

Quasi-regular sequences and optimal schedules for security games

David Kempe*

Leonard J. Schulman†

Omer Tamuz‡

October 31, 2017

We study *security games* in which a defender commits to a mixed strategy for protecting a finite set of targets of different values. An attacker, knowing the defender’s strategy, chooses which target to attack and for how long. If the attacker spends time t at a target i of value α_i , and if he leaves before the defender visits the target, his utility is $t \cdot \alpha_i$; if the defender visits before he leaves, his utility is 0. The defender’s goal is to minimize the attacker’s utility. The defender’s strategy consists of a schedule for visiting the targets; it takes her unit time to switch between targets. Such games are a simplified model of a number of real-world scenarios such as protecting computer networks from intruders, crops from thieves, etc.

We show that optimal defender play for such security games, although played in continuous time, reduces to the solution of a combinatorial question regarding the existence of infinite sequences over a finite alphabet, with the following properties for each symbol i : (1) i constitutes a prescribed limiting fraction p_i of the sequence. (2) The occurrences of i are spread apart close to evenly, in that the ratio of the longest to shortest interval between consecutive occurrences is bounded by a parameter K . We call such sequences K -quasi-regular; a 1-quasi-regular sequence is one in which the occurrences of each symbol form an arithmetic sequence.

As we show, a 1-quasi-regular sequence ensures an optimal defender strategy for these security games: the intuition for this fact lies in the famous “inspection paradox.” However, as we demonstrate, for $K < 2$ and general p_i , K -quasi-regular sequences may not exist. Fortunately, this does not turn out to be an obstruction: we show that, surprisingly, 2-quasi-regular sequences also suffice for optimal defender play. What is more, even randomized 2-quasi-regular sequences suffice for optimality. We show that such sequences always exist, and can be calculated efficiently. Thus, we can ensure optimal defender play for these security games.

The question of the least K for which deterministic K -quasi-regular sequences exist is fascinating. Using an ergodic theoretical approach, we proceed to show that deterministic 3-quasi-regular sequences always exist (and can be calculated efficiently). We also show that these deterministic 3-regular sequences give rise to a ≈ 1.006 -approximation algorithm for the defender’s optimal strategy. For $2 \leq K < 3$ we do not know whether deterministic K -quasi-regular sequences always exist; however, when the p_i are all small, improved bounds are possible, and in fact, $(1 + \epsilon)$ -quasi-regular deterministic sequences exist for any $\epsilon > 0$ for sufficiently small p_i .

*Department of Computer Science, University of Southern California. Supported in parts by NSF grants 1619458 and 1423618.

†Engineering and Applied Science, California Institute of Technology. Supported in part by NSF grants 1319745, 1618795, and a EURIAS Senior Fellowship co-funded by the Marie Skłodowska-Curie Actions under the 7th Framework Programme.

‡Departments of Economics and Mathematics, California Institute of Technology

1 Introduction

One of the most successful real-world applications at the intersection of game theory and computer science has been *security games*. Security games have been used recently to model and address problems including the protection of infrastructure (airports, seaports, flights), deterrence of fare evasion and smuggling, as well as the protection of wildlife and plants. The related model of *inspection games* [4] has been used to model interactions as varied as arms control, accounting and auditing, environmental controls, or data verification.

In general models of security games, there are n targets of various values that the defender tries to protect with her limited resources.¹ Different assumptions and scenarios can lead to different interesting combinatorial constraints; see [31] for an overview of much recent work.

In the present work, we are concerned with a defender who cannot switch instantaneously between different targets, which introduces a timing component and a scheduling problem. At a high level, such constraints arise in many natural security settings, including:

1. Protection of computer networks (with multiple databases or account holders) from infiltrators.
2. Protection of wildlife from poachers (e.g., [13, 12]), crops or other plants from thieves, or homes in a neighborhood from burglars.

Stripping away details, we propose the following simplified model for these types of settings: If the attacker has access to an unprotected target, he gains utility in proportion to the value of the target and to the time he spends at the target.² The game is zero-sum, i.e., the attacker’s gain is the defender’s loss. If the attack is interrupted by the defender at any time, both players receive utility 0. Due to physical distances between targets or switching costs between databases, the defender requires one unit of time to switch between any two targets. The problem of interest is to determine a schedule for the defender that will lead to minimum expected defender loss against a best-responding attacker.

More formally, each of the n targets has value $\alpha_i \geq 0$, scaled so that $\sum_i \alpha_i = 1$. We assume that no target is strictly more valuable than all other targets combined,³ so that $\alpha_i \leq 1/2$. Time is continuous, and the game has an infinite time horizon. A defender strategy is a schedule of which target is visited at each point in time, with some time spent in transit. An attacker strategy consists of a choice of a single time interval $[t_0, t_0 + t]$ and a target i to attack. If the defender does not visit target i during this interval, the attacker obtains a utility of $U = t \cdot \alpha_i$, and the defender receives $-U$. If the defender visits the target at any point during the interval, both players’ utilities are 0.

In our modeling of the security setting, an attack is a sufficiently disruptive event that it effectively ends the game, by which we mean that the defender will subsequently re-randomize her schedule. Thus, the entire game is concluded after one attack. We assume that the defender’s mixed strategy (but not randomness) is known to the attacker who will best-respond by choosing

¹For consistency, we always refer to the attacker with male pronouns and the defender with female pronouns.

²In the case of access to computer systems, this models a scenario observed in recent attacks where the attacker lurks — whether in order to monitor legitimate users, create ongoing damage, or because file sizes or bandwidth concerns make it impossible to download the entire database in a short amount of time.

³See Section 7 for a discussion of this choice; outside of this assumption is a different régime that requires different analysis.

an attack (i, t_0, t) , comprising the target i and the start time and duration⁴ of the attack, t_0 and t . Accordingly, the defender chooses a Min-Max strategy: a strategy which minimizes the maximum expected return of any attack (i, t_0, t) .

Although the game is played in continuous time, we show (in Section 2) that under the assumption that $\alpha_i \leq 1/2$ for all i , optimal defender strategies can be obtained as follows: choose a suitable (random) *discrete sequence* s_0, s_1, s_2, \dots of targets and a uniformly random $\tau \in [0, 1]$; then visit each target s_t at time $t + \tau \in \mathbb{R}^+$. At all other times, the defender is in transit, so that each target is visited instantaneously, with the next target visited exactly one time unit later.

It is not difficult to show that if each target i occupies a α_i fraction of the sequence and is exactly evenly spaced in the sequence, the resulting defender schedule is optimal. Some intuition is derived from the famous “Inspection Paradox” or “Waiting Time Problem:” passengers of a bus service which departs a station with perfect regularity (e.g., 15 minutes apart) wait on average half as long as passengers of a service with the same frequency of operation but Poisson departure times. In our case, higher variance in the defender’s interarrival times lengthens the expected time until the next defender visit, making longer attacks more attractive.

We call sequences in which each target i is exactly evenly spaced *regular*; generalizing this notion, we call a sequence s K -quasi-regular (with respect to frequencies p_i ; in our applications, we typically choose $p_i = \alpha_i$) if, as before, each target i takes up a p_i fraction of the sequence, and the ratio of the longest to shortest interval between consecutive occurrences of i in s is bounded by K . Our first result (Theorem 3.1) is that — surprisingly — it suffices for optimality that the sequence is *2-quasi-regular*.

1.1 Our Main Result

It is fairly straightforward to show that there is some vector $(p_i)_i$ such that there are no $(2-\epsilon)$ -quasi-regular sequences for any $\epsilon > 0$; we do this in Section 4. Our main result (Theorem 4.1 in Section 4) is that for any values p_i , there exists a 2-quasi-regular random sequence, which can furthermore be efficiently computed from the p_i . By the aforementioned Theorem 3.1, the corresponding defender mixed strategy is optimal.

1.2 Ergodic Schedules

Quasi-regular sequences are basic combinatorial objects, quite apart from our application of them. One limitation of our work (although it does not affect the application) is that the resulting schedules are not ergodic: they randomize between different schedules in which the targets have frequencies differing from the desired p_i . It is then a natural question whether 2-quasi-regular ergodic sequences can be obtained as well. This is related to the following combinatorial question: given densities p_i , does there always exist a 2-quasi-regular deterministic sequence?

We provide two partial answers to this question. In Section 5, we analyze a very simple schedule called the *Golden Ratio Schedule* (variations of which have been studied in the context of hashing [22, pp. 510,511,543], bandwidth sharing [21, 26] and elsewhere). This schedule is generated by the following random sequence: partition the circle of circumference 1 into intervals of size p_i

⁴One could consider an “adaptive” attacker, who initially only chooses the start time t_0 , and decides on t on the fly. The resulting model would be equivalent, as such an adaptive attacker at any time $t' > t_0$ has learned no new information aside from the fact that he has not yet been caught; any attacker who under these conditions will decide to wait for exactly t units of time (if not caught) is exactly equivalent to one who chooses the attack (i, t_0, t) at once.

corresponding to the targets i . Choose a uniformly random starting point on the circle. In each step, add φ to the current point; here, $\varphi = \frac{1}{2}(1 + \sqrt{5})$ is the Golden Ratio. In each time step, the defender visits the target i into whose interval the current point falls.

This random sequence is ergodic, and at worst 3-quasi-regular.⁵ Moreover, for any choice of the random starting point, the deterministic sequence is 3-quasi-regular. Thus we show that there always exist deterministic 3-quasi-regular sequences. We do not know if this is true for any $K < 3$.

It is interesting that such a simple schedule achieves constant quasi-regularity, but the bound is not strong enough to guarantee optimality of the schedule for the defender. However, we show that the schedule is nearly optimal for the defender: the attacker’s utility is within a factor of at most 1.006 of the minimum attacker utility (and this bound is tight). The proof of this approximation guarantee relies on a theorem of Slater about simple dynamical systems like the Golden Ratio shift, and a somewhat intricate analysis of the attacker’s response. We find it remarkable that such a simple policy comes provably within 0.6% of the optimum, in particular compared to another very simple policy: as we show in Appendix A, the simple *i.i.d. schedule*, which always chooses the next target i to visit with probability α_i , independent of the history, is only a $4/e$ -approximation.

As a second partial result towards obtaining an optimal ergodic schedule, in Section 6, we show a sufficient condition for the existence of $(1 + \epsilon)$ -quasi-regular sequences, for any $\epsilon > 0$. Specifically, let M be the smallest common denominator of all p_i . If $p_i = O(\frac{\epsilon}{\sqrt{n \log M}})$ for all i , then a $(1 + \epsilon)$ -quasi-regular periodic schedule exists and can be found efficiently using a randomized algorithm that succeeds with high probability.

The algorithm is based on placing points for target i at uniform distance proportional to $1/p_i$ on the unit circle, with independent uniformly random offsets. Points can only be matched to sufficiently close multiples of $1/M$. An application of Hall’s Theorem, similar to [32, 17], shows that under the conditions of the theorem, this algorithm succeeds with high probability in producing a $(1 + \epsilon)$ -quasi-regular sequence.

Related Work

Several notions of “sequences in which elements i are roughly regularly spaced, with given frequencies p_i ” have been studied in different contexts.

Dating back to the work of Tijdeman [32, 33], several papers [1, 17, 32, 33, 9] (see also an overview in [8]) have studied sequences with low *discrepancy* in the following sense: up to any time t , the number of occurrences of element i approximates $t \cdot p_i$ as closely as possible. For our application, the rate of convergence of the frequencies to p_i is not essential; but it is crucial that the defender’s interarrival times at each target be as regular as possible. Consequently, methods from this literature are not sufficient to optimally solve our problem.

In the *Pinwheel Problem* [18, 19, 6, 24, 14], one is given integers $n_1, n_2, \dots, n_k \geq 2$, whose density is defined by $\beta := \sum_i 1/n_i \leq 1$. The goal is to produce a sequence s of the elements $\{1, \dots, k\}$ such that for each i , each subsequence $(s_t, s_{t+1}, \dots, s_{t+n_i-1})$ of length n_i contains at least one occurrence of the symbol i . One of the main questions in this area is what values β are sufficient to guarantee the existence of such a sequence. Currently, it is known that $\beta \leq 5/6$ is necessary (i.e., there are examples with $\beta = 5/6 + \epsilon$ such that no s exists), and $\beta \leq 7/10$ is sufficient. In our case, we always have desired frequencies adding up to 1, and we have not only (looser) upper bounds, but also lower bounds on the distance between consecutive occurrences. While some of the basic techniques

⁵When all $p_i \leq 1 - 1/\varphi$, the quasi-regularity guarantee improves to $8/3$, and as $p_i \rightarrow 0$, it converges to φ^2 .

in the context of the Pinwheel Problem are similar to the ones we use (in Section 4), solutions for one problem do not imply solutions to the other.

In concurrent and independent work, Immorlica and Kleinberg [20] — motivated in part also by applications to preventing wildlife poaching — defined a “Recharging Bandits” problem in which the available reward at targets grows according to (unknown) concave functions. In the full-information setting (in which the functions are known), (near-)optimal solutions correspond to schedules that visit targets with given frequencies at roughly evenly spaced intervals. The precise definition of “roughly evenly spaced” differs from ours, and while some of the techniques used in [20] are similar to ours, optimality results and approximation guarantees do not imply each other in either direction.

Our work is related to the inspection games literature (see, e.g., [4, 35]). Specifically, several articles [11, 2, 3] consider models in which an inspectee indulges in illegal activity once or multiple times within a finite time interval. An inspector distributes optimally the times at which she performs a fixed number of inspections, and suffers a loss that is proportional to the time that has elapsed between the beginning of illegal activity and the next inspection. In these models, as in ours, the inspector wants to visit inspectees regularly while keeping the inspectee uncertain about visit times. The lack of travel time restrictions as well as the lack of a need to catch the inspectee at the time of his action make the models mathematically incomparable.

Finally, our work is also related to the literature on *patrol games*. Here, as in our model, a defender (or multiple cooperating defenders) must decide on a schedule of visits to targets facing attacks. However, unlike our model, the attacker observes the past locations of the defender(s) before deciding whether and where to attack (see, e.g., [5, 36]).

2 Preliminaries

The n targets have values $\alpha_i > 0$ for all i . Because the units in which target values are measured are irrelevant, we assume that $\sum_i \alpha_i = 1$. We assume that no target has value exceeding the sum of all other targets’ values, meaning (after normalization) that $\alpha_i \leq \frac{1}{2}$ for all i .

A *pure strategy* (*schedule*) for the defender is a measurable mapping $\ell : \mathbb{R}^{\geq 0} \rightarrow \{1, 2, \dots, n, \perp\}$, where \perp denotes that the defender is in transit. A schedule ℓ is *valid* if $\ell(t) = i$ and $\ell(t') = j \neq i$ implies that $|t' - t| \geq 1$. In other words, there is enough time for the defender to move from i to j (or from j to i). We use \mathcal{L} to denote the set of all valid pure defender strategies.

The defender moves first and chooses a mixed strategy, i.e., a distribution Λ over \mathcal{L} , or a random ℓ . Then, the attacker chooses an attack (i, t_0, t) consisting of a target i , a start time t_0 , and an attack duration t . Subsequently, a mapping ℓ is drawn from the defender’s distribution Λ . The attacker’s utility is

$$U(\ell, (i, t_0, t)) = \begin{cases} 0 & \text{if } \ell(\tau) = i \text{ for some } \tau \in [t_0, t_0 + t] \\ \alpha_i \cdot t & \text{otherwise.} \end{cases} \quad (1)$$

Since we are considering a zero-sum game (see Section 7 for a discussion), the defender’s utility is $-U(\ell, (i, t_0, t))$. Note that the attacker attacks only once.

A rational attacker will choose (i, t_0, t) so as to maximize $\mathbb{E}_{\ell \sim \Lambda} [U(\ell, (i, t_0, t))]$; therefore, the defender’s goal is to choose Λ to minimize

$$U(\Lambda) = \sup_{i, t_0, t} \mathbb{E}_{\ell \sim \Lambda} [U(\ell, (i, t_0, t))].$$

The fact that we assume an infinite time horizon is primarily an idealization, in order to avoid mathematical inconveniences at the end of the time horizon. The reader is encouraged to think of the “infinite” time horizon as one or a few days, long enough that a significant schedule needs to be planned and boundary effects can be ignored at small cost; but short enough that the attacker cannot observe early parts of the schedule to infer which schedule ℓ was drawn from Λ .

2.1 Canonical, Shift-Invariant, and Ergodic Schedules

The general definition of defender schedules allows for strange schedules that are clearly suboptimal. We would like to restrict our attention to “reasonable” schedules. In particular, we will assume the two following conditions, which we later show to hold without loss of generality. (Here, we will be slightly informal in our definitions. Precise definitions and constructions ensuring these properties are given in Appendix B.)

- Whenever the defender transitions from one target i to a target i' ($i' = i$ is possible), she spends exactly one time unit in transit. We call such schedules *canonical*.
- To the attacker, any two times t and t' “look the same,” in that for any $t, t', \tau \in \mathbb{R}^+$, the distributions of the defender’s schedule restricted to the time intervals $[t, t + \tau]$ and $[t', t' + \tau]$ are the same. We call such schedules *shift-invariant* or *stationary*.

Because the strategy spaces of both players are infinite, it is not clear a priori that a Min-Max schedule for the defender exists. In Appendix B, we show that an optimal mixed Min-Max defender strategy does exist, and is w.l.o.g. canonical and shift-invariant. Therefore, for the remainder of this paper, we will focus only on shift-invariant canonical schedules. When the defender’s strategy is shift-invariant, the start time t_0 of the attack does not matter, so shift-invariance allows us to implicitly assume that the attacker always attacks at time $t_0 = 0$. We then simply write $U(\Lambda, (i, t))$ for the attacker’s expected utility from attacking target i for t units of time.

One may additionally be interested in constructing *ergodic* shift-invariant mixed schedules: Λ is ergodic if Λ cannot be written as the convex combination $\Lambda = \lambda\Lambda_1 + (1 - \lambda)\Lambda_2$ of two different shift-invariant mixed schedules (see the formal discussion in Appendix B). While we are not aware of game-theoretical implications of ergodicity, the question is mathematically natural, and may be important in some extensions of our model.

2.2 Return Times and Target Visit Frequencies

For the following definitions, recall that we are focusing only on shift-invariant canonical schedules. Accordingly, we assume without loss of generality that the attacker starts his attack at time 0. Since the attacker chooses only the target i and the duration t of the attack, from his perspective, the property of the defender’s strategy that matters is the distribution of her next *return time* to target i , defined as $R_i = \min\{t \geq 0 \mid \ell(t) = i\}$. Given a target i and a defender strategy, let $F_i(t) = \mathbb{P}[R_i \leq t]$ denote the CDF of R_i . In particular, notice that $F_i(0)$ is the fraction of time the defender spends waiting at target i . In terms of the distribution of return times F_i , the attacker’s utility can be expressed as follows:

$$U(F_i, (i, t)) = \alpha_i \cdot t \cdot (1 - F_i(t)). \tag{2}$$

Next, we define a random variable B_i capturing the (random) duration between consecutive visits to the same target i . Notice the subtle difference between this quantity and the time from the defender’s perspective between leaving a target and returning to it. By comparison, the distribution of B_i should assign higher probability to higher values: as in the inspection paradox, at a random point in time, the attacker is more likely to find himself in a large gap. In other words, larger gaps are more likely to appear at a fixed time than on average over a long time stretch.

Defining the distribution for B_i precisely requires some care, and is done formally in Appendix B. The construction formalizes the following intuition: we can consider the limit as we shift the random schedule ℓ left by a real number $t \rightarrow \infty$. This extends a shift-invariant random schedule from the non-negative reals to all reals, and allows us to consider the random variable

$$B_i = (\inf\{t \in \mathbb{R}, t > 0 \mid \ell(t) = i\}) - (\sup\{t \in \mathbb{R}, t \leq 0 \mid \ell(t) = i\}).$$

The random variable B_i captures the time between consecutive visits before and after time 0 — by shift-invariance, the time 0 is arbitrary here. Our constructions will always ensure that B_i is finite.

We can write F_i in terms of B_i . First, we note that conditioned on $B_i = \tau$, the distribution of R_i is uniform between 0 and τ . Hence, its conditional CDF is $\mathbb{P}[R_i \leq t \mid B_i = \tau] = \min(1, t/\tau)$. The unconditional CDF of R_i for $t > 0$ is then

$$F_i(t) = \mathbb{E}_{B_i} [\min(1, t/B_i)] = \int_0^1 \mathbb{P}[t/B_i \geq \tau] d\tau = \int_0^1 \mathbb{P}[B_i \leq t/\tau] d\tau = t \cdot \int_t^\infty \frac{\mathbb{P}[B_i \leq \tau]}{\tau^2} d\tau; \quad (3)$$

the final equality uses a change of integration variable. For $t = 0$, we have that $F_i(0) = \mathbb{P}[B_i = 0]$ is the probability that the defender is at target i at time 0. Two key quantities for our analysis are

$$p_i = F_i(1), \quad T_i = \frac{1 - F_i(0)}{p_i - F_i(0)}.$$

In every canonical schedule, p_i is the fraction of time that the defender is either at target i or in transit to i . Thus, $\sum_i p_i = 1$ in every canonical schedule. T_i intuitively captures the “expected time” between consecutive visits to target i , as seen by the defender. However, this intuition formally holds only for periodic schedules, necessitating the preceding more complex definition for arbitrary schedules. The most useful facts about F_i are summarized by the following proposition:

Proposition 2.1 1. $F_i(\cdot)$ is concave.

2. $F_i(t) \leq F_i(0) + (p_i - F_i(0)) \cdot t$, with equality iff $\mathbb{P}[B_i = 0 \text{ or } B_i > t] = 1$.

Proof. 1. For each x , the function $\min(1, t/x)$ is concave in t . Hence, $F_i(t) = \mathbb{E}_x [\min(1, t/x)]$, being a convex combination of concave functions, is concave.

2. Rearranging and applying (3), we want to show that

$$t \cdot \int_1^\infty \frac{\mathbb{P}[B_i \leq \tau]}{\tau^2} d\tau - t \cdot \int_t^\infty \frac{\mathbb{P}[B_i \leq \tau]}{\tau^2} d\tau \geq (t - 1) \cdot F_i(0).$$

Combining the two integrals and lower-bounding $\mathbb{P}[B_i \leq \tau] \geq F_i(0)$ yields that the left hand side is at least

$$F_i(0) \cdot t \cdot \int_1^t \frac{1}{\tau^2} d\tau = F_i(0) \cdot (t - 1),$$

with the lower bound holding with equality iff $\mathbb{P}[B_i = 0 \text{ or } B_i > t] = 1$. ■

2.3 Schedules from Sequences

All the constructions of mixed defender schedules in this paper will have the property that the defender never waits at any target, instead traveling immediately to the next target. That such schedules are optimal (and hence the focus on such constructions is w.l.o.g.) is a consequence of our main Theorem 3.1, and hinges on the restriction that $\alpha_i \leq 1/2$ for all i . A brief discussion of what happens when this assumption is relaxed is given in Section 7.

Canonical schedules without waiting are readily identified with schedules defined only on integer times, since the defender must only choose, after visiting a target, which target she will visit next. We call such schedules *sequences*, defined as $s: \mathbb{N} \rightarrow \{1, \dots, n\}$. A sequence, together with a start time t_0 , naturally defines a canonical schedule, by setting $\ell(t) = s_{t-t_0}$ if $t - t_0 \in \mathbb{N}$, and $\ell(t) = \perp$ otherwise. Σ denotes a distribution over sequences, or the distribution of a random sequence s .

Shift-invariance (or stationarity) can be defined for random sequences as for (continuous) mixed schedules. When s is a periodic sequence, i.e., there is a k such that $s_{t+k} = s_t$ for all t , a shift-invariant random sequence can be obtained particularly easily, by choosing a uniformly random $\kappa \in \{0, \dots, k-1\}$, and defining s' via $s'_t = s_{t+\kappa}$; for aperiodic sequences, we give a construction in Appendix B. From a shift-invariant random sequence, we can obtain a shift-invariant mixed schedule straightforwardly, by choosing the start time $t_0 \in [0, 1]$ uniformly.

For the special case of random sequences, the definitions of p_i and T_i simplify to $p_i = \mathbb{P}[s_1 = i]$ (which is now exactly the fraction of target visits devoted to target i), and $T_i = 1/p_i$ (since $F_i(0) = 0$).

2.4 Regular and Quasi-Regular Sequences

We say that a shift-invariant random sequence s is K -quasi-regular (with respect to target frequencies p_i) if the following two hold for each target i :

1. $\mathbb{P}[s_1 = i] = p_i$.
2. There is some b_i such that $\mathbb{P}[b_i \leq B_i \leq K \cdot b_i] = 1$.

In other words, each target i is visited with frequency p_i , and the maximum gap for consecutive visits to target i is within a factor K of the minimum gap with probability 1. A random sequence is *regular* if it is 1-quasi-regular, meaning that all visits to target i are spaced exactly T_i apart. (All definitions extend directly to canonical, mixed, shift-invariant schedules.)

A particularly straightforward way to obtain a K -quasi-regular random sequence Σ is to consider the a subsequential limit of uniformly random shifts of a deterministic sequence s in which the gaps between consecutive visits to i are bounded between b_i and Kb_i , and the density of entries which are i is p_i .

3 The Attacker's Response, and Optimal Schedules

In this section, we show the following main theorem, a sufficient condition for a random sequence to be optimal for the defender.

Theorem 3.1 *Consider a random shift-invariant sequence such that the following two hold for each target i :*

- $T_i = 1/\alpha_i$.
- For each i , there exists an η_i such that $\mathbb{P}[\frac{\eta_i}{\eta_i+1}T_i \leq B_i \leq \eta_i T_i] = 1$.

Then, the associated mixed strategy is optimal for the defender.

In particular, these conditions hold for 2-quasi-regular random sequences with respect to α_i .

In Section 4, we show that there always exists a 2-quasi-regular sequence. With the eventual goal of proving Theorem 3.1, we fix a target i , and for now drop the subscript i , so that

$$p = p_i \qquad F(t) = F_i(t) \qquad T = T_i \qquad B = B_i.$$

We fix p and T and study which sequences — among all those with given p and T — are optimal for the defender. Our algorithmic constructions will choose $p_i = \alpha_i$ for all i ; however, in order to show the optimality of this choice, the following proposition and corollary are proved for general p, T .

Proposition 3.2 *Consider any canonical shift-invariant mixed defender schedule (over the non-negative real numbers). By choosing $t = T/2$, the attacker guarantees himself a utility of at least $\alpha \cdot \frac{1-F(0)}{4} \cdot T$.*

Proof. By Equation (2), the attacker’s utility at time $t = T/2$ is $\alpha \cdot (T/2) \cdot (1 - F(T/2))$. Using Proposition 2.1, we can bound

$$1 - F(T/2) \geq 1 - F(0) - (p - F(0)) \cdot (T/2) = (1 - F(0)) \cdot \left(1 - \frac{p - F(0)}{1 - F(0)} \cdot (T/2)\right) = \frac{1 - F(0)}{2}.$$

Hence, the attacker’s utility is at least $\alpha \cdot \frac{1-F(0)}{4} \cdot T$. ■

Recalling that a random *sequence* by definition does not involve waiting at any target, we obtain the following simple corollary about random sequences that are worst for the attacker:

Corollary 3.3 *Among random sequences with fixed T and p , any random sequence is optimal for the defender if the attacker’s payoff is upper-bounded by $\frac{1}{4} \cdot \alpha \cdot T$.*

The following corollary is particularly useful:

Corollary 3.4 *Fix T and p , and consider a random sequence in which the attacker’s optimal attack duration t satisfies $\mathbb{P}[B > t] = 1$. Then, this random sequence is optimal for the defender. Furthermore, in this case, w.l.o.g., $t = T/2$.*

Proof. By the assumption that $\mathbb{P}[B > t] = 1$ and Proposition 2.1, we have that $F(t) = p \cdot t$. Hence, the attacker’s utility is $\alpha \cdot (1 - p \cdot t) \cdot t = \alpha \cdot \frac{t(T-t)}{T} \leq \alpha \cdot \frac{T}{4}$. Now, the claim follows directly from Corollary 3.3. That $t = T/2$ is a best response follows from Proposition 3.2. ■

We can now apply these corollaries to show optimality for a single target for which the “quasi-regularity” of return times holds.

Proposition 3.5 Fix T and p , and consider a random sequence such that for some η ,

$$\mathbb{P}\left[\frac{\eta}{\eta+1}T \leq B \leq \eta T\right] = 1.$$

Then, this schedule is optimal for the defender among schedules with these T and p .

Proof. We write $\xi = \frac{\eta}{\eta+1}$. By Proposition 3.2, choosing $t = T/2$, the attacker can guarantee himself at least a utility of $\frac{1}{4} \cdot \alpha \cdot T$. We will show below that the attacker's utility for any attack duration $t \in [\xi T, \infty)$ is at most $\frac{1}{4} \cdot \alpha \cdot T$.

Hence, the attacker has an optimal attack duration $t \leq \xi T$ (either $t = T/2$ or a different t). By the assumption that $\mathbb{P}[\xi T \leq B] = 1$ and Proposition 2.1, $F(\xi T) = p \cdot \xi T$. Using the concavity of F , this implies that $F(t) = p \cdot t$ for all $t \leq \xi T$. Thus, whichever such t is optimal for the attacker, Corollary 3.4 implies that F is optimal for the defender, and furthermore, that $t = T/2$ is optimal for the attacker after all.

It remains to prove the upper bound for $t \geq \xi T$. For any $t \geq \eta T$, the assumption that $B \leq \eta T$ with probability 1 implies that $F(\eta T) = 1$, and hence a utility of 0 for the attacker. So we focus on $t \in [\xi T, \eta T]$, and show that in this range, the maximum utility of the attacker is at most $\alpha \cdot T/4$.

We proved above that $F(\xi T) = \xi$, and by the assumption $\mathbb{P}[B \leq \eta T] = 1$, we get that $F(\eta T) = 1$. Since F is concave by Proposition 2.1, for $t \in [\xi T, \eta T]$, F is bounded below by the line connecting $(\xi T, \xi)$ and $(\eta T, 1)$, so

$$F(t) \geq \xi + \frac{t - \xi T}{(\eta - \xi)T} \cdot (1 - \xi) = \xi \cdot \frac{\eta - 1}{\eta - \xi} + \frac{1 - \xi}{(\eta - \xi)T} \cdot t = \frac{\eta - 1}{\eta} + \frac{1}{\eta^2} \cdot \frac{t}{T}.$$

Hence, the attacker's utility is upper-bounded by $\alpha \cdot t \cdot (1 - F(t)) \leq \alpha \cdot t \cdot \left(\frac{1}{\eta} - \frac{1}{\eta^2} \cdot \frac{t}{T}\right)$. This is maximized at $t^* = \frac{\eta T}{2}$, so we obtain that $\alpha \cdot t \cdot (1 - F(t)) \leq \alpha \cdot t^* \cdot \left(\frac{1}{\eta} - \frac{1}{\eta^2} \cdot \frac{t^*}{T}\right) = \alpha \cdot \frac{T}{4}$. ■

Proof of Theorem 3.1. To complete the proof of Theorem 3.1, we now consider multiple targets i . By the assumptions of the theorem and Proposition 3.5, against the proposed class of random sequences, the attacker can obtain utility at most $\frac{1}{4}$, regardless of which target i he attacks and for how long; this follows by substituting $T_i = 1/p_i = 1/\alpha_i$.

We will show that no shift-invariant mixed defender schedule (now considered over the non-negative real numbers) can achieve an expected attacker payoff strictly smaller than $\frac{1}{4}$. Focus on a shift-invariant mixed defender schedule Λ . By Lemma B.2, we may assume that Λ is canonical.

Fix some index i such that $\alpha_i/p_i \geq 1$. Such an index must exist because $\sum_i \alpha_i = 1$ and $\sum_i p_i \leq 1$. Because we assumed that $\alpha_i \leq \frac{1}{2}$ for all i , this also implies that $p_i \leq \frac{1}{2}$.

By Proposition 3.2, attacking target i for $t = T_i/2$ units of time, the attacker can guarantee himself a utility of at least

$$\alpha_i \cdot \frac{1 - F_i(0)}{4} \cdot T_i = \alpha_i \cdot \frac{(1 - F_i(0))^2}{4(p_i - F_i(0))} \stackrel{p_i \leq \frac{1}{2}}{\geq} \alpha_i \cdot \frac{1}{4p_i} \geq \frac{1}{4},$$

where the final inequality followed because the chosen index i satisfied $\alpha_i/p_i \geq 1$. Hence, the attacker can guarantee himself a payoff of at least $\frac{1}{4}$ against any mixed defender schedule, proving optimality of the proposed class of random sequences.

Finally, we show that this applies to 2-quasi-regular random sequences. Assume that there exists a b such that $\mathbb{P}[b \leq B_i \leq 2b] = 1$, and define $\eta_i = 2b/T_i$. First, this definition directly implies that $B_i \leq \eta_i T_i$ with probability 1. Second, because $B_i \leq 2b$ with probability 1, we get that $T_i \leq 2b$, and hence $\frac{\eta_i}{\eta_i+1} T_i = \frac{2b}{2b/T_i+1} \leq b$. Hence, the fact that $B_i \geq b$ with probability 1 implies that $B_i \geq \frac{\eta_i}{\eta_i+1} T_i$ with probability 1, completing the proof. ■

4 An Optimal Defender Strategy

In this section, we present Algorithm 1, constructing a 2-quasi-regular random sequence. Such a random sequence is optimal for the defender by Theorem 3.1. We begin with the high-level algorithm, and fill in the details of the key steps below.

Algorithm 1 An optimal schedule for the defender

- 1: Let $p_i = \alpha_i$ for all i .
 - 2: For each i , let m_i be such that $2^{-m_i} \leq p_i < 2^{1-m_i}$. Let $I_i = [2^{-m_i}, 2^{1-m_i}]$.
 - 3: Use the algorithm from the proof of Lemma 4.3 for \mathbf{p} and the I_i to randomly round \mathbf{p} to a probability vector \mathbf{q} , such that all but at most one index i have $q_i = 2^{-m_i}$ or $q_i = 2^{1-m_i}$.
 - 4: Use the algorithm from the proof of Lemma 4.5 to produce a periodic sequence s .
 - 5: Return the random sequence obtained by choosing a uniform random shift of s .
-

Notice that the sequence produced by Algorithm 1 is shift-invariant by construction, but not ergodic, since it randomizes over different shift-invariant distributions.

Theorem 4.1 *The random sequence generated by Algorithm 1 is 2-quasi-regular, and hence optimal for the defender.*

We begin with a simple technical lemma.

Lemma 4.2 *Let S be a multiset of powers of 2, such that $\max_{p \in S} p \leq 2^{-k} \leq \sum_{p \in S} p$. Then, there exists a submultiset $T \subseteq S$ with $\sum_{p \in T} p = 2^{-k}$.*

Proof. We prove this claim by induction on $|S|$. The claim is trivial for $|S| = 1$. Consider $|S| \geq 2$, and distinguish two cases.

1. If S contains two copies of some number $p < 2^{-k}$, then construct S' by replacing these two copies with $p' = 2p$. By induction hypothesis, S' contains a subset T' adding up to 2^{-k} . If T' contained the newly constructed element p' , then replace it with the two copies of p . In either case, we have the desired set $T \subseteq S$.
2. Otherwise, S contains at most one copy of each number $p \leq 2^{-k}$. If S did not contain 2^{-k} , then $\sum_{p \in S} p < \sum_{i=1}^{\infty} 2^{-(k+i)} = 2^{-k}$, contradicting the assumptions of the lemma. Hence, S contains 2^{-k} , and the singleton set of that number is the desired subset. ■

Lemma 4.3 *Let $\mathbf{p} = (p_1, p_2, \dots, p_n)$ be a probability distribution. For each i , let $I_i = [\ell_i, r_i] \ni p_i$ be an interval. Then, there exists a distribution \mathcal{D} over probability distributions $\mathbf{q} = (q_1, q_2, \dots, q_n)$ such that:*

1. $\mathbb{E}[q_i] = p_i$ for all i ,
2. $q_i \in I_i$ for all \mathbf{q} in the support of \mathcal{D} , and
3. For each \mathbf{q} in the support of \mathcal{D} , all but at most one of the q_i are equal to ℓ_i or r_i .

Proof. We will give a randomized “rounding” procedure that starts with \mathbf{p} and produces a \mathbf{q} , satisfying all of the claimed properties, by making the p_i equal to ℓ_i or r_i one at a time. The randomized rounding bears similarity to dependent randomized rounding algorithms in the approximation algorithms literature (e.g., [7, 15, 30]), though we do not require concentration bounds, and allow one of the q_i to be an interior point of its interval. In the rounding, we always consider two indices i, j with $p_i = \ell_i + \epsilon_i, p_j = \ell_j + \epsilon_j$, such that $0 < \epsilon_i < r_i - \ell_i, 0 < \epsilon_j < r_j - \ell_j$. (That is, neither p_i nor p_j is on the boundary of its interval.) We probabilistically replace them with p'_i, p'_j , such that all of the following hold:

- At least one of p'_i, p'_j is at the boundary of its interval.
- $\ell_i \leq p'_i \leq r_i$ and $\ell_j \leq p'_j \leq r_j$.
- $p'_i + p'_j = p_i + p_j$.
- $\mathbb{E}[p'_i] = p_i$ and $\mathbb{E}[p'_j] = p_j$.

The rounding terminates when there is at most one p_i that is not at the boundary of its interval; let \mathbf{q} be the vector of probabilities at that point. By iterating expectations, we obtain that $\mathbb{E}[q_i] = p_i$ for all i . The upper and lower bounds on q_i are maintained inductively, and the termination condition ensures the third claimed property of \mathbf{q} .

So consider arbitrary p_i, p_j as above. Let $\delta_i = \min(\epsilon_i, r_j - \ell_j - \epsilon_j)$ and $\delta_j = \min(\epsilon_j, r_i - \ell_i - \epsilon_i)$. With probability $\frac{\delta_j}{\delta_i + \delta_j}$, round p_i to $p'_i = p_i - \delta_i$ and p_j to $p'_j = p_j + \delta_j$. With probability $1 - \frac{\delta_j}{\delta_i + \delta_j} = \frac{\delta_i}{\delta_i + \delta_j}$, round p_i to $p'_i = p_i + \delta_j$ and p_j to $p'_j = p_j - \delta_j$.

First, it is clear that $p'_i + p'_j = p_i + p_j$. Also, by definition of δ_i, δ_j , we get that $\ell_i \leq p'_i \leq r_i$ and $\ell_j \leq p'_j \leq r_j$. If we round according to the first case, then $p'_i = p_i - \delta_i$ and $p'_j = p_j + \delta_j$. If $\delta_i = \epsilon_i$, then we get that $p'_i = \ell_i$, while if $\delta_i = r_j - \ell_j - \epsilon_j$, then $p'_i = \ell_j + \epsilon_j + (r_j - \ell_j - \epsilon_j) = r_j$. Calculations are similar in the other case. Finally, $\mathbb{E}[p'_i] = \frac{\delta_j}{\delta_i + \delta_j} \cdot (p_i - \delta_i) + \frac{\delta_i}{\delta_i + \delta_j} \cdot (p_i + \delta_j) = p_i$. Hence, all the claimed properties hold in each step. \blacksquare

As a first step towards a 2-quasi-regular random sequence, we consider the case of probability vectors in which all probabilities are powers of 2.⁶

Lemma 4.4 *Assume that the probability vector \mathbf{p} is such that each $p_i = 2^{-m_i}$ is a power of 2. Then, there exists a regular sequence for \mathbf{p} .*

⁶Lemma 4.4 generalizes to powers of any integer, and in fact to any probabilities p_i such that for any i, j , we have that $p_i | p_j$ or $p_j | p_i$. The existence of a schedule for powers of 2 (and the generalization) has been previously observed in the context of the Pinwheel Problem in [18].

Proof. We will prove this claim by induction on the number of targets. If we have a single target, then its probability must be 1, so it is visited at intervals of 1 and we set s to be the constant sequence. Otherwise, the maximum probability of any target is $\frac{1}{2}$, and the sum of all probabilities is 1. Lemma 4.2 therefore guarantees the existence of a subset S whose probabilities add up to $\frac{1}{2}$.

Consider instances obtained from S and \bar{S} by scaling up all probabilities by a factor of 2, resulting in $p'_i = 2p_i$. By induction hypothesis, each of those instances can be scheduled such that each target i is visited every $1/p'_i = 1/(2p_i)$ time steps. Now alternate between the two sequences. In this new sequence, each target i is visited every $2/p'_i = 1/p_i$ steps, as desired. ■

Next, we show that sufficiently good sequences can also be achieved when at most one of the probabilities is not a power of 2.

Lemma 4.5 *Assume that the probability vector \mathbf{p} is such that each $p_i = 2^{-m_i}$ is a power of 2, except for (possibly) $p_1 = 2^{-m_1} - \epsilon$, with $0 \leq \epsilon < 2^{-(m_1+1)}$. Then, there exists a (non-random) periodic sequence s with the following properties:*

1. *The time between consecutive visits to target $i > 1$ is always exactly $1/p_i$.*
2. *The time between consecutive visits to target 1 is always either 2^{m_1} or 2^{m_1+1} .*
3. *The frequency of target i is p_i for all i .*

Proof. We distinguish two cases⁷:

1. If $p_1 \leq \frac{1}{2}$, then by Lemma 4.2, there exists a subset $S \subseteq \{2, \dots, n\}$ with $\sum_{i \in S} p_i = \frac{1}{2}$. Schedule the targets in S regularly in the odd time slots, and set $p'_i = 2p_i$ for all $i \notin S$. The p'_i satisfy the conditions of the lemma, so we inductively find a schedule for all the $i \notin S$ satisfying the conclusion of the lemma, then stretch this schedule by a factor of 2 and use all the even time slots for it.
2. If $p_1 > \frac{1}{2}$, then we schedule target 1 in all odd time slots, and set $p'_1 = 2(p_1 - \frac{1}{2})$ and $p'_i = 2p_i$ for all $i \geq 2$. The new p'_i sum to 1 and satisfy the conditions of the lemma. So we inductively find a schedule for the frequencies p'_i , stretch it by a factor of 2, and use all the even time slots for it. Notice that item 1 is scheduled at distance at most 2, and at least 1, while all other items are scheduled regularly.

That the resulting schedule is periodic is seen inductively over the applications of the two cases. ■

Proof of Theorem 4.1. Consider any target i . The rounding of Lemma 4.3 guarantees that $2^{-m_i} \leq q_i \leq 2^{1-m_i}$. Therefore, the algorithm of Lemma 4.5 produces a random sequence $\Sigma_{\mathbf{q}}$ in which the time intervals between consecutive occurrences of target i lie between 2^{m_i-1} and 2^{m_i} . Thus to verify that $\Sigma_{\mathbf{q}}$ is 2-quasi-regular it remains to show that the expected density of each target is equal to α_i . But this is guaranteed by (1) in Lemma 4.3. The optimality of $\Sigma_{\mathbf{q}}$ now follows from Theorem 3.1. ■

The second part of Theorem 4.1 shows that 2-quasi regular random sequences exist; here, we remark that this result cannot be improved, in the following sense (proved in Appendix C).

⁷We would like to thank an anonymous reviewer for suggesting this elegant proof.

Proposition 4.6 *Let $n = 3$ and $\alpha = (1/2, 1/3, 1/6)$. Then, for every $\epsilon > 0$, there are no $(2 - \epsilon)$ -quasi-regular random sequences.*

5 Golden Ratio Scheduling

In this section, we present a very simple ergodic random sequence. The associated schedule may in general be suboptimal, but we prove that it is within less than 0.6% of optimal.

Let $\varphi = \frac{1}{2}(1 + \sqrt{5})$ denote the Golden Ratio, solving $\varphi^2 = \varphi + 1$. Given a desired frequency vector \mathbf{p} (which will equal the targets' values, $p_i = \alpha_i$), we identify the unit circle with $[0, 1)$, and equip it with addition modulo 1. We define the function $h: [0, 1) \rightarrow \{1, \dots, n\}$ via $h^{-1}(i) = [\sum_{i' < i} p_{i'}, p_i + \sum_{i' < i} p_{i'})$; that is, we assign consecutive intervals of length p_i for the targets i .

Algorithm 2 The Golden Ratio Schedule

- 1: Let λ be uniformly random in $[0, 1)$.
 - 2: **for** $t = 0, 1, 2, \dots$ **do**
 - 3: In step t , set $s_t = h((\lambda + \varphi t) \bmod 1)$.
-

We can think of advancing a “dial” by φ (or $\varphi - 1$) at each step, and visiting the target whose interval the dial falls into. This algorithm is nearly identical to one previously proposed for hashing [22, pp. 510,511,543] and broadcast channel sharing [21, 26]. While the algorithm is simple, as stated, it seems to require precise arithmetic with real numbers. This issue is discussed in more detail in Appendix E.

That Algorithm 2 returns an ergodic random sequence follows from the classical fact that the action on the interval by an irrational rotation is ergodic. Our main theorem in this section is the following:

Theorem 5.1 *The Golden Ratio algorithm is a $\frac{2966-1290\sqrt{5}}{81} \approx 1.00583$ approximation for the defender.*

The underlying reason that this schedule performs so well, and the reason for choosing specifically the Golden Ratio, is related to the hardness of diophantine approximation of the Golden Ratio: it is an irrational number that is hardest to approximate by rational numbers (see, e.g., [16]).

Our analysis relies heavily on various properties of Fibonacci numbers. We denote the k^{th} Fibonacci number by f_k , indexed as $f_0 = 0, f_1 = 1$ and $f_{k+2} = f_k + f_{k+1}$. The following basic facts about Fibonacci numbers are well-known, and easily proved directly or by induction.

Lemma 5.2 1. *For any k , we have that $f_{k+2}f_k - f_{k+1}^2 = (-1)^{k+1}$.*

2. *For any k , we have that $f_k = \frac{\varphi^k - (-1/\varphi)^k}{\sqrt{5}}$.*

3. *For any odd k , we have that $f_{k+1}/f_k < \varphi$.*

4. *For any even k , we have that $f_{k+1}/f_k > \varphi$.*

To prove Theorem 5.1, we analyze the distribution of B_i for any target i . The proof of Theorem 5.1 consists of two parts. First, Theorem 5.3 precisely characterizes the distribution of B_i for every target i , i.e., it characterizes exactly, for each target i and τ , how frequently a visit to target

i is followed by another visit τ steps later. (We call such a τ a *return time*.) As a second part, we characterize the attacker's best response against this distribution, and calculate its cost to the defender.

Theorem 5.3 (Slater [28]) *Assume that $p \leq \frac{1}{2}$. Let k be smallest such that*

$$|f_{k+1} - \varphi f_k| \leq \varphi p \quad (4)$$

Then, the distribution of return times is

$$\begin{aligned} \mathbb{P}[B_i = f_{k+1}] &= f_{k+1} \cdot \left(p - (1/\varphi)^{k+1} \right), \\ \mathbb{P}[B_i = f_{k+2}] &= f_{k+2} \cdot \left(p - (1/\varphi)^{k+2} \right), \\ \mathbb{P}[B_i = f_{k+3}] &= f_{k+3} \cdot \left(-p + (1/\varphi)^k \right), \\ \mathbb{P}[B_i = t] &= 0 \text{ for all other } t. \end{aligned} \quad (5)$$

Theorem 5.3 shows, remarkably, that for each possible p , there are at most three possible return times, and they are three consecutive Fibonacci numbers. Theorem 5.3 is a special case of a theorem of Slater [28, Theorem 4] (see also [29]), which characterizes the distribution when the Golden Ratio φ is replaced by an arbitrary real number. We give a self-contained proof for the simpler case of the Golden Ratio in Appendix D.

As a direct corollary of Theorem 5.3, we obtain an upper bound on the quasi-regularity of the Golden Ratio schedule.

Corollary 5.4 *The Golden Ratio schedule is 3-quasi-regular. If $p_i \leq 1 - 1/\varphi$ for all i , then it is 8/3-quasi-regular. As the frequencies $p_i \rightarrow 0$, the regularity guarantee improves to φ^2 -quasi-regular.*

Proof. Consider one target i with desired frequency p_i , and define k as in Theorem 5.3. The schedule is at worst (f_{k+3}/f_{k+1}) -quasi-regular. For all k , we have the bound $f_{k+3}/f_{k+1} \leq 3$. If $p_i < 1 - 1/\varphi$, then $k \geq 2$, and for $k \geq 2$, the ratio f_{k+3}/f_{k+1} is upper-bounded by 8/3, converging to φ^2 as $k \rightarrow \infty$. ■

5.1 The Optimal Attacker Response to 3-Point Distributions

Next, we characterize the optimal attacker response to defender strategies in which the return time distribution to target i is supported on three points $1 \leq x_1 < x_2 < x_3$ only, so that with probability one, $B_i \in \{x_1, x_2, x_3\}$. In the remainder of this section, we omit the subscript i , as we will always analyze one target only.

Recall that $T = 1/F(1)$, and let $q_j = \mathbb{P}[B = x_j] \cdot \frac{T}{x_j}$. Note that $\sum_j q_j x_j = \sum_j \mathbb{P}[B = x_j] \cdot T = T$. Informally (but in a sense that can be formalized), the q_j 's are the return time probabilities from the point of view of the defender, rather than those of the attacker at a fixed time 0. (The attacker's distribution of B_i oversamples long return times, compared to the defender's distribution.) The attacker's response and utility are summarized by the following lemma.

Lemma 5.5 *Assume that $x_2 \leq 2x_1, x_3 \leq 2x_2$ (as is the case with Fibonacci numbers). Let $u_1^* = \frac{1}{4} \cdot \alpha \cdot T$, and $u_2^* = \frac{1}{4} \cdot \alpha \cdot \frac{(T - q_1 x_1)^2}{T(1 - q_1)}$. Then, against the given three-point distribution, the attacker's utility is at most $\max(u_1^*, u_2^*)$. Against any distribution with expected defender absence T , the attacker's utility is at least u_1^* .*

Proof. The lower bound of u_1^* is simply the statement of Corollary 3.3. So we focus on the upper bound of $\max(u_1^*, u_2^*)$ for the remainder of the proof.

From the attacker's perspective, when arriving at a target, by Equations (5) and (3), the CDF of the distribution of the defender's next return time is

$$F(t) = \begin{cases} \frac{1}{T} \cdot t & \text{for } t \leq x_1, \\ \frac{1}{T} \cdot (t(1 - q_1) + q_1 x_1) & \text{for } x_1 \leq t \leq x_2, \\ \frac{1}{T} \cdot (t(1 - q_1 - q_2) + q_1 x_1 + q_2 x_2) & \text{for } x_2 \leq t \leq x_3, \\ 1 & \text{for } t \geq x_3. \end{cases}$$

Since the attacker's utility for waiting for t steps is $t(1 - F(t))$, $t \geq x_3$ cannot be optimal for him. By taking derivatives with respect to t , we obtain the following local optima for the functions in the remaining three cases:

$$t_1^* = \frac{T}{2} \qquad t_2^* = \frac{T - q_1 x_1}{2(1 - q_1)} \qquad t_3^* = \frac{T - q_1 x_1 - q_2 x_2}{2(1 - q_1 - q_2)} = \frac{x_3}{2}.$$

These are all local maxima because the functions are concave. Under the assumption of the theorem, $t_3^* = x_3/2 \leq x_2$. Therefore, t_3^* does not lie in the interval it optimized for, and can never be optimal. As a result, the attacker's best response⁸ will always be t_1^* or t_2^* . The attacker's utility for these two attack times will be

$$\begin{aligned} u_1^* &= \alpha \cdot t_1^* \cdot (1 - F(t_1^*)) = \frac{\alpha T}{2} \cdot \frac{1}{2} = \frac{\alpha T}{4}, \\ u_2^* &= \alpha \cdot t_2^* \cdot (1 - F(t_2^*)) = \alpha \cdot \frac{T - q_1 x_1}{2(1 - q_1)} \cdot \left(1 - \frac{1}{T} \cdot \left(\frac{T - q_1 x_1}{2} + q_1 x_1\right)\right) \\ &= \frac{\alpha}{4} \cdot \frac{(T - q_1 x_1)^2}{T(1 - q_1)} = \frac{\alpha}{4} \cdot \frac{(q_2 x_2 + q_3 x_3)^2}{T(q_2 + q_3)}. \end{aligned}$$

The attacker's utility will thus be at most the maximum of u_1^*, u_2^* . ■

5.2 The Attacker's Response to the Golden Ratio Schedule

Proof of Theorem 5.1. Applying Lemma 5.5, it is our goal to upper-bound $\frac{\max(u_1^*, u_2^*)}{u_1^*}$. If $\max(u_2^*, u_1^*) = u_1^*$, the approximation ratio is 1; hence, it suffices to upper-bound u_2^*/u_1^* .

In applying Lemma 5.5, we have $T = 1/\alpha$, $x_1 = f_{k+1}$, $x_2 = f_{k+2}$, $x_3 = f_{k+3}$, and the q_j are given via Equation (5) and $q_j = \mathbb{P}[B_i = x_j] \cdot \frac{T}{x_j}$. Then, u_2^*/u_1^* can be written as follows:

⁸It is possible that one of t_1^*, t_2^* also lies outside its interval. But our goal here is only to derive an upper bound.

$$\begin{aligned}
\frac{(T - q_1 x_1)^2}{T^2(1 - q_1)} &= \alpha^2 \cdot \frac{(q_2 f_{k+2} + q_3 f_{k+3})^2}{q_2 + q_3} \\
&= \alpha \cdot \frac{((\alpha - (1/\varphi)^{k+2}) \cdot f_{k+2} + (-\alpha + (1/\varphi)^k) \cdot f_{k+3})^2}{(\alpha - (1/\varphi)^{k+2}) + (-\alpha + (1/\varphi)^k)} \\
&= \alpha \cdot \frac{((1/\varphi)^{k+2} \cdot (\varphi^2 - 1) \cdot f_{k+2} + (-\alpha + (1/\varphi)^k) \cdot f_{k+1})^2}{(1/\varphi)^{k+2} \cdot (\varphi^2 - 1)} \\
&= \alpha \cdot (1/\varphi)^{k+1} \cdot \left(f_{k+2} + \varphi f_{k+1} - \alpha \cdot \varphi^{k+1} \cdot f_{k+1} \right)^2. \tag{6}
\end{aligned}$$

Treating everything except α as a constant, the approximation ratio is thus of the form $g(\alpha) = a\alpha \cdot (c - b\alpha)^2$. g has a local maximum of $4ac^3/27b$ at $\alpha = c/(3b)$, a local minimum of 0 at $\alpha = c/b$, and goes to infinity as $\alpha \rightarrow \infty$. Thus, the two candidates for α that we need to check are (1) the largest α that is possible for a given k , and (2) the value $\alpha = c/(3b)$.

We therefore next calculate the largest possible α for a given k . By recalling the definition of k from Equation (4) (smallest such that $|f_{k+1}/\varphi - f_k| \leq \alpha$), and using Lemma D.1, we can solve for α to determine the range in which we obtain a particular k , giving us that $\alpha \in [(1/\varphi)^{k+1}, (1/\varphi)^k]$.

1. If we substitute the upper bound $\alpha = (1/\varphi)^k$, Equation (6) simplifies to

$$\begin{aligned}
&\alpha \cdot (1/\varphi)^{k+1} \cdot \left(f_{k+2} + \varphi f_{k+1} - (1/\varphi)^k \cdot \varphi^{k+1} \cdot f_{k+1} \right)^2 \\
&= (1/\varphi)^k \cdot (1/\varphi)^{k+1} \cdot (f_{k+2} + \varphi f_{k+1} - \varphi f_{k+1})^2 \\
&= \frac{1}{5} \cdot (1/\varphi)^{2k+1} \cdot \left(\varphi^{2k+4} - 2\varphi^{k+2}(-1/\varphi)^{k+2} + (-1/\varphi)^{2k+4} \right) \\
&= \frac{1}{5} \cdot (1/\varphi)^{2k+1} \cdot \left(\varphi^{2k+4} - 2(-1)^k + (1/\varphi)^{2k+4} \right) \\
&\leq \frac{1}{5} \cdot (\varphi^3 + 3/\varphi^{2k+1}) \leq \frac{1}{5} \cdot (\varphi^3 + 3/\varphi^3) < 1.
\end{aligned}$$

This shows that the attacker's utility cannot be maximized by waiting for more than x_1 steps when α is as large as it can be for a given k .

2. Next, we investigate the local maximum⁹ of Equation (6). Substituting $a = (1/\varphi)^{k+1}$, $b = \varphi^{k+1} \cdot f_{k+1}$, and $c = f_{k+2} + \varphi f_{k+1}$, the approximation ratio is

$$\begin{aligned}
&\frac{4}{27} \cdot (1/\varphi)^{2k+2} \cdot \frac{(f_{k+2} + \varphi f_{k+1})^3}{f_{k+1}} \\
&= \frac{4}{27} \cdot \frac{1}{5} \cdot (1/\varphi)^{2k+2} \cdot \frac{(\varphi^{k+2} - (-1/\varphi)^{k+2} + \varphi \cdot \varphi^{k+1} - \varphi \cdot (-1/\varphi)^{k+1})^3}{\varphi^{k+1} - (-1/\varphi)^{k+1}} \\
&= \frac{4}{27} \cdot \frac{1}{5} \cdot (1/\varphi)^{2k+2} \cdot \frac{(2\varphi^{k+2} - (-1/\varphi)^{k+1})^3}{\varphi^{k+1} - (-1/\varphi)^{k+1}}.
\end{aligned}$$

⁹This local maximum is indeed always a feasible choice for α for a given k , but since we are only interested in an upper bound, we omit the feasibility proof.

We will approximate the function $\frac{(2\varphi^{k+2}-(-1/\varphi)^{k+1})^3}{\varphi^{k+1}-(-1/\varphi)^{k+1}}$ by $8\varphi^{2k+5}$, its highest-order term. We therefore consider $\frac{(2\varphi^{k+2}-(-1/\varphi)^{k+1})^3}{\varphi^{k+1}-(-1/\varphi)^{k+1}}/(8\varphi^{2k+5})$. When k is even, this ratio is always upper-bounded by 1 (and increasing in k , converging to 1). When k is odd, this ratio is *lower-bounded* by 1, and decreasing in k , also converging to 1. Thus, it is maximized among feasible values of k for $k = 3$, where it equals $\frac{8\varphi^{15}-12\varphi^6+6\varphi^{-3}-\varphi^{-12}}{8\varphi^{15}-8\varphi^7}$. Overall, we get an upper bound on the attacker's utility of

$$\begin{aligned} & \frac{1}{5} \cdot \frac{4}{27} \cdot \frac{8\varphi^{15} - 12\varphi^6 + 6\varphi^{-3} - \varphi^{-12}}{8\varphi^{15} - 8\varphi^7} \cdot (1/\varphi)^{2k+2} \cdot (8\varphi^{2k+5}) \\ &= \frac{\varphi^3}{5} \cdot \frac{4}{27} \cdot \frac{8\varphi^{15} - 12\varphi^6 + 6\varphi^{-3} - \varphi^{-12}}{\varphi^{15} - \varphi^7}. \end{aligned}$$

To evaluate this ratio, we can repeatedly apply the fact that $\varphi^2 = 1 + \varphi$, then substitute that $\varphi = \frac{1+\sqrt{5}}{2}$, make the denominator rational, and cancel out common factors. This shows that $\frac{\varphi^3}{5} \cdot \frac{4}{27} \cdot \frac{8\varphi^{15}-12\varphi^6+6\varphi^{-3}-\varphi^{-12}}{\varphi^{15}-\varphi^7} = \frac{2966-1290\sqrt{5}}{81} \approx 1.00583$, completing the proof. ■

Remark 5.6 The analysis of Theorem 5.1 is actually tight. By choosing $k = 3$, $b = \varphi^4 \cdot f_4$, and $c = f_5 + \varphi f_4$, we obtain that the worst-case value of a target is $\alpha = c/(3b) = \frac{23}{18} - \frac{\sqrt{5}}{2} \approx 0.1597$. Substituting $k = 3$ into the attacker's utility in Case (2) (before the lower bound is applied) gives us exactly a ratio of $\frac{2966-1290\sqrt{5}}{81} \approx 1.00583$.

6 Scheduling via Matching

In this section, we show that for every $\epsilon > 0$, if the individual probabilities α_i are small enough (as a function of ϵ), then $(1 + \epsilon)$ -quasi-regular ergodic, periodic schedules exist.

In order to obtain a periodic strategy, it is clearly necessary for all target values α_i (equaling the visit frequencies) to be rational. Write $\alpha_i = a_i/b_i$, and let $M = \text{lcm}(b_1, \dots, b_n)$. Our algorithm is based on embedding M slots for visits evenly on the unit circle, and matching them with targets to visit. We identify the circle with the interval $[0, 1]$ and use the distance $d(x, y) = \min(|x - y|, 1 - |x - y|)$.

Theorem 6.1 *Fix $\epsilon > 0$, and assume that $\alpha_i \leq \frac{\epsilon}{4+2\epsilon} \cdot \sqrt{\frac{2}{n \log M}}$ for all i . Then, Algorithm 3 succeeds with high probability. Whenever Algorithm 3 succeeds, it produces a $(1 + \epsilon)$ -quasi-regular (and hence defender-optimal) sequence.*

We begin by proving the second part of the theorem. First, in a perfect matching, exactly A_i of the M slots, i.e., an α_i fraction, are scheduled for target i , giving that $p_i = a_i/b_i = \alpha_i$. Thus, $T_i = 1/\alpha_i$.

If t is matched with (i, j) , by definition of the edges, $d(y_{i,j}, t/M) \leq \delta$. Consider two occurrences j, j' of target i , and let t, t' be the slots they are matched to. Then, by triangle inequality,

$$d(t/M, t'/M) \geq d(y_{i,j}, y_{i,j'}) - 2\delta \geq \frac{1}{A_i} - 2\delta. \quad (7)$$

On the other hand, specifically for consecutive occurrences of target i , i.e., the slots matched to $y_{i,j}$ and $y_{i,j+1}$, we get

Algorithm 3 A matching-based algorithm for a periodic defender strategy

- 1: **for** each target i **do**
 - 2: Let $\theta_i \in [0, 1]$ independently uniformly at random.
 - 3: Let $A_i = M \cdot \alpha_i$.
 - 4: For $j = 0, \dots, A_i - 1$, let $y_{i,j} = (\theta_i + j/A_i) \bmod 1$.
 - 5: Let $\delta = \frac{1}{M} \cdot \sqrt{\frac{n \log M}{2}}$.
 - 6: Let $Z = \{0, 1, \dots, M-1\}$ be the set of *slots* and let $Y = \{(i, j) \mid 0 \leq j < A_i\}$. Define a bipartite graph G on $Z \cup Y$ by including an edge between $t \in Z$ and $(i, j) \in Y$ iff $d(y_{i,j}, t/M) \leq \delta$.
 - 7: **if** G contains a perfect matching \mathcal{M} **then**
 - 8: Define a sequence s with period M as follows: For each time t , set s_t to be the (unique) target i such that t is matched with (i, j) in \mathcal{M} for some j .
 - 9: **else**
 - 10: Start from the beginning.
-

$$d(t/M, t'/M) \leq d(y_{i,j}, y_{i,j'}) + 2\delta \leq \frac{1}{A_i} + 2\delta. \quad (8)$$

Using that $\delta = \frac{1}{M} \cdot \sqrt{\frac{n \log M}{2}}$, as well as $\frac{1}{A_i} = \frac{1}{M \cdot \alpha_i}$ and $\alpha_i \leq \frac{\epsilon}{4+2\epsilon} \cdot \sqrt{\frac{2}{n \log M}}$, we obtain that

$$\frac{\frac{1}{A_i} + 2\delta}{\frac{1}{A_i} - 2\delta} = \frac{\frac{1}{\alpha_i} + 2\sqrt{\frac{n \log M}{2}}}{\frac{1}{\alpha_i} - 2\sqrt{\frac{n \log M}{2}}} \leq \frac{\frac{4+2\epsilon}{\epsilon} \cdot \sqrt{\frac{n \log M}{2}} + 2\sqrt{\frac{n \log M}{2}}}{\frac{4+2\epsilon}{\epsilon} \cdot \sqrt{\frac{n \log M}{2}} - 2\sqrt{\frac{n \log M}{2}}} = \frac{\frac{4+2\epsilon}{\epsilon} + 2}{\frac{4+2\epsilon}{\epsilon} - 2} = 1 + \epsilon,$$

proving that the resulting sequence is $(1 + \epsilon)$ -quasi-regular.

To complete the proof, it remains to show that with high probability, the graph G contains a perfect matching. We will prove this using Hall's Theorem and a direct application of the Hoeffding Bound:

Lemma 6.2 (Hoeffding Bound) *Let X_i be independent random variables such that $a_i \leq X_i \leq b_i$ with probability 1. Let $X = \sum_i X_i$. Then, for all $t > 0$,*

$$\mathbb{P}[X < \mathbb{E}[X] - t], \mathbb{P}[X > \mathbb{E}[X] + t] < e^{-\frac{2t^2}{\sum_i (b_i - a_i)^2}}.$$

To establish the Hall condition of G , we begin with intervals $W \subseteq Z$ of slots, and then use the bounds for intervals to derive the condition for arbitrary sets of slots. A similar style of proof was used by Tijdeman [32] to construct a schedule with somewhat different specific combinatorial properties.

For any set $W \subseteq Z$ of slots, let $\Gamma(W)$ denote the neighborhood of W in G . Fix an interval $W = \{\ell, \ell + 1, \dots, r - 1\} \subseteq [0, M)$ with ℓ, r integers. Let the random variable $X_W = |\Gamma(W)|$ denote the number of neighbors in G of slots in the interval W . For each target i , let $X_{W,i}$ be the number of j such that $(i, j) \in \Gamma(W)$. Then, $X_W = \sum_i X_{W,i}$, and the $X_{W,i}$ are independent.

Lemma 6.3 *Fix a target i , and assume that $|W| \leq (1 - 2\delta)M$, and write $x_i = A_i \cdot (2\delta + (r - \ell)/M)$. Then, $\mathbb{E}[X_{W,i}] = x_i$ and $X_{W,i} \in \{\lfloor x_i \rfloor, \lfloor x_i \rfloor + 1\}$.*

Proof. For each slot $t \in W$, let J_t be the interval $[(t/M - \delta) \bmod 1, (t/M + \delta) \bmod 1]$. Then, (i, j) is adjacent to t iff $y_{i,j} \in J_t$. Define $J := \bigcup_{t \in W} J_t$; then, (i, j) is adjacent to a slot in W iff $y_{i,j} \in J$. Because $\delta = \frac{1}{M} \cdot \sqrt{\frac{n \log M}{2}} \geq 1/(2M)$, J is an interval, of the form $[(\ell/M - \delta) \bmod 1, (r/M + \delta) \bmod 1]$.

The length of the interval is $|J| = 2\delta + (r - \ell)/M$. Because each $y_{i,j}$ is uniformly random in $[0, 1]$, $\mathbb{E}[X_{W,i}] = A_i \cdot |J|$. Furthermore, because $d(y_{i,j}, y_{i,j+1}) = 1/A_i$, there can be no more than $1 + \lfloor \frac{|J|}{1/A_i} \rfloor$ pairs (i, j) with $y_{i,j} \in J$, and no fewer than $\lfloor \frac{|J|}{1/A_i} \rfloor$. Finally, note that $\frac{|J|}{1/A_i} = A_i(2\delta + (r - \ell)/M) = x_i$. ■

We use Lemma 6.3 to show that with high probability, G has a perfect matching.

Lemma 6.4 *Whenever $\alpha_i \leq \frac{\epsilon}{4+2\epsilon} \cdot \sqrt{\frac{2}{n \log M}}$ for all i , with probability at least $1 - 1/M^2$, G contains a perfect matching.*

Proof. First, we show that when the Hall condition holds for all intervals W of slots, it holds for all sets W . We prove this by induction on the number of disjoint intervals that W comprises. The base case of W being an interval is true by definition. For the induction step, suppose that $k \geq 2$ and $W = \bigcup_{j=1}^k W_j$, where the W_j are disjoint intervals.

If the neighborhoods of all the W_j are disjoint, then $|\Gamma(W)| = \sum_j |\Gamma(W_j)| \geq \sum_j |W_j| = |W|$, where the inequality was from the base case (intervals). Otherwise, w.l.o.g., $\Gamma(W_k) \cap \Gamma(W_{k-1}) \neq \emptyset$. Then, there exists an interval $I' \supset W_k \cup W_{k-1}$ with $\Gamma(I') = \Gamma(W_k) \cup \Gamma(W_{k-1})$. Let $W' = W \cup I'$. We get that $|\Gamma(W)| = |\Gamma(W')| \geq |W'| \geq |W|$, where the first inequality was by induction hypothesis (because W' has at least one less interval).

Next, we establish that the Hall Condition holds with high probability for all M^2 intervals. First, focus on one interval $W = [\ell, r)$, with $\ell, r \in \mathbb{N}$. If $|W| > (1 - 2\delta)M$, then $\Gamma(W)$ contains all pairs (i, j) , so the Hall Condition is satisfied. So focus on $|W| \leq (1 - 2\delta)M$. From Lemma 6.3, we get that

$$\mathbb{E}[X_W] = \sum_i A_i \cdot (2\delta + (r - \ell)/M) = 2\delta M + (r - \ell) = \sqrt{2n \log M} + (r - \ell).$$

Furthermore, X_W is the sum of independent random variables $X_{W,i}$ which each takes on one of two adjacent values. From the Hoeffding Bound (Lemma 6.2), we get that

$$\mathbb{P}[X_W < (r - \ell) + 2 \sum_i A_i \delta - \tau] < e^{-2\tau^2/n}.$$

Because $|W| = r - \ell$, choosing $\tau = 2 \sum_i A_i \delta = 2M\delta = \sqrt{2n \log M}$, we get that

$$|\Gamma(W)| = X_W \geq (r - \ell) + 2 \sum_i A_i \delta - \tau = r - \ell,$$

with probability at least $1 - e^{-4n \log M/n} = 1 - 1/M^4$.

Taking a union bound over all M^2 candidate intervals W , we obtain that the probability of having a perfect matching is at least $1 - 1/M^2$. Thus, with high probability, G contains a perfect matching. This completes the proof of Lemma 6.4 and thus also Theorem 6.1. ■

7 Future Work

Our work suggests a number of directions for future work. Most immediately, it suggests trying to find optimal *ergodic* schedules for all value vectors (not only those covered by Theorem 6.1). A promising approach toward this goal is to use the randomized rounding of Section 4, but re-round the probabilities every T steps, for some sufficiently large “epoch size” T . The difficulty with this approach is “stitching together” the schedules for different rounded frequencies at the boundary of epochs, without violating the conditions of Theorem 3.1.

Throughout, we assumed that no target had value more than the sum of all other targets’ values, i.e., $\alpha_i \leq \frac{1}{2}$ for all i . When this assumption is violated, the optimal schedule will wait at the highest-value target. In the specific case of two targets of values $\alpha_1 > \frac{1}{2}$ and $\alpha_2 = 1 - \alpha_1$, it is fairly straightforward to calculate that the wait time at target 1 is $2(\sqrt{\frac{\alpha_1}{\alpha_2}} - 1)$. We anticipate that a similar analysis will extend to more than two targets. The difficulty is that the waiting time at one target will result in qualitatively different schedules, likely to complicate the analysis.

We assumed here that the game is zero-sum. In general, the utilities of the attacker and defender may be different. A general treatment is likely quite difficult. One special case is motivated directly by the wildlife protection application, and appears quite amenable to analysis. Specifically, when a poacher kills animals (or chops down trees), even if the poacher is captured, the damage is not reversed. Thus, while the attacker’s utility is as before, the defender’s utility from visiting target i at time τ when the attacker intends to stay for t units of time is $-\alpha_i \cdot \min(\tau, t)$. Because the sum of utilities is thus not constant (and typically negative), the defender’s goal becomes more strongly that of deterring (rather than just capturing) the attacker. As a result, the ability to commit to a strategy first (i.e., treating the game as a *Stackelberg Game* [34, 10]) may now carry some advantage for the defender. One can show that in the model in this paragraph, whenever the attacker attacks target i for $t \leq T_i/2$ units of time, the defender’s utility is $-\frac{3}{2}U(F_i, (i, t))$. Since the optimal defender strategies of Section 4 and 6 ensure such a choice of t by the attacker, the algorithms in those sections are optimal in the non-zero sum model as well.

Among the other natural generalizations are the attacker’s (and defender’s) utility function and more complex constraints on the defender’s schedule. Throughout, we have assumed that the attacker’s utility grows linearly in the time spent at a target. The security game formulations studied in much of the prior work in the area [31] correspond to a step function at 0: when the attacker reaches an unprotected target, he immediately causes the maximum target-specific damage α_i (e.g., by blowing up the target). Other natural utility functions suggest themselves: if the resources to collect at targets are limited, the utility function would be linear with a cap. If a destructive attack takes a non-zero amount of time to set up, one obtains a step function at a time other than 0. The latter leads to a scheduling problem with a harder constraint on the inter-visit absence time from targets i — as in some of the prior security games literature, the defender may “sacrifice” some low-value targets to be able to fully protect the others.

The other natural generalization is to relax the assumption of uniform travel time between targets. If an arbitrary metric is defined between targets, the problem becomes significantly more complex: even if all targets have value 1, the attacker’s utility will be proportional to the cost of a minimum TSP tour, and thus the defender’s optimization problem is NP-hard. However, it is far from obvious how to adapt standard TSP approximation techniques to the general problem with non-uniform values: high-value targets should be visited more frequently, and TSP approximation algorithms are not suited to enforce constraints that these visits be spaced out over time.

As with TSP problems and past work on security games, a further natural generalization is to consider multiple defenders, as, e.g., in [23].

Acknowledgments

We would like to thank Omer Angel, Sami Assaf, Shaddin Dughmi, Fei Fang, Ron Graham, Bobby Kleinberg, Jim Propp, and Milind Tambe for useful discussions, and anonymous referees for useful feedback.

A Utility of an i.i.d. Defender

One of the most natural random sequences to consider is the *i.i.d.* one, in which at each step t , the defender visits target i with probability p_i , independent of any past choices. Intuitively, this strategy is suboptimal because it may visit a target i several times in close succession, or go for a long time without visiting target i . Here, we calculate the approximation ratio of this strategy, showing:

Proposition A.1 *The i.i.d. strategy is a $4/e$ -approximation for the defender, and this is tight.*

Proof. From the attacker’s viewpoint, the defender’s next arrival time at target i is the sum of two independent random variables $\text{geom}(p_i) + \text{unif}([0, 1])$. Given a t , the defender will return within at most t steps if and only if $\text{geom}(p_i) \leq \lfloor t \rfloor$ or $\text{geom}(p_i) = 1 + \lfloor t \rfloor$ and $\text{unif}([0, 1]) \leq (t \bmod 1)$. The two events are disjoint, the first one having probability $1 - (1 - p_i)^{\lfloor t \rfloor}$, and the second having probability $p_i \cdot (1 - p_i)^{\lfloor t \rfloor} \cdot (t \bmod 1)$. Hence, $F_i(t) = 1 - (1 - p_i \cdot (t \bmod 1)) \cdot (1 - p_i)^{\lfloor t \rfloor}$, and the attacker’s utility from attacking target i for t time units is

$$\alpha_i \cdot t \cdot (1 - F_i(t)) = \alpha_i \cdot t \cdot (1 - p_i \cdot (t \bmod 1)) \cdot (1 - p_i)^{\lfloor t \rfloor}.$$

Writing $t = x + k$ for an integer $k = \lfloor t \rfloor$ and $x = (t \bmod 1) \in [0, 1)$, a derivative test shows that the expression is monotone decreasing in x for any $k \geq 1$, whereas for $k = 0$, it has a local maximum at $x = \frac{1}{2p_i} \geq 1$. Because the latter is not feasible, we only need to consider the case $(t \bmod 1) = 0$ for the remainder, so the attacker’s utility simplifies to $\alpha_i \cdot t \cdot (1 - p_i)^{\lfloor t \rfloor}$.

Taking a derivative with respect to t and setting it to 0 gives us that the unique local extremum is at $t = \frac{-1}{\ln(1-p_i)}$, where the attacker’s utility is $\frac{p_i}{e \cdot \ln(1/(1-p_i))}$. This local extremum is a maximum because the attacker’s utility at $t = 0$ and $t = \infty$ is 0.

A derivative test and Taylor series bound shows that $\frac{p_i}{e \cdot \ln(1/(1-p_i))}$ is monotone decreasing in p_i , so it is maximized as $p_i \rightarrow 0$, where it converges to $1/e$. Notice that as $p_i \rightarrow 0$, there are infinitely many values of p_i for which $\frac{-1}{\ln(1-p_i)}$ is an integer, so the choice of t in our previous optimization is indeed valid.

Under an optimal schedule, the attacker’s expected utility is $\frac{1}{4}$, completing the proof of the approximation guarantee. \blacksquare

B Formalization of Notions about Schedules

B.1 Canonical Schedules

The general definition of defender schedules allows for strange schedules that are clearly suboptimal, such as the defender leaving a target i and returning to it shortly afterwards, or visiting a target infinitely often within a bounded time interval with shorter and shorter return times. For ease of notation and analysis, we would like to rule out such schedules. The following definition captures “reasonable” schedules.

Definition B.1 (canonical schedules) *We say that a valid schedule ℓ is canonical if \mathbb{R}^+ can be partitioned into countably many disjoint intervals I_1, I_2, I_3, \dots with the following properties:*

1. All odd intervals I_{2k-1} are open and of length exactly 1, and $\ell(t) = \perp$ if and only if $t \in \bigcup_k I_{2k-1}$.

2. All even intervals I_{2k} are closed. (Even intervals could consist of a single point.)

A defender mixed schedule Λ is canonical if it is a distribution over canonical deterministic schedules.

Note that it follows from validity that any canonical ℓ is constant on the even intervals.

Intuitively, a canonical schedule is one in which the defender travels as quickly as possible (in one unit of time) from one target to the next target, visits it for some (possibly zero) time, then travels to the next (necessarily different) target, etc. That we may focus on canonical schedules w.l.o.g. is captured by the following proposition:

Proposition B.2 *For each valid schedule ℓ , there exists a canonical schedule ℓ' that is at least as good for the defender, in the sense that for any choice (i, t_0, t) of the attacker,*

$$U(\ell', (i, t_0, t)) \leq U(\ell, (i, t_0, t)).$$

Proof. Given ℓ , define ℓ' as follows.

1. For every t with $\ell(t) \neq \perp$ let $\ell'(t) = \ell(t)$.
2. For every t with $\ell(t) = \perp$
 - (a) If t is in the closure of $\ell^{-1}(i)$, set $\ell'(t) = i$.
 - (b) Denote by $i(t)$ the last target visited before time t (setting $i(t) = 1$ if none exists) and by $j(t)$ the first target visited after time t (again setting $j(t) = 1$ if none exists). Note that $i(t)$ and $j(t)$ are well-defined because ℓ is valid; this would not in general be true for an arbitrary $\ell: \mathbb{R}^+ \rightarrow \{1, \dots, n, \perp\}$.
 - (c) If $i(t) = j(t)$ then set $\ell'(t) = i(t)$. That is, if in ℓ , the defender leaves a target i and then comes back to it without visiting another, then in ℓ' , the defender just stays at i . (In addition to the defining property of being canonical, this ensures that no target i is visited twice in a row.)
 - (d) If $i(t) \neq j(t)$ and the difference between t and $\inf\{\tau > t \mid \ell(\tau) = j(t)\}$ is at least 1, then set $\ell'(t) = i(t)$. That is, if the defender took more than one unit of time to reach target $j(t)$ from $i(t)$, then she might as well have stayed at $i(t)$ until one time unit before getting to $j(t)$.
 - (e) Otherwise, set $\ell'(t) = \perp$.

It is easy to verify that ℓ' is indeed canonical. Consider any choice of attack (i, t_0, t) . Because the preceding transformations only replaced \perp (i.e., transit) times with times at targets, whenever the attacker is not caught in ℓ' , he was not caught in ℓ , so his utility can only decrease:

$$U(\ell', (i, t_0, t)) \leq U(\ell, (i, t_0, t)). \blacksquare$$

B.2 Shift Invariance

To simplify the analysis, we would like to restrict our attention to *shift invariant* schedules for the defender: schedules such that the attacker's and defender's utilities depend only on the duration $t' - t$ of the attack, but not on the start time t . We formally define this notion as follows, and show that this restriction is without loss of generality, as there is always an optimal shift-invariant schedule. For each $\tau \in \mathbb{R}^+$, define the *shift operator* $M_\tau : \mathcal{L} \rightarrow \mathcal{L}$ by

$$[M_\tau(\ell)](t) = \ell(t + \tau).$$

That is, the pure schedule $M_\tau(\ell)$ is equal to ℓ , but leaves out the first τ time units of ℓ , shifting the remainder of the schedule forward in time. Note that

$$U(M_\tau(\ell), (i, t_0, t)) = U(\ell, (i, t_0 + \tau, t)). \quad (9)$$

The operator M_τ extends naturally to act on mixed schedules Λ .¹⁰ We say that a mixed schedule Λ is *shift-invariant* if $M_\tau(\Lambda) = \Lambda$ for all $\tau \in \mathbb{R}^+$. The following lemma shows that an optimal schedule for the defender exists, and that we may focus on shift-invariant schedules without loss of generality.

Lemma B.3 *The defender has an optimal mixed schedule that is shift-invariant.*

Proof. To prove this lemma, we introduce a natural topology on \mathcal{L} , the space of valid canonical pure strategies. This topology is related to the Skorohod topology [27]. Given a $\ell \in \mathcal{L}$, define $\bar{\ell} : \mathbb{R}^+ \rightarrow \{1, \dots, n\}$ by setting $\bar{\ell}(t)$ to be either $\ell(t)$, if $\ell(t) \neq \perp$, or else setting it to be the first target visited after time t . Thus, $\bar{\ell}(t)$ is the target visited at time t , or the target that the defender is en route to visiting. Note that $\bar{\ell}^{-1}(i)$ is the union of a countable set of intervals of length at least 1, each open on the left and closed on the right. Note also that the map $\ell \mapsto \bar{\ell}$ is “almost” invertible; since travel times are always 1, we know when each visit to each target began. The exception is the first visit, and so ℓ is determined by $\bar{\ell}$, together with the time of the beginning of the first target visit, which is always at most 1.

The topology on \mathcal{L} is the topology of convergence in \mathcal{L}^1 on compact sets. Specifically, for any $t_1, t_2 \in \mathbb{R}^+$, define $\Delta_{t_1, t_2}(\ell', \ell)$ to be the measure of the subset of $[t_1, t_2]$ on which at least one of the following two holds: (1) $\ell' \neq \ell$, or (2) $\bar{\ell}' \neq \bar{\ell}$. Then, we say that the limit of ℓ_m for $m \rightarrow \infty$ is equal to ℓ iff $\Delta_{t_1, t_2}(\ell_m, \ell) \rightarrow 0$ for all $t_1, t_2 \in \mathbb{R}^+$. It is straightforward to verify that this topology is compact and metrizable.¹¹ Hence the corresponding weak* topology on mixed strategies is also compact. Note also that the shift operator $M_\tau : S \rightarrow S$ is continuous in this topology.

Note that if $\ell_m \rightarrow_{m \rightarrow \infty} \ell$, and if target i is visited in $[t_0, t_0 + t]$ in every ℓ_m , then it is also visited in $[t_0, t_0 + t]$ in ℓ . Hence,

$$\lim_{m \rightarrow \infty} U(\ell_m, (i, t_0, t)) \geq U(\ell, (i, t_0, t)),$$

and so $U(\cdot, (i, t_0, t))$ is a lower semi-continuous map from \mathcal{L} to \mathbb{R}^+ . It follows that

$$\Lambda \mapsto \mathbb{E}_{\ell \sim \Lambda} [U(\ell, (i, t_0, t))]$$

¹⁰A measurable map $P : X \rightarrow X$ can be extended to a linear operator on probability measures on X as follows: For any measurable subset $A \subseteq X$, define $[P(\mu)](A) = \mu(P^{-1}(A))$. This defines a mapping $\mu \mapsto P(\mu)$.

¹¹The metric is $\sum_{m=1}^{\infty} 2^{-m} d_m(\ell_1, \ell_2)$, where $d_m(\ell_1, \ell_2)$ is the measure of the subset of $[0, m]$ in which either ℓ_1 and ℓ_2 differ, or $\bar{\ell}_1$ and $\bar{\ell}_2$ differ.

is lower semi-continuous as well. Hence

$$U(\Lambda) = \sup_{i,t_0,t} \mathbb{E}_{\ell \sim \Lambda} [U(\ell, (i, t_0, t))]$$

is also lower-semicontinuous, and thus attains a minimum on the compact space of mixed strategies. Thus we have shown that an optimal schedule exists.

When the attacker can obtain expected utility u against $M_\tau(\ell)$ by choosing (i, t_0, t) , he can obtain the same utility u against Λ by choosing $(i, t_0 + \tau, t)$. Therefore, the defender's utility is (weakly) monotone in τ , in the following sense:

$$U(M_\tau(\Lambda)) \leq U(\Lambda). \quad (10)$$

Let Λ_1 and Λ_2 be mixed strategies, and let $\Lambda = \beta\Lambda_1 + (1 - \beta)\Lambda_2$ be the schedule in which Λ_1 is carried out with probability β and Λ_2 with probability $1 - \beta$. Since suprema are subadditive, the attacker's utility is convex:

$$U(\Lambda) \leq \beta U(\Lambda_1) + (1 - \beta)U(\Lambda_2). \quad (11)$$

Let Λ be an optimal mixed schedule. For $m \in \mathbb{N}$ let

$$\Lambda_m = \frac{1}{m} \int_0^m M_\tau(\Lambda) d\tau.$$

By the monotonicity (Eq. (10)) and convexity (Eq. (11)) of $U(\Lambda)$, we have that $U(\Lambda_m) \leq U(\Lambda)$.

Since \mathcal{L} is compact, the sequence $(\Lambda_m)_m$ has a converging subsequence that converges to some Λ_∞ . By the lower semi-continuity of $U(\Lambda)$,

$$U(\Lambda_\infty) \leq \lim_{m \rightarrow \infty} U(\Lambda_m) \leq U(\Lambda);$$

therefore Λ_∞ is also optimal. Finally, Λ_∞ is by construction shift-invariant. ■

B.3 Transitive and Ergodic Schedules

We say that a shift-invariant mixed schedule Λ is *transitive* if almost every pure schedule ℓ_0 chosen from Λ is periodic with some period τ (i.e., $M_\tau(\ell_0) = \ell_0$) and

$$\Lambda = \frac{1}{\tau} \int_0^\tau \delta_{M_t(\ell_0)} dt,$$

where δ_ℓ is the point mass on ℓ . Intuitively, Λ simply repeats the same periodic schedule, with a phase chosen uniformly at random.

A weaker property of a shift-invariant mixed schedule Λ is *ergodicity*: Λ is ergodic if almost every pure schedule ℓ_0 chosen from Λ satisfies

$$\Lambda = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau \delta_{M_t(\ell_0)} dt.$$

In fact, this is not the usual definition of an ergodic measure, but the conclusion of the Ergodic Theorem. An equivalent property is that Λ cannot be written as the convex combination $\Lambda = \beta\Lambda_1 + (1 - \beta)\Lambda_2$ of two different shift-invariant measures. That is, Λ is an extremal point in the convex set (simplex, in fact) of shift-invariant measures.

B.4 Times Between Visits to Targets

We now more formally define the notion of the (random) time between visits to a target i . While the notion is intuitively clear, for arbitrary defender strategies Λ , a precise definition requires some subtlety. We give a general definition for arbitrary mixed schedules, not just random sequences.

Having defined schedules on \mathbb{R}^+ , we now extend the definition to schedules on $[-\tau, \infty)$ and eventually to \mathbb{R} , using a standard construction called the *bi-infinite extension*. We define a modified shift operator $\tilde{M}_\tau(\cdot)$, mapping schedules $(\ell: \mathbb{R}^+ \rightarrow \{1, \dots, n, \perp\})$ to τ -schedules $\ell': [-\tau, \infty) \rightarrow \{1, \dots, n, \perp\}$, via $[\tilde{M}_\tau(\ell)](t) = \ell(t + \tau)$. Thus, $\tilde{M}_\tau(\ell)$ is simply a version of ℓ shifted τ units to the left. The map $\tilde{M}_\tau(\cdot)$ extends to a map on mixed schedules in the obvious way. For any shift-invariant mixed schedule Λ , $\tilde{M}_\tau(\Lambda)$ is also shift-invariant, and furthermore, for any $\tau' < \tau$, the distribution $\tilde{M}_\tau(\Lambda)$, projected to $[-\tau', \infty)$, is the same distribution as $\tilde{M}_{\tau'}(\Lambda)$. It follows that

$$\Lambda_\infty = \lim_{\tau \rightarrow \infty} \tilde{M}_\tau(\Lambda) \quad (12)$$

is a well defined measure on pure schedules that are functions $\ell_\infty: \mathbb{R} \rightarrow \{1, \dots, n, \perp\}$. We call Λ_∞ the *bi-infinite extension* of Λ . It is straightforward to verify that it, too, is shift-invariant. Note that the distribution of the first visit to i at non-negative times, $R_i = \min\{t \geq 0 \mid \ell(t) = i\}$, has the same distribution under Λ_∞ as under Λ , since the restriction of Λ_∞ to non-negative times is equal to Λ .

Given a target i and a shift-invariant mixed schedule Λ , let $\tilde{\ell}: (-\infty, \infty) \rightarrow \{1, \dots, n, \perp\}$ be a random schedule with distribution Λ_∞ . Let B_i be the (random) time between the last visit to i before time zero, until the first visit to i after time zero:

$$B_i = (\inf\{t \geq 0 \mid \tilde{\ell}(t) = i\}) - (\sup\{t \leq 0 \mid \tilde{\ell}(t) = i\}).$$

The choice of time 0 here is immaterial because of shift invariance. B_i could be infinite, but this will never happen in an optimal Λ , because it would imply that the attacker's expected utility for choosing i is infinite; we hence assume henceforth that $\mathbb{P}[B_i = \infty] = 0$. Finally, contrary to what one might intuitively guess, even for transitive Λ , the distribution of B_i is *not* the same as the long-run empirical distribution of times between visits, as gaps are chosen at time 0 with probability proportional to their length. The same holds for general Λ .

C Tightness of the 2-Quasi-Regularity Result

In this section, we prove Proposition 4.6. For convenience, we restate the proposition here:

Proposition 4.6 *Let $n = 3$ and $\mathbf{p} = (1/2, 1/3, 1/6)$. Then, for every $\epsilon > 0$, there are no $(2 - \epsilon)$ -quasi-regular random sequences.*

Proof. Let s be a $(2 - \epsilon)$ -quasi-regular random sequence. We claim that $B_1 = 2$ with probability 1, and $B_2 = 3$ with probability 1. For suppose that with positive probability $B_1 \leq 1$. Then, because $T_1 = 2$, we also would have to have $B_1 \geq 3$ with positive probability, and vice versa. Similarly, $B_2 \leq 2$ with positive probability iff $B_2 \geq 4$ with positive probability. Either of those cases would lead to a ratio $(3/1$ or $4/2)$ larger than $2 - \epsilon$, violating $(2 - \epsilon)$ -quasi-regularity.

Hence, with probability one, target 1 appears in every other time period and target 2 appears in every third time period, which is impossible. ■

D Proof of Theorem 5.3

In this section, we prove Theorem 5.3, restated here for convenience:

Theorem 5.3 (Slater [28]) *Assume that $p \leq \frac{1}{2}$. Let k be smallest such that*

$$|f_{k+1}/\varphi - f_k| \leq p.$$

Then, the distribution of return times is

$$\begin{aligned} \mathbb{P}[B_i = f_{k+1}] &= f_{k+1} \cdot \left(p - (1/\varphi)^{k+1}\right), \\ \mathbb{P}[B_i = f_{k+2}] &= f_{k+2} \cdot \left(p - (1/\varphi)^{k+2}\right), \\ \mathbb{P}[B_i = f_{k+3}] &= f_{k+3} \cdot \left(-p + (1/\varphi)^k\right), \\ \mathbb{P}[B_i = t] &= 0 \text{ for all other } t. \end{aligned}$$

We begin with a few simple, but useful, technical lemmas. First, we give a closed form for expressions of the form $f_{k+1} - \varphi f_k$.

Lemma D.1 *For any k , we have that $f_{k+1} - \varphi f_k = (-1/\varphi)^k$.*

Proof. Using the closed-form expression for Fibonacci Numbers (Part 2 of Lemma 5.2), we can write

$$\begin{aligned} f_{k+1} - \varphi f_k &= \frac{(\varphi^{k+1} - (-1/\varphi)^{k+1}) - (\varphi^{k+1} - \varphi(-1/\varphi)^k)}{\sqrt{5}} \\ &= \frac{-(-1/\varphi)^{k+1} + \varphi(-1/\varphi)^k}{\sqrt{5}} \\ &= (-1/\varphi)^k \cdot \frac{1/\varphi + \varphi}{\sqrt{5}} = (-1/\varphi)^k \varphi. \blacksquare \end{aligned}$$

Lemma D.2 1. *For all $\delta \in (-\frac{1}{2}, \frac{1}{2}]$ and integers $N \geq 1$, the following two are equivalent:*

- $\delta = (\varphi N) \pmod{1}$.
- *There exists a positive integer D with $N/D - \varphi = \delta\varphi/D$.*

2. *Let $\delta = (\varphi f_k) \pmod{1}$ for $k \geq 2$ (where we consider the range of the \pmod operation to be $(-\frac{1}{2}, \frac{1}{2}]$). Then, $\delta = f_k/\varphi - f_{k-1}$.*

Proof. 1. Because $\varphi = 1 + 1/\varphi$, the first condition can be rewritten as $(1 + 1/\varphi)N = \delta + D'$ for some integer D' . When $\delta < 0$, we must have $D' > N$; when $\delta \geq 0$, because $N/\varphi > \frac{1}{2} \geq \delta$, we again have $D' > N$. The preceding equality can therefore be rearranged to $N/\varphi = \delta + (D' - N)$. Multiplying by $\varphi/(D' - N)$ now gives equivalence with the second condition, writing $D = D' - N$.

2. In the first part of the lemma, set $N = f_k$. Then, the condition is equivalent to the existence of a positive integer D with $f_k/D - \varphi = \delta\varphi/D$, implying that $|f_k - D\varphi| = |\delta|\varphi$. By choosing $D = f_{k-1}$, according to Lemma D.1, we get that

$$|f_k/\varphi - f_{k-1}| = (1/\varphi)^k \stackrel{k \geq 2}{\leq} \frac{1}{2}.$$

Therefore, for any $D \neq f_{k-1}$, we get that $|f_k/\varphi - D| > 1 - \frac{1}{2} = \frac{1}{2}$, meaning that no $D \neq f_{k-1}$ can satisfy $f_k/D - \varphi = \delta\varphi/D$. By substituting the unique choice $D = f_{k-1}$, we obtain the second part of the lemma. ■

Because the Fibonacci numbers are the convergents of the Golden Ratio, they provide the best rational approximation, in the following sense.

Theorem D.3 *Let $\hat{M} \geq 1$ be arbitrary. Let k be the largest even number with $f_k \leq \hat{M}$, and k' the largest odd number with $f_{k'} \leq \hat{M}$. Then, for all $M \leq \hat{M}$ and all N , we have the following:*

1. $N/M > \varphi$ implies $f_{k+1}/f_k \leq N/M$.
2. $N/M < \varphi$ implies $f_{k'+1}/f_{k'} \geq N/M$.

Theorem D.3 follows directly from standard results stating that the convergents provide the best approximation to real numbers (e.g., [25, p. 11], noting that the second (intermediate) case cannot happen for the Golden Ratio).

We are now ready to prove the characterization of the distribution of B_i under the Golden Ratio schedule.

Proof of Theorem 5.3. We begin by showing that the support of return times consists only of Fibonacci Numbers. Consider the interval $I = [0, p)$. Let $m \geq 1$ be a return time. Let $x \in I$ be arbitrary, and $y = (x + \varphi m) \bmod 1$, which is in I by assumption. Define $\delta = y - x$. Because both $x, y \in I$, we have that $\delta \in [-x, p - x)$. By Part 1 of Lemma D.2, there is a positive integer D such that

$$m/D - \varphi = \delta\varphi/D \in [-x\varphi/D, (p - x)\varphi/D).$$

We now distinguish two cases:

- If $\delta > 0$, then $m/D > \varphi$, so Case 1 of Theorem D.3 implies that the largest even j such that $f_j \leq D$ satisfies $f_{j+1}/f_j > \varphi$ and $f_{j+1}/f_j \leq m/D$. Thus, $(x + \varphi f_{j+1}) \bmod 1 \in I$, meaning that the defender returns to the target in f_{j+1} steps. Because $D \geq f_j$ and $m/D \geq f_{j+1}/f_j$, we get that $m \geq f_{j+1}$; unless $m = f_{j+1}$, this would contradict the definition of m as a return time, so we have shown that m is a Fibonacci number.
- Similarly, if $\delta < 0$, and thus $m/D - \varphi < 0$, then Case 2 of Theorem D.3 implies that the largest odd j such that $f_j \leq D$ satisfies $f_{j+1}/f_j - \varphi < 0$ and $f_{j+1}/f_j \geq m/D$. By the same argument, we obtain now that $m = f_{j+1}$.

Next, we prove the second part of the theorem. First, notice that the k defined in the theorem actually exists. By Lemma D.1, we get that $|f_{k+1} - \varphi f_k| = (1/\varphi)^k \rightarrow 0$ as $k \rightarrow \infty$, so there exists a k (and thus a smallest k) with $|f_{k+1} - \varphi f_k| \leq \varphi p$.

We show that there cannot be a return time $m < f_{k+1}$. If there were, then by the previous part of the proof, m would be a Fibonacci number, say, $m = f_\ell$. And because $m \geq 1$, we get that $\ell \geq 2$. By Part 2 of Lemma D.2, that means that $f_\ell/\varphi - f_{\ell-1} = y - x$, and hence $|f_\ell/\varphi - f_{\ell-1}| = |y - x| < p$, contradicting the definition of k as smallest with that property.

Consider a return to I within m steps, starting from $x \in I$ and ending at $y \in I$, so that $\delta_\ell := y - x$ satisfies $|\delta_\ell| < p$. By the preceding analysis, $m = f_\ell$ for some $\ell \geq k + 1$. Again, by Part 2 of Lemma D.2, we obtain that $\delta_\ell = f_\ell/\varphi - f_{\ell-1}$.

When $\delta_\ell < 0$, the $x \in I$ with $x + \delta_\ell \in I$ are exactly captured by the interval $J_\ell := [|\delta_\ell|, p]$, while for $\delta_\ell > 0$, they are exactly the interval $J_\ell := [0, p - \delta_\ell]$. In either case, the interval J_ℓ has size exactly $|J_\ell| = p - |\delta_\ell|$.

We will show that $J_{k+2} \cup J_{k+3} = I$. By Lemma 5.2, the signs of δ_ℓ are alternating, meaning that the intervals J_ℓ alternate being of the form $[0, x]$ and $[y, p]$. In particular, to show that $J_{k+2} \cup J_{k+3} = I$, it suffices to show that $|J_{k+2}| + |J_{k+3}| \geq p$. Because $|J_{k+2}| + |J_{k+3}| = 2p - |\delta_{k+2}| - |\delta_{k+3}|$, this is equivalent to showing that $|\delta_{k+2}| + |\delta_{k+3}| \leq p$. We distinguish two cases, based on whether k is even or odd.

- If k is even, then $\delta_{k+2} = f_{k+2}/\varphi - f_{k+1} < 0$ and $\delta_{k+3} = f_{k+3}/\varphi - f_{k+2} > 0$, so we obtain that
$$|\delta_{k+3}| + |\delta_{k+2}| = f_{k+3}/\varphi - f_{k+2} - f_{k+2}/\varphi + f_{k+1} = f_{k+1}/\varphi - f_k = |f_{k+1}/\varphi - f_k| \leq p,$$
by the definition of k .
- If k is odd, then $\delta_{k+2} = f_{k+2}/\varphi - f_{k+1} > 0$ and $\delta_{k+3} = f_{k+3}/\varphi - f_{k+2} < 0$, so we obtain that
$$|\delta_{k+2}| + |\delta_{k+3}| = f_{k+2}/\varphi - f_{k+1} - f_{k+3}/\varphi + f_{k+2} = f_k - f_{k+1}/\varphi = |f_{k+1}/\varphi - f_k| \leq p.$$

Thus, we have shown that the support of the distribution is indeed contained in $\{f_{k+1}, f_{k+2}, f_{k+3}\}$. Finally, we can work out the frequencies. Conditioned on being in the interval of size p , the probability of being in J_ℓ is $q_\ell = |J_\ell|/p$. To arrive at the attacker's observed distribution of B_i , we notice that the probability of time 0 being in an interval of length f_ℓ is

$$\frac{q_\ell f_\ell}{\sum_j q_j f_j} = \frac{q_\ell f_\ell}{1/p} = f_\ell \cdot |J_\ell|.$$

Thus, we obtain that

$$\begin{aligned} \mathbb{P}[B_i = f_{k+1}] &= f_{k+1} \cdot |J_{k+1}| = f_{k+1} \cdot (p - |f_{k+1}/\varphi - f_k|), \\ \mathbb{P}[B_i = f_{k+2}] &= f_{k+2} \cdot |J_{k+2}| = f_{k+2} \cdot (p - |f_{k+1} - f_{k+2}/\varphi|), \\ \mathbb{P}[B_i = f_{k+3}] &= 1 - q_1 - q_2 \\ &= f_{k+3} \cdot (-p + |f_{k+1}/\varphi - f_k + f_{k+1} - f_{k+2}/\varphi|) \\ &= f_{k+3} \cdot (-p + |f_{k-1} - f_k/\varphi|). \end{aligned}$$

Notice that we arranged the terms inside absolute values such that for even k , they are all positive, while for odd k , they are all negative. This allowed us to simply add inside the absolute value. Applying Lemma D.1 to all three terms now completes the proof. \blacksquare

E Computational Considerations for the Golden Ratio Schedule

As phrased, Algorithm 2 requires precise arithmetic on irrational numbers, and drawing a uniformly random number from $[0, 1]$. Here, we discuss how to implement the algorithm such that each target i visited in step t can be computed in time polynomial in the input size.

Let $p_i = a_i/b_i$ for each i , and write $M = \text{lcm}(b_1, \dots, b_m)$ for the common denominator. Notice that $\log M \leq \sum_i \log b_i$ is polynomial in the input size.

For each i , the number $P_i = \sum_{i' < i} p_{i'}$ is rational. To decide whether target i is visited in step t , the algorithm needs to decide if $(\lambda + t/\varphi) \bmod 1 \in [P_i, P_{i+1}]$, or — equivalently — if there is an integer D with $\lambda + t/\varphi \in [D + P_i, D + P_{i+1}]$. To decide whether $\lambda + t/\varphi < D + P_j$ or $\lambda + t/\varphi > D + P_j$ (for $j \in \{i, i + 1\}$), the algorithm needs to decide if $\varphi < \frac{t}{D + P_j - \lambda}$ or $\varphi > \frac{t}{D + P_j - \lambda}$. The key question is how many digits of φ the algorithm needs to evaluate for this decision, and how many digits of the uniformly random offset λ it needs to decide on.

Suppose that the algorithm has generated the first k random digits of λ , having committed to $\frac{\ell}{10^k} \leq \lambda < \frac{\ell+1}{10^k}$ for some $\ell \in \{0, 1, \dots, 10^k - 1\}$. Writing $P_j = N_j/M$ (using the denominator M defined above), a decision about target P_j can be made whenever $\varphi < \frac{tM10^k}{10^k \cdot MD + 10^k N_j - M \cdot \ell}$ or $\varphi > \frac{tM10^k}{10^k \cdot MD + 10^k N_j - M \cdot (\ell+1)}$. In both cases, the right-hand side is a rational approximation to φ with denominator bounded by $\hat{M} := 2 \cdot 10^k \cdot MD$.

It is well known (see, e.g., [16, Theorems 193–194]) that $|\varphi - \frac{\hat{N}}{\hat{M}}| \geq \frac{1}{(\sqrt{5}-\epsilon)\hat{M}^2}$ for all $\epsilon > 0$. In particular, this implies that evaluating φ to within $O(\log \hat{M}^2) = O(k + \log M + \log D)$ digits is sufficient to test whether $\varphi < \frac{tM10^k}{10^k \cdot MD + 10^k N_j - M \cdot \ell}$, and whether $\varphi > \frac{tM10^k}{10^k \cdot MD + 10^k N_j - M \cdot (\ell+1)}$. In either of these cases, the algorithm has resolved whether $\varphi < \frac{t}{D + P_j - \lambda}$.

The only case where the algorithm cannot resolve whether $\varphi < \frac{t}{D + P_j - \lambda}$ is when

$$\frac{tM10^k}{10^k \cdot MD + 10^k N_j - M \cdot \ell} < \varphi < \frac{tM10^k}{10^k \cdot MD + 10^k N_j - M \cdot (\ell + 1)}.$$

In this case, the number of digits for λ is insufficient. Notice that there is a unique value of ℓ for which this happens, so the probability of failure is at most 10^{-k} . Taking a union bound over all n interval boundaries and all t rounds, we see that in order to succeed with high probability, the number of digits of λ the algorithm needs to generate is $O(\log n + \log t)$.

In particular, the computation and required randomness are polynomial.

References

- [1] Omer Angel, Alexander E. Holroyd, James B. Martin, and James Propp. Discrete low-discrepancy sequences. <https://arxiv.org/abs/0910.1077>, 2010.
- [2] Rudolf Avenhaus and Morton John Canty. Playing for time: A sequential inspection game. *European Journal of Operational Research*, 167(2):475–492, 2005.
- [3] Rudolf Avenhaus and Thomas Krieger. Distributing inspections in space and time—proposed solution of a difficult problem. *European Journal of Operational Research*, 231(3):712–719, 2013.

- [4] Rudolf Avenhaus, Bernhard von Stengel, and Shmuel Zamir. Inspection games. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory, Vol. 3*, chapter 51, pages 1947–1987. North-Holland, Amsterdam, 2002.
- [5] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proc. 8th Intl. Conf. on Autonomous Agents and Multiagent Systems*, pages 57–64, 2009.
- [6] Mee Yee Chan and Francis Y. L. Chin. Schedulers for larger classes of pinwheel instances. *Algorithmica*, 9:425–462, 1993.
- [7] Chandra Chekuri, Jan Vondák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *Proc. 51st IEEE Symp. on Foundations of Computer Science*, pages 575–584, 2010.
- [8] William Chen, Anand Srivastav, and Giancarlo Travaglini, editors. *A Panorama of Discrepancy Theory*, volume 2107 of *Lecture Notes in Mathematics*. Springer, 2014.
- [9] Fan Chung and Ron Graham. On the discrepancy of circular sequences of reals. *Journal of Number Theory*, 164:52–65, 2016.
- [10] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *Proc. 7th ACM Conf. on Electronic Commerce*, pages 82–90, 2006.
- [11] Harvey Diamond. Minimax policies for unobservable inspections. *Mathematics of Operations Research*, 7(1):139–153, 1982.
- [12] Fei Fang, Thanh H. Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Brian C. Schwedock, Milind Tambe, and Andrew Lemieux. PAWS — a deployed game-theoretic application to combat poaching. *AI Magazine*, 2017. (to appear).
- [13] Fei Fang, Peter Stone, and Milind Tambe. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *Proc. 24th Intl. Joint Conf. on Artificial Intelligence*, pages 2589–2595, 2015.
- [14] Peter C. Fishburn and Jeffrey C. Lagarias. Pinwheel scheduling: achievable densities. *Algorithmica*, 22:14–38, 2002.
- [15] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM*, 53:324–360, 2006.
- [16] Godfrey H. Hardy and Edward M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, fourth edition, 1960.
- [17] Alexander E. Holroyd and James Propp. Rotor walks and markov chains. *Algorithmic probability and combinatorics*, 520:105–126, 2010.
- [18] Robert Holte, Al Mok, Louis Rosier, Igor Tulchinsky, and Donald Varvel. The pinwheel: a real-time scheduling problem. In *Proc. Twenty-Second Annual Hawaii International Conference on System Sciences. Vol.II: Software Track*, 1989.

- [19] Robert Holte, Louis Rosier, Igor Tulchinsky, and Donald Varvel. Pinwheel scheduling with two distinct numbers. *Theoretical Computer Science*, 100:105–135, 1992.
- [20] Nicole Immorlica and Robert D. Kleinberg. Recharging bandits. Manuscript. Available upon request, 2017.
- [21] Alon Itai and Zvi Rosberg. A golden ratio control policy for a multiple-access channel. *IEEE Transactions on Automatic Control*, 29(8):712–718, 1984.
- [22] Donald E. Knuth. *The Art of Computer Programming: Sorting and Searching*, volume 3. Pearson Education, 1998.
- [23] Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Security games with multiple attacker resources. In *Proc. 22nd Intl. Joint Conf. on Artificial Intelligence*, pages 273–279, 2011.
- [24] Shun-Shii Lin and Kwei-Jay Lin. A pinwheel scheduler for three distinct numbers with a tight schedulability bound. *Algorithmica*, 19:411–426, 1997.
- [25] László Lovász. *An Algorithmic Theory of Numbers, Graphs and Convexity*, volume 50 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, 1986.
- [26] Shivendra S. Panwar, Thomas K. Philips, and Mon Song Chen. Golden ratio scheduling for flow control with low buffer requirements. *IEEE Transactions on Communications*, 40(4):765–772, 1992.
- [27] Anatolii Volodimirovich Skorokhod. Limit theorems for stochastic processes with independent increments. *Theory of Probability & Its Applications*, 2(2):138–171, 1957.
- [28] Noel B. Slater. The distribution of the integers n for which $\{\theta n\} < \varphi$. *Proceedings of the Cambridge Philosophical Society*, 46(4):525–534, 1950.
- [29] Noel B. Slater. Gaps and steps for the sequence $n\theta \pmod 1$. *Mathematical Proceedings of the Cambridge Philosophical Society*, 63(04):1115–1123, 1967.
- [30] Aravind Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science*, pages 588–597, 2001.
- [31] Milind Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.
- [32] Robert Tijdeman. On a distribution problem in finite and countable sets. *Journal of Combinatorial Theory, Series A*, 15(2):129–137, 1973.
- [33] Robert Tijdeman. The chairman assignment problem. *Discrete Mathematics*, 32:323–330, 1980.
- [34] Heinrich von Stackelberg. *Marktform und Gleichgewicht*. Springer, 1934.
- [35] Bernhard von Stengel. Recursive inspection games. *Mathematics of Operations Research*, 41:935–952, 2016.
- [36] Yevgeniy Vorobeychik, Bo An, and Milind Tambe. Adversarial patrolling games. In *Proc. 11th Intl. Conf. on Autonomous Agents and Multiagent Systems*, pages 1307–1308, 2012.