Mitigating Covert Compromises
A Game-Theoretic Model of Targeted and Non-Targeted Covert Attacks

Aron Laszka[1,2], Benjamin Johnson[3], and Jens Grossklags[1]

[1] College of Information Sciences and Technology,
Pennsylvania State University, USA
[2] Department of Networked Systems and Services,
Budapest University of Technology and Economics, Hungary
[3] School of Information, University of California, Berkeley, USA

Abstract. Attackers of computing resources increasingly aim to keep security compromises hidden from defenders in order to extract more value over a longer period of time. These covert attacks come in multiple varieties, which can be categorized into two main types: targeted and non-targeted attacks. Targeted attacks include, for example, cyber-espionage, while non-targeted attacks include botnet recruitment.

We are concerned with the subclass of these attacks for which detection is too costly or technically infeasible given the capabilities of a typical organization. As a result, defenders have to mitigate potential damages under a regime of incomplete information. A primary mitigation strategy is to reset potentially compromised resources to a known safe state, for example, by reinstalling computer systems, and changing passwords or cryptographic private keys.

In a game-theoretic framework, we study the economically optimal mitigation strategies in the presence of targeted and non-targeted covert attacks. Our work has practical implications for the definition of security policies, in particular, for password and key renewal schedules.

Keywords: Game Theory, Computer Security, Covert Compromise, Targeted Attacks, Non-Targeted Attacks.

1 Introduction

Most organizations devote significant resources to prevent security compromises which may harm their financial bottomline or adversely affect their reputation. Security measures typically include technologies to detect known attack vectors. However, recent studies of anti-malware and anti-virus tools have demonstrated their ineffectiveness against novel attack approaches and even incrementally modified known malware.

At the same time, attackers prey upon opportunities to keep successful security compromises covert. The goal is to benefit from defenders' lack of awareness by exploiting resources, and extracting credentials and company secrets for as

Y. Chen and N. Immorlica (Eds.): WINE 2013, LNCS 8289, pp. 319–332, 2013.

long as possible. In contrast to non-covert attacks and compromises that focus on short-term benefits, these long-lasting and (for typical organizations) undetectable attacks pose specific challenges to system adminstrators and creators of security policies. Discoveries of such attacks by sophisticated security companies provide evidence for damage caused over many months or years.

CDorked, a highly advanced and stealthy backdoor, was discovered in April 2013 [5]. The malware uses compromised webservers to infect visitors with common system configurations. To stay covert, the malware uses a number of different techniques, for example, not delivering malicious content if the visitor's IP address is in a customized blacklist. The operation has been active since at least December 2012, and has infected more than 400 webservers, including 50 from Alexa's top 100,000 most popular websites.

Another example is Gauss, a complex, nation-state sponsored cyber-espionage toolkit, which is closely related to the notorious Stuxnet [1,10]. Gauss was designed to steal sensitive financial data from targets primarily located in the Middle East, and was active for at least 10 months before it was discovered.

Such recently-revealed attack vectors as well as the suspected number of unknown attacks highlight the importance of developing mitigation strategies to minimize the resulting expected losses. Potentially effective mitigation approaches include resetting of passwords, changing cryptographic private keys, reinstalling servers, or reinstantiating virtual servers. These approaches are often effective at resetting the resource to a known safe state, but they reveal little about past compromises. For example, if a server is reinstalled, knowledge of if and when a compromise occured may be lost. Likewise, resetting a password does not reveal any information about the confidentiality of previous passwords.

Covert (and non-covert) attacks can be distinguished in another dimension by the extent to which the attack is targeted (or customized) for a particular organization [4,7]. Approaches related to cyber-espionage are important examples of targeted attacks, and require a high effort level customized to a specific target [14]. A typical example of a non-targeted covert attack is the recruitment of a computer into a botnet via drive-by-download. Such attacks are relatively low effort, and do not require a specific target. Further, they can often be scaled to affect many users for marginal additional cost [7]. See Table 1 for a comparison between targeted and non-targeted attacks.

Table 1. Comparison of Targeted and Non-Targeted Attacks

	Targeted	Non-Targeted
Number of attackers	low	high
Number of targets	low	high
Effort required for each attack	high	low
Success probability of each attack	high	low

The targeted nature of an attack also matters to the defender, because targeted and non-targeted attacks do different types of damage. For example, a targeting attacker might use a compromised computer system to access an organization's secret e-mails, which may potentially cause enormous economic damage; while a non-targeting attacker might use the same compromised machine to send out spam, causing different types of damage.

The presence of both targeted and non-targeted covert attacks presents an interesting dilemma for a medium-profile target to choose a mitigation strategy against covert attacks. Strategies which are optimal against non-targeted attacks may not be the best choice against targeted attacks. At the same time, mitigation strategies against targeted attacks may not be economically cost-effective against only non-targeting attackers.

To address this dichotomy, we present a game in which a defender must vie for a contested resource that is subject to both targeted attacks from a strategic attacker, and non-targeted covert attacks from a large set of non-strategic attackers. We identify Nash equilibria in the simultaneous game, and subgame perfect equilibria in the sequential game with defender leading. The optimal mitigation strategies for the defender against these combined attacks lend insights to policy makers regarding renewal requirements for passwords and cryptographic keys.

The rest of the paper is organized as follows. In Section 2, we review related work. We define our game-theoretic model in Section 3, and we give analytical results for this model in Section 4. In Section 5, we present numerical and graphical observations; we conclude in Section 6.

2 Related Work

2.1 Security Economics and Games of Timing

Research studies on the economics of security decision-making primarily investigate the optimal or bounded rational choice between different canonical options to secure a resource (i.e., protection, mitigation, risk-transfer), or the determination of the optimal level of investment in one of these security dimensions. In our own work, we have frequently contributed to the exploration of these research objectives (see, for example, [6,9,8]). Further, these studies have been thoroughly summarized in a recent review effort [11].

Another critical decision dimension for successfully securing resources is the consideration of *when* to act to successfully thwart attacks. Scholars have studied such time-related aspects of tactical security choices since the cold war era by primarily focusing on zero-sum games called *games of timing* [2]. The theoretical contributions on some subclasses of these games have been surveyed by [17].

2.2 FlipIt: Modeling Targeted Attacks

Closely related to our study is the `FlipIt` model which identifies optimal timing-related security choices under targeted attacks [19]. In `FlipIt`, two players compete for a resource that generates a payoff to the current owner. Players can make costly moves (i.e., “flips”) to take ownership of the resource, however, they have to make moves under incomplete information about the current state of possession.

In the original FlipIt paper, equilibria and dominant strategies for simple cases of interaction are studied [19]. Other groups of researchers have worked on extensions [16,12]. For example, Laszka et al. extended the FlipIt game to the case with multiple resources. In addition, the usefulness of the FlipIt game has been investigated for various application scenarios [3,19]. We detail the difference of our model to FlipIt in Section 3.3. The current study generalizes our previous work which was restricted to exponential distributions for the attack time [13].

FlipIt has been studied in experiments in which human participants were matched with computerized opponents [15]. This work has also been extended to consider different interface feedback modalities [18]. The results complement the theoretical work by providing evidence for the difficulty to identify optimal choices when timing is the critical decision dimension.

3 Model Definition

We model the covert compromise scenario as a randomized, one-shot, non-zero-sum game. For a list of symbols used in our model, see Table 2. The player who is the rightful owner of the resource is called the defender, while the other players are called attackers. The game starts at time $t = 0$ with the resource being uncompromised, and it is played indefinitely as $t \to \infty$. We assume that time is continuous.

Table 2. List of Symbols

C_D	move cost for the defender
C_A	move cost for the targeting attacker
B_A	benefit received per unit of time for the targeting attacker
B_N	benefit received per unit of time for the non-targeting attackers
F_A	cumulative distribution function of the attack time for the targeting attacker
λ_N	rate of the non-targeted attacks' arrival

We let D, A, and N denote the defender, the targeting attacker, and the non-targeting attackers, respectively. At any time instance, player i may make a move, which costs her C_i. (Note that, for attackers, we will use the words attack and move synonymously). When the defender makes a move, the resource becomes uncompromised immediately for every attacker. When the targeting attacker makes a move, she starts her attack, which takes some random amount of time. If the defender makes a move while an attack is in progress, the attack fails. We assume that the time required by a targeted attack follows the same distribution every time. Its cumulative distribution function is denoted by F_A, and is subject to $F_{A_i}(0) = 0$. In practice, this distribution can be based on industry-wide beliefs, statistics of previous attacks, etc.

The attackers' moves are stealthy; i.e., the defender does not know when the resource became compromised or if it is compromised at all. On the other hand, the defender's moves are non-stealthy. In other words, the attackers learn immediately when the defender has made a move.

The cost rate for player i up to time t, denoted by $c_i(t)$, is the number of moves per unit of time made by player i up to time t, multiplied by the cost per move C_i.

For attacker $i \in \{A, N\}$, the benefit rate $b_i(t)$ up to time t is the fraction of time up to t that the resource has been compromised by i, multiplied by B_i. Note that if multiple attackers have compromised the resource, they all receive benefits until the defender's next move. For the defender D, the benefit rate $b_D(t)$ up to time t is $-\sum_{i\in\{A,N\}} b_i(t)$. The relation between defender and attacker benefits implies that the game would be zero-sum if we only considered the players' benefits. Because our players' payoffs also consider move costs, our game is *not* zero-sum. Player i's payoff is defined as

$$\liminf_{t\to\infty} \; b_i(t) - c_i(t) \; . \tag{1}$$

It is important to note that the asymptotic benefit rate $\liminf_{t\to\infty} \; b_i(t)$ of attacker i is equal to the probability that i has the resource compromised at a random time instance, multiplied by B_i. For a discussion on computing the payoffs for the key strategy profiles in this paper, see the extended version of this paper available on the authors' websites.

3.1 Types of Strategies for the Defender and the Targeting Attacker

Not Moving. A player can choose to *never move*. While this might seem counter-intuitive, it is actually a best-response if the expected benefit from making a move is always less than the cost of moving.

Adaptive Strategies for the Targeting Attacker. Let $\mathcal{T}(n) = \{T_0, T_1, \ldots, T_n\}$ denote the move times of the defender up to her nth move (or in the case of $T_0 = 0$, the start of the game). The attacker uses an *adaptive strategy* if she waits for $W(\mathcal{T}(n))$ time until making a move after the defender's nth move (or after the start of the game), where W is a non-deterministic function. If the defender makes her $n+1$st move before the chosen wait time is up, the attacker chooses a new wait time $W(\mathcal{T}(n+1))$, which also considers the new information that is the defender's $n+1$st move time. This class is a simple representation of all the rational strategies available to an attacker, since W can depend on all the information that the attacker has, and we do not have any constraints on W.

Renewal Strategies. Player i uses a *renewal strategy* if the time intervals between consecutive moves are identically distributed independent random variables, whose distribution is given by the cumulative function F_{R_i}. Renewal strategies are well-motivated by the fact that the defender is playing blindly; thus, she has the same information available after each move. So it makes sense to use a strategy which always chooses the time until her next flip according to the same distribution.

Periodic Strategies. Player i uses a *periodic strategy* if the time intervals between consecutive moves are identical. This period is denoted by δ_i. Periodic strategies are a special case of renewal strategies.

3.2 Non-Targeted Attacks

Suppose that there are N non-targeting attackers. In practice, N is very large, but the expected number of attacks in any time interval is finite. Hence, as N goes to infinity, the probability that a given non-targeting attacker targets the defender approaches zero. Since non-targeting attackers operate independently of each other, the number of successful attacks in any time interval depends solely on the length of the interval. Thus, the arrival of *successful non-targeted attacks* follows a *Poisson process.*

Furthermore, since the economic decisions of the non-targeting attackers depend on a very large pool of possible targets, the effect of the defender's behaviour on the non-targeting attackers' strategies is negligible. Thus, the non-targeting attackers' strategies (that is, the arrival rate of the Poisson process) can be considered exogenously given. We let λ_N denote the expected number of arrivals that occur per unit of time; and we model all the non-targeting attackers together as a single attacker whose benefit per unit of time is B_N.

3.3 Comparison to `FlipIt`

Even though our game-theoretic model resembles `FlipIt` in many ways, it differs in three key assumptions. First, we assume that the defender's moves are *non-stealthy.* The motivation for this is that, when the attacker receives benefits from continuously exploiting the compromised resource, she should know whether she has the resource compromised or not. For example, if the attacker uses the compromised password of an account to regularly spy on its e-mails, she will learn of a password reset immediately when she tries to access the account. Second, we assume that the targeting attacker's moves are *not instantaneous*, but take some time. The motivation for this is that an attack requires some effort to be carried out in practice. Furthermore, the time required for a successful attack may vary, which we model using a random variable for the attack time. Third, we assume that the defender faces *multiple attackers*, not only a single one.

Moreover, to the authors' best knowledge, papers published on `FlipIt` so far give analytical results only on a very restricted set of strategies. In contrast, we completely describe our game's equilibria and give optimal defender strategies based on very mild assumptions, which effectively do not limit the power of players (see the introduction of Section 4).

4 Analytical Results

In this section, we provide analytical results based on our model. We start with a discussion of the players' strategies.

Recall that the defender has to play blindly, which means that she has the same information available after each one of her moves. Consequently, it makes sense for her to choose the time until her next flip according to the same distribution each time. In other words, a rational defender can restrict herself to using only renewal strategies.

Now, if the defender uses a renewal strategy, the time of her next move depends only on the time elapsed since her last move T_n, and the times of her previous moves (including T_n) are irrelevant to the future of the game. Therefore, it is reasonable to assume that the attacker's response strategy to a renewal strategy also does not depend on $T_0, T_1, \ldots, T_n$. For the remainder of the paper, when the defender plays a renewal strategy, the attacker uses a fixed probability distribution – given by the density function f_W – over her wait times for when to begin her attack. Note that it is clear that there always exists a best-response strategy of this form for the attacker against a renewal strategy.

Since the attacker always waits an amount of time that is chosen according to a fixed probability distribution after the defender's each move, the amount of time until the resource would be successfully compromised after the defender's move also follows a fixed probability distribution. Let S be the random variable measuring the time after the defender has moved until the attacker's attack would finish. The probability density function f_S of S can be computed as

$$f_S(s) = \int_{w=0}^{s} f_W(w) \int_{a=0}^{(s-w)} f_A(a)\; da\; dw\ . \tag{2}$$

Finally, we let F_S denote the cumulative distribution function of S.

4.1 Best Responses

Defender's Best Response. We begin our analysis with finding the defender's best-response strategy.

Lemma 1. *Suppose that the non-targeted attacks arrive according to a Poisson process with rate λ_N, and the targeting attacker uses an adaptive strategy with a fixed wait time distribution F_W. Then,*

- *not moving is the only best response if $C_D = \mathcal{D}(l)$ has no solution for $l > 0$, where*

$$\mathcal{D}(l) = B_A \left(l F_S(l) - \int_{s=0}^{l} F_S(s)\; ds \right) + B_N \left(-l e^{-\lambda_N l} + \frac{1 - e^{-\lambda_N l}}{\lambda_N} \right)\ ; \tag{3}$$

- *the periodic strategy whose period is the unique solution to $C_D = \mathcal{D}(l)$ is the only best response otherwise.*

The proof is available in the paper's extended version on the authors' websites.

Even though we cannot express the solution of $C_D = \mathcal{D}(l)$ in closed form, it can be easily found using numerical methods, as the right hand side is continuous and increasing.[1] Note that all the equations presented in the subsequent lemmas and theorems of this paper can also be solved by applying numerical methods.

Lemma 2. *Suppose that the non-targeted attacks arrive according to a Poisson process with rate λ_N, and the targeting attacker never attacks. Then,*

[1] We show in the proof of the lemma that the right hand side is increasing in l.

- *not moving is the only best response if $C_D = \mathcal{D}^N(l)$ has no solution for $l > 0$, where*

$$\mathcal{D}^N(l) = B_N \left(-le^{-\lambda_N l} + \frac{1 - e^{-\lambda_N l}}{\lambda_N} \right) \; ; \tag{4}$$

- *the periodic strategy whose period is the unique solution to $C_D = \mathcal{D}^N(l)$ is the only best response otherwise.*

Proof. Follows readily from the proof of Lemma 1 with the terms belonging to the targeting attacker omitted everywhere. □

Observe that $\mathcal{D}(0) = \mathcal{D}^N(0) < 0$ and $\mathcal{D}(l) \geq \mathcal{D}^N(l)$. Consequently, $C_D = \mathcal{D}(l)$ has a solution whenever $C_D = \mathcal{D}^N(l)$ has one. Furthermore, if both have solutions, the solution of $C_D = \mathcal{D}(l)$ is less than or equal to the solution of $C_D = \mathcal{D}^N(l)$. In other words, the *defender is more likely to keep moving if there is a threat of targeted attacks*, and she *will move at least as frequently as she would if there was no targeting attacker.*

Attacker's Best Response. We continue our analysis with finding the attacker's best-response strategy.

Lemma 3. *Against a defender who uses a periodic strategy with period δ_D,*

- *never attacking is the only best response if $C_A > \mathcal{A}(\delta_D)$, where*

$$\mathcal{A}(\delta) = B_A \int_{a=0}^{\delta} F_A(a) da \; ; \tag{5}$$

- *attacking immediately after the defender has moved is the only best response if $C_A < \mathcal{A}(\delta_D)$;*
- *both not attacking and attacking immediately are best responses otherwise.*

The proof is available in the paper's extended version on the authors' websites.

The lemma shows that the targeting attacker should either attack immediately or not attack at all, but she should never wait to attack. For the never attack strategy, we already have the defender's best response from Lemma 2. For the attacking immediately strategy, the defender can determine the optimal period of her strategy *solely based on the distribution of A*, which is an exogenous parameter of the game. More formally, the defender's best response is not to move if $C_D = \mathcal{D}^A(l)$ has no solution, and it is a periodic strategy whose period is the unique solution to $C_D = \mathcal{D}^A(l)$ otherwise, where

$$\mathcal{D}^A(l) = B_A \left(lF_A(l) - \int_{a=0}^{l} F_A(a) \; da \right) + B_N \left(-le^{-\lambda_N l} + \frac{1 - e^{-\lambda_N l}}{\lambda_N} \right) \; . \tag{6}$$

This follows readily from Lemma 1 by substituting F_S for F_A.[2]

[2] Recall that S was defined as the sum of the waiting time W, which is always zero in this case, and the attack time A.

4.2 Nash Equilibria

Based on the previous lemmas, we can describe all the equilibria of the game (if there are any) as follows.

Theorem 1. *Suppose that the defender uses a renewal strategy, the targeting attacker uses an adaptive strategy, and the non-targeted attacks arrive according to a Poisson process. Then, the game's equilibria can be described as follows.*

1. *If $C_D = \mathcal{D}^A(l)$ does not have a solution for l, then there is a unique equilibrium in which the defender does not move and in which the targeting attacker moves once at the beginning of the game.*
2. *If $C_D = \mathcal{D}^A(l)$ does have a solution δ_D for l:*
   (a) *If $C_A \leq \mathcal{A}(\delta_D)$, then there is a unique equilibrium in which the defender plays a periodic strategy with period δ_D, and the targeting attacker moves immediately after each of the defender's moves.*
   (b) *If $C_A > \mathcal{A}(\delta_D)$,*
      i. *if $C_D = \mathcal{D}^N(l)$ has a solution δ'_D for l, and $C_A \geq \mathcal{A}(\delta'_D)$, then there is a unique equilibrium in which the defender plays a periodic strategy with period δ_D, and the targeting attacker never moves;*
      ii. *otherwise, there is no equilibrium.*

The proof is available in the paper's extended version on the authors' websites. For an illustration of the hierarchy of the theorem's criteria, see Figure 1. Finally, recall that a discussion on the payoffs can also be found in the extended version.

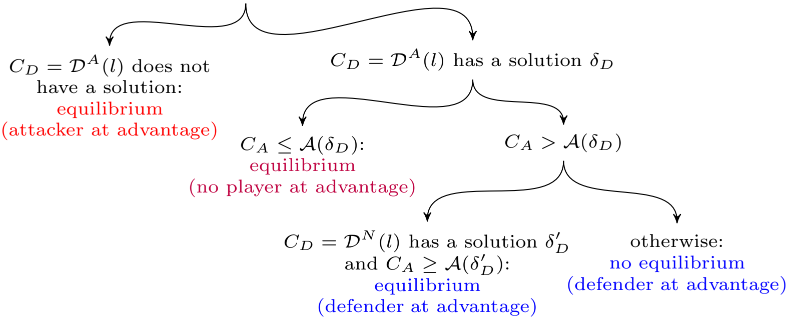

$C_D = \mathcal{D}^A(l)$ does not have a solution:
equilibrium
(attacker at advantage)
$C_D = \mathcal{D}^A(l)$ has a solution δ_D
$C_A \leq \mathcal{A}(\delta_D)$:
equilibrium
(no player at advantage)
$C_A > \mathcal{A}(\delta_D)$
$C_D = \mathcal{D}^N(l)$ has a solution δ'_D and $C_A \geq \mathcal{A}(\delta'_D)$:
equilibrium
(defender at advantage)
otherwise:
no equilibrium
(defender at advantage)

Fig. 1. Illustration for the hierarchy of criteria in Theorem 1

In the first case, the attacker is at an overwhelming advantage, as the relative cost of defending the resource is prohibitively high. Consequently, the defender simply "gives up," as any effort to protect the resource is not profitable, and the attacker will eventually have the resource compromised indefinitely (see Figure 2 for an illustration). In the second case, no player is at an overwhelming advantage. Both players are actively moving, and the resource gets compromised and uncompromised from time to time. In the third and fourth cases, the defender is at an overwhelming advantage. However, this does not necessarily lead to an equilibrium. If the defender moves with a sufficiently high rate, she makes moving unprofitable for the targeting attacker. But if the targeting attacker decides

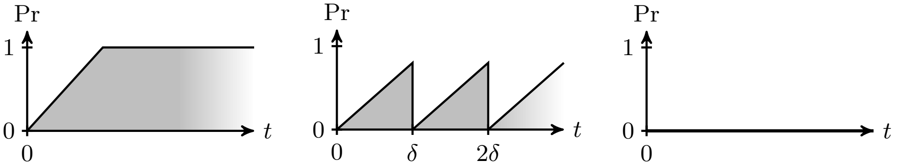

(a) Case 1. (b) Case 2. (a) (c) Case 2. (b) i.

Fig. 2. The probability that the targeting attacker has compromised the resource (vertical axis) as a function of time (horizontal axis) in various equilibria (see Theorem 1 for each case). Note that these are just examples, the actual shapes of the functions depend on F_A.

not to move, then the defender switches to a lower move rate, which is optimal against only non-targeted attacks. However, once the defender switches to the lower move rate, it might again be profitable for the targeting attacker to move, which would in turn trigger the defender to switch back to the higher move rate.

4.3 Sequential Game: Deterrence by Committing to a Strategy

So far, we have modeled the mitigation of covert compromises as a simultaneous game. This is realistic for scenarios where neither the defender nor the targeting attacker can learn the opponent's strategy choice in advance. However, in practice, the defender can easily let the targeting attacker know her strategy by publicly announcing it. Even though one of the key elements of security is confidentiality, the defender can actually gain from revealing her strategy – as we will show in Section 5 – since this allows her to deter the targeting attacker from moving.

In this section, we model the conflict as a sequential game, where the defender chooses her strategy before the targeting attacker does. We assume that the defender announces her strategy (e.g., publicly commits herself to a certain cryptographic-key update policy) and the targeting attacker chooses her best response based on this knowledge. Furthermore, in this section, we restrict the defender's strategy set to periodic strategies and not moving. The following theorem describes the defender's subgame perfect equilibrium strategies.

Theorem 2. *Let δ_1 be the solution of $C_D = \mathcal{D}^A(\delta)$ (if any), δ_2 be the maximal period δ for which $C_A = \mathcal{A}(\delta)$, and δ_3 be the solution of $C_D = \mathcal{D}^N(\delta)$ (if any). In a subgame perfect equilibrium, the defender's strategy is one of the following:*
- *not moving,*
- *periodic strategies with periods $\{\delta_1, \delta_2, \delta_3\}$.*

The proof is available in the paper's extended version on the authors' websites.

Based on the above theorem, one can easily find all subgame perfect equilibria by iterating over the above strategies and, for each strategy, computing the targeting attacker's best response using Lemma 3, and finally comparing the defender's payoffs to find her equilibrium strategy (or strategies). Note that, for each case of Theorem 1, the set of possible equilibrium strategies in Theorem 2 could be restricted further. For example, in Case 2. (b) i., the only subgame perfect equilibrium is the defender moving periodically with δ'_D and the targeting attacker never moving. We defer the remaining cases to future work.

5 Numerical Illustrations

In this section, we present numerical results on our game. For the illustrations, we instantiate our model with the *exponential distribution* as the distribution of the attack time. For rate parameter λ_A, the cumulative distribution function of the exponential distribution is $F_A(a) = 1 - e^{-\lambda_A a}$. For the remainder of this section, unless indicated otherwise, the parameters of the game are $C_D = C_A = B_A = \lambda_A = \lambda_N = 1$ and $B_N = 0.1$. Finally, we refer to the simultaneous-game Nash equilibria simply as equilibria, and we refer to the defender's subgame perfect equilibrium strategies as optimal strategies (because they maximize the defender's payoff given that the targeting attacker will play her best response).

First, in Figure 3, we study the effects of varying the value of the resource, that is, the unit benefit B_A received by the targeting attacker. Figure 3a shows the equilibrium payoffs as functions of B_A (the defender's period for the same

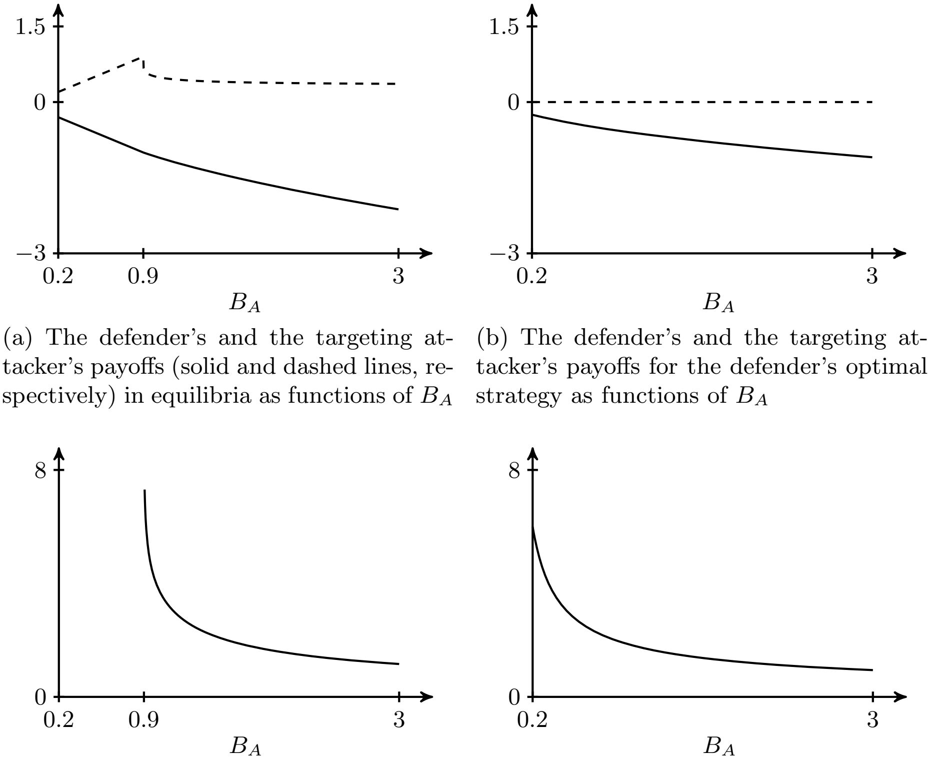

(a) The defender's and the targeting attacker's payoffs (solid and dashed lines, respectively) in equilibria as functions of B_A

(b) The defender's and the targeting attacker's payoffs for the defender's optimal strategy as functions of B_A

(c) The defender's equilibrium period as a function of B_A

(d) The defender's optimal period as a function of B_A

Fig. 3. The effects of varying the unit benefit B_A of the targeting attacker

setup is shown by Figure 3c). The defender's payoff is strictly decreasing, which is not surprising: the more valuable the resource is, the higher the cost of security is. The targeting attacker's payoff, on the other hand, starts growing linearly, but then suffers a sharp drop, and finally converges to a finite positive value.

For lower values ($B_A < 0.9$), the defender does not protect the resource, as it is not valuable enough. Accordingly, Figure 3c shows no period for this range, and the targeting attacker's payoff is the value of the resource B_A. However, once the value reaches about 0.9, the defender starts protecting the resource. Hence, the targeting attacker's payoff drops as she no longer always has the resource. For higher values, the defender balances between losses and costs, which means that the time the resource is compromised decreases as its value increases.

Figure 3b shows the payoffs for the defender's optimal strategy as functions of B_A (the optimal period is shown by Figure 3d). The figure shows that the defender's strategy for this range of B_A is always to deter the targeting attacker (hence, the targeting attacker's payoff is zero). To achieve this, the defender is using a strictly shorter period than her equilibrium period. Interestingly, the defender's payoff is much higher compared to her equilibrium payoff.

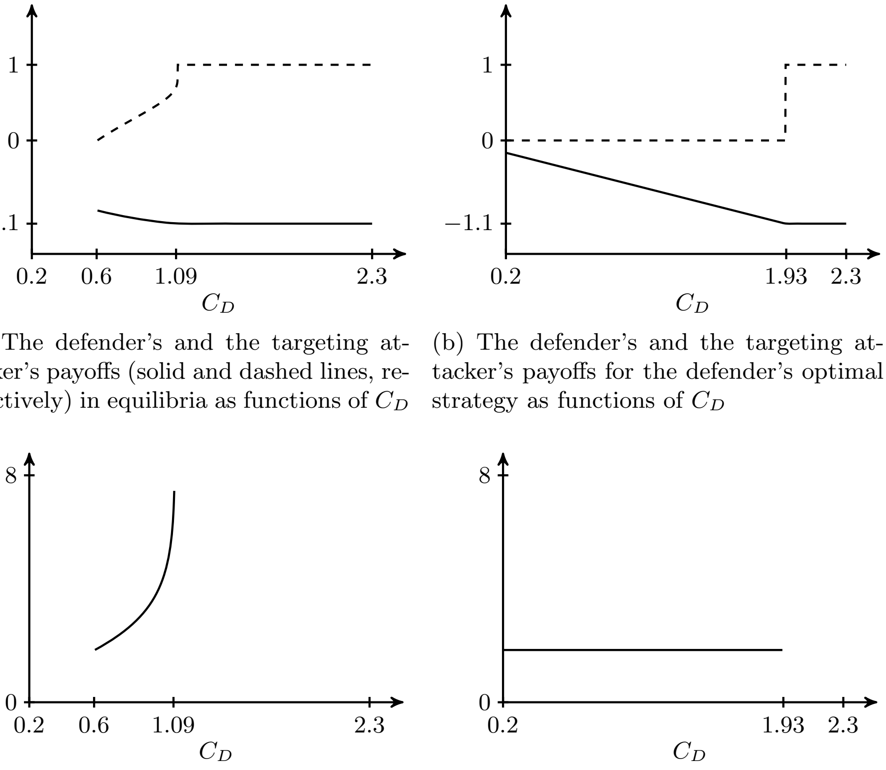

(a) The defender's and the targeting attacker's payoffs (solid and dashed lines, respectively) in equilibria as functions of C_D

(b) The defender's and the targeting attacker's payoffs for the defender's optimal strategy as functions of C_D

(c) The defender's equilibrium period as a function of C_D

(d) The defender's optimal period as a function of C_D

Fig. 4. The effects of varying the defender's move cost C_D

In Figure 4, we study the effects of varying the defender's move cost C_D. Figure 4a shows the equilibrium payoffs as functions of C_D (the defender's period for the same setup is shown by Figure 4c). The figure shows that the defender's payoff is decreasing, while the targeting attacker's payoff is increasing, which is unsurprising: the more costly it is to defend, the greater the attacker's advantage.

For lower costs ($C_D < 0.6$), the defender is at an overwhelming advantage, but there is no equilibrium (Case 2. (b) ii. of Theorem 1). For costs between 0.6 and 1.09, no player is at an overwhelming advantage; hence, both players move from time to time. For higher costs, the targeting attacker is at an overwhelming advantage. In this case, the defender never moves, while the attacker moves once. Hence, their payoffs are $B_A + B_N = -1.1$ and $B_A = 1$, respectively.

Figure 4b shows the payoffs for the defender's optimal strategy as a function of C_D (the optimal period is shown by Figure 4d). The defender's optimal strategy for move costs lower than 1.93 is to deter the targeting attacker. Hence, the targeting attacker's payoff is zero. The defender's payoff decreases linearly as the cost of deterrence increases. Again, we see that the defender's payoff is much higher than her equilibrium payoff. However, for higher move costs, she must give up defending the resource, as in her equilibrium strategy for this range.

6 Conclusions

In this paper, we studied the mitigation of both targeted and non-targeted covert attacks. As our main result, we found that periodic mitigation is the most effective strategy against both types of attacks and their combinations. Considering the simplicity of this strategy, our result can be surprising, but it also serves as a theoretical justification for the prevalent periodic password and cryptographic-key renewal practices. Moreover, this result contradicts the lesson learned from the `FlipIt` model [19], which suggests that a defender facing an adaptive attacker should use an unpredictable, randomized strategy.

Further, a defender is more willing to commit resources to defensive moves when being threatened by non-targeted and targeted attacks at the same time. This stands in contrast to the result that a high level of either threat type can force the defender to abandon defensive activities altogether.

Finally, we observed that there is an important difference between the defender's simultaneous and sequential (i.e., optimal) equilibrium strategies, both in the lengths of the periods and the resulting payoffs. Thus, a defender should not try to keep her strategy secret, but rather publicly commit to it.

Acknowledgements. We gratefully acknowledge the support of the Penn State Institute for CyberScience and the National Science Foundation under ITR award CCF-0424422 (TRUST). We also thank the reviewers for their comments on an earlier draft of the paper.

References

1. Bencsath, B., Pek, G., Buttyán, L., Felegyhazi, M.: The cousins of Stuxnet: Duqu, Flame, and Gauss. Future Internet 4(4), 971–1003 (2012)

2. Blackwell, D.: The noisy duel, one bullet each, arbitrary accuracy. Technical report, The RAND Corporation, D-442 (1949)
3. Bowers, K., van Dijk, M., Griffin, R., Juels, A., Oprea, A., Rivest, R., Triandopoulos, N.: Defending against the unknown enemy: Applying FLIPIT to system security. In: Grossklags, J., Walrand, J. (eds.) GameSec 2012. LNCS, vol. 7638, pp. 248–263. Springer, Heidelberg (2012)
4. Casey, E.: Determining intent - Opportunistic vs. targeted attacks. Computer Fraud & Security 2003(4), 8–11 (2003)
5. ESET Press Center. ESET and Sucuri uncover Linux/Cdorked.A: The most sophisticated Apache backdoor (2013), http://www.eset.com/int/about/press/articles/article/eset-and-sucuri-uncover-linuxcdorkeda-apache-webserver-backdoor-the-most-sophisticated-ever-affecting-thousands-of-web-sites/
6. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: Proc. of the 17th International World Wide Web Conference (WWW), pp. 209–218 (2008)
7. Herley, C.: The plight of the targeted attacker in a world of scale. In: 9th Workshop on the Economics of Information Security, WEIS (2010)
8. Johnson, B., Böhme, R., Grossklags, J.: Security games with market insurance. In: Baras, J.S., Katz, J., Altman, E. (eds.) GameSec 2011. LNCS, vol. 7037, pp. 117–130. Springer, Heidelberg (2011)
9. Johnson, B., Grossklags, J., Christin, N., Chuang, J.: Are security experts useful? Bayesian nash equilibria for network security games with limited information. In: Gritzalis, D., Preneel, B., Theoharidou, M. (eds.) ESORICS 2010. LNCS, vol. 6345, pp. 588–606. Springer, Heidelberg (2010)
10. Kaspersky Lab. Gauss (2012), http://www.kaspersky.com/gauss
11. Laszka, A., Felegyhazi, M., Buttyán, L.: A survey of interdependent security games. Technical Report CRYSYS-TR-2012-11-15, CrySyS Lab, Budapest University of Technology and Economics (November 2012)
12. Laszka, A., Horvath, G., Felegyhazi, M., Buttyan, L.: FlipThem: Modeling targeted attacks with FlipIt for multiple resources. Technical report, Budapest University of Technology and Economics (2013)
13. Laszka, A., Johnson, B., Grossklags, J.: Mitigation of targeted and non-targeted covert attacks as a timing game. In: Proc. of GameSec 2013 (2013)
14. Laszka, A., Johnson, B., Schöttle, P., Grossklags, J., Böhme, R.: Managing the weakest link: A game-theoretic approach for the mitigation of insider threats. In: Crampton, J., Jajodia, S., Mayes, K. (eds.) ESORICS 2013. LNCS, vol. 8134, pp. 273–290. Springer, Heidelberg (2013)
15. Nochenson, A., Grossklags, J.: A behavioral investigation of the FlipIt game. In: 12th Workshop on the Economics of Information Security, WEIS (2013)
16. Pham, V., Cid, C.: Are we compromised? Modelling security assessment games. In: Grossklags, J., Walrand, J. (eds.) GameSec 2012. LNCS, vol. 7638, pp. 234–247. Springer, Heidelberg (2012)
17. Radzik, T.: Results and problems in games of timing. In: Statistics, Probability and Game Theory: Papers in Honor of David Blackwell. Lecture Notes-Monograph Series, Statistics, vol. 30, pp. 269–292 (1996)
18. Reitter, D., Grossklags, J., Nochenson, A.: Risk-seeking in a continuous game of timing. In: Proc. of the 13th International Conference on Cognitive Modeling (ICCM), pp. 397–403 (2013)
19. van Dijk, M., Juels, A., Oprea, A., Rivest, R.: FlipIt: The game of "stealthy takeover". Journal of Cryptology 26, 655–713 (2013)