

# Comparing Decision Support Approaches for Cyber Security Investment

Andrew Fielder\*, Emmanouil Panaousis†, Pasquale Malacaria‡,  
Chris Hankin\*, and Fabrizio Smeraldi‡

\* Imperial College London, UK

† University of Brighton, UK

‡ Queen Mary University of London, UK

**Abstract**—When investing in cyber security resources, information security managers have to follow effective decision-making strategies. We refer to this as the cyber security investment challenge. In this paper, we consider three possible decision-support methodologies for security managers to tackle this challenge. We consider methods based on game theory, combinatorial optimisation and a hybrid of the two. Our modelling starts by building a framework where we can investigate the effectiveness of a cyber security control regarding the protection of different assets seen as targets in presence of commodity threats. In terms of game theory we consider a 2-person control game between the security manager who has to choose among different implementation levels of a cyber security control, and a commodity attacker who chooses among different targets to attack. The pure game theoretical methodology consists of a large game including all controls and all threats. In the hybrid methodology the game solutions of individual control-games along with their direct costs (e.g. financial) are combined with a knapsack algorithm to derive an optimal investment strategy. The combinatorial optimisation technique consists of a multi-objective multiple choice knapsack based strategy. We compare these approaches on a case study that was built on SANS top critical controls. The main achievements of this work is to highlight the weaknesses and strengths of different investment methodologies for cyber security, the benefit of their interaction, and the impact that indirect costs have on cyber security investment.

## I. INTRODUCTION

One of the biggest issues facing organisations today is how they are able to defend themselves from potential cyber attacks. The range and scope of these unknown attacks create the need for organisations to prioritise the manner in which they defend themselves. With this each organisation needs to consider the threats that they are most at risk from and act in such a way so as to reduce the vulnerability across as many relevant vulnerabilities as possible. This is a particularly difficult task that many Chief Information Security Officers (CISOs) are not confident in achieving, with only 24% of CISOs considering themselves very confident of preventing attacks according to Deloitte and NASICO [1]. In this report 86% of CISOs were concerned that the biggest issue facing their ability to successfully defend their systems was down to a “lack of sufficient funding”.

It is this perceived lack of sufficient funding that this work wishes to address. From our work with Small-Medium Enterprises (SMEs), we have identified that they are heavily restricted with the available funding for cyber security, gen-

erally working with a fixed budget with little to no additional funding being made available for cyber security purposes. It is generally perceived that this budget is insufficient for them to cover all of the vulnerabilities that their system may have. In this way organisations have to make trade-offs with regard to how they defend their systems.

When an organisation is making the decisions regarding the defence of their network, they generally have to consider two critical factors, the cost of implementing a particular defence and the impact that defence has on the business. The first of these has been discussed, stating that a company can only implement defences that are within their limited budget, considered the *Direct Cost* of the defence. However we question whether the apparently most optimal defence based solely on direct costs is the correct choice for an organisation. The reason behind this lies with the second criteria, such that the manner in which a defence is implemented will likely have some effect on either the operation of the system or the users of the system. These effects may cause a reduction in the speed that tasks can be performed by users or by a weakening of the defence caused by users circumventing the controls in order to more easily perform their required tasks. We consider that these factors create additional *indirect costs* for implementing a given defence. These two factors are at the core of our work into the decision support of *how to use the limited financial budget available to best protect against cyber attacks*.

### A. Contributions

This work proposes a two stage model designed to aid security managers with decisions regarding the optimal allocation of a cyber security budget. And it provides a detailed analysis of the model that was first proposed by Panaousis et al. [2].

We analyse the two stages of the model by first presenting an overview of the environment from which we define the problem of *cyber security investment*, identifying a unique manner for reasoning about the *targets* that a potential attacker has, and the defences associated with those targets. This is done by considering the physical location of a data asset, which needs to be protected, as well as the degree to which a particular defence, herein referred to as a *control*, is implemented.

We use the above environment to formulate control-games, which analyse how well each given control performs at

different *degrees of implementation* (i.e. levels). We compute the Nash Equilibrium condition in control-games, and we motivate the trade-offs required with the indirect costs. The Nash Equilibrium of a control-game dictates the most efficient manner, in which, a control should be implemented.

The solution to each control-game alone is insufficient in dictating the optimal allocation of an organisation's cyber security budget. So to identify the best way to allocate a budget, we formalise the problem as a *multi-objective multiple choice Knapsack problem*, as proposed in [3].

We motivate the use of this methodology by comparing the two-stage model to two alternative methods. Firstly, we model the scenario as an one shot game that aims to optimise the defense including *direct costs*, and secondly a Knapsack problem that considers only pure strategies for each control level including indirect costs. In both cases we highlight where our proposed method is able to outperform alternative methods.

## B. Outline

Section III-A introduces *targets* i.e. vulnerability and the associated potential loss, *controls*, effectiveness of controls on the potential loss as well as indirect costs associated to implement controls at different levels. In Section III-B those notions are used to build a game model, and sections III-C to III-H develop a basic analysis of these games. Sections III-I to III-K present a toy 2x2 game example with a single control with two implementation levels and two targets. This aims to provide a feel for these games and what elements determine the equilibria.

In Section IV we discuss general games where the defender has available a single control and the attacker can choose a set of vulnerabilities: we provide an interpretation of mixed strategies for these games and explain how these game solutions can be used with a Knapsack algorithm. Section IV-A illustrates why a single game comprising all possible controls doesn't provide good security guarantees and so is not a suitable investment methodology. Section IV-B introduces the particular Knapsack we think is relevant for this work, that is 0-1 Multiple Choice, Multi-Objective Knapsack. The items relevant to this Knapsack are game solutions: this constitutes the hybrid methodology presented in section IV-C.

In Section V we develop a case study based on the SANS top critical controls. We illustrate a mapping from the SANS controls' descriptions into our framework and based on this mapping we run a simulation to calculate the security effectiveness of the two Knapsack methodologies. The simulation shows that the hybrid solution is more flexible and provides higher security guarantee.

## II. RELATED WORK

The work of Anderson [4] considers the traps that defenders may fall into in finding bugs and protecting their systems, where it only needs to be a single unseen vulnerability that exposes the whole of a network. The approach taken in this work is to model attackers using *commodity* attacks against SMEs,

where the attacker is using commonly available attack vectors against known defendable vulnerabilities. While this doesn't negate the possibility of zero-day vulnerabilities, it removes the expectation that it is in the best interest of either player to invest heavily in order to either find a new vulnerability or be able to protect against these unknown vulnerabilities. In his work, Anderson considers that the management of information security is a more difficult problem than initially considered as there are often deeper issues, such as politics, that need to be addressed.

Important to the modelling is the concept that the defenders have to attempt to *defend everywhere*. This is due to the fact that attackers can strike anywhere they wish. We can highlight this observation by assuming that the defence provided by an optimal budget allocations can only be considered as strong as the defence of the *weakest target*. This is because the weakest target is at most risk from an attacker who can potentially attack anywhere. Our approach is quite different to Anderson's as we focus on developing *cyber security decision support tools* to assist security managers on how to spend a cyber security budget in terms of different controls acquisition and implementation.

Our work has been partially influenced by a recent contribution within the field of physical security [5], where the authors address the problem of finding an optimal defensive coverage. The latter is defined as the one maximising the worst-case payoff over the targets in the potential attack set. One of the main ideas of this work we adopt here is that *the more we defend the less rewards the attacker receives*.

Alpcan [6] (p. 134) discusses the importance of studying the quantitative aspects of risk assessment with regard to cyber security in order to better inform decisions makers. This kind of approach is taken in this work where we provide an analytical method for deciding the level of risk introduced by different vulnerabilities, and the impact that different security controls have in mitigating these risks. By studying the incentives for risk management, Alpcan [7] develops a game-theoretic approach that optimises the investment in security across different autonomous divisions of an organisation, where each of the divisions is seen as a greedy entity. Furthermore, Alpcan et al. examine in [8] security risk dependencies in organisations, and they propose a framework which ranks the risks by considering the different complex interactions. This rank is dictated by an equilibrium that is derived by a Risk-Rank algorithm.

Saad et al. [9] model cooperation among autonomous parts of an organisation that have dependent security assets, and vulnerabilities for reducing overall security risks, as a cooperative game. In [10] Bommannavar et al. capture risk management in a quantitative framework which aids decision makers upon allocation of security resources. The authors use a dynamic zero-sum game to model the interactions between attacking and defending players. A Markov model, in which states represent probabilistic risk regions and transitions, has been defined. The authors use Q-learning to cope with scenarios when players are not aware of the different Markov model

parameters.

Fielder et. al. [11] investigate *how to optimally allocate the time for security tasks for system administrators*. This work identifies how to allocate the limited amount of time, which a system administrator has, to work on the different security related tasks for an organisation's data assets.

One of the initial works studying the way to model investment in cyber security is published by Gordon and Loeb [12]. The authors consider the optimum level of investment given different levels of information security level. The authors propose a model in which for any given vulnerability there are different levels of information security that can be implemented, where a higher level of information security will cause the expected loss to that particular vulnerability to drop. This is modelled as a function of the security level's responsiveness to an increasing vulnerability in reducing loss. In our model, here, we consider a single value for a vulnerability, and then for each control there are a number of levels of implementation, which represent the information security levels proposed by Gordon and Loeb. The main message of this work is that to maximise the expected benefit from information security investment, an organisation should spend only a small fraction of the expected loss due to a security breach.

The work published in [13] examines the weakest target game which refers to the case where an attacker is always able to compromise the system target with the lowest level of defence, and not to cause any damage to the rest of the targets. The game-theoretic analysis, which the authors have undertaken, shows that the game leads to a conflict between pure economic interests and common social norms. While the former are concerned with the minimisation of cost for security investments, the latter imply that higher security levels are preferable. Cavusoglu et. al. [14] compare a decision theory based approach to game-theoretic approaches for investment in cyber security. Their work compares a decision theory model to both simultaneous and sequential games. The results show that the expected payoff from a sequential game is better than that of the decision theoretic method, however, a simultaneous game is not always better.

Recent work on cyber security spending has been published by Smeraldi and Malacaria [3]. The authors identify the optimum manner in which investments can be made in a cyber security scenario given that the budget allocation problem is most fittingly represented as a multi-objective Knapsack problem. Cremonini and Nizovtsev, in [15], have developed an analytical model of the attacker's behaviour by using cost-benefit analysis, and therefore considering rewards and costs of achieving different actions. One issue that we factor into this work is that *security comes at a cost that is greater than that of the price of implementing a policy*. Al-Humaigani and Dunn proposed a model of Return on Security Investment (ROSI) [16], where the authors define the return on investment of an attack as a function of eleven factors, which comprise of *direct costs* for implementing a security tool, *indirect costs* of having that *security tool* in place, as well as the cost to the company should there be a breach (i.e. *damage*). Wang

et al. note that game-theoretic approaches to cyber security suffer from the fact that "the rationality of hackers is hard to be captured by a model, because they may be motivated by different value systems" [17]. While the authors do not argue on the rationality of the attacker, but the idea that imposing on them a similar set of values as a defender is not adequate. Previous work we have conducted in this area notes that the reward for the attacker is in line with the loss of the defender by the way of an affine transformation [11]. This was done to represent the loss of value that an attacker gets from the data that has been stolen, when compared to the value to the defender.

Demetz and Bachlechner [18] provide a survey of models that have been proposed for the study of economic viability of tools for security policy and configuration. The authors identify a series of requirements that a security investment tool should contain. We compare our approach to the conditions set out by Demetz and Bachlechner's:

- Financial Measures - The optimisation method looks to take into consideration the financial constraints of the organisation and identifies what should be purchased given the range of possible budgets. An organisation is then able to select an appropriate set of controls given their financial constraints and threat tolerance levels.
- Non-Financial Measures - One of the key features of a multi-level model of a control implementation is that, it is possible to clearly identify non-financial measures such as System Performance or Staff Morale that will be impacted as a result of implementing certain controls or extending the reach that some controls have (such as surveillance).
- Support One-Time Costs and Benefits - One of the direct costs is Capital Costs, where the capital costs for an organisation incurred for implementing the control.
- Support Running Costs and Benefits - As with the One-Time costs, the model supports the inclusion of running costs into these direct costs as well, primarily as labour costs. Additionally any non-financial costs incurred in the actual implementation are represented in the indirect costs.
- Does not explicitly consider Attacks - The game-theoretic model that is presented here is defined in such a way that it allows for the representation of a control that is capable of mitigating any number of attacks. However, this goal focuses on policy and configuration that will not only protect against attacks, but will also work with security breaches that are not related to the attacks.
- Consider the Network Effects of Investments - Within the scope of the model considered, there is no direct consideration of the additional benefits to other organisations from the implementation of a given security policy given by our model.

While the work we provide covers many of the aspects designated to be an effective security tool, it notably lacks the aspects that relate to non-attack related issues of security

such as unintentional loss or network benefits.

### III. MODEL DEFINITION

In this section we use game theory to model the interactions between two players; Defender and Attacker. Defender is the cyber security manager in an SME, and her overall objective is to defend the organisation's assets from cyber theft, mitigate any potential business disruption, and maintain the organisation's reputation.

On the other hand, Attacker is a cyber hacker who tries to subvert the system to her own end, by launching commodity cyber-attacks against the organisation Defender is working for. Commodity cyber-attacks are based on capabilities and techniques that are available on the internet, where the attack tool can be purchased therefore the adversaries do not develop the attack themselves, and they can only configure the tool for their own use.

First, we present the cyber security model and environment of the game, and introduce the two players. The model of environment and game presented here were initially proposed in [2]. We then describe the sets of their pure strategies and the payoffs associated with each of them. We also discuss how mixed strategies are represented in this game, and we provide the expected payoffs of Defender and Attacker given any pair of mixed strategies. Fig. 1 illustrates our environment.

#### A. Environment

In our model, Defender work as a cyber security manager in an SME with an available cyber security budget  $B$ , and she wants to invest in implementing cyber security controls to protect the organisation's data assets against *commodity attacks*.

Each control can be implemented at a different level. Note that the higher the level the greater the degree to which the control is implemented. After its implementation, each control brings some security benefits to the system, but it is also associated with indirect and direct costs. The challenge Defender has to address is how to decide upon implementation of the different cyber security controls against *commodity attacks*, given a limited budget  $B$ , and other preferences the organisation has in terms of risks and indirect costs. In the following we discuss the different components of the model, and we define appropriate terminology and notations, which are consistent throughout this article.

1) *Asset Depth*: We define the *depth* of a data asset as the location of this asset within the organisation's structure following the rule: *the higher the depth is, the more confidential data the asset holds*. In other words, a depth determines the importance of the data asset that the organisation loses if a commodity attack (herein referred to as attack) is successful. In this paper, we specify that data assets that located at the same are depth, worth the same value to Defender's firm.

2) *Cyber Security Targets*: We denote the set of cyber security targets within an organisation by  $\mathcal{T} := \{t_i\}$ , the set of vulnerabilities threatened by commodity attacks by  $\mathcal{V} := \{v_z\}$ ,

and the set of depths by  $\mathcal{D} := \{d_x\}$ . A *cyber security target* is defined as a (*vulnerability, depth*) pair; formally

$$t_i := (v_z, d_x). \quad (1)$$

And it abstracts any data asset, located at  $d_x$ , that an attack threatens to compromise by exploiting  $v_z$ . We specify that data assets located at the same depth and having the same vulnerabilities are abstracted by the same target.

Each target is associated with an impact value which expresses the level of damage incurred to Defender's organisation when Attacker succeeds in their attack against that target. The different impact factors can be *data loss*, *business disruption*, and *reputation damage*. Each impact factor depends on the depth  $d_x$  that the attack targets.

Furthermore, there is a *threat value* for each target. This can account, for instance, for the frequency of attacks launched against that target. Each software weakness (we use the terms weakness and vulnerability interchangeably) has some factors that can determine an overall score.

Let  $I: \mathcal{T} \rightarrow \mathbb{Z}^+$  be the random variable which takes targets  $t_i$  to the impact value that the compromise of  $t_i$  will have to the organisation, and let  $T: \mathcal{T} \rightarrow \mathbb{Z}^+$ , be the random variable which takes target  $t_i \in \mathcal{T}$  to its threat value. Given Definition(1) note that  $I(t_i)$  depends on the depth  $d_x$ , and  $T(t_i)$  depends on the vulnerability  $v_z$ .

3) *Cyber Security Controls*: A *cyber security control* is the defensive mechanism that Defender can be put in place to alleviate the risk from one or more attacks by reducing the probability of these attacks successfully exploiting vulnerabilities.

Defender chooses to implement a control at a certain level for their organisation. We define the set of implementation levels of a control as  $\mathcal{L} := \{l_j\}$ . The higher the level the greater the degree to which the control is implemented.

Note that we abuse notation by setting  $l_j = l, t_i = t, v_z = v$ , and  $d_x = d$ .

An implementation level  $l$  has a *degree of vulnerability mitigation* on each target. This is determined by the efficacy, in terms of cyber defence, of  $l$  on this target. For a pair  $(l, t)$ , which represents the level of implementation of a particular control, we define the random variable  $E: \mathcal{L} \times \mathcal{T} \rightarrow [0, 1)$ , which takes a pair of  $(l, t)$  to the efficacy value of  $l$  on  $t$ . Here, we have postulated that  $E(l, t) \neq 1$  due to the existence of 0-day vulnerabilities that Attacker has the potential to exploit. Assume Defender implements a control at  $l$  that has efficacy  $E(l, t)$  on  $t$ .

We define the *cyber security loss* random variable

$$S(l, t) = I(t) T(t) [1 - E(l, t)].$$

This is the expected damage (e.g. losing some data asset) that Defender suffers when  $t$  is attacked and a control has implemented at level  $l$ .

While the implementation of a cyber security control strengthens the defence of Defender's organisation, it is associated with two types of costs namely; *indirect* and *direct*.




Fig. 1: Illustration of our environment [2].

Examples of indirect cost are System Performance Cost, Morale Cost, and Re-Training Cost.

For a level  $l$  we express its indirect cost by the random variable  $C: \mathcal{L} \rightarrow \mathbb{Z}^+$ . From the above we can derive the overall loss of Defender's organisation. This is equivalent to the sum of the security damages inflicted by Attacker and the indirect cost for implementing a control at a certain level. Formally, when Defender implements a control at some level  $i$  then the expected loss of their organisation is derived by

$$\sum_t S(l, t) - C(l).$$

The implementation of a control, at some level, has a direct cost which refers to the budget the organisation must spend to this implementation. For instance, we can split such direct cost into two categories, the Capital Cost and Labour Cost. We express the direct cost of an implementation level  $l$  by the random variable  $\Gamma: \mathcal{L} \rightarrow \mathbb{Z}^+$  that takes implementation levels to the monetary cost of the control implementation.

### B. Game Characterization

In this article we formulate a two-player non-cooperative static game. The players in our game are Defender (she represents any cyber security decision-maker) and Attacker (she represents any cyber hacker who uses commodity attacks). Defender defends their organisation's data assets by minimising expected cyber security losses with respect to the indirect costs, while the attacker Attacker aims at benefiting from compromising Defender's organisation data assets.

The Defender is choosing how to implement a cyber security control (i.e. at which level) and Attacker decides which target to attack to exploit its vulnerability at a certain depth. Since we consider a simultaneous game Attacker does not know the control implementation strategy and Defender does not know the attack strategy. We refer to our games as *control games* because the basis of our formulation is that Defender has one control at her disposal.

In this article we formulate a zero-sum game. This represents scenarios where Attacker aims at causing the maximum possible damage to Defender. We believe that if we consider a non-zero sum game then a specific threat model must be defined as well. Such a model could consider, for instance, some cost for Attacker when undertaking an attack. However cost in terms of cyber attacks is tightly coupled with the profile of the attacker. A consideration of a specific threat model would also have some influence on the way Attacker sees the different targets as she is after specific goals based on her motivation (i.e. cyber crime, hacktivism, cyber espionage). In this case, different Attacker profiles could have been investigated. In our work here, we have not investigated such profiles and our work is limited to a generic assumption that Attacker is taking advantage of commodity attacks that she can purchase from online sources. In other word, we have assumed a set of attack methods that Attacker can choose from but we have not postulated anything about their motivations.

### C. Pure-Strategy Sets

For a given cyber security control, Defender can choose to implement the control at level  $l \in \mathcal{L}$  and therefore her pure strategy set coincides with  $\mathcal{L}$ . The Attacker selects a vulnerability to exploit at a certain depth. Formally, Attacker chooses  $t = \langle v, d \rangle \in \mathcal{T}$ . Thus the pure strategy set of Attacker coincides with  $\mathcal{T}$ .

### D. Payoffs

Given that the pure strategy sets of the players are  $\mathcal{L}$  and  $\mathcal{T}$  then Defender has  $m$  pure strategies and Attacker has  $n$ , correspondingly. We denote by  $G := \langle \mathbb{A}, \mathbb{E} \rangle$  an  $m \times n$  bi-matrix cyber security game where Defender (i.e. row player) has a payoff matrix  $\mathbb{A} \in \mathbb{R}^{m \times n}$  and the payoff matrix of Attacker (i.e. the column player) is denoted by  $\mathbb{E} \in \mathbb{R}^{m \times n}$ .

Defender chooses as one of her pure strategies one of the rows of the payoff bi-matrix  $\langle \mathbb{A}, \mathbb{E} \rangle := [(a_{lt}, e_{lt})]_{l,t}$ . For any pair of strategies  $(l, t)$ , Defender and Attacker have payoff

TABLE I: Notation.

$\mathcal{T}$	set of cyber security targets	$\mathcal{L}$	set of implementation levels
$\mathcal{V}$	set of vulnerabilities	$\mathcal{D}$	set of depths
$T(t)$	threat value of target $t$	$I(t)$	impact of a successful attack against target $t$
$S(l, t)$	security loss when target $t$ is compromised	$E(l, t)$	effectiveness of $l$ on $t$
$\theta_j$	probability $t_j$ to be attacked	$\phi_j$	probability $l_j$ to be selected
$C(l)$	indirect cost of $l$	$\Gamma(l)$	total direct cost of $l$

values equivalent to  $a_{lt}$  and  $e_{lt}$ , given by

$$\begin{aligned} a_{lt} &:= S(l, t) - C(l) \\ e_{lt} &:= -S(l, t) + C(l). \end{aligned}$$

Tables II and III are the player's payoff matrices.

TABLE II: Defender's payoff matrix.

	$t$	$t'$
$l$	$S(l, t) - C(l)$	$S(l, t') - C(l)$
$l'$	$S(l', t) - C(l')$	$S(l', t') - C(l')$

TABLE III: Attacker's payoff matrix.

	$t$	$t'$
$l$	$-S(l, t) + C(l)$	$-S(l, t') + C(l)$
$l'$	$-S(l', t) + C(l')$	$-S(l', t') + C(l')$

### E. Representation of Mixed Strategies

A player's mixed strategy is a distribution over the set of their pure strategies. The representation of Defender's mixed strategy space is a finite probability distribution over the set of the different control implementation levels  $\{l_1, \dots, l_m\}$ . For Attacker, the representation of their mixed strategy space is a probability distribution over the different targets  $\{t_1, \dots, t_n\}$ .

In this paper we are interested in how different control implementation levels are combined in a proportional manner to give an implementation plan for this control. We call this a *cyber security plan*. This allows us to examine advanced ways of mitigating vulnerabilities. A cyber security plan is a probability distribution over different cyber security processes. When investing in cyber security we will be looking into the direct cost of each cyber security plan which is derived as a combination of the different costs of the cyber security processes that comprise this plan.

1) *Defender's Mixed Strategies*: We define Defender's mixed strategy as the probability distribution  $\Phi = [\phi_1, \dots, \phi_m]$ . This expresses a cyber security plan, where  $\phi_j$  is the probability of implementing the control at  $l_j$ . A cyber security plan can be realised as advice to Defender on how to implement a cyber security control by combining different implementation levels. Although this assumption complicates our analysis at the same time it allows us to reason about equilibria of the control game therefore providing a more effective strategy for Defender. We claim that our model is

flexible thus allowing Defender to interpret mixed strategies in different ways to satisfy their requirements.

2) *Attacker's Mixed Strategies*: A mixed strategy of Attacker is defined as a probability distribution over the set  $\{v_1, \dots, v_\nu\} \times \{d_1, \dots, d_\Delta\}$ . In a simpler form, the mixed strategy of Attacker is a probability distribution over the different targets and it is denoted by  $\Theta = [\theta_1, \dots, \theta_n]$ , where  $\theta_i$  is the probability of the adversary attacking  $t_i$ .

3) *Payoffs for Mixed Strategies*: When both players choose a pure strategy randomly according to the probability distributions determined by  $\Phi$  and  $\Theta$ , the expected payoffs to Defender and Attacker are

$$\begin{aligned} J_D(\Phi, \Theta) &:= \sum_{i=1}^n \sum_{j=1}^m \phi_j a_{ij} \theta_i \\ J_A(\Phi, \Theta) &:= \sum_{i=1}^n \sum_{j=1}^m \phi_j e_{ij} \theta_i. \end{aligned}$$

### F. Best Responses Analysis

For the remainder of this section, we analyse a specific control game. We assume that for a specific target  $t$ , Defender has only two possible levels at her disposal namely  $l$ , and  $l'$  (e.g. performing penetration testing rarely during a year or often), to implement a control. We define

$$\begin{aligned} \Delta S(t) &:= S(l', t) - S(l, t) \\ \Delta C &:= C(l') - C(l) \end{aligned}$$

$\Delta S(t)$  is the reduction in damage when  $l'$  is chosen, and  $\Delta C$  is the extra indirect cost of  $l'$  over  $l$ .

*Lemma 1*: When the reduction in damage achieved by  $l'$  over  $l$  is higher than the extra indirect cost that  $l'$  introduces, Defender chooses  $l'$ .

*Proof*: If the reduction in damage achieved by  $l'$  over  $l$  is higher than the extra indirect cost that  $l'$  then  $\Delta S(t) > \Delta C$ . This can be broken down as,  $S(l', t) - S(l, t) > C(l') - C(l) \Leftrightarrow S(l', t) - C(l') > S(l, t) - C(l) \Leftrightarrow a_{l't} > a_{lt}$ . Therefore, the Defender is incentivised to pick  $l'$  as it has a higher utility. ■

*Lemma 2*: If  $S(l, t) > S(l, t')$  then Attacker attacks target  $t$ .

*Proof*: For a specific control implementation  $l$  and two targets  $t, t'$ , Attacker's best response can be found by comparing  $e_{lt}, e_{lt'}$ . If  $e_{lt} > e_{lt'} \Leftrightarrow S(l, t) - C(l) > S(l, t') - C(l) \Leftrightarrow S(l, t) > S(l, t')$ , Attacker prefers to attack target  $t$ . Specifically we define this property as:

$$\Delta S(l) := S(l, t) - S(l, t')$$

Therefore, if  $S(l, t) > S(l, t') \Leftrightarrow S(l, t') - S(l, t) < 0 \Leftrightarrow \Delta S(l) < 0$ , Attacker chooses  $t$ . ■

### G. Saddle Points

Since we are investigating a two-person zero-sum game with finite number of actions for both players, and according to Nash [19] it admits at least a Nash Equilibrium (NE) in mixed strategies. Saddle-points correspond to Nash equilibria as discussed in [20]. The following result, from [21], establishes the existence of a saddle (equilibrium) solution in the games we examine and summarises their properties.

The investigated cybersecurity game admits a saddle point in mixed strategies,  $(\Phi^*, \Theta^*)$ , with the property

$$\begin{aligned}\Phi^* &= \arg \max_{\Phi} \min_{\Theta} J_U(\Phi, \Theta) \\ \Theta^* &= \arg \max_{\Theta} \min_{\Phi} J_A(\Phi, \Theta).\end{aligned}$$

*Corollary 1:* Regardless of the Attacker's strategy, the Nash Defender guarantees a minimum performance, that is an upper limit of expected damage.

*Proof:* The minimax theorem [22] states that for zero sum games NE, maxmin and minimax solutions coincide. Therefore  $\Phi^* = \arg \min_{\Phi} \max_{\Theta} J_A(\Phi, \Theta)$ . ■

### H. Non-zero Sum Games

Note here that, since in this work we consider zero sum games, two criticisms are possible:

*Remark 1:* The gain of the Attacker is not, in general, equal to the loss of the defender.

*Remark 2:* The Attacker's payoff is not related to the defender indirect costs.

We address both Remarks by noticing that a significant class of *realistic cybersecurity games* can be mathematically reduced to zero sums games. Remark 1 is addressed by the following lemma.

*Lemma 3:* The equilibrium  $(\Phi^*, \Theta^*)$  in our zero sum cybersecurity game  $G$  remains the same in the negative affine transformation of this game in which the Attacker's gain does not equal the Defender's loss.

*Proof:* We claim that a model of the Attacker where his payoffs are a negative affine transformation of the Defender loss is a reasonable model. For example by selling stolen data on the black market for only one tenth of the data's value.

A negative affine transformation of the Defender's  $\mathbb{A}$  matrix is defined as  $\omega \mathbb{A} + \psi$ , where  $\omega$  is a negative scalar, and  $\psi$  is a constant matrix of the same dimension as  $\mathbb{A}$ . Therefore, in addition to the cybersecurity game  $G = (\mathbb{A}, -\mathbb{A})$ , we intuitively define the negative affinity of this game as  $G^- = (\mathbb{A}, \omega \mathbb{A} + \psi)$ .

Suppose  $\Phi^*, \Theta^*$  are the equilibrium strategies in  $G$ . First, it is easy to see that  $\Phi^*$  is the Defender's equilibrium strategy in both  $G$  and  $G^-$  due to the Defender's game matrix remaining the same. Formally,  $\Phi \mathbb{A} \Theta^* \leq \Phi^* \mathbb{A} \Theta^*$ . Similarly, we prove that  $\Theta^*$  is Attacker's equilibrium strategy in both games. We have that  $\Phi^* (-\mathbb{A}) \Theta \leq \Phi^* (-\mathbb{A}) \Theta^* \Rightarrow \Phi^* \mathbb{A} \Theta \geq$

$\Phi^* \mathbb{A} \Theta^* \Rightarrow \Phi^* (\omega \mathbb{A} + \psi) \Theta \leq \Phi^* (\omega \mathbb{A} + \psi) \Theta^*$ . This means that equilibria are the same in both  $G, G^-$ . ■

*Lemma 4:* A game  $\hat{G}$  where the Defender's indirect cost  $C$  is a positive affine transformation of the direct cost  $S$ , has the same maxmin solution with  $G$ .

*Proof:* According to the Lemma we have that in  $\hat{G}$  Defender's payoff is given by

$$S - (\kappa S - \mu) = S(1 - \kappa) - \mu,$$

where  $\kappa, \mu$  are positive scalars. Assume that at the equilibrium of  $\hat{G}$  Defender's best response is  $\Phi^*$ . Then we have

$$\begin{aligned}\Phi [S(1 - \kappa) - \mu] \Theta^* &\leq \Phi^* [S(1 - \kappa) - \mu] \Theta^* \\ \Rightarrow \Phi (S - \kappa S - \mu) \Theta^* &\leq \Phi^* (S - \kappa S - \mu) \Theta^* \\ \Rightarrow \Phi (S - \mu) \Theta^* &\leq \Phi^* (S - \mu) \Theta^* \\ \xrightarrow{\mu=C} \Phi (S - C) \Theta^* &\leq \Phi^* (S - C) \Theta^* \Rightarrow \Phi \mathbb{A} \Theta^* \leq \Phi^* \mathbb{A} \Theta^*.\end{aligned}$$

Therefore  $G, \hat{G}$  have the same equilibria, and from Corollary 1 these are also maxmin solutions. ■

### I. A Small Game Example

To illustrate the game approach let's consider a toy example consisting of a 2-level, and 2-target control games, where Defender and Attacker make their decisions simultaneously, or, equivalently, independently of each other. The information sets associated with the the control game, investigated in this section, depicted in Fig.2; a dashed curve encircling the Attacker nodes has been drawn. This indicates that Attacker cannot distinguish between these two points. In other words, Attacker has to arrive at a decision without knowing what Defender has actually chosen.




Fig. 2: Game tree for the control game with 2 implementation levels and two targets.

Due to the game being zero-sum, we have kept only the payoffs of Defender at the game tree. We also defined the mixed strategy of Defender as the probability distribution  $[\phi, 1 - \phi]$ , where  $\phi$  is the probability of implementing the control at level  $l$ . Attacker's mixed strategy is denoted by  $[\theta, 1 - \theta]$ , where Attacker chooses to attack  $t$  with probability  $\theta$ . Table IV summarizes all possible best responses of the control game for the different conditions discussed in this section.

TABLE IV: Nash equilibria for the different conditions.

	$\Delta S(t') > \Delta C$		$\Delta S(t') < \Delta C$	
$\Delta S(t) > \Delta C$	$\Delta S(l') > 0$	$(l', t)$	$\Delta S(l') < 0$	$(\phi l', (1 - \theta)t)$
	$\Delta S(l') < 0$	$(l', t')$	$\Delta S(l) > 0$	$((1 - \phi)l, \theta t')$
$\Delta S(t) < \Delta C$	$\Delta S(l) < 0$	$((1 - \phi)l', (1 - \theta)t')$	$\Delta S(l) > 0$	$(l, t)$
	$\Delta S(l') > 0$	$(\phi l, \theta t)$	$\Delta S(l) < 0$	$(l, t')$

### J. Player Mixed Strategies

In a two target, two level control sub-game, it is possible to define the probabilities that each player plays a particular mixed strategy.

*Lemma 5:* The Nash equilibrium for a control sub-game for the Defender's, given by  $\phi^* \in [0, 1]$  is:

$$\phi^* = \frac{\Delta S(l')}{\Delta S(l') - \Delta S(l)}$$

*Proof:*

The Defender wants to make the Attacker indifferent to which target they should attack.

This is given by equalising the expected payoff of the Attacker, thus

$$\begin{aligned} A(t) &= \phi^* e_{lt} + (1 - \phi^*) e_{l't} \\ A(t') &= \phi^* e_{l't'} + (1 - \phi^*) e_{l't'} \end{aligned}$$

giving

$$\phi^* e_{lt} + (1 - \phi^*) e_{l't} = \phi^* e_{l't'} + (1 - \phi^*) e_{l't'}. \quad (2)$$

We can substitute terms such that Eq.(2) can be written in terms of  $e_{lt}$ , hence

$$\begin{aligned} e_{l't} &= e_{lt} - \Delta S(t) + \Delta C \\ e_{l't'} &= e_{lt} - \Delta S(l) \\ e_{l't} &= e_{lt} - \Delta S(t) + \Delta C - \Delta S(l') \end{aligned}$$

By substituting the equations above into Eq.(2) we get

$$\begin{aligned} \phi^* e_{lt} + (1 - \phi^*) (e_{lt} - \Delta S(t) + \Delta C) &= \phi^* (e_{lt} - \Delta S(l)) \\ + (1 - \phi^*) (e_{lt} - \Delta S(t) + \Delta C - \Delta S(l')). \end{aligned} \quad (3)$$

Eq.(3) can be expanded and reduced to

$$\Delta S(l') = \phi^* (\Delta S(l') - \Delta S(l)).$$

This then gives

$$\phi^* = \frac{\Delta S(l')}{\Delta S(l') - \Delta S(l)}$$

*Lemma 6:* The Nash strategy of the Attacker in a control sub-game, is given by

$$\theta^* = \frac{\Delta S(t) - \Delta C + \Delta S(l') - \Delta S(l)}{\Delta S(l') - \Delta S(l)}$$

*Proof:* At the equilibrium, the Attacker wants to make the Defender indifferent to which target they should attack. This is given by equalising the expected payoff of the Defender:

$$\begin{aligned} D(l) &= \theta^* a_{lt} + (1 - \theta^*) a_{l't} \\ D(l') &= \theta^* a_{l't} + (1 - \theta^*) a_{l't'}. \end{aligned}$$

Therefore

$$\theta^* a_{lt} + (1 - \theta^*) a_{l't} = \theta^* a_{l't} + (1 - \theta^*) a_{l't'}. \quad (4)$$

We can substitute terms such that Eq.(4) can be written in terms of  $a_{lt}$ , and therefore

$$\begin{aligned} a_{l't} &= a_{lt} + \Delta S(t) - \Delta C \\ a_{l't'} &= a_{lt} + \Delta S(l) \\ a_{l't} &= a_{lt} + \Delta S(t) - \Delta C + \Delta S(l') \end{aligned}$$

By substituting the equations above into Eq.(4) we get:

$$\begin{aligned} \theta^* a_{lt} + (1 - \theta^*) (a_{lt} + \Delta S(l)) &= \\ \theta^* (a_{lt} + \Delta S(t) - \Delta C) + (1 - \theta^*) (a_{lt} + \Delta S(t) - \Delta C + \Delta S(l')). \end{aligned}$$

The above equation can be expanded and reduced to:

$$a_{lt} + \Delta S(l) - \theta^* \Delta S(l) = a_{lt} + \Delta S(t) - \Delta C + \Delta S(l') - \theta^* \Delta S(l').$$

This then gives

$$\theta^* = \frac{\Delta S(t) - \Delta C + \Delta S(l') - \Delta S(l)}{\Delta S(l') - \Delta S(l)}.$$

### K. Numerical Illustration

We see that Defender's strategy is derived only from  $\Delta S(l)$  and  $\Delta S(l')$ . This is since Defender wants to make the Attacker indifferent to the target they want to attack at the equilibrium. In this way the aspects of  $\Delta S(t)$  and  $\Delta C$  are not represented as they do not impact Attacker.

In Fig. 3 we see the results with the inclusion of the pure strategy solutions. When  $\Delta S(l) = \Delta S(l')$  the solution is undefined, such that Defender has no incentive to play one defence over the other, as there is no  $\phi$  that will influence the outcome of the game. Additionally, we see that when  $\Delta S(l)$  does not equal  $\Delta S(l')$  and both  $\Delta S(l) \leq 0$  and  $\Delta S(l') \leq 0$  or  $\Delta S(l) \geq 0$  and  $\Delta S(l') \geq 0$  then we have a collapse onto a single dominant strategy. The single strategy, represented here






Fig. 3: Defender's Strategy Space

by  $\phi^* = 0$  or  $\phi^* = 1$ , is defined depending on which solution would make the attacker the most indifferent to which target they wish to attack.

When  $(\Delta S(l) > 0) \wedge (\Delta S(l') < 0)$  or  $(\Delta S(l) < 0) \wedge (\Delta S(l') > 0)$ , then we have the conditions for a mixed strategy. This is because there is no pure strategy that is dominant, and therefore the defender aims to make the attacker as indifferent as possible.

#### IV. CYBER SECURITY INVESTMENT

Thus far the analysis performed has considered a single-control, two-targets, two-levels game. Our plan for cyber investment is to solve a set of control sub-games with  $n$  targets and up to  $m$  control levels for each control. Given these game solutions we will then use a Knapsack algorithm to provide the general investment solution. The control-game solutions provide us with information regarding the way in which each security control is best implemented, so as to maximise the benefit of the control with regard to both the Attacker's strategy, and the indirect costs of the organisation.

It is easy to see that, in control sub-games, the games look only at the vulnerabilities that are directly relevant to the control being implemented. The cyber security investment problem expands to represent all of an organisation's vulnerabilities and selecting the best cyber security controls based on the outcomes of the control-games.

With regard to an implementation of cyber security processes based on the sub-game solutions, it is important to understand what a control game solution represents in the process of making those decisions. In particular this is about what a mixed strategy means in terms of control implementation.

We motivate the concept of a mixed strategy as a method for trying to define where in the system it is most effective to implement the control. Based on our interpretation of the structure of a network, this will generally involve protecting devices at the highest level with the strictest controls where possible, then assigning lower levels of controls to devices and users that operate at depths with less sensitive data.

This is performed by creating a logical ordering of the most important devices, based on the perceived risk of the device or the user. While there may be a logical ordering across an organisation for all controls, it often might make more sense to order users and devices specifically for each control based on vulnerability.

*Example 1:* To illustrate this we take for example a security control entitled *Vulnerability Scanning and Automated Patching*, and we assume 5 different implementation levels i.e.  $\{0, 1, 2, 3, 4\}$  where level 4 corresponds to *real-time scanning* while level 2 to *regular scanning*. We say that a mixed strategy  $[0, 0, 0.7, 0, 0.3]$  determines a cyber security plan that dictates the following:

- 0.3  $\mapsto$  real-time scanning for the 30% of the most important devices
- 0.7  $\mapsto$  regular scanning for the remaining 70% of devices

This mixed strategy can be realised more as an advice to a security manager on how to undertake different control implementations rather than a rigorous set of instructions related only to a time factor. We claim that our model is flexible thus allowing the defender to interpret mixed strategies in different ways to satisfy their requirements.

#### A. Full Game Representation

A Full Game representation considers the method of solving the investment problem by creating a strategic game containing the set of feasible choices available to both players. Defender's pure strategies are comprised of an implementation level for each of the controls, and Attacker's pure strategies consist of each target in the set of all possible targets. One of the considerations that needs to be made is with regards to the budget. A pure game-theoretic solution for the cyber investment problem would require modelling  $n$  targets,  $m$  control levels and  $c$  controls. A naive choice would be to consider  $c \times m \times n$  games. However it is not clear how to force these game solutions to satisfy budget constraints. A game model satisfying budget constraints could be built using the idea of "schedules" [23], i.e. a pure strategy is a tuple of  $c \times m$  bits where each bit represents the implementation of a control at a particular level, 1 stands for implemented and 0 for not implemented. The budget requirement can be easily imposed on such tuples, for example by only considering tuples whose costs do not exceed the budget. The problem with this proposal is that, in principle, there could be an exponential number of pure strategies, in the order of  $2^{(c \times m)}$ . Also it would be non-trivial to choose appropriate payoffs for such tuples. In this case, we restrict the combination of controls in the payoff matrix to only those that can be purchased based on the maximum amount of budget.

#### B. Pure Knapsack Representation

A Pure Knapsack representation considers the method of solving the investment problem given that Defender only considers the implementation of "whole" controls. We have chosen a 0-1 Multiple Choice, Multi-Objective Knapsack.

The choice of this type of Knapsack is motivated as follows: "0-1" because each level of implementation of a control is implemented or not implemented, "Multiple Choice" because only one solution for each control (the optimal one) ought to be chosen and "Multi-Objective" because each target represents an optimisation objective. More precisely we define

in this case as  $Q$  a control implementation, where  $Q_{jl}$  is the implementation for control  $j$  at level  $l$ .

We define a solution to the Knapsack problem as

$$\Psi = \{Q_{1l_1} \dots Q_{cl_c}\}, l_i \in \mathcal{L}.$$

A solution  $\Psi$  takes exactly one level for each control as a policy for implementation, where notice for each control there exists a solution  $Q_{j0}$ , which dictates that control  $j$  should not be used.

To represent the cyber security investment problem, we need to expand the definitions for both expected damage ( $S$ ) and effectiveness ( $E$ ) to incorporate multiple controls. So we expand  $S$  such that  $S(\Psi, t)$ , which is the expected damage on target  $t$  given the implementation of the levels associated with  $\Psi$ . Likewise, we expand the definition of the effectiveness of the implemented solution on a given target as  $E(\Psi, t)$ . Additionally, we consider  $\Gamma(Q_{jl})$  as the direct cost of implementing  $Q_{jl}$ , and the maximum of all  $\Gamma$ s cannot exceed the available cyber security budget  $B$ . Then we solve

$$\begin{aligned} \max_{\Psi} \min_{t_i} \quad & \{1 - \sum_{j=1}^c \sum_{l=0}^m E(Q_{jl}, t_i) z_{jl}\} S(\Psi, t_i) - C_j(l) \\ \text{s.t.} \quad & \sum_{j=1}^c \sum_{l=0}^m \Gamma(Q_{jl}) z_{jl} \leq B \\ & \sum_{l=0}^m z_{jl} = 1, z_{jl} \in \{0, 1\}, \forall j = 1, \dots, c \end{aligned}$$

where  $z_{jl}$  is the probability of implementing control  $j$  at level  $l$ , and  $B$  is the available cyber security budget.

### C. Hybrid Method

The Hybrid approach avoids the problems of the Full Game method by considering the particular game solutions for each control as part of an overall combinatorial optimisation problem which we also solve using a 0-1 Multiple Choice, Multi-Objective Knapsack.

In this case we define as  $Q$  a sub-game solution, where  $Q_{jl}$  is the sub-game solution for control  $j$  implemented at level  $l$ .

We represent the 0-1 Multiple Choice, Multi-Objective Knapsack Problem as:

$$\begin{aligned} \max_{\Psi} \min_{t_i} \quad & \{1 - \sum_{j=1}^c \sum_{l=0}^m E(Q_{jl}, t_i) z_{jl}\} S(\Psi, t_i) \\ \text{s.t.} \quad & \sum_{j=1}^c \sum_{l=0}^m \Gamma(Q_{jl}) z_{jl} \leq B \\ & \sum_{l=0}^m z_{jl} = 1, z_{jl} \in \{0, 1\}, \forall j = 1, \dots, c \end{aligned}$$

Notice that in the above formula, we do not have the factor of indirect cost ( $C_j(l)$ ), this is because indirect cost is taken into consideration in the control sub games.

In addition, we consider a tie-break condition in which if multiple solutions are viable, in terms of maximising the minimum, according to the above function we will select the

solution with the lowest cost. This ensures that an organisation is not advised to spend more on security controls than would produce the same net effect.

## V. CASE STUDY: SANS CRITICAL SECURITY CONTROLS, CWE TOP SOFTWARE VULNERABILITIES

To compare the Full Game, Hybrid and Pure Knapsack methods of decision support, we have developed a sample case study similar to one expected in a real environment.

In this work, we are interested in comparing two aspects of the solutions generated by the different methods. The first is the optimality of the solutions, and the other is their complexity. To this end, we consider the *optimality of the solution* to be the expected damage of the implemented set of controls at the weakest level, for a given budget.

The *complexity of the solution* provides a pivotal role in decision support with cyber security, where overly complex solutions are potentially difficult to implement and follow. This is relevant with mixed strategy equilibria.

### A. Modelling SANS Critical Controls and CWE Top Software Vulnerabilities

Our case study is created using a mapping from the SANS Critical Security Controls [24] combined with the CWE Top 25 Software Vulnerabilities[25]. In this mapping, we define a control as a collection of any of the associated processes defined by a single critical security control in the SANS top 20. Additionally, we consider a vulnerability as any of the software vulnerabilities that are defined in the CWE Top 25. Using data associated with these two sources we are able to build the core components of a case study to test our methodology.

From the SANS Critical Security controls, we define a level of Implementation for a given control as a single action point listed for the control. For each control, the control levels are considered in order, based on their position in the list. In some cases where there is significant overlap between control levels, levels can be combined. This is aimed at reducing the number of strategies and computational complexity of the problem.

Using the classifications provided by the CWE Top 25 Software vulnerabilities we are able to categorise the different classes of vulnerabilities that each of the controls is able to mitigate. CWE proposes three categories *Insecure Interactions Between Components*, *Risky Resource Management* and *Porous Defences*. A given vulnerability falls into one of the three categories and we consider that any control may cover the vulnerabilities associated with one or more of the categories.

To calculate part of the risk, we consider the threat value to be directly associated with each of the vulnerabilities. CWE defines a score out of 100 for a number of vulnerabilities, with those scoring highest published in their top 25. The damage values associated with each of the weaknesses have been scaled to fit within the range of the other values gathered for this case study.

The efficiency of a control level is considered to be a reduction in the effectiveness of a given control based on

the idea that a control should be effective at stopping an attack. In reality an Attacker may be able to circumvent these controls. We consider that their ability to do so is linked to the availability and ease with which information about vulnerabilities can be discovered. Then, we directly link this ability to the attack factors provided by CWE. To do this, we apply a weighted percentage of the values for the four key factors that CWE defines regarding a vulnerability capped according to the level. The four attack factors that CWE defines for a given vulnerability are *Prevalence*, *Attack Frequency*, *Ease of Detection* and *Attacker Awareness*.

CWE provides an expected cost value to repair the vulnerability for each weakness. For this mapping, we consider each cost separately for each vulnerability. More specifically, the direct cost of a control is given as the sum of all the costs for the vulnerabilities that it covers. This direct cost value is considered to be the cost of implementation at the highest level for the control. The direct costs for lower levels are scaled uniformly based on the number of levels the control has.

The mapping provided is able to cover the technical aspects of the controls and vulnerabilities, however, there are certain aspects unique to each organisation. We consider that both the impact of a successful attack and the indirect costs incurred need to be defined based on priorities and requirements of the organisation. The impact of a successful attack is given by not only the data loss to the organisation, but also by the loss of reputation, as well as other costs associated with a breach. The indirect costs are considered to be the importance that the organisation places on the day to day performance of the system, as well as the ability and willingness of staff to comply with any additional security policies.

### B. Values

The case study presented in this work considers a network separated into three different depths, consistent with Fig. 1, where Defender has seven different controls available to protect the network from twelve different vulnerabilities. For this example, we consider the levels available to Defender to consist of the quick win processes provided by SANS. The seven controls are shown in Table V and the twelve vulnerabilities are shown in Table VI. Based on the controls used, the budget at which all controls become available at the highest level equals 82.

### C. Optimality Comparison

In comparing the damage at the weakest target provided by the Full Game, Hybrid Method to the Knapsack Representation, we can see in Fig.4 that, in general, the Full Game Representation will provide a better defence to the weakest target for low budget levels. However, once the budget becomes larger, we see that the Hybrid Method is able to reach a level of coverage that will minimise the damage at each target, whereas neither the Full Game Representation nor the pure Knapsack Representation fully cover the weakest target, even with the maximum budget. This is owing to the impact that the indirect cost has on the decision-making process.

Where the Hybrid Method includes the impact of direct cost in the decisions, regarding the optimality of the deployment of the control at each level

- the pure Knapsack includes the indirect cost, as a whole, in the outcome of the optimisation, and
- the Full Game applies the indirect cost to each strategy in the payoff matrix.




Fig. 4: Case Study Results with Normally Expected Indirect Costs

We can also see that the impact of the indirect cost causes the Full Game Representation to become inefficient compared to the Hybrid and Knapsack, before reaching the maximum defence. This occurs in Fig. 4 within the budget range of 27 - 32.

Fig. 4 shows that with low indirect costs, the outcome of the control sub-games allow for the availability of better strategies with lower budgets than the Knapsack-only representation. This is due to the Hybrid Method being able to use a combination of packages that has a lower direct cost, but it provides the required coverage to the weakest target. Where a similar coverage of the weakest target is only available to the Knapsack when the pure strategies are made available.

It has also been seen that with higher indirect costs both the Full Game and pure Knapsack Representation will offer increasingly poor results when compared to the Hybrid Method. This is due to the point at which the cost of implementing controls outweighing the benefit being reached at a lower budget.

In Fig. 5, we see that when there are no indirect costs, the Hybrid Method and Knapsack Representation-only method have exactly the same output. This is due to the outcome of each control sub-game providing a pure strategy at the highest level of implementation, which would result in the Knapsack solver having identical strategies with regard to the optimisation.

With regard to the indirect cost, if there is no indirect cost to implementing the control then there is no trade-off to a higher defence. This means that providing that an appropriate budget is available, then the best defence will be purchased by all methods, as seen in Fig. 5. Most importantly, this means that the Full Game Representation provides solutions that are often

TABLE V: Case Study Controls.

Control	Levels
Inventory of Authorised and Unauthorised Devices (1)	3
Inventory of Authorised and Unauthorised Software (2)	3
Secure Configuration for Hardware and Software on Devices (3)	5
Continuous Vulnerability Assessment and Remediation (4)	4
Malware Defences (5)	6
Application Software Security (6)	2
Controlled Use of Administrative Privileges (12)	6

TABLE VI: Case Study Vulnerabilities.

$v_z$ : Vulnerability (CWE-code)	PR	AF	ED	AA	Vulnerability	PR	AF	ED	AA
$v_1$ : SQLi (89)	2	3	3	3	$v_7$ : Missing encryption (311)	2	2	3	2
$v_2$ : OS command injection (78)	1	3	3	3	$v_8$ : Unrestricted upload (434)	1	2	2	3
$v_3$ : Buffer overflow (120)	2	3	3	3	$v_9$ : Unnecessary privileges (250)	1	2	2	2
$v_4$ : XSS (79)	2	3	3	3	$v_{10}$ : CSRF (352)	2	3	2	3
$v_5$ : Missing authentication (306)	1	2	2	3	$v_{11}$ : Path traversal (22)	3	3	3	1
$v_6$ : Missing authorization (862)	2	3	2	2	$v_{12}$ : Unchecked code (494)	1	1	2	3

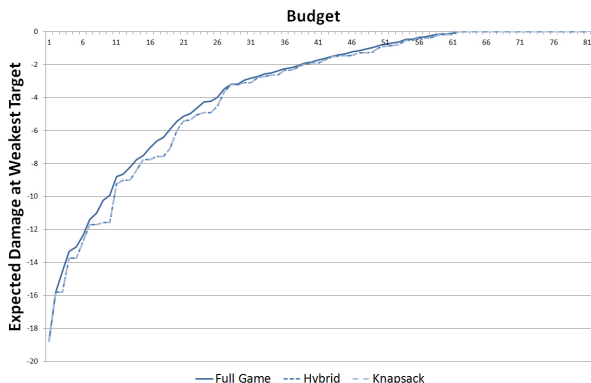


Fig. 5: Case Study Results with No Indirect Costs

more optimal, but at least no worse than those generated by the other methods. This is because the Full Game Representation has no drawbacks to the implementation of the best controls in the most optimal configuration, which is still a restriction on the two methods that implement the 0-1 restriction of the Knapsack.

#### D. Complexity Comparison

To identify how complex a solution is, we need to consider the composition of solutions to assess how complex the solutions are and if they advice could be reasonably followed by an individual on behalf of an organisation.

If we take for example a budget of 18, the solution provided by the Full Game Representation, comprises of a mixed strategy consisting of 4 packages. If we consider the three major strategies in Table VII, then we can see that they all suggest the use of control 5 at level 2. Additionally, we see

that at a minimum, control 6 should be implemented at level 1, with an increase to level 2 for 23% of the time. This suggests always having software versions supported by vendors, but only implementing web application firewalls on the top 23% of the system.

We can also see that the solution suggests using control 2 and control 7 with 41.3% and 76.8% respectively. Both of these controls determine restricting access, either through application whitelisting or a reduction in non-essential admin rights and activities. Control implementations can be considered in one of two ways, either the percentage relates to the number of devices that feature this control or the severity with which the control is implemented. In the case of application whitelisting an implementation level of 41.3% would allow a greater degree of software availability to users than a higher percentage.

Control 3 suggests using level 2 35.5% of the time, with level 1 not being suggested individually. This relates to using both secure configurations for each operating system and implementing automated patching tools. The easiest way to interpret this solution is to state that the top 35.5% of devices utilise both levels of this control, while all other devices don't.

While in this case the mixed strategies provided by the Full Game Representation do not represent vastly different strategies, the addition of more controls and vulnerabilities will increase the complexity of the solution space.

Additionally for a budget of 18, we found that the Hybrid Method suggests using the solution  $[0, 1, 1, 0, 4, 0, 1]$ , while the pure Knapsack solution suggests  $[0, 0, 2, 0, 2, 1, 0]$ . In this example the package suggested by the Hybrid Method is viable as a solution for Knapsack, as none of the sub-game solutions differ from the highest level pure strategy that would

TABLE VII: Full Game Solution for a Budget of 18.

Package	Probability
[0, 0, 0, 0, 2, 2, 0]	0.23
[0, 0, 2, 0, 2, 1, 1]	0.355
[0, 1, 0, 0, 2, 1, 1]	0.413
[0, 1, 1, 0, 0, 1, 1]	0.001

be available to the Defender. This further highlights that the indirect costs are pivotal in demonstrating the optimality of these results.

With a budget of 29, we see that both the Hybrid and the Pure Knapsack select the package [0, 1, 1, 0, 2, 2, 1]. While the Full Game Representation is able to use this package, it selects it with probability  $p = 0.001$ . This again shows the importance that indirect costs can have on the optimality of the solution, given that while feasible, the game-theoretic method considers it too costly to implement.

If we consider the case when the budget is 48, the Hybrid Method provides the solution [0, 2, 5, 0, 4, 2, 2], where for control 3 (Secure Configuration for Hardware and Software on Devices) the outcome is to use a mixed strategy. The mixed strategy suggests using level 4 with  $p = 0.609$  and level 5 with  $p = 0.391$ . At level 4 we consider the following of strict configuration management for creating secure images of each OS, and level 5 is concerned with the storage of these images. In this case we can consider that at all times the secure image is used, but the secure storage of master images is only considered for approximately 40% of the time.

The pure Knapsack has solutions that can be followed intuitively as they only ever consider a single level of implementation. We can also see that the Hybrid Method often uses pure strategies as in many cases the outcomes of the control sub-games lead to a single strategy at many levels. However, we find that there is an additional level of complexity in the comprehension of the strategies that are produced by the Full Game. Such complexity can potentially lead to strategies that can not easily be followed by a user to gain the most from the solution. In these cases, there is a risk that the solutions are not followed correctly and with security. This could lead to a potentially weaker defence over a seemingly weaker, but more easily interpreted solution.

## VI. CONCLUSIONS

In this paper we have presented an analysis of a hybrid game-theoretic and optimisation approach to the allocation of an organisation's cyber security budget. For this purpose, we have compared three different approaches to allocating this budget. We found that when there are no indirect costs to consider or the indirect costs have a minimal impact compared to the benefit, then the Full Game Representation gave solutions better than or equal to those of the other two methods. However, when an increase in security is matched by the indirect costs, then the Full Game Representation, is not able to overcome the addition of the indirect costs in favour of a stronger defence, in a similar way to the Pure Knapsack

Representation. The Hybrid Method, however, considers the indirect costs as part of each control game and therefore considers the optimality of each control first, and that an optimal solution is the best valid combination of the optimised controls.

In terms of understanding the solutions, we have found that with a relatively small case study the results can be interpreted in a relatively simple manner. However, we are concerned that for a larger case study the Full Game Representation would create solutions that are too complex to be interpreted in an accurate manner so that they could result in a weaker defence.

This work also highlights the impact, which the indirect costs have on the problem of cyber security budget allocation. Considering the downside that the implementation of a control may have on the organisation is important, since it can better capture the decision-making process required for investment. The results presented in this paper demonstrate how those indirect costs are able to influence the optimal decision in cyber security budget allocation.

We aim to use the work presented in this paper to inform models of attacks against a system. These games model the interactions between an attacker and defender at the *Point of Attack*, and during an ongoing attack. To do this we will consider multi-stage games which represent the stages of an attack and recovery in a system. The techniques presented in this paper should allow for the development of tools for better allocation of resources to help prevent successful attacks made against a system.

In addition, we aim to work with security practitioners in order to create a more detailed case study, and to highlight the operation of this method in a realistic setting. The objective of this is to better understand the parameters that organisations have, and how these can best be applied within our framework. The result of this analysis will lead to the development of a dedicated tool for cyber security budget allocation.

## REFERENCES

- [1] Deloitte, "Nascio cyber security study state governments at risk: a call for collaboration and compliance (2012)," <https://www.deloitte.com/assets/Dcom-UnitedStates/Localsecurity>
- [2] E. Panaousis, A. Fielder, P. Malacaria, C. Hankin, and F. Smeraldi, "Cyber security games and investments: A decision support approach," in *Proc. of the 5th International Conference on Decision and Game Theory for Security*, 2014.
- [3] F. Smeraldi and P. Malacaria, "How to spend it: Optimal investment for cyber security," in *Proc. of the 1st International Workshop on Agents and cyber security (ACySe)*, 2014.
- [4] R. Anderson, "Why information security is hard," in *Proc. of the 17th Annual Computer Security Applications Conference*, 2001.
- [5] C. Kiekintveld, T. Islam, and V. Kreinovich, "Security games with interval uncertainty," in *Proc. of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, 2013, pp. 231–238.
- [6] T. Alpcan and T. Basar, *Network Security: A Decision and Game-Theoretic Approach*. Cambridge University Press, 2010.
- [7] T. Alpcan, "Dynamic incentives for risk management," in *Proc. of the 5th IEEE International Conference on New Technologies, Mobility and Security (NTMS)*, 2012.
- [8] T. Alpcan and N. Bambos, "Modeling dependencies in security risk management," in *Proc. of the Fourth International Conference on Risks and Security of Internet and Systems (CRiSIS)*, 2009, pp. 113–116.

- [9] W. Saad, T. Alpcan, T. Basar, and A. Hjorungnes, "Coalitional game theory for security risk management," in *Proc. of the 5th International Conference on Internet Monitoring and Protection (ICIMP)*, 2010, pp. 35–40.
- [10] P. Bommannavar, T. Alpcan, and N. Bambos, "Security risk management via dynamic games with learning," in *Proc. of the 2011 IEEE International Conference on Communications (ICC)*, 2011, pp. 1–6.
- [11] A. Fielder, E. Panaousis, P. Malacaria, C. Hankin, and F. Smeraldi, "Game theory meets information security management," in *Proc. of the 29th IFIP International Information Security and Privacy Conference*, 2014.
- [12] L. Gordon and M. Loeb, "The economics of information security investment," *ACM Transactions on Information and System Security (TISSEC)*, 2002.
- [13] B. Johnson, J. Grossklags, N. Christin, and J. Chuang, "Nash equilibria for weakest target security games with heterogeneous agents," *Gamenets 2011 LNICST*, vol. 75, pp. 444–458, 2012.
- [14] H. Cavusoglu, R. Srinivasan, and T. Wei, "Decision-theoretic and game-theoretic approaches to it security investment," *Journal of Management Information Systems (ACySe)*, pp. 281–304, 2008.
- [15] M. Cremonini and D. Nizovtsev, "Understanding and influencing attackers' decisions: Implications for security investment strategies," *Working Paper*, 2006.
- [16] M. Al-Humaidani and D. Dunn, "A model of return on investment for information systems security," *argument 4*, 2003.
- [17] J. Wang, A. Chaudhury, and H. Raghav Rao, "Research note-a value-at-risk approach to information security investment," *Information Systems Research*, pp. 106–120, 2008.
- [18] L. Demetz and D. Bachlechner, "To invest or not to invest? assessing the economic viability of a policy and security configuration management tool," *The Economics of Information Security and Privacy*, pp. 25–47, 2013.
- [19] J. Nash, "Equilibrium points in n-person games." in *Proc. of the National Academy of Sciences*, 1950, pp. 48–49.
- [20] T. Alpcan and T. Basar, *Network Security: A Decision and Game-Theoretic Approach*. Cambridge University Press, 2010.
- [21] T. Basar and G. J. Olsder, *Dynamic noncooperative game theory*. London Academic press, 1995.
- [22] J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*, 2007.
- [23] J. Tsai, C. Kiekintveld, F. Ordonez, M. Tambe, and S. Rathi, "Iris-a tool for strategic security allocation in transportation networks," 2009.
- [24] SANS, "The critical security controls for effective cyber defense (version 5.0)," <http://www.counciloncybersecurity.org/attachments/article/12/CSC-MASTER-VER50-2-27-2014.pdf>, 2014.
- [25] CWE, "Cwe top 25 most dangerous software errors (2011)," ["http://cwe.mitre.org/top25/"](http://cwe.mitre.org/top25/), 2014.