

Cybersecurity Games and Investments: A Decision Support Approach

Emmanouil Panaousis¹, Andrew Fielder², Pasquale Malacaria¹,
Chris Hankin², and Fabrizio Smeraldi^{*1}

¹ Queen Mary University of London

{e.panaousis,p.malacaria,f.smeraldi}@qmul.ac.uk

² Imperial College London

{andrew.fielder,c.hankin}@imperial.ac.uk

Abstract. In this paper we investigate how to optimally invest in cybersecurity controls. We are particularly interested in examining cases where the organization suffers from an underinvestment problem or inefficient spending on cybersecurity. To this end, we first model the cybersecurity environment of an organization. We then model non-cooperative *cybersecurity control-games* between the *defender* which abstracts all defense mechanisms of the organization and the *attacker* which can exploit different vulnerabilities at different network locations. To implement our methodology we use the SANS Top 20 Critical Security Controls and the 2011 CWE/SANS top 25 most dangerous software errors. Based on the profile of an organization, which forms its preferences in terms of *indirect costs*, its concerns about different kinds of *threats* and the importance of the assets given their associated risks we derive the Nash Equilibria of a series of control-games. These game solutions are then handled by optimization techniques, in particular multi-objective, multiple choice Knapsack to determine the optimal cybersecurity investment. Our methodology provides security effective and cost efficient solutions especially against *commodity attacks*. We believe our work can be used to advise security managers on how they should spend an available cybersecurity budget given their *organization profile*.

Keywords: cybersecurity, game theory, optimization.

1 Introduction

One of the single largest concerns facing organizations today is how to protect themselves from cyber attacks whose prominence impose the need for organizations to prioritize their cybersecurity concerns with respect to their perceived threats. Organizations are then required to act in such a way so as to minimize their vulnerability to these possible threats. The report [4] published by Deloitte and NASCIO, points out that only 24% of Chief Information Security Officers (CISOs) are very confident in protecting their organization's assets against external threats. Another important finding in this report is that the biggest concern

* The authors are supported by the Project "Games and Abstraction: The Science of Cyber Security" funded by EPSRC, Grants: EP/K005820/1, EP/K005790/1.

CISOs face in addressing cybersecurity is a “Lack of sufficient funding” where 86% of respondents were concerned.

Most organizations will have a fixed budget for the protection of their systems. Generally this budget would not allow them to fully cover all of the vulnerabilities that their data assets are at risk from. As such an organization is interested in *how to use the limited financial budget available to best protect them* from various vulnerabilities given that the implementation of a cybersecurity control is associated with a *direct cost*.

Apart from the direct costs of controls, there are also *indirect costs* incurred by the implementation of these controls. From this point of view investing more in cybersecurity might not always be the most efficient approach that CISOs can follow. Therefore another dimension of the cybersecurity investment problem is “what is the optimal cybersecurity budget allocation given the importance that the organization places into its different assets, the system performance requirements, and the profile of employees and clients?”

1.1 Our Contributions

In this work we provide a methodology and a tool that can support security managers with decisions regarding the optimal allocation of their cybersecurity budgets.

We first motivate a method for the creation of an organization’s cybersecurity strategy (Section 3). This is achieved by performing a risk analysis of the data assets that an organization has, and analyzing the effectiveness of different security controls against different vulnerabilities. We then formulate control-games (Section 4) based on these risk assessments, in order to calculate the most effective way for an organization to implement each control. In a control-game the defender aims at reducing cybersecurity risks by implementing a control in a certain way dictated by the Nash Equilibrium (NE). In this way, the defender minimizes the maximum potential damage inflicted by the attacker. The solutions of the different control-games are handled by the optimization techniques of multi-objective, multiple choice Knapsack (Section 5) to decide upon an optimal allocation of a cybersecurity budget. We also present a case study (Section 6) which includes vulnerabilities (i.e. CWE) and cybersecurity controls published by the Council on CyberSecurity. We have implemented our methodology (Section 7) for this case study by computing games solutions and investments and measure its performance in terms of cybersecurity defense for different organization profiles.

To demonstrate the effectiveness of our methodology we have implemented part of the SANS Top 20 Critical Security Controls and the 2011 CWE/SANS top 25 most dangerous software errors. We present examples of investment strategies that our tool recommends and test their optimality by looking at alternatives to show that they are the best. In this way, our work is a step towards implementing a theoretical cybersecurity investment decision-making methodology into a realistic scenario.

2 Related Work

Anderson [1] first proposed the study of security from an economics perspective putting forward the idea that cybersecurity is bounded by other non-technical incentives. Anderson highlighted with an example that although some organizations spend less money on security they spend it more effectively therefore having put in place better cyber defenses. In our work we share Anderson's view. However our approach is quite different as we focus on developing cybersecurity decision support tools to assist security managers on how to spend a cybersecurity 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 [17], where the authors address the problem of finding an optimal defensive coverage. The latter is defined as the one maximizing 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 [5] (p.134) discusses the importance of studying the quantitative aspects of risk assessment with regards to cybersecurity 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 associated from different vulnerabilities and the impact that different security controls have in mitigating these risks. By studying the incentives for risk management Alpcan [6] developed a game theoretic approach that optimizes the investment in security across different autonomous divisions of an organization, where each of the divisions is seen as a greedy entity. Furthermore Alpcan et al. examine in [14] security risk dependencies in organizations and they propose a framework which ranks the risks by considering the different complex interactions. This rank is dictated by an equilibrium derived by a Risk-Rank algorithm. Saad et al. [12] model cooperation among autonomous parts of an organization that have dependent security assets and vulnerabilities for reducing overall security risks, as a cooperative game. In [13] 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 are using Q-learning to cope with scenarios when players are not aware of the different Markov model parameters. Previous work carried out by Fielder et. al. [9] considers *how to optimally allocate the time for security tasks for system administrators*. This work identifies how to allocate the limited amount of time that a system administrator has to work on the different security related tasks for an organization's data assets.

One of the initial works studying the way to model investment in cybersecurity was conducted by Gordon and Loeb [7]. The authors identify a method for determining the level of investment for the protection of individual targets, showing that the optimal level of investment should be related to the probability of a vulnerability occurring. The main message of this work is that to

maximize the expected benefit from information security investment, an organization should spend only a small fraction of the expected loss due to a security breach. The work published in [8] 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 defense and not to cause any damage to the rest of the targets. The game theoretic analysis 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 minimization of cost for security investments, the latter imply that higher security levels are preferable. Cavusoglu et. al. [11] compare a decision theory based approach to game theoretic approaches for investment in cybersecurity. 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 cybersecurity spending has been published by Smeraldi and Malacaria [10]. The authors identified the optimum manner in which investments can be made in a cybersecurity 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 behavior by using cost-benefit analysis therefore considering rewards and costs of achieving different actions. Lastly, Demetz and Bachlechner [16] have identified, analyzed and presented a set of approaches for supporting information security investment decisions. A limitation of this paper, as highlighted by the authors, is that they assume that sufficient money is available to make an investment although in reality cybersecurity budgets are limited.

3 Cybersecurity Model

In this section we describe our cybersecurity model to illustrate an organization’s network topology, systems and security components. The network architecture will determine how the different assets of an organization are interconnected. In this paper we follow the network architecture as proposed in the *SANS Critical Security Control 19-1* entitled “Secure Network Engineering” and published in [3]. This consists of three depths namely the demilitarized zone (*DMZ*), the *Middleware*, and the *Private Network*. An organization’s assets that can be accessed from the Internet are placed in the *DMZ*, and they should not contain any highly sensitive data. Any asset with highly sensitive data must be located at the *Private Network*, and communicate with the outside world only through a proxy which resides on the *Middleware*.

We define the *depth* of an asset, denoted by d , as the location of this asset within an organization’s network architecture. Depths are separated from each other by a set of network security software, e.g., firewalls, IDS. A depth determines (i) the level of security that needs to be breached or bypassed in order for an attack to successfully exploit a vulnerability at this depth, and (ii) the importance of the data asset compromised if an attack is successful.

We denote the set of all cybersecurity targets within an organization by T and the set of all vulnerabilities threatened by commodity attacks by \mathcal{V} .

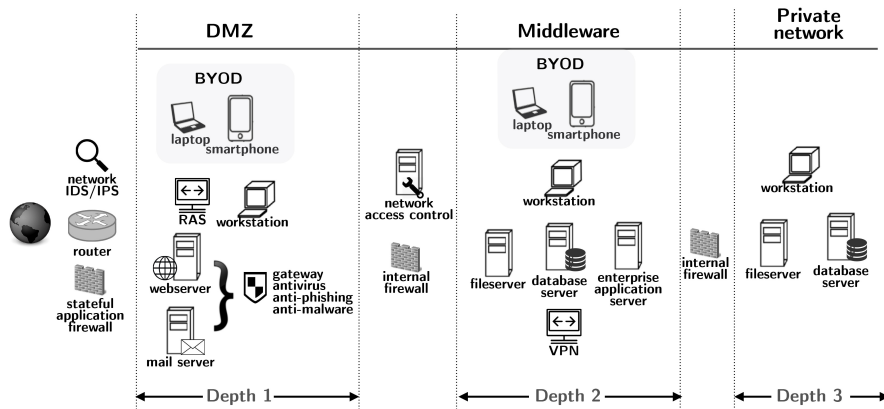


Fig. 1. Sample network architecture.

Definition 1. *Commodity attacks are attack methods where the attack tools can be purchased by a user, where the adversaries do not develop the attacks themselves, and only configure the tools for their own use.*

A *cybersecurity target* is defined as a $(vulnerability, depth)$ pair, i.e. $t_i = (v_z, d)$. A target abstracts any *data asset*, located at depth d , that an attack threatens to compromise by exploiting the vulnerability v_z .

We define the set of all targets as $T = \{(v_z, d) | v_z \in \mathcal{V}, d \in \{1, \dots, n\}\}$. We assume that each network architecture has its own set of targets however throughout this paper we consider the network architecture depicted in Fig. 1. In this paper, we specify that *data assets* located at the *same depth* and having the same *vulnerabilities* are abstracted by the *same target*, and they are worth the *same value* to the organization.

A *cybersecurity control* is the defensive mechanism that can be put in place to alleviate the risk from one or more attacks by reducing the probability of these attacks successfully exploiting vulnerabilities. The defender can choose to implement a control c_j at a certain level $l \in \{0, \dots, \mathcal{L}\}$. The higher the level the greater the degree to which the control is implemented.

Definition 2. *We define a cybersecurity process as implementation of a control at a certain level, and we denote by p_{jl} the cybersecurity process that implements the control c_j at level l .*

We define as $C = \{c_j\}$ the set of all cybersecurity controls the defender is able to implement to defend the system, and $P_j = \{p_{jl}\}$ the set of all cybersecurity processes associated with control c_j . A cybersecurity process p_{jl} has a *degree of mitigation* for each target t_i which equals the effectiveness of the cybersecurity process on this target, denoted by $e(t_i, p_{jl}) \in (0, 1]$. We also define $MITIGATION = e(t_i, p_{jl})$. In this paper we are interested in how cybersecurity processes are combined in a proportional manner to give an implementation plan for this control. We call this a *cybersecurity plan* which allows us to examine advanced ways of mitigating vulnerabilities.

Definition 3. *A cybersecurity plan is a probability distribution over different cybersecurity processes.*

In the following we describe the notions of *Risks*, *Indirect* and *Direct Costs* resulting from the implementation and purchase of a control and the *Organization Profile* which determines the preferences of an organization in terms of risks, indirect costs and how concerned the organization is about the different threats.

Risks. The *target risks* express the damage incurred to the defender when the attacker succeeds in compromising one or more targets. The different risks we consider are *Data Loss* (DL), *Business Disruption* (BD), and *Reputation* (RE). Each risk factor depends on the depth d that the attack targets; therefore we denote by DL_d , RE_d , and BD_d the risk values associated with a depth d .

Indirect Costs. For each cybersecurity process we consider three different types of *indirect costs*. The *System Performance Cost* (SPC) is associated with anything related to system performance being affected by a cybersecurity process (e.g., processing speed affected by anti-malware scanning). The *Morale Cost* (MOC) accounts for morale issues that higher levels of security can cause to users' happiness and job satisfaction. One negative implication of high MOC is that the stricter the security measures that an organization implements, the more likely an individual will want to circumvent them if possible. In these cases the attacker is able to take advantage of the reduced security from user actions. For example, having a control about different passwords for everything, might annoy users therefore increasing MOC. This might lead to circumvention of security by the user picking weak, memorable passwords which can often be cracked by dictionary or brute force attacks. Lastly, *Re-Training Cost* (RTC) refers to the cost for re-training users, including system administration, so they can either perform the cybersecurity process in the right way or be able to continue using all systems after a security update. We express the different indirect costs of a cybersecurity process p_{jl} by SPC_{jl} , RTC_{jl} , and MOC_{jl} .

Direct Costs. Each cybersecurity process has a direct cost which refers to the budget the organization must spend to implement the control c_j at a level l . The direct cost of a cybersecurity process is split into two categories, the *Capital Cost* (CAC) and the *Labour Cost* (LAC). CAC is related to hardware or software that must be purchased for the implementation of a control at some level. LAC is the direct cost for having system administrators implementing the control such as (hours spent) \times (cost/hour). When investing in cybersecurity we will be looking into the direct cost of each cybersecurity plan which is derived as a combination of the different costs of the cybersecurity processes that comprise this plan.

Vulnerability Factors. The Council on CyberSecurity has published in [2] software weaknesses (in this paper weakness and vulnerability are used interchangeably) and their factors. These factors are *Prevalence* (PR), *Attack Frequency* (AF), *Ease of Detection* (ED), and *Attacker Awareness* (AA). For a vulnerability v_z we denote the vulnerabilities factors by PR_z , AF_z , ED_z , AA_z . The level of a factor determines its contribution towards an overall vulnerability assessment score. For a commodity attack, one can argue that AA measures whether the average adversary would know that a malicious script is for sale, and ED is a measure of the computational cost of the attack discovery process. PR indicates

the number of times the weakness is found in the system (e.g., only 30% of windows systems ever downloaded a given patch), and AF dictates the number of times someone actually tries to exploit it (e.g., how many random SQL injection probes a day). We see PR and AF accounting for threats that are currently widespread (*current threats*) and ED and AA for threats that have the most potential for future attack vectors (*future potential threats*).

Organization Profile. To represent an *organization profile* we define a set $\{\mathcal{R}, \mathcal{K}, \mathcal{T}\}$ which dictates the preferences that an organization has with regards to risks, indirect costs and how concerned the organization is about well-known threats, respectively. These are given by the probability distributions $\mathcal{R} = [r_1, r_2, r_3]$, $\mathcal{K} = [k_1, k_2, k_3]$, and $\mathcal{T} = [\tau_1, \tau_2]$. The idea behind defining an organization profile is that a security manager can reason about the organization at a high level. This means whenever managers use our model they do not have to undertake some detailed security assessment, but only considers the high level needs of the organization.

The *Risk Profile*, denoted by \mathcal{R} , represents the importance that each of the potential areas of loss (DL,RE,BD) has to the organization. This is designed to prioritize the risk factors, such that each organization is able to identify the balance of the damage that they can expect from a successful attack. While the expectation is that data loss will be the predominant concern for most organizations, there are some that may consider that their reputation or the disruption to the operation of the business have a more significant impact. The most noticeable case for this would be organizations that predominantly deal with third party payment systems (e.g., Paypal), where the organization will hold relatively little data of value for their customers. For the *Risk Profile* weights we create the relation such that $r_1 \mapsto \text{DL}$, $r_2 \mapsto \text{RE}$, and $r_3 \mapsto \text{BD}$. We implicitly assume here that the organization’s risk profile remains the same at all depths. We then define $\text{RISKS} = r_1\text{DL}_d + r_2\text{RE}_d + r_3\text{BD}_d$.

The *Indirect Costs Profile* \mathcal{K} defines an importance for each of three different indirect cost factors SPC, RTC, and MOC. This is so that an organization can reason about the relative importance of indirect costs that it may incur when implementing a defense. The mapping of the different weights to costs are $k_1 \mapsto \text{SPC}$, $k_2 \mapsto \text{RTC}$, and $k_3 \mapsto \text{MOC}$. Therefore, $\text{IND_COSTS} = k_1\text{SPC}_{jl} + k_2\text{RTC}_{jl} + k_3\text{MOC}_{jl}$.

Lastly, *Threat Concern* \mathcal{T} is the level of importance that the business places on each of the threat factors. The main priority here is identifying whether the organization is concerned more about *current threats* or *future potential threats*. Therefore $\tau_1 \mapsto \text{current threats}$ and $\tau_2 \mapsto \text{future potential threats}$. We define $\text{THREAT} = \tau_1[(\text{PR}_z + \text{AF}_z)/2] + \tau_2[(\text{ED}_z + \text{AA}_z)/2]$.

4 Cybersecurity Control-Games

In this section we use game theory to model the interactions between two players; the *defender* and the *attacker*. The defender \mathcal{D} abstracts any cybersecurity decision-maker (e.g., security manager) which defends an organization’s data assets by minimizing cybersecurity risks with respect to the indirect costs of the cybersecurity processes while the attacker \mathcal{A} abstracts all adversaries that aim to

benefit from compromising the defender's data assets. The game we model here is a two-player game where there is a negative functional correlation between the attacker and the defender payoffs; the idea is that the more an attacker gains the more the defender loses. This means that equilibria in these games are minimax in an associated zero sum game. For any control c_j we define a control-subgame as follows.

Definition 4 (Control-subgame $\mathcal{G}_{j\lambda}$). 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 c_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.

\mathcal{D} 's finite strategy space is given by the set $\mathcal{A}^D = \{p_{jl}\}$. This means that \mathcal{D} 's actions are the different cybersecurity processes akin to implementations of a control c_j at different levels. The attacker can choose among different targets to attack therefore $\mathcal{A}^A = \{(v_z, d)\}$. We define \mathcal{D} 's mixed strategy 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 the control-subgame $\mathcal{G}_{j\lambda}$.

A mixed strategy of \mathcal{A} is defined as a probability distribution over the different targets and it is denoted by $H_{j\lambda} = [h_{j1}, \dots, h_{jn}]$, where h_{ji} is the probability of the adversary attacking target t_i when \mathcal{D} has only the control 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 Nash Equilibrium (NE) of $\mathcal{G}_{j\lambda}$ and it minimizes cybersecurity risks and indirect costs.

Example 1. In this example we consider 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 $[0, 0, \frac{7}{10}, 0, \frac{3}{10}]$ determines a cybersecurity plan that dictates the following:

$$\begin{aligned} \frac{7}{10} &\mapsto \text{regular scanning for the rest 70\% of devices} \\ \frac{3}{10} &\mapsto \text{real-time scanning for the 30\% of the most important devices} \end{aligned}$$

This mixed strategy can be realized more as 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.

We denote by $U_{\mathcal{D}}(p_{jl}, t_i)$ the utility of \mathcal{D} when target $t_i = \langle v_z, d \rangle$ is attacked, and the cybersecurity process p_{jl} has been selected to mitigate v_z at depth d , in general:

$$U_{\mathcal{D}}(p_{jl}, \langle v_z, d \rangle) := \text{RISKS} \times \text{THREAT} \times (1 - \text{MITIGATION}) + \text{IND_COSTS} \quad (1)$$

Theorem 1. The zero-sum cybersecurity control-subgame $\mathcal{G}_{j\lambda}$ admits an NE in mixed strategies, $(Q_{j\lambda}^*, H_{j\lambda}^*)$, with the property that

$$Q_{j\lambda}^* = \arg \max_{Q_{j\lambda}} \min_{H_{j\lambda}} U_{\mathcal{D}}(Q_{j\lambda}, H_{j\lambda}), \text{ and } H_{j\lambda}^* = \arg \max_{H_{j\lambda}} \min_{Q_{j\lambda}} U_{\mathcal{A}}(Q_{j\lambda}, H_{j\lambda})$$

The minimax theorem states that for zero sum games NE and minimax solution coincide. Therefore in $\mathcal{G}_{j\lambda}$ any Nash cybersecurity plan mini-maximizes the attacker's payoff. If any $\mathcal{G}_{j\lambda}$ admits multiple Nash cybersecurity plans they have the ordered interchangeability property which means that \mathcal{D} reaches the same level of defense independent from \mathcal{A} 's strategy, i.e.

$$Q_{j\lambda}^* = \arg \min_{Q_{j\lambda}} \max_{H_{j\lambda}} U_{\mathcal{A}}(Q_{j\lambda}, H_{j\lambda})$$

Definition 5. The non-zero sum control-subgame $\mathcal{G}'_{j\lambda} = \langle U_{\mathcal{D}}, U'_{\mathcal{A}} \rangle$ where $U'_{\mathcal{A}} = \alpha U_{\mathcal{A}} + \beta$, and α, β constants and $\alpha > 0$ is called a positive affine transformation (PAT) of the zero sum control-subgame $\mathcal{G}_{j\lambda} = \langle U_{\mathcal{D}}, U_{\mathcal{A}} \rangle$.

Proposition 1. If one of the game matrices of a control-subgame \mathcal{G}_j is a positive affine transformations (PAT) of a zero sum control-subgame \mathcal{G}'_j (and the other matrix is the same for both games) then the Nash equilibria of \mathcal{G}_j are minimax strategies. These also correspond to saddle-points [5] (p. 42).

In the rest of the paper we will restrict ourselves to control-subgames which are positive affine transformation of a zero sum control-subgame.

Definition 6 (Control-game \mathcal{G}_j). For any control c_j , with \mathcal{L} possible implementation levels, we define a control-game \mathcal{G}_j which consists of \mathcal{L} control-subgames, each of them denoted by $\mathcal{G}_{j\lambda}$, $\lambda \in \{0, 1, \dots, \mathcal{L}\}$.

In other words, a *control-game* is the collection of \mathcal{L} control-subgames for a specific control. The solution \mathcal{C}_j of a control-game for the defender is a set of Nash cybersecurity plans $\{Q_{jl}^*\}$, $\forall l \in \{0, \lambda\}$ each of them determined by the NE of each control-subgame. The set $\{\mathcal{C}_j\}$ for all controls $c_j \in C$ contains all sets of Nash cybersecurity plans one per control.

5 Cybersecurity Investment Optimization

In the previous section we were concerned with the implementation of a cybersecurity control. Nevertheless, organizations will generally implement more than one control. In this section we identify a method for combining these controls given that an organization's budget is constrained. More specifically, we describe how the control-game solutions are handled by optimization techniques to provide investment strategies. Each cybersecurity plan imposes its own direct costs including both CAC and LAC. Given a set $\{c_j\}$ of N controls each of them being associated with a set $\{\mathcal{C}_j\}$ of \mathcal{L} Nash cybersecurity plans, and an available budget B , in this section we examine how to optimally invest in the different plans by choosing at most one plan per control.

In relation to the cybersecurity investment problem we consider a 0-1 Knapsack problem similar to the cybersecurity budget allocation problem studied by Smeraldi and Malacaria [10]. In fact, in this paper we model this cybersecurity investment optimization problem as a *0-1 Multiple-Choice Multi-Objective Knapsack Problem*.

We assume that a plan can be effective in protecting more than one target and its benefit on a target is determined by the *expected damage* caused to the target when only this plan is purchased. The benefit of an investment solution on a target is determined by the sum of the benefits of the different plans on that

target where this sum never exceeds 1. Furthermore, *each investment solution has a score* determined by the maximum expected damage across all targets. When there are investment solutions with the same score we consider a *tie-break*. The question then arises, which one solution should one use? We consider that in the event of a tie-break, the solver uses the solution with the lowest cost. This tie-break makes sense as no-one would normally pay more for a defense that does no better.

The optimization method creates one objective function per target, which constraint is constrained by a common total budget B . Our goal is to derive the set of Nash cybersecurity plans (one per control) which minimizes the investment solution score. To derive the optimal investment solution, we compute the expected damage of each target for each possible set of plans. The weakest target is defined as the target that suffers the highest damage. In this way our method of evaluating the security of a system is to consider that “the security of a system is only as strong as it’s weakest point”. We then choose the set of plans that provides the minimum final expected damage among all highest expected damages.

Definition 7. *Defining the value of any target t_i as $\gamma_i = -\text{RISKS} \times \text{THREAT}$, considering N controls and assuming that each Nash cybersecurity plan $Q_{j\lambda}^*$ is associated with some benefit $b_{j\lambda}(t_i)$ ³ upon target t_i , and it has cost $\omega_{j\lambda}$, the defender seeks a cybersecurity investment \mathcal{I} such that*

$$\max_{\mathcal{I}} \min_{t_i} \left\{ 1 - \sum_{j=1}^N \sum_{\lambda=0}^{\mathcal{L}} b_{j\lambda}(t_i) x_{j\lambda} \right\} \gamma_i \quad (2)$$

subject to $\sum_{j=1}^N \sum_{\lambda=0}^{\mathcal{L}} \omega_{j\lambda} x_{j\lambda} \leq B$ and $\sum_{\lambda=0}^{\mathcal{L}} x_{j\lambda} = 1, x_{j\lambda} \in \{0, 1\}, \forall j = 1, \dots, N$

We denote by \mathcal{I} , the vector of cybersecurity plans (i.e. investment solution) purchased by solving the cybersecurity investment optimization problem for a *constant* number of targets. This approach aims to maximize the minimum amount of defense implemented across each target. In this way, the maximum damage across each target is minimized. So that the best solution represents the investment which has the lowest expected damage for the weakest target, which is the target that has the lowest overall cumulative benefit. This is subject to the condition that the budget (B) is not exceeded, where we consider if the control is used given by $x_{j\lambda}$ (either 0 or 1) and the cost of implementing it in the $\omega_{j\lambda}$. Additionally, we must satisfy that for each of the N controls that only a single sub game solution can and must be selected. Such that although they can only take a value of 0 or 1, the sum must equal 1, ensuring that a valid solution must select one sub game solution at level 1 and all others at level 0. It is important to note that we buy a single plan for each control (given by a control sub game solution). For example in Table 5. we buy the solution (represented by a value of 1 in the knapsack) for sub-game 3 of control 3, which might correspond to

³ we assume that $\sum_{j=1}^N \sum_{\lambda=0}^{\mathcal{L}} b_{j\lambda}(t_i) \leq 1$ achieved by normalized benefit values.

$[0,0,0.3,0.7,0]$, which suggests that the control is implemented at level 2 30 % of the time and at level 3 70% of the time, with a cost of 8.2. Such that we then select 0 for all other sub-game solutions for that control. For each control c_j there is a cybersecurity plan, denoted by Q_j^* that represents the optimal choice for the defender to purchase given some budget. Therefore $\mathcal{I} = [Q_1^*, Q_2^*, \dots, Q_N^*]$.

6 Case Study

In this section we describe the case study we use to implement our methodology. With regards to the organization size, we consider an SME with approximately 30 employees and we are interested in mitigating *commodity attacks*. This assumption allows us to have complete information in all control-games because the defender can be aware of the attacker’s payoff when it has been disclosed online. From the 2011 CWE/SANS top 25 most dangerous software errors aka vulnerabilities published in [2], we have considered 12 of those for the purposes of this case study as described in Table 1 along with their factors, and associated levels. For each vulnerability factor different levels are defined as in [2], and summarized in Table 2. Moreover, we have chosen 6 controls out of the The SANS 20 Critical Security Controls published by the Council on Cybersecurity in [3]. These are shown in Table 4 along with the different vulnerabilities that each control mitigates. As the same vulnerability can appear at different data assets at the same depth, we assume that the implementation of a control mitigates all occurrences of this vulnerability. Otherwise, the security of the system won’t increase because it is as strong as the weakest point. For a control, we assume five possible levels (i.e. 0-4) that the control can be implemented, where level 0 corresponds to no defense against the vulnerabilities and level 4 presents the highest possible level of control implementation with no regard for system operation. In Table 3, we highlight the indirect costs for all 6 controls considered in this case study. In the following we classify the controls into different implementation methods.

Depth Based Mitigation. This refers to controls that when applied at higher levels are used to cover additional depths within a system. This form of mitigation applies a system-wide control at level 1 and then applies more advanced countermeasures at additional depths, based on the level of implementation. The different levels are $\langle c, 0 \rangle$: no implementation, $\langle c, 1 \rangle$: all depths – basic, $\langle c, 2 \rangle$: depths 1,2 – basic & depth 3 – advanced, $\langle c, 3 \rangle$: depth 1 – basic & depths 2,3 – advanced, and $\langle c, 4 \rangle$: all depths – advanced. *Associated controls:* c_1, c_3, c_6 .

Table 1. Notation of 12 examined 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

level	PR	AF	ED	AA
3	Widespread	Often	Easy	High
2	High	Sometimes	Moderate	Medium
1	Common	Rarely	Difficult	Low

Table 2. Values of vulnerabilities factors published by CWE.

cybersecurity processes	SPC	MOC	RTC
p_{00}, \dots, p_{60}	0,0,0,0,0,0	0,0,0,0,0,0	0,0,0,0,0,0
p_{01}, \dots, p_{61}	1,1,1,1,2,2	1,1,1,0,1,1	0,0,0,2,1,0
p_{02}, \dots, p_{62}	2,2,1,2,2,2	2,1,1,0,2,1	0,0,0,2,1,0
p_{03}, \dots, p_{63}	2,3,2,3,2,2	4,1,1,0,3,3	0,0,0,2,1,1
p_{04}, \dots, p_{64}	3,3,2,4,2,2	5,2,2,0,4,3	0,0,0,2,2,2

Table 3. Indirect costs of the different cybersecurity processes.

Table 4. Vulnerabilities that each control mitigates.

	v_1	v_2	v_3	v_4	v_5	v_6	v_7	v_8	v_9	v_{10}	v_{11}	v_{12}
c_1 : Account Monitoring and Control	-	✓	-	-	-	✓	-	✓	✓	✓	-	-
c_2 : Continuous Vulnerability Assessment and Remediation	✓	✓	✓	-	✓	-	✓	-	-	-	-	-
c_3 : Malware Defenses	-	-	-	✓	-	-	-	✓	-	-	-	✓
c_4 : Penetration Tests and Red Team Exercises	✓	✓	✓	-	✓	✓	✓	-	✓	✓	✓	✓
c_5 : Controlled Use of Administrative Privileges	-	-	-	✓	-	-	-	-	✓	-	✓	-
c_6 : Data Loss Prevention	✓	-	-	✓	-	-	✓	✓	-	-	✓	-

Frequency Based Mitigation. This type of mitigation applies a control in a system-wide manner and higher levels of implementation reduce the time between scheduled performance of the mitigation. Low levels of a frequency based control may be performed as a one-off event or very infrequently, but this is then made more frequent at higher levels, where at the highest level these actions can be performed on demand. The different levels are $\langle c, 0 \rangle$: no implementation, $\langle c, 1 \rangle$: all depths – infrequent, $\langle c, 2 \rangle$: all depths – regular, $\langle c, 3 \rangle$: all depths – frequent, $\langle c, 4 \rangle$: all depths – real-time. *Associated controls:* c_2, c_4 .

Hybrid Mitigation. A *hybrid mitigation* control implements an approach to reducing the vulnerability of a system that acts with aspects of both depth based and frequency based controls. As such, these controls increase defense at lower depths as the control level increases, but additionally the frequency with which the schedule of the control at the other depths is also increased. The different levels are $\langle c, 0 \rangle$: no implementation, $\langle c, 1 \rangle$: all depths – basic & infrequent, $\langle c, 2 \rangle$: all depths – basic & regular, $\langle c, 3 \rangle$: all depths – basic & frequent, $\langle c, 4 \rangle$: all depths – advanced & real-time. *Associated control:* c_5 .

Each cybersecurity plan $Q_{j\lambda} = [q_{j0}, \dots, q_{j\lambda}]$ has a benefit, denoted by $b_{j\lambda}(t_i)$, on a target t_i . This is derived by the sum of the effectiveness values of the cybersecurity processes on t_i multiplied by the corresponding probability hence $b_{j\lambda}(t_i) = \sum_{l=0}^{\lambda} e(t_i, p_{jl})q_{j\lambda}$. A cybersecurity process p_{jl} has its own direct costs denoted by y_{jl} . Therefore the direct cost of a cybersecurity plan $Q_{j\lambda}$ is given by $\omega_{j\lambda} = \sum_{l=0}^{\lambda} y_{jl}q_{jl}$. In this work here we have defined the direct costs CAC and LAC per annum. Some of the controls have a one-off cost therefore any purchase can benefit the organization’s cybersecurity for the next years also. However, we examine the challenge of spending a cybersecurity budget annually assuming that in the worst case controls might need to be replaced or updated by spending an amount of money similar or even higher to the amount spent in the last year. In this paper we neither present the cybersecurity products we have chosen to implement the various controls nor their direct costs.

Lastly, the different risks values $\langle DL_d, RE_d, BD_d \rangle$ are defined as depth 1 $\mapsto \langle 2, 4, 3 \rangle$, depth 2 $\mapsto \langle 3, 2, 4 \rangle$, and depth 3 $\mapsto \langle 4, 3, 2 \rangle$. We have chosen the value

Table 5. Nash cybersecurity plans for the Case 1 with their associated direct costs.

c_j	Q_{j0}^*	Q_{j1}^*	Q_{j2}^*	Q_{j3}^*	Q_{j4}^*
c_1	[1,0,0,0,0] 0	[0,1,0,0,0] 9.7	[0,0.7,0.3,0,0] 9.8	[0,0.4,0.23,0.37,0] 10.7	[0,0,0.14,0.22,0.64] 12.4
c_2	[1,0,0,0,0] 0	[0,1,0,0,0] 1.7	[0,0.4,0.6,0,0] 2	[0,0,0.5,0.5,0] 5.1	[0,0,0.5,0.5,0] 5.1
c_3	[1,0,0,0,0] 0	[0,1,0,0,0] 7.1	[0,0,1,0,0] 7.3	[0,0,0.3,0.7,0] 8.2	[0,0,0.3,0.7,0] 8.2
c_4	[1,0,0,0,0] 0	[0,1,0,0,0] 4.2	[0,0,1,0,0] 8.3	[0,0,0,1,0] 16.7	[0,0,0,0,1] 33.4
c_5	[1,0,0,0,0] 0	[0,1,0,0,0] 4.1	[0,0.47,0.53,0,0] 4.1	[0,0,0.41,0.59,0] 4.1	[0,0,0,0.33,0.67] 5.4
c_6	[1,0,0,0,0] 0	[0,1,0,0,0] 6	[0,0,1,0,0] 7.4	[0,0,0.44,0.56,0] 12	[0,0,0.32,0.52,0.16] 13.6

of data loss to be the highest within the *Private Network*, because this depth will generally contain the most sensitive data. We have assessed the value of reputation loss RE independently of the value of data loss DL to show the impact that only RE has to the organization. We have set the highest value of RE to the DMZ because it contains the forward facing assets of the organization. For example when the organization’s website is defaced this can be seen by any potential user who visits the organization’s website and harm its reputation. As most of the organization’s workload is likely to be handled by devices in the *Middleware* we have assigned the highest BD value to this depth.

7 Games Solutions and Investments

This section explains the set of results we have retrieved for 3 different organization profiles (3 Cases). For each profile: (i) we solve a series of control-games therefore a set of control-subgames for each control to derive the Nash cybersecurity plans (in this section we use the terms Nash cybersecurity plans and plans interchangeably) and, (ii) we determine the optimal cybersecurity investment given a budget by using optimization techniques and the control-game solutions. In Cases 1, and 3 we consider an organization which places a high importance on its data assigning a value 0.8 to DL. RE, and BD are equally important taking the same value 0.1. We also consider here that the organization is equally interested in *current* and *potential future threats* in all cases. In Cases 1 and 2 the organization prioritizes system performance costs higher than re-training and morale costs by giving an SPC value twice that of RTC and MOC values. We have increased MOC in Case 3 making it twice as large as SPC to assess the impact of morale in cybersecurity strategies.

To derive the different Nash cybersecurity plans we have solved 24 different control-subgames (i.e. 6 control-games) for each organization profile. The game solutions were computed by using a minimax solver, implemented in Python. For simplicity reasons we have chosen the first equilibrium computed by our solver noting that all equilibria offer the same level of defense as we state in Theorem 1. In Table 5 we present the results of the control-subgames in Case 1, where the solution calculated for each control sub-game taken is the strategy with the smallest support.

The graphs presented in Figs. 2 and 3 are designed to show which plans should be chosen for each possible budget level. In other words, each graph shows the optimal investment \mathcal{I} which is the set of plans chosen for a certain budget. The graphs should be used to identify which are the most important plans for




Fig. 2. Case 1: $\mathcal{R} = [0.8, 0.1, 0.1]$, $\mathcal{T} = [0.5, 0.25, 0.25]$, $\mathcal{K} = [0.5, 0.5]$.

a given organization at an available budget. It is worth noting here that we have normalized the cost values such that the sum of the direct costs of all the controls implemented at the highest possible level (i.e. the most expensive possible cybersecurity plans) equals 100. Our methodology uses the optimization technique presented in Section 5 to compute the investment solution \mathcal{I} that has the highest score for a given budget. As we have discussed in Section 5, *each investment solution has a score* determined by the maximum expected damage across all targets. A question a security manager may ask is, how is the investment solution \mathcal{I} translated in terms of controls acquisition and how can someone describe that it is better, in terms of cyber defense, than alternative solutions where $\mathcal{I}' \neq \mathcal{I}$?

Example 2. From Fig. 2 we consider an available budget of 17. In this example our decision support methodology advises the security manager to implement $\mathcal{I} = [Q_{10}^*, Q_{21}^*, Q_{30}^*, Q_{42}^*, Q_{50}^*, Q_{61}^*]$ with a cost of 16.102. The above solution determines a set of plans to be selected for the implementation of the 6 SANS controls as defined in our case study. To be able to translate the solution into the implementation of the different available controls in Table 5 we present the Nash cybersecurity plans of Case 1. According to \mathcal{I} the controls that should be implemented and the manner in which they are implemented is listed as follows

- Q_{10}^* : With the given budget, *Account Monitoring and Control* (c_1) software should not be purchased, nor should system administrators spend time on activities to this control.
- Q_{21}^* : The organization must implement the *Continuous Vulnerability Assessment and Remediation* (c_2) control by purchasing a *vulnerability scanner and patch management* software. Additionally system administrators measuring the delay in patching new vulnerabilities and audit the results of vulnerability scans at all network depths *infrequently* (e.g, once per month).
- Q_{30}^* : The decision tool does not recommend the implementation of specific *Malware Defenses* (c_3) given the available budget.
- Q_{42}^* : The security manager is advised to schedule *regular* (e.g., twice a year) system-wide *Penetration Tests and Red Team Exercises* (c_4), with system updates being performed based on the results of the exercise.

- Q_{50}^* : The tool does not recommend the implementation of the *Controlled Use of Administrative Privileges* (c_5) control which means that neither enterprise password manager software must be purchased nor any password renewal policy must be enforced.
- Q_{61}^* : The tool recommends the implementation of the *Data Loss Prevention* (c_6) control system-wide and at a basic level (e.g., integrated services router with security, VPN).

By using Table 4 we see that with these controls all targets are covered to some degree. In the following we consider alternative cases to highlight the optimality of the solution. If we implement system-wide *Penetration Tests and Red Team Exercises* infrequently (e.g., once per year) (Q_{41}^*) instead of regularly (Q_{42}^*) then we release a budget of 4.174 therefore we can implement *Controlled Use of Administrative Privileges* by using an *enterprise password manager* software and *renew passwords* of all systems infrequently (e.g., annually) (Q_{51}^* with cost 4.153). This gives another investment $\mathcal{I}' = [Q_{10}^*, Q_{21}^*, Q_{30}^*, Q_{41}^*, Q_{51}^*, Q_{61}^*]$ with cost 16.081. Under \mathcal{I}' the *Controlled Use of Administrative Privileges* control improves the defense on targets associated with the following vulnerabilities; **XSS** (v_4), **Unnecessary privileges** (v_9), and **Path traversal** (v_{11}). But it then leaves worse off, due to the less frequent *Penetration Tests and Red Team Exercises*, 8 vulnerabilities namely; **SQLi** (v_1), **OS command injection** (v_2), **Buffer overflow** (v_3), **Missing authentication** (v_5), **Missing authorization** (v_6), **Missing encryption** (v_7), **CSRF** (v_{10}), and **Unchecked code** (v_{12}). Due to \mathcal{I} being the choice of the optimization the score achieved by \mathcal{I} is higher than \mathcal{I}' therefore the weakest target in \mathcal{I}' must appear in these 8 vulnerabilities and it must be weaker than the weakest target under \mathcal{I} . By saying weakest target we refer to the target with the maximum expected damage. Therefore our methodology advises the security manager to undertake *Penetration Tests and Red Team Exercises* regularly (e.g., twice a year, Q_{42}^*) without implementing *Controlled Use of Administrative Privileges* at all.

If we do not spend any money on *Penetration Tests and Red Team Exercises* we then have an available budget of 8.347 which can be spent in implementing *Malware Defenses* by installing a free anti-malware software with manual scheduled scans and database updates in all devices of the organization (Q_{31}^* , 7.095). Therefore another investment is $\mathcal{I}' = [Q_{10}^*, Q_{21}^*, Q_{31}^*, Q_{40}^*, Q_{50}^*, Q_{61}^*]$. Under \mathcal{I}' targets associated with **Missing authorization** (v_6) and **Unnecessary privileges** (v_9) are not covered by any control thus one of these becomes the weakest target under \mathcal{I}' . Due to \mathcal{I} being the optimal investment solution provided by our tool, the weakest target (not covered at all) under \mathcal{I}' must be weaker than the weakest (partially covered) target under \mathcal{I} . Therefore our solution recommends not to ignore (even at some basic level) the implementation of *Penetration Tests and Red Team Exercises* which can actually identifies if a user can access a given resource, despite not being authorized for that (v_6) and it can also mitigate v_9 by identifying cybersecurity processes that run with extra privileges, such as root or Administrator, and they can disable the normal security checks.

Finally, if we assume a slightly higher budget of 17.145 we can choose the investment strategy $\mathcal{I}' = [Q_{10}^*, Q_{21}^*, Q_{31}^*, Q_{42}^*, Q_{50}^*, Q_{60}^*]$ which does not implement the *Data Loss Prevention* control but it installs free anti-malware with manual scheduled scans and database updates system-wide (Q_{31}^*). This is not a better investment than \mathcal{I} despite being more expensive. Both *Data Loss Prevention* (in \mathcal{I}) and *Malware Defenses* (in \mathcal{I}') mitigate XSS (v_4) and **Unrestricted upload** (v_8) which *Penetration Tests and Red Team Exercises* does not. Thus from this point of view the replacement of *Data Loss Prevention* by *Malware Defenses* does not affect the effectiveness of the targets associated with XSS and **Unrestricted upload**. However due to \mathcal{I} being the choice of the optimization the target associated with **Path traversal** (v_{11}), which is mitigated only by *Penetration Tests and Red Team Exercises* under \mathcal{I}' , is weaker than a target associated with **CSRF** (v_{10}) or **Unchecked code** (v_{12}) mitigated only by *Penetration Tests and Red Team Exercises* under \mathcal{I} . In other words, according to the effectiveness values we have provided in our case study, **Path traversal** is not mitigated as much as **CSRF** and **Unchecked code**, which are both mitigated by *Penetration Tests and Red Team Exercises*, therefore \mathcal{I} is better than \mathcal{I}' .

Example 3. According to Fig. 2 for a budget of 28 our methodology gives the investment solution $\mathcal{I} = [Q_{13}^*, Q_{21}^*, Q_{31}^*, Q_{41}^*, Q_{52}^*, Q_{60}^*]$ with a total direct cost 27.80. This solution provides the following list of recommendations.

- Q_{13}^* : Implementation of *Account Monitoring and Control* (c_1) at a basic level (e.g., control built into OS and manually review all accounts or set files/folders auditing properties) in all devices in *DMZ*; in 63% of the devices in *Middleware*; and in 40% of the devices in *Private Network*. The control must be also implemented at an advanced level (e.g., vulnerability scanner and patch management software) in 37% of the devices in *Middleware* and 60% of the devices in *Private Network*.
- Q_{21}^* : System-wide *Continuous Vulnerability Assessment and Remediation* (c_2) must be implemented infrequently (e.g., once per month).
- Q_{31}^* : System-wide *Malware Defenses* (c_3) must be implemented at a basic level (e.g., free anti-malware with manual scheduled scans and database updates).
- Q_{41}^* : *Penetration Tests and Red Team Exercises* (c_4) to be undertaken infrequently (e.g., once per year).
- Q_{52}^* : *Controlled Use of Administrative Privileges* (c_5) to be implemented at a basic level (e.g., using an enterprise password manager software) with 47% of the devices to change passwords infrequently (e.g., once per year) and 53% regularly (e.g., every 4 months).
- Q_{60}^* : The purchase of a *Data Loss Prevention* control is not recommended.

To see how \mathcal{I} outperforms other investments we have considered some alternative investments for a budget of 28. As the first alternative investment \mathcal{I}' , we decide not to follow Q_{13}^* therefore saving 10.68. By doing that the targets associated with $v_2, v_6, v_8, v_9, v_{10}$ are now defended in a lower degree than in \mathcal{I} as the effectiveness of Q_{13}^* does not count in the sum of the benefits for these targets. Also under \mathcal{I}' the targets associated with v_2, v_9 , and v_{10} are defended

by two controls while the targets with v_6 , and v_8 are only defended by one control. With a budget of 10.68 available we can purchase c_6 and implement it according to plan Q_{62}^* with cost 7.408 or according to Q_{61}^* with cost 6.052. If we choose the former the control is implemented at an advanced level (e.g., drive encryption, system recovery) in the *Private network*, and at a basic level (e.g., integrated services router with security, VPN) in *DMZ* and *Middleware*. *Data Loss Prevention* improves the defense of targets with v_1, v_4, v_7, v_8 , and v_{11} . When implementing *Data Loss Prevention*, v_6 (**Missing authorization**) is mitigated only by one control therefore making any target with this vulnerability likely to be the weakest among all in \mathcal{I}' and weaker than the weakest target under \mathcal{I} . Our solution \mathcal{I} dictates that it is preferable for the security manager to purchase the *Account and Monitoring Control* as opposed to *Data Loss Prevention* to prevent unauthorized users accessing resources or data of the organization in the first case rather than allowing such access and hoping that data encryption and system recovery capabilities can discourage an adversary from attacking.

Next, we assume another variation \mathcal{I}' of our investment where *Malware Defenses* (c_3) is removed and 7.095 budget is available for spending in other controls. By not purchasing c_3 vulnerabilities v_4, v_8 and v_{10} are mitigated by one less control. With an available budget of 7.095 we can purchase c_6 and implementing it according to Q_{61}^* . The difference between \mathcal{I} and \mathcal{I}' is that v_{10} is mitigated by one more control in \mathcal{I} . Since the latter has been the result of optimization, any target with v_{10} (**CSRF**) is the weakest target and weaker than the weakest target under \mathcal{I} . In other words \mathcal{I} advises the security manager to purchase *Malware Defenses* rather than *Data Loss Prevention* to detect malware that can be installed when a **CSRF** attack is launched. Again here \mathcal{I} dictates that stopping the attack at a first infection stage is more important than guaranteeing that stolen data are encrypted thus unreadable. Besides the attacker's motivation might be to just corrupt or delete data which is something *Data Loss Prevention* can address only at high levels of implementation which require a higher budget.

Graph Trends. From the graphs in Fig. 2 we see the results level off (at around a budget of 45), when the perceived benefit from a combination of plans brings the expected damage down to a minimum, this is such that adding a new plan or a plan at a higher level won't improve the defense of the system. This is as a result of us capping the sum of improvements to 1, but would exist in any form of interdependent control methodology and only the point at which it levels off would change. Furthermore, this observation dictates that *cybersecurity does not get improved by investing more in cybersecurity plans*. With higher budgets it is much more feasible to reduce the damage of not just the weakest target, but other targets as well. A spike exists when there is a small budget range that opens up a number of new cybersecurity plans. In reality a number of the solutions in that range will have similar expected damage values, but we only see the best solution for that particular budget. For budgets 1-19 the progression is the same regardless of the plans. The reason for this is that at these levels only certain plan combinations are available and we want to ensure that as low an expected damage as possible is achieved. At these levels it is seen as most important to




Fig. 3. (i) Case 2: $\mathcal{R} = [0.6, 0.4, 0]$, $\mathcal{T} = [0.5, 0.25, 0.25]$, $\mathcal{K} = [0.5, 0.5]$, (ii) Case 3: $\mathcal{R} = [0.8, 0.1, 0.1]$, $\mathcal{T} = [0.3, 0.1, 0.6]$, $\mathcal{K} = [0.5, 0.5]$.

cover all of the targets with some form of plan, bringing down the system-wide expected damage.

While there is a consistent strategy for investment with budget levels for solutions up to a budget of 20, after this budget we see that different investment profiles are suggested by our methodology across the different organization profiles. From budgets 22 to 26 and from 36 to 38, there is no change in the solution. While alternative solutions may become available in these ranges, none of these solutions will improve on the security of the weakest target, which means that as with very low budgets there is no incentive to implement a more expensive plan combination that does not improve the effectiveness of defense on the weakest target. Between budgets 30 and 36 as the budget increases more, there are new combinations of plans that become available at each of these levels that will improve the overall defense of the system. However it can also be seen that in order to implement a different solution some components of the previous solution need to be removed in order to reduce the cost to fit within the budgetary constraint.

Sensitivity to Organizations' Profile Perturbations. One question that arises is how robust is the proposed approach to informing the way an organization should invest in cyber security? We have focused on the importance of the decisions made by the organization with regards to their profile. In this way we have looked at how small perturbations in a single case affect the allocation of investment. We consider the two cases of $[0.75, 0.125, 0.125]$ and $[0.85, 0.075, 0.075]$ for \mathcal{R} . Both alternative profiles have minor deviations from the original solution. Each of the deviations found would cause the solution to differ for up to a maximum of 3 consecutive budget levels, before the proposed controls would realign. Using the values of $[0.55, 0.45]$ and $[0.45, 0.55]$ for \mathcal{T} we find that, in the case of $[0.55, 0.45]$ there is a different investment strategy that is proposed between controls 5 and 6 for budgets between 13 and 17. This is the only case we have seen where there is a difference in the low budget strategies across all the cases tested for this work. With a value for \mathcal{T} of $[0.45, 0.55]$, we find that there is no change to the proposed investment plans. In the case of \mathcal{K} we consider $[0.45, 0.275, 0.275]$ and $[0.55, 0.225, 0.225]$, which for both values give us no change in the proposed

investment from the original case. Importantly in all of the cases tested we have seen that the stable investment solution for all of the results is the same as the case presented in Fig. 2.

8 Conclusions

This paper presents a cybersecurity decision support methodology for calculating the optimal security investment for an organization. This is formulated as a multiple choice and multi-objective Knapsack problem which handles the solutions of cybersecurity control-games. Our methodology creates strategies for each control at different levels of implementation and enforcement, where the combination of the most effective controls within a budget are suggested for implementation. The model supports the movement of human decision making from trying to analyze the explicit security requirements of the system to deciding upon an organization's priorities. The feature of the model that helps to create this movement is the *organization profile*, where a profile allows the model to reflect the individual nature of different organizations in the proposed investment. One of the most important factors that this work highlights is that it is important for an organization to know how to appropriately generate their profile. This is crucial because it influences the way an organization should invest in their cybersecurity defenses. From the results we have noticed that for similar organizations the best protection will be similar if not the same, because the results of the control-games will favor certain targets or controls.

In this paper we have assumed additive benefits for the different plans and the same target. One important future aim is to better understand the steps of the attacks, as such identifying steps in the chain will better inform the way in which different security controls interact in order to better cover different targets. This will better inform the way in which the subgames results are combined in the investment problem and reflect in a more realistic way how cyber defenses work.

At the moment our data is generated with the advice of a limited set of experts. In the future we aim to increase the number of experts involved to better understand their cyber environment needs. This will allow us to implement our methodology in a realistic environment. Additional limitations of our work that we wish to address in future work is to consider a higher number of available controls and continuous values for the levels of controls implementation. Moreover, at the moment our control-subgames are games of complete information. In the future we will examine incomplete information games where the defender is not aware of the attacker's payoff therefore any investment solution has to respect this uncertainty which highlights a situation very close to realistic environments that are prone to 0-days attacks and Advanced Persistent Threats (APTs). Finally, we do not see a strong case for using Stackelberg games in the case of commodity attacks where both players have publicly available information about attack types. The case would have been stronger if we were considering sophisticated cyber criminals or nation states where surveillance of the defender's actions prior to the attack would be important for the recognition of the defending mechanisms and the exploitation of one or more weak targets.

References

1. Anderson, R: Why Information Security is Hard. In Proc. of the 17th Annual Computer Security Applications Conference (2001)
2. CWE: 2011 CWE/SANS Top 25 Most Dangerous Software Errors. <http://cwe.mitre.org/top25/> (accessed May 2014)
3. Council on Cybersecurity: 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> (accessed May 2014)
4. 2012 Deloitte-NASCIO Cybersecurity Study State governments at risk: a call for collaboration and compliance. https://www.deloitte.com/assets/Dcom-UnitedStates/Local/%20Assets/Documents/AERS/us_aers_nascio/%20Cybersecurity/%20Study_10192012.pdf (accessed May 2014)
5. Alpcan, T., Basar, T.: Network Security: A Decision and Game-Theoretic Approach. Cambridge University Press (2010)
6. Alpcan, T.: Dynamic incentives for risk management. In: Proc. of the 5th IEEE International Conference on New Technologies, Mobility and Security (NTMS) (2012)
7. Gordon, L.A., Loeb, M.P.: The economics of information security investment. In: ACM Transactions on Information and System Security (TISSEC) (2002)
8. Johnson, B., Grossklags, J., Christin, N., Chuang, J.: Nash equilibria for weakest target security games with heterogeneous agents. In: Jain, R., Kannan, R. (eds.) *Gamenets 2011*. LNCS, vol. 75, pp. 444-458. Springer, Heidelberg (2012)
9. Fielder, A., Panaousis, E., Malacaria, P., Hankin, C., Smeraldi, F.: Game theory meets information security management. In: Proc. of the 29th IFIP International Information Security and Privacy Conference (2014)
10. Smeraldi, F., Malacaria, P.: How to Spend it: Optimal Investment for Cyber Security. In: Proc. of the 1st International Workshop on Agents and CyberSecurity (ACySe) (2014)
11. Cavusoglu, H., Srinivasan, R., Wei, T.Y.: Decision-theoretic and game-theoretic approaches to IT security investment. In: *Journal of Management Information Systems (ACySe)*, 25(2), pp.281-304 (2008)
12. Saad, W., Alpcan, T., Basar, T., Hjorungnes, A.: Coalitional game theory for security risk management. In: Proc. of the 5th International Conference on Internet Monitoring and Protection (ICIMP), pp. 35-40 (2010)
13. Bommannavar, P., Alpcan, T., Bambos, N.: Security risk management via dynamic games with learning. In: Proc. of the 2011 IEEE International Conference on Communications (ICC), pp. 1-6 (2011)
14. Alpcan, T., Bambos, N.: Modeling dependencies in security risk management. In: Proc. of the Fourth International Conference on Risks and Security of Internet and Systems (CRiSIS), pp. 113-116 (2009)
15. Cremonini, M., Nizovtsev, D.: Understanding and influencing attackers' decisions: Implications for security investment strategies
16. Demetz, L., Bachlechner, D.: 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*, Springer Berlin Heidelberg, pp. 25-47 (2013)
17. Kiekintveld, C., Islam, T., Kreinovich, V.: Security games with interval uncertainty. In: Proc. of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), pp. 231-238. International Foundation for Autonomous Agents and Multiagent Systems, Richland (2013)