

Decision support approaches for cyber security investment



Andrew Fielder^a, Emmanouil Panaousis^b, Pasquale Malacaria^c, Chris Hankin^a, Fabrizio Smeraldi^c

^aImperial College London, United Kingdom

^bUniversity of Brighton, United Kingdom

^cQueen Mary University of London, United Kingdom

ARTICLE INFO

Article history:

Received 5 March 2015

Received in revised form 21 December 2015

Accepted 23 February 2016

Available online 19 March 2016

Keywords:

Cyber security investments decision support

Game theory

Optimisation

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. As game theory captures the interaction between the endogenous organisation's and attackers' decisions, 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. To compare these approaches we built a decision support tool and a case study regarding current government guidelines. The endeavour 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. Going a step further in validating our work, we have shown that our decision support tool provides the same advice with the one advocated by the UK government with regard to the requirements for basic technical protection from cyber attacks in SMEs.

© 2016 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. 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 while in a report published by Deloitte and NASCIO [1], 75.5% of CISOs cited lack of sufficient budget as a top challenge. It is this perceived lack of sufficient funding that this work wishes to address. As approximately 72% of cyber breaches occur at Small-Medium Enterprises (SMEs) [2], we have decided to investigate cyber security investment decisions for SMEs. In addition to SMEs being attractive targets for cyber attackers, from

our work with local SMEs we have identified that they are heavily restricted with the available funding for cyber security, generally 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 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*.

The approach taken in this work is to model attackers using *commodity cyber threats* 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. The same approach has been taken by the UK government to provide cyber security advice to SMEs [3] and published in a report called "Cyber Essentials Scheme: Requirements for basic technical protection from cyber attacks".

The seminal 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. 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 optimal budget allocations can only be considered as strong as the defence of the *weakest target*, as defined in [5]. 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.

In a nutshell, this work proposes a two stage model designed to aid security managers with decisions regarding the optimal allocation of a cyber security budget. 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*. We motivate the use of this methodology by comparing the two-stage model to two alternative methods. Firstly, we model the scenario as a one shot game that aims to optimise the defence 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.

This paper significantly extends the results initially presented in [6]. Its additional contributions include: enriching the mathematical notation to represent more coherent game information; providing a more in-depth mathematical analysis of Control Games' equilibria; comparing the previously proposed, in [6], method of investment, which was based on both Control games and multi-objective

Table 1

Comparative analysis of major works that investigate allocation of a cyber security budget after conducting cyber security risk assessment.

Paper	Game th.	Optimis.	Real world data	Sec. controls sel.
[13,15,18]	x	✓	x	x
[10,11]	x	✓	✓	✓
[12]	x	✓	x	✓
[17]	✓	x	x	x
[19]	✓	✓	x	x
Our article	✓	✓	✓	✓

Knapsack optimisation, to (i) a pure multi-objective Knapsack optimisation, and (ii) a Full Game approach, where we consider all possible controls, levels, and targets under a single very large game; implementing a large scale case study using real world data from various reputable sources; and drawing thorough insights regarding the effectiveness of our cyber security investment method and highlighting how it is in line with [3].

The remainder of this paper is organised as follows. Section 2 summarises the most important related work at the intersection of cyber security investments and selection of security controls. It also highlights how our approach is different. In Section 3, we introduce the model components, which facilitate the risk assessment prior to selection of security controls and investment. Section 4 uses these components to build a game model and analyse a toy 2×2 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 5 we introduce the Control Subgames to support the analysis of games larger than 2×2 . In Section 6, we present three different cyber security investment approaches, which we have simulated by using a novel decision support tools developed for the purposes of this paper. In Section 7 we develop a real world case study based on the SANS Critical Security Controls and CWE Top Software Vulnerabilities. This case study has been used to compare our findings to the set of guidelines that the UK government has published in [3].

2. Related work

Our work has been partially inspired by a recent contribution within the field of physical security [7], 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*.

As the purpose of cyber security investments methodologies is to lead to the selection of a set of cyber security controls that maximise the benefit of an organisation with respect to some available budget, we find papers that investigate this optimal selection [8–12] as the most relevant to our work. In this section, we summarise the most prominent works that investigate allocation of a cyber security budget after conducting cyber security risk assessment. Their differences are briefly also summarised in Table 1.

One of the initial works studying the way to model investment in cyber security is published by Gordon and Loeb [13]. 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.

Inspired by [14], Lee et al. apply the profit-at-risk and operational risk modelling approaches to propose a model that facilitates optimal customer information security investments by providing undertaking trade-off analysis between risk and return [15]. The authors define a minimum information security protection level that must be achieved in order for investments in a customer privacy protection to be effective. Rakes et al. [10] extended previous mathematical models [16] to develop an integer programming model that optimises the selection of a subset of security controls to mitigate certain threat level profiles. Authors assessed their model under expected and worst-case threat levels towards deriving tradeoffs for optimal security planning between these two threat levels. They also demonstrated budget-dependent risk curves giving emphasis in showing how perturbed budget levels affect the aforesaid tradeoffs. In a similar vein, Viduto et al. [11] formulate a multi-objective optimisation problem to select security controls in a cost-effective manner taking into account both financial cost and security risk. Inspired by [10], Sawik [12] applies two popular in financial engineering (e.g. in portfolio management) measures of risk: value-at-risk and conditional value-at-risk.

In [17], Nagurney et al. propose a supply chain network game theoretic model in which retailers may be subject to a cyberattack and seek to maximise their expected profits by selecting their optimal product transactions and cybersecurity levels. A successful attack likelihood depends not only on the security level of the retailer per se, but also on that of the other retailers. Authors also show how cyber security investment cost functions vary according to consumers' preferences for the product, which, in turn, depends on both the demand and the security level. Srinidhi et al. [18] propose an optimisation model to reason about the allocation of cyber security resources to assets that have inherent strength against cyber attack and security-enhancing assets (i.e. security controls). They also investigate the role of cyber insurance in mitigating the effects of breach costs as well as the incentives that both managers and investors in spending upon cyber security products given that the first (i.e. managers) are more concerned with potential financial losses while the second (i.e. investors) are reluctant to spend more in strengthening the firm's security due to spreading their risks by investing in different firms. Lastly, Cavusoglu et al. [19] compare a decision-theoretic approach to game-theoretic approaches for investment in cyber security. Authors neither use real world data to undertake their risk assessment nor do they investigate the optimal selection of security controls.

3. Model definition

In this section we use game theory to model the interactions between two players; the Defender (\mathcal{D}) and the Attacker (\mathcal{A}). The Defender might be 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, \mathcal{A} is a cyber hacker who tries to subvert the system to her own end, by launching commodity cyber attacks against the organisation \mathcal{D} 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.

In our model, \mathcal{D} has 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 \mathcal{D} 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. Our work is based on quantitative risk assessment prior to deciding upon spending a cyber security budget. Alpcan [20] (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. In the following we discuss the different components of the model, and we define appropriate terminology and notations, which are consistent throughout this article.

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 are located at the same are depth worth the same value to \mathcal{D} 's firm.

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 $t_i := (v_z, d_x)$. This 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 \mathcal{D} 's organisation when \mathcal{A} 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. Aligned with the definition of target, $I(t_i)$ depends on the depth d_x , and $T(t_i)$ depends on the vulnerability v_z .

A *cyber security control* is the defensive mechanism that \mathcal{D} can put in place to alleviate the risk from one or more attacks by reducing the probability of these attacks successfully exploiting vulnerabilities. \mathcal{D} chooses to implement a control at a certain level for their organisation. We define the set of implementation levels of a control as $\mathcal{L} := \{l_\lambda\}$. The higher the level the greater the degree to which the control is implemented¹. 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 \mathcal{A} has the potential to exploit. Assume \mathcal{D} 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 \mathcal{D} suffers when t is attacked and a control has implemented at level l . This definition of loss is in line with the well-known formula, risk = expected damage ($I(t)$) \times probability of occurrence ($T(t)$) [21].

¹ Note that we abuse notation by setting $l_\lambda = l, t_i = t, v_z = v$, and $d_x = d$.

While the implementation of a cyber security control strengthens the defence of \mathcal{D} 's organisation, it is associated with two types of costs namely; *indirect* and *direct*. 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 \mathcal{D} 's organisation. This is equivalent to the sum of the security damages inflicted by \mathcal{A} and the indirect cost for implementing a control at a certain level. Formally, when \mathcal{D} implements a control at some level i then the expected loss of their organisation is derived by $\sum_t S(i, t) - C(i)$. 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.

4. Cyber security control games

In this article we formulate two-player non-cooperative static games, called *Control Games*. The players in a Control Game are the Defender (\mathcal{D}), which represents any cyber security decision-maker, and \mathcal{A} , which represents any cyber hacker who uses commodity attacks. The former defends their organisation's data assets by minimising expected cyber security losses with respect to some indirect costs, while \mathcal{A} aims at benefiting from compromising these assets. \mathcal{D} is choosing how to implement a cyber security control (i.e. at which level) and \mathcal{A} decides which target to attack to exploit its vulnerability at a certain depth. Since we consider a simultaneous game, \mathcal{A} does not know the control implementation strategy and \mathcal{D} does not know the attack strategy. We refer to our games as *Control Games* because the basis of our formulation is that \mathcal{D} has one control at her disposal.

First, we formulate zero-sum Control Games. These represent scenarios where \mathcal{A} aims at causing the maximum possible damage to \mathcal{D} . 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 \mathcal{A} when undertaking an attack. However cost in terms of cyber attacks is tightly coupled with the adversary's profile. A consideration of a specific threat model would also have some influence on the way \mathcal{A} 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 \mathcal{A} 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 \mathcal{A} 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 \mathcal{A} can choose from but we have not postulated anything about their motivations.

For a given cyber security control, \mathcal{D} can choose to implement the control at level $l \in \mathcal{L}$ and therefore her pure strategy set coincides with \mathcal{L} . \mathcal{A} selects a vulnerability to exploit at a certain depth. Formally, \mathcal{A} chooses $t = \langle v, d \rangle \in \mathcal{T}$. Thus the pure strategy set of \mathcal{A} coincides with \mathcal{T} . Given that the pure strategy sets of the players are \mathcal{L} and \mathcal{T} then \mathcal{D} has m pure strategies and \mathcal{A} has n , correspondingly. We denote by $G := \langle \mathbb{A}, \mathbb{E} \rangle$ an $m \times n$ bi-matrix cyber security game where \mathcal{D} (i.e. row player) has a payoff matrix $\mathbb{A} \in \mathbb{R}^{m \times n}$ and the payoff matrix of \mathcal{A} (i.e. the column player) is denoted by $\mathbb{E} \in \mathbb{R}^{m \times n}$. \mathcal{D} 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) , \mathcal{D} and \mathcal{A}

have payoff values equivalent to a_{lt} and e_{lt} , given by $a_{lt} := S(l, t) - C(l)$ and $e_{lt} := -S(l, t) + C(l)$. Table 2 is the game matrix presenting player's payoffs for the different pure strategy profiles.

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

We occasionally refer to the implementation of a control at a certain level as a *cyber security process*. We can then define the *cyber security plan* as the 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.

We define \mathcal{D} 's mixed strategy as the probability distribution $\Phi = [\phi_1, \dots, \phi_m]$. This expresses a cyber security plan, where ϕ_λ is the probability of implementing the control at l_λ . A cyber security plan can be realised as advice to \mathcal{D} 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 \mathcal{D} . We claim that our model is flexible thus allowing \mathcal{D} to interpret mixed strategies in different ways to satisfy their requirements. A mixed strategy of \mathcal{A} is a probability distribution over the different targets and it is denoted by $\Theta = [\theta_1, \dots, \theta_n]$, where θ_i is the probability of the adversary attacking t_i . When both players choose a pure strategy randomly according to the probability distributions determined by Φ and Θ , the expected payoffs to \mathcal{D} and \mathcal{A} are $J_D(\Phi, \Theta) := \sum_{i=1}^n \sum_{\lambda=1}^m \phi_\lambda a_{i\lambda} \theta_i$, and $J_A(\Phi, \Theta) := \sum_{i=1}^n \sum_{\lambda=1}^m \phi_\lambda e_{i\lambda} \theta_i$.

For the remainder of this section, we analyse a specific Control Game. We assume that for a specific target t , \mathcal{D} 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 $\Delta S(t) := S(l', t) - S(l, t)$ and $\Delta C := C(l') - C(l)$. $\Delta S(t)$ is the reduction in damage when l' is chosen, and ΔC 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, \mathcal{D} 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, $S(l', t) - S(l, t) > C(l') - C(l) \iff S(l', t) - C(l') > S(l, t) - C(l) \iff a_{l't} > a_{lt}$. Therefore, the \mathcal{D} is incentivised to pick l' as it has a higher utility. \square

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' , \mathcal{A} 's best response can be found by comparing $e_{lt}, e_{l't}$. If $e_{lt} > e_{l't} \iff S(l, t) - C(l) > S(l, t') - C(l) \iff S(l, t) > S(l, t')$, \mathcal{A} prefers to attack target t . Specifically we define this property as $\Delta S(l) := S(l, t') - S(l, t)$. Therefore, if $S(l, t) > S(l, t') \iff S(l, t') - S(l, t) < 0 \iff \Delta S(l) < 0$, \mathcal{A} chooses t . \square

Since we are investigating a two-person zero-sum game with finite number of actions for both players, and according to Nash [22] it admits at least a Nash Equilibrium (NE) in mixed strategies. Saddle-points correspond to Nash equilibria as discussed in [20].

Table 2
Game matrix.

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

The following result, from [23], establishes the existence of a saddle (equilibrium) solution in the games we examine and summarises their properties.

The investigated cyber security game admits a saddle point in mixed strategies, (Φ^*, Θ^*) , with the properties $\Phi^* = \arg \max_{\Phi} \min_{\Theta} J_U(\Phi, \Theta)$, $\Theta^* = \arg \max_{\Theta} \min_{\Phi} J_A(\Phi, \Theta)$.

Corollary 3. *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 states that for zero sum games NE, maxmin and minimax solutions coincide. Therefore $\Phi^* = \arg \min_{\Phi} \max_{\Theta} J_A(\Phi, \Theta)$. \square

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 cyber security games can be mathematically reduced to zero sums games. Remark 1 is addressed by the following lemma.

Lemma 4. *The equilibrium (Φ^*, Θ^*) in our zero sum cyber security 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 \mathcal{A} where his payoffs are a negative affine transformation of the \mathcal{D} 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 ω is a negative scalar, and ψ is a constant matrix of the same dimension as \mathbb{A} . Therefore, in addition to the cyber security game $G = (\mathbb{A}, -\mathbb{A})$, we intuitively define the negative affinity of this game as $G^- = (\mathbb{A}, \omega \mathbb{A} + \psi)$.

Suppose Φ^*, Θ^* are the equilibrium strategies in G . First, it is easy to see that Φ^* 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 Θ^* 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^- . \square

Lemma 5. *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} 's payoff is given by $S - (\kappa S - \mu) = S(1 - \kappa) - \mu$, where κ, μ are positive scalars. Assume that at the equilibrium of \hat{G} 's best response is Φ^* . Then we have $\Phi [S(1 - \kappa) - \mu] \Theta^* \leq \Phi^* [S(1 - \kappa) - \mu] \Theta^* \Rightarrow \Phi (S - \kappa S - \mu) \Theta^* \leq \Phi^* (S - \kappa S - \mu) \Theta^* \Rightarrow \Phi (S - \mu) \Theta^* \leq \Phi^* (S - \mu) \Theta^* \xrightarrow{H=C} \Phi (S - C) \Theta^* \leq \Phi^* (S - C) \Theta^* \Rightarrow \Phi \mathbb{A} \Theta^* \leq \Phi^* \mathbb{A} \Theta^*$. Therefore G, \hat{G} have the same equilibria, and from Corollary 3 these are also maxmin solutions. \square




Fig. 1. Game tree for the control game with 2 implementation levels and two targets.

To illustrate the game approach let's consider a toy example consisting of a 2-level, and 2-target control games, where \mathcal{D} and \mathcal{A} make their decisions simultaneously, or, equivalently, independently of each other. The information sets associated with the control game, investigated in this section, depicted in Fig. 1; a dashed curve encircling the \mathcal{A} nodes has been drawn. This indicates that \mathcal{A} cannot distinguish between these two points. In other words, \mathcal{A} has to arrive at a decision without knowing what \mathcal{D} has actually chosen.

Due to the game being zero-sum, we have kept only the payoffs of \mathcal{D} at the game tree. We also defined the mixed strategy of \mathcal{D} as the probability distribution $[\phi, 1 - \phi]$, where ϕ is the probability of implementing the control at level l . \mathcal{A} 's mixed strategy is denoted by $[\theta, 1 - \theta]$, where \mathcal{A} chooses to attack t with probability θ . Table 3 summarizes all possible best responses of the control game for the different conditions discussed in this section.

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

Lemma 6. *The Nash equilibrium for a control sub-game for the \mathcal{D} 's, given by $\phi^* \in [0, 1]$ is $\phi^* = \frac{\Delta S(l')}{\Delta S(l') - \Delta S(l)}$.*

Proof. The \mathcal{D} wants to make the \mathcal{A} indifferent to which target they should attack. This is given by equalising the expected payoff of the \mathcal{A} , thus $A(t) = \phi^* e_{lt} + (1 - \phi^*) e_{l't}$ and $A(t') = \phi^* e_{lt'} + (1 - \phi^*) e_{l't'}$, giving

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

We can substitute terms such that Eq. (1) can be written in terms of e_{lt} , hence $e_{l't} = e_{lt} - \Delta S(t) + \Delta C$, $e_{l't'} = e_{lt} - \Delta S(l)$, and $e_{l't} = e_{lt} - \Delta S(t) + \Delta C - \Delta S(l')$. By substituting the above equations into Eq. (1) we get $\phi^* e_{lt} + (1 - \phi^*) (e_{lt} - \Delta S(t) + \Delta C) = \phi^* (e_{lt} - \Delta S(l)) + (1 - \phi^*) (e_{lt} - \Delta S(t) + \Delta C - \Delta S(l')) \Rightarrow \Delta S(l') = \phi^* (\Delta S(l') - \Delta S(l)) \Rightarrow \phi^* = \frac{\Delta S(l')}{\Delta S(l') - \Delta S(l)}$. \square

Lemma 7. *The Nash strategy of the \mathcal{A} in a control sub-game, is given by $\theta^* = \frac{\Delta S(t) - \Delta C + \Delta S(l') - \Delta S(l)}{\Delta S(l') - \Delta S(l)}$.*

Table 3
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')

Proof. At the equilibrium, the \mathcal{A} wants to make the \mathcal{D} indifferent to which target they should attack. By equalising the expected payoff of the \mathcal{D} we have that $D(l) = \theta^* a_{lt} + (1 - \theta^*) a_{lt'}$ must equal $D(l') = \theta^* a_{l't} + (1 - \theta^*) a_{l't'}$. Therefore $\theta^* a_{lt} + (1 - \theta^*) a_{lt'} = \theta^* a_{l't} + (1 - \theta^*) a_{l't'}$. We can substitute terms such that the above equation can be written in terms of a_{lt} , $a_{l't} = a_{lt} + \Delta S(t) - \Delta C$, $a_{lt'} = a_{lt} + \Delta S(l)$, and $a_{l't'} = a_{l't} + \Delta S(t) - \Delta C + \Delta S(l')$, and by substituting these equations into the former we get $\theta^* a_{lt} + (1 - \theta^*) (a_{lt} + \Delta S(l)) = \theta^* (a_{lt} + \Delta S(t) - \Delta C) + (1 - \theta^*) (a_{lt} + \Delta S(t) - \Delta C + \Delta S(l')) \Rightarrow a_{lt} + \Delta S(l) - \theta^* \Delta S(l) = a_{lt} + \Delta S(t) - \Delta C + \Delta S(l') - \theta^* \Delta S(l') \Rightarrow \theta^* = \frac{\Delta S(t) - \Delta C + \Delta S(l') - \Delta S(l)}{\Delta S(l') - \Delta S(l)}$. \square

5. Cyber security control subgames

When looking into investing in cyber security one might face the challenge of not having a necessary financial budget to implement the equilibria of a cyber security Control Game. To tackle this challenge we define cyber security Control Subgames, which constitute a Control Game by gradually increasing the available implementation levels of the control. In this way, we can derive a number of equilibria that can satisfy a wider range of financial capacity. A Control Subgame $\mathcal{G}_{j\lambda}$ is a game where (i) \mathcal{D} 's pure strategies correspond to consecutive implementation levels of the control j starting always from 0 (i.e. fictitious control-game) and including all levels up to λ and, (ii) \mathcal{A} 's pure strategies are the different targets akin to pairs of vulnerabilities and depths.

We represent \mathcal{D} 's mixed strategy, in $\mathcal{G}_{j\lambda}$, as the probability distribution $Q_{j\lambda} = [q_{j0}, \dots, q_{j\lambda}]$. This expresses a *cybersecurity plan*, where q_{jl} is the probability of implementing c_j at level l in $\mathcal{G}_{j\lambda}$. A mixed strategy of \mathcal{A} is defined as a probability distribution over the different targets, in $\mathcal{G}_{j\lambda}$, and it is denoted by $H_{j\lambda} = [h_{j1}, \dots, h_{jn}]$, where h_{ji} is the probability of the adversary attacking t_i when \mathcal{D} has only c_j in their possession. \mathcal{D} 's aim in a Control Subgame is to choose the *Nash cybersecurity plan* $Q_{j\lambda}^* = [q_{j0}^*, \dots, q_{j\lambda}^*]$. This consists of λ cybersecurity processes chosen probabilistically as determined by the NE of $\mathcal{G}_{j\lambda}$.

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.

6. Cyber security investment

The analysis performed in Section 4 has considered a single-control, two-targets, two-levels game. When having c cyber security controls, our plan for cyber investment is to solve c Control Games by splitting each of them up to a set of $m - 1$ Control Subgames with n targets and up to λ implementation levels for each control, where $\lambda \in \{1, m\}$. For a Control Game the Control Subgame equilibria constitute the Control Game solution.

Given the Control Subgame equilibria we then use a Knapsack algorithm to provide the general investment solution. The equilibria 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 \mathcal{A} 's strategy, and the indirect costs of the organisation.

It is easy to see that, Control Subgames (and consequently Control Games) look only at 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 select 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 Control Subgame 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 depth 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, as part of a risk assessment methodology. 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.

6.1. 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. \mathcal{D} 's pure strategies are comprised of an implementation level for each of the controls, and \mathcal{A} '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" [24], 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, and 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 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.

6.2. Hybrid method

The Hybrid method avoids the problems of the Full Game method by considering the particular game solutions of each Control Game (and consequently the game solutions of all Control Subgames that comprise this Control Game) as part of an overall combinatorial optimisation problem which we will solve using 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.

For convenience, we denote the Control Subgame solution by the maximum level of implementation available. For instance, for c_j and the solution of Control Subgame $\mathcal{G}_{j\lambda}$ is denoted by $Q_{j\lambda}^*$. Let us assume that for control j the equilibria of all Control Subgames are given by

the set $\{Q_{j0}^*, \dots, Q_{jm}^*\}$. For each control there exists a unique Control Subgame solution Q_{j0} , which dictates that control j should not be used.

We define an optimal solution to the Knapsack problem as $\Psi = \{Q_{j\lambda}^*, \forall j \in \{1, \dots, c\}, \forall \lambda \in \{1, \dots, m\}\}$. A solution Ψ takes exactly one solution (i.e. equilibrium or cyber security plan) for each control as a policy for implementation. To represent the cyber security investment problem, we need to expand the definitions for both expected damage S and effectiveness E to incorporate the Control Subgame solutions. Hence, we expand S such that $S(Q_{j\lambda}, t)$, which is the expected damage on target t given the implementation of $Q_{j\lambda}$. Likewise, we expand the definition of the effectiveness of the implemented solution on a given target as $E(Q_{j\lambda}, t)$. Additionally, we consider $\Gamma(Q_{j\lambda})$ as the direct cost of implementing $Q_{j\lambda}$.

If we represent the solution Ψ by the bitvector z , we can then represent the 0–1 Multiple Choice, Multi-Objective Knapsack Problem as:

$$\begin{aligned} & \max_z \min_{t_i} \left\{ \left[1 - \sum_{j=1}^c \sum_{\lambda=0}^m E(Q_{j\lambda}, t_i) z_{j\lambda} \right] I(t_i) T(t_i) \right\} \\ & \text{s.t. } \sum_{j=1}^c \sum_{\lambda=0}^m \Gamma(Q_{j\lambda}) z_{j\lambda} \leq B \\ & \sum_{\lambda=0}^m z_{j\lambda} = 1, z_{j\lambda} \in \{0, 1\}, \forall j = 1, \dots, c. \end{aligned}$$

where B is the available cyber security budget, and $z_{j\lambda} = 1$ when $Q_{j\lambda}^* \in \Psi$. 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 than would produce the same net effect.

6.3. Pure Knapsack representation

A Pure Knapsack representation considers the method of solving the investment problem given that \mathcal{D} only considers the implementation of “whole” controls. This is to say that the solutions supplied to the Knapsack solver are representative of pure strategies solutions to the Control Subgames. To do this in a fair manner, we need to include the indirect costs of each cyber security plan (i.e. Control Subgame solution) into the calculation of benefit from each target. This is because the Hybrid representation has taken into account the impact of the indirect costs in the Control Subgames. We first extend the definition of indirect cost to incorporate Control Subgame solutions. Thus, we expand C such that $C(Q_{j\lambda})$, which is the indirect cost of $Q_{j\lambda}$. Thereafter, we can extend the representation of the Knapsack problem to include the indirect costs as follows:

$$\begin{aligned} & \max_z \min_{t_i} \left\{ \left[1 - \sum_{j=1}^c \sum_{\lambda=0}^m E(Q_{j\lambda}, t_i) z_{j\lambda} \right] I(t_i) T(t_i) - \sum_{j=1}^c \sum_{\lambda=0}^m C(Q_{j\lambda}) z_{j\lambda} \right\} \\ & \text{s.t. } \sum_{j=1}^c \sum_{\lambda=0}^m \Gamma(Q_{j\lambda}) z_{j\lambda} \leq B \\ & \sum_{\lambda=0}^m z_{j\lambda} = 1, z_{j\lambda} \in \{0, 1\}, \forall j = 1, \dots, c. \end{aligned}$$

6.4. Comparison of methods

To compare the Full Game, Hybrid and Pure Knapsack methods of decision support, we have developed a small case study that represents a small defence decision making problem that might be seen




Fig. 2. Comparison of Full Game, Hybrid and Pure Knapsack Methods of Decision Support.

by system administrators. The problem creates an example with 7 controls and 13 vulnerabilities, created using a mapping from the SANS Critical Security Controls [25] combined with the CWE Top 25 Software Vulnerabilities [26]. The case study presented in this work considers a network separated into three different depths (i.e. Demilitarized Zone, Intranet, and Private Subnet). For this example, we consider the levels available to \mathcal{D} to consist of the quick win processes provided by SANS.

In comparing the damage at the weakest target provided by the Full Game and Hybrid method to the Knapsack representation, we can see in Fig. 2 that, in general, the Full Game representation will provide a better defence to the weakest target for low budget levels. This is due to the Full Game representation being able to combine all controls in a more flexible manner than either the Hybrid or Pure Knapsack, 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. Additionally, the Hybrid is occasionally able to offer a better solution than the Pure Knapsack, because the mixed strategies allow for certain control combinations to be used at a lower budget.

Each of the methods eventually reach a similar stable value, where although there is still damage expected from attacks against the system, the additional cost to the performance of the system and users do not outweigh the benefit of implementing the additional control. This is owing to the impact that the indirect cost has on the decision-making process, where the cost is added to the damage to create the utility.

In the case where there are no indirect costs to implementing each control there is no trade-off to achieve a higher defence. This means that providing that an appropriate budget is available, the best defence will be purchased by all methods. In this case the solutions to the Control Subgames, in the Hybrid method, become the same as the pure strategies used in the Pure Knapsack and the resulting optimal solutions are the same.

We have found in terms of complexity of solutions provided, that 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 Subgames lead to a single strategy at many levels. However, we find that there is an additional

Table 4
Case study vulnerability type distribution.

Vulnerability type	No.	Vulnerability type	No.
Software errors: data interaction	8	Software errors: resource management	8
Software errors: defence flaws	6	Social engineering: targeted	3
Social engineering: untargeted	4	Network vulnerabilities: direct	3
Network vulnerabilities: indirect	4	–	–

level of complexity in the comprehension of the strategies that are produced by the Full Game. Such complexity can potentially lead to strategies that cannot 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.

In addition to the comparison above, we tested the example case presented by Rakes et al. [10] to ensure that the optimisation algorithms used for solving the Hybrid and Knapsack problem were acting correctly. The study could be rebuilt using a reduced set of values from our model, which included the removal of indirect costs and the assumption of only one level of implementation for each control. In this example, we are able to obtain the same optimal set of countermeasures as the authors present in their work with a higher than 95% success rate on tested cases.

While we have seen that the Full Game representation of the problem is the most optimal on a small scale, the practicalities of operating such a system on a larger scale is not possible. The next section details a more realistic case study consisting of 27 different controls acting on 36 different attacks, which is not feasibly solved by a Full Game representation. The challenge behind the Full Game representation of such a large case study, is that with multiple levels the number of pure strategies is of the magnitude 10^{15} ; however, this is not a challenge that is faced by the Hybrid representation, which does not need to represent each pure strategy to calculate a solution, only up to the maximum number of levels in a Control Subgame.

7. Case study

We have further used the ideas presented to develop a decision support tool that is capable of working on more realistic scenarios. The role of this tool is to be able to offer realistic actionable advice to organisations. The following represents a case study based on the design of a typical SME network, with the data used in the representation of the attacks, controls and costs for this case study available online [27].

7.1. Case study composition

The attacks have been built from a subset of the CVE [28] and CWE, which a conventional networked system would be expected to face, as well as a number of social engineering based attacks. The distribution of attacks amongst certain kinds can be seen in Table 4. The factors that are associated with the CWSS have been used to formulate the basis of the values for the vulnerabilities. There are two differences between the CWSS scoring system and the one used for this study. The first is the isolation of threat factors, since we are interested in the ability of an organisation to be able to identify their own concerns regarding the impact of a successful attack.

While a number of factors have been removed, a number of additional factors have been included to better differentiate different attacks. This has also provided a more generic insight into the decision making process of the attacker. Critically this involves the identification, availability and ease of the attacks for them to perform. This is done to indicate a partial reduction in risk of certain attacks, while making those that are easier to perform more enticing

for the attacker. These are designed in such a way as to work not only with the attacker decision making in the Control Subgames, but also affect the designation of the potential weakest target in the optimisation. This aids in shifting some risk to the requirement of attacker capability.

The controls used in this case study are a set of actionable controls that a system administrator can implement to improve the security of their network. The controls have been derived from the SANS Top 20 critical security controls, separating the overarching control advice, to better reflect a single point of investment. The controls cover a variety of types of defensive strategy, the distribution of which can be seen in Table 5.

The CVSS and CWSS both contain details regarding the efficiency of controls for protecting against a particular vulnerability. This internal definition of control effectiveness against each attack does not support our model for optimal defence spending. As such the effectiveness was redesigned to identify which controls can mitigate which vulnerabilities, spread the efficiency amongst the viable levels, and interpret the viability of the control over the life of the solution, based either on complexity or frequency.

Each organisation is likely to have different configurations of systems and sizes and this makes defining the costs, in terms of a direct financial value, difficult. An over specialised budget requirement would make using the tool infeasible in the real world. To remedy this we have normalised the direct costs of the control, such that the implementation of a number of controls operating at a conventional level from the advice has a cost of 1; an example of this is weekly patching.

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. To do this, we define the indirect cost as an expected level of additional disruption caused in one of three categories: *System Performance*, any reduction in the speed and capability of the system to perform the related business tasks; *User Morale*, the impact of the control on the behaviour of the system users; and *Re-Training*, the additional requirements for users of the system to be able to use the new control.

7.2. Experimental comparison

The UK has published a set of guidelines that organisations, similar to the one in the case study, should comply with in order to reduce the risk of damage from cyber attacks [3]. The document called Cyber Essentials suggests a number of basic controls that organisations should implement to protect themselves from cyber attacks. The controls considered by Cyber Essentials are the use of firewalls and gateways, user access control, secure configuration, malware protection and patch management.

Table 5
Case study control type distribution.

Control type	No.	Control type	No.
Security software	4	Network security tools	7
System configuration	8	Administration tools	2
Policy development	4	Education and training	2

To perform this comparison, we have set a number of budget points in our control data, which represent different levels of investment. The first test case presents the scenario which accounts for the lowest price that allows for the implementation of Cyber Essentials at the lowest level. The second budget value allows for each of the controls from Cyber Essentials to be implemented at the highest level, while the final budget considers the availability of a higher level of investment beyond the advice offered by Cyber Essentials.

We are interested in investigating if controls in positions, 2, 4, 5, 7, 9, 20 and 23 are recommended, as they relate to those suggested by Cyber Essentials. Both controls 5 and 7 relate to the implementation of firewalls, where control 5 is for network firewalls and control 7 relates to web application firewalls. Likewise controls 20 and 23 relate to User Access Controls, with User Access Controls representing the access policies concerning the network and devices in control 20 and Account Management Control tools for users implemented in control 23.

7.3. Results

The following section describes the results obtained from calculating the optimal defence strategy for the case study outlined. The results shown here are obtained using an implementation of the hybrid model solved using a genetic algorithm.

The lowest budget shown in Table 6 chooses to implement four out of the five controls outlined in Cyber Essentials, with two of the controls being preferred at a higher level than the implementation of an additional control. The optimal solution suggests implementing Patch Management and Network Firewalls at level two, with Anti-Malware and Secure Configurations both suggested at the most basic level. Under this limited budget, it is not suggested to have either Web Application Firewalls or User Access Controls, over the higher levels of other controls. In addition to the Cyber Essentials controls, the optimal solution suggests the implementation of an Incident Response Policy, which covers predominantly social engineering targets, with some minimal effect, but has a small cost.

When the budget is increased to 8, we see that the initial four controls in Cyber Essentials represented in the previous optimal solution are still represented, but with three controls recommended at a higher level. The optimal solution is to perform Patch Management at the highest level, such that it should be performed on demand. In this context it means that patches should be checked on a daily basis and implemented as soon as possible.

This budget brings in Account Management Control as a recommendation, which represents the last of the controls recommended by Cyber Essentials. The implementation recommends a strict account management control system, limiting the potential use of accounts, which reduces the risk of attacks being able to escalate privileges or access sensitive files with hijacked low level accounts. Additionally, we see the introduction of Web Application Firewalls in addition to the Network Firewalls that were suggested before, with a strict level of implementation. Since more of the vulnerabilities are covered to a higher level, the impact of the relatively cheap Incident Report Handling control has been removed as its addition has too minimal an impact to justify the cost. The controls outside of Cyber Essentials that are now considered are Automated Inventory Scanning and Management and basic Intrusion Detection Systems.

For a budget of 16, given by the solution 16a in Table 6, the solution differs from the previous budgets only in a few aspects. In terms of Cyber Essentials the optimal solution recommends implementing Network Firewalls at a lower level, however the solution still maintains the strict implementation of Web Application Firewalls. The other controls proposed by Cyber Essentials that had been recommended remain implemented at the same level as with a budget of 8. Inventory Management tools, are now implemented at a higher level than previously seen, moving from yearly inventory logging to weekly. The optimal solution now recommends the use of Intrusion Prevention Systems instead of Intrusion Detection Systems, which operate to cover more vulnerabilities than Network Firewalls, but were more costly to implement, making them less viable at lower budgets. Another additional control suggested at this budget level is the inclusion of yearly User Education and Training, which is used to improve a number of social engineering based attacks.

7.4. Discussion

The results from the experimentation show with some consistency that the controls associated with Cyber Essentials are appropriate defensive measures for this kind of network. At low budgets, the system recommends implementing a number of controls that are suggested by Cyber Essentials, but not all of them, preferring to offer a more stable configuration of these controls over adding additional controls. At a higher budget, we see that the remaining controls are considered, with them being used beyond a basic level of implementation.

In all the cases presented the implementation of a rigorous Patching policy is recommended where possible, as well as the presence of some Anti-malware, Firewalls and Secure Configuration. The main thing that can be observed from the data is that a combination of all of these four controls covers each of the vulnerabilities in the case study to some degree. This means that by increasing the level of any of these controls, there is guaranteed to be some observable reduction in damage on the system.

To follow this, one of the observations made throughout the experimentations the impact that the indirect costs have on the decision to implement certain security controls. The crucial component, is that as has been noted, a set of four controls are able to cover all the vulnerabilities to some level of efficiency; this means that the implementation of an additional control will only serve to reduce the impact of a vulnerability by a fraction of its maximum efficiency, while the costs for that control remain the same. As such there is a diminishing return on each control that you add to the system, which means that after certain values, it makes it more costly to the organisation to implement the control against the additional risk that they might mitigate.

Having seen how important the indirect costs are to calculating the optimality of the advice given, we have looked at the impact of a reduction in indirect costs. For this we have taken the highest budget level, which in the previous example was not using the whole of the budget due to indirect costs and have reduced the impact of indirect costs to 0.1 of their previous values.

The suggested implementation of controls, given by the solution 16b in Table 6, changes the optimal strategy to introduce a series of new controls in addition to those seen previously. Even with a lower importance on indirect costs, we see that the optimal solution still recommends the implementation of the Cyber Essentials controls suggested in the initial tests.

Additional government advice suggest the use of Whitelisting, which is not seen in the initial solutions. While whitelisting of both applications (control 19) and websites (control 21) are able to prevent a number of cyber attacks by preventing access, they have a high negative impact on the daily operations of the organisation. This results in a high indirect cost, which reflects their exclusion

Table 6
Case study results.

Budget	Solution
3	[0, 2, 0, 1, 2, 0, 0, 0, 1, 0, 1, 0]
8	[1, 5, 0, 4, 2, 0, 3, 0, 3, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 3, 0, 0, 0, 0, 0]
16a	[4, 5, 0, 4, 1, 3, 3, 0, 3, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 3, 1, 0, 0, 0, 0]
16b	[2, 5, 0, 3, 2, 2, 3, 0, 3, 0, 0, 0, 0, 0, 0, 1, 0, 1, 4, 0, 2, 0, 3, 1, 0, 0, 1]

in the previous optimal solutions. However, with relaxed indirect costs, the negative impact now no longer outweighed by the benefit, which now represents their inclusion in the optimal solution. The same advice recommends penetration testing if possible if you are expecting to be at risk of more long term attacks, this control is also recommended in the solution with revised indirect costs.

8. Conclusions

In this paper we have presented an analysis of a hybrid game-theoretic and optimisation approach to the allocation of an SME's cyber security budget. For this purpose, we have compared three different approaches to allocating this budget by using a decision support tool. 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. In addition, we aim to investigate cyber security investments by following a multidisciplinary approach that combines economic, behavioural, societal and engineering insights. Our end goal is to achieve increased societal resilience to cyber security risks through more efficient and effective institutional and incentives structures. Last but not least, in future work we aim to investigate how cyber insurance [2] can influence cyber security investment decisions.

References

- [1] 2014 Deloitte NASCIO Cybersecurity Study, 2014, http://www.nascio.org/publications/documents/Deloitte-NASCIOCybersecurityStudy_2014.pdf.
- [2] P.H. Meland, I.A. Tondel, B. Solhaug, Mitigating risk with cyberinsurance, *IEEE Security Privacy* 13 (6) (2015) 38–43.
- [3] H. Government, Cyber Essentials Scheme Requirements for Basic Technical Protection from Cyber Attacks, accessed: 2015-12-19 (June 2014), https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/317481/Cyber_Essentials_Requirements.Pdf.
- [4] R. Anderson, Why information security is hard, in: Proceedings of the 17th Annual Computer Security Applications Conference, New Orleans, Louisiana, pp. 2001.
- [5] J. Grossklags, N. Christin, J. Chuang, Secure Or insure?: a game-theoretic analysis of information security games, Proceedings of the 17th International Conference on World Wide Web, ACM, Beijing, China, 2008, pp. 209–218.
- [6] F.Smeraldi, Cybersecurity games and investments: a decision support approach, Decision and game theory for security, Springer 2014, pp. 266–286.
- [7] M.Tambe, Stackelberg vs. nash in security games, an extended investigation of interchangeability, equivalence, and uniqueness, *J. Artif. Intell. Res.* 41 (2011) 297–327.
- [8] J.Chi, Matching information security vulnerabilities to organizational security profiles: a genetic algorithm approach, *Decis. Support. Syst.* 41 (3) (2006) 592–603.
- [9] L.P. Rees, J.K. Deane, T.R. Rakes, W.H. Baker, Decision support for cybersecurity risk planning, *Decis. Support. Syst.* 51 (3) (2011) 493–505.
- [10] T.R. Rakes, J.K. Deane, L.P. Rees, IT Security planning under uncertainty for high-impact events, *Omega: International Journal of Management Science* 40 (1) (2012) 79–88.
- [11] D.López-Peréz, A novel risk assessment and optimisation model for a multi-objective network security countermeasure selection problem, *Decis. Support. Syst.* 53 (3) (2012) 599–610.
- [12] T. Sawik, Selection of optimal countermeasure portfolio in it security planning, *Decis. Support. Syst.* 55 (1) (2013) 156–164.
- [13] L.A. Gordon, M.P. Loeb, The economics of information security investment, *ACM Trans. Inf. Syst. Secur.* 5 (4) (2002) 438–457.
- [14] R.Sougstad, R.J. Kauffman, Risk management of contract portfolios in it services: the profit-at-risk approach, *J. Manag. Inf. Syst.* 25 (1) (2008) 17–48.
- [15] R.Sougstad, Y.J. Lee, R.J. Kauffman, Profit-maximizing firm investments in customer information security, *Decis. Support. Syst.* 51 (4) (2011) 904–920.
- [16] J.K. Deane, C.T. Ragsdale, T.R. Rakes, L.P. Rees, Managing supply chain risk and disruption from IT security incidents, *Oper. Manag. Res.* 2 (1–4) (2009) 4–12.
- [17] S.Shukla, A Supply Chain Game Theory Framework for Cybersecurity Investments Under Network Vulnerability, Springer International Publishing 2015, pp. 381–398.
- [18] J.Yan, G.K. Tayi, Allocation of resources to cyber-security: the effect of misalignment of interest between managers and investors, *Decis. Support. Syst.* 75 (2015) 49–62.
- [19] T.Wei, Decision-theoretic and game-theoretic approaches to it security investment, *J. Manag. Inf. Syst.* (2008) 281–304.
- [20] T.Basar, *Network Security: A Decision and Game-theoretic Approach*, Cambridge University Press 2010.
- [21] R.Oppliger, Quantitative risk analysis in information security management: a modern fairy tale, *IEEE Security Privacy* 13 (6) (2015) 18–21.
- [22] J.Nash, Equilibrium points in N-person games, *Proc. of the National Academy of Sciences*, 1950, pp. 48–49.
- [23] T.Basar, G.J. Olsder, *Dynamic Noncooperative Game Theory*, 1995.
- [24] J.Tsai, C.Kiekintveld, F.Ordonez, M.Tambe, S.Rathi, IRIS—a tool for strategic security allocation in transportation networks, in: Proc. of 8th International Conference on Autonomous Agents and Multiagent Systems, Budapest, Hungary, pp. 2009.
- [25] SANS, The critical security controls for effective cyber defense (version 5.0), accessed: 2015-12-19, <http://www.counciloncybersecurity.org/attachments/article/12/CSC-MASTER-VER50-2-27-2014.pdf>.
- [26] CWE, Cwe Top 25 Most Dangerous Software Errors, 2011, accessed: 2015-12-19, <http://cwe.mitre.org/top25/>.
- [27] Decision Support Approaches for Cyber Security Investment:Data for Cyber Essentials Case Study, 2015, <http://www.panaousis.com/papers/casestudy.pdf>.
- [28] Common Vulnerabilities and Exposures:The Standard for Information Security Vulnerability Names, accessed: 2015-12-19, <https://cve.mitre.org/>.

Andrew Fielder received a B.Sc. degree in Computer Science and an M.Sc.in Intelligent Systems Engineering from the University of Birmingham in 2007 and 2008 respectively. In 2013 he achieved a Ph.D. degree in Computer Science from the University of Birmingham in the field of evolutionary computation and market modelling. He is currently a Postdoctoral Research Associate with the Institute for Security Science and Technology at Imperial.

Emmanouil Panaousis is a Senior Lecturer at the School of Computing, Engineering and Mathematics at University of Brighton, UK. He received the B.Sc. degree in Informatics and Telecommunications from the University of Athens, Athens, Greece, in 2006 and the M.S. degree in Computer Science from the Department of Informatics of the Athens University of Economics and Business, Athens, Greece in 2008 and Ph.D. degree in Mobile Communications Security from Kingston University, London, UK in 2012. He is also affiliated with the Institute for Security Science and Technology, Imperial College London, London, UK. His research interests are in the fields of cyber security, privacy, and algorithmic decision making. He is a member of the IEEE and a professional member of the ACM.

Pasquale Malacaria is a Professor at the School of Electronic Engineering and Computer Science at Queen Mary University of London. His research focuses on quantitative analysis of security with recent applications to quantifying confidentiality leaks in Linux Kernel functions and, sponsored by Google, developing a plug-in for NASA model checker Java PathFinder to quantify confidentiality leaks in Java bytecode.

In 1996 he was awarded an EPSRC Advanced Fellowship to further his work on game semantics. He worked with Chris Hankin at Imperial College on applications of game semantics to program analysis and security, collaboration which brought the publication of several algorithms for program analysis and security based on abstractions of game semantics.

He is an associate editor for the ACM Computing Surveys and is currently serving as program committee member for several international conferences and workshops. He is an academic visitor with the Institute for Security Science and Technology at Imperial.

Malacaria has held several EPSRC grants, and currently is principal investigator for the Queen Mary parts of Games and Abstraction and Compositional Security Analysis for Binaries. The former is part of the GCHQ/EPSC Research Institute in Science of Cyber Security and the latter is part of the GCHQ/EPSC Research Institute in Automated program Analysis and Verification.

Chris Hankin is a Professor of Computing Science at Imperial College London and Director of the Institute for Security Science and Technology (ISST) at Imperial. He has worked on program analysis, particularly abstract interpretation, for about thirty years. Notable achievements include the first fully higher-order analysis for functional programs, the development of a number of efficient algorithms for computing fixed points, a fast pointer analysis (implemented in the GNU C compiler) and the first

flow-sensitive analysis for information flow. He led a European Future and Emerging Technologies (FET) project which investigated security analysis of Javacard bytecode programs. More recently, he has been applying the same techniques to the analysis of large datasets that arise in criminal investigations; a recent contribution has been a new multi-scale algorithm for detecting sub-communities in networks.

Whilst his early work concentrated on the use of discrete mathematics, his recent work has focussed on the analysis of probabilistic systems. He worked with Malacaria on the first uses of game semantics in program analysis. He is currently working with Malacaria on the use of game theory to provide better decision support in protecting systems against cyber attacks.

He has held several European and EPSRC grants. He is currently Director of the CPNI/EPSC Research Institute on Trustworthy Industrial Control Systems and principal investigator on the GCHQ/EPSC grant Games and Abstraction which is part of the Research Institute in Science of Cyber Security. He also receives support from the MoD, US Department of Homeland Security and Department of Defense. He was principal investigator of the CPNI/EPSC project Making Sense.

He is president of the Scientific Council of INRIA, Vice Chair of the European Commission's DG Connect Advisory Forum (CAF) and, until recently, Editor-in-Chief of ACM Computing Surveys.

Fabrizio Smeraldi received an M.Sc. in Physics from the University of Genoa, Italy, in 1996 and a Ph.D. in Science from EPFL, Lausanne, in 2000. He is currently a Lecturer in the Risk and Information Management group of the School of Electronic Engineering and Computer Science at Queen Mary, University of London. His research interests are in applied statistical learning, with particular reference to pattern recognition, cybersecurity and bioinformatics.