(\boldsymbol{\delta}^{\text{NE}}, \boldsymbol{\alpha}^{\text{NE}})$ of Γ is a mixed NE if and only if*

1. $\boldsymbol{\delta}^{\text{NE}} \succsim \boldsymbol{\delta}, \forall \boldsymbol{\delta} \in \Pi_{\mathcal{L}}$, when the Attacker chooses $\boldsymbol{\alpha}^{\text{NE}}$, i.e.

$$U_d(\boldsymbol{\delta}^{\text{NE}}, \boldsymbol{\alpha}^{\text{NE}}) \geq_{\forall \boldsymbol{\delta} \in \Pi_{\mathcal{L}}} U_d(\boldsymbol{\delta}, \boldsymbol{\alpha}^{\text{NE}}); \quad (9)$$

2. $\alpha^{\text{NE}} \succsim \alpha, \forall \alpha \in \Pi_{\mathcal{U}}$, when the Defender chooses δ^{NE} , i.e.

$$U_a(\delta^{\text{NE}}, \alpha^{\text{NE}}) \geq_{\forall \alpha \in \Pi_{\mathcal{U}}} U_a(\delta^{\text{NE}}, \alpha). \quad (10)$$

Definition 5. The Nash Safeguards Plan (NSP), denoted by δ^{NE} , is a probability distribution over the different levels, as determined by the NE of the CSG.

Example 1. For a safeguard with 3 application levels including level 0, which corresponds to not applying the safeguard at all, an NSP (0, 0.2, 0.8) dictates that 20% of the users will be strengthened (e.g., trained) at $j = 1$ (e.g., once when they join the organization), while 80% of the users will be applied a higher level of the safeguard $j = 2$ (e.g., attending training once per year).

3.4 Optimality analysis

We model *complete information* Nash CSGs, according to which both players know the game matrix, which contains the utilities of both players for each pure strategy profile. The utility function of the Defender is determined by the probability of failing to protect a user group and the indirect costs of the chosen application levels. We consider a *zero-sum* CSG, where the Attacker's utility is the opposite of the Defender's utility. The rationale behind the zero-sum CSG is that when the Defender is uncertain about the Attacker type, she considers the *worst case scenario*, which can be formulated by a zero-sum game where the Attacker can cause her *maximum damage*. The idea behind a zero-sum game like this is that the Attacker focuses on causing maximum corruption to cyberspace, while the Defender aims at minimizing the damage. Due to the Attacker's goal being conflicting to the Defender's objective, the application of game theory to study the selection of safeguards application levels is convenient. While in most security situations the interests of the players are neither in strong conflict nor in complete identity, the zero-sum game provides important insights into the notion of "optimal play", which is closely related to the *minimax theorem* [26].

In the zero-sum CSG,

$$\Gamma_0 = \langle \{d, a\}, \mathcal{L} \times \mathcal{U}, \{U_d, -U_d\} \rangle, \quad (11)$$

the Attacker's gain is equal to the Defender's security loss, and vice versa. We define the utility of the Defender in Γ_0 as

$$U_d^{\Gamma_0}(j, i) := -w_L L(j, i) - w_C C(j, i). \quad (12)$$

The first term of (12) is the expected loss of the Defender inflicted by the Attacker when attempting to compromise user group i , while the second term expresses the aggregated indirect cost of the safeguard application irrespective of the attacking strategy. Let $w_L, w_C \in [0, 1]$ are importance weights, which can facilitate the Defender with setting her preferences in terms of security loss, and indirect cost, accordingly.

For a mixed profile $(\boldsymbol{\delta}, \boldsymbol{\alpha})$, the utility of the Defender equals

$$\begin{aligned}
U_d^{I_0}(\boldsymbol{\delta}, \boldsymbol{\alpha}) &\stackrel{(7)}{=} \sum_{j \in \mathcal{L}} \sum_{i \in \mathcal{U}} U_d^{I_0}(j, i) \boldsymbol{\delta}(j) \boldsymbol{\alpha}(i) \\
&\stackrel{(12)}{=} \sum_{j \in \mathcal{L}} \sum_{i \in \mathcal{U}} [-w_L L(j, i) - w_C C(j)] \boldsymbol{\delta}(j) \boldsymbol{\alpha}(i) \\
&= -w_L \sum_{j \in \mathcal{L}} \sum_{i \in \mathcal{U}} L(j, i) \boldsymbol{\delta}(j) \boldsymbol{\alpha}(i) - w_C \sum_{j \in \mathcal{L}} C(j, i) \boldsymbol{\delta}(j).
\end{aligned} \tag{13}$$

As I_0 is a zero-sum game, the Attacker's utility is given by $U_a^{I_0}(\boldsymbol{\delta}, \boldsymbol{\alpha}) = -U_d^{I_0}(\boldsymbol{\delta}, \boldsymbol{\alpha})$. Since the Defender's equilibrium strategies maximize her utility, given that the Attacker maximizes her own utility, we will refer to them as *optimal strategies*.

As I_0 is a two-person zero-sum game with a finite number of actions for both players, according to Nash [27], it admits at least a NE in mixed strategies and saddle-points correspond to Nash equilibria as discussed in [28] (p. 42). The following result from [29], establishes the existence of a saddle (equilibrium) solution in the games, we examine and summarizes their properties.

Definition 6 (Saddle point of the CSG). *The I_0 Cyber Safeguard Game (CSG) admits a saddle point in mixed strategies, $(\boldsymbol{\delta}_{I_0}^{\text{NE}}, \boldsymbol{\alpha}_{I_0}^{\text{NE}})$, with the property that*

$$\begin{aligned}
- \boldsymbol{\delta}_{I_0}^{\text{NE}} &= \arg \max_{\boldsymbol{\delta} \in \Delta_{\mathcal{L}}} \min_{\boldsymbol{\alpha} \in \Delta_{\mathcal{U}}} U_d^{I_0}(\boldsymbol{\delta}, \boldsymbol{\alpha}), \forall \boldsymbol{\alpha}, \text{ and} \\
- \boldsymbol{\alpha}_{I_0}^{\text{NE}} &= \arg \max_{\boldsymbol{\alpha} \in \Delta_{\mathcal{U}}} \min_{\boldsymbol{\delta} \in \Delta_{\mathcal{L}}} U_a^{I_0}(\boldsymbol{\delta}, \boldsymbol{\alpha}), \forall \boldsymbol{\delta}.
\end{aligned}$$

Then, due to the zero-sum nature of the game, the minimax theorem [26] holds, i.e. $\max_{\boldsymbol{\delta} \in \Delta_{\mathcal{L}}} \min_{\boldsymbol{\alpha} \in \Delta_{\mathcal{U}}} U_d^{I_0}(\boldsymbol{\delta}, \boldsymbol{\alpha}) = \min_{\boldsymbol{\alpha} \in \Delta_{\mathcal{U}}} \max_{\boldsymbol{\delta} \in \Delta_{\mathcal{L}}} U_d^{I_0}(\boldsymbol{\delta}, \boldsymbol{\alpha})$.

The pair of saddle point strategies $(\boldsymbol{\delta}_{I_0}^{\text{NE}}, \boldsymbol{\alpha}_{I_0}^{\text{NE}})$ are at the same time security strategies for the players, i.e. they ensure a minimum performance regardless of the actions of the other. Furthermore, if the game admits multiple saddle points (and strategies), they have the ordered interchangeability property, i.e. the player achieves the same performance level independent from the other player's choice of saddle point strategy.

The minimax theorem [26] states that for zero-sum games, NE and minimax solutions coincide. Therefore, $\boldsymbol{\delta}_{I_0}^{\text{NE}} = \arg \min_{\boldsymbol{\delta} \in \Delta_{\mathcal{L}}} \max_{\boldsymbol{\alpha} \in \Delta_{\mathcal{U}}} U_a^{I_0}(\boldsymbol{\delta}, \boldsymbol{\alpha})$. This means that regardless of the strategy the Attacker chooses, NSP is the Defender's security strategy that guarantees a minimum performance.

Formally, the Defender seeks to solve the following LP:

$$\begin{aligned}
&\max_{\boldsymbol{\delta} \in \Delta_{\mathcal{L}}} \min_{\boldsymbol{\alpha} \in \Delta_{\mathcal{U}}} U_d^{I_0}(\boldsymbol{\delta}, \hat{i}) \\
&\text{subject to } \begin{cases} U_d^{I_0}(\boldsymbol{\delta}, 1) - \min_{\boldsymbol{\alpha} \in \Delta_{\mathcal{U}}} U_d^{I_0}(\boldsymbol{\delta}, \hat{i}) e \geq 0 \\ \vdots \\ U_d^{I_0}(\boldsymbol{\delta}, |\mathcal{U}|) - \min_{\boldsymbol{\alpha} \in \Delta_{\mathcal{U}}} U_d^{I_0}(\boldsymbol{\delta}, \hat{i}) e \geq 0 \\ \boldsymbol{\delta} e = 1 \\ \boldsymbol{\delta} \geq 0. \end{cases} \tag{14}
\end{aligned}$$

In this problem, e is a vector of ones of size $|\mathcal{U}|$.

3.5 Multiple Games Per Safeguard

Given that we have to allocate a budget in applying different safeguards, we may come across the challenge of not having enough monetary resources to select some of the equilibria of the CSG. Therefore, one has to derive the financial cost of equilibrium and assess its feasibility by comparing its financial cost to the available remaining budget. We refer to “remaining” budget as we expect that the Defender will have to select among a number of equilibria, one per safeguard, as we show later in this section.

To provide to the Defender a wider variety, in terms of financial cost, of equilibria per safeguard, we define a number of CSGs per safeguard. Each of these games has a different number of application levels available to the Defender. Aligned with [3], for each safeguard σ , we study $|\mathcal{L}|$ CSGs.

Definition 7. *To differentiate among different safeguards and implementations levels, we denote the CSG by $\Gamma_{\sigma,\lambda}$, where the safeguard σ can be applied up to $\lambda \in [0, |\mathcal{L}|]$.*

Note that we allow $\lambda = 0$ so that the Defender has the option to avoid selecting a safeguard should this violate some budget constraints. A Knapsack optimisation is used in the second phase of the model to select the equilibria, at most one per safeguard. In this way, we manage to have $|\mathcal{L}|$ NSPs per safeguard, each of a different financial cost. Each $\Gamma_{\sigma,\lambda}$ is a game where (i) Defender’s pure strategies correspond to consecutive application levels of safeguard σ starting always from 0 and including all levels up to λ and, (ii) Attacker’s pure strategies are the different targets akin to user groups. Figure 1 illustrates the different Cybersecurity Safeguards Games along with the utilities of the Defender.

Let $\delta_{\sigma,\lambda}^{NE}$ be the equilibrium of $\Gamma_{\sigma,\lambda}$ then

$$\delta_{\sigma,\lambda}^{NE} = [\delta_{\sigma,0}^{NE}, \delta_{\sigma,1}^{NE}, \dots, \delta_{\sigma,\lambda}^{NE}]. \quad (15)$$

Let $F(\delta_{\sigma,\lambda})$ be the financial cost of the safeguards plan $\delta_{\sigma,\lambda}$ which can be derived by summing the financial costs of all application levels $j \in \{1, 2, \dots, \lambda\}$ for safeguard σ contributed proportionally by using the corresponding probability from $\delta_{\sigma,\lambda}$, i.e., $\delta_{\sigma,j}$. Let $F(\sigma, j)$ denote the financial cost of safeguard σ then

$$F(\delta_{\sigma,\lambda}) = \sum_{j \in \{1, 2, \dots, \lambda\}} \delta_{\sigma,j} F(\sigma, j). \quad (16)$$

3.6 Investment in Nash Safeguards Plans

Let \mathcal{S} be the set of all available safeguards to the Defender. We can solve all $|\mathcal{S}| \times |\mathcal{L}|$ CSGs and derive a set of equilibria per safeguard σ represented as follows

$$\{\delta_{\sigma,1}^{NE}, \delta_{\sigma,2}^{NE}, \dots, \delta_{\sigma,|\mathcal{L}|}^{NE}\}. \quad (17)$$

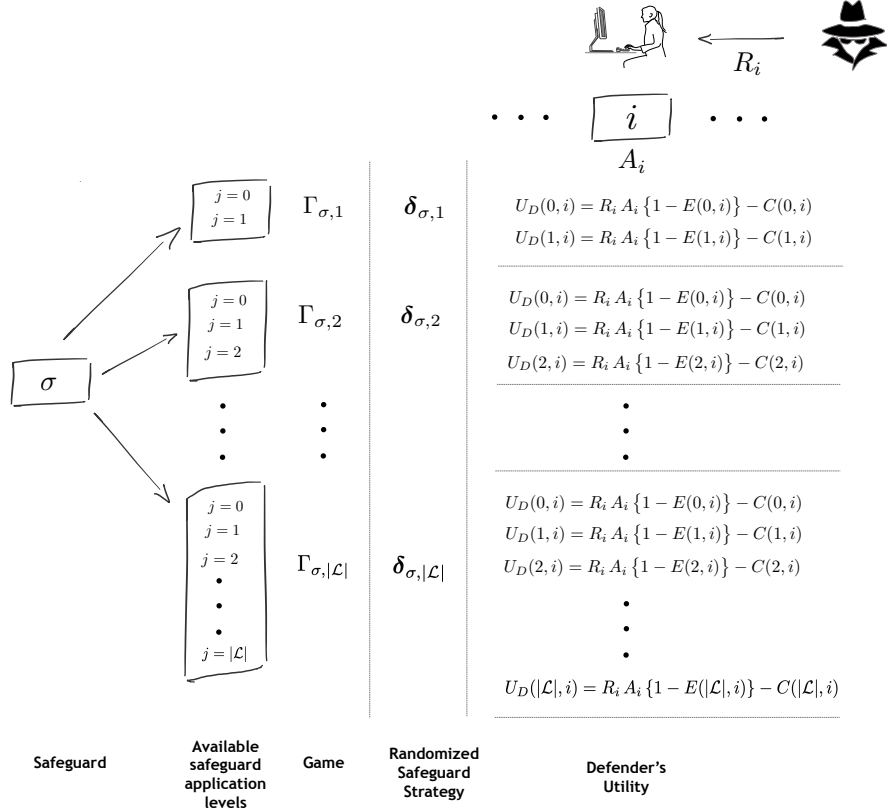


Fig. 1: Illustration of the safeguard-centered model of OST used to devise game-theoretic strategies for the Defender.

For all safeguards $\{1, 2, \dots, |\mathcal{S}|\}$ the following set of sets of equilibria, i.e., NSPs, is available

$$\left\{ \{ \delta_{1,0}^{NE}, \delta_{1,1}^{NE}, \dots, \delta_{1,|\mathcal{L}|}^{NE} \}, \{ \delta_{2,0}^{NE}, \delta_{2,1}^{NE}, \dots, \delta_{2,|\mathcal{L}|}^{NE} \}, \dots, \{ \delta_{|\mathcal{S}|,0}^{NE}, \delta_{|\mathcal{S}|,1}^{NE}, \dots, \delta_{|\mathcal{S}|,|\mathcal{L}|}^{NE} \} \right\}. \quad (18)$$

Optimal budget allocation in cybersecurity can be tackled by combinatorial optimization as previously investigated by Smeraldi and Malacaria [25]. We are concerned with the challenge of protecting multiple targets, in our case user groups, with the use of a number of NSPs that interact between them in different ways. In the following, we model the challenge of investing in these different NSPs in a way that at most one NSP per safeguard is chosen and the sum of financial costs of these NSPs fits an available cybersecurity budget. We have used 0-1 Knapsack Optimization to solve this problem. As opposed to the solution

provided in [3], we have chosen the objective function of the Defender to consider the sum of expected losses incurred from the different user groups being attacked. This is not to say that the proposed weakest-link model in [3] is not relevant anymore but we realize the potential risk to have all user groups targeted by the Attacker with the goal to maximize the collective damage over a number of assets rather than trying to compromise the most precious asset. We argue that such a goal to maximize the aggregated damage is more applicable in attacks like Advanced Persistent Threat, where the goal is to maximize the Defender's overall loss in a number of different ways.

The Knapsack Problem (KP) is an NP-hard problem [30]. There are several applications of KP such as resource distribution, investment decision making and budget controlling. In our model, we define KP as: Assuming that there is a knapsack with a maximum capacity of B , which represents the budget of the Defender. Given the set of all possible $|\mathcal{S}| \times |\mathcal{L}|$ NSPs shown in (18), each Knapsack candidate solution consists of at most $|\mathcal{S}|$ NSPs, one per each safeguard. Each NSP reduces, to some degree, the overall cyber risk of the organization as a result of reducing the individual risk on each user group. The problem is to select a subset of NSPs that maximize the knapsack profit without exceeding the maximum capacity of the knapsack. We define an optimal solution to our KP as $\Psi = \{\delta_{\sigma,\lambda}^{NE}\}, \forall \sigma \in \mathcal{S}, \forall \lambda \in \mathcal{L}$. A solution Ψ takes exactly one solution (i.e., equilibrium or cybersecurity plan) for each safeguard as a *policy for implementation/application*. To represent the cyber security investment problem, we need to expand the definitions for both expected loss L and effectiveness E to incorporate the solutions of the different CSGs. Hence, we expand L such that $L(\delta_{\sigma,\lambda}, i)$ is the expected loss inflicted by compromising user group i given the application of the plan $\delta_{\sigma,\lambda}$. We also expand E such that $E(\delta_{\sigma,\lambda}, i)$ is the efficacy that $\delta_{\sigma,\lambda}$ brings when applied to user group i . From equation (2) the expected loss on user group i when NSP $\delta_{\sigma,\lambda}$ is applied is given by

$$L(\delta_{\sigma,\lambda}, i) = R_i A_i [1 - E(\delta_{\sigma,\lambda}, i)] \quad (19)$$

A natural approach is the KP to seek a set of NSPs that minimize the aggregated expected risks across all user groups. We assume that each NSP may protect more than one user groups. We then seek optimal safeguards allocation for a series of user groups each of which can be protected by a different set of NSPs. The latter may not necessarily have an additive efficacy. The following illustrated example considers two NSPs and explains how we have decided to combine their efficacy in a single formula that we then use in KP formulation.

Example 2. By slightly abusing notation, assume two NSPs δ, δ' that mitigate 20% and 30% of the same user group risk, respectively. If the NSPs had additive efficacy the total expected loss on user group i when applying both δ, δ' equals $R_i A_i \{1 - \{E(\delta, i) + E(\delta', i)\}\} = R_i A_i (1 - 0.2 - 0.3) = 0.5 R_i A_i$. In this paper, we assume a more conservative expected loss mitigation function when combining two or more NSPs as follows $R_i A_i \left\{ \{1 - E(\delta, i)\} \{1 - E(\delta', i)\} \right\} = R_i A_i \cdot (1 - 0.2) (1 - 0.3) = R_i A_i \cdot 0.8 \cdot 0.7 = 0.56 R_i A_i$.

Given the above, if we represent the solution Ψ by the bitvector \mathbf{z} , we can then represent the 0-1 KP as

$$\begin{aligned} \min_{\mathbf{z}} \quad & \sum_{i=0}^{|\mathcal{U}|} A_i R_i \left\{ \prod_{\sigma=1}^{|\mathcal{S}|} \left\{ 1 - \sum_{j=0}^{\lambda} E(\delta_{\sigma,j}^{NE}, i) z_{\sigma,\lambda} \right\} \right\} \\ \text{s.t.} \quad & \sum_{\sigma=1}^{|\mathcal{U}|} \sum_{\lambda=0}^{|\mathcal{L}|} F(\delta_{\sigma,\lambda}) z_{\sigma,\lambda} \leq B, \\ & \sum_{\lambda=0}^{|\mathcal{L}|} z_{\sigma,\lambda} = 1, z_{\sigma,\lambda} \in \{0, 1\}, \forall \sigma = 1, 2, \dots, |\mathcal{S}|. \end{aligned} \quad (20)$$

where B is the available budget of the Defender to be spent in cyber safeguards and $z_{\sigma,\lambda} = 1$ holds when $\delta_{\sigma,\lambda}^{NE} \in \Psi$. Among KP solutions that all maximize the overall expected loss, we choose the solution with the lowest financial cost as this will be, in overall, the best advice to the defender producing same benefit for lower price.

4 Model Evaluation

We have developed the proposed models as part of the Optimal Safeguards Tool (OST) proposed in [1]. OST computes Nash Safeguards Plans as well as the Knapsack solutions. OST aims at offering realistic actionable advice to healthcare organizations. The following represents a case study based on Critical Internet Security (CIS) 17 Control “Implement a Security Awareness and Training Program”.

4.1 Use Case

User groups. Here, we assume a representative (non-exhaustive) set of three user groups, denoted by i , in decreasing order of *access privileges*:

- $i = 1$; **ICT**: The information and communication technology professionals responsible for the systems, networks and software. They set up digital systems, support staff who use them, diagnose and address faults, as well as set up and maintain security provisions. In addition to the ICT infrastructure, they may also interact with medical devices and electronic healthcare record systems. We consider the value of corresponding assets that can be affected by an attack on this group to be the highest possible, $A_1 = 100$ (e.g., \$100k). At the same time, due to limited interaction with the public, this is the group with the lowest visibility to attacks targeting the human, and as such we can consider it as lower risk, $R_1 = 0.2$.
- $i = 2$; **Clinical**: Nurses, doctors and other clinical staff have access to medical devices and electronic healthcare records. We consider the value of corresponding assets that can be affected by an attack on this group to be $A_2 = 50$ (e.g., \$50k). As a result of visibility due to interaction with the patients and presence on the hospital’s website, this group has a moderate risk, $R_2 = 0.5$.

- $i = 3$; **Administration**: Receptionists, medical secretaries and other administration roles involve access to electronic healthcare records. We consider the value of corresponding assets that can be affected by an attack on this group to be $A_3 = 25$ (e.g., \$25k). This group of users may have high interaction with the public and volume of email traffic (e.g., appointment requests) and as such high risk, $R_3 = 0.8$.

Control level \ Role	ICT Clinical Administration			
Low (once per year)	E	0.35	0.3	0.3
	C	1	30	10
Medium (twice per year)	E	0.6	0.5	0.5
	C	2	60	20
High (once per month)	E	0.8	0.7	0.7
	C	12	360	120

Table 2: Evaluation parameters for control CIS-17.4.

We have assumed a user group ratio of size 1:30:10 that loosely follows the corresponding breakdown of hospital workforce in the United States¹³: 81,790 computer, information system and security managers and analysts; 2,437,540 healthcare practitioners; 737,750 receptionists, healthcare record information clerks and other office and administrative support staff.

Control level \ Role	ICT Clinical Administration			
Low (Tests)	E	0.25	0.2	0.2
	C	1	30	10
Medium (Videos)	E	0.7	0.6	0.6
	C	2	60	20
High (Games)	E	0.6	0.5	0.5
	C	4	120	40

Table 3: Evaluation parameters for control CIS-17.6.

Safeguards. As safeguards, we have considered a representative pair from the SANS institute’s CIS-17 group of critical security controls¹⁴: CIS-17.4 “Update Awareness Content Frequently” and CIS-17.6 “Train Workforce on Identifying Social Engineering Attacks”. All values used in this case study, for these two safeguards, are presented in Tables 2 and 3.

¹³ https://www.bls.gov/oes/current/naics3_622000.htm.

¹⁴ <https://www.cisecurity.org/controls/implement-a-security-awareness-and-training-program/>.

For CIS-17.4, we set the frequency of completion of the updated training (once per year, twice per year, or once per month - i.e., 12 times per year) as the level of control. As indirect cost $C(j, i)$, we consider the total time spent in training by the employees in group i at application level j (in this case is *frequency*), which is proportionate to the size of the group and the frequency of the training. This time can be translated to some financial cost (in \$) resulting from loss of productive working hours. In this way, the indirect cost can be subtracted from the expected loss comprising the final utility value of the Defender in each cell of the game utility matrix.

For CIS-17.6, we set the nature of the work-based training (tests, videos, games) as the levels of control. Further, we set the corresponding efficacy values for each type roughly equivalent to their importance in helping predict user susceptibility to semantic social engineering attacks. Specifically, [15] has identified work-based security training with videos as the best predictor out of the three. In terms of efficacy values, we have differentiated slightly between groups based on our perceived rate of adoption of controls in each one. Specifically, we assume that adoption is greater for ICT than for clinical and administration employees. This is only for illustration purposes, so that the model can also take into account the group at each level of control. We also assume, further, that the primary indirect cost is employee time required, with a ratio of 1:2:4 for the three control levels.

4.2 Comparison with Alternative Defense Strategies

In the following, we analyze the proposed model in two phases; (i) the game-theoretic; and (ii) the 0-1 Knapsack optimization. The *first phase* evaluates different cybersecurity safeguard selection strategies using the utility table of the investigated Cybersecurity Safeguards Games (CSG) based on the use case discussed in the previous section. To evaluate our approach, we have created a simulated environment in Python which performs the attack sampling. For all comparisons performed, a sample size of 1,000 attacks was used. Such a sample is referred to as a *run* in the results. In the following, we present the results, where 25 runs have been performed in each case and the average Defender Utility (in \$) seen across the runs have been plotted.

More specifically, we have simulated $\Gamma_{\sigma,2}$ and $\Gamma_{\sigma,3}$ (please see Table 1 for the notation) for the two different safeguards presented in the use case, i.e., CIS 17.4 (denoted as $\sigma = 1$) and 17.6 (denoted as $\sigma = 2$). The games $\Gamma_{1,2}$, $\Gamma_{2,2}$ exhibit maximum safeguard application level of 2 (Medium), while the games $\Gamma_{1,3}$, $\Gamma_{2,3}$ are investigated up to application level 3 (High). Each CSG generates a utility table that we use to derive three different Defender application level selection strategies:

- *Nash* Safeguard Strategy (NSS), as described in Section 3 and computed using the the open source *Nashpy* Python library¹⁵.

¹⁵ <https://nashpy.readthedocs.io/en/stable/index.html>

- the *Weighted* Safeguard Strategy (WSS), which distributes the choice of a safeguard level over the weighted expected utility of the CSG by computing probability $\delta_{\sigma,j}$ of choosing application level j of safeguard σ as follows:

$$\delta_{\sigma,j} := \frac{\sum_{i=1}^{|\mathcal{U}|} U_d(j, i)}{\sum_{j=1}^{|\mathcal{L}|} \sum_{i=1}^{|\mathcal{U}|} U_d(j, i)}$$

- the *Cautious* Safeguard Strategy (CSS), which always prefers the *highest* application level of a safeguard.

Regarding adversarial strategies, we consider three profiles:

- the *Nash* Attacker who plays the Nash Attacking Strategy (NAS), presented in Section 3 and computed using the *Nashpy* Python library.
- a *Weighted* Attacker who plays the Weighted Attacking Strategy (WAS) by attacking a user group i with probability $\frac{A_i}{\sum_{i \in \mathcal{U}} A_i}$, i.e. the Attacker attacks the different user groups proportionally based on the asset values they have access to.
- the *Opportunistic* Attacker who uniformly chooses the different user groups to attack.

Figure 2 illustrates the performance of NSS against WSS and CSS in terms of average Defender’s utility over the 1,000 attacks for 25 runs. In all cases, we contrast between Attackers who follow NAS and WAS.

Nash Attacker. The results, in Figure 2(a), show that NSS outperforms both WSS and CSS when the Attacker chooses NAS. More specifically, the percentage improvement values, seen when choosing NSS, in comparison to WSS for the different games $[I_{1,2}, I_{1,3}, I_{2,2}, I_{2,3}]$ are [20.2%, 79.78%, 16%, 52.12%], respectively. Likewise, when choosing NSS over CSS, we observe improvement values of [34.48%, 87.07%, 28.57%, 62.26%] for the different games $[I_{1,2}, I_{1,3}, I_{2,2}, I_{2,3}]$, respectively.

Remark 3. These results demonstrate an average improvement of approximately 42% of NSS over WSS and 53% over CSS.

Comparably, the smallest average improvement for NSS over WSS is around 16% when playing Control 17.6 at the maximum application level of 2 ($\lambda = 2$). Likewise, the minimum improvement of NSS over CSS, approximately equal to 28%, is for the same control and $\lambda = 2$. On the other hand, the maximum improvement seen in NSS over CSS is approximately 87%, where the maximum improvement over CSS does not exceed 80%, for Control 17.4. and $\lambda = 3$.

One of the primary reasons why naive-deterministic safeguard selection approaches perform poorly against the Nash Defending strategy is that they fail to incorporate the opponent’s strategies. At the same time, we have considered CSG as a zero-sum game. The class of zero-sum games offers a degree of freedom as it can be shown that assuming that the adversary’s intentions are exactly opposite to the defender’s assets, i.e., the Attacker seeks to cause maximum damage, any other incentive of the Attacker can only improve the Defender’s situation [31].

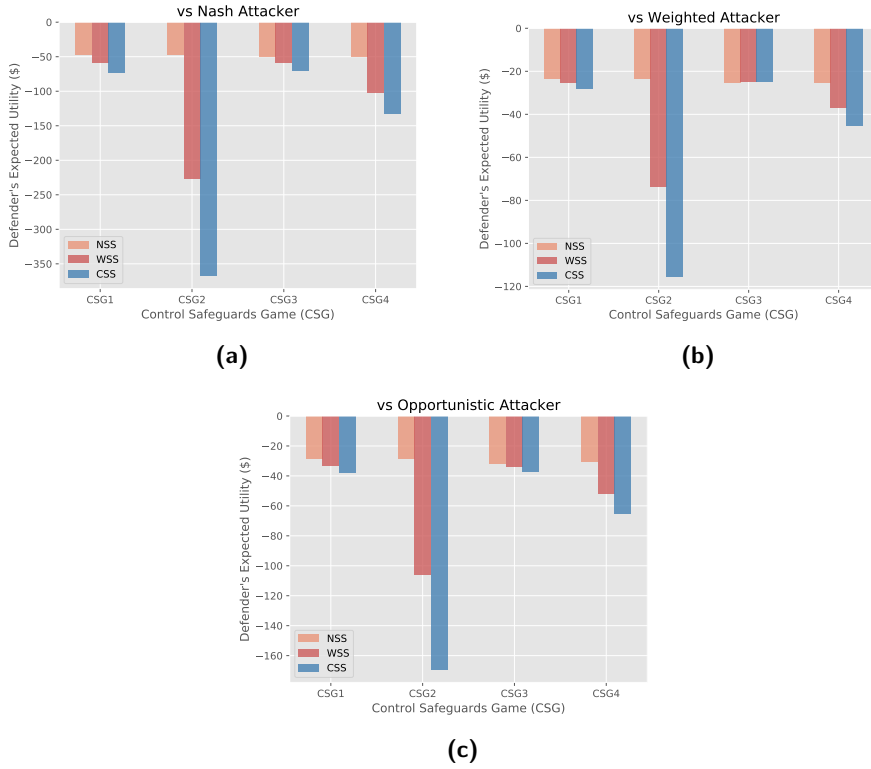


Fig. 2: Game-theoretic optimization results: Average Utility of the Defender over 1,000 attacks for 25 runs. for various CSGs.

Weighted Attacker. When the Weighted Attacking Strategy is simulated, the results demonstrate that NSS has higher efficacy over WSS and CSS apart from one game $I_{2,2}$ in which both WSS and CSS perform approximately 2% better than NSS (Figure 2(b)). This difference is negligible making NSS being at least as good as the rest of the Defending strategies in all investigated games. Despite the performance of NSS in $I_{2,2}$, for the rest of the games, NSS performs significantly better than WSS and CSS. The percentage improvement values, seen when choosing NSS, in comparison to WSS and CSS for $[I_{1,2}, I_{1,3}, I_{2,2}, I_{2,3}]$ are $[7.34\%, 70.25\%, -2.25\%, \text{ and } 32.29\%]$ and $[15.1\%, 80.44\%, -2.1\%, \text{ and } 44.79\%]$, respectively.

Remark 4. These results demonstrate an average improvement of approximately 28% of NSS over WSS and 34% over CSS.

The smallest average improvements for NSS over WSS and CSS are approximately 7% (in $I_{1,2}$) and 15% ($I_{1,2}$), respectively, and the maximum average improvement values are 70% (in $I_{1,3}$) and 80% (in $I_{1,3}$).

Opportunistic Attacker. Finally, when the Opportunistic Attacking Strategy is simulated, the results demonstrate that NSS has higher efficacy over WSS and CSS (Figure 2(c)). The percentage improvement values, seen when choosing NSS, in comparison to WSS and CSS for $[I_{1,2}, I_{1,3}, I_{2,2}, I_{2,3}]$ are $[13.3\%, 74.24\%, 5.4\%, \text{ and } 40.8\%]$ and $[23.73\%, 83.4\%, 12.33\%, \text{ and } 52.51\%]$, respectively.

Remark 5. These results demonstrate an average improvement of approximately 33% of NSS over WSS and 43% over CSS.

The smallest average improvements for NSS over WSS and CSS are approximately 5% (in $I_{2,2}$) and 12% ($I_{2,2}$), respectively, and the maximum average improvement values are 74% (in $I_{1,3}$) and 83% (in $I_{1,3}$).

We notice that the highest improvements among the three different Attacking strategies are introduced by the first scenario where Nash Attacker is simulated. This was anticipated as at the Nash Equilibrium the Defender does the best against a rational Attacker. Between the results for Weighted and Opportunistic Attacker, NSS is more efficient against an Opportunistic Attacker than a Weighted one.

4.3 Analysis of the Investment Problem

The Knapsack optimization phase investigates the optimal investment in Nash Safeguards Plans (NSPs) given a budget B (for details refer to section 3.6). The Knapsack takes as input every NSP generated in the game-theoretic phase and recommends a single solution which minimizes the aggregated risk of all the user groups while satisfying the investment budget constraint. This is different to the weakest-link model investigated by [3].

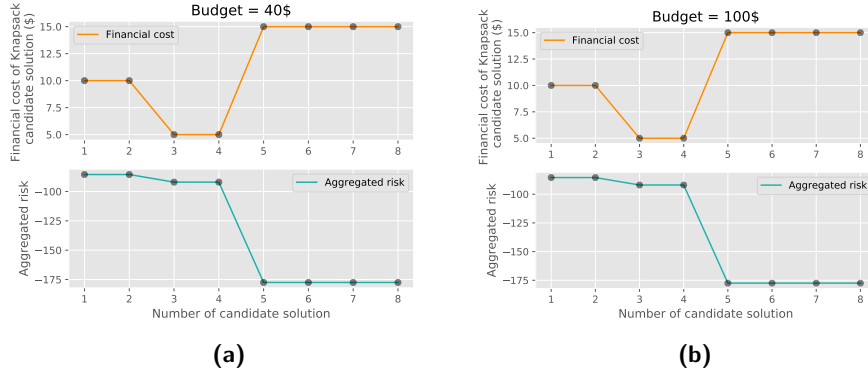


Fig. 3: Knapsack selection over available candidate solutions.

Figure 3 presents the financial cost and aggregated risk overall users for each Knapsack candidate solution, i.e., a combination of NSPs for two different

available budget values. We notice that there are multiple Knapsack optimal solutions, which are candidate solutions number 5, 6, 7 and 8. In the presence of multiple optimal solutions, the Knapsack solver, we have implemented, chooses the first option. For both budgets 40 and 100, the Knapsack optimization recommends investing in both CIS controls 17.4 and 17.6 at application level 1 i.e., Low (once per year) and Low (Tests), respectively. Note here that the small size of the use case effectively prohibits high variability of the parametric values, which led to the selection of only two types.

Note that the plots in Knapsack optimization only present the candidate solutions for the Nash Defender against Nash Attacker, in contrast to the plots in game-theoretic phase, (Figure 2), which presents all three Defender strategies. This choice was made due to the Knapsack optimization not involving the notion of CSG. As a result of this, it does not optimize the overall indirect cost of safeguards when choosing NSPs which has been done in the previous phase. In addition, Knapsack does not consider the behavior of the Attacker characterizing all adversarial strategies as irrelevant to the Knapsack objective function.

5 Conclusions

In this paper, we have presented an approach, extending the previous work [3], which implements a cybersecurity safeguards selection model along with game-theoretic and Knapsack optimization tools. We have evaluated our model in a healthcare use case using the CIS group 17 controls which attend to implementation of security awareness and training programs for employees. The simulation results demonstrate that the Nash Safeguard Strategy comfortably outperforms common-sense selection strategies, such as the Weighted and Cautious, in terms of Defender’s expected utility over a large number of attacks. This work is our step towards integrating the developed Optimal Safeguards Tool (OST) within cybersecurity risk management and investment environments.

An interesting extension to this work would be to capture the real-world uncertainty about an Attacker’s type, for example considering a Bayesian game of application level selection. Furthermore, we plan to bring together several objective functions for Knapsack to compare the performance of the investment strategies. As the next steps, we aim at creating a use case with greater size of safeguards in collaboration with healthcare organizations. We also aim at using the well-known repository of cybersecurity safeguards like the 20 CIS controls or a list of Privacy Enhancing Technologies (PETs) to support our research.

6 Notes

Caxton told us that the implementation of some sub-controls satisfies the implementation of other sub-controls of the same control. In OST terms, this will mean that an equilibrium (solution to a sub-control game) can cover more than one sub-controls. For instance $Q_{1.1.1}$ can replace all equilibria of sub-controls that can be implemented by sub-control 1.1. This gives us less equilibria per

control than the number of sub-controls which makes the possible Knapsack solutions fewer than before.

Another issue is the following. Having a predetermined budget per control and then use this budget to solve the OST framework is fine. However, the question is “How do we compute this budget?” Let us assume that we have an overall budget B . How do we split this budget among the different controls so that we can then apply OST? Given that CIS advises to use all controls for IG3 (hospital) one approach could be the following:

Let us assume that there are \mathcal{C} controls and each of them has a number of sub-controls denoted by $|\mathcal{C}_\sigma|$ for control σ . For each subcontrol, we define a game Γ

7 Acknowledgments

We thank the reviewers for their valuable feedback and comments.

Emmanouil Panaousis is partially supported by the European Commission as part of the CUREX project (H2020-SC1-FA-DTS-2018-1 under grant agreement No 826404). The work of Christos Laoudias has been partially supported by the CUREX project (under grant agreement No 826404), by the European Union’s Horizon 2020 research and innovation programme (under grant agreement No 739551 (KIOS CoE)), and from the Republic of Cyprus through the Directorate General for European Programmes, Coordination and Development.

References

1. Farnaz Mohammadi, Angeliki Panou, Christoforos Ntantogian, Eirini Karapistoli, Emmanouil Panaousis, and Christos Xenakis. CUREX: seCUre and pRivate hEalth data eXchange. In *IEEE/WIC/ACM International Conference on Web Intelligence*, volume 24800, pages 263–268, 2019.
2. Arun Vishwanath, Loo Seng Neo, Pamela Goh, Seyoung Lee, Majeed Khader, Gabriel Ong, and Jeffery Chin. Cyber hygiene: The concept, its measure, and its initial tests. *Decision Support Systems*, page 113160, 2019.
3. Andrew Fielder, Emmanouil Panaousis, Pasquale Malacaria, Chris Hankin, and Fabrizio Smeraldi. Decision support approaches for cyber security investment. *Decision Support Systems*, 86:13–23, 2016.
4. Clemens Scott Kruse, Benjamin Frederick, Taylor Jacobson, and D Kyle Monticone. Cybersecurity in healthcare: A systematic review of modern threats and trends. *Technology and Health Care*, 25(1):1–10, 2017.
5. Òscar Solans Fernández, Carlos Gallego Pérez, Francesc García-Cuyàs, Núria Abdón Giménez, Manel Berruezo Gallego, Adrià Garcia Font, Miquel González Quintana, Sara Hernández Corbacho, and Ester Sarquella Casellas. Shared medical record, personal health folder and health and social integrated care in catalonia: ICT services for integrated care. In *New Perspectives in Medical Records*, pages 49–64. Springer, 2017.
6. Lynne Coventry and Dawn Branley. Cybersecurity in healthcare: A narrative review of trends, threats and ways forward. *Maturitas*, 113:48–52, 2018.

7. David Kotz, Carl A Gunter, Santosh Kumar, and Jonathan P Weiner. Privacy and security in mobile health: a research agenda. *Computer*, 49(6):22–30, 2016.
8. George Loukas. *Cyber-physical attacks: A growing invisible threat*. Butterworth-Heinemann, 2015.
9. Luanne Billingsley and Shawn A McKee. Cybersecurity in the clinical setting: Nurses’ role in the expanding internet of things. *The Journal of Continuing Education in Nursing*, 47(8):347–349, 2016.
10. Ryan Heartfield and George Loukas. Detecting semantic social engineering attacks with the weakest link: Implementation and empirical evaluation of a human-as-a-security-sensor framework. *Computers & Security*, 76:101–127, 2018.
11. Jose M Such, Pierre Ciholas, Awais Rashid, John Vidler, and Timothy Seabrook. Basic cyber hygiene: Does it work? *Computer*, 52(4):21–31, 2019.
12. Leming Zhou, Bambang Parmanto, Zakiy Alfikri, and Jie Bao. A mobile app for assisting users to make informed selections in security settings for protecting personal health data: development and feasibility study. *JMIR mHealth and uHealth*, 6(12):e11210, 2018.
13. Steven Furnell, Peter Sanders, and Matthew Warren. Addressing information security training and awareness within the european healthcare community. *Studies in health technology and informatics*, 43:707–711, 1997.
14. Ryan Heartfield and George Loukas. A taxonomy of attacks and a survey of defence mechanisms for semantic social engineering attacks. *ACM Computing Surveys*, 48(3):37, 2016.
15. Ryan Heartfield, George Loukas, and Diane Gan. You are probably not the weakest link: Towards practical prediction of susceptibility to semantic social engineering attacks. *IEEE Access*, 4:6910–6928, 2016.
16. Rick Wash and Molly M Cooper. Who provides phishing training?: Facts, stories, and people like me. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 492. ACM, 2018.
17. Lawrence A Gordon and Martin P Loeb. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)*, 5(4):438–457, 2002.
18. Andrew Fielder, Sandra König, Emmanouil Panaousis, Stefan Schauer, and Stefan Rass. Risk assessment uncertainties in cybersecurity investments. *Games*, 9(2):34, 2018.
19. Andrew Fielder, Emmanouil Panaousis, Pasquale Malacaria, Chris Hankin, and Fabrizio Smeraldi. Game theory meets information security management. In *IFIP International Information Security Conference*, pages 15–29. Springer, 2014.
20. Anna Nagurney, Patrizia Daniele, and Shivani Shukla. A supply chain network game theory model of cybersecurity investments with nonlinear budget constraints. *Annals of operations research*, 248(1-2):405–427, 2017.
21. Shaun S Wang. Integrated framework for information security investment and cyber insurance. *Pacific-Basin Finance Journal*, 57:101173, 2019.
22. Michail Chronopoulos, Emmanouil Panaousis, and Jens Grossklags. An options approach to cybersecurity investment. *IEEE Access*, 6:12175–12186, 2017.
23. Fabio Martinelli, Ganbayar Uganbayar, and Artsiom Yautsiukhin. Optimal security configuration for cyber insurance. In *IFIP International Conference on ICT Systems Security and Privacy Protection*, pages 187–200. Springer, 2018.
24. Michael E Whitman and Herbert J Mattord. *Principles of information security*. Cengage Learning, 2011.

25. Fabrizio Smeraldi and Pasquale Malacaria. How to spend it: optimal investment for cyber security. In *Proceedings of the 1st International Workshop on Agents and CyberSecurity*, page 8. ACM, 2014.
26. J. Von Neumann and O. Morgenstern. *Theory of games and economic behavior (60th anniversary commemorative edition)*. Princeton university press, 2007.
27. J.F. Nash. Equilibrium points in n-person games. In *Proc. of the National Academy of Sciences*, pages 48–49, 1950.
28. T. Alpcan and T. Basar. *Network security: a decision and game-theoretic approach*. Cambridge University Press, 2010.
29. T. Basar and G. J. Olsder. *Dynamic noncooperative game theory*. London Academic press, 1995.
30. David Pisinger. Where are the hard knapsack problems? *Computers & Operations Research*, 32(9):2271–2284, 2005.
31. Stefan Rass and Sandra König. Password security as a game of entropies. *Entropy*, 20(5):312, 2018.