Mitigation of Targeted and Non-Targeted Covert Attacks as a Timing Game

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

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

Abstract. We consider a strategic game in which a defender wants to maintain control over a resource that is subject to both targeted and non-targeted covert attacks. Because the attacks are covert, the defender must choose to secure the resource in real time without knowing who controls it. Each move by the defender to secure the resource has a one-time cost and these defending moves are not covert, so that a targeted attacker may time her attacks based on the defender's moves. The time between when a targeted attack starts and when it succeeds is given by an exponentially distributed random variable with a known rate. Non-targeted attackers are modeled together as a single attacker whose attacks arrive following a Poisson process. We find that in this regime, the optimal moving strategy for the defender is a periodic strategy, so that the time intervals between consecutive moves are constant.

Keywords: Game Theory, Computer Security, Games of Timing, Covert Compromise, Targeted Attacks, Non-Targeted Attacks

1 Introduction

A growing trend in computer security is the prevalence of continuous covert attacks on networked resources. In contrast to one-time attacks with immediate benefit, such as initiating a wire transfer from a compromised bank account, a covert attack seeks to maintain control of a resource while keeping the compromise a secret. This type of attack is ubiquitous in the formation of botnets, as individual computer owners rarely know that their computer is a botnet member. Routers that are used to conduct man-in-the-middle attacks are also typically covertly compromised; and when web servers are used to compromise client's computers, the initial infection is typically covert.

In light of the prevalence of covert attacks, it behooves the user to consider what mitigation strategies can be taken to minimize the losses resulting from such attacks. Mitigation strategies include resetting passwords, changing private keys, re-installing servers, or re-instantiating virtual servers. Such strategies have notable characteristics in that they are often effective at securing the resource,

but they reveal little about past attacks or compromises. For example, if a server is re-installed, knowledge of when the server was compromised may be lost. Similarly, resetting a password does not reveal any information about the integrity of the previous password.

A second dimension of the attack space is the extent to which an attack is targeted or customized for a particular user [4,2]. DoS attacks and incidents of cyber-espionage are examples of targeted attacks. Typical examples of non-targeted attacks include spam and phishing. The dichotomy between targeted and non-targeted attacks is explained by Cormac Herley as a consequence of economic considerations of the attacker [4]. In that framework, an outsized number of users are both susceptible to and subject to scalable attacks which compromise their computer systems, but most are never targeted simply because they cannot be distinguished from low value targets. 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

Whether or not an attack is targeted is also important for the defender, because targeted and non-targeted attacks do different types of damage. For example, targeted attackers might read all of an organization's secret e-mails, causing economic damages of one type, while a non-targeted attacker might use the same compromised machine to send out spam, causing reputation loss, or machine blacklisting, or another separate type of damage. This dichotomy suggests that damages resulting from targeted and non-targeted attacks should be modeled additively.

The presence of both targeted and non-targeted covert attacks presents an interesting dilemma for a common user to choose a mitigation strategy against covert attacks. Strategies which are optimal against non-targeted covert 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-targeted attackers.

This paper fills the research gap induced by the aforementioned dichotomy, by considering the strategy spaces of users who may be subject to both targeted and non-targeted attacks. In our game, a defender must vie for a contested resource that is subject to the risk of compromise from both targeted and non-

targeted covert attacks. We explore the strategy space to find good mitigation strategies against this combination.

2 Related Work

2.1 Games of timing

Cybersecurity economics has been concerned with how to reduce the impact of the actions of financially or politically motivated adversaries who threaten computing resources and their users. Previous research in this domain has primarily focused on the choice between different canonical actions to prevent, deter or otherwise mitigate harm (e.g., [3,5,6]).

However, being successful in dynamic environments shifts the focus from selecting the most suitable option from a pool of alternatives to a decision problem of *when* to act to get an advantage over an opponent. For example, in tactical security scenarios it is important to jump to action at the right time to avoid a loss of money or even human life (see, for example, timing of interventions in international conflicts). To understand these scenarios, so-called games of timing have been studied with the tools of non-cooperative game theory since the cold war era (see, for example, [11,14]). For a detailed survey and summary of the theoretical contributions in this area, we refer the interested reader to [10].

2.2 `FlipIt`: Modeling Targeted Attacks

In response to recent high-profile stealthy attacks, researchers at RSA proposed the `FlipIt` model [13] to study such scenarios. In the original model, there are two players, a defender and an attacker, and a resource that they are both interested in maintaining control of. For each unit of time that a player is controlling the resource, she gains a fixed amount of benefit. Conversely, when a player is not in control, she gains no benefit from the resource. At any time instance, either player may "flip" the resource to gain control of it for some cost. Flipping while in control does not give the opponent control of the resource, therefore the players have to be careful not to make too many unnecessary flips to keep their costs low. This game can model, for example, the case of a password-protected account. Benefit is derived from using the account, and flipping the resource is analogous to the defender resetting the password or the attacker compromising it.

In the original `FlipIt` paper, dominant strategies and equilibria are studied for some simple cases [13]. Other researchers have worked on extensions [9,7]. For example, Laszka et al. extended the `FlipIt` game to the case of multiple resources. In addition, the usefulness of the `FlipIt` game has been investigated for various application scenarios [1,13].

In comparison to previous work, the `FlipIt` game is of interest because it combines a number of important decision-making factors [8]. First, it covers aspects of uncertainty about the game status by assuming that moves by the

players are "stealthy". Second, the game is played in continuous time and asynchronous fashion. Hence, ex-post the game appears to be divided in multiple periods of uneven length. Similarly, the number of actions that can be taken by the players is quasi-unlimited (if agents have an unrestricted budget). Third, action have a cost. That is, players do not only value the time in which they have possession of the board, but they also have to balance these benefits with the cost of gaining possession of the board.

The original `FlipIt` game has also been studied in an experiment with human subjects [8]. In that paper, the experimenters matched human participants with computerized opponents in several fast-paced rounds of the `FlipIt` game. The results indicate that participant performance improves over time; but that it is dependent on age, gender, and a number of individual difference variables. The researchers also show that human participants generally perform better when they have more information about the strategy of the computerized player; i.e., they are able to make use of such game-relevant information. This experimental work was extended to also include different visual presentation modalities for the available feedback during the experiment [12].

3 Model Definition

We model the covert compromise scenario as a non-zero-sum game. The player who is the rightful owner of the resource is called the defender, while the other players are called the attackers. The game starts at time $t = 0$ with the defender in control of the resource, and it is played indefinitely as $t \to \infty$. We assume that time is continuous.

We let D, A, and N denote the defender, the targeted attacker, and the non-targeted attackers respectively. At any time instance, player i may make a move, which costs her C_i. When the defender makes a move, the resource immediately becomes uncompromised for every attacker. When the targeted 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 the attack follows an exponential distribution. Formally, the probability that the attack has successfully finished in a amount of time is $1 - e^{-\lambda_A a}$, where λ_A is the rate parameter of the targeted attacker's attack time.

The attackers' moves are stealthy; i.e., the defender does not know when the resource got 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 player 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

benefit until the defender's next move. For the defender D, the benefit rate $b_D(t)$ up to time t is defined to be $-\sum_{i\in\{A,N\}} b_i(t)$ (i.e., what has been lost to the attackers). The relation between the defender's and attackers' 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

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

Table 2: List of Symbols

C_D	move cost for the defender
C_A	move cost for the targeted attacker
B_A	benefit received per unit of time for the targeted attacker
B_N	benefit received per unit of time for the non-targeted attackers
λ_A	rate of the targeted attacker's attack time
λ_N	rate of the non-targeted attacks' arrival

3.1 Types of Strategies for the Defender and the Targeted Attacker

Adaptive Strategies for Attackers 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 the function W depends on all the information that the attacker has, and we don't 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 Note that every renewal strategy is a special case of an adaptive strategy.

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

3.2 Non-Targeted Attacks

Suppose that there are N non-targeted attackers. In practice, N is very large, but the expected number of successful compromises is finite. As N goes to infinity, the probability that a given non-targeted attacker targets the defender approaches zero. Since the non-targeted attackers operate independently, *successful* non-targeted attacks arrive following a Poisson process. Furthermore, as the economic decisions of the non-targeted attackers depend on a very large pool of possible targets, the defender's effect on the decisions is negligible. Thus, the non-targeted attackers' strategies (that is, the attack rate) 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-targeted 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 is in many ways similar to `FlipIt`, it differs in three key assumptions. First, we assume that the defender's moves are *not stealthy*. The motivation for this is that an attacker must know whether she is in control of a resource if she receives benefits from it continuously. 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 the next time she tries to log in. Second, we assume that the targeted attacker's moves are *not instantaneous*, but take some time. The motivation for this is that an attack requires some time and 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 give analytical results on the game. We first consider the special case of a targeted attacker only (i.e., $\lambda_N = 0$), and then the general case of both targeted and non-targeted attackers.

We start with a discussion on the players' strategies. First, 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 use a renewal strategy.

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 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 for the attacker of this form 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

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

We let F_S denote the cumulative distribution function of S. Since $\lambda_A e^{-\lambda_A a} > 0$ for every $a \in \mathbb{R}_{\geq 0}$, if there exists an s for which $F_S(s) > 0$, then F_S is strictly increasing on $[s, \infty)$.

4.1 Nash Equilibrium for Targeted Attacker and Renewal Defender

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

Lemma 1. *Suppose that the attacker uses an adaptive strategy with a fixed probability distribution for choosing the time to wait until starting the attack. Then,*

- *not moving is the only best response if*

$$\frac{C_D}{B_A} = l F_S(l) - \int_{s=0}^{l} F_S(s) \, ds \tag{3}$$

  has no solution for l;
- *a periodic strategy whose period is the unique solution of Equation* (3) *is the only best response otherwise.*

Even though we cannot express the solution of Equation (3) in closed form, it can be easily found using numerical methods, as the right hand side is continuous and increasing.[4] Note that the equations presented in the subsequent lemmas and theorems of this paper can also be solved using numerical methods.

[4] We show that the right hand side is continuous and increasing in the proof of the lemma.

Proof. When playing a renewal strategy, the defender randomly selects the intervals between her consecutive moves according to the distribution generating the renewal strategy. In a best response, her strategy and, hence, every interval length in the support of the strategy's distribution has to minimize the defender's loss per unit of time. The defender's expected loss per unit of time for an interval of length l is

$$\frac{1}{l}\left(B_A \int_{s=0}^{l} f_S(s)(l-s)\ ds + C_D\right) \tag{4}$$

$$=\frac{1}{l}\left(B_A \left([F_S(s)(l-s)]_{s=0}^{l} - \int_{s=0}^{l} F_S(s)\cdot(-1)\ ds\right) + C_D\right) \tag{5}$$

$$=\frac{1}{l}\left(B_A \left(0 + \int_{s=0}^{l} F_S(s)\ ds\right) + C_D\right) \tag{6}$$

$$=\frac{1}{l}\left(B_A \int_{s=0}^{l} F_S(s)\ ds + C_D\right)\ . \tag{7}$$

To find the minimizing interval lengths (if there exists any), we take the derivative of (7) and solve it for equality with 0 as follows:

$$0 = \frac{d}{dl}\left[\frac{1}{l}\left(B_A \int_{s=0}^{l} F_S(s)\ ds + C_D\right)\right] \tag{8}$$

$$0 = -\frac{1}{l^2}\left(B_A \int_{s=0}^{l} F_S(s)\ ds + C_D\right) + \frac{1}{l} B_A F_S(l) \tag{9}$$

$$\int_{s=0}^{l} F_S(s)\ ds + \frac{C_D}{B_A} = l F_S(l) \tag{10}$$

$$\frac{C_D}{B_A} = l F_S(l) - \int_{s=0}^{l} F_S(s)\ ds\ . \tag{11}$$

Suppose that l^* is the least number for which this equation is satisfied. Then $l^* > 0$, and also $F(l^*) > 0$. This in turn implies that F_S is strictly increasing on $[l^*, \infty)$; and thus also the right hand side of the above equation is strictly increasing as a function of l on $[l^*, \infty)$. Therefore, if there is any solution to the above equation, then it is unique. Furthermore, this value of l is a minimizing value for the expected loss per unit of time as the second derivative at this

minimizing l^* is greater than zero:

$$\frac{d}{dl}\left[-\frac{1}{l^2}\left(B_A\int_{s=0}^{l}F_S(s)\;ds+C_D\right)+\frac{1}{l}B_AF_S(l)\right] \tag{12}$$

$$=\frac{2}{l^3}\left(B_A\int_{s=0}^{l}F_S(s)\;ds+C_D\right)+\left(-\frac{1}{l^2}\right)B_AF_S(l)$$

$$+\left(-\frac{1}{l^2}\right)B_AF_S(l)+\frac{1}{l}B_Af_S(l) \tag{13}$$

$$=\frac{2}{l^3}\left(B_A\int_{s=0}^{l}F_S(s)\;ds+C_D\right)+\left(-\frac{2}{l^2}\right)B_AF_S(l)+\frac{1}{l}B_Af_S(l)\;. \tag{14}$$

We care about the value of this expression when the first derivative is zero. Using this constraint, we obtain

$$\frac{2}{l^3}\left(B_A\int_{s=0}^{l}F_S(s)\;ds+C_D\right)+\left(-\frac{2}{l^2}\right)B_AF_S(l)+\frac{1}{l}B_Af_S(l) \tag{15}$$

$$=-\frac{2}{l}\left(-\frac{1}{l^2}\left(B_A\int_{s=0}^{l}F_S(s)\;ds+C_D\right)+\frac{1}{l}B_AF_S(l)\right)+\frac{1}{l}B_Af_S(l) \tag{16}$$

$$=-\frac{2}{l}(0)+\frac{1}{l}B_Af_S(l)>0\;. \tag{17}$$

Consequently, the only best response is the periodic strategy with the minimizing l^* as the period.

On the other hand, if Equation (11) is not satisfiable for l, then the only best response for the defender is to never move. When $l\to\infty$, the defender's expected loss per unit of time approaches B_A, which is equal to her loss for never moving. When $l\to 0$, her expected loss per unit of time goes to infinity due to the ever increasing costs. Consequently, if the expected loss per unit of time does not have a minimizing l, then it is always greater than B_A. □

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

Lemma 2. *Against a defender who uses a periodic strategy with period δ_D,*

- *never attacking is the only best response if*

$$\frac{C_A}{B_A}>\frac{e^{-\delta_D\lambda_A}-1}{\lambda_A}+\delta_D\;; \tag{18}$$

- *attacking immediately after the defender moved is the only best response if*

$$\frac{C_A}{B_A}<\frac{e^{-\delta_D\lambda_A}-1}{\lambda_A}+\delta_D\;; \tag{19}$$

– *both not attacking and attacking immediately are best responses otherwise.*

The lemma shows that the attacker should either attack immediately or not attack at all, but she should never wait to attack. Consequently, if the attacker uses her best response strategy, the defender can determine the optimal period of her strategy *solely based on the distribution of* A, which is an exponential distribution with parameter λ_A. This observation will be of key importance for characterizing the game's equilibria.

Proof. First, assume that the attacker does attack. Given that the attacker waits $w < \delta_D$ time before making her move, the expected amount of time she has the resource compromised until the defender's next move is

$$\int_{a=0}^{\delta_D - w} \lambda_A e^{-\lambda_A a}(\delta_D - w - a)da \ . \tag{20}$$

It is easy to see that the maximum of this equation is attained for $w = 0$. Therefore, if the attacker does attack, she attacks immediately. The expected amount of time she has the resource compromised until the defender's next move is

$$\int_{a=0}^{\delta_D} \lambda_A e^{-\lambda_A a}(\delta_D - a)da \tag{21}$$

$$= \left[\left(1 - e^{-\lambda_A a}\right)(\delta_D - a)\right]_{a=0}^{\delta_D} - \int_{a=0}^{\delta_D} \left(1 - e^{-\lambda_A a}\right)(-1)da \tag{22}$$

$$= \left(1 - e^{-\lambda_A \delta_D}\right) \underbrace{(\delta_D - \delta_D)}_{0} - \underbrace{\left(1 - e^{-\lambda_A 0}\right)}_{0}(\delta_D - 0) + \int_{a=0}^{\delta_D} 1 - e^{-\lambda_A a}\ da \tag{23}$$

$$= \int_{a=0}^{\delta_D} 1 - e^{-\lambda_A a}\ da = \delta_D - \left[-\frac{e^{-\lambda_A a}}{\lambda_A}\right]_{a=0}^{\delta_D} = \frac{e^{-\delta_D \lambda_A} - 1}{\lambda_A} + \delta_D \ . \tag{24}$$

Therefore, if the attacker does attack, her asymptotic benefit rate is

$$B_A \frac{\frac{e^{-\delta_D \lambda_A} - 1}{\lambda_A} + \delta_D}{\delta_D} \ , \tag{25}$$

and her payoff is

$$B_A \frac{\frac{e^{-\delta_D \lambda_A} - 1}{\lambda_A} + \delta_D}{\delta_D} - \frac{C_A}{\delta_D} \ . \tag{26}$$

Thus, when the above value is less than or equal to zero, never attacking is a best-response strategy; when the above value is greater than or equal to zero, always attacking immediately is a best-response strategy. When the above value is equal to zero, the attacker can decide whether to attack immediately or to not attack at all after each move of the defender. □

Equilibrium Based on the above 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 and the attacker uses an adaptive strategy. Then the game's equilibria can be described as follows.*

1. *If* $\frac{C_D}{B_A} = -le^{-\lambda_A l} + \frac{1-e^{-\lambda_A l}}{\lambda_A}$ *does not have a solution for* l*, then there is a unique equilibrium in which the defender does not move and in which the attacker attacks exactly once at the beginning of the game.*
2. *If* $\frac{C_D}{B_A} = -le^{-\lambda_A l} + \frac{1-e^{-\lambda_A l}}{\lambda_A}$ *does have a solution* δ_D *for* l*, then*
   (a) *if* $\frac{C_A}{B_A} \leq \frac{e^{-\delta_D \lambda_A} - 1}{\lambda_A} + \delta_D$*, then there is a unique equilibrium in which the defender plays a periodic strategy with period* δ_D*, and the attacker attacks immediately after the defender's each move;*
   (b) *if* $\frac{C_A}{B_A} > \frac{e^{-\delta_D \lambda_A} - 1}{\lambda_A} + \delta_D$*, then there is no equilibrium.*

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" the game since any effort to gain control of the resource is not profitable for her, and the attacker will have control of the resource all the time. In the second case, no player is at an overwhelming advantage. Both the defender and the attacker are actively trying to gain control of the resource, and both succeed from time to time. In the third case, the defender is at an overwhelming advantage. However, this does not lead to an equilibrium. If the defender moves with a sufficiently high rate, she makes moving unprofitable for the attacker. But if the attacker decides not to move, the defender is also better off not moving, as this decreases her cost. However, once the defender stops moving, it is again profitable for the attacker to move, which in turn triggers the defender to start moving.

Proof. First, we have from Lemma 1 that in any equilibrium, the defender either never moves or uses a periodic strategy. If the defender never moves, then the best strategy for the attacker is to attack immediately after the game starts. Now, if the defender moves using a periodic strategy, we have from Lemma 2 that the attacker either never attacks or attacks immediately. This leaves us with two strategies for defender and two strategies for attacker from which all equilibria must be composed.

Second, we show that there is no equilibrium in which the attacker never attacks. To see this, suppose that the attacker never attacks. Then the defender's best response is to never move, because this preserves control of the resource while minimizing the defender's cost. But if the defender never moves, then it is advantageous for the attacker to compromise the resource immediately after the start of the game. So this situation is not an equilibrium.

Next, we analyze the situation where a defender never moves. In this circumstance, the attacker attacks once and controls the resource for the duration of

the game. From Lemma 1, we see that this is indeed a unique equilibrium if

$$\frac{C_D}{B_A} = lF_S(l) - \int_{s=0}^{l} F_S(s)\ ds \tag{27}$$

$$= l\left(1 - e^{-\lambda_A l}\right) - \int_{s=0}^{l} 1 - e^{-\lambda_A s}\ ds \tag{28}$$

$$= l - le^{-\lambda_A l} - \frac{e^{-\lambda_A l} - 1}{\lambda_A} - l \tag{29}$$

$$= -le^{-\lambda_A l} + \frac{1 - e^{-\lambda_A l}}{\lambda_A} \tag{30}$$

does not have a solution in $\mathbb{R}_{\geq 0}$ for l.

Finally, we consider the scenario where the defender plays a periodic strategy with period δ_D. In this case, Lemma 2 gives conditions for the best response of the attacker. Either the attacker never moves or the attacker attacks immediately. Since we know that there is no equilibrium in which an attacker never moves, we concern ourselves in the theorem only with the circumstances under which the attacker has a reason to attack immediately. From Lemma 2, the condition for this is $\frac{C_A}{B_A} \leq \frac{e^{-\delta_D \lambda_A} - 1}{\lambda_A} + \delta_D$. □

4.2 Equilibrium for Both Targeted and Non-Targeted Attackers

Defender's Best Response Again, we begin our analysis by finding the defender's best response strategy.

Lemma 3. *Suppose that the non-targeted attacks arrive according to a Poisson process with rate λ_N, and the targeted attacker uses an adaptive strategy with a fixed wait time distribution given by the cumulative function F_S. Then,*

- *not moving is the only best response if*

$$C_D = B_A \left(lF_S(l) - \int_{s=0}^{l} F_S(s)\ ds \right) + B_N \left(-le^{-\lambda_N l} + \frac{1 - e^{-\lambda_N l}}{\lambda_N} \right) \tag{31}$$

  has no solution for l;
- *a periodic strategy whose period is the solution to Equation* (31) *is the only best response otherwise.*

Proof. The outline of the proof is similar to that of Lemma 1.

The defender's expected loss per unit of time for an interval of length l is

$$\frac{1}{l}\left(B_A\int_{s=0}^{l} f_S(s)(l-s)\ ds + B_N\int_{a=0}^{l}(l-a)\lambda_N e^{-\lambda_N a}da + C_D\right) \tag{32}$$

$$=\frac{1}{l}\left(B_A\left([F_S(s)(l-s)]_{s=0}^{l}\int_{s=0}^{l}F_S(s)\ ds\right) + B_N\left(\frac{e^{-\lambda_N l}-1}{\lambda_N}+l\right)+C_D\right) \tag{33}$$

$$=\frac{1}{l}\left(B_A\int_{s=0}^{l}F_S(s)\ ds + B_N\left(\frac{e^{-\lambda_N l}-1}{\lambda_N}+l\right)+C_D\right)\ . \tag{34}$$

To find the minimizing interval lengths (if there exists any), we take the derivative of (34) and solve it for equality with 0 as follows:

$$0=\frac{d}{dl}\left[\frac{1}{l}\left(B_A\int_{s=0}^{l}F_S(s)\ ds + B_N\left(\frac{e^{-\lambda_N l}-1}{\lambda_N}+l\right)+C_D\right)\right] \tag{35}$$

$$0=-\frac{1}{l^2}\Bigg(\quad B_A\left(\int_{s=0}^{l}F_S(s)\ ds - lF_S(l)\right) + B_N\frac{e^{-\lambda_N l}(\lambda_N l - e^{\lambda_N l}+1)}{\lambda_N}+C_D\Bigg) \tag{36}$$

$$C_D=B_A\left(lF_S(l)-\int_{s=0}^{l}F_S(s)\ ds\right)+B_N\left(-le^{-\lambda_N l}+\frac{1-e^{-\lambda_N l}}{\lambda_N}\right)\ . \tag{37}$$

From the proof of Lemma 1, we have that the first term of the right hand side is monotonically increasing. Furthermore, the second term is strictly increasing, as its derivate is $\lambda_N l e^{-\lambda_N l} > 0$. Thus, the right hand side is strictly increasing, which implies that if there is an l^* for which the equality holds, it has to be unique. Furthermore, this l^* is a minimizing value as the second derivative is greater than zero:

$$\frac{d}{dl}\Bigg[-\frac{1}{l^2}\Bigg(\quad B_A\left(\int_{s=0}^{l}F_S(s)\ ds - lF_S(l)\right) + B_N\frac{e^{-\lambda_N l}(\lambda_N l - e^{\lambda_N l}+1)}{\lambda_N}+C_D\Bigg)\Bigg] \tag{38}$$

$$=\frac{1}{l^3}\Bigg(\quad B_A\left(2\int_{s=0}^{l}F_S(s)\ ds - 2lF_S(l)+l^2 f_S(l)\right) + B_N\frac{e^{\lambda_N l}(\lambda_N^2 l^2+2\lambda_N l - 2e^{\lambda_N l}+2)}{\lambda_N}+2C_D\Bigg)\ . \tag{39}$$

We care about the value of this expression when the first derivative is zero. Using this constraint, we obtain

$$\frac{1}{l^3}\Bigg(\quad B_A \left(2\int_{s=0}^{l} F_S(s)\ ds - 2lF_S(l) + l^2 f_S(l) \right) + B_N \frac{e^{\lambda_N l}(\lambda_N^2 l^2 + 2\lambda_N l - 2e^{\lambda_N l} + 2)}{\lambda_N} + 2C_D \Bigg) \tag{40}$$

$$= -\frac{2}{l}\left(B_A \left(\int_{s=0}^{l} F_S(s)\ ds - lF_S(l) \right) + B_N \frac{e^{-\lambda_N l}(\lambda_N l - e^{\lambda_N l} + 1)}{\lambda_N} + C_D \right) + \frac{1}{l}\left(B_A f_S(l) + B_N e^{-\lambda_N l}\lambda_N \right) \tag{41}$$

$$= -\frac{2}{l}\Big(0\Big) + \frac{1}{l}\left(B_A f_S(l) + B_N e^{-\lambda_N l}\lambda_N \right) > 0\ . \tag{42}$$

Consequently, the only best response is the periodic strategy with the minimizing l^* as the period.

On the other hand, if Equation (11) is not satisfiable for l, then the only best response for the defender is to never move. When $l \to \infty$, the defender's expected loss per unit of time approaches $B_A + B_N$, which is equal to her loss for never moving. When $l \to 0$, her expected loss per unit of time goes to infinity due to the ever increasing costs. Therefore, if there is no minimizing l, then the expected loss per unit of time is always greater than $B_A + B_N$. □

Equilibrium Since the targeted attacker's payoff and, consequently, best response are not directly affected by the presence of non-targeted attackers, we can use Lemma 2 and the above lemma to describe the equilibria of the game.

Theorem 2. *Suppose that the defender uses a renewal strategy, the targeted attacker uses an adaptive strategy, and the non-targeted attacks arrive according to a Poisson process with rate λ_N. Then the game's equilibria can be described as follows.*

1. *If $C_D = B_A\left(-le^{-\lambda_A l} + \frac{1-e^{-\lambda_A l}}{\lambda_A}\right) + B_N\left(-le^{-\lambda_N l} + \frac{1-e^{-\lambda_N l}}{\lambda_N}\right)$ does not have a solution for l, then there is a unique equilibrium in which the defender does not move and in which the attacker attacks exactly once at the beginning of the game.*
2. *If $C_D = B_A\left(-le^{-\lambda_A l} + \frac{1-e^{-\lambda_A l}}{\lambda_A}\right) + B_N\left(-le^{-\lambda_N l} + \frac{1-e^{-\lambda_N l}}{\lambda_N}\right)$ does have a solution δ_D for l, then:*
   (a) *If $\frac{C_A}{B_A} \leq \frac{e^{-\delta_D \lambda_A}-1}{\lambda_A} + \delta_D$, then there is a unique equilibrium in which the defender plays a periodic strategy with period δ_D, and the targeted attacker moves immediately after the defender's each move.*
   (b) *If $\frac{C_A}{B_A} > \frac{e^{-\delta_D \lambda_A}-1}{\lambda_A} + \delta_D$, then*

- *if* $C_D = B_N\left(-le^{-\lambda_N l} + \frac{1-e^{-\lambda_N l}}{\lambda_N}\right)$ *has a solution* δ'_D *for* l*, and* $\frac{C_A}{B_A} \geq \frac{e^{-\delta'_D \lambda_A}-1}{\lambda_A} + \delta'_D$*, then there is a unique equilibrium in which the defender plays a periodic strategy with period* δ'_D *and the targeted attacker never moves;*
- *otherwise, there is no equilibrium.*

By comparing the equation determining the defender's strategy in the theorem above to the equation in Theorem 1, we see that the parameter values B_A and C_D for which there is a solution is larger in the theorem above. Thus, the defender is more likely to move instead of giving it up when there is a threat of non-targeted attacks.

Proof. Cases 1. and 2. (a) follow from Lemma 2 and Lemma 3 using the argument as the proof of Theorem 1.

In Case 2. (b), there could be no equilibrium when the defender faced only a targeted attacker (Theorem 1), since the defender had no incentives to move if the targeted attacker did not move. However, when there are non-targeted attacker present as well, the defender moving periodically and the targeted attacker never moving can be an equilibrium. The necessary and sufficient conditions for this are that moving periodically is a best response for the defender against non-targeted attackers only (the existence of δ'_D) and that never attacking is a best-response for the targeted attacker against this period δ'_D. □

5 Numerical Illustrations

In this section, we present numerical results on our game.

First, in Figure 1, we study the effects of varying the value of the resource, that is, the unit benefit B_A received by the targeted attacker. Figure 1a shows both players' payoffs for various values of B_A (the defender's periods for the same setup are shown by Figure 1b). The figure shows that the defender's payoff is strictly decreasing, which is not surprising: the more valuable the resource is, the higher the cost of security is for the defender. The 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 ($B_A < 1$), the defender does not protect the resource, as it is not valuable enough to defend. Accordingly, Figure 1b shows no period for this region. In this case, the attacker's payoff is equal to simply the value of the resource. However, once the value of the resource reaches 1, the defender starts protecting it. At this point, the attacker's payoff drops as she no longer has the resource compromised all the time. For higher values, the defender balances between losses due to compromise and moving costs, which means that the time the resource is compromised decreases steadily as its value increases.

In Figure 2, we study the effects of varying the defender's move cost C_D. Figure 2a shows both players' payoffs for various values of C_D (the defender's periods for the same setup are shown by Figure 2b). The figure shows that the

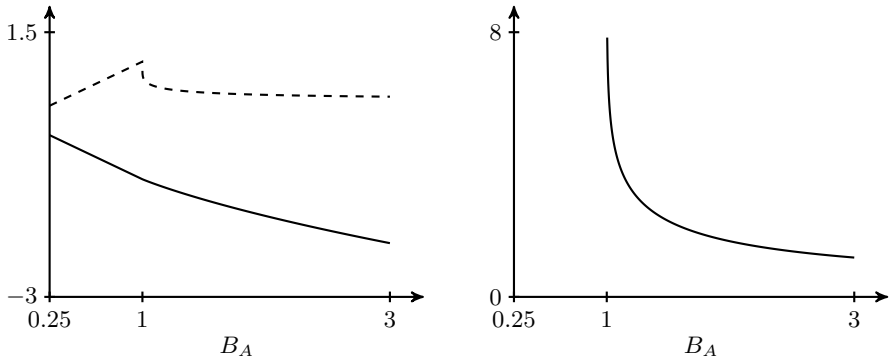

(a) The defender's and the targeted attacker's payoffs (solid and dashed lines, respectively) as a function of B_A.

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

Fig. 1: The effects of varying the unit benefit B_A received by the targeted attacker.

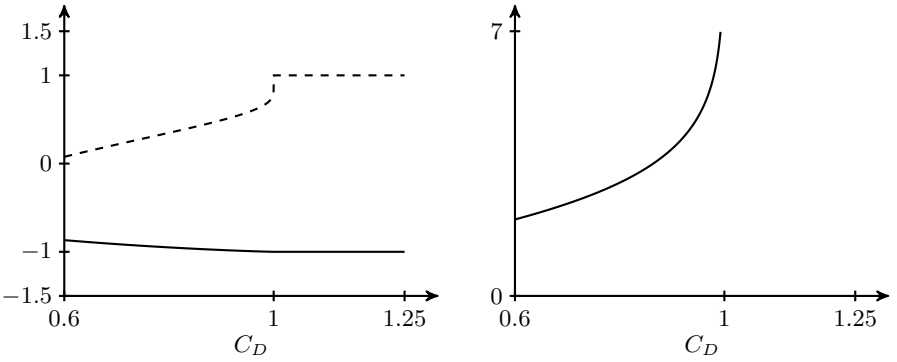

(a) The defender's and the targeted attacker's payoffs (solid and dashed lines, respectively) as a function of C_D.

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

Fig. 2: The effects of varying the defender's move cost C_D.

defender's payoff is decreasing, while the attacker's payoff is increasing, which is again not surprising: the more costly it is to defend the resource, the greater the attacker's advantage is.

For lower costs, no player is at an overwhelming advantage, as both players try to control the resource and succeed from time to time. As the cost increases, the defender's payoff steadily decreases, while the attacker's payoff steadily increases. For higher costs, the attacker is at an overwhelming advantage. In this case, the defender never moves, while the attacker moves once. Hence, their payoffs are -1 and 1, respectively.

6 Conclusions

Targeted and non-targeted attacks are born of different motivations and have different types of consequences. In this paper, we modeled a regime in which a defender must vie for a contested resource against both targeted and non-targeted covert attacks.

As a principal result, we found that the most effective strategy against both types of attacks (and also against their combination) is the periodic strategy. This result can be surprising considering the simplicity of this strategy, but it also serves as a theoretical justification of the periodic password and cryptographic key renewal practices. Furthermore, this contradicts the lesson learned from the `FlipIt` model [13], which suggests that a defender playing against an adaptive attacker should use an unpredictable strategy.

We also found that a defender is more likely to stay in play and bear the costs of periodic risk mitigation if she is threatened by non-targeted attacks. While this result seems very intuitive, it is not obvious, as we also demonstrated that a very high level of either threat type can force the defender to abandon all hope and stop moving.

Our work can be extended in multiple directions. First, even though the exponential attack time distribution can be well-motivated for a number of resources, it would be worthwhile to extend our model to general distributions with some mild assumptions only. Second, our model focuses on medium-profile targets that are susceptible to both targeted and non-targeted attacks, but it could be easily extended to a broader range by having a susceptibility probability for each type.

Acknowledgements

We gratefully acknowledge the support of the Penn State Institute for CyberScience. We also thank the reviewers for their comments on an earlier draft of the paper.

References

1. Kevin Bowers, Marten Dijk, Robert Griffin, Ari Juels, Alina Oprea, Ronald Rivest, and Nikos Triandopoulos. Defending against the unknown enemy: Applying FlipIt to system security. In Jens Grossklags and Jean Walrand, editors, *Decision and Game Theory for Security*, volume 7638 of *Lecture Notes in Computer Science*, pages 248–263. Springer, 2012.
2. Eoghan Casey. Determining intent - opportunistic vs targeted attacks. *Computer Fraud & Security*, 2003(4):8–11, 2003.
3. Jens Grossklags, Nicolas Christin, and John Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proceedings of the 17th International World Wide Web Conference (WWW)*, pages 209–218, 2008.

4. Cormac Herley. The plight of the targeted attacker in a world of scale. In *Proceedings of the 9th Workshop on the Economics of Information Security (WEIS)*, 2010.
5. Benjamin Johnson, Rainer Böhme, and Jens Grossklags. Security games with market insurance. In John Baras, Jonathan Katz, and Eitan Altman, editors, *Decision and Game Theory for Security*, volume 7037 of *Lecture Notes in Computer Science*, pages 117–130. Springer, 2011.
6. Aron Laszka, Mark Felegyhazi, and Levente Buttyán. A survey of interdependent security games. Technical Report CRYSYS-TR-2012-11-15, CrySyS Lab, Budapest University of Technology and Economics, Nov 2012.
7. Aron Laszka, Gabor Horvath, Mark Felegyhazi, and Levente Buttyan. FlipThem: Modeling targeted attacks with FlipIt for multiple resources. Technical report, Budapest University of Technology and Economics, 2013.
8. Alan Nochenson and Jens Grossklags. A behavioral investigation of the FlipIt game. In *Proceedings of the 12th Workshop on the Economics of Information Security (WEIS)*, 2013.
9. Viet Pham and Carlos Cid. Are we compromised? Modelling security assessment games. In Jens Grossklags and Jean Walrand, editors, *Decision and Game Theory for Security*, volume 7638 of *Lecture Notes in Computer Science*, pages 234–247. Springer, 2012.
10. Tadeusz Radzik. Results and problems in games of timing. *Lecture Notes-Monograph Series, Statistics, Probability and Game Theory: Papers in Honor of David Blackwell*, 30:269–292, 1996.
11. Tadeusz Radzik and Krzysztof Orlowski. A mixed game of timing: Investigation of strategies. *Zastosowania Matematyki*, 17(3):409–430, 1982.
12. David Reitter, Jens Grossklags, and Alan Nochenson. Risk-seeking in a continuous game of timing. In *Proceedings of the 13th International Conference on Cognitive Modeling (ICCM)*, pages 397–403, 2013.
13. Marten van Dijk, Ari Juels, Alina Oprea, and Ronald Rivest. FlipIt: The game of "stealthy takeover". *Journal of Cryptology*, 26:655–713, October 2013.
14. Vitaliy Zhadan. Noisy duels with arbitrary accuracy functions. *Issledovanye Operacity*, 5:156–177, 1976.