Perfect Correlated Equilibria in Stopping Games

Yuval Heller

School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv 69978, Israel.
Phone: 972-3-640-5386. Fax: 972-3-640-9357.
Email: helleryu@post.tau.ac.il

Abstract

We prove that every undiscounted multi-player stopping game in discrete time admits an approximate correlated equilibrium. Moreover, the equilibrium has three appealing properties: "trembling-hand" perfectness - players do not use non-credible threats; normal-form correlation - communication is required only before the game starts; uniformness - it is an approximate equilibrium in any long enough finite-horizon game and in any discounted game with high enough discount factor.

1 Introduction

Stopping games have been introduced by Dynkin ([6]) as a generalization of optimal stopping problems, and later used in several models in economics, management science and biology, such as job search, research and development (see e.g., Fudenberg and Tirole [9] and Mamer [12]), the analysis of strategic exit (see e.g., Fudenberg and Tirole [10], Ghemawat and Nalebuff [11]), and the war of attrition (see e.g., Nalebuff and Riley [17]).

In this paper we focus on (undiscounted) multi-player stopping games in discrete time. The game is played by a finite set of players. There is an unknown state variable, on which players receive a symmetric partial information along the game. At stage 1 all the players are active. At every stage n, each active player declares, independently of the others, whether he stops or continues. A player that stops at stage n, becomes passive for the rest of the game. The

[1] This work is in partial fulfillment of the requirements for the Ph.D. in mathematics at Tel-Aviv University. I would like to thank Eilon Solan for his careful supervision, for the continuous help he offered, and for many insightful discussions.

payoff of a player depends on the history of players' actions while he has been active and on the state variable.

Much work has been devoted to the study of 2-player zero-sum stopping games in discrete time. Dynkin ([6]) proved that this game has a value under an assumption that at any stage only one of the two players is allowed to stop, and Neveu ([19]) proved the existence of the value under an assumption that at each stage each player prefers the other player to be the stopping player. Rosenberg, Solan and Vieille ([24]) allowed the players to use randomized strategies, and proved that the game has a value, assuming only integrability of the payoffs.

The 2-player nonzero-sum problem in discrete time when the payoffs have a special structure was studied, among others, by Mamer ([12]), Morimoto ([14]), Ohtsubo ([21,22]), Nowak and Szajowski ([20]) and Neumann, Ramsey and Szajowski ([18]) and the references therein. Those authors provided various sufficient conditions under which (Nash) ϵ-equilibria exist. Recently, Shmaya and Solan ([27]) have proved the existence of (Nash) ϵ-equilibria assuming only integrability of the payoffs. In contrast with the 2-player case, there is no existence result for ϵ-equilibria in multiplayer stopping games.

The equilibrium path of Nash equilibrium may be sustained by "non-credible" threats of punishment. Since by punishing a deviator, some of the punishing players may receive low payoff (lower than if they do not punish the deviator), it is not clear whether one should expect players to follow such an equilibrium. Thus, a few papers study the stronger concept of perfect equilibrium (Selten [25,26]) in 2-player stopping games, such as: Fine and Li ([7]), Dutta ([5]) and Mashiah ([13]).

Aumann ([1]) defined the concept of correlated equilibrium: a correlated equilibrium in a finite normal-form game is a Nash equilibrium in an extended game that includes a correlation device, which sends to each player, before the start of play, a private signal; the strategy of each player can then depend on the private signal that he received. Correlated equilibria have a number of appealing properties. They are computationally tractable. Existence is verified by checking a system of linear inequalities rather than a fixed point. The set of correlated equilibria is closed and convex. Aumann ([2]) argues that it is the solution concept consistent with the Bayesian perspective on decision making.

For sequential games, two main versions of correlated equilibrium have been studied (see e.g., [8]): normal-form correlated equilibrium, in which each player receives only private signal before the game starts, and extensive-form correlated equilibrium, in which each player receives a private signal at each stage of the game. Note that every normal-form correlated equilibrium is an extensive-form correlated equilibrium, but the converse is not true.

Communication between the players, that can lead to correlation of strategies, is natural in many setups, for example: countries negotiate about their actions to each other and to other countries; firms decide on their strategies based on common information such as past behavior of the market; and a manager coordinates the actions taken by his subordinates. In some situations players may coordinate before the play starts, but coordination along the play is costly or impossible, and only the notion of normal-form correlated equilibrium is appropriate. Two examples of such situations are:

- War of attrition in nature, which is commonly modeled as a stopping game (see e.g., [17]), where normal-form (but not extensive-form) correlation devices are implemented by evolution of phenotype roles (see e.g., Shmida and Peleg [28]).
- Brokers of a certain firm who act in different stock-exchange markets. The brokers can correlate their moves before the commerce begins, but due to the need to make many actions in a short period of time, the ability to communicate during the commerce is limited.

A few papers have defined and studied the properties of perfect correlated equilibria in finite games, see e.g., Myerson ([15,16]) and Dhillon and Mertens ([4]). Generalizing the definition of the last paper, we define a "trembling-hand" perfect correlated (δ, ϵ)-equilibrium, where $\delta > 0$ is an upper bound for the probabilities of the following: an event E in the probability space that determines the state variable, and an event in which the correlation device sends signals in some set M'. The parameter $\epsilon > 0$ is the maximal profit a player can earn by deviating at any stage of the game and after any history of play, conditioned on that the state variable is not in E and the signal profile is not in M'. We hope that the definition of an approximate perfect correlated equilibrium may be useful in future study of other dynamic games.

Our main result shows that for every $\delta, \epsilon > 0$, a multi-player stopping game admits a normal-form uniform perfect correlated (δ, ϵ)-equilibrium. This implies the existence of a uniform perfect correlated equilibrium payoff. Due to the uniformness property, the (δ, ϵ)-equilibrium is also an approximate equilibrium in any long enough finite-horizon stopping game and in any discounted stopping game with high enough discount factor. Arguments in favor of the notion of uniform equilibrium can be found in Aumann and Mashcler ([3]).

The proof relies on two reductions: we first define terminating games, as stopping games that immediately end as soon as any player stops, and reduce the problem of existence of equilibrium from general stopping games to terminating games. [2] This reduction requires us to use correlation devices that are

[2] In other papers, both games (terminating and stopping) are referred to as stopping

"universal" (depend only on ϵ and the number of players) and "unrevealing" (The expected payoff of a player almost does not change when he receives his signal). Next, we use a stochastic variation of Ramsey's theorem ([27]) to further reduce the problem to that of studying the properties of correlated ϵ-equilibria in multiplayer absorbing games [3], by adapting the methods of Solan and Vohra [31] who prove that any multiplayer absorbing game admits a correlated ϵ-equilibrium.

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 make reductions from existence of equilibria in general stopping games to existence of equilibria in terminating games with special properties. In Section 5 we define the notion of games played on a finite tree and study some of their properties. In Section 6 we use the stochastic variation of Ramsey's theorem, which allows us to construct a perfect correlated (δ, ϵ)-equilibrium in Section 7.

2 Model and Main Result

Definition 1 *A* (multi-player) stopping game *(in discrete time) is a 6-tuple* $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ *where:*

- I is a finite set of players;
- $(\Omega, \mathcal{A}, p)$ is a probability space (the state space);
- $\mathcal{F} = (\mathcal{F}_n)_{n\geq 0}$ is a filtration over $(\Omega, \mathcal{A}, p)$;
- $R = (R_n)_{n\geq 0} \bigcup R_\infty$ is an $\mathcal{F}$-adapted process:
  - · Let H_n^S denote the set of all possible histories of realized actions (stop or continue) before stage n in which the members of S always continue. The coordinates of R_n are denoted by $R^i_{S,n,h_n^S,}$ where $n \in \mathbf{N}$, $i \in S \subseteq I$ is the set of players that stop at stage n and $h_n^S \in H_n^S$ is the history of realized actions of each player before stage n.
  - · Let H_∞^S denote the set of all possible infinite histories of realized actions in which the members of S always continue and all the members of $I \backslash S$ have stopped. Given $h_\infty^S \in H_\infty^S$, let $n_{h_\infty^S}$ be the last stage in which a player stops in h_∞^S. The coordinates of R_∞ are denoted by $R^i_{S,\infty,h_\infty^S}$ where $S \subseteq I$ is the set of players who have never stopped in the entire game, and $h_\infty^S \in H_\infty^S$ is an history of realized actions in which all the players in $I \backslash S$ (and only them) have stopped. we require that $R^i_{S,\infty,h_\infty^S}$ is measurable in $\mathcal{F}_{n_{h_\infty^S}}$. The

games. We have chosen to call them by a different name, as the reduction of existence of correlated equilibrium from stopping games to terminating games is not trivial due to our requirement of normal-form correlation.

[3] An absorbing game is a stochastic game with a single non-absorbing state.

last requirement means that the payoff of a player who never stops changes only when other players stop.

A stopping game is played as follows. At stage 1 all the players are active. At each stage n, each active player is informed about $F_n(\omega)$, the minimal set in $\mathcal{F}_n$ that includes the state $\omega \in \Omega$, and declares, independently of the others, whether he stops or continues. An active player i that stops, becomes passive for the rest of the game, and his payoff is given by R^i_{S,n,h^S_n}, where $i \in S \subseteq I$ is the set of active players who stop at stage n, and $h^S_n \in H^S_n$ is the history of realized actions until stage n. If a player i never stops, his payoff is $R^i_{S,\infty,h^S_\infty}$ where $i \in S \subseteq I$ is the set of players who never stop, and h^S_∞ is the infinite realized history of actions.

Definition 2 A *(normal-form) correlation device* is a pair $\mathcal{D} = (M, \mu)$ where:

(1) $M = (M^i)_{i \in I}$, where M^i is a finite space of signals the device can send player i.
(2) $\mu \in \triangle(M)$ is the probability distribution according to which the device sends the signals to the players before the stopping game starts.

Given a correlation device $\mathcal{D}$, we define an extended game $G(\mathcal{D})$. The game $G(\mathcal{D})$ is played exactly as the game G, except that before the game starts, a signal combination $m = (m^i)_{i \in I}$ is drawn according to μ, and each player is informed of m^i. Then, each player may base his strategy on his signal. When $|M| = 1$ we say that $\mathcal{D}$ is trivial, and in that case $G(\mathcal{D})$ is equivalent to G.

For simplicity of notation, let the singleton coalition $\{i\}$ be denoted as i, and let $-i = \{I \backslash i\}$ denote the coalition of all the players besides player i. A (behavioral) *strategy* for player $i \in I$ in $G(\mathcal{D})$ is an $\mathcal{F}$-adapted process $x^i = (x^i_n)_{n \geq 0}$, where $x^i_n : (\Omega \times M^i \times H^i_n) \to [0,1]$. The interpretation is that $x^i_n(\omega, m^i, h^i_n)$ is the probability by which an active player i stops at stage n after an history of play h^i_n.

A (behavioral) strategy profile $x = (x^i)_{i \in I}$ is *completely mixed* if at each stage, given any history of play, each player has a positive probability to stop and a positive probability to continue. Formally: for each player $i \in I$, message $m^i \in M^i$, stage $n \in \mathbf{N}$, and history $h^i_n \in H^i_n$: $0 < x^i_n(\omega, m^i, h^i_n) < 1$

Let θ_i be the stage in which player i stops and let $\theta_i = \infty$ if player i never stops. If $\theta < \infty$ let $i \in S \subseteq I$ be the coalition that stops at stage θ, and if $\theta = \infty$ let $i \in S \subseteq I$ be the coalition that never stop in the game. Let h^i_θ the history of realized actions until stage θ_i. The payoff to player i is $R^i_{S,\theta_i,h^S_{\theta_i}}$. The expected payoff under the strategy profile $x = (x^i)_{i \in I}$ is given by: $\gamma^i(x) = \mathbf{E}_x\left(R^i_{S,\theta_i,h^S_{\theta_i}}\right)$ where the expectation $\mathbf{E}_x$ is with respect to (w.r.t.) the distribution $\mathbf{P}_x$ over plays induced by x. Given an event $E \subseteq \Omega$, the expected payoff conditioned on $\Omega \backslash E$ is: $\gamma^i\left(x \mid (\Omega \backslash E)\right) = \mathbf{E}_x\left(R^i_{S,\theta_i,h^S_{\theta_i}} \mid (\Omega \backslash E)\right)$.

The strategy x^i is ϵ-best reply (ϵ-best reply conditioned on $\Omega \backslash E$) for player i when all his opponents follow x^{-i} if for every strategy of player i, y^i: $\gamma^i(x) \geq \gamma^i(x^{-i}, y^i) - \epsilon$ ($\gamma^{\mathbf{i}}(x \mid (\Omega \backslash E)) \geq \gamma^{\mathbf{i}}(x^{-i}, y^i \mid (\Omega \backslash E)) - \epsilon$). Note that when x^i is ϵ-best reply conditioned on $\Omega \backslash E$, player i assumes in his evaluation of his expected payoff that $\omega \in \Omega \backslash E$. Let $H_n = \bigcup_{i \in I} H^i_n$ denote the set of all histories of realized actions before stage n, and Let $\hat{\mathcal{F}}_n \subseteq \mathcal{F}_n$ the set of all minimal sets in $\mathcal{F}_n$:

$$\hat{\mathcal{F}}_n = \left\{F_n \in \mathcal{F}_n | \neg\exists\emptyset \neq \hat{F}_n \in \mathcal{F}_n,\ s.t.\ \hat{F}_n \subsetneq F_n\right\}$$

Let $G(h_n, F_n, \mathcal{D}, m)$ be the induced stopping game that begins at stage n after a signal m^i has been sent to each player i, an history of play h_n has been played, and the players are informed that $\omega \in F_n \subseteq \hat{\mathcal{F}}_n$. The active players when the game $G(h_n, F_n, \mathcal{D}, m)$ starts, are those who have not stopped in h_n. For simplicity of notation, we denote by x also the induced strategy profile in $G(h_n, F_n, \mathcal{D}, m)$. We now define a perfect correlated (δ, ϵ)-equilibrium, generalizing the definition of perfect correlated equilibria in finite games ([4]).

Definition 3 Let $G(\mathcal{D})$ be a stopping game, let $E \subseteq \Omega$ be an event, let $M' \subseteq M$ be a set of signal profiles of the correlation device, and let $\epsilon > 0$. A strategy profile $x = (x^i)_{i \in I}$ is a *perfect ϵ-equilibrium of $G(\mathcal{D})$ conditioned on $\Omega \backslash E$ and given $M \backslash M'$, if* there exists a sequence $(y_k)_{k \in \mathbf{N}} = (y^i_k)_{k \in \mathbf{N}, i \in I}$ of completely mixed strategy profiles in $G(\mathcal{D})$, and a sequence $(\epsilon_k)_{k \in \mathbf{N}}$ $(0 < \epsilon_k < 1)$ converging to 0, such that for all $i \in I$, $m \in M$, $n \in \mathbf{N}$, $h^i_n \in H^i_n$, $F_n \in \hat{\mathcal{F}}_n$ satisfying $F_n \notin E$, x^i is ϵ-best reply for player $i \in I$ in the induced game $G(h_n, F_n, \mathcal{D}, m)$ conditioned on $\Omega \backslash E$, when all his opponents $j \in -i$ use $(1 - \epsilon_k) x^j + \epsilon_k y^j_k$.

Definition 4 Let $G(\mathcal{D})$ be a stopping game and let $\delta, \epsilon > 0$. A strategy profile $x = (x^i)_{i \in I}$ is a *perfect* (δ, ϵ)-equilibrium *of* $G(\mathcal{D})$ if there exists an event $E \subseteq \Omega$ and a set of signal profiles $M' \subseteq M$, such that $p(E) < \delta$, $\mu(M') < \delta$, and x is a perfect ϵ-equilibrium of $G(\mathcal{D})$ conditioned on $\Omega \backslash E$ and given $M \backslash M'$.

Definition 5 Let G be a stopping game and let $\delta, \epsilon > 0$. A *perfect correlated (δ, ϵ)-equilibrium* is a pair $(\mathcal{D}, x)$ where $\mathcal{D}$ is a correlation device and x is a *perfect (δ, ϵ)-equilibrium* in the extended game $G(\mathcal{D})$.

Our main Result is the following:

Theorem 6 Let $\delta, \epsilon > 0$ and let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a multi-player stopping game such that $\sup_{n \in \mathbf{N} \bigcup \infty} \|R_n\|_\infty \in L^1(p)$. *Then for every $\delta, \epsilon > 0$,*

G has a prefect correlated (δ, ϵ)-equilibrium. Moreover, the correlation device $\mathcal{D} = \mathcal{D}(\epsilon)$ is universal: it depends only on ϵ and $|I|$.

Remark 7 The perfect correlated (ϵ, δ)-equilibrium that we construct is uniform in a strong sense: it is a $(\delta, 3\epsilon)$-equilibrium in every finite n-stage game, provided that n is sufficiently large. This can be seen by the construction itself (Prop. 30) or by applying a general observation made by [29, Prop. 2.13].

Definition 8 A payoff vector $r \in R^{|I|}$is a (uniform) *perfect correlated payoff* if for every $\epsilon, \delta, \epsilon' > 0$ there is a perfect correlated (ϵ, δ)-equilibrium x with a payoff $r - \epsilon' \leq \gamma(x) \leq r + \epsilon'$.

Corollary 9 let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a multi-player stopping game such that $\sup\limits_{n \geq 0} \|R_n\|_\infty \in L^1(p)$. Then G admits a (uniform) perfect correlated payoff.

3 Sketch of the Proof

In this section we provide the main ideas of the proof. Let a *terminating game* be a stopping game in which as soon as any player stops, the payoffs to all the players are determined. Let G be a terminating game. To simplify the presentation, assume that $\mathcal{F}_n$ is trivial for every n, so that the payoff process is deterministic, and that payoffs are uniformly bounded by 1. For every two natural numbers $k < l$, define the periodic game $G(k, l)$ to be 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 terminal payoff at stage $k + (n \bmod l - k)$ in G.

This periodic game is equivalent to an absorbing game, where each round of T corresponds to a single stage of the absorbing game (a stochastic game with a single non-absorbing state). Moreover, $G(k, l)$ has two special properties: It is recursive (the payoff in the non-absorbing state is 0), and there is a single action profile with a zero absorbing probability.

Solan and Vohra ([31]) proved a classification result for absorbing games (Prop. 4.10). Applying it to the two special properties of $G(k, l)$ yields that $G(k, l)$ has one of the following: (1) A perfect stationary absorbing equilibrium. (2) A perfect stationary non-absorbing equilibrium. (3) A correlated distribution η over the set of action profiles in which a single player stops. The special properties of η allows the construction of a perfect correlated ϵ-equilibrium in $G(k, l)$.

Assign to each pair of non-negative integers $k < l$ an element from a finite set of *colors* $c(k, l)$ that denotes which case of the classification result holds and the (ϵ)-approximation of the equilibrium payoff.[4] A consequence of Ramsey's theorem ([23]) is that there is an increasing sequence of integers $0 \le k_1 < k_2 < ...$ such that $c(k_1, k_2) = c(k_j, k_{j+1})$ for every j.

Assume first that $k_1 = 0$. This allows to construct a perfect correlated ϵ-equilibrium for G. The constriction depends on the case indicated by $c(k_1, k_2)$. If the case is 1 or 2, then between stages k_j and k_{j+1} the players follow a periodic (δ, ϵ)-equilibrium in the game $G(k_j, k_{j+1})$ with a payoff in an ϵ neighborhood of the payoff indicated by $c(k_1, k_2)$. For this concatenated strategy to indeed be a 3ϵ-equilibrium in G in case 1, it is needed to verify that the game is absorbed with probability 1. This is done by giving appropriate lower bounds to the stopping probability of each $G(k_j, k_{j+1})$ in the first round. These bounds are adaptations to the multi-player case of the bounds given for 2-player games in Shmaya and Solan ([27]).

If the indicated case $c(k_1, k_2)$ is 3, then we adept the procedure presented by Solan and Vohra ([30, Section 4.2]) to the requirement of perfection and to the use of a universal correlation device. Originally, their procedure allows the construction of a correlated ϵ-equilibrium in *quitting games* - stationary terminating games where the payoff matrix is the same at all stages. As part of the adaptation, we verify that at stage k_1, with high probability the signal a player receives does not affect his expected payoff by more than ϵ.

If $k_1 > 0$, then Between stages 0 and k_1, the players follow an equilibrium in the k_1-stage game with the terminal payoff that is implied by $c(k_1, k_2)$. From stage k_1 and on, the players follow the strategy described above. It is easy to verify that this strategy profile forms a 5ϵ-equilibrium.

We now consider a general stopping game. Assume by induction that any m-player stopping game admits a perfect correlated payoff vector. Given a stopping game G with $m + 1$ players we construct an auxiliary terminating game G' with $m + 1$ players by setting the payoff of a player $i \notin S$ when the non-empty coalition S stops at stage n, as his perfect correlated payoff in the induced $(m + 1 - |S|)$-player game that begins at stage $n + 1$. The perfect correlated (δ, ϵ)-equilibrium in G' implies naturally a perfect correlated (δ, ϵ)-equilibrium in G.

When the payoff process is general, 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. We analyze this kind of periodic games, by adapting the methods presented in [27] for 2-player stopping games, and by using their

[4] If more than one case holds, or there is more than one profile in one of the cases, then we choose arbitrarily according to some lexicographic order. In case 3 the color indicates the ϵ-approximations of two payoff vectors: the payoff under the distribution η, and the maximal payoff of each player when he stops alone.

stochastic version of Ramsey's theorem.

4 Reductions

In this section we make three reductions to the problem of existence of perfect correlated (δ,ϵ)-equilibrium in stopping games:

(1) We reduce the problem to that of existence of perfect correlated (δ,ϵ)-equilibrium in terminating games (Subseq. 4.1).
(2) We further reduce it to the problem of existence of such equilibrium in tree-like terminating games (Subseq. 4.2), by relying on [27, Sec. 6].
(3) We make a last reduction to the problem of existence of such equilibrium in an induced terminating game G_F, deep enough in the original game-tree, where with high probability each approximate matrix payoff occurs infinitely often or does not occur at all (Subseq. 4.4).

4.1 Terminating games

Definition 10 *A* terminating game *is a 6-tuple* $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ *where:*

- I is a finite set of players;
- $(\Omega, \mathcal{A}, p)$ 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 an $\mathcal{F}$-adapted $\mathbf{R}^{|I|\cdot(2^{|I|}-1)}$-valued process. The coordinates of R_n are denoted by $R^i_{S,n}$where $i \in I$ and $\emptyset \neq S \subseteq N$.

A terminating game is played as follows. At each stage $n \in \mathbf{N}$, each player is informed about $F_n(\omega)$, the minimal set in $\mathcal{F}_n$ that includes ω, and declares, independently of the others, whether he stops or continues. If all players continue the game continues to the next stage. If at-least one player stops, say a coalition $S \subseteq I$, the game terminates, and the payoff to player i is $R^i_{S,n}$. If no player ever stops, the payoff to everyone is normalized to zero.

A (behavioral) *strategy* for player $i \in I$ in $G\,(\mathcal{D})$ is an $\mathcal{F}$-adapted process $x^i = (x^i_n)_{n\geq 0}$, where $x^i_n : (\Omega \times M^i) \to [0,1]$. The interpretation is that $x^i_n\,(\omega, m^i)$ is the probability by which player i stops at stage n, provided the game has not stopped before that stage. *A* perfect correlated (δ,ϵ)-equilibrium and a perfect correlated payoff vector are defined in an analog way to Def. 4, 5 and 8.

Proposition 11 Assume that every terminating game with bounded payoffs $(\sup_{n\in\mathbf{N}} \|R_n\|_\infty \in L^1(p))$ admits a perfect correlated (δ,ϵ)-equilibrium for every

$\delta, \epsilon > 0$, and that the correlation device is universal (depends only on ϵ and $|I|$). Let $\delta, \epsilon > 0$ and let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ a stopping game with bounded payoffs. Then G admits a prefect correlated (δ, ϵ)-equilibrium.

PROOF. We prove the proposition by induction on the number of players. Let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a stopping game with $m = |I|$ players. By the induction hypothesis every stopping game with $k < m$ players has a perfect correlated (δ, ϵ)-equilibrium with a universal correlation device $\mathcal{D}_{\epsilon,k}$. For each induced stopping game $G(h_n, F_n, \mathcal{D}_{\epsilon,k})$ with k players, let $x_{h_n,F_n,\mathcal{D}_{\epsilon,k}}$ be a perfect correlated (δ, ϵ)-equilibrium with a payoff of $v_{h_n,F_n,\mathcal{D}_{\epsilon,k}}$. We define an auxiliary terminating game $G' = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R')$, where the payoff process $R' = \left(R'^i_{S,n}\right)_{i \in I.S \subseteq I, n \in \mathbf{N}}$ is defined as follows for each $n \in \mathbf{N}$ and $F_n \in \hat{\mathcal{F}}_n$:

- For each $i \in S \subseteq I$: $R'^i_{S,n}(F_n) = R^i_{S,n,h^I_n} - R^i_{I,\infty,h^I_\infty}$, where h^I_n is the history of realized actions, in which all players continue at all stages before stage n.
- For each $i \notin S \subset I$: $R'^i_{S,n}(F_n) = v^i_{h^{(I\setminus S)*}_{n+1},F_n,\mathcal{D}_{\epsilon,|I\setminus S|}} - R^i_{I,\infty,h^I_\infty}$, where $h^{(I\setminus S)*}_{n+1}$ is the history of realized actions, in which all the players continue at all stages before stage n, and the players in S stop at stage n.

The terminating game G' has a perfect correlated (δ, ϵ)-equilibrium (x', D') according the assumption of Prop. 11. Let $\mathcal{D}_\epsilon = D' \times \prod_{k<m} D_{\epsilon,k}$, and let the strategy x in $G(\mathcal{D})$ be as follows: $x = x'$ as long as no player stops, and $x = x_{h^{S*}_{n+1},F_n,\mathcal{D}_{\epsilon,|I\setminus S|}}$ after a coalition $S \subseteq I$ stops at stage n. The construction of x implies that it is a perfect correlated $\left(2^{|I|} \cdot \delta, \epsilon\right)$-equilibrium in G. QED.

Thus, in the rest of this paper, we focus only on terminating games.

4.2 Tree-like stopping game

Definition 12 A terminating game $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ is *tree-like* if for every $n \in N$, $|\mathcal{F}_n| < \infty$.

Shmaya and Solan prove ([27, Sec. 6]) that any 2-player terminating game can be approximated by a tree-like terminating game such that any approximate equilibrium of the tree-like game is also an approximate equilibrium of the original game. With minor changes, the proof can be adapted for multi-player terminating games, and for perfect correlated equilibria. This implies the following lemma (the proof is omitted):

Lemma 13 Assume that any tree-like terminating game with bounded payoffs admits a perfect correlated (δ, ϵ)-equilibrium for every $\delta, \epsilon > 0$. Let $\epsilon, \epsilon > 0$

and let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ a terminating game with bounded payoffs . Then G admits a perfect correlated (δ, ϵ)-equilibrium.

Thus, in the rest of this paper, we assume without loss of generality (w.l.o.g.) that the terminating game is tree-like.

4.3 Preliminaries

The definitions imply that for every two payoff processes R and $\tilde{R}$ such that $\mathbf{E}\left(\sup_{n\geq 0}\left\|R_n - \tilde{R}_n\right\|_\infty\right) < \epsilon$, every perfect correlated (δ, ϵ)-equilibrium in the terminating game $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ is a $(\delta, 3\epsilon)$-equilibrium in the terminating game $\tilde{G} = \left(I, \Omega, \mathcal{A}, p, \mathcal{F}, \tilde{R}\right)$. Hence we can assume w.l.o.g. that the payoff process R is uniformly bounded and that its range is finite. Actually, we assume that for some $K \in \mathbf{N}$, $R^i_{S,n} \in \left\{0, \pm\frac{1}{K}, \pm\frac{2}{K}, ..., \pm\frac{K}{K}\right\}$ for every $n \in \mathbf{N}$. Let $D = \prod_{i\in I, \emptyset\neq S\subseteq I}\left\{0, \pm\frac{1}{K}, \pm\frac{2}{K}, ..., \pm\frac{K}{K}\right\}$ be the set of all possible one-stage payoff matrices of the terminating game G. Let $R_n(\omega)$ be the payoff matrix at stage n. Let $\tau : \Omega \to \mathbf{N}$ a bounded terminating time. The partition $\hat{\mathcal{F}}_\tau$ is:

$$\hat{\mathcal{F}}_\tau = \bigcup_{n\in\mathbf{N}}\left\{F_n \in \hat{\mathcal{F}}_n | \exists\omega,\ s.t.\ F_n(\omega) = F_n, \tau(\omega) = n\right\}$$

Given any payoff matrix $d \in D$, let $A_d \subseteq \bigvee_{n\in\mathbf{N}}\mathcal{F}_n$ be the event that d occurs infinitely often: $A_d = \{\omega \in \Omega | i.o.\ R_n(\omega) = d\}$, and let $B_{d,k} \subseteq \bigvee_{n\in\mathbf{N}}\mathcal{F}_n$ be the event that d never occurs after stage k: $B_{d,k} = \{\omega \in \Omega | \forall n \geq k,\ R_n(\omega) \neq d\}$. Since all A_d and $B_{d,k}$ are in $\bigvee_{n\in\mathbf{N}}\mathcal{F}_n$, there exist $N_0 \in \mathbf{N}$ and sets $\left(\bar{A}_d, \bar{B}_d\right)_{d\in D} \in \mathcal{F}_{N_0}$ such that:

(1) For each $d \in D$: $\bar{A}_d \bigcap \bar{B}_d = \emptyset$ and $\left(\bar{A}_d \bigcup \bar{B}_d\right) = \Omega$.
(2) $\forall d \in D$, $p\left(A_d | \bar{A}_d\right) \geq 1 - \frac{\epsilon}{4\cdot|D|}$
(3) $\forall d \in D$, $p\left(B_{d,N_0} | \bar{B}_d\right) \geq 1 - \frac{\epsilon}{4\cdot|D|}$

Let $E = \bigcup_{d\in D}\left(\left\{\omega \in \bar{A}_d | \omega \notin A_d\right\} \bigcup \left\{\omega \in \bar{B}_d | \omega \notin B_{d,N_0}\right\}\right)$. Observe that $p(E) < \frac{\epsilon}{2}$. For any $F \in \mathcal{F}$ let $D_F = \left\{d \in D | F \in \bar{A}_d\right\}$, and let $\alpha^i_F = \max\left(d^i_{\{i\}} | d \in d_F\right)$.

The induced game G_F is the terminating game that begins at stage τ when the players know that $F(\omega) \subseteq \hat{\mathcal{F}}_\tau$.

Definition 14 Let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a terminating game, let $N_0 \in \mathbf{N}$ be as defined in the last subsection, $\tau > N_0$ a bounded terminating time, and $F \subseteq \hat{\mathcal{F}}_\tau$. The game G_F is the terminating game that is restricted to F and starts at stage τ: $G_F = \left(I, F, \mathcal{A}_F, p_{|F}, (\mathcal{F}_{\tau+k})_{k\geq 0}, (R_{\tau+k})_{k\geq 0}\right)$ where: $\mathcal{A}_F$ is the σ-algebra over F induced by $\mathcal{A}$, and $p_{|F}$ is the probability distribution p conditioned on F.

A strategy profile x in $G(\mathcal{D})$ is ϵ-unrevealing if when each player obtains his message $m^i \in M^i$, his expected payoff is changed by at most ϵ.

Definition 15 Let G be a terminating game, $\epsilon > 0$, $\mathcal{D} = (M, \mu)$ a correlation device, and x a profile in $G(\mathcal{D})$. *The profile x is* ϵ-unrevealing *if there is a set $M' \subseteq M$ satisfying $\mu(M') \leq \epsilon$, such that for every player $i \in I$ and every message $m^i \in (M \backslash M')^i$: $|\gamma^i(x|m^i) - \gamma^i(x)| \leq \epsilon$, where $\gamma^i(x|m^i) = \mathbf{E}_x\left(R^i_{S_*,\theta}\mathbf{1}_{\{\theta<\infty\}}|m^i\right)$ is the expected payoff of player i where the players play according to x, conditioned on receiving a message m^i.*

The following lemma is standard (An extension of Lemma 7.3 in [27]).

Lemma 16 Let G be a terminating game, $\delta, \epsilon > 0$, *τ a bounded stopping time, and $E \subseteq \Omega$ an event with $p(E) < \delta$. Assume that for every $F \in \mathcal{F}_\tau$ satisfying $F \notin E$, there is a a correlation device $\mathcal{D}_F = (M_F, \mu_F)$, a set of signals $M'_F \subseteq M_F$ satisfying $\mu_F(M'_F) \leq \delta$ and a perfect correlated* ϵ-unrevealing *ϵ-equilibrium x_F of $G_F(\mathcal{D}_F)$ conditioned on $\Omega \backslash E$ and given $M_F \backslash M'_F$. Moreover, assume that the correlation device $\mathcal{D}_F = (M_F, \mu_F)$ depends only on ϵ and D_F (the set of matrix payoffs that occur i.o.). Then the game G admits a perfect correlated $\left(2^{|D|} \cdot \delta, 3 \cdot \epsilon\right)$-equilibrium with a universal correlation device .*

PROOF. As each M_F and μ_F depend only on D_F and ϵ, we identify M_F with $M_{D_F,\epsilon}$ and μ_F with $\mu_{D_F,\epsilon}$. Let $M = \prod_{D'\subseteq D} M_{D',\epsilon}$, $\mu = \prod_{D'\subseteq D} \mu_{D',\epsilon}$ and $\mathcal{D} = (M, \mu)$. Let $M' = \bigcup_{D'\subseteq D} \left\{m \in M | m_{D',\epsilon} \in M'_{D',\epsilon}\right\}$. Note that $\mu(M') \leq 2^{|D|} \cdot \delta$. It is well known that any finite-stage game admits a 0-equilibrium (see, e.g., [24, Prop. 3.1]). Since τ is bounded, $p(E) \leq \epsilon$ and $\mu(M') \leq 2^{|D|} \cdot \delta$, the following strategy profile x is a $\left(2^{|D|} \cdot \delta, 3 \cdot \epsilon\right)$-equilibrium in $G(\mathcal{D})$:

- Until stage τ, play a 0-equilibrium in the game that terminates at τ, if no player stops before that stage, with a terminal payoff $\gamma^i(x_F)$ where $F =$

$F_{\tau(\omega)} \in \hat{\mathcal{F}}_\tau$.
- If the game has not terminated by stage τ, play from that stage on the profile x_F in G_F.

Thus in order to prove Theorem 6, it remains to show that there exists a bounded terminating time $\tau \geq N_0$, such that for every $\epsilon > 0$ and for every $F \in \hat{\mathcal{F}}_\tau$, there is a correlation device $\mathcal{D}_F = (M_F, \mu_F)$ that depends only on D_F and ϵ, a set of signals $M'_F \subseteq M_F$ satisfying $\mu_F(M'_F) \leq \delta$ and a perfect correlated ϵ-*unrevealing* ϵ-equilibrium x_F of $G_F(\mathcal{D}_F)$ conditioned on $\Omega \backslash E$ and given $M_F \backslash M'_F$.

5 Terminating Games on Finite trees

An important building block in our analysis is terminating games that are played on finite trees. In the present subsection we define these games. discuss their equivalence with absorbing games, and study some of their properties.

5.1 Finite trees

Definition 17 *A terminating game on a finite tree* (or simply a *game on a tree*) is a tuple $T = \left(I, V, V_{leaf}, r, V_{stop}, (C_v, p_v, R_v)_{v \in V \backslash V_{leaf}}\right)$, where:

- I is a finite non-empty set of players.
- $\left(V, V_{leaf}, r, (C_v)_{v \in V \backslash V_{leaf}}\right)$ is a tree, V is a nonempty finite set of nodes, $V_{leaf} \subseteq V$ is a nonempty set of *leaves*, $r \in V$ is the root, and for each $v \in V \backslash V_{leaf}$, $C_v \subseteq V \backslash \{r\}$ is the nonempty set of children of v. We denote by $V_0 = V \backslash V_{leaf}$ the set of nodes which are not leaves.
- $V_{stop} \subseteq V_0$ is the set of nodes the players can choose to stop in. Observe that players can not stop at the leaves.

and for every $v \in V_0$:

- p_v is a probability distribution over C_V; We assume that $\forall \tilde{v} \in C_v$: $p_v(\tilde{v}) > 0$.
- $R_v = \left(R^i_{v,S}\right)_{i \in I, \emptyset \neq S \subseteq I} \in D$ is the payoff matrix at v if a nonempty set of players stops at that node.

A terminating game on a finite tree starts at the root and is played in stages. Given the current node $v \in V_{stop}$, and the sequence of nodes already visited, the players decide, simultaneously and independently, whether to stop or to continue. Let S be the set of players that decides to stop. If $S \neq \emptyset$, the play terminates and the terminal payoff to each player i is $R^i_{v,S}$. If $S = \emptyset$, a new node $v \in C_V$ is chosen according to p_s . The process now repeats itself, with v being the current node. If $v \in V \backslash V_{stop}$ then the players can not stop at that stage, and a new node $v \in C_V$ is chosen according to p_v. If $v \in V_{leaf}$ then the new current node is the root r. The game on the tree is essentially played in rounds, where each round starts at the root and ends once it reaches a leaf.

A *stationary strategy* of player i is a function $x^i : V_{Stop} \rightarrow [0,1]$; $x^i(v)$ is the probability that player 1 stops at v. Let $x = (x^i)_{i \in I}$ be a stationary strategy profile. Let c^i be the stationary strategy of player i that never stops, and let $c = (c^i)_{i \in I}$. Denote by $\gamma_T^i(x) = \gamma^i(x)$ the expected payoff under x, and denote by $\pi_T(x) = \pi(x)$ the probability the game that the game is stopped at the first round (before reaching a leaf).

Definition 18 A profile of stationary strategies $x = (x_i)_{i \in I}$ is an ϵ-equilibrium of the game on a tree T if, for each player $i \in I$, and for each strategy y_i, $\gamma^i(x) > \gamma^i(x^{-i}, y^i) - \epsilon$.

Assuming no player ever stops, the collection $(p_v)_{v \in V_0}$ of probability distributions at the nodes induces a probability distribution over the set V_{leaf} of leaves or, equivalently, over the set of branches that connect the root to the leaves. For each set $\hat{V} \subseteq V_0$, we denote by $p_{\hat{V}}$ the probability that the chosen branch passes through $\hat{V}$. For each $v \in V$, we denote by F_v the event that the chosen branch passes through v.

We finish this this subsection by defining the game on a finite tree $T_{n,\sigma}(F)$. The game begins at stage n, when $\omega \in F \subseteq \hat{\mathcal{F}}_n$ is randomly chosen (according to $p_{|F}$). If the game has not absorbed before reaching stage $\tau(n)$, the game restarts at stage n again (and a new $\omega \in F \subseteq \hat{\mathcal{F}}_n$ is randomly chosen).

Definition 19 Let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a tree-like terminating game, $n \in \mathbf{N}$ a number, $n < \tau$ a bounded terminating time, and $F \in \hat{\mathcal{F}}_n$. The terminating game on the finite tree $T_{n,\sigma}(F)$ is $\left(I, V, V_{leaf}, r, V_{stop}, (C_v, p_v, R_v)_{v \in V \backslash V_{leaf}}\right)$ where:

- $V = \bigcup_{\substack{\omega \in F \\ n \leq k \leq \tau(\omega)}} \{F_k(\omega)\}$, $V_{leaf} = \bigcup_{\omega \in F} \{F_\tau(\omega)\}$, $r = F$
- R_v,C_v, p_v are defined by induction. Assume that $v \in V \backslash V_{leaf}$ and $v \in \hat{\mathcal{F}}_k$ for some $n \leq k$, then: $R_v = R_n(v)$, $C_v = \left\{F_{k+1} \in \hat{\mathcal{F}}_{k+1} | F_{k+1} \subseteq v\right\}$, and $p_v(F_{k+1}) = p\left(F_{k+1} | v\right)$.
- $V_{stop} = \{v \in V | d_v \in D_F\}$.

A terminating game on a finite tree T is equivalent to an absorbing game, where each round of T corresponds to a single stage of the absorbing game. An absorbing game is a stochastic game with a single non-absorbing state. Solan and Vohra ([31]) proved that every absorbing game admits a correlated ϵ-equilibrium for every $\epsilon > 0$.

As an absorbing game, the game T has two special properties:

- It is a recursive game: the payoff in the non-absorbing state is 0.
- There is a single action profile that is non-absorbing. In all other action profiles the game has a positive probability to be absorbed.

Adapting [31]'s Prop. 4.10 to the two special properties described above gives the following proposition:

Definition 20 Let T be a game on a tree, and $i \in I$ a player. $g^i = \max_{v \in V_{stop}} \left(R^i_{i,v}\right)$ is the maximal payoff a player can get in T by terminating alone. Let $\tilde{v}^i$ be a node that maximizes the last expression, and let $d_{\tilde{v}^i} \in D$ be the payoff matrix in that stage. [5]

Proposition 21 (an adaptation of Prop. 4.10 from [31]). Let T be a game on a finite tree. Then T has one of the following:

(1) A stationary absorbing equilibrium $x \neq c$.
(2) For each player $i \in I$ and for each node $v \in V_{stop}$, : $R^i_{i,v} \leq 0$. This implies that c is a stationary equilibrium.
(3) There is a distribution $\eta \in \Delta(I \times \{\tilde{v}^i\})$ such that:
  (a) $\sum_{i \in I} \mathbf{P}_\eta(\tilde{v}^i, i) = 1$.
  (b) For each player $j \in I$: $\sum_{i \in I} \mathbf{P}_\eta(\tilde{v}^i, i) \cdot R^j_{\{i\},\tilde{v}^i} \geq g^j$.
  (c) Let the players $i \in I$ that satisfy $\mathbf{P}_\eta(\tilde{v}^i, i) > 0$ be denoted as the *stopping players*. For every stopping player $i \in I$ there exists a player $j_i \neq i$, the punisher of i, such that: $g^i \geq R^i_{\{j_i\},\tilde{v}^{j_i}}$.

When we want to emphasize the dependency of these variables on the game T, we write $g^i_T, \tilde{v}^i_T, \eta_T, x_T$.

In the original prop. 21, the equilibrium in case 1 may not be perfect, as players

[5] Originally part 3 of Prop. 21 requires that every player would have a unique pure action that maximizes his payoff, conditioned on that the other players always continue. This can be achieved by small ($o(\epsilon)$) perturbations on the payoffs, such that $R^i_{i,\tilde{v}^i}$ is strictly larger than any other payoff $R^i_{i,v}$ where $v \in V_{stop}$.

may use non-credible threats after a node where the game is terminated with probability 1. The following lemma asserts that a perfect ϵ-equilibrium exists.

Lemma 22 *In case 1 of prop. 21, the game admits a stationary absorbing perfect ϵ-equilibrium $x \neq c$.*

PROOF. Let T_ϵ be a perturbed version of the game on a tree T: In T_ϵ when a non-empty coalition stops at some node, there is a probability ϵ^2 that the stopping is ignored, and the game continues to the next stage as if no player has stopped. In T_ϵ under any profile x, any node is reached with a positive probability, thus non-credible threats cannot be used in an equilibrium. Thus if case 1 of prop. 21 applies, then the game T_ϵ admits a perfect equilibrium x_ϵ, and x_ϵ is a perfect stationary absorbing ϵ-equilibrium in T.

5.3 Limits on Per-Round Probability of Termination

In this subsection we bound the probability of termination in a single round when a stationary equilibrium $x \neq c$ exists (case 1 of Prop. 21), by adapting to the multi-player case the methods presented in [27, Subsec. 5.2] for two players. We first bound the probability of termination in a single round when the ϵ-equilibrium payoff is low for at least one player. The lemma is an adaptation of Lemma 5.3 in [27]. The proof is omitted as the changes compared with [27] are minor.

Lemma 23 *(An adaptation of Lemma 5.3 in [27]) Let* G *be a terminating game, $n \in \mathbf{N}$, $\sigma > n$ a bounded stopping time, $F \in \hat{\mathcal{F}}_n$, and $\epsilon > 0$. Let $x \neq c$ be a stationary $\frac{\epsilon}{2}$-equilibrium in* $T_{n,\sigma}(F)$ such that there exists a player $i \in I$ with a low payoff: $\gamma^i(x) \leq \alpha^i_F - \epsilon$. Then $\pi(c^i, x^{-i}) \geq \frac{\epsilon}{6} \cdot q^i$, where $q^i = q^i_T = p\left(\bigcup_{v \in V_{stop}} \left\{F_v | R^i_{\{i\},v} = \alpha^i_F\right\}\right)$ is the probability that if all the players never stop, the game visits a node $v \in V_{stop}$ with $R^i_{\{i\},v} = \alpha^i_F$ in the first round.

Definition 24 Let $T = \left(I, V, V_{leaf}, r, V_{stop}, (C_v, p_v, R_v)_{v \in V_0}\right)$ and let $T' = \left(I, V', V'_{leaf}, r', V'_{stop}, (C'_v, p'_v, R'_v)_{v \in V'_0}\right)$ be two games on trees. We say that T' is a *subgame* of T if: $V' \subseteq V$, $V'_{stop} = V_{stop} \bigcap V'$, $r' = r$, and for every $v \in V'_0$, $C'_v = C_v$, $p'_v = p_v$ and $R'_v = R_v$.

In words, T' is a subgame of T if we remove all the descendants (in the strict sense) of several nodes from the tree $\left(V, V_{leaf}, r, (C_v)_{v \in V_0}\right)$ and keep all other parameters fixed. Observe that this notion is different from the standard definition of a subgame in game theory.

Let T be a game on a tree. For each subset $D \subseteq V_0$, we denote by T_D the subgame of T generated by trimming T from D downward. Thus, all strict descendants of nodes in D are removed. For every subgame T' of T and every subgame T'' of T', let $p_{T'',T'} = p_{V''_{leaf},V'_{leaf}}$ be the probability that the chosen branch in T passes through a leaf of T'' strictly before it passes through a leaf of T'.

The following definition divides the sets in $\hat{\mathcal{F}}_n$ into 2: *simple* and *complicated.*

Definition 25 Let G be a terminating game, $\epsilon > 0$, and $N_0 \le n \in \mathbf{N}$. The set $F \in \hat{\mathcal{F}}_n$ is *ϵ-simple* if one of the following holds:

(1) For every $i \in I$: $\alpha^i_F < 0$. or

(2) There is a distribution $\theta \in \Delta(D_F \times I)$ such that for each player $i \in I$:

 (a) $\theta(d,i) > 0 \Rightarrow R^i_{\{i\},d} = \alpha^i_F$. and

 (b) $\alpha^i_F + \epsilon \ge \sum_{j \in I,\, d \in D_F} \theta(d,j) \cdot R^i_{\{j\},d} \ge \alpha^i_F - \epsilon$.

F is simple if it is *ϵ-simple* for every $\epsilon > 0$. F is complicated if it is not simple, i.e.: there is an $\epsilon_0 > 0$ such that F is not ϵ_0-simple. In that case we say that F is complicated w.r.t. ϵ_0. *Observe that* $F_n \in \hat{\mathcal{F}}_n$ *is ϵ-simple if and only if* $F_{N_0} \in \hat{\mathcal{F}}_{N_0}$ *is ϵ-simple.*

The next proposition analyzes stationary $\epsilon-$equilibria that yield a high payoff to all the players. The proposition is an adaptation of Prop. 5.5 in [27, Sec. 8]. The proof is omitted as the changes compared with [27] are minor.

Proposition 26 *Let* G *be a terminating game,* $N_0 \le n \in \mathbf{N}$, $\sigma > n$ *a bounded stopping time,* $F \in \hat{\mathcal{F}}_n$ *a complicated set (w.r.t.* ϵ_0*),* $\epsilon << \frac{\epsilon_0}{|I| \cdot K^2}$ *, and for each* $i \in I$ *let* $a^i \ge \alpha^i_F - \epsilon$*. Then there exists a set* $U \subseteq V_0$ *of nodes and a strategy profile* x *in* $T = T_{n,\sigma}(F)$ *such that:*

(1) No subgame of T_U has an ϵ-equilibrium with a corresponding payoff in $\prod_{i \in I} [a^i, a^i + \epsilon]$

(2) Either: (a) $U = \emptyset$ (so that $T_U = T$) or (b) x is a 9ϵ-equilibrium in T, and for every $i \in I$ and for every strategy y^i: $a^i - \epsilon \le \gamma^i(x)$, $\gamma^i(x^{-i}, y^i) \le a^i + 8\epsilon$, and $\pi(x) \ge \epsilon^2 \cdot p_U$.

6 The Use of Ramsey Theorem

In this section we use a stochastic variation of Ramsey theorem ([27]), to disassemble an infinite terminating game into games on finite trees with special properties.

We begin by defining an $\mathcal{F}$-consistent C-valued NT-function.

Definition 27 An NT-function is a function that assigns to every integer $n > 0$ and every bounded stopping time τ an $\mathcal{F}_n$-measurable r.v. that is defined over the set $\{\tau > n\}$. We say that an NT-function f is C-valued, for some finite set C, if the r.v. $f_{n,\tau}$ is C-valued, for every $n > 0$ and every bounded stopping time τ.

For every $A, B \in \mathcal{A}$, A holds on B if and only if $p(A^c \bigcap B) = 0$.

Definition 28 An NT-function f is *$\mathcal{F}$-consistent* if for every $n > 0$, every $\mathcal{F}_n$-measurable set F, and every two stopping times τ_1, τ_2, we have: $\tau_1 = \tau_2 > n$ on F implies $f_{n,\sigma_1} = f_{n,\sigma_2}$ on F.

When f is an NT-function, and $\tau_1 < \tau_2$ are two bounded stopping times we denote $f_{\tau_1,\tau_2}(\omega) = f_{\tau_1(\omega),\tau_2}(\omega)$. Thus f_{τ_1,τ_2} is an $\mathcal{F}_n$-measurable r.v.

The following proposition was proved by Shmaya and Solan ([27, Theorem 4.3]):

Proposition 29 For every finite set C, every C-valued $\mathcal{F}$-consistent NT-function f, and every $\epsilon > 0$, there exists an increasing sequence of bounded stopping times $0 < \sigma_1 < \sigma_2 < \sigma_3 < ...$ such that: $p\left(f_{\sigma_1,\sigma_2} = f_{\sigma_2,\sigma_3} = ...\right) > 1-\epsilon$.

In the rest of this section we provide an algorithm that attaches a color $c_{n,\sigma}(F)$ and several numbers $\left(\lambda_{j,n,\sigma}(F)\right)_j$ for every $F \in \hat{\mathcal{F}}_n$, s.t. $c_{n,\sigma}(F)$ is a C-valued $\mathcal{F}$-consistent NT-function.

If $F \in \hat{\mathcal{F}}_{N_0}$ is complicated, let $\epsilon_0(F) > 0$ satisfies that F is complicated w.r.t. $\epsilon_0(F)$. Otherwise let $\epsilon_0(F) = 1$. From now on we fix $0 < \epsilon << \min_{F \in \hat{\mathcal{F}}_{N_0}} \frac{\epsilon_0(F)}{|I| \cdot K^2}$.

A hyper-rectangle $([a^i, a^i + \epsilon])_{i \in I}$ is *bad* if for every $i \in I$ $\alpha^i_F - \epsilon \leq a^i$. It is *good* if there exists a player $i \in I$ such that $a^i + \epsilon \leq \alpha^i_F - \epsilon$.

Let M be a finite covering of $[-1,1]^{|I|}$ with (not necessarily disjoint) hype-rectangles $([a^i, a^i + \epsilon])_{i \in I}$, all of which are either good or bad. Thus, for every $u \in [-1,1]^{|I|}$ there is a rectangle $m \in M$ such that $u \in m$. We denote by $B = \{b_1, b_2, ..., b_J\}$ the set of bad rectangles in M and denote by $O = \{o_1, o_2, ..., o_W\}$ the set of good rectangles in M.

Set $C = (simple \bigcup allbad \bigcup \{1 \times O\} \bigcup \{2\} \bigcup \{3 \times M \times M\})$

Let G be a terminating game, $n \in \mathbf{N}$, $\sigma > n$ a bounded stopping time, and $F \in \hat{\mathcal{F}}_n$. If F is *ϵ-complicated* then the color $c_{n,\sigma}(F)$ is determined by the following procedure [6] :

- Set $T^{(0)} = T_{n,\sigma}(F)$.
- For $1 \leq j \leq J$ apply Prop. 21 to $T^{(j-1)}$ and the bad rectangle $h_j = \prod_{i \in I} \left[a_j^i, a_j^i + \epsilon\right]$ to obtain a subgame $T^{(j)}$ of $T^{(j-1)}$ and strategy profile x_j in $T^{(j)}$ such that:

(1) No subgame of $T^{(j)}$ has a stationary ϵ-equilibrium with a corresponding payoff in h_j.

(2) Either $T^{(j)} = T^{(j-1)}$ or the following three conditions hold:

  (a) For every $i \in I$, $a_j^i - \epsilon \leq \gamma^i(x_j)$.

  (b) For every $i \in I$ and every strategy y^i: $\gamma^i(x_j^{-i}, y^i) \leq a_j^i + 8\epsilon$.

  (c) We have $\pi\left(x_j\right) \geq \epsilon^2 \times p_{T^{(J)},T^{(j-1)}}$, where $p_{T^{(J)},T^{(j-1)}}$ is the probability that a randomly chosen branch passes through a leaf of $T^{(J)}$ which is not a leaf of $T^{(J-1)}$.

- If $T^{(J)}$ is trivial (i.e., the only node is the root), set $c_{n,\sigma}(F) = allbad$; otherwise due to Prop. 21 and our procedure one of the following must hold:

(1) $T^{(J)}$ has a perfect stationary absorbing ϵ-equilibrium x, with a payoff $\gamma(x)$ in one of the good hyper-rectangles. Let $c_{n,\sigma}(F) = (1, o_l)$, where o_l is the good rectangle that includes γ_x .

(2) $T^{(J)}$ has a perfect stationary non-absorbing equilibrium c, with a payoff 0. Let $c_{n,\sigma}(F) = (2)$.

(3) There is a correlated strategy profile $\eta \in \Delta(A)$ in $T^{(J)}$ that satisfies 3(a)+3(b)+3(c) in Prop. 21. Let $c_{n,\sigma}(F) = (3, m_1, m_2)$ where m_1 is the hyper-rectangle that includes $\gamma_{T^{(J)}}(\eta)$, and m_2 is the hyper-rectangle that includes $g(T^{(J)})$.

The strategy profile x_J, as given by Prop. 21, are strategies in $T^{(j-1)}$. We consider them as strategies in T by letting them continue from the leaves of $T^{(j-1)}$ downward.

We also define, for every $j \in J$, $\lambda_{j,n,\sigma}(F) = p_{T^{(j)},T^{(j-1)}}$. Observe that due to Prop. 21: $\pi(x^{(j)}) \geq \epsilon^2 \times \lambda_{j,n,\sigma}(F)$.

If F is ϵ-simple we let $c_{n,\sigma}(F) = simple$.

By Prop. 29 there exists an increasing sequence of bounded stopping times $0 < \sigma_1 < \sigma_2 < \sigma_3 < ...$ such that: $p\left(c_{\sigma_1,\sigma_2} = c_{\sigma_2,\sigma_3} = ...\right) > 1 - \epsilon$. For every $F \in \hat{\mathcal{F}}_{\sigma_1}$, let $c_F = c_{\sigma_1,\sigma_2}(F)$.

Let $(A_{\epsilon,j}, A_{\infty,j})_{j \in J} \in \bigvee_{n=1..\infty} \mathcal{F}_n$ be: $A_{\infty,j} = \left\{w \in \Omega | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}}\left(F_{\sigma_k(\omega)}\right) = \infty\right\}$,

[6] The procedure is an adaptation of the 2-player procedure described in [27, Sec. 5]

$A_{\epsilon,j} = \left\{ w \in \Omega | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}} \left(F_{\sigma_k(\omega)} \right) \leq \frac{\epsilon}{|J|} \right\}$. As $(A_{\epsilon,j}, A_{\infty,j})_{j \in J} \in \bigvee_{n=1..\infty} \mathcal{F}_n$, there is large enough $N_1 \geq N_0$ and sets $\left(\bar{A}_{\epsilon,j}, \bar{A}_{\infty,j} \right)_{j \in J} \in \mathcal{F}_{N_1}$ such that:

(1) For each $j \in J$:$\bar{A}_{\epsilon,j} \bigcap \bar{A}_{\infty,j} = \emptyset$ and$\left(\bar{A}_{\epsilon,j} \bigcup \bar{A}_{\infty,j} \right) = \Omega$.
(2) $p \left(A_{\epsilon,j} | \bar{A}_{\epsilon,j} \right) \geq 1 - \frac{\epsilon}{2 \cdot |J|}$
(3) $p \left(A_{\infty,j} | \bar{A}_{\infty,j} \right) \geq 1 - \frac{\epsilon}{2 \cdot |J|}$

Let E' be defined as follows (where E was defined in Subseq. 4.3):

$$E' = E \bigcup_{j \in J} \left\{ \omega \in \bar{A}_{\epsilon,j} | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}} \left(F_{\sigma_k(\omega)} \right) > \frac{\epsilon}{|J|} \right\}$$
$$\bigcup_{j \in J} \left\{ \omega \in \bar{A}_{\infty,j} | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}} \left(F_{\sigma_k(\omega)} \right) < \infty \right\}$$

Observe that $p(E') \leq \epsilon$. We assume w.l.o.g. that $\sigma_1 \geq N_1$.

7 ϵ-Unrevealing Correlated 9ϵ- Equilibria in G_F

In this subsection we finish the proof of the main theorem by the following proposition:

Proposition 30 *Let G be a tree-like terminating game, let the event $E' \subseteq \Omega$ and the stopping time σ_0 be defined as in the last subsection, and let $F \in \hat{\mathcal{F}}_{\sigma_1}$. Then there is a correlation device $\mathcal{D}_F = (M_F, \mu_F)$ and a strategy profile x_F in the game $G_F(\mathcal{D}_F)$, such that:*

- *$\mathcal{D}_F$ depends only on ϵ and the set of payoffs D_F.*
- *The profile x_F is a perfect correlated ϵ-unrevealing ϵ-equilibrium in the game $G_F(\mathcal{D}_F)$ conditioned on $\Omega \backslash E$ and given $M \backslash M'$.*

PROOF. The proof is divided to a few cases according to the color of c_F and to whether $F \in \bar{A}_{\infty,j}$. The proof in the first 3 cases is an adaptation of [27, Sec.7].

7.1 There exists $j \in J$ s.t. $F \in \bar{A}_{\infty,j}$

Let $1 \leq j \leq J$ be the smallest index such that $F \in \bar{A}_{\infty,j}$. Let $x_{j,\sigma_k,\sigma_{k+1}}$ be the j^{th} profile in the procedure described in Section 6, when applied to $T_{\sigma_k,\sigma_{k+1}}$.

Let x_F be the following strategy profile in G_F: between σ_k and σ_{k+1} play according to $x_{j,\sigma_k,\sigma_{k+1}}$. The procedure of Section 6 implies the following:

- Conditioned on that the game was absorbed between σ_k and σ_{k+1} the profile $x_{j,\sigma_k,\sigma_{k+1}}$ gives each player a payoff: $a^i_j - \epsilon \leq \gamma^i_{\sigma_k,\sigma_{k+1}}(x_j) \leq a^i_j + 8\epsilon$.
- For each player $i \in I$ and for each strategy y^i in $T_{\sigma_k,\sigma_{k+1}}$:
  - $\gamma^i_{\sigma_k,\sigma_{k+1}}(x^{-i}_j, y^i) \leq a^i_j + 8\epsilon$.
  - $\pi_{\sigma_k,\sigma_{k+1}}(x_j) \geq \epsilon^2 \times \lambda_j(T_{\sigma_k,\sigma_{k+1}})$

The fact that $F \in \bar{A}_{\infty,j}$ implies that outside E' the game is absorbed with probability 1. All those facts imply that x_F is a 11ϵ-equilibrium in $\Omega \backslash E'$. Observe that $c_F = allbed$ implies that there exists $j \in J$ such that $F \in \bar{A}_{\infty,j}$.

7.2 $F \in \bigcap\limits_{j \in J} \bar{A}_{\epsilon,j}$ and $c_F = 2$

Let x_F be the profile in which everyone continues. It is implied that no player can profit more than ϵ by deviating, conditioned on $\omega \in \Omega \backslash E'$.

7.3 $F \in \bigcap\limits_{j \in J} \bar{A}_{\epsilon,j}$ and $c_F = (1, o_w) \in (1 \times O)$

Let $x_{\sigma_k,\sigma_{k+1}}$ be a stationary absorbing equilibrium in $T^{(J)}$ with a payoff $\gamma_{\sigma_k,\sigma_{k+1}}$ in the good hyper-rectangle o_w: $\prod\limits_{i \in I} [a^i_w, a^i_w + \epsilon]$. As o_w is good, there is a player $i \in I$ s.t.: $a^i_w \leq \alpha^i_F - 2\epsilon$. Let x_F be the following strategy profile in G_F: between σ_k and σ_{k+1} play according to $x_{\sigma_k,\sigma_{k+1}}$. Lemma 23 implies that $\pi(c^i, x^{-i}_{\sigma_k,\sigma_{k+1}}) \geq \frac{\epsilon}{6} \cdot q^i_{\sigma_k,\sigma_{k+1}}$, where $q^i_{\sigma_k,\sigma_{k+1}} = p(\exists \sigma_k \leq n < \sigma_{k+1}, R^i_{i,n} = \alpha^i_F, R^i_{i,n} \in D_F)$. Outside E', $R^i_{i,n} = \alpha^i_F$ infinitely often and $\sum\limits_{j=1..J} \sum\limits_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}} < \frac{\epsilon}{2}$. This implies that under x_F the game is absorbed with probability 1, and that x_F is a 4ϵ-equilibrium in G, conditioned on $\Omega \backslash E'$.

7.4 $F \in \bigcap\limits_{j \in J} \bar{A}_{\epsilon,j}$ and $c_F = (1, m_w, m_{w'}) \in (1 \times M \times M)$

The construction in this case is as an adaptation of the procedure of [30], which deals with quitting games - terminating games where the payoff matrix is the same in all the stages.

Let $\eta = \eta_{\sigma_1,\sigma_2}$ be a correlated strategy profile in T_{σ_1,σ_2} that satisfies 3(a), 3(b) and 3(c) in Prop. 21. The definition of α^i_F implies that $\alpha^i_F = g^i(\eta) \in m^i_{w'}$.

This implies that there is a distribution $\theta = \theta(\eta) \in \Delta(D_F \times I)$ such that for each player $i \in I$:

(1) $\theta(d,i) > 0 \Rightarrow R^i_{\{i\},d} = \alpha^i_F, \ \forall d' \in D_F \ \theta(d',i) = 0$. Let $d(i) \in D_F$ be the single matrix payoff that satisfies $\theta(d_i,i) > 0$. If no such payoff exists, let $d(i) = \emptyset$.
(2) $\sum_{j \in I, d \in D_F} \theta(d,j) \cdot R^i_{\{j\},d} \geq \alpha^i_F$
(3) If there is $d \in D_F$ such that $\theta(d,i) > 0$, then there exists a punisher $j_i \in I$ such that: $d(j_i) \neq \emptyset$ and $d(j_i)^i_{\{j_i\}} \leq \alpha^i_F$.

Let $\zeta \in \Delta(I)$ be: $\zeta(i) = \eta(d(i),i)$. Let $(\tau^i_k)_{i \in I. k=1..\infty}$ be an increasing sequence of stopping times defined by induction: $\tau^{i_0}_1$ is the first stage n such that $R_n = d(i_0)$. $\tau^{i_0}_{n+1}$ is the first stage $n > \max_{i \in I}(\tau^i_n)$ such that $R_n = d(i_0)$. Observe that in $\Omega \backslash E'$ each $\tau^i_k < \infty$.

We now describe the correlation device $\mathcal{D}_F = (M_F, \mu_F)$. Let $M^i_F = \{1, ..., \hat{T} + T + 1\}$, where $T \in \mathbf{N}$ is sufficiently large, and $\hat{T} >> T$. Let μ_F, the probability according to which the signals are sent to the players, be as follows:

(1) A number $\hat{l} \in \mathbf{N}$ is chosen uniformly over $\left\{1, \hat{T}\right\}$.
(2) The quitter $i \in I$ is randomly chosen according to ζ. Player i receives the signal $\hat{l}$.
(3) A number $l \in \mathbf{N}$ is chosen uniformly over $\left\{\hat{l}+1, \hat{l}+T\right\}$
(4) Let player j be the punisher of player i. Player j receives the signal l.
(5) Each other player $\tilde{i} \neq i, j$ receives the signal $l+1$.

Observe that $\mathcal{D}_F$ is universal: it depends only on $D_F(\epsilon)$ and the number of players. Let $M'_F \subseteq M_F$ be those signal profiles in which some of the players receive an "extreme" signal: relative close to 1 or to $\hat{T} + T$. If $T, \hat{T}$ are large enough, we can assume that $\mu(M') \leq \delta$. Define now the following strategy x^i_F for each player $i \in I$: let m_i be the signal of player i. Player i stops in stage τ_{m_i}, and continues in all other stages.

If the players follow the strategy profile x_F then the game is absorbed with probability 1 in $\Omega \backslash E'$ and the expected payoff satisfies $\alpha^i_F \leq \gamma^i_F(x) \in m^i_w$. Moreover, if $\hat{T} >> T$, then immediately after receiving his signal m_i (assuming $m \in M \backslash M'$) no player can infer from his signal whether or not he is the quitter, thus x_F is ϵ-unrevealing w.r.t. $\mathcal{D}_F$.

We now verify that if $T, \hat{T}$ are sufficiently large, no player can gain too much by deviating conditioned on that $\omega \in \Omega \backslash E'$ and given $m \in M \backslash M'$. First, the probability the quitter $i \in I$ correctly guesses the punishment stage is very low, and thus he cannot profit too much by deviating. Similarly, any other player $(j \neq i \in I)$ has a low probability to correctly guess $\tau^i_{\hat{l}}$, the stage the quitter stops . Moreover, if T is sufficiently large, then, with high probability, player j does not know whether he is the punisher or not, and thus he cannot infer which of the other players is more likely to be the quitter. Therefore, player j can't earn much by stopping before stage $\hat{l}$. Observe that when the quitter deviates and does not stop. his punisher, say player i, does not know that he is a punisher, but when he has to stop, he believes that he is the quitter (assuming $m \in M \backslash M'$). This implies that the players ϵ-best-respond while punishing, and the ϵ-equilibrium is in-deed perfect. Thus it is implied that x_F is a perfect ϵ-equilibrium in $G_F(\mathcal{D}_F)$ conditioned on $\omega \in \Omega \backslash E'$ and given $m \in M \backslash M'$.

7.5 $c_F = simple$

If for every $i \in I$: $\alpha_F^i \leq 0$, then the profile in which all the players always continue is an equilibrium in $\Omega \backslash E'$. Otherwise, the fact that $c_F = simple$ implies that there is a distribution $\theta \in \Delta(D_F \times I)$ such that for each player $i \in I$:

(1) $\theta(d,i) > 0 \Rightarrow R^i_{\{i\},d} = \alpha_F^i$
(2) $\alpha_F^i + \epsilon \geq \sum\limits_{j \in I,\, d \in D_F} \theta(d,j) \cdot R^i_{\{j\},d} \geq \alpha_F^i - \epsilon$

Thus we can use a simpler version of the procedure of the previous case. Let $(\tau^i_k)_{i \in I.k=1..\infty}$ be as defined earlier. We now describe the correlation device $\mathcal{D}_F = (M_F, p_F)$. Let $M_F^i = \left\{1, ..., \hat{T} + T + 1\right\} \times D_F$, where $T \in \mathbf{N}$ is sufficiently large, and where $\hat{T} >> T$. Let μ_F, the probability according to which the signals are sent to the players before the game starts, be as follows:

(1) A number $\hat{l} \in \mathbf{N}$ is chosen uniformly over $\left\{1, \hat{T}\right\}$.
(2) The couple $(d,i) \in D_F \times I$ is randomly chosen according to θ. Player i receives the signal $\left(\hat{l}, d\right)$.
(3) A number $l \in \mathbf{N}$ is chosen uniformly over $\left\{\hat{l}+1, \hat{l}+T\right\}$
(4) The couple $(d',j) \in D_F \times I$ is randomly chosen according to θ, conditioned on that $j \neq i$. Player j receives the signal (l, d').
(5) For each other player $\tilde{i} \neq i, j$ s.t., $\sum\limits_{d \in D_F} \theta(d, \tilde{i}) > 0$: $d_{\tilde{i}}$ is randomly chosen according to θ, conditioned on that the chosen player is $\tilde{i}$. Player $\tilde{i}$ receives the signal $(l+1, d_{\tilde{i}})$.

M'_F is defined as in the previous case. Define now the following strategy x_F^i for each player $i \in I$: let m_i be the signal of player i. Player i stops in stage τ_{m_i}, and continues in all other stages. It is straightforward to see that x_F^i is ϵ-unrevealing w.r.t. $\mathcal{D}_F$, and that x_F^i is a perfect ϵ-equilibrium in $G_F(\mathcal{D}_F)$ conditioned on $\omega \in \Omega \backslash E'$ and given $m \in M \backslash M'$.

References

[1] R.J. Aumann, Subjectivity and Correlation in Randomized Strategies. J. Math. Econ., 1 (1974), 67–96.

[2] R.J. Aumann, Correlated equilibrium as an expression of Bayesian rationality, Econometrica 55 (1987), 1-18.

[3] R.J. Aumann, M. Maschler, Repeated games with incomplete information. The MIT press (1995).

[4] Dhillon A, Mertens JF, Perfect correlated equilibria. Journal of Economic Theory 68 (1996), 279–302.

[5] Dutta P, A Folk Theorem for Stochastic Games. Journal of Economic Theory 66 (1995), 1–32.

[6] E. B. Dynkin, Game variant of a problem on optimal stopping. Soviet Math. Dokl. 10 (1969), 270-274.

[7] C.H. Fine, L. Li, Equilibrium exit in stochastically declining industries. Games Econ. Behav. 1 (1989), 40-59.

[8] Forges, F. (1986). An Approach to Communication Equilibria, Econometrica 54, 1375–1385.

[9] Fudenberg D, Tirole J (1985), Preemption and Rent Equalization in the Adoption of New Technology. Rev. Econ. Stud., LII, 383-401.

[10] Fudenberg D, Tirole J (1986), A Theory of Exit in Duopoly. Econometrica, 54, 943-960.

[11] P. Ghemawat, B. Nalebuff. Exit. RAND J. Econ. 16 (1985),184-194.

[12] J.W. Mamer. Monotone stopping games. J. Appl. Prob. 24 (1987), 386-401.

[13] A. Mashiah. Subgame Perfect Equilibria in Stopping Games. mimeo.

[14] H. Morimoto. Non-zero–sum discrete parameter stochastic games with stopping times. Probab. Theory Related Fields 72 (1986), 155–160.

[15] Myerson, R. B. (1986a), Multistage Games with Communication, Econometrica 54, 323–358.

[16] Myerson R (1986b), Acceptable and predominant correlated equilibria, Int. J. Game Theory 15 (3), 133-154.

[17] Nalebuff, B, Riley JG (1985), Asymmetric Equilibria in the War of Attrition. Journal of Theoretical Biology, 113: 517-27.

[18] P. Neumann, D. RAMSEY, and K. Szajowski. Randomized stopping times in Dynkin games. Z. Angew. Math. Mech. 82 (2002), 811–819.

[19] J. Neveu, Discrete-parameters Martingales. Borth-Holland, Amsterdam (1975).

[20] A.S. Nowak and K. Szajowski. Nonzero-sum stochastic games. In Stochastic and Differential Games (M. Bardi, T. E. S. Raghavan and T. Parthasarathy, eds.) 297–342 (1999). BirkhÃďuser, Boston.

[21] Y. Ohtsubo. A nonzero–sum extension of Dynkin's stopping problem. Math. Oper. Res. 12 (1987), 277–296.

[22] Y. Ohtsubo. On a discrete-time non-zero–sum Dynkin problem with monotonicity. J. Appl. Probab. 28 (1991), 466–472.

[23] F. P. Ramsey, On a problem of formal logic. Proc. London Math. Soc. 30 (1930), 264-286.

[24] D. Rosenberg, E. Solan, N. Vielle, Stopping games with randomized strategies. Probab. Theory Related Fields 119 (2001), 433-451.

[25] R. Selten, Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit. Zeitschrift für die gesamte Staatswissenschaft 121(1965), 301—324.

[26] Selten R, Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory 4 (1975), 25-55.

[27] E. Shmaya, E. Solan, Two-player nonzero-sum stopping games in discrete time, Annals of Probability 32 (2004), 2733-2764.

[28] A. Shmida, B. Peleg, Strict and Symmetric Correlated Equilibria Are the Distributions of the ESS's of Biological Conflicts with Asymmetric Roles, in Understanding Strategic Interaction, ed. by W. Albers, W. GÃijth, P. Hammerstein, B. Moldovanu, and E. van Damme. Berlin: Springer-Verlag (1997), 149–170.

[29] E. Solan, N. Vieille, Quitting games, Mathematics of Operations Research 26 (2001), 265-285.

[30] E. Solan, V. Vohra, Correlated Equilibrium in Quitting Games, Mathematics of Operations Research 26 (2001), 601-610.

[31] E. Solan, V. Vohra, Correlated equilibrium payoffs and public signaling in absorbing games, Int J Game Theory 31 (2002), 91-121.