Two Player Non Zero-sum Stopping Games in Discrete Time

Eran Shmaya* and Eilon Solan†‡

May 15, 2002

Abstract

We prove that every two player non zero-sum stopping game in discrete time admits an ϵ-equilibrium in randomized strategies, for every $\epsilon > 0$.

We use a stochastic variation of Ramsey Theorem, which enables us to reduce the problem to that of studying properties of ϵ-equilibria in a simple class of stochastic games with finite state space.

Keywords: Stopping games, Dynkin games, stochastic games.

MSC 2000 subject classification: primary - 60G40, secondary - 91A15, 91A05.

*The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. e-mail: gawain@post.tau.ac.il

†Kellogg School of Management, Northwestern University, *and* the School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. e-mail: eilons@post.tau.ac.il

‡The research of the second author was supported by the Israel Science Foundation (grant No. 03620191).

1 Introduction

Consider the following optimization problem, that was presented by Dynkin (1969). Two players observe a realization of two real-valued processes (x_n) and (r_n). Player 1 can *stop* whenever $x_n \geq 0$, and player 2 can *stop* whenever $x_n < 0$. At the first stage θ in which one of the players stops, player 2 pays player 1 the amount r_θ, and the process terminates. If no player ever stops, player 2 does not pay anything.

A *strategy* of player 1 is a stopping time μ that satisfies $\{\mu = n\} \subseteq \{x_n \geq 0\}$. A strategy ν of player 2 is defined analogously. The termination stage is simply $\theta = \min\{\mu, \nu\}$. For a given pair (μ, ν) of strategies, denote by

$$\gamma(\mu, \nu) = \mathbf{E}[1_{\{\theta < +\infty\}} r_\theta]$$

the expected payoff of player 1.

Dynkin (1969) proved that if $\sup_{n \geq 0} |r_n| \in L_1$ this problem has a value v; that is

$$v = \sup_\mu \inf_\nu \gamma(\mu, \nu) = \inf_\nu \sup_\mu \gamma(\mu, \nu).$$

He moreover characterized ϵ-optimal strategies; that is, strategies μ (*resp.* ν) that achieve the supremum (*resp.* the infimum) up to ϵ.

Neveu (1975) generalized this problem by allowing both players to stop at every stage, and by introducing 3 real valued processes $(r_{\{1\},n})$,$(r_{\{2\},n})$ and $(r_{\{1,2\},n})$. The payoff player 2 pays player 1 is defined by

$$\gamma(\mu, \nu) = \mathbf{E}[1_{\{\mu < \nu\}} r_{\{1\},\mu} + 1_{\{\mu > \nu\}} r_{\{2\},\nu} + 1_{\{\mu = \nu < +\infty\}} r_{\{1,2\},\mu}].$$

He then proved that this problem has a value, provided (a) $\sup_{n \geq 0} \max\{|r_{\{1\},n}|, |r_{\{2\},n}|, |r_{\{1,2\},n}|\} \in L_1$, and (b) $r_{\{1\},n} = r_{\{1,2\},n} \leq r_{\{2\},n}$.

Recently Rosenberg et al (2001) studied games in Neveu's setup, but they allowed the players to use *randomized* stopping times; a strategy is a $[0, 1]$-valued process, that dictates the probability by which the player stops at every stage. They prove that the problem has a value, assuming only condition (a).

A broad literature provides sufficient conditions for the existence of the value in continuous time (see, e.g., Bismuth (1979), Alario-Nazaret, Lepeltier and Marchal (1982), Lepeltier and Maingueneau (1984) , Touzi and Vieille (2002)). Some authors work in the diffusion case, see e.g. Cvitanić and Karatzas (1996).

The non zero-sum problem in discrete time was studied, amongst others, by Mamer (1987), Morimoto (1986) and Ohtsubo (1987, 1991). In the non zero-sum case, the processes $(r_{\{1\},n})$,$(r_{\{2\},n})$ and $(r_{\{1,2\},n})$ are $\mathbf{R}^2$-valued, and the expected payoff of player i, $i = 1, 2$, is

$$\gamma^i(\mu,\nu) = \mathbf{E}_{\mu,\nu}[1_{\{\mu<\nu\}} r^i_{\{1\},\mu} + 1_{\{\mu>\nu\}} r^i_{\{2\},\nu} + 1_{\{\mu=\nu<+\infty\}} r^i_{\{1,2\},\mu}].$$

The goal of each player is to maximize his own expected payoff. Given $\epsilon > 0$, a pair of stopping times (μ, ν) is an *ϵ-equilibrium* if for every pair of stopping times (μ', ν'),

$$\gamma^1(\mu,\nu) \geq \gamma^1(\mu',\nu) - \epsilon, \text{ and } \gamma^2(\mu,\nu) \geq \gamma^2(\mu,\nu') - \epsilon.$$

The above mentioned authors provide various sufficient conditions under which ϵ-equilibria exist.

In the present paper we study two player non zero-sum problems in discrete time with randomized stopping times, and we prove the existence of an ϵ-equilibrium for every $\epsilon > 0$, under condition (a). Our technique is based on a stochastic variation of Ramsey Theorem (Ramsey (1930)), which states that for every coloring of a complete infinite graph by finitely many colors there is a complete infinite monochromatic subgraph. This variation serves as a substitute for a fixed point argument, which is usually used to prove existence of an equilibrium. It allows us to reduce the problem of existence of ϵ-equilibrium in a general stopping game to that of studying properties of ϵ-equilibria in a simple class of stochastic games with finite state space.

The paper is arranged as follows. In section 2 we provide the model and the main result. A sketch of the proof appears in section 3. In section 4 we present a stochastic variation of Ramsey Theorem. In section 5 we show that to prove existence of ϵ-equilibria in a general stopping game, it is sufficient to consider a restricted class of stopping games. In section 6 we define the notion of games played on a finite tree, and we study some of their properties. In section 7 we construct an ϵ-equilibrium. We end by discussing extensions to more than two players in section 8.

2 The Model and the Main Result

A *two-player non zero-sum stopping game* is a 5-tuple $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ where:

- $(\Omega, \mathcal{A}, p)$ is a probability space.
- $\mathcal{F} = (\mathcal{F}_n)_{n\geq 0}$ is a filtration over $(\Omega, \mathcal{A}, p)$.
- $R = (R_n)_{n\geq 0}$ is a $\mathcal{F}$-adapted $\mathbf{R}^6$-valued process. The coordinates of R_n are denoted by $R^i_{Q,n}$, $i = 1, 2$, $\phi \neq Q \subseteq \{1, 2\}$.

A (behavior) *strategy* for player 1 (*resp.* player 2) is a $[0, 1]$-valued $\mathcal{F}$-adapted process $x = (x_n)_{n\geq 0}$ (*resp.* $y = (y_n)_{n\geq 0}$). The interpretation is that x_n (*resp.* y_n) is the probability by which player 1 (*resp.* player 2) stops at stage n.

Let θ be the first stage, possibly infinite, in which at least one of the players stops, and let $\phi \neq Q \subseteq \{1, 2\}$ be the set of players that stop at stage θ (provided $\theta < \infty$.) The expected payoff under (x, y) is given by

$$\gamma^i(x, y) = \mathbf{E}_{x,y}[R^i_{Q,\theta} 1_{\{\theta<\infty\}}], \tag{1}$$

where the expectation $\mathbf{E}_{x,y}$ is w.r.t the distribution $\mathbf{P}_{x,y}$ over plays induced by (x, y).

Definition 2.1. Let $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a non zero-sum stopping game, and let $\epsilon > 0$. A pair of strategies (x^*, y^*) is an *ϵ-equilibrium* if $\gamma^1(x, y^*) \leq \gamma^1(x^*, y^*) + \epsilon$ and $\gamma^2(x^*, y) \leq \gamma^2(x^*, y^*) + \epsilon$, for every x and y.

The main result of the paper is the following.

Theorem 2.2. *Let $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a two-player stopping game such that $\sup_{n\geq 0} \|R_n\|_\infty \in L^1(p)$. Then for every $\epsilon > 0$, the game admits an ϵ-equilibrium.*

3 Sketch of the Proof

In the present section we provide the main ideas of the proof. Let Γ be a stopping game. To simplify the presentation, assume that $\mathcal{F}_n$ is trivial for every n, so that the payoff process is deterministic. Assume also that payoffs are uniformly bounded by 1.

Given $\epsilon > 0$, fix a finite covering M of the space of payoffs $[-1, 1]^2$ by sets with diameter smaller than ϵ.

For every two non negative integers $k < l$ define the periodic game $G(k, l)$ as the game that starts at stage k, and, if not stopped earlier, restarts at stage l. Formally, $G(k,l)$ is a stopping game in which the terminal payoff at stage n is equal to the payoff at stage $k+(n \mod l)$ in Γ.

This periodic game is a simple stochastic game (see, e.g., Shapley (1953), or Flesch et al (1996)), and is known to admit an ϵ-equilibrium in periodic strategies. Assign to each pair of non negative integers $k < l$ an element $m(k,l) \in M$ which contains a periodic ϵ-equilibrium payoff of the periodic game $G(k,l)$.

Thus, we assigned to each $k < l$ a color $m(k,l) \in M$. A consequence of Ramsey Theorem is that there is an increasing sequence of integers $0 \leq k_1 < k_2 < \cdots$ such that $m(k_1,k_2) = m(k_n,k_{n+1})$ for every n.

Assume first that $k_1 = 0$. A naive candidate for a 4ϵ-equilibrium suggests itself: between stages k_n and k_{n+1}, the players follow a periodic ϵ-equilibrium in the game $G(k_n,k_{n+1})$ with corresponding payoff in the set $m(k_1,k_2)$.

So that this strategy pair is indeed 4ϵ-equilibrium, one has to study properties of the ϵ-equilibria in periodic games. The complete solution of this case appears in Shmaya et al (2002), where it is observed that in each periodic game $G(k,l)$ there exists a periodic ϵ-equilibrium that satisfies one of the following conditions. (i) In this ϵ-equilibrium, no player ever stops. (ii) Under this ϵ-equilibrium, both players receive non negative payoff, and termination occurs in each period with probability at least ϵ^2. (iii) If some player receives in this ϵ-equilibrium a negative payoff, then his opponent stops in each period with probability at least ϵ^2. The fact that at least one of these conditions hold is sufficient to prove that the concatenation described above is a 4ϵ-equilibrium, with corresponding payoff in the convex hull of $m(k_1,k_2)$.

If $k_1 > 0$, choose an arbitrary $m \in m(k_1,k_2)$. Between stages 0 and k_1, the players follow an equilibrium in the k_1-stage game with terminal payoff m; that is, if no player ever stops before stage k_1, the payoff is m. From stage k_1 and on, the players follow the strategy described above. It is easy to verify that this strategy pair forms a 5ϵ-equilibrium.

When the payoff process is general, few difficulties appear. First, a periodic game is defined now by two stopping times $\mu_1 < \mu_2$; μ_1 indicates the initial stage, and μ_2 indicates when the game restarts. So that we can analyze this periodic game, we have to reduce the problem to the case where the σ-algebras $\mathcal{F}_{\mu_1}, \mathcal{F}_{\mu_1+1}, \ldots, \mathcal{F}_{\mu_2}$ are finite. This is done in section 7.

Second, we have to generalize Ramsey Theorem to this stochastic setup. This is done in section 4.

Third, we have to study properties of ϵ-equilibria in these periodic games, so that a proper concatenation of ϵ-equilibria in the different periodic games would generate a 4ϵ-equilibrium in the original game. This is done in section 6.

4 A Stochastic Variation of Ramsey Theorem

In the present section we provide a stochastic variation of Ramsey Theorem. Let $(\Omega, \mathcal{A}, p)$ be a probability space and $\mathcal{F} = (\mathcal{F}_n)_{n\geq 0}$ a filtration. All stopping times that appear in the sequel are $\mathcal{F}$-adapted. For every set $A \subseteq \Omega$, $A^c = \Omega \setminus A$ is the complement of A. For every $A, B \in \mathcal{A}$, A *holds on* B if and only if $p(A^c \cap B) = 0$.

Definition 4.1. A *NT-function* is a function that assigns to every integer $n \geq 0$ and every bounded stopping time τ a $\mathcal{F}_n$-measurable r.v. that is defined over the set $\{\tau > n\}$. We say that a NT-function f is C*-valued*, for some set C, if the r.v. $f_{n,\tau}$ is C-valued, for every $n \geq 0$ and every τ.

Definition 4.2. A NT-function f is $\mathcal{F}$*-consistent* if for every $n \geq 0$, every $\mathcal{F}_n$-measurable set F, and every two bounded stopping times τ_1, τ_2, we have

$$\tau_1 = \tau_2 > n \text{ on } F \text{ implies } f_{n,\tau_1} = f_{n,\tau_2} \text{ on } F.$$

When f is a NT-function, and $\sigma < \tau$ are two bounded stopping times, we denote $f_{\sigma,\tau}(\omega) = f_{\sigma(\omega),\tau}(\omega)$. Thus $f_{\sigma,\tau}$ is a $\mathcal{F}_\sigma$-measurable r.v.

The main result of this section is the following.

Theorem 4.3. *For every finite set C of colors, every C-valued $\mathcal{F}$-consistent NT-function c, and every $\epsilon > 0$, there exists a sequence of bounded stopping times $0 \leq \theta_0 < \theta_1 < \theta_2 < \ldots$ such that $p(c_{\theta_0,\theta_1} = c_{\theta_1,\theta_2} = c_{\theta_2,\theta_3} = \ldots) > 1 - \epsilon$.*

Comment: The natural stochastic generalization of Ramsey Theorem requires the stronger condition $p(c_{\theta_0,\theta_1} = c_{\theta_i,\theta_j} \quad \forall 0 \leq i < j) \geq 1 - \epsilon$. We do not know whether this generalization is correct.

The following example shows that a sequence of stopping times $\theta_0 < \theta_1 < \theta_2 < \theta_3 < \ldots$ such that $p(c_{\theta_0,\theta_1} = c_{\theta_1,\theta_2} = \ldots) = 1$ need not exist even without the boundedness condition.

Example 4.4. Let X_n be a biased random walk on the integers, $X_0 = 0$ and $p(X_{n+1} = X_n + 1) = 1 - p(X_{n+1} = X_n - 1) = 3/4$. Let $\mathcal{F}_n = \sigma(X_0, X_1, \ldots, X_n)$. For every $n \geq 0$ let $R_n = \cup_{1\leq k\leq n}\{X_k = -1\}$; for every finite (but not necessarily bounded) stopping time τ define $c_{n,\tau}$ = Red on $R_n \cap \{\tau > n\}$ and $c_{n,\tau}$ = Blue on $R_n^c \cap \{\tau > n\}$. Since $p(\cup_{n\geq 0} R_n) < 1$, whereas for every finite stopping time θ and every $B \in \mathcal{F}_\theta$ $p(\cup_{n\geq 0} R_n \mid B) > 0$, it follows that for every sequence $\theta_0 < \theta_1 < \ldots$ of finite stopping times $p(c_{\theta_0,\theta_1} = \text{Blue}) > 0$ whereas $p(c_{\theta_0,\theta_1} = c_{\theta_1,\theta_2} = \ldots = \text{Blue} \mid c_{\theta_0,\theta_1} = \text{Blue}) < 1$.

We start by proving a slightly stronger version of Theorem 4.3 when $|C| = 2$.

Lemma 4.5. *Let $C = \{\text{Blue}, \text{Red}\}$, and let c be a C-valued $\mathcal{F}$-consistent NT-function. For every $\epsilon > 0$ there exist $N \in \mathbf{N}$, two sets $\bar{R}, \bar{B} \in \mathcal{F}_N$, and a sequence $N \leq \tau_0 < \tau_1 < \tau_2 < \ldots$ of bounded stopping times, such that:*

a) $\bar{R} = \bar{B}^c$.

b) $p(c_{\tau_0,\tau_1} = c_{\tau_1,\tau_2} = \ldots = \text{Red} \mid \bar{R}) > 1 - \epsilon$.

c) $p(c_{\tau_k,\tau_l} = \text{Blue} \quad \forall k, l \mid \bar{B}) > 1 - \epsilon$.

Proof. We claim first that for every $n \in \mathbf{N}$ one can find two sets $R_n, B_n \in \mathcal{F}_n$ and a bounded stopping time σ_n such that:

1. $p(R_n \cup B_n) > 1 - \frac{1}{2^n}$.
2. $\{\sigma_n > n\}$ on R_n and c_{n,σ_n} = Red on R_n.
3. For every bounded stopping time τ, $c_{n,\tau}$ = Blue on $B_n \cap \{\tau > n\}$.

To see this, fix $n \in \mathbf{N}$. Call a set $F \in \mathcal{F}_n$ *red* if there exists a bounded stopping time σ_F such that on F both $\sigma_F > n$ and c_{n,σ_F} = Red. Observe that since c is $\mathcal{F}$-consistent, if $F, G \in \mathcal{F}_n$ are red, then so is $F \cup G$. Let $\alpha = \sup_F \{p(F), F \in \mathcal{F}_n \text{ is red}\}$. For every $k \geq 1$ let $F_k \in \mathcal{F}_n$ be a red set such that $p(F_k) > \alpha - \frac{1}{k}$. Let $F_* = \cup_{k\geq 1} F_k$. Observe that $F_* \in \mathcal{F}_n$ and $p(F_*) = \alpha$. Moreover, no subset of F_*^c with positive probability is red. Let $R_n = F_{2^n}$, let σ_n be a bounded stopping time such that on R_n $\sigma_n > n$ and c_{n,σ_n} = Red, and let $B_n = F_*^c$. This concludes the proof of the claim.

Let $B = \{B_n \quad \text{i.o.}\}$, and set $R = B^c$. Since $R, B \in \bigvee_n \mathcal{F}_n$, there is $N \in \mathbf{N}$ and two sets $\bar{B}, \bar{R} \in \mathcal{F}_N$ such that (i) $\bar{R} = \bar{B}^c$, (ii) $p(B \mid \bar{B}) > 1 - \epsilon$, and (iii) $p(R \mid \bar{R}) > 1 - \epsilon$. On R, and therefore also on $R \cap \overline{R}$, both B_n and

$(B_n \cup R_n)^c$ occur only finitely many times. By sufficiently increasing N we assume w.l.o.g. that $p(\cap_{n\geq N} R_n \mid R \cap \bar{R}) > 1-\epsilon$. In particular,

$$p(\cap_{n\geq N} R_n \mid \bar{R}) > 1-2\epsilon. \tag{2}$$

Let $N = n_0 < n_1 < n_2 < \cdots$ be a sequence of integers such that, for every $k \geq 0$, $p(T_k \mid B \cap \bar{B}) > 1 - \frac{\epsilon}{2^{k+1}}$, where $T_k = \cup_{n_k \leq n < n_{i+k}} B_n$. Then $p(\cap_{k\geq 0} T_k \mid B \cap \bar{B}) > 1-\epsilon$, and therefore

$$p(\cap_{k\geq 0} T_k \mid \bar{B}) > 1-2\epsilon. \tag{3}$$

We now define the sequence $(\tau_k)_{k\geq 0}$ inductively, working separately on $\bar{R}$ and $\bar{B}$. Consider first the set $\bar{R}$. Define $\tau_0 = N$. Given τ_k, define $\tau_{k+1} = \Sigma_{n\in\mathbf{N}} \sigma_n 1_{\{\tau_k = n\}\cap R_n \cap \bar{R}}$ on $\bar{R} \cap \bigcup_n (\{\tau_k = n\} \cap R_n)$. Since τ_k and $(\sigma_n)_{n\geq 0}$ are bounded, τ_{k+1} can be extended to a bounded stopping time on $\bar{R}$. By definition $c_{\tau_0,\tau_1} = c_{\tau_1,\tau_2} = \ldots =$ Red on $\bar{R} \cap (\cap_{n\geq N} R_n)$, and it follows from (2) that $p(c_{\tau_0,\tau_1} = c_{\tau_1,\tau_2} = \ldots = \text{Red} \mid \bar{R}) \geq 1-2\epsilon$.

Consider now the set $\bar{B}$. Define $\tau_0 = N$. Define $\tau_{k+1}(w) = \min\{n_k \leq n < n_{k+1}, w \in B_n\}$ on $\bar{B} \cap T_k$, and $\tau_{k+1} = n_{k+1} - 1$ on $\bar{B} \setminus T_k$. By (c), for every $k, l \in \mathbf{N}$, $c_{\tau_k,\tau_l} =$ Blue on $\bar{B} \cap (\cap_{k\geq 0} T_k)$, and it follows from (3) that $p(c_{\tau_k,\tau_l} = \text{Blue} \quad \forall k, l \mid \bar{B}) > 1-2\epsilon$. □

Proof of Theorem 4.3

We prove the Theorem by induction on $|C|$. The case $|C| = 2$ follows from Lemma 4.5. Assume we have already proven the lemma whenever $|C| = r$ and assume $|C| = r+1$. Let Red be a color in C.

By considering all colors different from Red as a single color, and applying Lemma 4.5 there exist $N \in \mathbf{N}$, two sets $\bar{R}, \bar{B} \in \mathcal{F}_N$, and a sequence of stopping times $N \leq \tau_0 < \tau_1 < \ldots$ such that: (i) $\bar{R} = \bar{B}^c$, (ii) $p(c_{\tau_0,\tau_1} = c_{\tau_1,\tau_2} = \ldots = \text{Red} \mid \bar{R}) > 1-\epsilon/2$, and (iii) $p(c_{\tau_k,\tau_l} \neq \text{Red} \quad \forall k, l \mid \bar{B}) > 1-\epsilon/2$. We define θ_i separately on $\bar{R}$ and $\bar{B}$.

On $\bar{R}$, we let $\theta_i = \tau_i$.

We now restrict ourselves to the space $(\bar{B}, \mathcal{A}_{\bar{B}}, p_{\bar{B}})$ with the filtration $\mathcal{G}_n = \mathcal{F}_{\tau_n} \cap \bar{B}$. Let $\tilde{c}$ be the C-valued NT function over $\mathcal{G}$ defined by $\tilde{c}_{n,\beta} = c_{\tau_n,\tau_\beta}$ for every $\mathcal{G}$-adapted stopping time β, where $\tau_\beta = \sum_n \tau_n 1_{\{\beta=n\}}$ is an $\mathcal{F}$-adapted stopping time. Let c' be the coloring that is obtained from $\widetilde{c}$ by changing the color Red with another color in C, say Green:

$$c'_{n,\beta} = \begin{cases} \tilde{c}_{n,\beta}, & \text{if } \tilde{c}_{n,\beta} \neq \text{Red} \\ \text{Green}, & \text{if } \tilde{c}_{n,\beta} = \text{Red} \end{cases}.$$

As c' is a $C \setminus \{\text{Red}\}$-valued $\mathcal{G}$-consistent NT-function, one can apply the induction hypothesis, and obtain a sequence of $\mathcal{G}$-adapted stopping times $0 \le \beta_0 < \beta_1 < \beta_2 < \ldots$ such that

$$p(c'_{\beta_0,\beta_1} = c'_{\beta_1,\beta_2} = \ldots \mid \bar{B}) > 1 - \epsilon/2. \tag{4}$$

By (4) and (iii) it follows that $p(\tilde{c}_{\beta_0,\beta_1} = \tilde{c}_{\beta_1,\beta_2} = \ldots \mid \bar{B}) > 1 - \epsilon$. We define $\theta_i = \tau_{\beta_i}$ on $\bar{B}$. Thus

$$p(c_{\theta_0,\theta_1} = c_{\theta_1,\theta_2} = \ldots \mid \bar{B}) > 1 - \epsilon. \tag{5}$$

Combining (ii) and (5) we get $p(c_{\theta_0,\theta_1} = c_{\theta_1,\theta_2} = \ldots) > 1 - \epsilon$, as desired.

5 Restricting the Class of Games

In the present section we show that to prove Theorem 2.2 it is sufficient to consider a restricted class of stopping games.

Definition 5.1. Let $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a stopping game and $B \in \mathcal{A}$ with $p(B) > 0$. The *game restricted to* B is the stopping game $\Gamma_B = (B, \mathcal{A}_B, p_B, \mathcal{F}_B, R)$, where $\mathcal{A}_B = \{A \cap B, B \in \mathcal{A}\}$, $p_B(A) = p(A \mid B)$ for every $A \in \mathcal{A}_B$, and $\mathcal{F}_{B,n} = \{F \cap B, F \in \mathcal{F}_n\}$ for every $n \ge 0$.

The following lemma is standard.

Lemma 5.2. *Let $(\Omega, \mathcal{A})$ be a measurable space and let $B \subseteq \Omega$. Let $\mathcal{A}_B = \{A \cap B, A \in \mathcal{A}\}$. Then for every $\mathcal{A}_B$-measurable function x on B there exists a $\mathcal{A}$-measurable function x^* on Ω such that $x^* = x$ on B.*

Set $m = \sup_{n \ge 0} \max\{|R^i_{Q,n}|, i = 1, 2, \emptyset \subset Q \subseteq \{1, 2\}\})$. Since $m \in L^1(p)$, for every $\epsilon > 0$ there exists $N \in \mathbf{N}$ such that

$$\mathbf{E}[m 1_{\{m > N\}}] < \epsilon. \tag{6}$$

Set $B = \{m \le N\}$. By Lemma 5.2 and (6), any ϵ-equilibrium in Γ_B can be extended to a 3ϵ-equilibrium in Γ. In particular, it is sufficient to consider games in which the payoff process is uniformly bounded. We further assume w.l.o.g. that the payoff process is uniformly bounded by 1.

Lemma 5.2 gives us the following.

Lemma 5.3. *Let $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a stopping game with payoffs bounded by 1, let $B \in \mathcal{A}$ such that $p(B) > 1-\delta$, and let $\epsilon > 0$. Assume that the game Γ_B admits an ϵ-equilibrium. Then the game Γ admits an $\epsilon + 2\delta$-equilibrium.*

Definition 5.4. Let $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a stopping game, and let τ be a bounded $\mathcal{F}$-adapted stopping time. The game that *starts at* τ is defined by $\Gamma_\tau = (\Omega, \mathcal{A}, p, \tilde{\mathcal{F}}, \tilde{R})$, where $\tilde{\mathcal{F}}_n = \mathcal{F}_{\tau+n}$ and $\tilde{R}_n = R_{\tau+n}$ for every $n \geq 0$.

In particular for every bounded $\mathcal{F}$-adapted stopping time τ, and every $B \in \mathcal{A}$, $\Gamma_{B,\tau}$ is the game restricted to B that starts at τ.

Lemma 5.5. *Let $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a stopping game, τ a bounded $\mathcal{F}$-adapted stopping time, and $\epsilon > 0$. Let $(B_1, \ldots, B_k)$ be a finite $\mathcal{F}_\tau$-measurable partition of Ω. Suppose that the games $\Gamma_{B_i,\tau}$, $1 \leq i \leq k$, admit ϵ-equilibria. Then the game Γ admits an ϵ-equilibrium.*

Proof. Let (x_i, y_i) be an ϵ-equilibrium in $\Gamma_{B_i,\tau}$, $1 \leq i \leq k$. Consider the strategy profile (x, y) for the game Γ_τ defined by $x = x_i$ and $y = y_i$ on B_i. Then (x, y) is an ϵ-equilibrium in Γ_τ.

For $i = 1, 2$, let $\gamma^i = \mathbf{E}_{x,y}[R^i_{Q,\theta} 1_{\{\theta < \infty\}} \mid \mathcal{F}_\tau]$ be the payoff to player i in Γ_τ conditioned on $\mathcal{F}_\tau$. By standard tools in dynamic programming, the finite stage game which, if no player stops before stage τ, terminates at stage τ with terminal payoff $\gamma = (\gamma^1, \gamma^2)$, has an equilibrium $(\bar{x}, \bar{y})$. Following $(\bar{x}, \bar{y})$ up to stage τ and (x, y) afterwards forms an ϵ-equilibrium in Γ. $\square$

The main result of this section is the following.

Proposition 5.6. *Suppose that every stopping game $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ that satisfies* **A.1**-**A.6** *below admits an ϵ-equilibrium, for every $\epsilon > 0$. Then Theorem 2.2 holds.*

A.1: *There exists $K \in \mathbf{N}$ such that for every $n \geq 0$, $R_n \in \{0, \pm\frac{1}{K}, \pm\frac{2}{K}, \ldots, \pm\frac{K}{K}\}^6$.*

A.2: $R^1 := \limsup_{n\to\infty} R^1_{\{1\},n}$ *is constant, and* $R^1_{\{1\},n} \leq R^1$ *for every* $n \geq 0$.

A.3: $R^2 := \limsup_{n\to\infty} R^2_{\{2\},n}$ *is constant, and* $R^2_{\{2\},n} \leq R^2$ *for every* $n \geq 0$.

A.4: $R^1 > 0$ *or* $R^2 > 0$.

A.5: $R^2_{\{1\},n} < R^2$ *whenever* $R^1_{\{1\},n} = R^1$.

A.6: $R^1_{\{2\},n} < R^1$ *whenever* $R^2_{\{2\},n} = R^2$.

Proof. Assume that every stopping game $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ that satisfies **A.1-A.6** admits an ϵ-equilibrium, for every $\epsilon > 0$. Let $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a stopping game with payoffs uniformly bounded by 1, and let $\epsilon > 0$. We prove that Γ admits a $C\epsilon$-equilibrium for some $C > 0$ by gradually reducing Γ to a game that satisfies **A.1-A.6**

If (x, y) is an ϵ-equilibrium in some stopping game $\Gamma = (\Omega, \mathcal{A}, p, \mathcal{F}, R)$, it is a 3ϵ-equilibrium in every stopping game $\Gamma' = (\Omega, \mathcal{A}, p, \mathcal{F}, R')$ such that $\|R'_n - R_n\|_\infty < \epsilon$ for every $n \geq 0$. By choosing $K > 1/\epsilon$, one can therefore assume **A.1**.

Let $R^1 = \limsup R^1_{\{1\},n}$. R^1 is a $\bigvee_{n \geq 0} \mathcal{F}_n$-measurable r.v. that by **A.1** admits finitely many values. Let N be large enough, and let $\{\bar{B}_t\}_{t \in \{0, \pm\frac{1}{K}, \pm\frac{2}{K}, \ldots, \pm\frac{K}{K}\}}$ be a $\mathcal{F}_N$-measurable partition of Ω such that $p(R^1 = t \mid \bar{B}_t) > 1 - \epsilon$.[1] By sufficiently increasing N we assume w.l.o.g that, for each $t \in \{0, \pm\frac{1}{K}, \pm\frac{2}{K}, \ldots, \pm\frac{K}{K}\}$, $p(\bigcap_{n \geq N}\{R^1_{\{1\},n} \leq t\} \mid \{R^1 = t\} \cap \bar{B}_t) > 1 - \epsilon$. Then $p(B_t \mid \bar{B}_t) > 1 - 2\epsilon$, where

$$B_t = \bar{B}_t \cap \{R^1 = t\} \cap \bigcap_{n \geq N} \{R^1_{\{1\},n} \leq t\}.$$

Let $B = \cup_t B_t$. Then

$$p(B) > 1 - 2\epsilon. \tag{7}$$

Suppose that the games $\Gamma_{B_t,N}$ admit ϵ-equilibria for $t \in \{0, \pm\frac{1}{K}, \pm\frac{2}{K}, \ldots, \pm\frac{K}{K}\}$. By Lemma 5.5 the game Γ_B admits an ϵ-equilibrium. By (7) and Lemma 5.3 it follows that Γ admits a 5ϵ-equilibrium. Therefore it is sufficient to prove the existence of an ϵ-equilibrium in the games $\Gamma_{B_t,N}$, so that one can assume that **A.2** (and analogously **A.3**) holds

Using similar arguments we can assume that $T^2 = \limsup\{R^2_{\{1\},n} \mid R^1_{\{1\},n} = R^1\}$ and $T^1 = \limsup\{R^1_{\{2\},n} \mid R^2_{\{2\},n} = R^2\}$ are constant. One can furthermore assume that $R^2_{\{1\},n} \leq T^2$ whenever $R^1_{\{1\},n} = R^1$, and that $R^1_{\{2\},n} \leq T^1$ whenever $R^2_{\{2\},n} = R^2$.

We now show that if at least one of **A.4-A.6** is not satisfied, the game admits an ϵ-equilibrium, for every $\epsilon > 0$.

If $R^1 \leq 0$ and $R^2 \leq 0$ then the strategy pair (x, y) that is defined by $x_n = y_n = 0$ (always continue) is an equilibrium. We can thus assume that **A.4** is satisfied. Assume w.l.o.g. that $R^1 > 0$.

If $T^2 \geq R^2$ then by **A.4** the following strategy (x, y) is an ϵ-equilibrium: $x_n = \epsilon \cdot 1_{\{R^1_{\{1\},n} = R^1\}}$ and $y_n = 0$. If $T^1 \geq R^1$ and $T^2 < R^2$ then the

[1] By convention $p(\phi \mid \phi) = 1$.

following strategy (x,y) is an ϵ-equilibrium: $y_n = \epsilon \cdot 1_{\{R^2_{\{2\},n}=R^2\}}$, $x_n = \left\{ \begin{array}{ll} 0 & n \leq N \\ \epsilon \cdot 1_{\{R^1_{\{1\},n}=R^1\}} & n > N \end{array} \right\}$, where N is large enough so that $\mathbf{P}_{0,y}(\theta < N) > 1-\epsilon$, and 0 is the strategy that never stops.

The remaining case is $T^2 < R^2$ and $T^1 < R^1$, so that **A.5** and **A.6** are satisfied. $\square$

6 Stopping Games on Finite Trees

An important building block in our analysis are stopping games that are played on a finite tree. In the present section we define these games and study some of their properties.

6.1 The Model and the Main Result

Definition 6.1. A *stopping game on a finite tree* (or simply a *game on a tree*) is a tuple $T = (S, S_1, r, (C_s, p_s, R_s)_{s \in S \setminus S_1})$, where

- $(S, S_1, r, (C_s)_{s\in S\setminus S_1})$ is a tree; S is a non empty finite set of *nodes*, $S_1 \subseteq S$ is a finite set of *leaves*, $r \in S$ is the root, and for each $s \in S \setminus S_1$, $C_s \subseteq S \setminus \{r\}$ is the set of *children* of s. We denote by $S_0 = S \setminus S_1$ the set of nodes which are not leaves. For every $s \in S$, depth(s) is the depth of s - the length of the path that connects the root to s.

For every $s \in S_0$,

- p_s is a probability distribution over C_s.
- $R_s \in \mathbf{R}^6$ is the *payoff* at s. The coordinates of R_s are denoted $(R^i_{Q,s})_{i=1,2,\phi\neq Q\subseteq\{1,2\}}$.

Throughout this section we consider games on trees whose payoffs $(R_s)_{s\in S_0}$ satisfy the following conditions for every $i = 1,2$, every $\emptyset \subset Q \subseteq \{1,2\}$, and every $s \in S_0$, (**B.1**) $R^i_{Q,s} \in \{0, \pm\frac{1}{K}, \ldots, \pm\frac{K}{K}\}$ for some $K \in \mathbf{N}$, (**B.2**) $R^i_{\{i\},s} \leq R^i$, where $R^1, R^2 \in \mathbf{R}$, and at least one of them is positive, (**B.3**) $R^2_{\{1\},s} < R^2$ whenever $R^1_{\{1\},s} = R^1$, and (**B.4**) $R^1_{\{2\},s} < R^1$ whenever $R^2_{\{2\},s} = R^2$. Observe the similarity between these conditions and conditions **A.1**-**A.6**.

A stopping game on a finite tree starts at the root and is played in stages. Given the current node $s \in S_0$, and the sequence of nodes that were already visited, both players decide, simultaneously and independently, whether to stop or to continue. Let Q be the set of players that decide to stop. If $Q \neq \phi$, the play terminates, and the terminal payoff to player i is $R^i_{Q,s}$. If $Q = \phi$, a new node s' in C_s is chosen according to p_s. The process now repeats itself, with s' being the current node. If $s' \in S_1$ the new current node is the root r. Thus, players cannot stop at leaves.

The game on the tree is essentially played in rounds. The round starts at the root, and ends once it reaches a leaf. The collection $(p_s)_{s \in S_0}$ of probability distributions induces a probability distribution over the set of leaves S_1, or, equivalently, over the set of branches that connect the root to the leaves. For each set $D \subseteq S_0$, we denote by p_D the probability that the chosen branch passes through D. For each $s \in S$ we denote by F_s the event that the chosen branch passes through s.

Definition 6.2. Let $T = (S, S_1, r, (C_s, p_s, R_s)_{s \in S_0})$ and $T' = (S', S'_1, r', (C'_s, p'_s, R'_s)_{s \in S'_0})$ be two games on trees. T' is a *subgame* of T if (i) $S' \subseteq S$, (ii) $r' = r$, and (iii) for every $s \in S'_0$, $C'_s = C_s$, $p'_s = p_s$ and $R'_s = R_s$.

In words, T' is a subgame of T if one removes all the descendents (in the strict sense) of several nodes from the tree $(S, S_1, r, (C_s)_{s \in S_0})$, and keep all other parameters fixed. Observe that this notion is different than the standard definition of a subgame in game theory.

Let $T = (S, S_1, r, (C_s, p_s, R_s)_{s \in S_0})$ be a game on a tree. For each subset $D \subseteq S_0$ we denote by T_D the subgame of T generated by trimming T from D downward. Thus, all strict descendents of nodes in D are removed.

For every subgame T' of T, and every subgame T'' of T', let $p_{T'',T'} = p_{S''_1 \setminus S'_1}$ be the probability that the chosen branch in T passes through a leaf of T'' strictly before it passes through a leaf of T'.[2]

Consider the first round of the game. Let t be the stopping stage. If no termination occurred in the first round $t = \infty$. If $t < \infty$ let s be the node (of depth t) in which termination occurred, and let Q be the set of players that stop at stage t. The r.v. $r^i = R^i_{Q,s} 1_{\{t<\infty\}}$ is the payoff to player i in the first round.

A *stationary strategy* of player 1 (*resp.* player 2) is a function $x : S_0 \to [0,1]$ (*resp.* $y : S_0 \to [0,1]$): $x(s)$ is the probability that player 1 stops at s. Denote by $\mathbf{P}_{x,y}$ the distribution over plays induced by (x,y), and by $\mathbf{E}_{x,y}$ the corresponding expectation operator.

[2]Here, S'_1 (*resp.* S''_1) is the set of leaves of T' (*resp.* T'').

For every pair of stationary strategies (x,y) we denote by $\pi(x,y) = \mathbf{P}_{x,y}(t < \infty)$ the probability that under (x,y) the game terminates in the first round of the game; that is, the probability that the root is visited only once along the play. We denote by $\rho^i(x,y) = \mathbf{E}_{x,y}[r^i]$, $i = 1,2$, the expected payoff of player i in a single round. Finally, we set $\gamma^i(x,y) = \rho^i(x,y)/\pi(x,y)$.[3] This is the expected payoff under (x,y). In particular,

$$\pi(x,y) \cdot \gamma^i(x,y) = \rho^i(x,y). \tag{8}$$

When we want to emphasize the dependency of these variables on the game T, we will write $\pi(x,y;T)$, $\rho^i(x,y;T)$ and $\gamma^i(x,y;T)$.

Observe that for every pair of stationary strategies (x,y)

$$\pi(x,0) + \pi(0,y) \geq \pi(x,y), \tag{9}$$

where 0 is the strategy that never stop; that is, $0(s) = 0$ for every s.

Definition 6.3. A pair of stationary strategies (x,y) is an ϵ-*equilibrium* of the game T if, for each pair of strategies (x',y'), $\gamma^1(x',y) \leq \gamma^1(x,y) + \epsilon$ and $\gamma^2(x,y') \leq \gamma^2(x,y) + \epsilon$.

Comment: A stopping game on a finite tree T is equivalent to a recursive absorbing game, where each round of the game T corresponds to a single stage of the recursive absorbing game. A recursive absorbing game is a stochastic game with a single non absorbing state, in which the payoff in non absorbing states is 0. Flesch et al (1996) proved that every recursive absorbing game admits an ϵ-equilibrium in stationary strategies. This result also follows from the analysis of Vrieze and Thuijsman (1989). However, there is no bound on the per-round probability of termination under this ϵ-equilibrium, and we need to bound this quantity.

The main result of this section is the following.

Proposition 6.4. *For every stopping game on a finite tree T, every $\epsilon > 0$ sufficiently small, and every a_1, a_2, b_1, b_2 that satisfy $R^i - \epsilon \leq a_i < b_i$ for $i \in \{1,2\}$, there exist a set $D \subseteq S_0$ of nodes and a strategy pair (x,y) in T such that:*

[3] By convention, $\frac{0}{0} = 0$.

1. *In every subgame T' of T_D there are no ϵ-equilibria in T' with corresponding payoffs in $[a_1, b_1] \times [a_2, b_2]$.*
2. *Either $D = \phi$ (so that $T_D = T$), or the following three conditions hold:*
   (a) *$a_1 - \epsilon \leq \gamma^1(x, y)$ and $a_2 - \epsilon \leq \gamma^2(x, y)$.*
   (b) *For every pair (x', y') of strategies, $\gamma^1(x', y) \leq b_1 + 7\epsilon$ and $\gamma^2(x, y') \leq b_2 + 7\epsilon$.*
   (c) *$\pi(x, y) \geq \epsilon^2 \times p_D$.*

Observe that 2(a) and 2(b) imply that if $b_i - a_i \leq \epsilon$ then (x, y) is a 9ϵ-equilibrium in T with corresponding payoffs in $[a_1 - \epsilon, b_1 + 7\epsilon] \times [a_2 - \epsilon, b_2 + 7\epsilon]$.

6.2 Union of Strategies

Given n stationary strategies $x_1, x_2, \ldots, x_n$, we define their *union* x by $x(s) = 1 - \Pi_{1 \leq k \leq n}(1 - x_k(s))$. The probability that the union strategy continues at each node is the probability that *all* of its components continue. We denote $x = x_1 \dot{+} x_2 \dot{+} \ldots \dot{+} x_n$. Given n pairs of stationary strategies $\alpha_k = (x_k, y_k)$, $1 \leq k \leq n$, we denote by $\alpha_1 \dot{+} \ldots \dot{+} \alpha_n$ the stationary strategy pair (x, y) that is defined by $x = x_1 \dot{+} \ldots \dot{+} x_n$, $y = y_1 \dot{+} \ldots \dot{+} y_n$.

Consider now n copies of the game that are played simultaneously, such that the choice of a new node is the same across the copies; that is, all copies that have not terminated at stage t are at the same node. Nevertheless, the lotteries made by the players concerning the decision whether to stop or not are independent. Let $\alpha_k = (x_k, y_k)$, $1 \leq k \leq n$, be the stationary strategy pair used in copy k and let $\alpha = \alpha_1 \dot{+} \ldots \dot{+} \alpha_n$.

We consider the first round of the game. Let t_k be the stopping stage in copy k, let s_k be the node in which termination occurred, let Q_k be the set of players that stop at stage t_k and let $r_k^i = R^i_{Q_k, s_k} 1_{\{t_k < \infty\}}$. Then $\pi_k = \pi(x_k, y_k) = \mathbf{P}(t_k < \infty)$ and $\rho_k^i = \rho^i(x_k, y_k) = \mathbf{E}[r_k^i]$.

Let t,r,ρ,π be the analog quantities w.r.t. α: $t = \min\{t_k, 1 \leq k \leq n\}$, $r^i = R^i_{Q,s} 1_{\{t < \infty\}}$, where $s = s_k$ for which t_k is minimal, and $Q = \bigcup_{k | s_k = s} Q_k$. Then $\rho^i = \rho^i(x, y) = \mathbf{E}[r^i]$ and $\pi = \pi(x, y) = \mathbf{P}(t < \infty)$.

Let $\gamma_k = \gamma(x_k, y_k)$ be the expected payoff under $\alpha_k = (x_k, y_k)$, and $\gamma = \gamma(x, y)$ be the corresponding quantity under α.

The following lemma follows from the independence of the plays given the branch.

Lemma 6.5. *Let $s \in S_0$ be a node of depth j. Then, for every $1 \le k, l \le n$, $l \neq k$, the event $\{t_k \le j\}$ and the random variable $t_k 1_{\{t_k \le j\}}$ are independent of t_l given F_s.*

Lemma 6.6. *Let $N = \sum_{k=1}^n 1_{\{t_k < \infty\}}$ be the number of copies that terminate in the first round. Then*

1. $\sum_{k=1}^n \pi_k - \mathbf{E}[N 1_{\{N \ge 2\}}] \le \pi \le \sum_{k=1}^n \pi_k$, *and*
2. $\sum_{k=1}^n \rho_k^i - \mathbf{E}[(N+1) 1_{\{N \ge 2\}}] \le \rho^i \le \sum_{k=1}^n \rho_k^i + \mathbf{E}[(N+1) 1_{\{N \ge 2\}}]$ *for each player $i \in \{1, 2\}$.*

Proof. Observe that

$$N - N 1_{\{N \ge 2\}} = 1_{\{N=1\}} \le 1_{\{N \ge 1\}} \le N = \sum_{k=1}^n 1_{\{t_k < \infty\}}.$$

The first result follows by taking expectations.

For the second result, note that

$$\sum_{k=1}^n r_k^i - (N+1) 1_{\{N \ge 2\}} \le r^i \le \sum_{k=1}^n r_k^i + (N+1) 1_{\{N \ge 2\}}. \tag{10}$$

Indeed, on $\{N \le 1\}$ (10) holds with equality, and on $\{N \ge 2\}$ the left hand side is at most -1, whereas the right hand side is at least $+1$. The result follows by taking expectations. □

6.3 Heavy and Light Nodes

Definition 6.7. Let $\sigma = (x, y)$ be a pair of stationary strategies and let $\delta > 0$. A node $s \in S_0$ is *δ-heavy* with respect to σ if $\mathbf{P}_\sigma(t < \infty \mid F_s) \ge \delta$; that is the probability of termination in the first round given that the chosen branch passes through s is at least δ. The node s is *δ-light* w.r.t. σ if $\mathbf{P}_\sigma(t < \infty \mid F_s) < \delta$.

For a fixed δ, we denote by $H_\delta(\sigma)$ the set of δ-heavy nodes w.r.t. σ. Two simple implications of this definition are the following.

Fact 1 $H_\delta(\alpha_1) \subseteq H_\delta(\alpha_1 \dot{+} \alpha_2)$.

Fact 2 $H_{\delta_1}(\sigma) \subseteq H_{\delta_2}(\sigma)$ whenever $\delta_1 \geq \delta_2$.

Lemma 6.8. *Let $\epsilon > 0$ be sufficiently small, and let (x, y) be a stationary ϵ-equilibrium such that $R^i - \epsilon \leq \gamma^i(x, y)$, $i = 1, 2$. Then $H_\epsilon(x, y) \neq \phi$. In particular, by Fact 2, $H_{\epsilon^2}(x, y) \neq \phi$.*

Comment: The proof hinges on the assumption that $R^2_{\{1\},s} < R^2$ whenever $R^1_{\{1\},s} = R^1$. As a counter example when this condition does not hold, take a game in which (a) $R^i_{Q,s} = 1$ for every i, Q and s, and (b) $R^1 = R^2 = 1$. Then any stationary strategy pair which stops with positive probability is a 0-equilibrium.

Proof. Assume w.l.o.g that $\pi(x, 0) \geq \pi(0, y)$. Let r be the probability that, under $(x, 0)$ termination occurs at a node s in which $R^1_{\{1\},s} < R^1$. Since payoffs are discrete,

$$\rho^1(x, 0) \leq \pi(x, 0) \cdot ((1 - r)R^1 + r(R^1 - \frac{1}{K})) = \pi(x, 0) \cdot (R^1 - \frac{r}{K}). \quad (11)$$

Assume to the contrary that $H_\epsilon(x, y) = \phi$. Then, in particular, $H_\epsilon(0, y) = \phi$. It follows that the sequence $\{(0, y), (x, 0)\}$ is ϵ-orthogonal. By Lemma 6.14, (8), (11) and sinc (x, y) is an ϵ-equilibrium,

$$\begin{aligned}
&(\pi(0, y) + \pi(x, 0)) \cdot \gamma^1(x, y) \\
&\quad \leq \rho^1(0, y) + \rho^1(x, 0) + 6\epsilon(\pi(0, y) + \pi(x, 0)) \\
&\quad \leq \pi(0, y) \cdot \gamma^1(0, y) + \pi(x, 0) \cdot (R^1 - \frac{r}{K}) + 6\epsilon(\pi(0, y) + \pi(x, 0)) \\
&\quad \leq \pi(0, y) \cdot (\gamma^1(x, y) + \epsilon) + \pi(x, 0) \cdot (R^1 - \frac{r}{K}) + 6\epsilon(\pi(0, y) + \pi(x, 0)).
\end{aligned}$$

It follows that

$$\pi(x, 0) \cdot \gamma^1(x, y) \leq \pi(x, 0)(R^1 - \frac{r}{K}) + 7\epsilon \cdot \pi(0, y) + 6\epsilon \cdot \pi(x, 0).$$

Since $R^1 - \epsilon \leq \gamma^1(x, y)$,

$$\pi(x, 0) \cdot (R^1 - \epsilon) \leq \pi(x, 0)(R^1 - \frac{r}{K}) + 7\epsilon \cdot \pi(0, y) + 6\epsilon \cdot \pi(x, 0),$$

which implies

$$\pi(x,0)r \leq 7\epsilon K \cdot (\pi(0,y)+\pi(x,0)). \tag{12}$$

Since payoffs are discrete and bounded by 1, since $R^2_{\{1\},s} < R^2$ whenever $R^1_{\{1\},s} = R^1$, and by (12),

$$\begin{aligned}
\rho^2(x,0) &\leq \pi(x,0)\left((1-r)(R^2-\frac{1}{K})+r\cdot 1\right)\\
&= \pi(x,0)\left((R^2-\frac{1}{K})+r(1-R^2+\frac{1}{K})\right) \qquad (13)\\
&\leq \pi(x,0)(R^2-\frac{1}{K}+3r)\\
&\leq \pi(x,0)\cdot(R^2-\frac{1}{K})+21\epsilon K(\pi(0,y)+\pi(x,0)).
\end{aligned}$$

By Lemma 6.14, since $\rho^2(0,y) \leq \pi(0,y)\cdot R^2$, (13), and since $\pi(x,0) \geq \pi(0,y)$

$$\begin{aligned}
&(\pi(0,y)+\pi(x,0))\gamma^2(x,y) \leq\\
&\quad \leq \rho^2(0,y)+\rho^2(x,0)+6\epsilon(\pi(0,y)+\pi(x,0))\\
&\quad \leq \pi(0,y)\cdot R^2+\pi(x,0)\cdot(R^2-\frac{1}{K})+(6\epsilon+21\epsilon K)\cdot(\pi(0,y)+\pi(x,0))\\
&\quad = (\pi(0,y)+\pi(x,0))\cdot(R^2+6\epsilon+21\epsilon K)-\pi(x,0)\cdot\frac{1}{K}\\
&\quad \leq (\pi(0,y)+\pi(x,0))\cdot(R^2+6\epsilon+21\epsilon K-\frac{1}{2K}).
\end{aligned} \tag{14}$$

Since ϵ is sufficiently small, and $R^i - \epsilon \leq \gamma^i(x,y)$, it follows from **A.4** that $\gamma^i(x,y) > 0$ for $i=1$ or $i=2$. In particular, it follows that $\pi(x,y) > 0$, so that by (9) $\pi(x,0)+\pi(0,y) \geq \pi(x,y) > 0$. It follows that, for sufficiently small ϵ (e.g. $\epsilon < \frac{1}{2K(7+21K)}$),

$$\gamma^2(x,y) \leq R^2+6\epsilon+21\epsilon K-\frac{1}{2K} \leq R^2-\epsilon,$$

a contradiction. □

6.4 Orthogonal Strategies

Definition 6.9. Let $\delta > 0$. A sequence $(\alpha_1, \alpha_2, \ldots, \alpha_n)$ of stationary strategy pairs is *δ-orthogonal* if $\alpha_{k+1}(s) = (0,0)$ for every $1 \leq k \leq n-1$ and every node $s \in H_\delta(\alpha_1 \dot{+} \ldots \dot{+} \alpha_k)$; that is α_{k+1} continues on δ-heavy nodes of $\alpha_1 \dot{+} \ldots \dot{+} \alpha_k$.

Lemma 6.10. *Let $\delta > 0$, let $(\alpha_1, \ldots, \alpha_n)$ be a δ-orthogonal sequence of stationary strategy pairs, let $k \in \{1, \ldots, n\}$, and let $s \in S$ be a node of depth j. Then*

$$\mathbf{P}(\{j \le t_k < \infty\} \cap (\cup_{l<k}\{t_l < \infty\}) \mid F_s) \le \delta \cdot \mathbf{P}(j \le t_k < \infty \mid F_s). \quad (15)$$

Proof. Fix $k \in \{1, \ldots, n\}$. We prove the lemma by induction on the nodes of T, starting from the leaves and climbing up to the root.

Let $s \in S_1$ be a leaf of T. Since s is a leaf, $\mathbf{P}(j \le t_k < \infty) = 0$ and (15) is satisfied trivially.

Assume now that $s \in S_0$. Then:

$$\begin{aligned}\mathbf{P}(\{j \le t_k < \infty\} \cap (\cup_{l<k}\{t_l < \infty\}) \mid F_s) &= \mathbf{P}(\{t_k = j\} \cap (\cup_{l<k}\{t_l < \infty\}) \mid F_s) + \\ \sum_{s' \in C_s} p_s[s'] \cdot \mathbf{P}(\{j+1 \le t_k < \infty\} &\cap (\cup_{l<k}\{t_l < \infty\}) \mid F_{s'}). \quad (16)\end{aligned}$$

By the induction hypothesis, for every child $s' \in C_s$,

$$\mathbf{P}(\{j+1 \le t_k < \infty\} \cap (\cup_{l<k}\{t_l < \infty\}) \mid F_{s'}) \le \delta \cdot \mathbf{P}(j+1 \le t_k < \infty \mid F_{s'}). \quad (17)$$

By Lemma 6.5 $\{t_k = j\}$ and $\cup_{l<k}\{t_l < \infty\}$ are independent given F_s. Therefore

$$\mathbf{P}(\{t_k = j\} \cap (\cup_{l<k}\{t_l < \infty\}) \mid F_s) = \mathbf{P}(t_k = j \mid F_s) \cdot \mathbf{P}(\cup_{l<k}\{t_l < \infty\} \mid F_s).$$

If s is δ-light w.r.t $\alpha_1 \dot{+} \ldots \dot{+} \alpha_{k-1}$ then $\mathbf{P}(\cup_{l<k}\{t_l < \infty\} \mid F_s) < \delta$ while if s is δ-heavy then $\mathbf{P}(t_k = j \mid F_s) = 0$ according to the definition of orthogonality. In particular,

$$\mathbf{P}(\{t_k = j\} \cap (\cup_{l<k}\{t_l < \infty\}) \mid F_s) \le \delta \cdot \mathbf{P}(t_k = j \mid F_s). \quad (18)$$

Eqs. (16),(17) and (18) yield

$$\begin{aligned}&\mathbf{P}(\{j \le t_k < \infty\} \cap (\cup_{l<k}\{t_l < \infty\})) \\ &\quad \le \delta \cdot \mathbf{P}(t_k = j \mid F_s) + \delta \cdot \sum_{s' \in C_s} p_s[s'] \cdot \mathbf{P}(j+1 \le t_k < \infty \mid F_{s'}) \\ &\quad = \delta \cdot \mathbf{P}(j \le t_k < \infty \mid F_s),\end{aligned}$$

as desired. ∎

Applying Lemma 6.10 to the root we get:

Corollary 6.11. *Let $\delta > 0$, and let $(\alpha_1, \ldots, \alpha_n)$ be a δ-orthogonal sequence of stationary strategy pairs. For every $k \in \{1, \ldots, n\}$,*

$$\mathbf{P}(\{t_k < \infty\} \cap (\cup_{l<k}\{t_l < \infty\})) \le \delta \cdot \mathbf{P}(\{t_k < \infty\}) = \delta\pi_k.$$

Lemma 6.12. *Let $\delta > 0$, let $(\alpha_1, \ldots, \alpha_n)$ be a δ-orthogonal sequence of stationary strategy pairs, and let $N = \sum_{k=1}^n 1_{\{t_k<\infty\}}$. Then $\mathbf{E}[(N+1)1_{\{N\ge 2\}}] \le 3\delta(\pi_1 + \pi_2 + \ldots + \pi_n)$.*

Proof. Observe that $N + 1 \le 3(N - 1)$ on $\{N \ge 2\}$, and $(N - 1)1_{\{N\ge 2\}} = \sum_{k=1}^n 1_{\{t_k<\infty\}\cap(\cup_{l<k}\{t_l<\infty\})}$. Therefore

$$\mathbf{E}[(N+1)1_{\{N\ge 2\}}] \le 3\mathbf{E}[(N-1)1_{\{N\ge 2\}}] = 3\sum_{k=1}^n \mathbf{P}(\{t_k < \infty\}\cap(\cup_{l<k}\{t_l < \infty\})).$$

The result follows by Corollary 6.11. □

From Lemma 6.6 and Lemma 6.12 we get the following.

Corollary 6.13. *Let $\delta > 0$, and let $(\alpha_1, \ldots \alpha_n)$ be a δ-orthogonal sequence of strategy pairs. Denote $\alpha = \alpha_1 \dot{+} \ldots \dot{+} \alpha_n$. Then for $i = 1, 2$*

1. $(1 - 3\delta)\sum_{k=1}^n \pi_k \le \pi \le \sum_{k=1}^n \pi_k$.
2. $\sum_{k=1}^n \rho_k^i - 3\delta\sum_{k=1}^n \pi_k \le \rho^i \le \sum_{k=1}^n \rho_k^i + 3\delta\sum_{k=1}^n \pi_k$.

Lemma 6.14. *Let $\delta > 0$, and let $(\alpha_1, \ldots \alpha_n)$ be a δ-orthogonal sequence of stationary strategy pairs. Denote $\alpha = \alpha_1 \dot{+} \ldots \dot{+} \alpha_n$. Then for $i = 1, 2$*

$$\sum_{k=1}^n \rho_k^i - 6\delta\sum_{k=1}^n \pi_k \le \gamma^i \cdot \sum_{k=1}^n \pi_k \le \sum_{k=1}^n \rho_k^i + 6\delta\sum_{k=1}^n \pi_k.$$

Proof. By Corollary 6.13 and (8)

$$\sum_{k=1}^n \rho_k^i - 3\delta\sum_{k=1}^n \pi_k \le \rho^i = \gamma^i \cdot \pi \le \begin{cases} \gamma^i \cdot \sum_{k=1}^n \pi_k, & \text{if } \gamma^i > 0 \\ \gamma^i(1 - 3\delta)\sum_{k=1}^n \pi_k, & \text{if } -1 \le \gamma^i \le 0 \end{cases}.$$

In both case, the right-hand side is bounded by $\gamma^i \cdot \sum_{k=1}^n \pi_k + 3\delta\sum_{k=1}^n \pi_k$, so that

$$\sum_{k=1}^n \rho_k^i - 6\delta\sum_{k=1}^n \pi_k \le \gamma^i \cdot \sum_{k=1}^n \pi_k.$$

The proof of the right-hand side inequality is similar. □

From Lemma 6.14 and (8) we get:

Corollary 6.15. *Let $\delta > 0$, and let $(\alpha_1, \ldots \alpha_n)$ be a δ-orthogonal sequence of stationary strategy pairs. Denote $\alpha = \alpha_1 \dot{+} \ldots \dot{+} \alpha_n$. Let $-1 \leq u, v \leq 1$.*

1. *If $u \leq \gamma_k^i$ for each $k \in \{1, \ldots, n\}$, then $u - 6\delta \leq \gamma^i$.*
2. *If $\gamma_k^i \leq v$ for each $k \in \{1, \ldots, n\}$, then $\gamma^i \leq v + 6\delta$.*

6.5 Strong Orthogonality

In the present section we define a stronger notion of orthogonality and study its properties.

Definition 6.16. Let $\delta > 0$. A sequence $(\alpha_1, \alpha_2, \ldots, \alpha_n)$ of stationary strategy pairs is *δ-strongly orthogonal* if, for every $k \in \{1, \ldots, n-1\}$ and every node $s \in H_\delta(\alpha_1 \dot{+} \ldots \dot{+} \alpha_k)$, $\alpha_{k+1}(s') = (0, 0)$ for $s' = s$ and for every descendent s' of s; that is α_{k+1} continues from s onwards.

The following lemma suggests a way to construct ϵ-orthogonal sequences of strategy pairs from a single ϵ^2-strongly orthogonal sequence.

Lemma 6.17. *Let $\epsilon > 0$ and let $y_1, y_2, \ldots, y_n$ be stationary strategies of player 2 such that the sequence $((0, y_1), \ldots, (0, y_n))$ is ϵ^2-strongly orthogonal. Let $\bar{x}$ be any pure stationary strategy of player 1 that does not stop twice on the same branch; that is, if $\bar{x}(s) = 1$ then $\bar{x}(s') = 0$ for every descendant s' of s. Define strategies $(\bar{x}_k)_{k=1}^n$ of player 1 in the following way: for each $s \in S$ such that $\bar{x}(s) = 1$ let $\bar{x}_k(s) = 1$, where $k \leq n$ is the greatest index for which $s \notin H_\epsilon((0, y_1) \dot{+} \ldots \dot{+} (0, y_{k-1}))$. Define $\bar{x}_k(s) = 0$ otherwise.*
Let $\bar{\alpha}_k = (\bar{x}_k, y_k)$. Then the sequence $(\bar{\alpha}_1, \ldots, \bar{\alpha}_n)$ is ϵ-orthogonal.

Proof. By the definition of $(\bar{x}_k)_{1 \leq k \leq n}$ and Fact 1, we get, for every $l \in \{1, \ldots, n-1\}$

$$\begin{array}{l} \text{If } \bar{x}_l(s) = 1 \text{ then } s \in H_\epsilon((0, y_1) \dot{+} \ldots \dot{+} (0, y_l)). \\ \text{If } \bar{x}_{l+1}(s) = 1 \text{ then } s \notin H_\epsilon((0, y_1) \dot{+} \ldots \dot{+} (0, y_l)). \end{array} \qquad (19)$$

Let $l \in \{1, \ldots, n-1\}$, and let $s \in S$ be ϵ-heavy with respect to $\bar{\sigma}_l = \bar{\alpha}_1 \dot{+} \ldots \dot{+} \bar{\alpha}_l$. We prove that $\bar{x}_{l+1}(s) = y_{l+1}(s) = 0$.

We first prove that $\bar{x}_{l+1}(s) = 0$. Since s is ϵ-heavy w.r.t. $\bar{\sigma}_l = \bar{\alpha}_1 \dot{+} \ldots \dot{+} \bar{\alpha}_l$, $\mathbf{P}_{\bar{\sigma}_l}(t < \infty \mid F_s) \geq \epsilon$. Assume to the contrary that $\bar{x}_{l+1}(s) = 1$. By (19)

s is ϵ-light w.r.t. $(0,y_1)\dot{+}\ldots\dot{+}(0,y_l)$, and therefore $\mathbf{P}_{(0,y_1)\dot{+}\ldots\dot{+}(0,y_l)}(t<\infty \mid F_s)<\epsilon$. It follows that $\mathbf{P}_{(\bar{x}_1,0)\dot{+}\ldots\dot{+}(\bar{x}_l,0)}(t<\infty \mid F_s)>0$, a contradiction to the assumption that $\bar{x}$ does not stop twice on the same branch.

We proceed to prove that $y_{l+1}(s)=0$. Assume first that there exists an ancestor s' of s such that $\bar{x}_1(s')+\ldots+\bar{x}_l(s')=1$. By (19) and Fact 1 $s'\in H_\epsilon((0,y_1)\dot{+}\ldots\dot{+}(0,y_l))$. Since $((0,y_1),\ldots,(0,y_n))$ is ϵ-strongly orthogonal, $y_{l+1}(s)=0$.

We assume now that $\bar{x}_1(s')+\ldots+\bar{x}_l(s')=0$ for every ancestor s' of s. Let $\tilde{D}$ be the (possibly empty) set of s's descendants d that are ϵ-heavy w.r.t. $(0,y_1)\dot{+}\ldots\dot{+}(0,y_l)$, and let D be the set that is obtained by removing from $\tilde{D}$ all nodes that have strict ancestor in $\tilde{D}$. By the definition of D, $\mathbf{P}_{(0,y_1)\dot{+}\ldots\dot{+}(0,y_l)}(t<\infty \mid F_d)\geq\epsilon$ for every $d\in D$. Let $Y=\cup_{d\in D}F_d$. Since this is a mutually disjoint union, it follows that if $Y\neq\phi$ then

$$\mathbf{P}_{(0,y_1)\dot{+}\ldots\dot{+}(0,y_l)}(t<\infty \mid Y)\geq\epsilon\geq\epsilon\cdot\mathbf{P}_{\bar{\sigma}_l}(t<\infty \mid Y).$$

By (19) and the definition of $(\bar{x}_k)_{1\leq k\leq n}$ it follows that

$$\mathbf{P}_{(0,y_1)\dot{+}\ldots\dot{+}(0,y_l)}(t<\infty \mid Y^c\cap F_s)=\mathbf{P}_{\bar{\sigma}_l}(t<\infty \mid Y^c\cap F_s)\geq\epsilon\cdot\mathbf{P}_{\bar{\sigma}_l}(t<\infty \mid Y^c\cap F_s).$$

Combining the last two inequalities, and observing that $Y\subseteq F_s$, we get

$$\mathbf{P}_{(0,y_1)\dot{+}\ldots\dot{+}(0,y_l)}(t<\infty \mid F_s)\geq\epsilon\cdot\mathbf{P}_{\bar{\sigma}_l}(t<\infty \mid F_s)\geq\epsilon^2.$$

Thus s is ϵ^2-heavy with respect to $(0,y_1)\dot{+}\ldots\dot{+}(0,y_l)$ and, as the sequence $((0,y_1),\ldots,(0,y_n))$ is ϵ^2-orthogonal, $y_{l+1}(s)=0$. □

Lemma 6.18. *Let $\epsilon>0$ be sufficiently small and let a_1,b_1,a_2,b_2 satisfy $a_i<b_i$ for $i\in\{1,2\}$. Let $(\alpha_1,\ldots,\alpha_n)$ be an ϵ^2-strongly orthogonal sequence of stationary strategy pairs such that α_k is an ϵ-equilibrium for each $k=1,\ldots,n$. Assume that for each k $\gamma_k\in[a_1,b_1]\times[a_2,b_2]$, where γ_k is the payoff that corresponds to α_k. Let $\alpha=\alpha_1\dot{+}\ldots\dot{+}\alpha_n=(x,y)$. Then*

- *a) $a_i-\epsilon\leq\gamma^i(x,y)$.*
- *b) For each pair (x',y') of stationary strategies, $\gamma^1(x',y)\leq b_1+7\epsilon$ and $\gamma^2(x,y')\leq b_2+7\epsilon$.*

Proof. Denote $\alpha_k=(x_k,y_k)$. We prove the result only for player 1.

We first prove (a). Since $a_1\leq\gamma^1_k(x_k,y_k)$ for each $1\leq k\leq n$, it follows from Corollary 6.15, and since ϵ is sufficiently small, that $a_1-\epsilon\leq a_1-6\epsilon^2\leq\gamma^1(x,y)$.

We now prove (b). Let $\bar{x}$ be a stationary strategy that maximizes player 1's payoff against y: $\gamma^1(\bar{x}, y) = \max_{x'} \gamma^1(x', y)$. Fixing y, the game reduces to a Markov decision process, hence such an $\bar{x}$ exists. Moreover, there exists such a strategy $\bar{x}$ that is pure (that is, $\bar{x}(s) \in \{0, 1\}$ for every s) and stops at most once in every branch. Observe that since the sequence $(\alpha_1, \ldots, \alpha_n)$ is ϵ^2-strongly orthogonal, so is the sequence $((0, y_1), \ldots, (0, y_n))$. Let $\bar{x}_1, \ldots, \bar{x}_k$ be the strategies defined in Lemma 6.17 w.r.t. $\bar{x}$ and $y_1, \ldots, y_n$. Then $\bar{x} = \bar{x}_1 \dot{+} \ldots \dot{+} \bar{x}_n$, and $(\bar{\alpha}_1, \ldots, \bar{\alpha}_n)$ is ϵ-orthogonal, where $\bar{\alpha}_k = (\bar{x}_k, y_k)$.

For each k, (x_k, y_k) is an ϵ-equilibrium, and therefore $\gamma^1(\bar{x}_k, y_k) \leq b_1 + \epsilon$. By Corollary 6.15 and the definition of $\bar{x}$, for every x' one has $\gamma^1(x', y) \leq \gamma^1(\bar{x}, y) \leq b_1 + \epsilon + 6\epsilon = b_1 + 7\epsilon$. □

6.6 Proof of Proposition 6.4

We now prove Proposition 6.4. Consider the following recursive procedure:

1. Initialization: Start with the game $\tilde{T} = T$, the strategy pair $\sigma_0 = (0, 0)$ (always continue) and $k = 0$.
2. If there exists a stationary ϵ-equilibrium in a subgame T' of $\tilde{T}$ with corresponding payoff in $[a_1, b_1] \times [a_2, b_2]$:
   (a) Set $k = k + 1$ and let $\alpha_k = (x_k, y_k)$ be any such ϵ-equilibrium. Extend x_k and y_k to strategies on T by setting $x_k(s) = y_k(s) = 0$ for every node $s \in S_0 \setminus T'$.
   (b) Set $\sigma_k = \sigma_{k-1} \dot{+} \alpha_k$.
   (c) Let $H_k = H_{\epsilon^2}(\sigma_k)$ be the set of ϵ^2-heavy nodes of σ_k (by Fact 1 $H_{k-1} \subseteq H_k$.) Set $\tilde{T} = T_{H_k}$.
   (d) Start stage 2 all over.
3. If, for all subgames T' of $\tilde{T}$, there are no ϵ-equilibria in T' with corresponding payoff in $[a_1, b_1] \times [a_2, b_2]$, set $n = k$, $x = x_1 \dot{+} \ldots \dot{+} x_n$, $y = y_1 \dot{+} \ldots \dot{+} y_n$, and $D = H_n$.

The idea is to keep adding strongly orthogonal ϵ-equilibria as long as we can. The procedure continues until there is no ϵ-equilibrium in any subgame of $\tilde{T}$ with payoffs in $[a_1, b_1] \times [a_2, b_2]$. The termination of the procedure follows from Lemma 6.8.

The first part of Proposition 6.4 is an immediate consequence of the termination of the procedure. We now prove that $\sigma_n = (x, y)$ satisfies the requirements of the second part. Since $D = H_n$ is the set of ϵ^2-heavy nodes of (x, y), claim 2c in Proposition 6.4 follows. For every $1 \leq k \leq n$, $\gamma^i(x_k, y_k) \geq R^i - \epsilon$, so that (x_k, y_k) is an ϵ-equilibrium in T. Thus $((x_1, y_1), \ldots, (x_n, y_n))$ is an ϵ^2-strongly orthogonal sequence of stationary ϵ-equilibria. The remaining claims of Proposition 6.4 follow from Lemma 6.18.

6.7 Equilibria with Low Payoff

In Proposition 6.4 we consider ϵ-equilibria with corresponding payoffs (u^1, u^2) such that $u^i \geq R^i - \epsilon$. We now deal with the case in which one of the players (w.l.o.g. player 1) gets low payoff.

Lemma 6.19. *Let $\epsilon > 0$, and let (x, y) be a stationary $\frac{\epsilon}{2}$-equilibrium in T such that $\gamma^1(x, y) \leq R^1 - \epsilon$. Then $\pi(0, y) \geq \frac{\epsilon}{4} \cdot r_1$, where $r_1 = p(\cup\{F_s, R^1_{\{1\},s} = R^1\})$ is the probability that, if both players never stop, the game visits a node s with $R^1_{\{1\},s} = R^1$ in the first round.*

Proof. Consider the following strategy z of player 1: $z_s = \begin{cases} 1, & \text{if } R^1_{\{1\},s} = R^1 \\ 0, & \text{otherwise} \end{cases}$.
By the definition of z, and since payoffs are bounded by 1,

$$\begin{aligned}
\rho^1(z, y) &= \mathbf{E}_{z,y}[R^1_{Q,s} 1_{\{t<\infty\}}] \\
&= \mathbf{E}_{z,y}[R^1_{Q,s} 1_{\{t<\infty, Q=\{1\}\}}] + \mathbf{E}_{z,y}[R^1_{Q,s} 1_{\{t<\infty, 2\in Q\}}] \\
&= R^1 \cdot \mathbf{P}_{z,y}(t < \infty, Q = \{1\}) + \mathbf{E}_{z,y}[R^1_{Q,s} 1_{\{t<\infty, 2\in Q\}}] \\
&\geq R^1 \cdot \mathbf{P}_{z,y}(t < \infty, Q = \{1\}) - \mathbf{P}_{z,y}(t < \infty, 2 \in Q).
\end{aligned} \tag{20}$$

Since (x, y) is an $\frac{\epsilon}{2}$-equilibrium, it follows that

$$\gamma^1(z, y) \leq \gamma^1(x, y) + \frac{\epsilon}{2} \leq R^1 - \frac{\epsilon}{2}. \tag{21}$$

Since $\pi(z, y) = \mathbf{P}_{z,y}(t < \infty, Q = \{1\}) + \mathbf{P}_{z,y}(t < \infty, 2 \in Q)$, and by (21), (8) and (20) we get:

$$\begin{aligned}
&(\mathbf{P}_{z,y}(t < \infty, Q = \{1\}) + \mathbf{P}_{z,y}(t < \infty, 2 \in Q)) \cdot R^1 - \pi(z, y) \cdot \frac{\epsilon}{2} = \\
&= \pi(z, y) \cdot (R^1 - \frac{\epsilon}{2}) \geq \pi(z, y) \cdot \gamma^1(z, y) = \rho^1(z, y) \\
&\geq R^1 \cdot \mathbf{P}_{z,y}(t < \infty, Q = \{1\}) - \mathbf{P}_{z,y}(t < \infty, 2 \in Q).
\end{aligned}$$

In particular,

$$\frac{\epsilon}{2}\pi(z,y) \leq (1+R^1)\cdot \mathbf{P}_{z,y}(t<\infty, 2\in Q) \leq 2\cdot \mathbf{P}_{z,y}(t<\infty, 2\in Q).$$

As $\pi(z,0) \leq \pi(z,y)$ and $\mathbf{P}_{z,y}(t<\infty, 2\in Q) \leq \pi(0,y)$ one has

$$\pi(0,y) \geq \mathbf{P}_{z,y}(t<\infty, 2\in Q) \geq \frac{\epsilon}{4}\pi(z,y) \geq \frac{\epsilon}{4}\pi(z,0) = \frac{\epsilon}{4}\cdot r_1,$$

as desired. □

7 Constructing an ϵ-equilibrium

In the present section we use all the tools we have developed so far to construct an ϵ-equilibrium. In section 7.1 we define a procedure that attaches for every finite tree T a color. In section 7.2 we explain the main ideas of the construction. We then proceed with the formal proof.

We fix throughout a stopping game that satisfies conditions **A.1-A.6** in Proposition 5.6. In particular, the constants R^1 and R^2 are fixed. We also fix $\epsilon > 0$ sufficiently small.

7.1 Coloring a Finite Tree

Definition 7.1. Let $a_1 < b_1$ and $a_2 < b_2$. A rectangle $[a_1,b_1]\times[a_2,b_2]$ is *bad* if $R^1-\epsilon \leq a_1$ and $R^2-\epsilon \leq a_2$. It is *good* if $b_1 \leq R^1-\epsilon$ or $b_2 \leq R^2-\epsilon$.

Let M be a finite covering of $[-1,1]^2$ with (not necessarily disjoint) rectangles $[a_1,b_1]\times[a_2,b_2]$ such that $b_1-a_1<\epsilon$ and $b_2-a_2<\epsilon$, all of which are either good or bad. Thus, for every $u\in[-1,1]^2$ there is a rectangle $m\in M$ such that $u\in m$. We denote by $H=\{h_1,h_2,\ldots,h_J\}$ the set of bad rectangles in M, and by $G=\{g_1,g_2,\ldots,g_V\}$ the set of good rectangles in M.

Set $C=G\cup\{\xi\}$. This set is composed of the set G of good rectangles together with another symbol ξ. For every game on a tree T consider the following procedure which attaches an element $c\in C$ to T:

- Set $T^{(0)}=T$.

- For $1 \leq j \leq J$ apply Proposition 6.4 to $T^{(j-1)}$ and $h_j = [a_{j,1}, b_{j,1}] \times [a_{j,2}, b_{j,2}]$, to obtain a subgame $T^{(j)}$ of $T^{(j-1)}$ and strategies $(x_T^{(j)}, y_T^{(j)})$ in T such that[4]

  1. No strict subgame of $T^{(j)}$ has an ϵ-equilibrium with corresponding payoffs in h_j.
  2. Either $T^{(j)} = T^{(j-1)}$ or the following three conditions hold.
     (a) $a_{j,i} - \epsilon \leq \gamma^i(x_T^{(j)}, y_T^{(j)})$ for $i \in \{1, 2\}$.
     (b) For every pair (x', y'), $\gamma^1(x', y_T^{(j)}) \leq b_{j,1} + 7\epsilon$ and $\gamma^2(x_T^{(j)}, y') \leq b_{j,2} + 7\epsilon$.
     (c) $\pi(x_T^{(j)}, y_T^{(j)}) \geq \epsilon^2 \cdot p_{T^{(j)},T^{(j-1)}}$, where $p_{T^{(j)},T^{(j-1)}}$ is defined in section 6.1.

- If $T^{(J)}$ is trivial (that is, the only node is the root,) set $c(T) = \xi$. Otherwise choose a stationary $\frac{\epsilon}{2}$-equilibrium $(x^{(0)}, y^{(0)})$ of $T^{(J)}$. By construction, the corresponding $\frac{\epsilon}{2}$-equilibrium payoff lies in a good rectangle $g \in G$. Set $c(T) = g$.

7.2 The Main Idea of the Construction

Before formally constructing a $K\epsilon$-equilibrium strategy pair for some fixed $K > 0$, we explain the basic idea of the construction.

Assume for simplicity that all the σ-algebras $\mathcal{F}_n$ are finite. In this case, every $n \geq 0$, every $\omega \in \Omega$ and every stopping time τ such that $\tau(\omega) > n$ define naturally a game $\Gamma_{n,\tau}(\omega)$ on a tree; the root is the atom F of $\mathcal{F}_n$ that contains ω, the nodes are all atoms $F' \in \cup_{m \geq n} \mathcal{F}_m$ that satisfy[5] (i) $F' \subseteq F$, and (ii) if $F' \in \mathcal{F}_m$ then $\tau \geq m$ on F'. All atoms F' where there is an equality in (ii) are leaves.

In section 7.1 we attached to each such triplet an element from a finite set C - a color. By Theorem 4.3, there is a sequence of bounded stopping times $0 \leq \tau_0 \leq \tau_1 \leq \cdots$ such that $p(c_{\tau_0,\tau_1} = c_{\tau_j,\tau_{j+1}} \quad \forall j > 0) \geq 1 - \epsilon$.

Fix for a moment $l \geq 0$. In section 7.1 we constructed for each one of the finitely many trees $T = \Gamma_{\tau_l(\omega),\tau_{l+1}}(\omega)$ and each $j = 0, \ldots, J$ a subtree

[4] $(x_T^{(j)}, y_T^{(j)})$ as given by Proposition 6.4 are strategies in $T^{(j-1)}$. We extend them to strategies in T by letting them continue from the leaves of $T^{(j-1)}$ downward.

[5] In this union, a set F which is an atom of *several* $\mathcal{F}_m$'s is counted several times. Thus, the union is actually a union of pairs $\{(m, F), F \text{ is an atom of } \mathcal{F}_m\}$.

$T^{(j)}$ and a pair of stationary strategies $(x_l^{(j)}, y_l^{(j)})$. The leaves of the subtrees define naturally a stopping time $\tau_l^{(j)}$. Thus, we obtain a sequence of stopping times $\tau_l \leq \tau_l^{(J)} \leq \tau_l^{(J-1)} \leq \ldots \leq \tau_l^{(1)} \leq \tau_l^{(0)} = \tau_{l+1}$. Let $(x_l^{(j)}, y_l^{(j)})_{j=0}^J$ be the collection of strategy pairs that was generated during the coloring procedure.

For $1 \leq j \leq J$, let $I_j = \{\tau_l^{(j)} < \tau_l^{(j-1)} \text{ i.o.}\}$. Set $G = (\cup_j I_j)^c$. Then, on G, $\tau_l^{(J)} < \tau_{l+1}$ only finitely many times. In particular, there is $L \geq 0$ sufficiently large such that $p\left(\tau_l^{(J)} < \tau_{l+1} \text{ for some } l \geq L \mid G\right) < \epsilon$. Assume w.l.o.g. that $L = 0$. Set $G_v = G \cap \{\tau_l^{(J)} = \tau_{l+1} \quad \forall l\} \cap \{c_{\tau_l, \tau_{l+1}} = g_v \quad \forall l \geq 0\}$, for every $v = 1, \ldots, V$.

Modulo punishment strategies, on I_j, the $K\epsilon$-equilibrium strategy pair coincides with the concatenation of the strategy pairs $(x_l^{(j)}, y_l^{(j)})$. It yields payoff in the rectangle h_j. The condition $\{\tau_l^{(j)} < \tau_l^{(j-1)} \text{ i.o.}\}$ ensures that under the concatenation the game will eventually terminate with probability 1. On G_v, the $K\epsilon$-equilibrium strategy pair coincides with the concatenation of the strategy pairs $(x_l^{(0)}, y_l^{(0)})$.

When the filtration is general, one needs to approximate the $\mathcal{F}_n$'s by finite sub-σ-algebras. This fact introduces some technical difficulties, but do not alter the general idea.

Adding a threat of punishment might be necessary as the following example shows.

Example 7.2. Consider a game with deterministic payoffs: $R_{\{1\},n} = (-1, 2)$, $R_{\{2\},n} = (-2, 1)$, and $R_{\{1,2\},n} = (0, -3)$. We first argue that all ϵ-equilibrium payoffs are close to $(-1, 2)$.

Given a strategy x of player 1, player 2 can always wait until the probability of termination under x is exhausted, and then stop. Therefore, in any ϵ-equilibrium, the probability of termination is at least $1 - \epsilon$, and the corresponding payoff is close to the convex hull of $(-1, 2)$ and $(-2, 1)$. Since player 1 can always guarantee -1 by stopping at the first stage, the claim follows.

However, in every ϵ-equilibrium (x, y), we must have $\mathbf{P}_{0,y}(\theta < \infty) \geq 1/2$, otherwise player 1 receives more than -1 by never stopping.

Thus, an ϵ-equilibrium will have the following structure, for some integer N. Player 1 stops with probability at least $1 - \epsilon$ before stage N, and with probability at most ϵ after that stage; player 2 stops with probability at most ϵ before stage N, and with probability at least $1/2$ after that stage.

The strategy of player 2 serves as a threat of punishment: if player 1 does not stop before stage N, he will be punished in subsequent stages.

7.3 Notations

Denote $\delta_n = \epsilon^2/2^{n+2}$ for each $n \geq 0$. Set $\Delta_n = \sum_{k \geq n} \delta_k = \epsilon^2/2^{n+1}$, so that $\sum_{n \geq 0} \Delta_n = \epsilon^2$.

For every $i \geq 0$ and every $n \in \mathbf{N}$ we choose once and for all a partition $\mathcal{B}_i^n$ of the $n-1$-dimensional simplex $\Delta(n) = \{x \in \mathbf{R}^n \mid \sum_{j=1}^n x_j = 1, x_j \geq 0 \ \ \forall j\}$ such that the diameter of each element in $\mathcal{B}_i^n$ is less than δ_i in the norm $\|\cdot\|_1$. We furthermore choose once and for all for each $B \in \mathcal{B}_i^n$ an element $q_B \in B$.

Definition 7.3. Let $\mathcal{F} = (\mathcal{F}_n, \mathcal{F}_{n+1}, \ldots, \mathcal{F}_M)$ be a sequence of σ-algebras. A $\mathcal{F}$*-strategy* x for player 1 is a collection $x = (x_i)_{i=n}^M$, where for each i, x_i is a $\mathcal{F}_i$-measurable $[0,1]$-valued r.v. $\mathcal{F}$-strategies y of player 2 are defined analogously.

Given a pair (x, y) of $\mathcal{F}$-strategies and a $\mathcal{F}$-adapted stopping time $\tau > n$, we denote by $\pi(x, y; \mathcal{F}, n, \tau)$ the conditional probability under (x, y) that the game that start at stage n ends before τ, and by $\rho(x, y; \mathcal{F}, n, \tau)$ the corresponding expected payoff. We define $\gamma(x, y; \mathcal{F}, n, \tau) = \frac{\rho(x,y;\mathcal{F},n,\tau)}{\pi(x,y;\mathcal{F},n,\tau)}$. These are $\mathcal{F}_n$-measurable r.v.s.

In the sequel, the sequence $(\mathcal{F}_n, \mathcal{F}_{n+1}, \ldots, \mathcal{F}_M)$ in Definition 7.3 will either coincide with the filtration of the game, or be a sequence of finite sub-σ-algebras that, in some sense, approximate the filtration.

7.4 Close Games

Let T be a stopping game on a finite tree with payoffs bounded by 1. Recall that S_0 is the set of nodes which are not leaves, and $(p_s)_{s \in S_0}$ are the probability distributions over children.

Let $\widetilde{T}$ be a game that coincides with T, except for the probability distributions over children $(\widetilde{p}_s)_{s \in S_0}$ which satisfy

$$\|p_s - \widetilde{p}_s\|_1 \leq \eta_{\text{depth}(s)},$$

where $(\eta_j)_{j \geq 0}$ is a sequence of positive reals. Observe that the set of strategies of the two players in T and in $\widetilde{T}$ coincide.

Under these notations we have the following estimates.

Lemma 7.4. *For every pair of stationary strategies (x, y) in $\Gamma(T)$ (or, equivalently, in $\Gamma(\widetilde{T})$) (a) $|\pi(x, y; T) - \pi(x, y; \widetilde{T})| \leq \sum_{j\geq 0} \eta_j$, and (b) for $i = 1, 2$, $|\rho^i(x, y; T) - \rho^i(x, y; \widetilde{T})| \leq \sum_{j\geq 0} \eta_j$, (c) for every subtree T' of T $|p_{T',T} - p_{\widetilde{T}',\widetilde{T}}| \leq \sum_{j\geq 0} \eta_j$, where $\widetilde{T}'$ is the subtree in $\widetilde{T}$ that corresponds to T', and (d) $|r_1(T) - r_1(\widetilde{T})| < \sum_{j\geq 0} \eta_j$ where $r_1(T)$ is the quantity defined in Lemma 6.19.*

Corollary 7.5. *Set $\eta^* = \sum_{j\geq 0} \eta_j$, and let $\epsilon > 0$. Let x be a strategy for player 1. Then for every strategy z of player 2 such that $\pi(x, z; \widetilde{T}) > \eta^*/\epsilon$ we have $\gamma^2(x, z; \widetilde{T}) \leq \gamma^2(x, z; T) + 2\epsilon$.*

Proof. By Lemma 7.4(a,b)

$$\gamma^2(x, z; \widetilde{T}) = \frac{\rho^2(x, z; \widetilde{T})}{\pi(x, z; \widetilde{T})} \leq \frac{\rho^2(x, z; T) + \eta^*}{\pi(x, z; \widetilde{T})} = \frac{\pi(x, z; T)\gamma^2(x, z; T) + \eta^*}{\pi(x, z; \widetilde{T})}$$
$$\leq \frac{\pi(x, z; \widetilde{T})\gamma^2(x, z; T) + 2\eta^*}{\pi(x, z; \widetilde{T})} \leq \gamma^2(x, z; T) + 2\epsilon. \quad (22)$$

□

7.5 From Games on Trees to Stopping Games

In this section we provide several constructions that relate a stopping game to games on trees.

Let $\mathcal{G} = (\mathcal{G}_n)_{n\geq 0}$ be a sequence of finite σ-algebras of $\mathcal{A}$ such that for every $n \geq 0$ (i) $\mathcal{G}_n \subseteq \mathcal{F}_n$, and (ii) R_n is $\mathcal{G}_n$-measurable. Let τ be a $\mathcal{G}$-adapted stopping time. Assume that moreover for every $n \geq 0$ and every atom $F \in \mathcal{G}_n$ we are given a probability distribution q_F over the atoms of $\mathcal{G}_{n+1}$ which are subsets of F. One can define naturally for every $\omega \in \Omega$ and every $n < \tau(\omega)$ a game on a tree $T(n, \tau, \omega)$ as follows.

- The root is the atom F of $\mathcal{G}_n$ that contains ω.
- The nodes are all atoms $F' \in \cup_{m\geq n}\mathcal{G}_m$ such that (a) $F' \subseteq F$, and (b) if $F' \in \mathcal{G}_m$, then $\tau \geq m$ on F'.
- The leaves are all atoms $F' \in \cup_{m\geq n}\mathcal{G}_m$ where there is equality in (b).

- Payoff is given by $(R_m)_{n\leq m\leq\tau}$ (recall that R_n is $\mathcal{G}_n$-measurable for every n).
- Transition from any node F' is given by $q_{F'}$.

Every $\mathcal{G}$-strategy x induces naturally a strategy in $T(n,\tau,\omega)$: take into account the behavior of x only at nodes which are not leaves.

Let $\sigma<\tau$ be two bounded $\mathcal{G}$-adapted stopping times. Then $\mathcal{G}_\sigma$ is a finite σ-algebra, and the set $\{T(\sigma(\omega),\tau,\omega),\omega\in\Omega\}$ is finite. Assume that for *every* T in this set we are given a subgame T'. That is, we are given a $\mathcal{G}_\sigma$-measurable function T' such that for every $\omega\in\Omega$, $T'(\omega)$ is a subgame of $T(\sigma(\omega),\tau,\omega)$. The leaves of all the subgames define naturally a stopping time ν in the following way.

$$\nu(\omega)=m \quad \Leftrightarrow \quad \text{The leaf of } T'(\omega) \text{ that contains } \omega \text{ is an atom of } \mathcal{G}_m.$$

Let $0=\tau_0<\tau_1<\cdots$ be an increasing sequence of bounded stopping times. Assume that for every $l\geq 0$ and every $\omega\in\Omega$ we are given a strategy $x(l,\omega)$ in the game on a tree $T(\tau_l(\omega),\tau_{l+1},\omega)$, and that the function $\omega\mapsto x(l,\omega)$ is $\mathcal{G}_{\tau_l}$-measurable.

One can define naturally a strategy x in the stopping game Γ by concatenating the strategies $(x(l,\cdot))_{l\geq 0}$.

Conversely, every $\mathcal{G}$-measurable strategy x in the stopping game Γ induces a strategy $x(l,\omega)$ in the game $T(\tau_l(\omega),\tau_{l+1},\omega)$, for every $l\geq 0$ and every $\omega\in\Omega$.

7.6 Representative Approximations

Throughout this subsection we fix two integers $0\leq n<M$, and an increasing sequence $\mathcal{G}=(\mathcal{G}_n,\ldots,\mathcal{G}_M)$ of finite partitions of Ω, such that for each $i=n,\ldots,M$, (i) $\mathcal{G}_i\subseteq\mathcal{F}_i$, and (ii) R_i is $\mathcal{G}_i$-measurable.

Definition 7.6. We say that $\mathcal{G}$ *δ-approximates* $\mathcal{F}$ on $n,\ldots,M$ if for every $i=n,\ldots,M-1$, $\sum_{G'\in\mathcal{G}_{i+1}}|\mathbf{P}(G'\mid\mathcal{F}_i)-\mathbf{P}(G'\mid\mathcal{G}_i)|\leq\delta_i$ a.e.

Alternatively, $\mathcal{G}$ δ-approximates $\mathcal{F}$, if for every $i=n,\ldots,M-1$ and every $\mathcal{G}_{i+1}$-measurable function h such that $|h|\leq 1$, $|\mathbf{E}(h|\mathcal{F}_i)-\mathbf{E}(h|\mathcal{G}_i)|\leq\delta_i$.

Two simple yet important properties of δ-approximating games are the following.

Lemma 7.7. *Assume that the sequence $\mathcal{G}$ δ-approximates $\mathcal{F}$, and let τ be a $\mathcal{G}$ adapted stopping time. Let (x, y) be a pair of $\mathcal{G}$-strategies. Then*

1. $|\pi(x, y; \mathcal{G}, n, \tau) - \pi(x, y; \mathcal{F}, n, \tau)| \leq \sum_{j \geq n} \delta_j = \Delta_n$ *a.e.*

2. $|\rho^i(x, y; \mathcal{G}, n, \tau) - \rho^i(x, y; \mathcal{F}, n, \tau)| \leq \sum_{j \geq n} \delta_j = \Delta_n$ *a.e.*, $i = 1, 2$.

The following Lemma states that if $\mathcal{G}$ δ-approximates $\mathcal{F}$, and if the opponent plays a $\mathcal{G}$-strategy, then a player does not lose much by considering only $\mathcal{G}$-strategies.

Lemma 7.8. *Assume that the sequence $\mathcal{G}$ δ-approximates $\mathcal{F}$, and let $\tau > n$ be a $\mathcal{G}$-adapted stopping time. Let x be a $\mathcal{G}$-strategy for player 1, and set $\gamma = \text{esssup}\{\gamma^2(x, y; \mathcal{G}, n, \tau), y$ is a $\mathcal{G}$-strategy$\}$. Then, for every $\mathcal{F}$-strategy y,*

$$\rho^2(x, y; \mathcal{F}, n, \tau) \leq \gamma \cdot \pi(x, y; \mathcal{F}, n, \tau) + \Delta_n \text{ a.e.}$$

Proof. Let $\alpha(\mathcal{G}) = \text{esssup}\{\rho^2(x, y; \mathcal{G}, n, \tau) + \gamma \cdot (1 - \pi(x, y; \mathcal{G}, n, \tau)), y \text{ is a } \mathcal{G}\text{-strategy}\}$. $\alpha(\mathcal{G})$ is the best possible payoff for player 2 in the game that starts at stage n and, if no player stopped before stage τ, terminates with payoff γ. From the definition of γ it follows that $\alpha(\mathcal{G}) \leq \gamma$. Plainly $\alpha(\mathcal{G}) = \alpha(n, \mathcal{G})$, where $(\alpha(i, \mathcal{G}))_{i=n}^M$ are given by

$$\alpha(i, \mathcal{G}) = \begin{cases} \gamma & i \geq \tau \\ \max\{\mathbf{E}(\alpha(i+1, \mathcal{G})|\mathcal{G}_i), x_i \cdot R^2_{i,\{1,2\}} + (1 - x_i) \cdot R^2_{i,\{2\}}\} & i < \tau \end{cases}. \tag{23}$$

Similarly, let $\alpha(\mathcal{F}) = \text{esssup}\{\rho^2(x, y; \mathcal{F}, n, \tau) + \gamma \cdot (1 - \pi(x, y; \mathcal{F}, n, \tau)), y \text{ is a } \mathcal{F}\text{-strategy}\}$. Then $\alpha(\mathcal{F}) = \alpha(n, \mathcal{F})$, where $(\alpha(i, \mathcal{F}))_{i=n}^M$ are given by

$$\alpha(i, \mathcal{F}) = \begin{cases} \gamma & i \geq \tau \\ \max\{\mathbf{E}(\alpha(i+1, \mathcal{F})|\mathcal{F}_i), x_i \cdot R^2_{i,\{1,2\}} + (1 - x_i) \cdot R^2_{i,\{2\}}\} & i < \tau \end{cases}. \tag{24}$$

Since x_i and R^2_i are $\mathcal{G}_i$-measurable, it follows from (23), (24) and the remark that follows Definition 7.6, that $\alpha(i, \mathcal{F}) \leq \alpha(i, \mathcal{G}) + \sum_{j=i}^M \delta_j$ for $i = n, \ldots, M$. In particular $\alpha(\mathcal{F}) = \alpha(n, \mathcal{F}) \leq \alpha(n, \mathcal{G}) + \Delta_n \leq \gamma + \Delta_n$. It follows that for every $\mathcal{F}$-strategy y, $\rho^2(x, y; \mathcal{F}, n, \tau) + \gamma \cdot (1 - \pi(x, y; \mathcal{F}, n, \tau)) \leq \gamma + \Delta_n$, which implies $\rho^2(x, y; \mathcal{F}, n, \tau) \leq \gamma \cdot \pi(x, y; \mathcal{F}, n, \tau) + \Delta_n$. $\square$

7.7 Constructing Approximating Games

We now define a NT-function Γ.[6] The range of the r.v. $\Gamma_{n,\tau}$ is games on trees. In the next subsection we show that our construction is consistent.

Fix a non-negative integer $n \geq 0$ and a bounded stopping time τ.

Define a r.v. $K_n^\tau = \min\{k \geq n, \mathbf{P}(\tau \leq k \mid \mathcal{F}_n) = 1\}$. Since τ is bounded, K_n^τ is bounded as well, and by definition it is $\mathcal{F}_n$-measurable.

For every $k \geq n$, every $m \in \{n, n+1, \ldots, k\}$, and every $\widehat{R} = (\widehat{R}_i)_{i=n}^m \in \mathcal{R}^{m-n+1}$, define

$$A_{k,n}^\tau(m, \widehat{R}) = \{K_n^\tau = k, \tau = m, R_i = \widehat{R}_i \quad \forall i \in \{n, n+1, \ldots, m\}\} \in \mathcal{F}_m.$$

Set

$$\mathcal{A}_{k,n,m}^\tau = \{A_{k,n}^\tau(m, \widehat{R}), \widehat{R} \in \mathcal{R}^{m-n+1}\}.$$

Define

$$T_{k,m}^\tau = \{K_n^\tau = k, \tau < m\}.$$

Observe that $\mathcal{A}_{k,n,m}^\tau \cup \{T_{k,m}^\tau\}$ is a partition of $\{K_n^\tau = k, \tau \leq m\}$.

We fix $k \geq n$, and restrict ourselves to the $\mathcal{F}_n$-measurable set $\{K_n^\tau = k\}$.

We now construct an increasing sequence of finite partitions $\widehat{\mathcal{F}}_{k,n}^\tau, \ldots, \widehat{\mathcal{F}}_{k,k}^\tau$ of the set $\{K_n^\tau = k\}$, and for every $m \in \{n, n+1, \ldots, k-1\}$ and every $F \in \widehat{\mathcal{F}}_{k,m}$ a probability distribution $q_F \in \Delta(|\widehat{\mathcal{F}}_{k,m+1}^\tau|)$, that satisfy the following properties for every $m \in \{n, n+1, \ldots, k-1\}$.

1. R_m is $\widehat{\mathcal{F}}_{k,m}^\tau$-measurable.

2. For every $F \in \widehat{\mathcal{F}}_{k,m}$, $\|\mathbf{P}(\cdot \mid \widehat{\mathcal{F}}_{k,m}^\tau) - q_F\|_1 < \delta_m$ on F.

Set

$$\widehat{\mathcal{F}}_{k,k}^\tau = \mathcal{A}_{k,n,k}^\tau \cup \{T_{k,k}^\tau\}.$$

Assume we have already defined $\widehat{\mathcal{F}}_{k,m+1}^\tau, \ldots, \widehat{\mathcal{F}}_{k,k}^\tau$. Recall that $\mathcal{B} = \mathcal{B}_m^{|\widehat{\mathcal{F}}_{k,m+1}^\tau|}$ is a partition of the set $\Delta(|\widehat{\mathcal{F}}_{k,m+1}^\tau|)$ into sets with diameter smaller than δ_m.

Define a function $g_{k,m}^\tau : \Omega \to \Delta(|\widehat{\mathcal{F}}_{k,m+1}^\tau|)$ by

$$g_{k,m}^\tau[A] = \mathbf{P}(A \mid \mathcal{F}_m), \quad \forall A \in \widehat{\mathcal{F}}_{k,m+1}^\tau.$$

[6] Recall that NT-functions are defined in Definition 4.1.

Let $\mathcal{G}^{\tau}_{k,m}$ be the inverse image of $\mathcal{B}$ under $g^{\tau}_{k,m}$. For each $G \in \mathcal{G}^{\tau}_{k,m}$ assign the element q_B, where $B \in \mathcal{B}$ is the unique set that satisfies $g^{\tau}_{k,m}(G) \subseteq B$.

Finally, define

$$\widehat{\mathcal{F}}^{\tau}_{k,m} = \mathcal{G}^{\tau}_{k,m} \cup \mathcal{A}^{\tau}_{k,n,m} \cup \{T^{\tau}_{k,m}\}.$$

Though the partition $\widehat{\mathcal{F}}^{\tau}_{k,m}$ depends on τ, the number of elements in this partition is independent of τ, and for every two bounded stopping times $\tau_1, \tau_2 > n$ there is a natural 1-1 mapping from $\widehat{\mathcal{F}}^{\tau_1}_{k,m}$ to $\widehat{\mathcal{F}}^{\tau_2}_{k,m}$:

$$\begin{array}{rcl} A^{\tau_1}_{k,n}(m, \widehat{R}) & \mapsto & A^{\tau_2}_{k,n}(m, \widehat{R}), \\ T^{\tau_1}_{k,m} & \mapsto & T^{\tau_2}_{k,m}, \text{ and} \\ (g^{\tau_1}_{k,m})^{-1}(B) & \mapsto & (g^{\tau_2}_{k,m})^{-1}(B), \quad \forall B \in \mathcal{B}_m^{|\widehat{\mathcal{B}}^{\tau}_{k,m+1}|}. \end{array}$$

As the sequence $(\widehat{\mathcal{F}}^{\tau}_{k,m})_m$ is not increasing, we replace $\widehat{\mathcal{F}}^{\tau}_{k,m}$ with $\widehat{\mathcal{F}}^{\tau}_{k,n} \bigvee \cdots \bigvee \widehat{\mathcal{F}}^{\tau}_{k,m}$.

The sequence $\widehat{\mathcal{F}}^{\tau}_{k,n}, \ldots, \widehat{\mathcal{F}}^{\tau}_{k,k}$, the collection $(q_F)_{F \in \widehat{\mathcal{F}}^{\tau}_{k,m}, n \leq m < k}$, and every $\omega \in \Omega$, define naturally a game played on a finite tree, as explained in section 7.5. We define $\Gamma_{n,\tau}(\omega)$ to be this game.

7.8 The Construction is $\mathcal{F}$-consistent

We here prove that Γ is $\mathcal{F}$-consistent.

Fix $n \geq 0$, a $\mathcal{F}_n$-measurable set F, and two bounded stopping times τ_1, τ_2 that satisfy (a) $\tau_1, \tau_2 > n$ on F, and (b) $\tau_1 = \tau_2$ on F.

Since F is $\mathcal{F}_n$-measurable, and since K^{τ}_n is $\mathcal{F}_n$-measurable for any bounded stopping time τ, we have the following.

Lemma 7.9. *$K^{\tau_1}_n = K^{\tau_2}_n$ on F.*

For the rest of the section we fix $k \geq n$, and we restrict ourselves to the set $F_k = F \cap \{K^{\tau_1}_n = k\} = F \cap \{K^{\tau_2}_n = k\}$.

The following lemma holds since $\tau_1 = \tau_2$ on F_k.

Lemma 7.10. *For every $m \in \{n+1, \ldots, k\}$ the following three assertions hold.*

a) $A^{\tau_1}_{k,n}(m, \widehat{R}) \cap F_k = A^{\tau_2}_{k,n}(m, \widehat{R}) \cap F_k$, for every $\widehat{R} \in \mathcal{R}^{m-n+1}$.

b) $T^{\tau_1}_{k,m} \cap F_k = T^{\tau_2}_{k,m} \cap F_k$.

c) $(g^{\tau_1}_{k,m})^{-1}(B) \cap F_k = (g^{\tau_2}_{k,m})^{-1}(B) \cap F_k$ *for every* $B \in \mathcal{B}_n^{\widehat{\mathcal{F}}^{\tau}_{k,m}}$.

The claim follows by the following Lemma.

Lemma 7.11. *For every* $m \in \{n+1, \ldots, k\}$ *the following three assertions hold.*

a) $\mathbf{P}(A^{\tau_1}_{k,n}(m, \widehat{R}) \mid \mathcal{F}_{m-1}) = \mathbf{P}(A^{\tau_2}_{k,n}(m, \widehat{R}) \mid \mathcal{F}_{m-1})$ *on* F_k, *for every* $\widehat{R} \in \mathcal{R}^{m-n+1}$.

b) $\mathbf{P}(T^{\tau_1}_{k,m} \mid \mathcal{F}_{m-1}) = \mathbf{P}(T^{\tau_2}_{k,m} \mid \mathcal{F}_{m-1})$ *on* F_k.

c) $\mathbf{P}\left((g^{\tau_1}_{k,m})^{-1}(B) \mid \mathcal{F}_{m-1}\right) = \mathbf{P}\left((g^{\tau_2}_{k,m})^{-1}(B) \mid \mathcal{F}_{m-1}\right)$ *for every* $B \in \mathcal{B}_n^{|\widehat{\mathcal{F}}^{\tau}_{k,m}|}$.

Proof. The proof follows by Lemma 7.10 and by the following simple fact. If $\mathcal{G}$ is a sub-σ-field of $\mathcal{F}$, $A \in \mathcal{F}$ and $B \in \mathcal{G}$, then $\mathbf{P}(A \mid \mathcal{G}) = \mathbf{P}(A \cap B \mid \mathcal{G})$ a.e. on B. □

7.9 Applying Theorem 4.3

For every finite tree T apply the procedure presented in section 7.1. This procedure yields (i) a sequence $(T^{(j)})_{j=1}^{k}$ of subtrees of T, (ii) a sequence of stationary strategy pairs $(x_T^{(j)}, y_T^{(j)})_{j=0}^{J}$, (iii) an element $c(T) \in C$, and (iv) if $c(T) \neq \xi$, a stationary $\frac{\epsilon}{2}$-equilibrium $(x_T^{(0)}, y_T^{(0)})$ in $T^{(J)}$.

We set $\chi^1(T) = 1$ if $\pi(x_T^{(0)}, 0; T^{(J)}) \geq \frac{\Delta_n}{\epsilon} + \Delta_n$, and 0 otherwise. We set $\chi^2(T) = 1$ if $\pi(0, y_T^{(0)}; T^{(J)}) \geq \frac{\Delta_n}{\epsilon} + \Delta_n$, and 0 otherwise.

Set $C^* = C \times \{0,1\}^2$. We now define a C^*-valued $\mathcal{F}$-consistent NT-function c.

For every $n \geq 0$ and every bounded $\mathcal{F}$-adapted stopping time τ set

$$c^*_{n,\tau} = (c(\Gamma_{n,\tau}), \chi^1(\Gamma_{n,\tau}), \chi^2(\Gamma_{n,\tau})).$$

Since the r.v. $\Gamma_{n,\tau}$ is $\mathcal{F}_n$-measurable and has finite range, c^* is $\mathcal{F}_n$-measurable. Since Γ is $\mathcal{F}$-consistent, so is c^*.

By Theorem 4.3 there is an increasing sequence of bounded stopping times $0 \leq \tau_0 < \tau_1 \leq \cdots$ such that

$$p(c^*_{\tau_0,\tau_1} = c^*_{\tau_l,\tau_{l+1}} \quad \forall l \geq 0) \geq 1 - \epsilon. \tag{25}$$

7.10 The Relation Between the Finite Games and the Original Game

To simplify computations, it is convenient to assume that the players continue to play even if the game is stopped. That is, at stage $\theta+1$ the players keep on playing, as if no player stopped at stage θ. The payoff, however, does not depend on the play after stage θ. We denote by θ_l the first stage bigger or equal to τ_l in which at least one player stops. We denote by Q_l the subset of players that stop at stage θ_l.

For every fixed $l \geq 0$ the range of $\Gamma_{\tau_l,\tau_{l+1}}$ is a finite set $\{T_1,\ldots,T_U\}$. For each $u=1,\ldots,U$, denote by $T_u^{(j)}$ the j'th subgame of T_u generated in the coloring procedure for $T=T_u$ presented in section 7.1.

For $1 \leq j \leq J$, the leaves of $(T_u^{(j)})_{u=1}^U$ define, as explained in section 7.5, a stopping time $\tau_l^{(j)}$, $\tau_l \leq \tau_l^{(j)} \leq \tau_{l+1}$. Thus, one obtains a sequence of increasing stopping times $\tau_l \leq \tau_l^{(J)} \leq \tau_l^{(J-1)} \leq \cdots \leq \tau_l^{(0)} = \tau_{l+1}$.

Let $\widehat{\mathcal{F}}_{\tau_l}$ be the partition that contains for every $\omega \in \Omega$ the atom of $\widehat{\mathcal{F}}^{\tau_{l+1}}_{K^{\tau_{l+1}}_{\tau_l(\omega)}(\omega),\tau_l(\omega)}$ that contains ω. Observe that $\widehat{\mathcal{F}}_{\tau_l}$ is a finite partition of Ω, but the sequence $(\widehat{\mathcal{F}}_{\tau_l})_{l\geq 0}$ is not increasing. Thus, $\Gamma_{\tau_l,\tau_{l+1}}$ is $\widehat{\mathcal{F}}_{\tau_l}$-measurable.

As explained in section 7.5, each pair of strategies (x,y) in Γ induce a pair of strategies $(x(\Gamma_{\tau_l,\tau_{l+1}}(\omega)), y(\Gamma_{\tau_l,\tau_{l+1}}(\omega)))$ in the game $\Gamma_{\tau_l,\tau_{l+1}}(\omega)$, for every $l \geq 0$ and every $\omega \in \Omega$.

For every pair (x,y) of strategies denote by $\pi(x,y;\Gamma_{\tau_l,\tau_{l+1}})$ the probability of termination in the first round of the game $\Gamma_{\tau_l,\tau_{l+1}}$ under $(x(\Gamma(\tau_l,\tau_{l+1})(\omega)), y(\Gamma(\tau_l,\tau_{l+1})(\omega)))$, and by $\rho(x,y;\Gamma_{\tau_l,\tau_{l+1}})$ the expected payoff in the first round. The r.v.s $\pi(x,y;\Gamma_{\tau_l,\tau_{l+1}})$ and $\rho(x,y;\Gamma_{\tau_l,\tau_{l+1}})$ are $\widehat{\mathcal{F}}_{\tau_l}$-measurable.

By Lemma 7.7 one has almost everywhere

$$\begin{aligned}
|\mathbf{P}_{x,y}(\theta_l < \tau_{l+1} \mid \widehat{\mathcal{F}}_{\tau_l}) - \mathbf{P}_{x,y}(\theta_l < \tau_{l+1} \mid \mathcal{F}_{\tau_l})| &< \Delta_{\tau_l}, \text{ and} \\
|\mathbf{E}_{x,y}(R^i_{Q_l,\theta_l} 1_{\{\theta_l<\tau_{l+1}\}} \mid \widehat{\mathcal{F}}_{\tau_l}) - \mathbf{E}_{x,y}(R^i_{Q_l,\theta_l} 1_{\{\theta_l<\tau_{l+1}\}} \mid \mathcal{F}_{\tau_l})| &< \Delta_{\tau_l}, \text{ for } i=1,2,
\end{aligned} \tag{26}$$

whereas by Lemma 7.4

$$\begin{aligned}
|\mathbf{P}_{x,y}(\theta_l < \tau_{l+1} \mid \widehat{\mathcal{F}}_{\tau_l}) - \pi(x,y;\Gamma_{\tau_l,\tau_{l+1}})| &< \Delta_{\tau_l}, \text{ and} \\
|\mathbf{E}_{x,y}(R^i_{Q_l,\theta_l} 1_{\{\theta_l<\tau_{l+1}\}} \mid \widehat{\mathcal{F}}_{\tau_l}) - \rho^i(x,y;\Gamma_{\tau_l,\tau_{l+1}})| &< \Delta_{\tau_l}, \text{ for } i=1,2.
\end{aligned} \tag{27}$$

By (26) and (27) one has almost everywhere

$$|\pi(x,y;\Gamma_{\tau_l,\tau_{l+1}}) - \mathbf{P}_{x,y}(\theta_l < \tau_{l+1} \mid \mathcal{F}_{\tau_l})| < 2\Delta_{\tau_l}, \text{ and} \tag{28}$$

$$|\rho^i(x,y;\Gamma_{\tau_l,\tau_{l+1}}) - \mathbf{E}_{x,y}(R^i_{Q_l,\theta_l} 1_{\{\theta_l<\tau_{l+1}\}} \mid \mathcal{F}_{\tau_l})| < 2\Delta_{\tau_l} \text{ for } i = 1,2. \tag{29}$$

For $1 \leq j \leq J$, let $I_j = \{\tau_l^{(j)} < \tau_l^{(j-1)} \text{ i.o.}\}$. Set $G = (\cup_j I_j)^c$, and, for $1 \leq v \leq V$, $G_v = G \cap \{g_v = c_{\tau_0,\tau_1} = c_{\tau_1,\tau_2} = \ldots\}$. Note that $\{\xi = c_{\tau_0,\tau_1} = c_{\tau_1,\tau_2} = \ldots\} = \cap_{l\geq 0}\{\tau_l = \tau_l^{(k)}\} \subseteq \cap_{l\geq 0}\{\tau_l^{(k)} < \tau_{l+1}\} \subseteq I_1 \cup \ldots \cup I_k$. By (25) $p(\bigcup_j I_j \cup \bigcup_v G_v) > 1-\epsilon$. Let $(\bar{I}_j)_{1\leq j\leq J}, (\bar{G}_v)_{1\leq v\leq V} \in \bigcup_{n\geq 0}\mathcal{F}_n$ be mutually disjoint sets such that:

$$p\left(\bigcup_{v=1}^{V} \bar{G}_v \cup \bigcup_{j=1}^{J} \bar{I}_j\right) > 1-2\epsilon, \tag{30}$$

$$p(I_j \mid \bar{I}_j) > 1-\epsilon, 1 \leq j \leq J, \text{ and} \tag{31}$$

$$p\big(G_v \mid \bar{G}_v\big) > 1-\epsilon, 1 \leq v \leq V. \tag{32}$$

We assume also w.l.o.g that $\bar{G}_v, \bar{I}_j \in \mathcal{F}_{\tau_0}$ for $v = 1,\ldots,V$ and $j = 1,\ldots,J$, and that

$$p\big(\cap_{l\geq 0}\{\tau_l^{(J)} = \tau_{l+1}\} \mid G_v \cap \bar{G}_v\big) > 1-\epsilon; \tag{33}$$

if necessary, start with τ_L instead of τ_0 for a sufficiently large $L \in \mathbf{N}$.

By Lemma 5.5 it is sufficient to prove that the games $\Gamma_{\bar{I}_j,\tau_0}$ (the game restricted to $\bar{I}_j$ and starting from τ_0) and $\Gamma_{\bar{G}_v,\tau_0}$ admit ϵ-equilibria. We therefore assume w.l.o.g that $\tau_0 = 0$ and deal separately with the games restricted to $\bar{I}_j$ and $\bar{G}_v$.

7.11 The Game Restricted to $\bar{I}_j$

We here consider the game restricted to $\bar{I}_j$, for some $j = 1,\ldots,J$. Denote $h_j = [a_1,b_1] \times [a_2,b_2]$. Let (x,y) be the concatenation of the strategies $(x^{(j)}_{\Gamma_{\tau_l,\tau_{l+1}}}, y^{(j)}_{\Gamma_{\tau_l,\tau_{l+1}}})$.

We first prove that

$$\mathbf{P}_{x,y}(\theta < \infty \mid I_j) = 1. \tag{34}$$

Indeed, by (28), the construction of $(\tau_l^{(j)})_{l\geq 0}$, Proposition 6.4(2c), and

Lemma 7.4(c), for every $l \geq 0$,

$$\begin{aligned}
\mathbf{P}_{x,y}(\theta_l < \tau_{l+1} \mid \mathcal{F}_{\tau_l}) &\geq \pi(x,y;\Gamma_{\tau_l,\tau_{l+1}}) - 2\Delta_{\tau_l} \\
&\geq \epsilon^2 \cdot p(\tau_l^{(j)} < \tau_l^{(j-1)} \mid \widehat{\mathcal{F}}_{\tau_l}) - 3\Delta_{\tau_l} \\
&\geq \epsilon^2 \cdot p(\tau_l^{(j)} < \tau_l^{(j-1)} \mid \mathcal{F}_{\tau_l}) - 4\Delta_{\tau_l}.
\end{aligned}$$

Since $p\left(\tau_l^{(j)} < \tau_l^{(j-1)} \text{ i.o.} \mid I_j\right) = 1$, whereas $\sum_{l\geq 0} \Delta_l = \epsilon^2$, it follows that $\mathbf{P}_{x,y}(\theta < \infty \mid I_j) = 1$, proving (34).

Next we prove that for every $L \geq 0$

$$\mathbf{E}_{x,y}[R^1_{Q,\theta} 1_{\{\theta<\tau_L\}}] \geq (a_1 - \epsilon)\mathbf{P}_{x,y}(\theta < \tau_L) - \epsilon. \tag{35}$$

Indeed, by (29), Proposition 6.4(2a), and (28)

$$\begin{aligned}
\mathbf{E}_{x,y}[R^1_{Q_l,\theta_l} 1_{\{\theta_l<\tau_{l+1}\}} \mid \mathcal{F}_{\tau_l}] &\geq \rho^1(x,y;\Gamma_{\tau_l,\tau_{l+1}}) - 2\Delta_{\tau_l} \\
&\geq (a_1 - \epsilon) \cdot \pi(x,y;\Gamma_{\tau_l,\tau_{l+1}}) - 2\Delta_{\tau_l} \\
&\geq (a_1 - \epsilon) \cdot \mathbf{P}_{x,y}(\theta_l < \tau_{l+1} \mid \mathcal{F}_{\tau_l}) - 4\Delta_{\tau_l}.
\end{aligned} \tag{36}$$

Since $\{\tau_l \leq \theta\} \in \mathcal{F}_{\tau_l}$ it follows from (36) that

$$\mathbf{E}_{x,y}[R^1_{Q_l,\theta_l} 1_{\{\tau_l\leq\theta<\tau_{l+1}\}}] \geq (a_1 - \epsilon) \cdot \mathbf{P}_{x,y}(\tau_l \leq \theta < \tau_{l+1}) - 4\Delta_{\tau_l}. \tag{37}$$

One obtains (35) by summing (37) over $0 \leq l \leq L$. In particular, it follows from (35), (34) and (31) that $\gamma^1(x,y) \geq a_1 - 3\epsilon$.

We now prove that for every strategy x' of player 1 and every $L \geq 0$

$$\mathbf{E}_{x',y}[R^1_{Q,\theta} 1_{\{\theta<\tau_L\}}] \leq (b_1 + 9\epsilon)\mathbf{P}_{x',y}(\theta < \tau_L) + 2\epsilon. \tag{38}$$

Indeed, let $0 \leq l < L$. If $\mathbf{P}_{0,y}(\theta_l < \tau_{l+1} \mid \widehat{\mathcal{F}}_{\tau_l}) > \frac{\Delta_{\tau_l}}{\epsilon}$ then by Corollary 7.5 and Proposition 6.4(2b), for every $\widehat{\mathcal{F}}$-strategy x',

$$\gamma^1(x',y;\widehat{\mathcal{F}},\tau_l,\tau_{l+1}) \leq \gamma^1(x',y;\Gamma_{\tau_l,\tau_{l+1}}) + 2\epsilon \leq b_1 + 9\epsilon. \tag{39}$$

By Lemma 7.8 it follows that in this case, for every $\mathcal{F}$-strategy x',

$$\mathbf{E}_{x',y}(R^1_{Q_l,\theta_l} 1_{\{\theta_l<\tau_{l+1}\}} \mid \mathcal{F}_{\tau_l}) \leq (b_1 + 9\epsilon)\mathbf{P}_{x',y}(\theta_l < \tau_{l+1} \mid \mathcal{F}_{\tau_l}) + \Delta_{\tau_l}. \tag{40}$$

If, on the other hand, $\mathbf{P}_{0,y}(\theta_l < \tau_{l+1} \mid \widehat{\mathcal{F}}_{\tau_l}) < \frac{\Delta_{\tau_l}}{\epsilon}$, then one has

$$\begin{aligned}
\mathbf{E}_{x',y}(R^1_{Q_l,\theta_l} 1_{\{\theta_l<\tau_{l+1}\}} \mid \mathcal{F}_{\tau_l}) &\leq R^1 \mathbf{P}_{x',y}(\theta_l < \tau_{l+1} \mid \mathcal{F}_{\tau_l}) + 2\frac{\Delta_{\tau_l}}{\epsilon} \\
&\leq (b_1 + \epsilon)\mathbf{P}_{x',y}(\theta_l < \tau_{l+1} \mid \mathcal{F}_{\tau_l}) + 2\frac{\Delta_{\tau_l}}{\epsilon}.
\end{aligned} \tag{41}$$

Eq. (38) follows by summing (40) and (41) over $l = 0, \ldots, L-1$, and taking expectation.

In particular, it follows from (38) that for every strategy x' of player 1 such that $\mathbf{P}_{x',y}(\theta < \infty) = 1$, $\gamma(x', y) \leq b_1 + 11\epsilon$. Thus, player 1 cannot profit much by deviating with a strategy that eventually stops. If $R^1 < 0$ it may still be the case that he can profit by never stopping (see Example 7.2). To overcome this difficulty we add a punishment strategy to y. Namely, we augment y by the following construction. Let $L \in \mathbf{N}$ be sufficiently large so that $\mathbf{P}_{x,y}(\theta < \tau_L) > 1-2\epsilon$. Let y^* be the strategy that follows y up to stage L, and from that stage on stops at each stage n with probability $\epsilon \cdot 1_{\{R^2_{\{2\},n} = R^2\}}$. That is, player 2 stops with small probability whenever $R^2_{\{2\},n} = R^2$.

Since $R^2_{\{2\},n} = R^2$ infinitely often, $\mathbf{P}_{0,y^*}(\theta < \infty) = 1$. Since $\mathbf{P}_{x,y}(\theta < \tau_L) > 1-2\epsilon$, $|\gamma^2(x, y^*) - \gamma^2(x, y)| \leq 4\epsilon$. By (38), **A.6**, and since $b_1 \geq R^1 - \epsilon$, one has for every x'

$$\gamma^1(x', y^*) \leq \mathbf{E}_{x',y}[R^1_{Q,\theta} 1_{\{\theta < \tau_L\}}] + (R^1 + 2\epsilon)\mathbf{P}_{x',y}(\theta \geq \tau_L) \leq b_1 + 11\epsilon.$$

We augment x in a similar fashion to obtain a strategy x^* of player 1. The pair (x^*, y^*) is then a 19ϵ-equilibrium.

7.12 The Game Restricted to $\bar{G}_v$

We here consider the game restricted to $\bar{G}_v$, for some $v = 1, \ldots, V$. Denote $g_v = [a_1, b_1] \times [a_2, b_2]$. Let (x, y) be the concatenation of the strategies $(x^{(0)}_{\Gamma_{\tau_l,\tau_{l+1}}}, y^{(0)}_{\Gamma_{\tau_l,\tau_{l+1}}})$.

We first claim that

$$\begin{array}{l} \text{If } b_1 \leq R^1 - \epsilon \text{ then } \mathbf{P}_{0,y}(\theta < \infty \mid G_v) = 1. \\ \text{If } b_2 \leq R^2 - \epsilon \text{ then } \mathbf{P}_{x,0}(\theta < \infty \mid G_v) = 1. \end{array} \tag{42}$$

We prove the first inequality. By (28), Lemma 6.19, Lemma 7.4(d), and (28) again, for every $l \geq 0$,

$$\mathbf{P}_{0,y}(\theta_l < \tau_{l+1} \mid \mathcal{F}_{\tau_l}) \geq \frac{\epsilon}{4}\mathbf{P}_{0,y}(\cup_{n=\tau_l}^{\tau_l^{(J)}}\{R^1_{\{1\},n} = R^1\} \mid \mathcal{F}_{\tau_l}) - 4\Delta_{\tau_l}.$$

On G_v, $R^1_{\{1\},n} = R^1$ infinitely often, whereas only finitely many times $\tau_l^{(J)} < \tau_{l+1}$. Therefore $\mathbf{P}_{0,y}(\theta < \infty \mid G_v) = 1$, proving (42).

We now claim that if $b_1 \leq R^1 - \varepsilon$ then $\chi^2(\Gamma_{\tau_l,\tau_{l+1}}) = 1$ for every $l \geq 0$ on G_v. Indeed, on G_v one has $c^*_{\tau_0,\tau_1} = c^*_{\tau_l,\tau_{l+1}}$ for every $l \geq 0$. Hence, if the claim does not hold then $\chi^2(\Gamma_{\tau_l,\tau_{l+1}}) = 0$ for every $l \geq 0$ on G_v. By definition this implies that $\pi(0, y^{(0)}_{\Gamma_{\tau_l,\tau_{l+1}}}; \Gamma_{\tau_l,\tau_{l+1}}) < \frac{\Delta_{\tau_l}}{\varepsilon} + \Delta_{\tau_l}$. Hence $\mathbf{P}_{0,y}(\theta < \infty \mid G_v) < 1$, a contradiction to (42). Thus, by Lemma 7.4, we get the following.

$$\text{If } b_1 \leq R^1 - \epsilon \text{ then } \mathbf{P}_{0,y}(\theta_l < \tau_l^{(J)} \mid \widehat{\mathcal{F}}_{\tau_l}) \geq \frac{\Delta_{\tau_l}}{\epsilon} \text{ on } G_v. \tag{43}$$

Next we prove that for every $L \geq 0$,

$$\mathbf{E}_{x,y}[R^1_{Q,\theta} 1_{\{\theta < \tau_L\}}] \geq a_1 \cdot \mathbf{P}_{x,y}(\{\theta < \tau_L\}) - 5\epsilon. \tag{44}$$

Indeed, by (29), the construction of (x, y), and (28)

$$\mathbf{E}_{x,y}(R^1_{Q_l,\theta_l} 1_{\{\theta_l < \tau_l^{(J)}\} \cap \{c_{\tau_l,\tau_{l+1}} = g_v\}} \mid \mathcal{F}_{\tau_l}) \\ \geq a_1 \cdot \mathbf{P}_{x,y}(\{\theta_l < \tau_l^{(J)}\} \cap \{c_{\tau_l,\tau_{l+1}} = g_v\} \mid \mathcal{F}_{\tau_l}) - 4\Delta_{\tau_l}. \tag{45}$$

Summing (45) over $l = 0, ..., L-1$, and taking expectation, we get:

$$\mathbf{E}_{x,y}\left[R^1_{Q,\theta} 1_{\bigcup_{0 \leq l < L}(\{\theta_l < \tau_l^{(J)}\} \cap \{c_{\tau_l,\tau_{l+1}} = g_v\})}\right] \\ \geq a_1 \cdot \mathbf{P}_{x,y}\left(\bigcup_{0 \leq l < L} (\{\theta_l < \tau_l^{(J)}\} \cap \{c_{\tau_l,\tau_{l+1}} = g_v\})\right) - \epsilon. \tag{46}$$

Let $G^*_v = \bar{G}_v \cap \bigcup_{0 \leq l < L}(\{\tau_l^{(J)} < \tau_{l+1}\} \cup \{c_{\tau_l,\tau_{l+1}} \neq g_v\})$. From (32) and (33) it follows that:

$$p(G^*_v) < 2\epsilon. \tag{47}$$

Since $\{\theta < \tau_L\} \subseteq G^*_v \cup \bigcup_{0 \leq l < L}(\{\theta_l < \tau_l^{(J)}\} \cap \{c_{\tau_l,\tau_{l+1}} = g_v\})$, (44) follows from (47) and (46).

Next we claim that for every $L \geq 0$, and every strategy x' of player 1,

$$\mathbf{E}_{x',y}[R^1_{Q,\theta} 1_{\{\theta < \tau_L\}}] \leq (b_1 + 2\epsilon) \cdot \mathbf{P}_{x',y}(\{\theta < \tau_L\}) + 6\epsilon. \tag{48}$$

Indeed, the same argument used to prove Eq. (38) proves, using Corollary 7.5, the definition of (x, y), (43) and Lemma 7.8, that

$$\mathbf{E}_{x',y}[R^1_{Q,\theta} 1_{\bigcup_{0 \leq l < L}(\{\theta_l < \tau_l^{(J)}\} \cap \{c_{\tau_l,\tau_{l+1}} = g_v\})}] \\ \leq (b_1 + 2\epsilon) \cdot \mathbf{P}_{x',y}\left(\bigcup_{0 \leq l < L} (\{\theta_l < \tau_l^{(J)}\} \cap \{c_{\tau_l,\tau_{l+1}} = g_v\})\right) + 2\epsilon.$$

Eq. (48) follows using (47).

In particular, it follows from (38) that for every strategy x' of player 1 such that $\mathbf{P}_{x',y}(\theta < \infty) = 1$, $\gamma(x', y) \leq b_1 + 8\epsilon$. Thus, player 1 cannot profit much by deviating with a strategy that eventually stops. If $b_1 < R^1 - \epsilon$ then by (42) and (32) $\mathbf{P}_{x',y}(\theta < \infty) \geq \mathbf{P}_{0,y}(\theta < \infty) \geq 1 - \epsilon$ for every x'. If $b_1 \geq R_1 - \epsilon$ one should augment y by adding a punishment strategy as in section 7.11.

8 More than Two Players

When there are more than two players, it is no longer true that the game on a tree admits a stationary ϵ-equilibrium. An example of a three-player game where this phenomenon happens was first found by Flesch et al (1997). Nevertheless, a consequence of Solan (1999) is that any three-player game on a tree admits a periodic ϵ-equilibrium, but the period may be long. We do not know whether one can use this result to generalize Proposition 6.4 for three-player games.

When there are at least four players, existence of ϵ-equilibria in stopping games on finite trees is still an open problem, even in the deterministic case; that is, when every node in the tree has at most a single child. For more details the reader is referred to Solan and Vieille (2001).

References

[1] Alario-Nazaret M., Lepeltier J.P. and Marchal B. (1982) Dynkin Games, *Stochastic Differential Systems (Bad Honnef)*, 23-32, *Lecture notes in Control and Information Sciences*, **43**, Springer Verlag

[2] Bismut J.M. (1977) Sur un Problème de Dynkin. *Z. Warsch. V. Geb.*, **39**,31-53

[3] Cvitanić J. and Karatzas I. (1996) Backward Stochastic Differential Equations with Reflection and Dynkin Games. *Ann. Prob.*, **24**, 2024-2056

[4] Dynkin E.B. (1969) Game Variant of a Problem on Optimal Stopping, *Soviet Math. Dokl.*, **10**, 270-274

[5] Flesch J., Thuijsman F., and Vrieze O.J. (1996) Recursive Repeated Games with Absorbing States, *Math. Oper. Res.*, **21**, 1016-1022

[6] Flesch J., Thuijsman F. and Vrieze O.J. (1997) Cyclic Markov Equilibria in Stochastic Games, *Int. J. Game Th.*, **26**, 303-314

[7] Lepeltier J.P. and Maingueneau M.A. (1984) Le Jeu de Dynkin en Thèorie Gènèrale sans l'Hypothèse de Mokobodski, *Stochastics*, **13**, 25-44.

[8] Mamer J.W. (1987) Monotone Stopping Games, *J. Appl. Prob.*, **24**, 386-401

[9] Morimoto H. (1986) Non-Zero-Sum Discrete Parameter Stochastic Games with Stopping Times, *Prob. Th. Related Fields*, **72**, 155-160

[10] Neveu J. (1975) Discrete-Parameter Martingales, North-Holland, Amsterdam

[11] Ohtsubo Y. (1987) A Nonzero-Sum Extension of Dynkin's Stopping Problem, *Math. Oper. Res.*, **12**, 277-296

[12] Ohtsubo Y. (1991) On a Discrete-Time Non-Zero-Sum Dynkin Problem with Monotonicity, *J. Appl. Prob.*, **28**, 466-472

[13] Ramsey F.P. (1930) On a Problem of Formal Logic, *Proc. London Math. Society*, **30**, 264-286

[14] Rosenberg D., Solan E. and Vieille N. (2001) Stopping Games with Randomized Strategies, *Prob. Th. Related Fields*, **119**, 433-451

[15] Shapley L.S. (1953) Stochastic Games, *Proc. Nat. Acad. Sci. U.S.A.*, **39**, 1095-1100

[16] Shmaya E., Solan E. and Vieille N. (2002) An Application of Ramsey Theorem to Stopping Games, *Games Econ. Behavior*, forthcoming.

[17] Solan E. (1999) Three-Person Absorbing Games, *Math. Oper. Res.*, **24**, 669-698

[18] Solan E. and Vieille N. (2001), Quitting Games, *Math. Oper. Res.*, **26**, 265-285

[19] Touzi N. and Vieille N. (2002) Continuous-Time Dynkin Games with Mixed Strategies. *SIAM J. Cont. Opt.*, forthcoming.

[20] Vrieze O.J. and Thuijsman F. (1989) On Equilibria in Repeated Games with Absorbing States, *Int. J. Game Th.*, **18**, 293-310