2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

On the Nash Equilibria of a Timed Asymmetric Skirmish

Kyle Treleaven, Kevin Spieser, and Emilio Frazzoli

***Abstract*— In this work, we present a Nash equilibrium solution for a timed, asymmetric skirmish between two agents: an attacker, and a defender. We derive a solution by focusing on strategy profiles in which both the attacker and defender randomize their actions, which correspond to times, over a common atomic support. We show this class of strategies admits a unique mixed-strategy Nash equilibrium and give an algorithm for its computation. A numerical example highlights interesting features of a typical equilibrium strategy profile.**

I. INTRODUCTION

In this work, we consider a simple two-player competitive game, inspired by a skirmish scenario, with a focus on strategic timing. One player is the *attacker*; the attacker's goal is to penetrate the opponent's defenses and seize control of a valuable resource. The second player is the *defender*; the defender tries to retain control of the resource by executing a strategically-timed defensive maneuver (e.g., using an expensive defensive stance or formation). Attacker and defender strategies consist of times, on $\mathbb{R}_{\geq 0}$, at which to attack and defend, respectively. With an interest in the emergence of temporally-based tactics, we explore solution methods for a two-player, simple-timed game in which players are allowed a single action and may act asynchronously.

Given the actions are times at which to attack and defend, the skirmish we have described belongs to the family of *timed games*. Among the more well-known members of this family is the War of Attrition (WoA) game [1]. In WoA, each player, i, has a valuation, V_i, for a common item, and bids a costly waiting time t_i in an attempt to obtain the good. The player with the lowest bid receives utility $V_i - t_i$. The remaining players each receive zero utility. The WoA is also popularly referred to as the second-price all-pay auction. A deeper exposition of timed games can be found in [2].

In this work, we analyze the structure of equilibria in a temporal game that is similar to, albeit more complex than, the War of Attrition game. Our model is an abstraction capturing the key strategic elements of a common attacker-defender stand-off found in a number of applications, including, for example, the duel and searchlight games in [3]. For example, the inspiration of this work was a base-defense scenario in which a defender must anticipate an enemy attack. There are also representative examples of our skirmish scenario in the financial world. For example, an aggressive firm (attacker) may attempt a hostile takeover of a competing business. With this threat looming, the targeted firm (defender) may have a single opportunity to muster their resources, e.g., call an emergency shareholder's meeting, in hopes of avoiding the takeover. Since we believe the model represents the underlying mechanics in a number of domains, we shall henceforth refer to this scenario simply as the Redhands game. This name was chosen in recognition of a popular schoolyard game in which one child attempts to slap the hands of a second child (turning them red). The second child tries to avoid the slap by quickly pulling their hands away at the last instant; however, the child is heavily penalized for removing their hands prematurely.

The game Redhands is more complex than many typical examples in classical game theory. It can be formulated naturally as a game on timed-automata. Ultimately, we show that Redhands is simple among this extremely complex family of games; however, its namesake (the multiple-round schoolyard version) demonstrates a mechanism which may require a more sophisticated approach. The control of timed automata generally falls within two competing frameworks: *supervisory control theory* (SCT) [4], and *forced-event* control [5]. There is also some existing work in the literature for dynamical timed games, e.g. [6], [7]. However, in [6], the authors are concerned with games-of-type solution concepts, e.g. the existence of policies that satisfy some formally expressed winning condition. In [7], the authors consider a worst-case optimal control problem on timed automata (restricting the results, in some sense, to zero-sum games). Introducing scoring functions over strategies, and removing zero-sum assumptions (as we do in this work), has the potential to vastly increases the complexity of such problems. Due to this complexity, to our knowledge, there is no fully game-theoretic treatment of timed automata. By analyzing the Redhands game, we hope instead to gain insight into the structure of solutions that can emerge in simple but relevant examples within this family.

The contributions of this work are as follows. First, we discuss the inability of standard results to guarantee the existence of a Nash equilibrium in the Redhands game. Nevertheless, we show that Redhands is amenable to analysis using classical techniques, i.e. strategy synthesis from a best-response characterization. We describe a class of strategies—namely, those with an atomic, uniformly-spaced support—within which we can guarantee existence of a mixed-strategy Nash equilibrium. Moreover, we show that the equilibrium strategy is unique within this class, and we synthesize an algorithm to compute it. One limitation of the work is that we leave open the issue of equilibrium uniqueness within the full space of mixed strategies. Also, this work does not consider the use of sensors to detect enemy attacks; though, the utility

This research is partially supported by the ONR, grant #00014-09-1-0751.

The authors are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, 02139. {ktreleav,kspieser,frazzoli}@mit.edu

978-1-4577-0079-8/11/$26.00 ©2011 AACC

of sensors is unclear if players may employ deceitful tactics (e.g. in the schoolyard game).

The remainder of the paper is organized as follows. Section II describes the problem formulation: a game model of the Redhands game. In Section III, we discuss existence guarantees for Nash equilibria in Redhands, and then characterize attacker and defender best responses. We characterize an atomic, mixed-strategy Nash equilibrium in Section IV, and we prove its uniqueness within an expedient class of atomic strategy profiles. In Section V, we provide a numerical example for a game played with sample parameter values. Section VI closes with conclusions and future directions.

II. Problem Formulation

In this section, we present our formulation of the Redhands game and motivate the analysis of its solution. The game is comprised of two players, player A and D. Henceforth, A and D will refer to the attacker and defender, respectively.

Player A's strategy is the time, $t_A \in \mathbb{R}_{\geq 0}$, at which he chooses to initiate an attack, ending the game. If A launches an attack at time t_A, then he will either register a "hit" or record a "miss". In the event of a hit, A seizes some time-discounted reward $V_A e^{-\beta_A t_A}$, and inflicts a punishment c_D on D. Here, $V_A > 0$ is the valuation of player A, and the discounting term in A's payoff reflects the fact that time is valuable, providing an incentive to attack early. The constant $\beta_A > 0$ is called the *impatience* of player A; the larger the value of β_A, the greater A's disposition to attack early. The punishment term $c_D > 0$ is used to model some damage (physical or economical) that is inflicted on the defender. In the event A misses at time t_A, he receives nothing, while D escapes punishment and receives the reward of value $V_D e^{-\beta_D t_A}$.

Player D's strategy is the time, $t_D \in \mathbb{R}_{\geq 0}$, at which he chooses to deploy a single defensive "guard". Deploying his guard improves D's ability to deflect an attack over the time interval $[t_D, t_D + \tau_D)$, but he may only do this once, and therefore must use it wisely. Specifically, if A attacks at a time when D is not guarding, then A's chance of registering a hit is $\bar{p}$. On the other hand, if A attacks when D is guarding, then A's chance of registering a hit is $\underline{p} < \bar{p}$. The hit probability can be expressed succinctly in terms of t_A and t_D as follows:

$$p_{\mathrm{H}}(t_A, t_D) = \begin{cases} \underline{p} & \text{if } t_A \in [t_D, t_D + \tau_D) \\ \bar{p} & \text{otherwise.} \end{cases} \tag{1}$$

We assume players are fully aware of the value τ_D. This would be the case, for example, if the guard used a well-known technology (e.g., forcefield) or if the game had been played many times in the past. We encode these mechanisms in the players' utility functions, which express the expected reward of A and D, respectively, under the action pair (t_A, t_D):

$$\bar{u}_A(t_A, t_D) = V_A e^{-\beta_A t_A} p_{\mathrm{H}}(t_A, t_D), \tag{2}$$

$$\bar{u}_D(t_A, t_D) = V_D e^{-\beta_D t_A} [1 - p_{\mathrm{H}}(t_A, t_D)] - c_D p_{\mathrm{H}}(t_A, t_D). \tag{3}$$

In the present work, we pursue a solution to the Redhands game in the classical sense of Nash equilibria. We define this notion, in the present context, as follows:

Definition 2.1 (Pure-strategy Nash equilibrium): A strategy profile $(t_A,\ t_D)$ is said to be a *pure-strategy Nash equilibrium* (PSNE), if

$$\begin{aligned} &\bar{u}_A(t_A, t_D) \geq \bar{u}_A(t_A{}', t_D) && \text{for all } t_A{}' \geq 0, \\ \text{and} \quad &\bar{u}_D(t_A, t_D) \geq \bar{u}_D(t_A, t_D{}') && \text{for all } t_D{}' \geq 0. \end{aligned}$$

In this work, we allow the attacker and defender to use randomized strategies. In particular, we say that player A (D) plays strategy f_A (f_D) if his action t_A (t_D) is randomized according to the distribution f_A (f_D). We refer to the pair (f_A, f_D) as a *mixed-strategy profile*. In the analysis of this work, we will not distinguish between pure and mixed strategies as arguments to the utility functions. We will simply maintain an equivalence through expectation, in the sense that

$$\bar{u}(f_A, \cdot) \doteq \mathbb{E}_{f_A}\ \bar{u}(t_A, \cdot), \text{ and } \bar{u}(\cdot, f_D) \doteq \mathbb{E}_{f_D}\ \bar{u}(\cdot, t_D).$$

Definition 2.2 (Mixed-strategy Nash equilibrium): A mixed strategy profile (f_A, f_D) is said to be a *mixed-strategy Nash equilibrium* (MSNE), if

$$\begin{aligned} &\bar{u}_A(f_A, f_D) \geq \bar{u}_A(f_A{}', f_D) && \text{for all } f_A{}' \text{ in } \Sigma_{\mathbb{R}_{\geq 0}}, \\ \text{and} \quad &\bar{u}_D(f_A, f_D) \geq \bar{u}_D(f_A, f_D{}') && \text{for all } f_D{}' \text{ in } \Sigma_{\mathbb{R}_{\geq 0}}, \end{aligned}$$

where $\Sigma_{\mathbb{R}_{\geq 0}}$ denotes the set of probability distributions over the support $\mathbb{R}_{\geq 0}$.

As it applies to the present problem, a Nash equilibrium, be it pure or mixed, is a strategy profile in which neither the attacker nor defender can strictly improve their utility by unilaterally changing their strategy. The notion of both pure and mixed-strategy Nash equilibrium in the Redhands game are explored in the next section.

III. Equilibrium Structure of the Redhands Game

In this section, we explore the issue of existence for both PSNE and MSNE. We show that the existence of a PSNE depends on the patience of the attacker and that a well-known existence result for discontinuous games does not guarantee that the Redhands game has a MSNE. Given this result, we proceed to explore MSNE in the Redhands game from first principles, by characterizing the best response of the attacker and defender to given opponent strategies.

A. Existence of PSNE

We begin by showing that the existence of a pure-strategy Nash equilibrium is critically linked to the patience of the attacker through β_A.

Lemma 3.1: The Redhands game has unique pure-strategy Nash equilibrium $(t_A, t_D) = (0, 0)$ if and only if

$$\beta_A \geq \log(\bar{p}/\underline{p})/\tau_D. \tag{4}$$

Proof: Because $\underline{p} < \bar{p}$, the defender has an incentive from (3) to deviate from a strategy (t_A, t_D) unless $t_D \leq t_A < t_D + \tau_D$ (i.e. unless the guard interval contains the attack).

Moreover, we reason that $t_D = 0$ in any equilibrium: if $t_D > 0$ (and so $t_A \geq t_D > 0$), we have by (2) that $\bar{u}_A(t_A', t_D) > \bar{u}_A(t_A, t_D)$ for any $t_A' < t_D$. Applying the condition $t_D = 0$ to (2) we have

$$\bar{u}_A(t_A, 0) = V_A e^{-\beta_A t_A} \begin{cases} \underline{p} & \text{for } t_A \in [0, \tau_D), \\ \bar{p} & \text{otherwise.} \end{cases}$$

We observe that $\bar{u}_A(t_A, 0)$ is strictly decreasing over the interval $[0, \tau_D)$, and again over $[\tau_D, \infty)$. Thus, the attacker may play $t_A = 0$ as a best response if $\beta_A \geq (1/\tau_D)\log(\bar{p}/\underline{p})$, or he may play $t_A = \tau_D$ for $\beta_A \leq (1/\tau_D)\log(\bar{p}/\underline{p})$. We prove the lemma by eliminating the PSNE candidate $(\tau_D, 0)$, since its guard interval does not contain the attack. ■

Remark 3.2: In the case of a *perfect guard*, i.e. $\underline{p} = 0$, there can never exist a PSNE.

Before embarking on a characterization of possible MSNE in the Redhands game, we first remark on the inapplicability of the well-known Dasgupta and Maskin [8] existence result for mixed-equilibria in discontinuous games, as applied to the Redhands game.

Theorem 3.3: [MSNE in discontinuous games [8]] Let t_A (t_D) lie in a closed interval of $\mathbb{R}$. Assume $u_A(t_A, t_D)$ ($u_D(t_A, t_D)$) is continuous, except possibly on a finite number of lower-dimensional, continuous manifolds. Assume also that $u_A(t_A, t_D) + u_D(t_A, t_D)$ is everywhere upper semicontinuous and that $u_A(t_A, t_D)$ ($u_D(t_A, t_D)$) is bounded and weakly lower semicontinuous at all points of discontinuity. Then the game has a mixed-stategy Nash equilibrium.

Proof: The proof can be found in [8]. ■

The result of theorem 3.3 requires the technical condition that $u_A(t_A, t_D) + u_D(t_A, t_D)$ be upper semicontinuous at all points of discontinuity. The utility functions of A and D, are discontinuous on the rays $t_A = t_D$ and $t_A = t_D + \tau_D$. It can be shown that upper semicontinuity is not satisfied on both rays, and so 3.3 does not apply.

B. Best Response Characterization of MSNE

In Section III-A we argued that the Dasgupta and Maskin existence theorem does not apply to the Redhands game, and so we cannot yet claim existence of Nash equilibria. Undeterred, we continue to pursue a MSNE by deriving the attacker's and defender's best-responses and invoking conditions necessary for equilibrium play. We begin with the definition of best-response strategies.

Definition 3.4 (Best-response strategy): A strategy t_A is said to be a best-response to strategy f_D of player D if

$$\bar{u}_A(t_A, f_D) \geq \bar{u}_A(t_A', f_D) \qquad \text{for all } t_A' \in \mathbb{R}_{\geq 0}\ .$$

A strategy t_D is said to be a best-response to strategy f_A of player A if

$$\bar{u}_D(f_A, t_D) \geq \bar{u}_D(f_A, t_D') \qquad \text{for all } t_D' \in \mathbb{R}_{\geq 0}\ .$$

Closely associated with best-responses is the notion of incentive compatibility.

Definition 3.5 (Incentive Compatibility): A strategy f_A is said to be *incentive compatible* with f_D, if

$$\operatorname{supp} f_D \subseteq B_D(f_A),$$

where $B_D(\cdot)$ is a relation returning the set of best-responses of player D, and $\operatorname{supp} f_D$ denotes the set of values with positive probability in f_D. A similar definition holds for f_D.

In a MSNE, both strategies necessarily satisfy incentive compatibility; as a corollary, all strategies in the support achieve the same (maximum) utility.

We characterize the sets of best-response strategies of A and D, respectively. We may express player A's best response to f_D as:

$$\begin{aligned} B_A(f_D) &\doteq \arg\max_{t_A} \int_{t=0}^{\infty} \left[V_A e^{-\beta_A t_A} p_{\mathrm{H}}(t_A, t)\right] f_D(t)\ dt \\ &= \arg\max_{t_A}\ V_A e^{-\beta_A t_A} \left[\bar{p} \int_{t=0}^{t_A - \tau_D} f_D(t)\ dt \right. \\ &\qquad \left. + \underline{p} \int_{t=t_A-\tau_D}^{t_A} f_D(t)\ dt + \bar{p} \int_{t_A}^{\infty} f_D(t)\ dt \right] \\ &= \arg\max_{t_A}\ V_A e^{-\beta_A t_A} \{\bar{p} \\ &\qquad -(\bar{p} - \underline{p})\left[F_D(t_A) - F_D(t_A - \tau_D)\right]\}, \end{aligned} \quad (5)$$

where, for convenience, we let $F_D(t) = \int_{t'=-\infty}^{t} f_D(t')dt'$ denote the cumulative density function of f_D.

Due to the incentive compatibility requirement for a MSNE, it follows that for all $t_A \in \operatorname{supp}(f_A)$, there exists a constant v_A such that

$$F_D(t_A) - F_D(t_A - \tau_D) = \frac{1}{\bar{p} - \underline{p}} \left(\bar{p} - \frac{v_A}{V_A} e^{\beta_A t_A}\right). \quad (6)$$

Player D's best response to strategy f_A is given by

$$\begin{aligned} B_D(f_A) &\doteq \arg\max_{t_D} \int_{t=0}^{\infty} \left[e^{-\beta_D t} V_D[1 - p_{\mathrm{H}}(t, t_D)] - \right. \\ &\qquad \left. c_D p_{\mathrm{H}}(t, t_D)\right] f_A(t)\ dt_A \\ &= \arg\max_{t_D}\ (\bar{p} - \underline{p}) \int_{t=t_D}^{t_D+\tau_D} \left[e^{-\beta_D t} V_D + c_D\right] f_A(t)\ dt \\ &\qquad + \kappa(f_A), \end{aligned} \quad (7)$$

where, in the interest of brevity, we have introduced

$$\kappa(f_A) = \int_{t=0}^{\infty} \left[e^{-\beta_D t} V_D(1 - \bar{p}) - c_D \bar{p}\right] f_A(t)\ dt. \quad (8)$$

We remark that τ_D features prominently in both (5) and (7). In the remaining development, we pursue a solution for f_A and f_D under the assumption that both players randomize over a finite set of common action times.

IV. Atomic MSNE of the Redhands Game

The following definition details strategies in which both players use strategies that are atomic over a common set of support points.

Definition 4.1: For the Redhands game in which A attacks and D defends, we say that $f=(f_A,f_D)$ is a $(K+1)$-atom, τ_D-spaced strategy profile if

$$f_A(t)=\sum_{k=0}^{K}\alpha_{K,k}\delta(t-k\tau_D) \tag{9}$$

$$f_D(t)=\sum_{k=0}^{K}\pi_{K,k}\delta(t-k\tau_D), \tag{10}$$

where $\alpha_{K,k}$ and $\pi_{K,k}$ are the coefficients of the respective strategies.

In the rest of this section, we restrict attention to the Redhands game in which A and D randomize according to (9) and (10), respectively. We proceed by considering the implications using atomic strategies has on the best response function of the attacker, before performing a similar analysis in the case of the defender.

A. Player A

From (5) and (9), we see that when D plays according to (10) with weights $\pi_K=\{\pi_{K,k}\}_{k=0}^{K}$, then A's utility is given (abusing the argument) by

$$\bar{u}_A(t,\pi_K)=V_Ae^{-\beta_At}\begin{cases}\bar{p}-(\bar{p}-\underline{p})\pi_{K,k} & t\in[k\tau_D,(k+1)\tau_D),\\ \bar{p} & \text{else.}\end{cases} \tag{11}$$

The payoff function in (11) is strictly decreasing between atoms. Therefore, incentive compatability (from definition 3.5) only requires that all times in A's support have the same (maximizing) payoff, i.e., (i) for all $k=0,\dots,K$,

$$\bar{u}_A(k\tau_D,\pi_K)=V_A\gamma_A^{-k}\left[\bar{p}-(\bar{p}-\underline{p})\pi_{K,k}\right]=v_A(K) \tag{12}$$

where $v_A(K)$ is the payoff in question, which we call the *conformity payoff*, and (ii) the *maximum deviation payoff* $v_{A,\text{dev}}$ satisfies

$$v_{A,\text{dev}}(K)\doteq\bar{u}_A((K+1)\tau_D,\pi_K)=V_A\gamma_A{}^{-(K+1)}\bar{p}\le v_A(K). \tag{13}$$

As the name implies, the maximum deviation payoff is the largest payoff A can receive if he deviates from playing according to (9). In (12) and (13), we have introduced the *attacker impatience factor* $\gamma_A=e^{\beta_A\tau_D}$, for brevity. We will denote the *defender impatience factor* $\gamma_D=e^{\beta_D\tau_D}$ (note both are expressed with the *defender's* guard duration τ_D). We use (12) along with the constraint $\mathbf{1}^T\pi_K=1$, i.e., the normality constraint for probability densities, to determine $v_A(K)$ according to

$$v_A(K)=V_A\left[K\bar{p}+\underline{p}\right]\left(\sum_{k=0}^{K}\gamma_A{}^k\right)^{-1}=V_A\left[K\bar{p}+\underline{p}\right]\Gamma_{A,K}^{-1}, \tag{14}$$

where we let

$$\Gamma_{A,K}\doteq\sum_{k=0}^{K}\gamma_A{}^k=\frac{\gamma_A{}^{K+1}-1}{\gamma_A-1}.$$

Finally, we can use (12) and (14) to solve for π_K, obtaining

$$\pi_{K,k}=\frac{\bar{p}-\left[K\bar{p}+\underline{p}\right]\Gamma_{A,K}^{-1}\gamma_A{}^k}{\bar{p}-\underline{p}},\ k=0,\dots,K. \tag{15}$$

B. Player D

We now perform a similar analysis for player D. Specifically, from (7) and (10), we may write D's utility when A plays according to (9) with the weights $\alpha_K=\{\alpha_{K,k}\}_{k=0}^{K}$ as

$$u_D(t,\alpha_K)=\kappa(\alpha_K)+\begin{cases}(\bar{p}-\underline{p})\left(c_D+V_D\gamma_D{}^{-k}\right)\alpha_{K,k}, & \\ \quad\text{if }(k-1)\tau_D<t\le k\tau_D & \\ 0 \quad \text{else.}\end{cases} \tag{16}$$

This payoff function is piece-wise constant, so for some *conformity payoff* $v_D(K)$, incentive compatibility requires

$$u_D(k\tau_D,\alpha_K)=\kappa(\alpha_K)+(\bar{p}-\underline{p})\left(c_D+V_D\gamma_D{}^{-k}\right)\alpha_{K,k}=v_D(K) \tag{17}$$

for all $k=0,\dots,K$. Note the term $\kappa(\alpha_K)$ is a function of α_K alone and not of t. Therefore, from the structure of (16), initiating a guard at any time outside the support will not be strictly profitable to the defender. It can be shown, again using normality requirements for probability densities, that

$$\alpha_{K,k}=\left(c_D+V_D\gamma_D^{-k}\right)^{-1}/N(\alpha_K), \tag{18}$$

where $N(\alpha_K)=\sum_{k=0}^{K}\left(c_D+V_D\gamma_D^{-k}\right)^{-1}$ is the normalization factor for the profile α_K. Finally, we complete the characterization of D's strategy by noting that D's conformity payoff, $v_D(K)$, may be expressed as

$$v_D(K)=\kappa(\alpha_K)+(\bar{p}-\underline{p})/N(\alpha_K). \tag{19}$$

C. An Algorithm to Compute an Atomic MSNE

From the preceding discussion, if we assume that A and D play atomically over the support $\{0,\tau_D,\dots,K\tau_D\}$ (note K is specified), then determining the associated weights, α_K and π_K reduces to solving a linear system. However, to establish the existence of a MSNE, it remains to check that for some K the solution is both (i) realizable, in the sense that the atomic weights π_K are positive, and (ii) incentive compatible for the attacker, i.e., A cannot strictly improve his payoff by playing somewhere outside the support.

Below, we propose an enumerative algorithm to determine an atomic MSNE for the Redhands game. Verification of the existence of an atomic equilibria, its uniqueness, and the algorithm's correctness is provided in the next section.

Algorithm *RedhandsNash*
Output: atomic weights of MSNE for the Redhands game

```
1.  K ← 0, nashFlag ← false
2.  while nashFlag ≠ true
3.        solve for v_A(K) using (14)
4.        if v_A(K) ≥ v_A,dev(K)
5.          then found K for Nash
6.                 nashFlag ← true
7.          else K ← K+1
8.  solve for π_K,k using (15)
9.  solve for α_K,k using (18)
10. return α_K and π_K
```

Given the atomic weights supplied by *RedhandsNash*, the probability distributions $f_A(t_A)$ and $f_D(t_D)$ are readily computed from (9) and (10), respectively.

D. Existence and Uniqueness of Atomic MSNE

In Section III-B, we determined the conditions a MSNE must satisfy. We subsequently specified these results for the particular case of atomic strategy profiles, and presented the *RedhandsNash* algorithm to compute an atomic MSNE.

Now we undertake the unification of our previous efforts. Specifically, we show that (i) there exists an atomic MSNE, and (ii) it is generally unique. Establishing these results guarantees that *RedhandsNash*—by way of enumeration—computes the atomic MSNE and terminates, thereby certifying completeness of the algorithm.

A Nash candidate, in the class under consideration, consists of a pair of atomic strategies, f_A and f_D, where each strategy is *realizable* (i.e., a valid pdf) and incentive compatible with its opponent strategy. In what follows, we verify that these conditions are uniquely satisfiable.

We observe, from (18), that A's weights are guaranteed to be positive. Also, from (17), D has no incentive to deviate from (10), as there is no point in time, outside of his support, that provides a strictly better payoff. Therefore, it remains only to ensure the incentive compatibility condition for A, and that all $\pi_{K,k} \geq 0$. Regarding the latter condition, we focus on the sign of $\pi_{K,K}$, the smallest atom in the defender's strategy. From (15), we have

$$\pi_{K,K} = \frac{\bar{p} - [K\bar{p} + \underline{p}]\,\Gamma_{A,K}^{-1}{\gamma_A}^K}{\bar{p} - \underline{p}}. \tag{20}$$

The conditions we need to check for K to form a MSNE are then (i) $\pi_{K,K} \geq 0$ and (ii) $v_A(K) \geq v_{A,\text{dev}}(K)$.

Lemma 4.2: There exists $K^* \geq 0$ such that the strategy π_K is realizable, i.e. $\pi_{K,K} \geq 0$, for all $K \leq K^*$, and for no $K > K^*$.

Proof: The proof is based on an argument showing that the weight $\pi_{K,K}$ of the last atom of the defender's strategy is strictly decreasing in K. It follows that a sequence which is strictly decreasing from $\pi_{0,0} = 1$, and is unbounded below, must have exactly one zero-crossing. Using (20), we have

$$\begin{aligned}
(\bar{p} - \underline{p})&\pi_{K+1,K+1} \\
&= \bar{p} - [(K+1)\,\bar{p} + \underline{p}]\,\Gamma_{A,K+1}^{-1}{\gamma_A}^{K+1} \\
&= \bar{p} - [K\bar{p} + \underline{p}]\,\Gamma_{A,K}^{-1}{\gamma_A}^K \\
&\quad + [K\bar{p} + \underline{p}]\left(\Gamma_{A,K}^{-1}{\gamma_A}^K - \Gamma_{A,K+1}^{-1}{\gamma_A}^{K+1}\right) \\
&\quad - \bar{p}\Gamma_{A,K+1}^{-1}{\gamma_A}^{K+1}.
\end{aligned}$$

Using (20) again, and recognizing that $\underline{p} < \bar{p}$, we have

$$\begin{aligned}
(\bar{p} - \underline{p})&\pi_{K+1,K+1} \\
&< (\bar{p} - \underline{p})\pi_{K,K} \\
&\quad + \bar{p}\left[(K+1)\Gamma_{A,K}^{-1}{\gamma_A}^K - (K+2)\Gamma_{A,K+1}^{-1}{\gamma_A}^{K+1}\right] \\
&= (\bar{p} - \underline{p})\pi_{K,K} \\
&\quad + \bar{p}\Gamma_{A,K}^{-1}{\gamma_A}^K\left[(K+1) - (K+2)\frac{\gamma_A\Gamma_{A,K}}{\Gamma_{A,K+1}}\right].
\end{aligned}$$

We then substitute $\gamma_A\Gamma_{A,K} = \Gamma_{A,K+1} - 1$, to obtain

$$\begin{aligned}
(\bar{p} - \underline{p})&\pi_{K+1,K+1} \\
&< (\bar{p} - \underline{p})\pi_{K,K} \\
&\quad + \bar{p}\Gamma_{A,K}^{-1}{\gamma_A}^K\left\{(K+1) - (K+2)\left[1 - \Gamma_{A,K+1}^{-1}\right]\right\} \\
&= (\bar{p} - \underline{p})\pi_{K,K} \\
&\quad + \bar{p}\Gamma_{A,K}^{-1}{\gamma_A}^K\left\{(K+2)\Gamma_{A,K+1}^{-1} - 1\right\}.
\end{aligned}$$

Finally, with the observation that

$$\Gamma_{A,K+1} = \sum_{k=0}^{K+1}{\gamma_A}^k > \sum_{k=0}^{K+1} 1 = K+2,$$

we obtain

$$\pi_{K+1,K+1} < \pi_{K,K}.$$

This proves the lemma. ■

Lemma 4.3: The atomic strategy of length K is realizable with $\pi_{K,K} > 0$ if and only if the atomic strategy of length $K-1$ is not incentive compatible.

Proof: This proof results from the rearrangement of the incentive compatibility condition of (13). If $K-1$ is incentive compatible, then from (12) and (13)

$$\begin{aligned}
\bar{p} &\leq [(K-1)\,\bar{p} + \underline{p}]\,\frac{{\gamma_A}^{K+1} - {\gamma_A}^K}{{\gamma_A}^K - 1} \\
&\implies \bar{p}\left({\gamma_A}^K - 1\right) \leq [K\bar{p} + \underline{p}]\left({\gamma_A}^{K+1} - {\gamma_A}^K\right) \\
&\qquad\qquad - \bar{p}\left({\gamma_A}^{K+1} - {\gamma_A}^K\right) \\
&\implies \bar{p}\left({\gamma_A}^{K+1} - 1\right) \leq [K\bar{p} + \underline{p}]\left({\gamma_A}^{K+1} - {\gamma_A}^K\right) \\
&\implies \bar{p} \leq [K\bar{p} + \underline{p}]\,\frac{{\gamma_A}^{K+1} - {\gamma_A}^K}{{\gamma_A}^{K+1} - 1}
\end{aligned}$$

Meanwhile, if $\pi_{K,K} > 0$ then from (20) we can write

$$\bar{p} > [K\bar{p} + \underline{p}]\,\frac{{\gamma_A}^{K+1} - {\gamma_A}^K}{{\gamma_A}^{K+1} - 1}.$$

We observe that these inequalities are strongly complementary, proving the lemma. ■

The following proposition ties together the previous results to verify that *RedhandsNash* returns a MSNE, thereby establishing the existence of a MSNE for the Redhands game.

Proposition 4.4: There exists K such that the output of *RedhandsNash* is a MSNE.

Proof: We consider the *last* realizable K strategy, which exists by Lemma 4.2. By definition, the $K+1$ strategy is not realizable, and so the K strategy is incentive compatible by Lemma 4.3. There is therefore a finite K and MSNE of the form (9) and (10) for which *RedhandsNash* returns the associated atomic weights α_K and π_K. ■

The following proposition addresses the issue of uniqueness of the MSNE returned by the algorithm *RedhandsNash*.

Proposition 4.5: The strategy profile returned by the algorithm *RedhandsNash* is a *unique* τ-spaced atomic Nash equilibrium (almost certainly).

Proof: As seen in Proposition 4.4, for some K, the $K-1$ incentive compatibility and the K realizability (strict) conditions are strongly complementary. Therefore, if K induces a Nash equilibrium, then for all positive m, the $K-m$ strategies are not incentive compatible, and likewise $K+m$ strategies are not realizable.
Caveat: In the event $\pi_{K,K} \geq 0$ holds with equality for some K, then $K+1$ (the *second* incentive compatible solution) will also be a MSNE. ■

V. Numeric Example

In this section, we specify a particular set of player parameters and illustrate the characteristics of the unique atomic MSNE for the associated Redhands game. The relevant parameters used are listed as follows: $\bar{p}=1$, $\underline{p}=0$, $c_D=10$, $\tau_D=1$, $V_A=V_D=1$, $\beta_A=0.1$, and $\beta_D=1$. Referring to Figure 1(a), the blue plots illustrate the weights used by each player in the unique atomic equilibrium. The green curves in Figure 1(a) demonstrate that neither player can improve their utility by unilaterally deviating. The curves in Figure 1(b) reveal the unique atomic MSNE has each player mixing over a support comprised of five points. For completeness, the weights associated with the MSNE are:

$$\alpha_K = \{0.187, 0.199, 0.203, 0.205, 0.206\}$$
$$\pi_K = \{0.352, 0.283, 0.208, 0.125, 0.033\}.$$

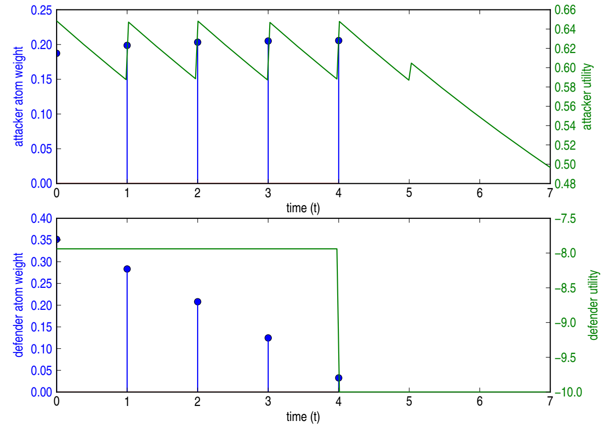

(a) An equilibrium for a representative version of the Redhands game. Note each's players utility is maximized at each of the time points in their support.

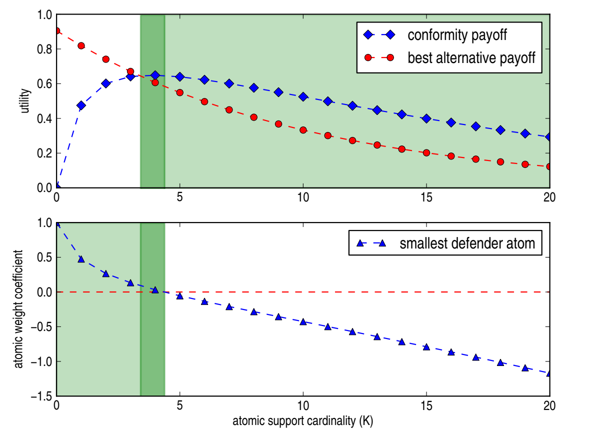

(b) Graphical evidence demonstrating the existence and uniqueness of the atomic MSNE.

Fig. 1. Simulation Results

VI. Conclusions

A. Summary

This work has explored equilibria in the Redhands game, a timed game in which an attacker and a defender compete in an adversarial setting with the goal of maximizing their respective utility. We used best-response functions to characterize a mixed-strategy Nash equilibrium in which both players mix over a common atomic support. This equilibrium was shown to be unique among the class, and an enumerative algorithm for computing the solution was provided. Numeric results highlighted interesting aspects of equilibrium strategies for a representative game instance.

B. Future Work

There are a number of directions in which to expand the research discussed in this paper. It remains a foremost goal to consider a dynamic version of the Redhands game, where players may take timed actions in sequence, as a study of games on timed automata.

In a similar vein, a valuable investigation would be to augment Redhands with observers, e.g. costly sensors, spies, or early-warning alarms (perhaps noisy), and consider explicitly the tradeoff between measured data and prior knowledge about players' predilections. One might argue that with good sensors, prior knowledge of the opponent's private information (which is generally considered difficult to measure) may be rendered unimportant. However, the relative value of measurements versus prior knowledge is less clear if players may employ deceitful tactics.

Finally, we remark that, although we demonstrated the uniqueness of the uniformly-space atomic MSNE within its class, we have left open whether MSNE may exist outside this class. It would be natural, therefore, to attempt either to certifying uniqueness of the proposed MSNE, or construct others.

References

[1] D. T. Bishop and C. Cannings, "A generalized war of attrition," *Journal of Theoretical Biology*, vol. 70, no. 1, pp. 85–124, 1978.
[2] D. Fudenberg and J. Tirole, *Game theory. 1991*. MIT Press.
[3] T. Başar and G. Olsder, *Dynamic Noncooperative Game Theory. 1995*. SIAM.
[4] P. Ramadge and W. Wonham, "Supervisory control of a class of discrete event processes," in *Analysis and Optimization of Systems*, 1984, pp. 475–498.
[5] K. M. Passino and P. J. Antsaklis, "On the optimal control of discrete event systems," in *Proceedings of the 28th IEEE Conference on Decision and Control*, 1989, pp. 2713–2718.
[6] L. Alfaro, M. Faella, T. A. Henzinger, R. Majumdar, and M. Stoelinga, "The element of surprise in timed games," *CONCUR 2003-Concurrency Theory*, pp. 144–158, 2003.
[7] P. Bouyer, F. Cassez, E. Fleury, and K. G. Larsen, "Optimal strategies in priced timed game automata."
[8] P. Dasgupta and E. Maskin, "The existence of equilibria in discontinuous economic games, i: Theory," *SIAM Review of Economic Studies*, vol. 53, no. 1, pp. 1–26, 1986.