

Exponential Discounting in Security Games of Timing

Jonathan Merlevede
imec-DistriNet
KU Leuven
jonathan.merlevede@cs.kuleuven.be

Jens Grossklags
Chair for Cyber Trust
Technical University of Munich
jens.grossklags@in.tum.de

Benjamin Johnson
Chair for Cyber Trust
Technical University of Munich
benjamin.johnson@tum.de

Tom Holvoet
imec-DistriNet
KU Leuven
tom.holvoet@cs.kuleuven.be

Abstract—Strategic game models of defense against stealthy, targeted attacks that cannot be prevented but only mitigated are the subject of a significant body of recent research, often in the context of advanced persistent threats (APTs). In these game models, the timing of attack and defense moves plays a central role. A common assumption, in this literature, is that players are indifferent between costs and gains now and those in the distant future, which conflicts with the widely accepted treatment of intertemporal choice across economic contexts.

This paper investigates the significance of this assumption by studying changes in optimal player behavior when introducing time discounting. Specifically, we adapt a popular model in the games of timing literature, the FlipIt model, by allowing for exponential discounting of gains and costs over time. We investigate changes of best responses and the location of Nash equilibria through analysis of two well-known classes of player strategies: those where the time between players' moves is constant, and a second class where the time between players' moves is stochastic and exponentially distributed.

By introducing time discounting in the framework of games of timing, we increase its level of realism as well as applicability to organizational security management, which is in dire need of sound theoretic work to respond to sophisticated, stealthy attack vectors.

Index Terms—games of timing, game theory, APT, discounting

1. Introduction

Over the last years many successful high-profile targeted attacks against supposedly well-protected, secure targets have been documented such as Iranian industrial control systems [1], the U.S. Office of Personnel Management [2], [3], the U.S. Internal Revenue Service [4], large telecommunication providers [5], [6] and major health care insurance companies [7]. The perpetrators of these attacks are commonly referred to as advanced persistent threats (sAPTs) and their attacks often remain unnoticed for a number of months, years, or possibly forever. In fact, analysis of data about reported security incidents

shows that organizations need on average about 200 days to detect successful attacks [8].

It is apparent that defense against APTs through perfectly effective preventative investments is an impossible goal in most contexts [9], [10], which places additional emphasis on mitigation strategies, in particular, against stealthy threats. Such strategies may include in-depth security audits, or resetting of key parts of the IT infrastructure to disrupt successful security compromises. However, due to the sparsity of data about such attacks, a principled theoretical approach is needed to understand the economic rationale for such mitigation strategies as part of a science-based security management approach.

The consideration of *time* has to be a central strategic element of economic decision-making to account for stealthy threats. Research in this context can build on a comprehensive body of literature of *games of timing*, which emerged during the cold war period (see, for example, [11]). More recently, this area of research received the addition of a series of papers about the so-called FlipIt game beginning with research by van Dijk et al. [12], which captures the competitive dynamics to control a contested resource under the assumption of attack stealthiness and canonical player strategies.

While this line of work has explored diverse facets of security decision-making, a crucial limitation of these studies is that players are indifferent between costs and gains now in comparison to those in the (distant) future, which conflicts with the widely accepted treatment of intertemporal choice across economic contexts by applying some notion of *discounting*. In particular, discounting of economic costs and gains with a monetary equivalent is seen as uncontroversial. Importantly, also non-monetary factors such as health benefits are typically associated with a positive discount factor [13]; even though the debate about the concrete specifics is more lively. Similarly, security decision-making may involve reasoning about monetary and non-monetary factors.

Economists distinguish between two high-level forms of discounting: *time discounting* and time preference. The former refers to *any* form of consideration to care less about an uncertain future in economic terms, while the latter refers to the robust empirical finding of the psychological tendency to prefer immediate utility over delayed

outcomes [14]. Our work focuses on the former broad definition.

As a matter of public policy and organizational decision-making it is critical to deliberate about the appropriate consideration of discounting. For example, in the security context, underappreciation of long-term consequences can falsely justify inadequate security measures at the expense of future competitiveness (see the problem context of cyberespionage). In our assessment, the debate of appropriately accounting for intertemporal decision-making and discounting in the context of security is still in its infancy, which stands in contrast to other disciplines – such as public health or environmental policy – where it has been central (see, for example, [15], [16]). With our work, we want to shed light on this challenge in the particularly troubling context of stealthy security threat mitigation.

In this work, we present and analyze a timing game similar to FlipIt [12] in which two players vie for control of a single resource. The players in our model apply exponential discounting to both the resource value and cost to flip ownership. Our model exhibits a number of features that differentiate it substantially from FlipIt. For example, the total cost of the resource is finite; the defender has a unique advantage due to the value being higher when time begins. Periodic strategies are also not necessarily optimal. Our penultimate result is a full characterization of the game’s Nash Equilibria under the assumption that both players choose strategies from the same restricted strategy space, where the space is among exponential strategies or periodic strategies. These results have practical implications for both the incentives of time-based security decisions, as well as their corresponding outcomes.

The remainder of the paper is organized as follows. In Section 2, we discuss related work. Next, in Section 3, we define our game-theoretic model. We analyze our model in Section 4, and discuss implications of our findings in Section 5. We offer concluding remarks in Section 6.

2. Related work

The past decade has seen a rapid increase in the use of game theory as a tool for studying security-related decision-making. Manshaei, Zhu, Alpcan, *et al.* [17] and Laszka, Felegyhazi, and Buttyán [18] give an excellent overview of the achievements in the space of general security games. We restrict our discussion here to work that is directly related to games of timing, and specifically to FlipIt, a game that was invented by researchers at RSA [12], [19] with the intent of capturing the essential nature of advanced, stealthy attacks such as those performed by entities that represent an APT. FlipIt has since received significant attention from the research community, such that there now exist many adapted and extended versions of the original game. In the following, we classify this literature by key distinguishing aspects.

Player types. Laszka, Johnson, and Grossklags [20], [21] have investigated the influence of including non-targeted attackers in the FlipIt model. Feng, Zheng, Hu, *et al.* [22] and Hu, Li, Fu, *et al.* [23] modified the game by considering insider threat actors. Feng, Zheng, Hu, *et al.* [22] accomplish this by adding a third player,

an insider, to the model. The insider derives benefits from the resource, when it is under control of the defender, by selling information to the attacker who will learn about ways to decrease the cost of attacks.

Attack model. In the basic FlipIt game, moves by both the attacker and defender are assumed to be instantaneous and always successful. Farhang and Grossklags [24] introduce the idea of imperfect defensive moves with a quality level $\alpha \in [0, 1]$ that expresses the fraction of the resource that remains under the control of the attacker after a flip by the defender. Zhang, Zheng, and Shroff [25] and Laszka, Johnson, and Grossklags [20], [21] capture the realistic notion that attacks are complex and take a random amount of time before taking effect. Johnson, Laszka, and Grossklags [26] redefine the probability of success of an attack as a function of time. They also consider that the cost of flipping may be time-dependent.

Game duration. In the basic FlipIt game, the game has an infinite time horizon and players compete for the resource forever. Zhang, Zheng, and Shroff [25] and Johnson, Laszka, and Grossklags [26] assume that the game ends at a fixed pre-defined point in time. Pham and Cid [27] propose a variation of the FlipIt game in which each action makes it more costly for the opponent to take over the resource again; effectively reducing the game to a finite version.

Multiple resources. Laszka, Horvath, Felegyhazi, *et al.* [28] consider two ways of composing resources: one where the attacker receives gain when she is in control of at least one resource (OR-model) and one where she receives gain only when in control of all of the resources (AND-model). Leslie, Sherfield, and Smart [29] generalize this to a model where the attacker has to compromise a threshold fraction of the defender’s resources before receiving any gain. Zhang, Zheng, and Shroff [25] also consider multiple resources, but model no interaction between them except through a resource constraint imposed on players in the form of a maximum play frequency that is shared across resources.

Stealthiness. Much of the follow-up work on FlipIt has made changes to the assumption of perfect stealthiness. Often the defender is assumed to be completely overt [20], [21], [24], [25]. Besides the conceptual difference, overttness also allows for a different characterization of the FlipIt game as a convex optimization problem [25]. Pham and Cid [27] add a new audit action to the game, which allows a player to query the current owner of the resource. The insider player introduced by Hu, Li, Fu, *et al.* [23] is at risk of being caught when selling information, which is integrated in her utility function.

Other Changes. Johnson, Laszka, and Grossklags [26] consider a discretized version of a timing game similar to FlipIt, in which players are only allowed to make decisions at discrete points in time. Discretization of time is especially relevant for defender moves, which often have to be performed according to some schedule so as not to interrupt business operations (e.g. only at night) [30]. Zhang, Zheng, and Shroff [25] impose budget constraints on players that limit the maximum flip frequency, a practical consideration that is ignored in other treatments of FlipIt. Pawlick, Farhang, and Zhu [31] define a meta-game that consists of a signalling game and a FlipIt game. The

parameters of the FlipIt game are defined by the outcome of the signalling game and vice versa.

Behavioral studies. Nochenson and Grossklags [32], Reitter, Grossklags, and Nochenson [33] and Grossklags and Reitter [34] study the behavior of human participants when playing FlipIt-like games against computerized opponents.

We are unaware of any studies that investigate the impact of discounting in FlipIt-like game models.

3. Model definition

This section introduces our model for stealthy timing-based security games with discounting. In this two-player game, a defender (D) and an attacker (A) are vying for control over a central resource in an interaction that starts at a set point in time $t = 0$ and that goes on forever. Whoever is in control receives value from the resource at a well-defined, constant rate. To obtain control, either player can choose to pay a fixed cost to execute an instantaneous move at any point in (continuous) time. The resource is always controlled by the last player to execute such a move; simultaneous moves cancel out and do not cause the owner of the resource to change. We also refer to moving as performing a ‘flip’. Neither player can observe when the other player moves or when she is in control of the resource. Both the value generated by the resource and the cost of performing a move decrease exponentially over time.

Figure 1 graphically depicts an instance of our game. Blue and red dots on the time axis indicate defender and attacker moves, respectively. Moves performed while already in control of the resource have no effect (e.g. by the defender at time $t = 1.5$), while moves performed while the other player is in control of the resource trigger a change of ownership (e.g. by the attacker at time $t = 3$). The total income received by the defender for being in control of the resource is proportional to the blue area; for the attacker this is the red area. We refer to a player’s total income as her *gain*. The discounted cost of a move is proportional to the length of the dotted line coming from its corresponding dot. We refer to the sum of all these discounted costs as the player’s *cost*. Total player *utility* is equal to the difference between gains and costs. Note that the defender starts out in control of the resource; in terms of gains, this gives her an advantage over the attacker that we refer to as the *defender advantage*. The defender advantage is proportional to the area colored in a darker shade of blue. Note that on the figure, the defender is more impatient than the attacker: her valuation of future gains and costs decreases faster with the passing of time.

The remainder of this section formally defines all the elements of our model and introduces the periodic and exponential strategy spaces. Our model is equal to the FlipIt model presented by van Dijk, Juels, Oprea, *et al.* [12] with more general definitions for gains and costs that allow discounting over time. Readers that are familiar with FlipIt can limit their attention to Section 3.4 and Section 3.5.

TABLE 1. TABLE OF SYMBOLS

Symbol	Description
D	The defending, “good” player (Section 3.1)
A	The attacking, “bad” player (Section 3.1)
i	An arbitrary player, $i \in \{D, A\}$ (Section 3.1)
j	The other player, $j \in \{D, A\} \setminus \{i\}$ (Section 3.1)
\mathbf{t}_i	Sequence of times at which player i moves (Section 3.2)
$t_{i,n}$	Time of the n th move by player i (Section 3.2)
\mathbf{t}	Sequence of times of moves by either player (Section 3.3)
t_n	Time of the n th move made by either player (Section 3.3)
$GS(t)$	The game state function (Section 3.3)
λ_i	Player i ’s discount factor for gains (Section 3.1)
Λ_i	Player i ’s discount factor for costs (Section 3.1)
c_i	Player i ’s (instantaneous) move cost (Section 3.1)
G_i	Player i ’s gain (Section 3.4)
G_D^+	The defender advantage (Section 3.4)
\bar{G}_i	The anonymous gain of player i (Section 3.4)
$\bar{G}_{i \lambda}$	Anonymous gain of player i for discount rate λ (Section 3.4)
C_i	Player i ’s cost (Section 3.5)
u_i	Player i ’s benefit or utility (Section 3.6)
η_i	Play rate of i playing an exponential strategy (Sect. 3.7.1)
φ_i	Phase of i playing a periodic strategy (Sect. 3.7.2)
δ_i	Period of i playing a periodic strategy (Sect. 3.7.2)
α_i	Play rate of i playing a periodic strategy (Sect. 3.7.2)

3.1. Players

There are two players: a defender D and an attacker A . We will also often refer to player i ; this is an arbitrary player and can be either D or A . Player j then refers to the other player.

Each player i is characterized by three real parameters: her (instantaneous) *move* or *flip cost* $c_i > 0$, her *discount factor* for gains $\lambda_i \in (0, 1)$, and her discount factor for costs $\Lambda_i \in (0, 1)$.

These discount factors indicate how she values instantaneous gains and costs compared to gains and costs in the future. As λ_i and Λ_i come closer to 1, player i places a higher value on the future, or equivalently that she discounts at a lower rate. Mathematically, we define the *discount rate* as $-\ln \lambda_i$. Note that we do not consider the possibility of players valuing the future higher than the present.

3.2. Player strategies

A player’s strategy determines when that player performs moves. Because we assume that moves by players are not observable (perfectly stealthy) and that players also do not obtain any other information about the game as it progresses, we can think of a player strategy simply as a probability distribution over sets of times at which the player moves. We can draw a set of moves from the distribution before the game starts.

Formally, let

$$\mathbf{t}_i = (t_{i,0}, t_{i,1}, t_{i,2}, \dots)$$

be a strictly increasing sequence of times at which player i moves. It is not allowed for a player to move twice at exactly the same point in time. The length of \mathbf{t}_i can be finite or infinite. However, we require that there are always a finite number of members of \mathbf{t}_i falling within any finite time interval. At its most general, a player strategy is simply a probability distribution over a set of possible \mathbf{t}_i .




Figure 1. Instance of a discounted security game of timing with periodic play ($\varphi_D = 1.5$, $\delta_D = 3$, $\varphi_A = 3$, $\delta_A = 4$). Each tick mark denotes the passing of one unit of time.

When deriving player gains and costs, we will not reason in terms of distributions over \mathbf{t}_i . Instead, we limit our attention to interesting subsets of the general strategy space that have real-world relevance as well as a much more convenient mathematical description. We introduce these strategy spaces in Section 3.7.

3.3. Game state

Let

$$\mathbf{t} = (\mathbf{t}_D \cup \mathbf{t}_A) \setminus (\mathbf{t}_D \cap \mathbf{t}_A) = (t_0, t_1, t_2, \dots)$$

be the strictly increasing sequence of times of moves made by either player, leaving out the points in time where both players move simultaneously – moves made at the same point in time cancel out and do not change the game state. Let $\text{Actor}: \mathbf{t} \rightarrow \{D, A\}$ identify who is making a move at a certain time:

$$\text{Actor}: t \mapsto \begin{cases} D & \text{if } t \in \mathbf{t}_D \text{ and} \\ A & \text{if } t \in \mathbf{t}_A. \end{cases} \quad (1)$$

We can then define the *game state function* $\text{GS}: \mathbb{R}^+ \rightarrow \{D, A\}$, which indicates who was the last to flip and is therefore deriving value from the resource at time t :

$$\text{GS}: t \mapsto \begin{cases} \text{Actor}(\max(\mathbf{t} \cap [0, t])) & \text{if } \mathbf{t} \cap [0, t] \neq \emptyset \text{ and} \\ D & \text{otherwise.} \end{cases} \quad (2)$$

Note that the defender starts out in control of the resource: $\text{GS}(t)$ always equals D at any time t before the first flip t_0 . If the attacker never moves, player D remains in control forever.

As we have seen, strategies determine distributions over \mathbf{t}_i . We can also treat the \mathbf{t}_i as random variables and GS as a stochastic process determined by the players' strategies.

3.4. Player gains

The (discounted) gain of player i is the value that she derives from being in control of the resource. Player gains form the positive part of the expressions for player utilities.

Let $\text{GS}_i: \mathbb{R}^+ \rightarrow \{0, 1\}$ be the stochastic process indicating whether player i is in control of the resource at a certain time:

$$\text{GS}_i: t \mapsto \begin{cases} 1 & \text{if } \text{GS}(t) \text{ equals } i \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Player i 's (discounted) *gain* is then just the expected exponentially weighted integral of $\text{GS}_i(t)$ over all of time, normalized with respect to the total (discounted) value of the resource:

$$\begin{aligned} G_i &= \frac{\mathbb{E} \left[\int_{\tau=0}^{+\infty} \text{GS}_i(\tau) \lambda_i^\tau d\tau \right]}{\int_{\tau=0}^{+\infty} \lambda_i^\tau d\tau} \\ &= (-\ln \lambda_i) \mathbb{E} \left[\int_{\tau=0}^{+\infty} \text{GS}_i(\tau) \lambda_i^\tau d\tau \right]. \end{aligned} \quad (3)$$

As we have seen in Section 3.1, $\lambda_i \in (0, 1)$ is the discount factor of player i for gains. We use $\mathbb{E}[\cdot]$ to denote the expected value of the expression \cdot . The expectation in equation (3) is with respect to the stochastic process $\text{GS}_i(t)$, which is determined by both players' strategies.

Scaling by the total value of the resource makes it possible to interpret the gain as the exponentially weighted “average” time that the player is in control of the resource, and enables comparing gains and costs of players for different discount rates. Without the scaling factor, players with discount rates closer to one would usually have higher gains than players with discount rates closer to zero. With the scaling factor, gains are always values between zero and one.

Note that if $\lambda_D = \lambda_A$, then the sum of the players' gains is one. To see this, first observe that for every real number τ , we have $\text{GS}_i(\tau) + \text{GS}_j(\tau) = 1$. If we also have $\lambda_i = \lambda_j$, then adding $G_i + G_j$ using Equation 3 allows us to do the sum inside the integral, in which case it simplifies to 1.

3.5. Player costs

When players perform a move, this comes at a fixed *instantaneous cost* of $c_i > 0$. Costs are exponentially discounted just like gains, but we allow players to discount

costs using a discount factor Λ_i that is different from λ_i . Instantaneous costs are unit-less and normalized with respect to the value of the resource: when keeping the rate at which the resource generates value fixed, an instantaneous cost of one corresponds to the value of being in control of the resource for one time unit. As an example: if a player values control of the resource twice as high as another player and can take control of the resource at the same nominal cost, her value for c_i will only be half that of the other player.

Player i 's (discounted) *cost* is equal to the sum of all discounted costs made by i , normalized with respect to the total (discounted) value of the resource:

$$C_i = \frac{\mathbb{E} \left[\sum_{\tau \in t_i} c_i \Lambda_i^\tau \right]}{\int_{\tau=0}^{+\infty} \lambda_i^\tau d\tau} = (-\ln \lambda_i) \mathbb{E} \left[\sum_{\tau \in t_i} c_i \Lambda_i^\tau \right]. \quad (4)$$

The expectation is taken over player i 's strategy. Scaling again enables us to compare costs of players for different discount rates λ_i . The absolute value of the cost can also be easily interpreted as a fraction of the maximally achievable gain. Executing a strategy with a cost greater than one can never be better than dropping out, so the cost of any reasonable strategy lies between zero and one.

3.6. Player utilities

Player i 's *utility* is simply the difference between her gain and her cost:

$$u_i = G_i - C_i. \quad (5)$$

In the literature related to FlipIt, utility is also often referred to as *benefit*.

Note that the scaling of gains and costs as explained in the previous subsections corresponds to an affine transformation of utility values. Scaling therefore has no impact on the behavior of rational players.

3.7. Restricted strategies

The description of player strategies as distributions over time sequences as in Section 3.2 is useful for defining the game state, but does not allow us to easily characterize good strategies, or to more generally conduct a fruitful analysis. Certain subsets of the general strategy space have an elegant description and warrant special attention because of their real-world relevance. This section introduces the two strategy spaces that we will discuss in this paper: the exponential strategies, and the periodic strategies.

3.7.1. Exponential strategies. An *exponential strategy* is characterized by having its flip inter-arrival times (the time between flips) drawn from the same exponential distribution. Exponential distributions have a probability density function (PDF)

$$f(t) = \eta e^{-\eta t}.$$

We refer to η as the flip rate, move rate or *play rate* of the exponential strategy. The expected time between two moves is equal to $1/\eta$. An exponential strategy is fully characterized by its play rate.

Exponential strategies are of special interest for two reasons. Firstly, they have an elegant mathematical description. When specifying a strategy from a continuum of choices, it is essential for such a strategy to have a short specification. Exponential strategies are a good example of this, being both mathematically rich and easy to specify.

Secondly, the memorylessness property of the exponential strategy makes exponential strategies robust with respect to information leakage. The stochastic process corresponding to an exponential strategy is completely independent of time, which makes certain information about a player who uses this strategy worthless. Specifically, information that says when player i last moved or has moved in the past does not allow player j to improve her response. Exponential play is therefore, in a sense, a robust strategy and a reasonable choice when playing against players who might receive information of this kind.

3.7.2. Periodic strategies with random phase. A strategy is a *periodic strategy* iff the time between consecutive moves is constant. The time between moves is the *period* of the strategy and we denote it by δ . The inverse of the period, $1/\delta$, is the strategy's *play rate* and we denote it by α . In the context of periodic strategies, we refer to the time of the first flip, $t_{i,0}$, as the *phase* and denote it as φ_i . *Periodic strategies with random phase* are those periodic strategies whose phase is drawn uniformly random from the positive values smaller than its period. A periodic strategy with random phase is fully characterized by the single real number δ . Formally,

$$t_{i,0} = \varphi_i \sim U[0, \delta] \text{ and} \\ t_{i,n+1} - t_{i,n} = \Delta t_{i,n} = \delta \text{ for all } n \geq 0,$$

where $U[0, \delta]$ denotes the uniform distribution between 0 and δ .

The strategies played by both the defender and attacker in Figure 1 are periodic with $\varphi_D = 1.5$, $\delta_D = 3$, $\varphi_A = 3$ and $\delta_A = 4$.

Periodic strategies have received a lot of attention in the literature [12], [19], [24], [27], [28], [31]. They are of special interest for several main reasons. The first one is that decision-makers implement them in real-world systems, an example being Microsoft's "Patch Tuesday". We refer to Farhang and Grossklags [24] for more examples. Another reason is that when moves are stealthy, periodic strategies tend to be good performers. In fact, van Dijk, Juels, Oprea, *et al.* [12] showed that in the absence of discounting and as long as moves are completely stealthy, the class of periodic strategies strictly dominates the class of non-arithmetic renewal strategies. Specifically, periodic strategies with period δ strictly dominate all non-arithmetic renewal strategies with the same average inter-arrival time, when playing against opponents that play non-arithmetic renewal strategies or periodic strategies. Finally, similar to exponential strategies, periodic strategies also exhibit an elegant behavior with only a short description.

4. Model analysis

Our analysis of the model proceeds as follows. First, in Subsection 4.1, we derive structural properties of the

function describing the player gains. This allows us to exhibit a number of symmetries that will simplify our proofs and make the subsequent results easier to state. Then, in Subsection 4.2, we use these symmetries to derive closed-form solutions for the function describing player utilities. Using this formalism, in Subsection 4.3, we investigate how player utilities change with respect to player actions, by deriving key properties of the utility functions' partial derivatives. In Subsection 4.4, we use these incentive properties to characterize the set of strategic best responses for each player, under both strategic regimes. Finally, in Subsection 4.5, we characterize all Nash equilibrium configurations for the game.

4.1. Player gains

We begin our analysis with structural investigation of the player gain function G_i , whose definition in terms of the discount factor λ_i and game state GS is given by Equation (3).

Let us partition time in two intervals: the time before the first flip, $[0, t_0)$, and the time after the first flip, $[t_0, +\infty)$. In the case where neither player ever flips, the first interval constitutes all of time and the second interval is empty. We can split up the total expected gain as the sum of the expected gains over both intervals.

Defender advantage. Over the period $[0, t_0)$, the defender owns the resource simply because she starts out in control. The expected gain of the attacker over $[0, t_0)$ is zero. We refer to the expected gain that the defender receives over $[0, t_0)$ as the *defender advantage* and denote it as G_D^+ :

$$G_D^+ = (-\ln \lambda_D) \mathbb{E} \left[\int_{\tau=0}^{t_0} \lambda_D^\tau d\tau \right] = 1 - \mathbb{E} \left[\lambda_D^{t_0} \right]. \quad (6)$$

In Figure 1, the defender advantage is drawn in a darker shade of blue. Note that without discounting (that is, for $\lambda_D \rightarrow 1$), the defender advantage is zero as long as at some point, at least one player has moved with nonzero probability (the probability of $t_0 < t$ goes to one for $t \rightarrow +\infty$). If neither player ever moves, the defender advantage is always equal to one.

Anonymous gain. Over the course of the interval $[t_0, +\infty)$, GS_i is independent of the identity of i (whether i is D or A). The gain depends only on the player strategies. We refer to the expected gain that player i accrues over $[t_0, +\infty)$ as player i 's *anonymous gain*, denoted by

$$\bar{G}_i = \mathbb{E} \left[\int_{\tau=t_0}^{+\infty} GS_i(\tau) \lambda_i^\tau d\tau \right]. \quad (7)$$

We can express attacker and defender gains in terms of the anonymous gain function and the defender advantage. The defender obtains the defender advantage in addition to anonymous gain, whereas the attacker only receives anonymous gain:

$$G_D = G_D^+ + \bar{G}_D \quad (8)$$

$$G_A = \bar{G}_A. \quad (9)$$

The fact that if the discount rates are the same, the defender and attacker gain sum to one also allows an alternative expression of the defender advantage and defender gain in terms of just the anonymous gain function:

If we denote the anonymous gain that player i would get if her discount rate were λ by $\bar{G}_{i|\lambda}$, we may write

$$G_D^+ = 1 - \bar{G}_D - \bar{G}_{A|\lambda_D} \quad (10)$$

$$G_D = 1 - \bar{G}_{A|\lambda_D}, \quad (11)$$

In other words, the function $\bar{G}_{i|\lambda}$ fully determines defender and attacker gains by Equations (9) and (11).

4.2. Player utilities

This subsection lists closed-form expressions for player utilities when the attacker and defender both play an exponential strategy (Sec. 4.2.1) as well as when they both play periodically with random phase (Sec. 4.2.2).

4.2.1. Utilities for exponential play. We begin by considering the scenario in which the defender plays exponentially with rate parameter η_D and the attacker plays exponentially with rate parameter η_A .

Lemma 1 (Costs for exponential play). *To player i , the cost of executing the exponential strategy with move rate η_i is:*

$$C_i = \frac{\ln \lambda_i}{\ln \Lambda_i} c_i \eta_i, \quad (12)$$

where Λ_i denotes player i 's discount factor for costs, λ_i denotes player i 's discount factor for gains and c_i denotes player i 's immediate move cost.

Proof. We can compute the total discounted cost of flips by player i directly as:

$$C_i = -(\ln \lambda_i) c_i \int_{\tau=0}^{+\infty} \Lambda_i^\tau \eta_i d\tau = \frac{\ln \lambda_i}{\ln \Lambda_i} c_i \eta_i,$$

where c_i is the (normalized) cost of a flip by player i . \square

If a player discounts gains and costs at the same rate ($\lambda_i = \Lambda_i$), the formula for the cost is the same when discounting as when not discounting. Notice that having different values for λ_i and Λ_i has the same effect as increasing the immediate flip cost (if gains are discounted faster than costs) or decreasing the immediate flip cost (if costs are discounted faster than gains) in a particular, well-defined way.

Lemma 2 (Anonymous gain for exponential play). *In a discounted game of timing where both the defender and the attacker are playing an exponential strategy, the anonymous gain function is given by:*

$$\bar{G}_i = \frac{\eta_i}{\eta_i + \eta_j - \ln \lambda_i}, \quad (13)$$

where η_i denotes player i 's move rate and λ_i denotes player i 's discount factor for gains.

Proof outline. Since we are deriving the anonymous gain function, we are interested only in the time interval starting at the first move: $[t_0, +\infty)$. At any time in this

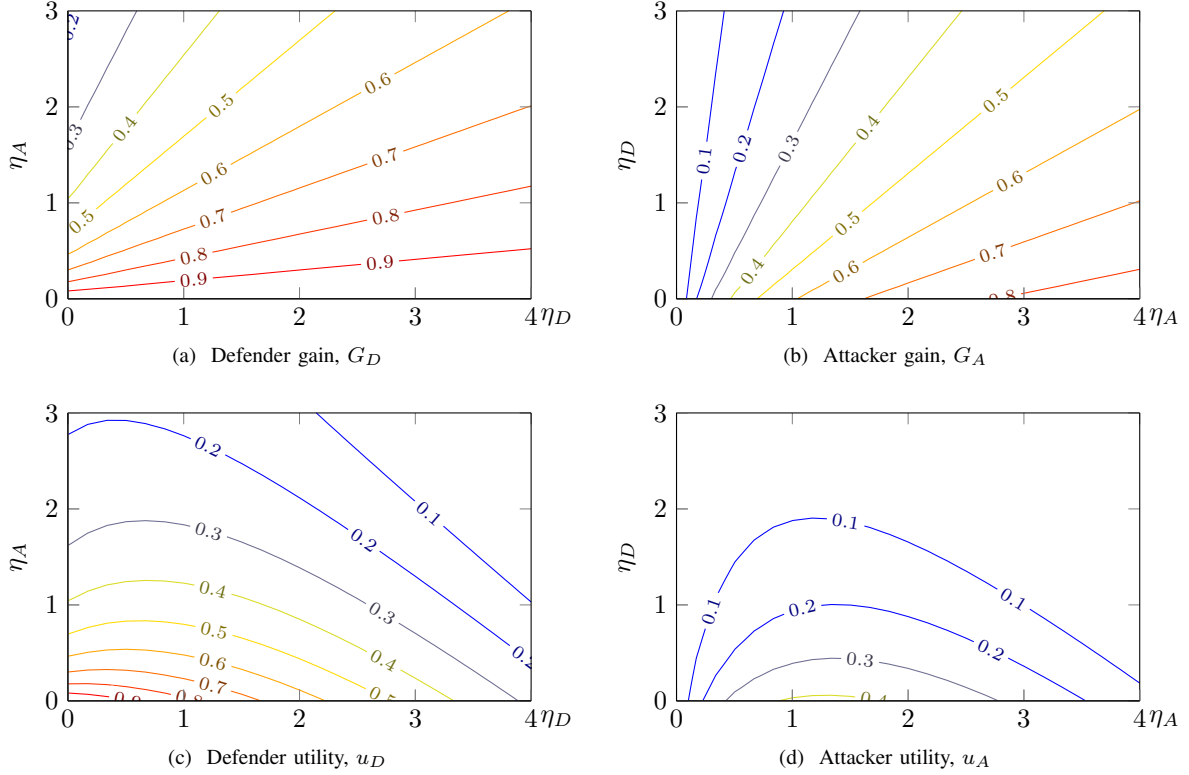


Figure 2. Contour lines of player gains and utilities for exponential play and $\lambda_D = \lambda_A = \Lambda_D = \Lambda_A = 0.5$ and $c_D = c_A = 0.18$.

interval, the memorylessness property of the exponential distribution implies that the probability p_i of player i being the last player to move remains constant over time. We can state p_i in terms of η_i and η_j . Multiplying p_i by the gain player i would obtain over the interval $[t_0, +\infty)$, if she were in control of the resource the entire time, yields the anonymous gain function.

The full derivation is in Appendix A. \square

The players' utility functions are easily derived from the cost function and the anonymous gain function.

Theorem 1 (Utilities for exponential play). *In a discounted game of timing where both the defender and the attacker are playing an exponential strategy, player utilities are given by:*

$$u_D = 1 - \frac{\eta_A}{\eta_A + \eta_D - \ln \lambda_D} - \frac{\ln \lambda_D}{\ln \Lambda_D} c_D \quad (14)$$

$$u_A = \frac{\eta_A}{\eta_A + \eta_D - \ln \lambda_A} - \frac{\ln \lambda_A}{\ln \Lambda_A} c_A, \quad (15)$$

where η_i denotes player i 's move rate, λ_i denotes player i 's discount factor for gains, Λ_i denotes player i 's discount factor for costs and c_i denotes player i 's immediate move cost.

Proof. This follows directly from Lemmas 1 and 2, the definition of utility as the difference between gains and costs (see Equation (5)), and Equations (9) and (11) that express the utility of the attacker and the defender in terms of the anonymous gain function. \square

Figure 2 illustrates the gains and utilities of defender and attacker for exponential play. We see that generally,

increasing the play rate η_i causes player i 's gain to rise. Increases in η_j cause player i 's gain to decrease. The difference between G_D and G_A illustrates that the defender advantage can be very significant, at least for fairly high discount rates like in the figure. The defender always achieves high benefit for attacker move rates that are close to zero; while the attacker always has to move at significant rates if she wants to obtain significant gain, even if the defender does not move at all.

Notice that for $\lambda_i \rightarrow 1$ the formulae for defender and attacker gains become equal to:

$$\lim_{\lambda_D \rightarrow 1} G_D = 1 - \lim_{\lambda_D \rightarrow 1} \bar{G}_{A|\lambda_D} = \frac{\eta_D}{\eta_D + \eta_A}$$

$$\lim_{\lambda_A \rightarrow 1} G_A = \lim_{\lambda_A \rightarrow 1} \bar{G}_{A|\lambda_A} = \frac{\eta_A}{\eta_D + \eta_A}.$$

These equations confirm the results presented in van Dijk, Juels, Oprea, *et al.* [12] for undiscounted games of timing.

4.2.2. Utilities for periodic play. We next consider the scenario in which both players play a periodic strategy with random phase.

Lemma 3 (Costs for periodic play). *To player i , the cost of executing the periodic strategy with period δ_i is:*

$$C_i = \frac{\ln \lambda_i c_i}{\ln \Lambda_i \delta_i}, \quad (16)$$

where Λ_i denotes player i 's discount factor for costs, λ_i denotes player i 's discount factor for gains and c_i denotes player i 's immediate move cost.

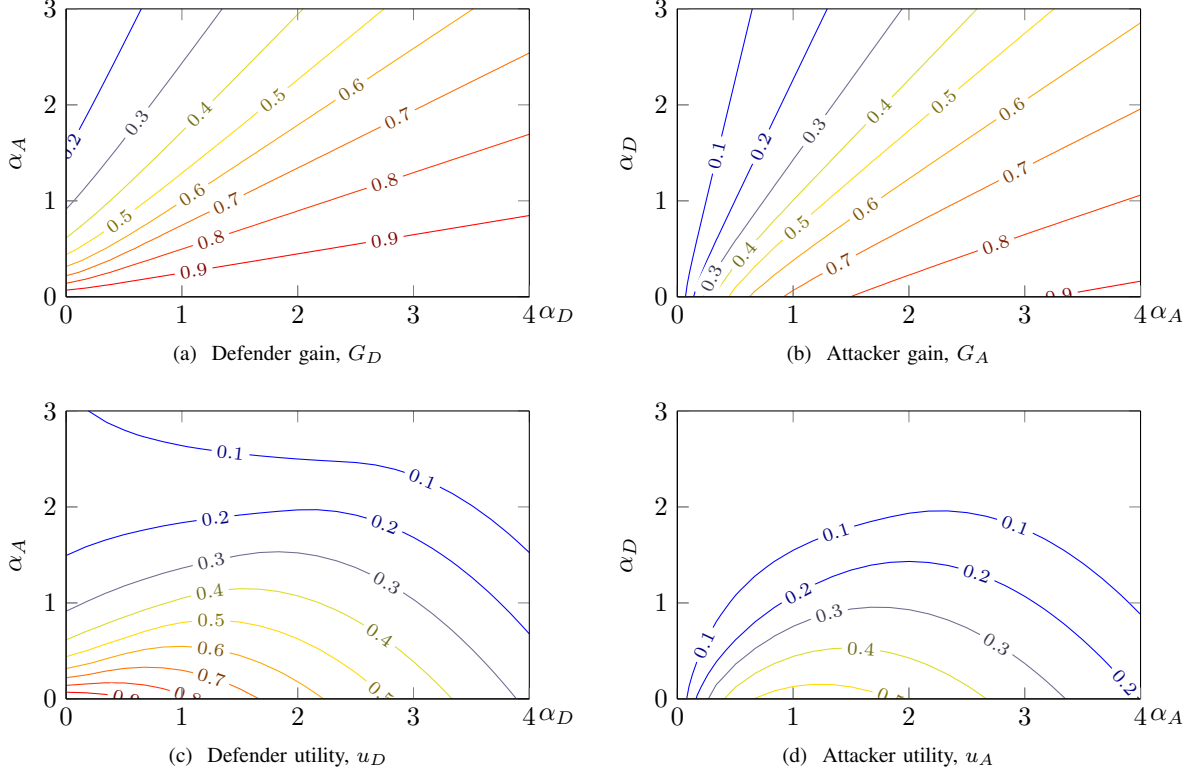


Figure 3. Contour lines of player gains and utilities for periodic play and $\lambda_D = \lambda_A = \Lambda_D = \Lambda_A = 0.5$ and $c_D = c_A = 0.18$.

Proof. The total cost of flips for a player i , who plays periodically with period δ_i and phase φ_i , is given by:

$$C_i|_{\varphi_i} = -\ln \lambda_i \sum_{i=0}^{\infty} c_i \Lambda^{\varphi_i + n\delta_i} = -\ln \lambda_i \frac{c_i \Lambda^{\varphi_i}}{1 - \Lambda^{\delta_i}}.$$

The total cost of moves by player i is then the expected value of $C_i|_{\varphi_i}$, taking expectations over the phases φ_i in player i 's strategy. Phases are drawn uniformly from the interval $[0, \delta_i)$, yielding:

$$C_i = \int_0^{\delta_i} \frac{1}{\delta_i} C_i|_{\varphi_i} d\varphi_i = -\ln \lambda_i \frac{c_i}{\delta_i} \frac{1}{1 - \Lambda^{\delta_i}} \int_0^{\delta_i} \Lambda^{\varphi_i} d\varphi_i,$$

which is equal to Equation (16). \square

Note that the formula for the cost is the same as for the exponential strategy with the same move rate. We can make the same remark about the impact of the discount factor for gains (λ_i) and the discount factor for costs (Λ_i). Since using different discount rates has the same effect as changing the immediate flip cost for periodic as well as exponential play, we will assume that gains and costs are discounted at the same rate in the remainder of this paper.

Lemma 4 (Anonymous gain for periodic play). *In a discounted game of timing where both the defender and the attacker are playing a periodic strategy, the anonymous gain function is given by:*

$$\bar{G}_i = \begin{cases} \bar{G}_f & \text{if } \delta_i \leq \delta_j \text{ and} \\ \bar{G}_s & \text{otherwise,} \end{cases} \quad (17)$$

where subscripts f and s refer to “fast” and “slow”, implying $\delta_f \leq \delta_s$, and

$$\bar{G}_f = \frac{-1}{\delta_i \delta_j \ln \lambda_i} \left(\delta_j - \frac{\lambda_i^{\delta_j} - 1}{\ln \lambda_i} - (\delta_j - \delta_i) \lambda_i^{\delta_i} \right) \quad (18)$$

$$\bar{G}_s = \frac{-1}{\delta_i \delta_j \ln \lambda_i} \left(\delta_j - \frac{\lambda_i^{\delta_j} - 1}{\ln \lambda_i} \right), \quad (19)$$

and where δ_i denotes player i 's move rate and λ_i denotes player i 's discount factor for gains.

Proof outline. Players only receive anonymous gain after their moves. We can divide time in intervals that are defined by a player's moves and express the total anonymous gain of a player as the sum of the gains she obtains over the course of those intervals. By linearity of expectation, a player's expected gain equals the sum of expected gains between moves.

Consider player i and define interval I_n as the time between her n th and $(n+1)$ th move. We make the following observations:

- At the beginning of the interval, player i is always in control of the resource. If player j moves during I_n , then player j takes control of the resource starting from her first move until the end of the interval. If player j does not move during the interval, then player i retains control of the resource for the entire interval.
- Depending on whether player i is the faster or the slower player, player j will always or sometimes move over the course of I_n . If player j moves faster, she flips at least once between any two flips

by the slower player. If player j moves slower, she moves once during I_n with probability δ_i/δ_j .

- If player j moves, she moves at a uniformly random time between the start of the interval and the point in time that is $\min\{\delta_i, \delta_j\}$ later.

We formalize these observations and obtain expressions for player i 's expected gain over I_n given a phase φ_i . We then take expectations over φ_i to obtain the expected gain over the course of I_n . The sum of the expected gains over all intervals forms a geometric series, yielding Equations (18) and (19).

The full derivation is in Appendix A. \square

Notice that for $\lambda_i \rightarrow 1$, the formulae for the anonymous player gains simplify significantly:

$$\lim_{\lambda_f \rightarrow 1} \bar{G}_f = 1 - \frac{\delta_i}{2\delta_j}$$

$$\lim_{\lambda_s \rightarrow 1} \bar{G}_s = \frac{\delta_j}{2\delta_i}.$$

These equations confirm the results presented in van Dijk, Juels, Oprea, *et al.* [12].

Theorem 2 (Utilities for periodic play). *In a discounted game of timing where both the defender and the attacker are playing a periodic strategy, player utilities are given by:*

$$u_D = 1 - \bar{G}_{A|\lambda_D} - \frac{\ln \lambda_D}{\ln \Lambda_D} c_D \quad (20)$$

$$u_A = \bar{G}_A - \frac{\ln \lambda_A}{\ln \Lambda_A} c_A, \quad (21)$$

where $\bar{G}_{i|\lambda}$ is the anonymous gain function for periodic play (Lemma 4), λ_i is player i 's discount factor for gains, Λ_i is player i 's discount factor for costs and c_i is player i 's instantaneous move cost.

Proof. This follows directly from Lemmas 3 and 4, the definition of utility as the difference between gains and costs (see Equation (5)), and Equations (9) and (11) that express the utility of the attacker and the defender in terms of the anonymous gain function. \square

Figure 3 illustrates the gains and utilities of defenders and attackers for periodic play. For similar attack rates, the periodic strategies allow the attacker to obtain the resource significantly faster than the exponential strategies. Consequently, the defender advantage is generally smaller for periodic play.

4.3. Player incentives

This subsection provides several results which describe how player utilities change with respect to their actions. We derive formulae for and properties of the *player incentives*, which we define as the partial derivatives of player utilities to their play rates. We consider first exponential play and then move on to periodic play.

The results of this section will prove useful in our discussion of best responses. Roots of the player incentive function are locally optimal player strategies. Best responses, which are the globally optimal player strategies within players' strategy spaces, must therefore be at such roots.

4.3.1. Player incentives for exponential play. Here we suppose that the attacker and defender both play exponential strategies with choice parameters η_A and η_D respectively.

Lemma 5. *Players' incentives are strictly decreasing in their play rates.*

Proof. The partial derivatives of player utilities to their play rate are given by:

$$\frac{\partial u_D}{\partial \eta_D} = \frac{\eta_A}{(\eta_D + \eta_A - \ln \lambda_A)^2} - c_D \quad (22)$$

$$\frac{\partial u_A}{\partial \eta_A} = \frac{\eta_D - \ln \lambda_A}{(\eta_D + \eta_A - \ln \lambda_A)^2} - c_A. \quad (23)$$

Clearly, $\frac{\partial u_i}{\partial \eta_i}$ is strictly decreasing in η_i , since η_i is only in its denominator and play rates cannot be negative. \square

4.3.2. Player incentives for periodic play. Here we suppose that the defender and attacker both play periodic strategies with play rates α_D and α_A respectively. Expressions for player incentives and proofs of all of the lemmas in this section can be found in Appendix B.

In the periodic strategy regime, the formula expressing a player's utility depends on whether she is the faster or the slower player. Nevertheless, utility functions are still continuous as a function of the players' play rates. The same turns out to be true for the player incentives.

Lemma 6. *Players' incentives are finite and incentive functions are continuous in all their arguments, including the players' play rates.*

Although players' incentives are not strictly decreasing in their play rates like incentives for exponential play, we can show that they are also non-increasing.

Lemma 7. *Players' incentives are independent of their play rate if they are the slower moving player, and strictly decreasing in their play rate otherwise.*

Given an opponent's play rate, a player's incentive for any play rate of hers is therefore upper bounded by her incentive when not playing. We will refer to this incentive as her *base incentive* and to $\frac{\partial u_i}{\partial \alpha_i} \Big|_{\alpha_i \leq \alpha_j}$ as player i 's base incentive function. We refer to the α_j for which player i 's base incentive is zero as the *roots* of player i 's base incentive function.

Figure 4 illustrates the defender's base incentive for different values of λ_D and indicates the values of α_D for which her incentive is negative if $\lambda_D = 0.6$ (—) by green and red areas.

For periodic play, we are also interested in the change in incentive with respect to the opposing player's play rate. This change depends on whether the defender is the faster or the slower moving player.

Lemma 8. *If the defender is the faster moving player ($\alpha_D > \alpha_A$), then*

- *the defender's incentive is strictly increasing in α_A , and*
- *the attacker's (base) incentive is strictly decreasing in α_D .*




Figure 4. Plot of $\frac{\partial G_D}{\partial \alpha_D} \Big|_{\alpha_D \leq \alpha_A}$ and $\frac{d\alpha_D}{d\alpha_A}$ for different λ_D .

Lemma 9. *If the defender is the slower moving player ($\alpha_D < \alpha_A$), then*

- *the defender's (base) incentive is first strictly increasing, then strictly decreasing in α_A , and*
- *the attacker's incentive is decreasing, independent of or increasing in α_D , depending on α_A .*

We can make the following statements as corollaries of Lemmas 8 and 9.

Corollary 1. *The defender's base incentive function has zero, one or two roots. It is positive only for attacker play rates that lie between these roots.*

Corollary 2. *The attacker's base incentive function has zero roots or one root. It is positive only for defender play rates that lie between zero and its root.*

4.4. Best responses

Building on the results of the previous subsection, this subsection characterizes the best response strategies for the attacker and defender. We begin with a discussion of non-participatory responses and characterize when they are optimal. These results apply equally to both strategy regimes. We then characterize participatory best responses for the exponential strategy regime. Finally we characterize participatory best responses for the periodic strategy regime.

4.4.1. Best responses to a non-participating player.

Here we derive best strategies for attackers and defenders in response to a non-participating opponent.

We start by noting that while his opponent's absence makes the attacker likely to play, the opposite is true for the defender.

Lemma 10. *The unique best response of the defender to a non-participating attacker is always not to play.*

Proof. Suppose that the attacker does not play. If the defender were to play at any strictly positive rate, she would incur a strictly positive cost, without receiving any increase in gain. It is therefore a strictly best response for the defender not to play. \square

Lemma 11. *If the attacker's best response to a non-participating defender is not to play, not playing is a best response of his to any play rate by the defender.*

Proof. If the defender does not participate, then the first flip by the attacker results in full control of the resource from that flip until the end of time. Supposing that the attacker evaluates this first flip not to be a worthwhile investment, then it is still not a worthwhile investment when the defender is playing. In that case, the cost to the attacker is the same as before, while the total value that can be accumulated is still bounded by the total value that can be accumulated after his first flip. Therefore, not playing is still a best response for the attacker if the defender is participating. \square

Further, we note that not participating being a best response is related to both a player's immediate move cost and her discount rate.

Lemma 12. *If player i 's move cost is high and/or her valuation of future gains is low such that*

$$c_i \geq \frac{-1}{\ln \lambda_i}, \quad (24)$$

then not moving is a best response. If the inequality is strict, then there cannot be other best responses.

Proof. A flip at any time t costs the player $(-\ln \lambda_i)\lambda_i^t c_i$. Irrespective of player i 's strategy space, the upper bound on the gain that can result because of this flip is

$$(-\ln \lambda_i) \int_{\tau=t}^{+\infty} \lambda_i^\tau d\tau = \lambda_i^t.$$

Since the cost is greater than or equal to the potentially-achievable gain, not participating is a best response. In the case of a strict inequality, the cost is strictly greater than the potentially-achievable gain and not participating is the only best response. \square

Lemma 13. *For periodic or exponential play, not playing being a best response for the attacker implies*

$$c_A \geq \frac{-1}{\ln \lambda_A}.$$

Proof. A flip at time t costs $(-\ln \lambda_A)\lambda_A^t c_A$. We will show that by playing very infrequently, the total increase in gain that can be attributed to every flip reaches λ_A^t , where the gain attributed to a flip is the gain that a player loses if he does not perform that flip. Therefore, not playing being a best response implies $c_A \geq \frac{-1}{\ln \lambda_A}$.

Formally, for any duration Δ and probability p , there exists a play rate such that the probability of two flips happening within Δ of each other is smaller than p . By taking $\Delta \rightarrow +\infty$ and $p \rightarrow 0$, a flip at time t causes the attacker's gain to increase by a value that approaches λ_A^t . \square

Note that Lemma 12 applies independent of any restrictions on player strategies. Lemma 13 similarly applies to many general strategic environments, including the environment with no restrictions on player strategies. A restrictive space that includes either all exponential or all periodic strategies is more than sufficient. The Lemma

requires roughly that attacker be able to play arbitrarily infrequently.

While outside the scope of this work, there do exist more restrictive conditions under which the lemma would not apply, such as variations requiring only periodic strategies with a fixed maximum period.

4.4.2. Best responses for exponential play. Here we consider the regime where both players are playing an exponential strategy and at least one player is moving at a non-zero rate. Our first result says that in this regime, each player always has a unique best response to the action of the other player.

Corollary 3. *A player's best response to any play rate by the opposing player is always single-valued, and situated at the (positive) root of her incentive function. If there is no such root, her best response is to not move.*

Proof. This follows from Lemma 5: since players' incentives are strictly decreasing in their own play rates, the local optima of u_D and u_A are global optima. \square

The roots of the partial derivatives of players' utilities for exponential play to their play rates (Equations (22) and (23)) have closed-form solutions and allow a fairly simple analytic description of the best-response functions for exponential play. We denote the best response of player i to the strategy η_j by $BR_i(\eta_j)$.

Lemma 14 (Best-response functions for exponential play). *The attacker's and defender's best-response functions for discounted exponential play are as follows:*

$$BR_D(\eta_A) = \max \left\{ 0, -\eta_A + \ln \lambda_D + \sqrt{\frac{\eta_A}{c_D}} \right\} \text{ and} \quad (25)$$

$$BR_A(\eta_D) = \max \left\{ 0, -\eta_D + \ln \lambda_A + \sqrt{\frac{\eta_D - \ln \lambda_A}{c_A}} \right\}. \quad (26)$$


Proof. The best-response function follows directly from Corollary 3. Solving Equation (22) for zeroes yields two roots, one of which is always negative and the other one equal to:

$$-\eta_A + \ln \lambda_D + \sqrt{\frac{\eta_A}{c_D}}.$$

Solving Equation (23) for roots yields the attacker's best-response function. \square


Figure 5 shows an example of best-response curves for parameters that yield three Nash equilibria. By inspecting Equation (26) we can see that a decrease of λ_A simply causes the attacker's best-response curve to shift to the left. Looking at Figure 5, we see that for large enough λ_A the attacker's best response first increases and then decreases for increasing η_D . Decreasing λ_A shifts the red curve to the left; once λ_A becomes small enough, the attacker's best-response curve becomes strictly decreasing in η_D . For the defender's best response, decreasing λ_D increases the length of the vertical line at the bottom of her best-response curve, increasing the minimum attack rate that incites the defender to defend her resource.

$$c_A = 0.75, \lambda_A = 0.995, c_D = 0.16, \lambda_D = 0.42$$



(a) On a small scale

$$c_A = 0.75, \lambda_A = 0.995, c_D = 0.16, \lambda_D = 0.42$$



(b) Zoomed in on the origin

Figure 5. Example of a set of attacker and defender best-response curves for exponential play and parameters that yield three Nash equilibria (●).

4.4.3. Best responses for periodic play. Here, we consider the regime in which attacker and defender each play a periodic strategy with play rates α_A and α_D respectively.

We know player i 's best response must always be a root of the partial derivative of her benefit to her play rate. From the properties stated in Lemma 7, we can characterize the best-response function in terms of player incentives.

Corollary 4. *Player i 's best response to an opponent play rate $\bar{\alpha}_j$ can be characterized in terms of her base incentive as follows.*

- *If her base incentive is strictly negative, then her unique best response is not to play.*
- *If her base incentive is zero, then moving at any play rate $\alpha_i \in [0, \bar{\alpha}_j]$ is a best response.*
- *If her base incentive is strictly positive, then her unique best response is to play at the rate $\alpha_i > \bar{\alpha}_j$ for which her incentive is zero.*

Proof. Player i 's utility when moving at rate $\bar{\alpha}_i$ against a player moving at rate $\bar{\alpha}_j$ is equal to

$$\int_{\alpha=0}^{\bar{\alpha}_i} \frac{\partial u_i}{\partial \alpha_i} \Big|_{\substack{\alpha_i=\alpha \\ \alpha_j=\bar{\alpha}_j}} d\alpha.$$

The first two implications therefore directly follow from Lemma 7. The third implication also follows from Lemma 7, provided that the player's incentive becomes negative for some play rate. This can easily be verified by inspection of the incentive functions. \square

Corollary 4 and Lemmas 7 to 9 together indicate that the players' best-response curves are shaped like one of the curves in Figures 6 and 7.

Corollary 5. *The attacker's best response to a defender who does not move, $BR_A(0)$, is always single-valued.*

Proof. This follows from Lemma 7, which states that $\frac{\partial u_A}{\partial \alpha_A} \Big|_{\alpha_A \geq \alpha_D}$ (and therefore $\frac{\partial u_A}{\partial \alpha_A} \Big|_{\alpha_D=0}$) is strictly decreasing in α_A . \square

Corollary 6. *$BR_A(0)$ is strictly positive iff $c_A < \frac{-1}{\ln \lambda_A}$.*

Proof. This follows from Corollary 5 and Lemmas 12 and 13, which state the conditions under which not moving is a best response for the attacker. \square

Figure 8 offers insight into the properties of $BR_A(0)$. The function illustrates the attacker's incentive when he is the faster player and the defender does not move. Consistent with Lemma 7, we see that the attacker's incentive is strictly decreasing in α_A , attaining its maximum value of $-1/\log \lambda_A$ for $\alpha_A \rightarrow 0$. Circles (● ● ●) indicate non-zero $BR_A(0)$ at the intersection of the derivative of the gain and the derivative of the cost. We see that higher costs always lead to lower best responses until at some point the best response is to drop out, as is the case for $\lambda_A = 0.3$ (—) and $c_A = 1.0$. Higher or lower λ_A do not always correspond to higher or lower best responses: for $c_A = 0.4$ the best response is faster for $\lambda_A = 0.3$ (—) than for $\lambda_A = 0.6$ (—), but for $c_A = 0.7$ the opposite is true.

There is no closed-form expression for $BR_A(0)$. However, locating its roots numerically can be very fast as the player incentive function is easy to evaluate and well-behaved.

4.5. Nash equilibria

4.5.1. Nash equilibria with non-participating player(s).

In this section we characterize all Nash equilibria in which at least one player is not participating.

Theorem 3. *If the attacker is playing a periodic or exponential strategy, there is a Nash equilibrium in which neither player moves iff*

$$c_A \geq \frac{-1}{\ln \lambda_A}.$$

Proof. This follows from the fact that an exponential or periodic attacker's best response to a non-participating defender is not to play iff $c_A \geq -1/\ln \lambda_A$ (Lemmas 12 and 13), and the fact that the unique best response of the defender to an absent attacker is not to play (regardless of her strategy space) (Lemma 10). \square

Theorem 4. *There is never a Nash equilibrium in which the defender plays, but the attacker does not.*

Proof. This follows immediately from Lemma 10. \square

The two previous theorems continue to hold for more general strategy spaces, as mentioned after the proofs of the relevant lemmas. However, for the remaining equilibria, our characterizations require that both players adhere to the same restricted strategy space, either exponential or periodic.

Theorem 5. *For exponential play, there exists at most one equilibrium in which only the attacker moves at non-zero rates. It exists iff $c_A < -1/\ln \lambda_A$ and*

$$-\eta_A^* + \ln \lambda_D + \sqrt{\frac{\eta_A^*}{c_D}} \leq 0,$$

where

$$\eta_A^* = -\ln \lambda_A + \sqrt{\frac{-\ln \lambda_A}{c_A}}.$$

Proof. This follows from the best response functions for exponential play (Lemma 14). \square

Theorem 6. *For periodic play, there exists at most one equilibrium in which only the attacker moves at non-zero rates. It exists iff $BR_A(0) > 0$ and $BR_A(0)$ is not strictly between the roots of the defender's base incentive function.*

Proof. This follows from the fact that the defender's best response to $BR_A(0)$ is not to play iff her base incentive is negative (Corollary 4) and the fact that this function is only negative for attacker play rates that are not between its roots (Corollary 1). \square

4.5.2. Nash equilibria with exponential strategies.

Now, we turn our attention to the equilibria in which both players move at non-zero rates, beginning with Nash equilibria in the exponential strategy regime.

Theorem 7. *If η_D^* is a strictly positive root of $BR_D(BR_A(\eta_D))$, and if $BR_A(\eta_D^*)$ is strictly positive, then $(\eta_D^*, \eta_A^*) := (\eta_D^*, BR_A(\eta_D^*))$ is a Nash equilibrium in which both players play. There are up to two such η_D^* . There are no other Nash equilibria in which both players play.*

Proof. The fact that (η_D^*, η_A^*) is a Nash equilibrium and that there exist no other equilibria follows from the definition of a Nash equilibrium and the single-valuedness of BR_D and BR_A (Corollary 3). Looking at $BR_D(BR_A(\eta_D))$, which under the assumption that η_D^* and η_A^* are strictly positive is equal to

$$\begin{aligned} \ln \lambda_D - \ln \lambda_A - \sqrt{\frac{\eta_D - \ln \lambda_A}{c_A}} \\ + \sqrt{\frac{\ln \lambda_A - \eta_D + \sqrt{\frac{\eta_D - \ln \lambda_A}{c_A}}}{c_A}}, \end{aligned} \quad (27)$$

we can see that finding the roots boils down to finding the roots of a quadratic equation while keeping track of some additional constraints that ensure the flipping rate is a real number. This equation has up to two positive roots. \square

There is a closed-form solution for the values of η_D^* , i.e. the roots of Equation (27). It is unsightly so we do not list it here. Also note that the choice in Theorem 7 to look for the roots of $BR_D(BR_A(\eta_D))$ is arbitrary; we




Figure 6. Defender best-response curves, defined by λ_D and c_D . Jumps in the curve are situated at the roots of $\frac{\partial u_D}{\partial \alpha_D} \Big|_{\alpha_D \leq \alpha_A}$.

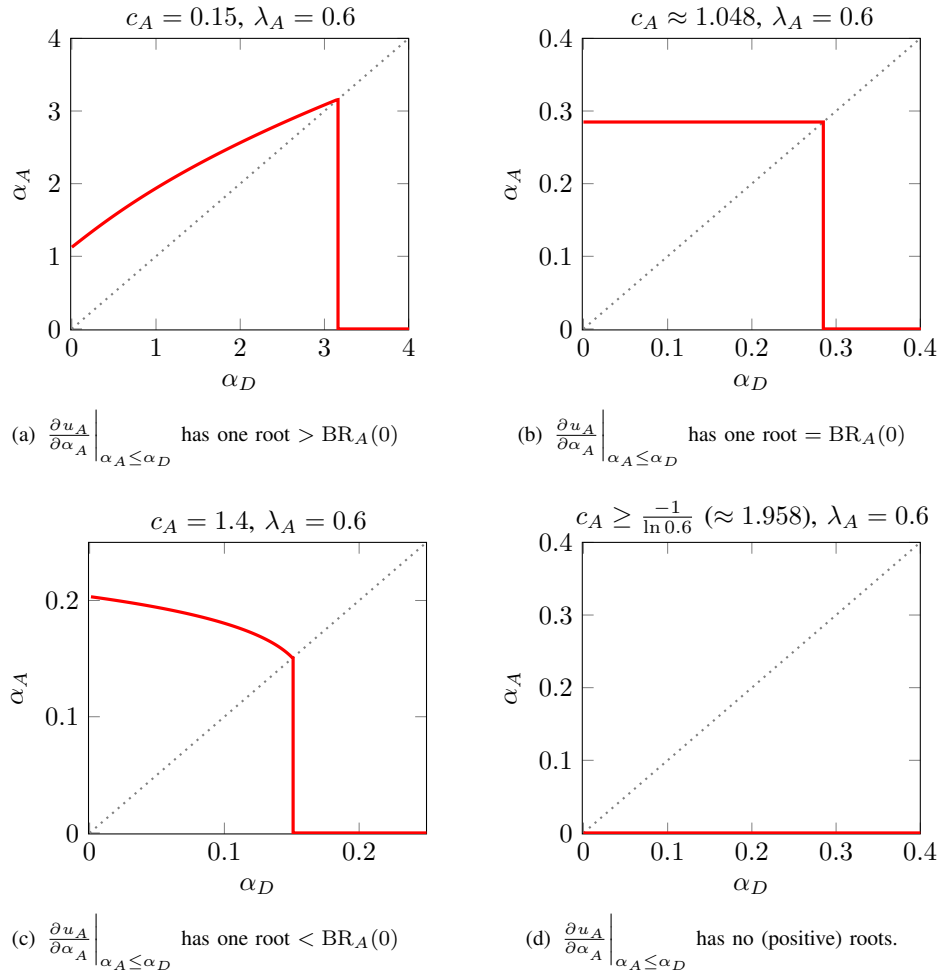


Figure 7. Attacker best-response curves, defined by λ_A and c_A . The jump in the curve is situated at the root of $\frac{\partial u_A}{\partial \alpha_A} \Big|_{\alpha_A \leq \alpha_D}$.




Figure 8. Plot of $\frac{\partial G_A}{\partial \alpha_A} \Big|_{\alpha_D=0}$ and $\frac{dC_A}{d\alpha_A}$ for different values of λ_A and c_A .




Figure 9. Example of a set of attacker and defender best-response curves for periodic play and parameters that yield three Nash equilibria (●).

could also look for the roots η_A^* of $BR_A(BR_D(\eta_A))$ for which $\eta_D^* = BR_D(\eta_A^*)$ is strictly positive. This yields the same result.

4.5.3. Nash equilibria with periodic strategies. Here we address Nash equilibria in the periodic strategy regime.

Figure 9 gives an example of a set of best-response curves for periodic play corresponding to parameters yielding three Nash equilibria: one in which the attacker moves and the defender does not and two where both players move, one where the defender moves faster and one where the attacker moves faster.

Characterizing all equilibria comes down to looking for conditions for which the defender's and the attacker's best response curves intersect. Some possibilities for intersections were discussed in Section 4.5.1.

- Theorem 3 gave the condition under which there is a Nash equilibrium in which neither player moves. In terms of Figure 7, Theorem 3 describes when the attacker's best-response curve looks like Figure 7d. It also shows that there is an intersection at $(0,0)$ iff it looks like Figure 7d, irrespective of what the defender's best-response curve looks like.

- Theorem 6 gave the conditions for an equilibrium where only the attacker moves. In terms of the best-response curves of Figure 6 and Figure 7, this corresponds to the case where $BR_A(0)$ — the leftmost point on the attacker's curve — is located on one of the vertical parts of the defender's best-response curve.

As a step towards characterizing the equilibria in which both players move, observe that moving at a rate that is non-zero but still slower than the other player's play rate α_j is only a best response for very specific values of α_j — those corresponding to the horizontal lines of Figure 6 and the vertical line of Figure 7.

Corollary 7. *In any Nash equilibrium where both players play at non-zero rates, the faster player f plays at a rate α_f that is a root of the slower player's base incentive function. Player s is then indifferent between playing at any rate in $[0, \alpha_f]$.*

Proof. Let (α_i, α_j) be a Nash equilibrium. By Corollary 4, playing at a lower rate than α_f is a best response for s only if she drops out or if player s 's base incentive is zero for α_f . The first possibility is ruled out by the condition that both players move at non-zero rates. \square

Since the roots of the players' base incentive functions are the only play rates that can potentially be part of an equilibrium, all that remains is to characterize the conditions under which one of the slower player rates α_s makes playing α_f a best response for player f . It turns out that these conditions depend on whether player f is the defender or the attacker.

Lemma 15 (Best-response with faster defender). *Let $\bar{\alpha}_D$ be any strictly positive defender play rate. Then the following statements are equivalent:*

- *There exists a unique attacker play rate $\bar{\alpha}_A \leq \bar{\alpha}_D$ to which $\bar{\alpha}_D$ is a best response.*
- *The defender's base incentive function is positive for attack rate $\alpha_A = \bar{\alpha}_D$.*

Proof. If $\bar{\alpha}_D$ is a best response to $\bar{\alpha}_A$ then

$$0 = \frac{\partial u_D}{\partial \alpha_D} \Big|_{\substack{\alpha_D = \bar{\alpha}_D \\ \alpha_A = \bar{\alpha}_A}} \leq \frac{\partial u_D}{\partial \alpha_D} \Big|_{\substack{\alpha_D = \bar{\alpha}_D \\ \alpha_A = \bar{\alpha}_D}} = \frac{\partial u_D}{\partial \alpha_D} \Big|_{\substack{\alpha_D \leq \alpha_A \\ \alpha_A = \bar{\alpha}_D}},$$

that is, the defender's base incentive function is positive for attack rate $\alpha_A = \bar{\alpha}_D$. The first step follows from the definition of a best-response and the second step follows from the fact that the defender's incentive for defender faster play is strictly increasing in the attacker's play rate (Lemma 8).

Now, assume that the attacker's base incentive function is positive for attack rate $\alpha_A = \bar{\alpha}_D$. We then know

$$0 \leq \frac{\partial u_D}{\partial \alpha_D} \Big|_{\substack{\alpha_D \leq \alpha_A \\ \alpha_A = \bar{\alpha}_D}} = \frac{\partial u_D}{\partial \alpha_D} \Big|_{\substack{\alpha_D = \bar{\alpha}_D \\ \alpha_A = \bar{\alpha}_D}},$$

that is, the attacker's incentive is positive when both players move at rate $\bar{\alpha}_D$. We also know that

$$0 \leq \frac{\partial u_D}{\partial \alpha_D} \Big|_{\substack{\alpha_D = 0 \\ \alpha_A = 0}} \leq \frac{\partial u_D}{\partial \alpha_D} \Big|_{\substack{\alpha_D = \bar{\alpha}_D \\ \alpha_A = 0}},$$

that is, for an attack rate of $\alpha_A = 0$, the defender's incentive is strictly negative when she moves at rate $\bar{\alpha}_D$. The first step is due to not playing being the defender's best response to an attacker who is not playing (Lemma 10) and the second step is due to the defender's incentive for faster play being strictly decreasing in α_D (Lemma 8). The defender's incentive when playing at rate $\alpha_D = \bar{\alpha}_D$ is positive for an attack rate of $\alpha_A = \bar{\alpha}_D$ and negative for an attack rate of $\alpha_A = 0$. Continuity and strict monotonicity of her incentive function in α_A (Lemmas 6 and 8) then implies that her incentive is zero for some unique $\bar{\alpha}_A \in [0, \bar{\alpha}_D]$, implying that $\bar{\alpha}_D$ is a best response to this unique $\bar{\alpha}_A$. \square

Lemma 16 (Best-response with faster attacker). *Let $\bar{\alpha}_A$ be any strictly positive attacker play rate. Then the following statements are equivalent:*

- *There exists a defender play rate $\bar{\alpha}_D \leq \bar{\alpha}_A$ to which $\bar{\alpha}_A$ is a best response.*
- *The attacker's base incentive function has a positive root α_D^* , and*

$$\bar{\alpha}_A \in [\min\{\text{BR}_A(0), \alpha_D^*\}, \max\{\text{BR}_A(0), \alpha_D^*\}].$$

If $\alpha_D^ = \text{BR}_A(0)$, then $\bar{\alpha}_D$ can be any value in $[0, \bar{\alpha}_A]$. Otherwise, $\bar{\alpha}_D$ is unique.*

Proof outline. The proof of Lemma 16 is similar to the proof of Lemma 15, but requires taking cases on the incentive for faster play being strictly increasing, independent of or strictly decreasing in the defender's play rate. Which of these cases applies is determined by the relative size of $\text{BR}_A(0)$ to that of α_D^* . Appendix B lists the full proof. \square

The following theorems can be stated as corollaries to Corollary 7 and Lemmas 15 and 16.

Theorem 8 (Equilibrium with faster defender). *If the attacker's base incentive function has a root $\bar{\alpha}_D > 0$ and the defender's base incentive function is positive for attack rate $\alpha_A = \bar{\alpha}_D$, then there exists a unique $\alpha_A^* \leq \bar{\alpha}_D$ for which $(\bar{\alpha}_D, \alpha_A^*)$ is an equilibrium. There are no other equilibria in which both players move at non-zero rates, and where the defender is the faster-moving player.*

Theorem 9 (Equilibria with faster attacker). *If the attacker's base incentive function has a strictly positive root $\bar{\alpha}_D$, then for every strictly positive root α_A^* of the defender's base incentive function, the following implications hold:*

- *If $\text{BR}_A(0) = \bar{\alpha}_D = \alpha_A^*$, then all strategy profiles in $\{(i, \alpha_A^*)\}_{i \leq \alpha_A^*}$ are Nash equilibria.*
- *If $\text{BR}_A(0) < \bar{\alpha}_D$, then $\alpha_A^* \in [\text{BR}_A(0), \bar{\alpha}_D]$ implies that there exists a unique $\bar{\alpha}_D \leq \alpha_A^*$ such that $(\bar{\alpha}_D, \alpha_A^*)$ is a Nash equilibrium.*
- *If $\text{BR}_A(0) > \bar{\alpha}_D$, then $\alpha_A^* \in [\bar{\alpha}_D, \text{BR}_A(0)]$ implies that there exists a unique $\bar{\alpha}_D \leq \alpha_A^*$ such that $(\bar{\alpha}_D, \alpha_A^*)$ is a Nash equilibrium.*

There are no other equilibria in which both players move at non-zero rates, and where the attacker is the faster-moving player.

5. Discussion

This section discusses the results from the previous sections and their impact. Specifically, we look at how discounting affects optimal player behavior (Section 5.1), the challenges associated with selecting one equilibrium out of many (Section 5.2), the possibility of achieving perfect security by raising an attacker's costs (Section 5.3), and some advantages and disadvantages for adopting a periodic strategy in a timing game. (Section 5.4).

5.1. The impact of discounting

For both the defender and the attacker, increasing impatience has the effect of increasing the 'effective' cost of moving, because move costs have to be paid up front whereas gains are accrued over time. Discounting also breaks the symmetry between attacker and defender behavior, because the defender is always the player who starts out in control of the resource. As such, the defender always receives gain when the resource is most valuable even if she does not move.


In contrast, the attacker has to attack at reasonably high rates if he wants to obtain the resource near the beginning of the game, when it is most valuable. This was illustrated by Figures 2 and 3. For equal move costs ($c_A = c_D$) and impatience ($\lambda_A = \lambda_D$), the best response of the defender to a certain attack rate α is therefore always strictly lower than the best response of the attacker to the same rate of play α by the defender.

Figure 10 illustrates the effect of increasing impatience on optimal attacker and defender behavior for periodic play. Discounting always causes the defenders to move more slowly. Both the increase in effective flip cost and the fact that she starts out in control cause the best-response attack rate to decrease when the discount rate increases. She starts playing only at higher attack rates, drops out at lower attack rates and optimal flip rates are lower for all attack rates in between. For attackers, the effect is more subtle. The increase in effective flip cost and the need to obtain the resource while it still has value work against each other. For low attack rates by the defender, decreasing λ_A can increase the best response; this is also illustrated by Figure 8. For some higher attack rate by the defender, the increased flip cost effect is always the stronger effect. The best response then always decreases with increased impatience. The same observations hold for exponential play.

5.2. The equilibrium selection problem

In undiscounted FlipIt, there is exactly one equilibrium solution for periodic and for exponential play. When discounting, there can be up to three equilibria, giving rise to the *equilibrium selection problem* [35]. As an example, Table 2 and Table 3 give the Nash equilibrium payoffs for the equilibria for periodic and exponential play that are visible on Figure 9 and Figure 5.

For periodic play, Table 2 shows that equilibrium #1 is the preferred outcome for the attacker. Equilibrium #2 is the preferred outcome for the defender. Equilibrium #3 is not preferred by either player; in fact, both players prefer equilibrium #2 over equilibrium #3. For exponential play,



(a) Defender best-response curves for $c_D = 0.3$ and $\Lambda_D = \lambda_D$.

(b) Attacker best-response curves for $c_A = 0.3$ and $\Lambda_A = \lambda_A$.

Figure 10. Attacker and defender best-response curves for periodic play for varying values of λ_i .

TABLE 2. NASH EQUILIBRIA AND THE CORRESPONDING ATTACKER AND DEFENDER BENEFITS FOR PERIODIC PLAY, $c_A = 0.18$, $\lambda_A = 0.99$, $c_D = 0.101$ AND $\lambda_D = 0.069$.

#	α_D	α_A	Defender utility	Attacker utility
1	0	0.16	0	0.94
2	0.30	0.92	0.69	0.67
3	2.30	2.52	0.38	0.09

TABLE 3. NASH EQUILIBRIA AND THE CORRESPONDING ATTACKER AND DEFENDER BENEFITS FOR EXPONENTIAL PLAY, $c_A = 0.75$, $\lambda_A = 0.995$, $c_D = 0.16$ AND $\lambda_D = 0.42$.

#	$\eta_D \cdot 100$	$\eta_A \cdot 100$	Defender utility	Attacker utility
1	0	7.67	91.87/100	88.11/100
2	3.54	19.17	81.92/100	68.21/100
3	23.31	32.54	73.45/100	33.34/100

Table 3 shows that the equilibrium in which the defender does not move is preferred by both players.

Notice that especially for the parameters for exponential play, the attacker and the defender can both obtain a good outcome in equilibrium (considering that a benefit of 1 corresponds to the player being in charge of the resource 100% of the time at no cost, a benefit of around 0.9 is very high). This is because competition for the resource is limited, since the players value the resource differently over time ($\lambda_D \ll \lambda_A$).

Although multiple equilibria can exist, most combinations of λ_D , λ_A , c_D and c_A yield just a single equilibrium. Notably, if there are multiple equilibria, there is at least one equilibrium in which the defender does not move. This can be verified visually by looking at the best-response curves of Figure 9; if $BR_A(0)$ does not fall on one of the vertical parts of the defender's best-response curve, the attacker's and the defender's curve only intersect at a single point.

This does not necessarily mean that the possibility for multiple equilibria is generally unimportant. The situation where the defender does not move is quite common in real-world systems. While this could be an equilibrium outcome, there might be other equilibria and it is unlikely to be the optimal outcome from the defender's point of view; it is only a good outcome for both players when λ_A and λ_D are very different.

5.3. Raising costs for perfect security

One of the main results in timing-based security games is that it is possible to disincentivize attacking the resource to such an extent that perfect security is achieved by performing defensive moves so quickly that the attacker drops out. Discounting adds to this the possibility that the resource's defenses (proportional to c_A) are strong enough to dissuade attackers even if the resource is not otherwise defended – giving rise to an equilibrium at $(0, 0)$.

An advantage of such a defense is that it is very strong. Firstly, if there is an equilibrium at $(0, 0)$, it is always the only one. Secondly, not playing is the strictly dominant strategy for the attacker even when considering a wide array of sensible strategies, not just periodic or exponential or even renewal strategies. Thirdly and lastly, the equilibrium at $(0, 0)$ remains even if the attacker were to obtain information during the game. The only caveat is that c_A is not only proportional to the (expected) cost to breach the resource's defenses, but is also inversely proportional to the value that the resource has to the attacker.

In our opinion, the existence of equilibria at $(0, 0)$ is not just a gimmick but an essential property of a model that claims to be representative of real-world security interactions. In practice, most attackers do not attack most systems, e.g. because they do not consider them valuable enough. This situation is captured when discounting: an uninteresting resource would give rise to a high c_A and an equilibrium at $(0, 0)$. Without discounting an attacker will always attack any undefended resource even if doing so is costly and the resource has low value, albeit at low rates.

Even if c_A is not high enough to induce the equilibrium $(0, 0)$, a higher c_A will increase the defender's control over the resource and/or decrease the need for her to perform defensive moves. Discounting limits the total value of the resource and player benefits to finite numbers that allow economic interpretation: what value can we expect to extract from the resource in terms of present-day units of value. This allows the exploration of whether an initial security investment that increases the intrinsic security of the resource (for example, installing a firewall or having an external firm audit a system's security settings) makes economic sense or not – such an investment would increase c_A , which can increase the

expected fraction of time during which the defender is in control of the resource or can decrease the need for her to perform defensive moves. Similarly, an initial investment in a more streamlined security mechanism might make it easier to perform a defensive move, decreasing c_D .

If we only allow the defender to invest, we can optimize her benefit by using her equilibrium benefits as the outcome of an investment level that determines c_A . Considering an initial investment phase for both the defender and the attacker gives rise to a two-stage game, in which players first pick their investment levels, fixing the parameters c_i for the timing game that follows.

5.4. Advantages of periodic play

One of the major theoretical results for games of timing without discounting is that the class of periodic strategies strictly dominates the class of non-arithmetic renewal strategies (see Sect. 3.7.2). This theoretical result has been a major motivator for the interest in periodic strategies. In this section, we discuss this result in the context of discounted games. We conclude that although the periodic strategy still performs well, the result as stated for games without discounting no longer holds and the good performance of the periodic strategy is partly due to an “unfair” advantage it holds over the exponential strategy regarding the timing of the first move.

We start by noting that the periodic strategy still seems to perform well compared to many renewal strategies, among them the exponential strategy. By a derivation similar to the proof of Theorem 2, we find that the discounted gain of a periodic attacker with period δ_A facing an exponential defender with play rate η_D is equal to:

$$G_A = \frac{1 - e^{-\delta_A \eta_D} \lambda_A^{\delta_A}}{\delta_A \eta_D - \delta_A \ln \lambda_A}. \quad (28)$$

Comparison of the periodic attacker and the benefit of the exponential attacker (Theorem 1) shows that the benefit of the periodic attacker is always strictly higher, for any play rate.

Theorem 10. *An attacker with discount rate λ_A who is playing an exponential strategy with play rate η_A that is facing a defender who is playing an exponential strategy with play rate η_D , can strictly increase her benefit by playing a periodic strategy with play rate $\alpha_A = \eta_A$ instead.*

Proof. The gain of the exponential attacker is equal to the anonymous gain given by Theorem 1. The gains for the periodic attacker facing an exponential attacker are given by Equation (28). For equal play rates, the cost of an exponential strategy and a periodic strategy are the same. Proving the theorem comes down to showing that the gain for the periodic attacker is strictly greater than the gain for the exponential attacker:

$$\frac{\alpha_A}{\eta_D + \alpha_A - \ln \lambda_A} < \alpha_A \frac{1 - e^{-\eta_D / \alpha_A} \lambda_A^{1/\alpha_A}}{\eta_D - \ln \lambda_A}.$$

By moving the first term of the right hand side to the left hand side, multiplying both sides by $-(\eta_D - \ln \lambda_A)$ and

substituting $x = \frac{1}{\alpha_A}(\eta_D - \ln \lambda_A)$, the inequality becomes $\frac{1}{1+x} > e^{-x}$, which strictly holds since $x > 0$. \square

We do not, however, think this result indicates that the periodic strategy is really inherently superior to the exponential strategy (or other renewal strategies). The periodic strategy differs from the exponential strategy in that the phase (the time of the first move) is not drawn from the same distribution as the inter-arrival times (see Section 3.7). In fact, for the periodic strategy the average time until the first move is only half the average inter-arrival time. This is an obvious advantage for the attacker, who wants to take control of the resource as fast as possible. If we lift the restriction on the first move for the exponential distribution and allow the attacker to move at $t = 0$, the attacker can essentially “switch positions” with the defender, obtaining the defender advantage at a fixed cost of c_A . Depending on the game parameters, such an “immediate exponential strategy” can again yield higher gains than the periodic strategy of Equation (28). Picking the phase from a uniform distribution $U[0, 2\delta]$ “fixes” the average time until the first move, but it is unclear why such an artificial change — which would decrease the attacker’s benefit — is the right one to make. Why would the attacker play according to this changed periodic strategy if she can increase her gain by playing the unchanged periodic strategy? Changing the distribution of the phase would also mean that the periodic and exponential strategies with the same move rates no longer cost the same to execute.

When discounting, allowing both players to play strategies like “flip at $t = 0$ ” would prevent us from defining a sensible outcome for all strategy combinations; and this motivates our use of randomness for the attacker’s first move. Because it is unclear what is the “right” way to flip periodically or restrict strategies, we chose to work with versions of periodic and exponential strategies that seemed most consistent with those introduced in the original FlipIt paper, and used in numerous extensions. Future work can motivate alternative restrictions on strategy spaces within the context of discounting.

6. Conclusions

The timing of security decisions is an aspect of security policy-making that is generally under-appreciated. The models that do consider the timing of offensive and defensive actions, do not usually consider the impact that the passing of time can have on the valuation of a resource. In this paper, we have extended games of timing to allow for exponential discounting of gains and costs over time. We derived formulae for player gains and benefits for exponential and periodic player strategies and characterized the Nash equilibria for exponential and periodic play.

Discounting allows the interpretation of gains and costs over time in terms of present-day units of value, allowing a more straight-forward interpretation of gains and costs. We have shown that discounting significantly influences optimal player behavior, and that it impacts defenders and attackers in different ways, breaking the symmetry between them. Discounting allows the model to predict a very common equilibrium outcome, in which

the attacker does not attack an uninteresting resource even if the defender does not perform defensive moves. Somewhat surprisingly, we have also seen that discounting can give rise to solutions with multiple equilibria.

In future work, we plan to analyze other, descriptive ways of discounting as well as classes of strategies that can change over time.

7. Acknowledgments

We thank the anonymous reviewers for their constructive comments and feedback. This work was partially supported by the German Institute for Trust and Safety on the Internet (DIVSI).

References

- [1] D. Kushner, "The Real Story of Stuxnet," *IEEE Spectrum*, Feb. 2013.
- [2] D. Barrett, D. Yadron, and D. Paletta, "U.S. Suspects Hackers in China Breached About 4 Million People's Records, Officials Say," *The Wall Street Journal*, Jul. 2015.
- [3] E. Perez and S. Prokupecz, "U.S. Data Hack may be 4 Times Larger than the Government Originally Said," *CNN*, Jul. 2015.
- [4] J. Hirschfeld Davis, "Hacking of Government Computers Exposed 21.5 Million People," *New York Times*, Jul. 2015.
- [5] R. Gallagher, "The Inside Story of How British Spies Hacked Belgium's Largest Telco," *The Intercept*, Dec. 2014.
- [6] R. Price, "TALKTALK HACKED: 4 million customers affected, stock plummeting, 'Russian jihadist hackers' claim responsibility," *Business Insider UK*, Oct. 2015.
- [7] Reuters, "Premiera Blue Cross Says Data Breach Exposed Medical Data," *New York Times*, Apr. 2015.
- [8] S. Farhang and J. Grossklags, "When to invest in security? Empirical evidence and a game-theoretic approach for time-based security," in *16th Workshop on the Economics of Information Security (WEIS)*, Jun. 2017.
- [9] E. Cole, "Preface," in *Advanced Persistent Threat: Understanding the Danger and How to Protect Your Organization*, 2013.
- [10] S. Nadela. (2015). Enterprise Security in a Mobile-first, Cloud-first World, [Online]. Available: <http://news.microsoft.com/security2015/> (visited on 02/19/2019).
- [11] T. Radzik, "Results and problems in games of timing," *Lecture Notes-Monograph Series*, 1996.
- [12] M. van Dijk, A. Juels, A. Oprea, and R. L. Rivest, "FlipIt: The Game of 'Stealthy Takeover'," *Journal of Cryptology*, vol. 26, Oct. 2012.
- [13] D. Smith and H. Gravelle, "The practice of discounting in economic evaluations of healthcare interventions," *International Journal of Technology Assessment in Health Care*, vol. 17, 2001.
- [14] S. Frederick, G. Loewenstein, and T. O'donoghue, "Time discounting and time preference: A critical review," *Journal of Economic Literature*, vol. 40, 2002.
- [15] T. Sheldon, "Discounting in health care decision-making: Time for a change?" *Journal of Public Health*, vol. 14, 1992.
- [16] R. Carson and B. Roth Tran, "Discounting behavior and environmental decisions," *Journal of Neuroscience, Psychology, and Economics*, Nov. 2009.
- [17] M. Manshaei, Q. Zhu, T. Alpcan, T. Başar, and J.-P. Hubaux, "Game Theory Meets Network Security and Privacy," *ACM Computing Surveys (CSUR)*, vol. 45, 2013.
- [18] A. Laszka, M. Felegyhazi, and L. Buttyán, "A Survey of Interdependent Security Games," *ACM Computing Surveys (CSUR)*, vol. 47, 2014.
- [19] K. D. Bowers, M. van Dijk, R. Griffin, A. Juels, A. Oprea, R. L. Rivest, and N. Triandopoulos, "Defending against the Unknown Enemy: Applying FlipIt to System Security," in *Decision and Game Theory for Security*, Nov. 2012.
- [20] A. Laszka, B. Johnson, and J. Grossklags, "Mitigating Covert Compromises," in *Web and Internet Economics*, Dec. 2013.
- [21] A. Laszka, B. Johnson, and J. Grossklags, "Mitigation of Targeted and Non-targeted Covert Attacks as a Timing Game," in *Decision and Game Theory for Security*, Nov. 2013.
- [22] X. Feng, Z. Zheng, P. Hu, D. Cansever, and P. Mohapatra, "Stealthy attacks meets insider threats: A three-player game model," in *MILCOM 2015 - 2015 IEEE Military Communications Conference*, Oct. 2015.
- [23] P. Hu, H. Li, H. Fu, D. Cansever, and P. Mohapatra, "Dynamic defense strategy against advanced persistent threat with insiders," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, Apr. 2015.
- [24] S. Farhang and J. Grossklags, "FlipLeakage: A Game-Theoretic Approach to Protect Against Stealthy Attackers in the Presence of Information Leakage," in *Decision and Game Theory for Security*, Nov. 2016.
- [25] M. Zhang, Z. Zheng, and N. B. Shroff, "A Game Theoretic Model for Defending Against Stealthy Attacks with Limited Resources," in *Decision and Game Theory for Security*, Nov. 2015.
- [26] B. Johnson, A. Laszka, and J. Grossklags, "Games of Timing for Security in Dynamic Environments," in *Decision and Game Theory for Security*, Nov. 2015.
- [27] V. Pham and C. Cid, "Are We Compromised? Modelling Security Assessment Games," in *Decision and Game Theory for Security*, Nov. 2012.
- [28] A. Laszka, G. Horvath, M. Felegyhazi, and L. Buttyán, "FlipThem: Modeling Targeted Attacks with FlipIt for Multiple Resources," in *Decision and Game Theory for Security*, Nov. 2014.
- [29] D. Leslie, C. Sherfield, and N. P. Smart, "Threshold FlipThem: When the Winner Does Not Need to Take All," in *Decision and Game Theory for Security*, Nov. 2015.

- [30] S. Rass, S. König, and S. Schauer, “Defending Against Advanced Persistent Threats Using Game-Theory,” *PLOS ONE*, vol. 12, Jan. 2017.
- [31] J. Pawlick, S. Farhang, and Q. Zhu, “Flip the Cloud: Cyber-Physical Signaling Games in the Presence of Advanced Persistent Threats,” in *Decision and Game Theory for Security*, Nov. 2015.
- [32] A. Nochenson and J. Grossklags, “A behavioral investigation of the FlipIt game,” in *12th Workshop on the Economics of Information Security (WEIS)*, Jun. 2013.
- [33] D. Reitter, J. Grossklags, and A. Nochenson, “Risk-seeking in a continuous game of timing,” in *Proc. 12th International Conference on Cognitive Modelling*, Jul. 2013.
- [34] J. Grossklags and D. Reitter, “How Task Familiarity and Cognitive Predispositions Impact Behavior in a Security Game of Timing,” in *2014 IEEE 27th Computer Security Foundations Symposium*, Jul. 2014.
- [35] J. Harsanyi and R. Selten, *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.

Appendix A.

Derivation of anonymous player gains

This appendix gives detailed derivations of the anonymous player gains for exponential and periodic play.

A.1. Exponential play

Lemma 2 (Anonymous gain for exponential play). *In a discounted game of timing where both the defender and the attacker are playing an exponential strategy, the anonymous gain function is given by:*

$$\bar{G}_i = \frac{\eta_i}{\eta_i + \eta_j - \ln \lambda_i}, \quad (13)$$

where η_i denotes player i 's move rate and λ_i denotes player i 's discount factor for gains.

Proof. Remember that t_0 denotes the time at which the first move by either player happens (Section 3.3) and that the anonymous gain of player i is the expected value of i 's gain over the interval $[t_0, +\infty)$.

The probability of player i being the first player to flip is equal to the probability of player j performing her flip after the first flip of player i , computable as:

$$\begin{aligned} p_i &= \int_{\tau_i=0}^{+\infty} \eta_i e^{-\eta_i \tau_i} \int_{\tau_j=\tau_i}^{+\infty} \eta_j e^{-\eta_j \tau_j} d\tau_j d\tau_i \\ &= \frac{\eta_i}{\eta_i + \eta_j}. \end{aligned}$$

Because the exponential distribution is memoryless, this probability is equal to the chance with which i was the last player to execute a flip at any point in time following the first flip. Consequently, at any point in time in the interval $[t_0, +\infty)$, the probability of player i being in control of the resource is equal to p_i .

Player i 's gain when in control during the entire interval $[t_0, +\infty)$ is equal to

$$\begin{aligned} & -\ln \lambda_i \int_{t_D=0}^{+\infty} \eta_D e^{-\eta_D t_D} \int_{t_A=0}^{+\infty} \eta_A e^{-\eta_A t_A} \\ & \int_{\tau=\min\{t_D, t_A\}}^{+\infty} \lambda_i^\tau d\tau dt_A dt_D \\ &= -(\ln \lambda_i) \eta_D \eta_A \int_{t_D=0}^{+\infty} \left(\int_{t_A=0}^{t_D} e^{-\eta_A t_A - \eta_D t_D} \int_{\tau=t_A}^{+\infty} \lambda_i^\tau d\tau dt_A \right. \\ & \quad \left. + \int_{t_A=t_D}^{+\infty} e^{-\eta_A t_A - \eta_D t_D} \int_{\tau=t_D}^{+\infty} \lambda_i^\tau d\tau dt_A \right) dt_D \\ &= \frac{\eta_A + \eta_D}{\eta_A + \eta_D - \ln \lambda_i}. \quad (29) \end{aligned}$$

Multiplying p_i and Equation (29) yields player i 's anonymous gain:

$$\bar{G}_i = \frac{\eta_i}{\eta_i + \eta_j - \ln \lambda_i}. \quad (30)$$

□

A.2. Periodic play

For periodic play, we derive the gain of the slower and the faster player separately. In the proofs below, s refers to the slower moving player and f to the faster moving player, by which we mean that $\delta_f \leq \delta_s$.

Lemma 17. *For periodic play, the expected anonymous gain of the slower player is given by:*

$$\frac{-1}{\delta_s \delta_f \ln \lambda_s} \left(\delta_f - \frac{\lambda_s^{\delta_f} - 1}{\ln \lambda_s} \right).$$

Proof. We can express the total anonymous gain of the slower player as the sum of the gains that she gets between her moves (she does not get any anonymous gain before her first move). By linearity of expectation, we can therefore derive the slower player's expected total anonymous gain as the sum of expected gains between her moves.

Consider any interval between two subsequent moves of the slower player. At the beginning of the interval, the slower player is always in control. At some point during the interval, the faster player moves, and the slower player loses control. Assuming phase φ_s for the slower player and assuming that it takes the faster player Δ_n units of time to re-take control after the slower player's n th move ($\Delta_n \leq \delta_s$), the gain obtained by the slower player between her n th and $(n+1)$ th move equals:

$$-\ln(\lambda_s) \int_{\tau=\varphi_s+n\delta_s}^{\varphi_s+n\delta_s+\Delta_n} \lambda_s^\tau d\tau = \lambda_s^{\varphi_s+n\delta_s} (1 - \lambda_s^{\Delta_n}).$$

Since the faster player moves periodically with random phase, Δ_n is a random variable that is uniformly distributed between 0 and the δ_f . Assuming phase φ_s , the expected gain obtained by the slower player between her n th and $(n+1)$ th move is therefore:

$$\begin{aligned} & \frac{1}{\delta_f} \int_{\Delta_n=0}^{\delta_f} \lambda_s^{\varphi_s+n\delta_s} (1 - \lambda^{\Delta_n}) d\Delta_n \\ &= \lambda_s^{\varphi_s+n\delta_s} \left(1 - \frac{\lambda^{\delta_f} - 1}{\delta_f \ln(\lambda_s)} \right). \end{aligned}$$

Taking expectations over φ_s yields the expected gain of the slower player between her n th and $(n+1)$ th move:

$$\begin{aligned} & \frac{1}{\delta_s} \int_{\varphi_s=0}^{\delta_s} \lambda_s^{\varphi_s+n\delta_s} \left(1 - \frac{\lambda_s^{\delta_f} - 1}{\delta_f \ln(\lambda_s)} \right) d\varphi_s \\ &= \frac{\lambda_s^{n\delta_s}}{\delta_s} \left(1 - \frac{\lambda_s^{\delta_f} - 1}{\delta_f \ln \lambda_s} \right) \int_{\varphi_s=0}^{\delta_s} \lambda_s^{\varphi_s} d\varphi_s \\ &= \frac{\lambda_s^{n\delta_s}}{\delta_s \delta_f \ln \lambda_s} \left(\delta_f - \frac{\lambda_s^{\delta_f} - 1}{\ln \lambda_s} \right) (\lambda_s^{\delta_s} - 1) \end{aligned}$$

Finally, compute the total expected gain of the slower player as a geometric series by summing up the expected gains of all intervals:

$$\begin{aligned} & \sum_{n=0}^{+\infty} \frac{\lambda_s^{n\delta_s}}{\delta_s \delta_f \ln \lambda_s} \left(\delta_f - \frac{\lambda_s^{\delta_f} - 1}{\ln \lambda_s} \right) (\lambda_s^{\delta_s} - 1) \\ &= \frac{1}{\delta_s \delta_f \ln \lambda_s} \left(\delta_f - \frac{\lambda_s^{\delta_f} - 1}{\ln \lambda_s} \right) (\lambda_s^{\delta_s} - 1) \sum_{n=0}^{+\infty} \lambda_s^{n\delta_s} \\ &= \frac{-1}{\delta_s \delta_f \ln \lambda_s} \left(\delta_f - \frac{\lambda_s^{\delta_f} - 1}{\ln \lambda_s} \right). \end{aligned}$$

□

Lemma 18. For periodic play, the expected anonymous gain of the faster player is given by:

$$\frac{-1}{\delta_s \delta_f \ln \lambda_f} \left(\delta_s - (\delta_s - \delta_f) \lambda_f^{\delta_f} - \frac{\lambda_f^{\delta_f} - 1}{\ln \lambda_f} \right).$$

Proof. This proof follows along the lines of the proof of Lemma 17: we compute the faster player's expected gain between her n th and $(n+1)$ th flip and her total gain as the sum over the expectations of all such intervals.

Consider any interval between two subsequent moves of the faster player. At the beginning of the interval, the faster player is always in control. During some intervals, the slower player moves once; during others she does not move. If the slower player does not move during the interval, the faster player is in control of the resource for its entire duration. Assuming phase φ_s , she then obtains a gain of:

$$-\ln \lambda_f \int_{\tau=\varphi_f+n\delta_f}^{\varphi_f+(n+1)\delta_f} \lambda_f^\tau d\tau = \lambda_f^{\varphi_f+n\delta_f} (1 - \lambda_f^{\delta_f}).$$

Assuming that the slower player does move and that she does so Δ_n units of time after the beginning of the interval, the faster player obtains a gain of:

$$-\ln \lambda_f \int_{\tau=\varphi_f+n\delta_f}^{\varphi_f+n\delta_f+\Delta_n} \lambda_f^\tau d\tau = \lambda_f^{\varphi_f+n\delta_f} (1 - \lambda_f^{\Delta_n}).$$

Taking expectations over the slower player's strategy, we see that she does not move during an interval with probability $(1 - \frac{\delta_f}{\delta_s})$. If she does move, she does so exactly once at a point in time that is distributed uniformly random over the interval. Assuming phase φ_f , the expected gain obtained by the faster player between her n th and $(n+1)$ th move is therefore:

$$\begin{aligned} & \lambda_f^{\varphi_f+n\delta_f} \left(\left(1 - \frac{\delta_f}{\delta_s} \right) (1 - \lambda_f^{\delta_f}) + \frac{\delta_f}{\delta_s} \int_{\Delta_n=0}^{\delta_f} \frac{1 - \lambda_f^{\Delta_n}}{\delta_f} d\Delta_n \right) \\ &= \lambda_f^{\varphi_f+n\delta_f} \left(\left(1 - \frac{\delta_f}{\delta_s} \right) (1 - \lambda_f^{\delta_f}) + \frac{\delta_f}{\delta_s} \left(1 - \frac{\lambda_f^{\delta_f} - 1}{\delta_f \ln \lambda_f} \right) \right) \\ &= \lambda_f^{\varphi_f+n\delta_f} \left(1 - \left(1 - \frac{\delta_f}{\delta_s} \right) \lambda_f^{\delta_f} - \frac{1}{\delta_s} \frac{\lambda_f^{\delta_f} - 1}{\ln \lambda_f} \right) \\ &= \frac{\lambda_f^{\varphi_f+n\delta_f}}{\delta_s} \left(\delta_s - (\delta_s - \delta_f) \lambda_f^{\delta_f} - \frac{\lambda_f^{\delta_f} - 1}{\ln \lambda_f} \right). \end{aligned}$$

Take expectations over φ_s to obtain the expected gain of the faster player between her n th and $(n+1)$ th move:

$$\begin{aligned} & \frac{\lambda_f^{n\delta_f}}{\delta_s} \left(\delta_s - (\delta_s - \delta_f) \lambda_f^{\delta_f} - \frac{\lambda_f^{\delta_f} - 1}{\ln \lambda_f} \right) \int_{\varphi_f=0}^{\delta_f} \frac{\lambda_f^{\varphi_f}}{\delta_f} d\varphi_f \\ &= \frac{\lambda_f^{n\delta_f}}{\delta_s \delta_f \ln \lambda_f} \left(\delta_s - (\delta_s - \delta_f) \lambda_f^{\delta_f} - \frac{\lambda_f^{\delta_f} - 1}{\ln \lambda_f} \right) (\lambda_f^{\delta_f} - 1). \end{aligned}$$

Finally, compute the total anonymous gain of the faster player as a geometric series by summing up the expected gains of all intervals:

$$\begin{aligned} & \sum_{n=0}^{+\infty} \lambda_f^{n\delta_f} \frac{\delta_s - (\delta_s - \delta_f) \lambda_f^{\delta_f} - \frac{\lambda_f^{\delta_f} - 1}{\ln \lambda_f}}{\delta_s \delta_f \ln \lambda_f} (\lambda_f^{\delta_f} - 1) \\ &= \frac{\delta_s - (\delta_s - \delta_f) \lambda_f^{\delta_f} - \frac{\lambda_f^{\delta_f} - 1}{\ln \lambda_f}}{\delta_s \delta_f \ln \lambda_f} (\lambda_f^{\delta_f} - 1) \sum_{n=0}^{+\infty} \lambda_f^{n\delta_f} \\ &= - \frac{\delta_s - (\delta_s - \delta_f) \lambda_f^{\delta_f} - \frac{\lambda_f^{\delta_f} - 1}{\ln \lambda_f}}{\delta_s \delta_f \ln \lambda_f} \end{aligned}$$

□

Lemma 4 simply combines Lemmas 17 and 18.

Appendix B. Equilibria for periodic play

B.1. Player incentives

This section lists the player incentive functions for periodic play and provides proofs for the properties that are listed in Section 4.3.2.

B.1.1. Incentive functions. Remember that we defined a player's incentive as the partial derivative of her utility to her play rate.

Defender incentives. The defender incentives for both faster and slower play are as follows:

$$\left. \frac{\partial u_D}{\partial \alpha_D} \right|_{\alpha_D \geq \alpha_A} = \frac{-\alpha_A}{\ln \lambda_D} h(\alpha_D) - c_D \quad (31)$$

$$\left. \frac{\partial u_D}{\partial \alpha_D} \right|_{\alpha_D \leq \alpha_A} = \frac{-\alpha_A}{\ln \lambda_D} h(\alpha_A) - c_D \quad (32)$$

where

$$h(\alpha) = \frac{\lambda_D^{1/\alpha} - 1}{\ln \lambda_D} - \frac{\lambda_D^{1/\alpha}}{\alpha}.$$

Attacker incentives. The attacker incentives both faster and slower play are as follows:

$$\left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\alpha_A \geq \alpha_D} = \frac{-\alpha_D}{\ln \lambda_A} \left(\frac{1}{\alpha_D} - f(\alpha_A) \right) - c_A \quad (33)$$

$$+ \frac{(\alpha_A - \alpha_D)(\alpha_A - \ln \lambda_A)}{\alpha_A^2 \ln \lambda_A} \lambda_A^{1/\alpha_A}$$

$$\left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\alpha_A \leq \alpha_D} = \frac{-\alpha_D}{\ln \lambda_A} \left(\frac{1}{\alpha_D} - f(\alpha_D) \right) - c_A, \quad (34)$$

where

$$f(\alpha) = \frac{\lambda_A^{1/\alpha} - 1}{\ln \lambda_A}.$$

From these expressions, continuity and finiteness follows directly.

Lemma 6. *Players' incentives are finite and incentive functions are continuous in all their arguments, including the players' play rates.*

Proof. Simply verify that verify that Equations (33) and (34) yield finite values and the same value for $\alpha_A = \alpha_D$, and that the same goes for Equations (31) and (32). \square

B.1.2. Incentive changes with respect to own play rate.

Lemma 7. *Players' incentives are independent of their play rate if they are the slower moving player, and strictly decreasing in their play rate otherwise.*

Proof. To see that the defender's and attacker's incentive for slower play is constant with respect to α_D and α_A respectively, simply note that α_D and α_A are not part of their respective function definitions.

To see that the defender's incentive for faster play is strictly decreasing in α_D when she is the faster player, compute its partial derivative to α_D :

$$\left. \frac{\partial^2 u_D}{\partial \alpha_D^2} \right|_{\alpha_D > \alpha_A} = -\frac{\alpha_A}{\alpha_D^3} \lambda_D^{1/\alpha_D}.$$

This function is strictly negative.

To see that the attacker's incentive for faster play is strictly decreasing in α_A , compute its partial derivative to α_A :

$$\left. \frac{\partial^2 u_A}{\partial \alpha_A^2} \right|_{\alpha_A > \alpha_D} = -\frac{\lambda_A^{1/\alpha_A}}{\alpha_A^4} (\alpha_A \alpha_D - (\alpha_A - \alpha_D) \ln \lambda_A).$$

This function is strictly negative. \square

B.1.3. Incentive changes with respect to opponent play rate. This section lists the properties of incentive changes with respect to the opponent's play rate.

Remark 1 (Constant total gain). The resource is always under the control of either the defender or the attacker. If both players have the same discount rates, the sum of player gains is therefore constant. Because we normalize gains, this constant is equal to one. Writing i 's and j 's gain as a function of play rates and discounting parameters, we have:

$$G_i(\alpha_i, \alpha_j, \lambda) + G_j(\alpha_i, \alpha_j, \lambda) = 1,$$

for arbitrary λ . Specifically, this also means that we can write i 's gain as a function of j 's gain:

$$G_i(\alpha_i, \alpha_j, \lambda_i) = 1 - G_j(\alpha_i, \alpha_j, \lambda_i).$$

In the remainder of this section, we will simply state that player gains sum to one under the understanding that G_j should be evaluated using i 's discounting parameters.

Remark 2 (Opposite incentive changes). Given equal discounting parameters of the defender and the attacker, the sum of player gains is constant (see Remark 1). Therefore, a change of one player's gain then causes an equal but opposite change of the other player's gain. In other words, we can write:

$$\left. \frac{\partial u_i}{\partial \alpha_j} \right|_{\alpha_i \geq \alpha_j} = - \left. \frac{\partial u_j}{\partial \alpha_j} \right|_{\alpha_j \leq \alpha_i} + c_j - c_i.$$

We can consequently also write the change of player i 's incentive with respect to player j 's move rate as:

$$\left. \frac{\partial^2 u_i}{\partial \alpha_j \partial \alpha_i} \right|_{\alpha_i \geq \alpha_j} = \left. \frac{\partial^2 u_i}{\partial \alpha_i \partial \alpha_j} \right|_{\alpha_i \geq \alpha_j} = - \left. \frac{\partial^2 u_j}{\partial \alpha_i \partial \alpha_j} \right|_{\alpha_j \leq \alpha_i}.$$

Note that this, together with Lemma 7, immediately implies that the rate of change of player i 's incentive with respect to player j 's play rate, is constant with respect to player j 's play rate.

Lemma 8. *If the defender is the faster moving player ($\alpha_D > \alpha_A$), then*

- *the defender's incentive is strictly increasing in α_A , and*
- *the attacker's (base) incentive is strictly decreasing in α_D .*

Proof. To see that the attacker's base incentive is strictly decreasing in α_D , compute its first partial derivative to α_D :

$$\left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\alpha_A < \alpha_D} = -\frac{1 + g(\alpha_D, \lambda_A)}{(\ln \lambda_A)^2},$$

where

$$g : \begin{cases} \mathbb{R}_0^+ \times (0, 1) \rightarrow (-1, 0) \\ (\alpha, \lambda) \mapsto \left(\frac{\ln \lambda}{\alpha} - 1 \right) \lambda^{1/\alpha}. \end{cases} \quad (35)$$

The partial derivative of Equation (34) is strictly negative, because g is bounded between -1 and 0. One way to see this is by verifying that $g(\alpha, \lambda)$ is strictly decreasing in α ,¹ limited by 0 for $\alpha_D \rightarrow 0^+$ and limited by -1 for $\alpha_D \rightarrow +\infty$.

The statement about the defender's incentive function follows from Remark 2. \square

Lemma 9. *If the defender is the slower moving player ($\alpha_D < \alpha_A$), then*

- the defender's (base) incentive is first strictly increasing, then strictly decreasing in α_A , and
- the attacker's incentive is decreasing, independent of or increasing in α_D , depending on α_A .

Proof. To see that the defender's base incentive is first strictly increasing, then strictly decreasing in α_A , compute its first and second partial derivative to α_A :

$$\left. \frac{\partial^2 u_D}{\partial \alpha_A \partial \alpha_D} \right|_{\alpha_D < \alpha_A} = \frac{1 + g(\alpha_A, \lambda_D)}{(\ln \lambda_D)^2} - \frac{\lambda_D^{1/\alpha_A}}{\alpha_A^2} \quad (36)$$

$$\left. \frac{\partial^3 u_D}{\partial \alpha_A^2 \partial \alpha_D} \right|_{\alpha_D < \alpha_A} = \frac{\alpha_A + \ln \lambda_D}{\alpha_A^4} \lambda_D^{1/\alpha_A}. \quad (37)$$

In Equation (36), g is as in Equation (35). The second partial derivative (Equation (37)) starts out negative, has a single root at $\alpha_A = -\ln \lambda_D$ and then becomes positive, or equivalently the first partial derivative (Equation (36)) is decreasing in α_A for $\alpha_A \in [0, -\ln \lambda_D)$ and increasing in α_A for $\alpha_A \in (-\ln \lambda_D, +\infty)$. To see that the first derivative is first positive and then negative, compute its value for $\alpha_A \rightarrow 0$, at $\alpha_A = -\ln \lambda_D$ and for $\alpha_A \rightarrow +\infty$:

$$\begin{aligned} \lim_{\alpha_A \rightarrow 0^+} \left. \frac{\partial^2 u_D}{\partial \alpha_A \partial \alpha_D} \right|_{\alpha_D < \alpha_A} &= \frac{1}{(\ln \lambda_D)^2} > 0 \\ \left. \frac{\partial^2 u_D}{\partial \alpha_A \partial \alpha_D} \right|_{\substack{\alpha_A = -\ln \lambda_D \\ \alpha_D < \alpha_A}} &= \frac{e-3}{e} \frac{1}{(\ln \lambda_D)^2} < 0 \\ \lim_{\alpha_A \rightarrow +\infty} \left. \frac{\partial^2 u_D}{\partial \alpha_A \partial \alpha_D} \right|_{\alpha_D < \alpha_A} &= 0. \end{aligned}$$

The first derivative of defender's base incentive starts out positive, decreases until it becomes strictly negative, then starts increasing again from $\ln \lambda_D$ at such a slow rate

1. It's partial derivative to α ,

$$\frac{\partial g}{\partial \alpha}(\alpha, \lambda) = -\frac{(\ln \lambda)^2}{\alpha^3} \lambda^{1/\alpha}$$

is strictly negative.

that it does not become positive for any finite α_A . In conclusion, the defender's base incentive is first strictly increasing, then forever strictly decreasing in α_A .

The statement about the attacker's incentive function follows from Remark 2. \square

B.2. Periodic best responses for faster play

Lemma 19. *Assume that $\text{BR}_A(0) > 0$ and let α_D^* be the root of the attacker's base incentive function. Then for all play rates*

$$\alpha_A \in [\min\{\text{BR}_A(0), \alpha_D^*\}, \max\{\text{BR}_A(0), \alpha_D^*\}] \text{ and } \alpha_D < \alpha_A,$$

the following implications hold:

- If $\text{BR}_A(0) < \alpha_D^*$, then the attacker's incentive is strictly increasing in α_D .
- If $\text{BR}_A(0) = \alpha_D^*$, then the attacker's incentive is zero for all $\alpha_D \in [0, \alpha_D^*]$.
- If $\text{BR}_A(0) > \alpha_D^*$, then the attacker's incentive is strictly decreasing in α_D .

Proof. If $\text{BR}_A(0) \leq \alpha_D^*$, we have from the definitions of α_D^* and $\text{BR}_A(0)$:

$$0 = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A \leq \alpha_D \\ \alpha_D = \alpha_D^*}} = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D = \alpha_D^*}} = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D = 0}},$$

implying:

$$\int_{\bar{\alpha}_D=0}^{\alpha_D^*} \left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\substack{\alpha_D = \bar{\alpha}_D \\ \alpha_A = \text{BR}_A(0)}} d\bar{\alpha}_D = 0.$$

Because $\left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\alpha_A \leq \alpha_D}$ and $\left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\alpha_A \geq \alpha_D}$ are both independent from α_D (Remark 2), this is equal to:

$$\begin{aligned} \text{BR}_A(0) \left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\substack{\alpha_A \geq \alpha_D \\ \alpha_A = \text{BR}_A(0)}} + \quad (38) \\ (\alpha_D^* - \text{BR}_A(0)) \left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\substack{\alpha_A \leq \alpha_D \\ \alpha_A = \text{BR}_A(0)}} = 0. \end{aligned}$$

If $\alpha_D^* = \text{BR}_A(0)$, Equation (38) implies that:

$$\left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\substack{\alpha_A \geq \alpha_D \\ \alpha_A = \text{BR}_A(0)}} = 0,$$

and consequently that the attacker's incentive when playing at rate $\alpha_A = \alpha_D^* = \text{BR}_A(0)$ is equal to zero for all $\alpha_D \in [0, \alpha_D^*]$.

If $\alpha_D^* > \text{BR}_A(0)$, Equation (38) and the fact that the attacker incentive for slower play is strictly decreasing in α_D together imply

$$\left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\substack{\alpha_A \geq \alpha_D \\ \alpha_A = \text{BR}_A(0)}} > 0,$$

and consequently that the attacker's incentive is always strictly increasing in α_D if $\alpha_A \geq \text{BR}_A(0)$ and $\alpha_D < \alpha_A$.

Lastly, if $\text{BR}_A(0) > \alpha_D^*$, we have from the definition of α_D^* and the fact that the incentive of the faster player is decreasing in her play rate that:

$$0 = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A \leq \alpha_D \\ \alpha_D = \alpha_D^*}} = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \alpha_D^* \\ \alpha_D = \alpha_D^*}} > \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D = \alpha_D^*}},$$

implying:

$$\left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D = 0}} > \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D = \alpha_D^*}},$$

which in turn implies:

$$\left. \frac{\partial^2 u_A}{\partial \alpha_D \partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D \leq \alpha_A}} < 0,$$

and consequently that the attacker's incentive is strictly decreasing in α_D for all $\alpha_D \leq \text{BR}_A(0)$. \square

Lemma 16 (Best-response with faster attacker). *Let $\bar{\alpha}_A$ be any strictly positive attacker play rate. Then the following statements are equivalent:*

- *There exists a defender play rate $\bar{\alpha}_D \leq \bar{\alpha}_A$ to which $\bar{\alpha}_A$ is a best response.*
- *The attacker's base incentive function has a positive root α_D^* , and*

$$\bar{\alpha}_A \in [\min\{\text{BR}_A(0), \alpha_D^*\}, \max\{\text{BR}_A(0), \alpha_D^*\}].$$

If $\alpha_D^ = \text{BR}_A(0)$, then $\bar{\alpha}_D$ can be any value in $[0, \bar{\alpha}_A]$. Otherwise, $\bar{\alpha}_D$ is unique.*

Proof. Assume that $\bar{\alpha}_A$ is a best response to $\bar{\alpha}_D$. Then, since the attacker's incentive for faster play is strictly decreasing in α_A :

$$0 = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \bar{\alpha}_A \\ \alpha_D = \bar{\alpha}_D}} \leq \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \bar{\alpha}_D \\ \alpha_D = \bar{\alpha}_D}} = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A \leq \alpha_D \\ \alpha_D = \bar{\alpha}_D}}.$$

Because the attacker's base incentive is strictly decreasing in α_D , this implies that the attacker's base incentive is zero for some $\alpha_D^* \geq \bar{\alpha}_D$.

Now, assume that $\text{BR}_A(0) \leq \alpha_D^*$. Then, from the definition of $\text{BR}_A(0)$ and the fact that the attacker's incentive is non-decreasing in α_D by Lemma 19, we have:

$$0 = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D = 0}} \leq \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D = \bar{\alpha}_D}}. \quad (39)$$

Because the attacker incentive for faster play is strictly decreasing in his play rate, we have:

$$0 = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A \leq \alpha_D \\ \alpha_D = \alpha_D^*}} = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \alpha_D^* \\ \alpha_D = \alpha_D^*}} \geq \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \alpha_D^* \\ \alpha_D = \bar{\alpha}_D}}. \quad (40)$$

Since the attacker's incentive is non-increasing in her own play rate, these two inequalities imply that the attacker's incentive is zero for some $\bar{\alpha}_A \in [\text{BR}_A(0), \bar{\alpha}_D]$. For $\text{BR}_A(0) \geq \alpha_D^*$, the reasoning is the same but the inequalities in Equations (39) and (40) are reversed, so $\bar{\alpha}_A \in [\bar{\alpha}_D, \text{BR}_A(0)]$.

Now assume that there exists a $\alpha_D^* \geq 0$ for which the attacker's base incentive is zero and that $\bar{\alpha}_A \in [\text{BR}_A(0), \alpha_D^*]$.

Assume $\text{BR}_A(0) = \alpha_D^*$. The fact that any $\bar{\alpha}_A$ is a best response to any α_D in $[0, \bar{\alpha}_A]$ follows directly from Lemma 19.

Assume $\text{BR}_A(0) < \alpha_D^*$. Since the attacker's incentive for faster play is strictly decreasing in her play rate, we have:

$$0 = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \text{BR}_A(0) \\ \alpha_D = 0}} \leq \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \bar{\alpha}_A \\ \alpha_D = 0}}.$$

Since his base incentive is decreasing in α_A , we also have:

$$0 = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A \leq \alpha_D \\ \alpha_D = \alpha_D^*}} = \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \alpha_D^* \\ \alpha_D = \alpha_D^*}} \geq \left. \frac{\partial u_A}{\partial \alpha_A} \right|_{\substack{\alpha_A = \bar{\alpha}_A \\ \alpha_D = \alpha_D^*}}.$$

From Lemma 19, we know that the attacker's incentive is strictly increasing in α_D . Therefore, there is precisely one value $\bar{\alpha}_D \leq \bar{\alpha}_A$ for which the attacker's incentive when playing $\bar{\alpha}_A$ is zero and to which $\bar{\alpha}_A$ consequently is a best response. For $\text{BR}_A(0) > \alpha_D^*$, the reasoning is identical but the inequalities above are reversed and the attacker's incentive is strictly decreasing in α_D . \square