Perfect correlated equilibria in stopping games

Yuval Heller

School of Mathematical Sciences, Tel-Aviv University

10. June 2009

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 five appealing properties: (1) "Trembling-hand" perfectness - players do not use non-credible threats; (2) Normal-form correlation - communication is required only before the game starts; (3) Uniformness - it is an approximate equilibrium in any long enough finite-horizon game and in any discounted game with high enough discount factor; (4) Universal correlation device -the device does not depend on the specific parameters of the game. (5) Canonical - the signal each player receives is equivalent to the strategy he plays in equilibrium.

1 Introduction

Stopping games have been introduced by Dynkin ([7]) as a generalization of optimal stopping problems, and later used in several models in economics, management science, political science and biology, such as research and development (see e.g., Fudenberg and Tirole [10] and Mamer [14]), struggle of survival among firms in a declining market (see e.g., Fudenberg and Tirole [11], Ghemawat and Nalebuff [12]), auctions (see e.g., Krishna and Morgan [13]), lobbying (see e.g., Bulow and Klemperer [4]), and conflict among animals (see e.g., Nalebuff and Riley [19]).

[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.

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 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 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 stopping games in discrete time. This problem, when the payoffs have a special structure, was studied, among others, by Neveu ([21]), Mamer ([14]), Morimoto ([16]), Ohtsubo ([23]), Nowak and Szajowski ([22]), Rosenberg, Solan and Vieille ([25]), and Neumann, Ramsey and Szajowski ([20]). Those authors provided various sufficient conditions under which (Nash) ϵ-equilibria exist. Recently, Shmaya and Solan ([28]) 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 multi-player 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 [26,27]) in 2-player stopping games (see for example, Fine and Li [8]).

Aumann ([1]) defined the concept of correlated equilibrium in a finite normal-form game as 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., Forges [9]): 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:

- News playing among day traders - An announcement of macroeconomic news is expected at a certain time. Empirical studies (see for example, Christie-David, Chaudhry and KhanEconometrica [5]) show that several minutes elapse before financial instruments adjust to such announcements. This gap of time may provide a chance for substantial profit for quick trading. In this setup, the traders of a financial institution can coordinate their actions in advance. For example, they may decide that if the announcement is of type a and the price of a certain financial instrument increases by more than b in the following 10 minutes, then certain buy and sell orders should be made quickly. On the contrary, coordination along the play is costly due to the time limit. Note that these traders may have different payoff functions: each trader may be interested not only in the firm's profit, but in the part of the profit that is made in financial instruments that are under his responsibility.
- War of attrition in nature, which is commonly modeled as a stopping game, where normal-form (but not extensive-form) correlation devices are implemented by evolution of phenotype roles (see e.g., Shmida and Peleg [29]).

A few papers have defined and studied the properties of perfect correlated equilibria in finite games, see e.g., Myerson ([17,18]) and Dhillon and Mertens ([6]). Generalizing the definition of the last paper, we define a ("trembling-hand") perfect correlated (δ, ϵ)-equilibrium, as a profile where with probability of at-least $1-\delta$, no player can earn more than ϵ by deviating at any stage of the game. [2] We hope that this definition, which has been adapted from may be useful in future study of other dynamic games. [3]

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. Due to the uniformness property, this equilibrium is an approximate equilibrium in any long enough finite-horizon stopping game and in any discounted stopping

[2] More formally, $\delta > 0$ is an upper bound for the probability that the correlation device sends signals in some set M' and for the probability that some event E occurs, and $\epsilon > 0$ 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'.

[3] Our definition is similar to the notion of (sub-game) perfect (δ, ϵ)-equilibrium presented in Mashiah-Yaakovi ([15]), where it is proven that such equilibrium exists in multi-player stopping games where at any stage a single player is allowed to stop.

game with high enough discount factor. [4] Moreover, the correlation device in this equilibrium has two appealing properties: (1) *Universality* - the device depends only on ϵ and on the number of players. (2) *Canonical* - the signal sent to each player is equivalent to the strategy he uses in equilibrium. [5] When the stopping game has special properties, we can further characterize the approximate equilibrium, as discusses in Sect. 8.

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. [6] This reduction requires us to use a universal correlation device that is (δ, ϵ)-*constant-expectation* - 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 ([28]) to further reduce the problem to that of studying the properties of correlated ϵ-equilibria in multi-player absorbing games [7]. The study uses the result of Solan and Vohra [32] that any multi-player absorbing game admits a correlated ϵ-equilibrium.

The paper is arranged as follows. Section 2 presents the model and the result. A sketch of the proof appears in Section 3. In Section 4 we reduce the problem to existence of perfect correlated (δ, ϵ)-equilibrium in terminating games with special properties. Section 5 studies games played on finite trees. 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. In Sect. 8 we discuss special properties of the approximate equilibrium in specific kinds of stopping games.

[4] Arguments in favor of the notion of uniform equilibrium can be found in Aumann and Mashcler ([3]).

[5] In sect. 2 we define a correlation device with a finite signal space, while the the strategy space is infinite. Thus the correlation device is not exactly canonical, but it is closely-related to canonical representation: The signal informs each player at which stages he should stop, conditioned on the information he has on the state of nature and on the history of play, in any subtree where the players play a correlated profile. All of our results remain the same if one would use a canonical correlation device with infinite signal space.

[6] In other papers, both games are referred to as stopping games. We have denoted them by a different name, because the reduction from stopping games to terminating games is not trivial in our setup due to the requirement of normal-form correlation.

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

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 histories of realized actions of each player (stop or continue) until stage n, under the constraint that the members of S always continue. The coordinates of R_n are denoted by R^i_{S,n,h_n^S}, where $n \in \mathbf{N}$, $S \subseteq I$ is the set of players that stop at stage n, $i \in S$ is a player, and $h_n^S \in H_n^S$ is the history of realized actions before stage n.
  - · Let H_∞^S denote the set of all infinite histories of realized actions, in which the members of S always continue and all the members of $I \backslash S$ have stopped. 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, and $h_\infty^S \in H_\infty^S$ is the infinite history of realized actions. Let $n_{h_\infty^S}$ be the last stage in which a player stops in h_∞^S. We require that $R^i_{S,\infty,h_\infty^S}$ is measurable in $\mathcal{F}_{n_{h_\infty^S}}$, i.e., 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 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_n^S}, where $i \in S \subseteq I$ is the set of active players who stop at stage n, and $h_n^S \in H_n^S$ is the history of realized actions until stage n. If player i never stops, his payoff is $R^i_{S,\infty,h_\infty^S}$ 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 $\mathcal{D} = (M, \mu)$: (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.

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

Let θ_i be the stage in which player i stops and let $\theta_i = \infty$ if player i never stops. If $\theta_i < \infty$ let $i \in S_{\theta_i} \subseteq I$ be the coalition that stops at stage θ_i, and if $\theta_i = \infty$ let $i \in S_{\theta_i} \subseteq I$ be the coalition that never stops in the game. Let h_{θ_i} be the history of realized actions until stage θ_i. The expected payoff of player i under the strategy profile $x = (x^i)_{i\in I}$ is given by: $\gamma^i(x) = \mathbf{E}_x\left(R^i_{S_{\theta_i},\theta_i,h_{\theta_i}}\right)$ 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 $E \subseteq \Omega$, let $\gamma^i(x|(E))$ be the expected payoff conditioned on $\Omega \backslash E$: $\gamma^i(x|(E)) = \mathbf{E}_x\left(R^i_{S,\theta_i,h^S_{\theta_i}}|(E)\right)$.

The strategy x^i is ϵ-*best reply* for player i when all his opponents follow x^{-i} if for every strategy of player i, y^i: $\gamma^i(x) \geq \gamma^i(x^{-i}, y^i) - \epsilon$. Similarly, x^i is ϵ-best reply conditioned on E if $\gamma^{\mathbf{i}}(x|(E)) \geq \gamma^{\mathbf{i}}(x^{-i}, y^i|(E)) - \epsilon$. Let H_n denote the set of all histories of realized actions before stage n, and Let $\hat{\mathcal{F}}_n \subseteq \mathcal{F}_n$ denote the 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 $G(h_n, F_n, \mathcal{D})$ be the induced stopping game that begins at stage n after, an history of play h_n has been played, and when the players are informed that $\omega \in F_n \subseteq \hat{\mathcal{F}}_n$. The active players when $G(h_n, F_n, \mathcal{D})$ starts, are those who have not stopped in h_n. For simplicity of notation, we use the same notation for a strategy profile in $G(\mathcal{D})$ and for the induced strategy profile in $G(h_n, F_n, \mathcal{D})$. We now define a perfect correlated (δ, ϵ)-equilibrium.

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, 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 $i \in I$, $m \in M\backslash M'$, $n \in \mathbf{N}$, $h^i_n \in H^i_n$, $F_n \in \hat{\mathcal{F}}_n$ satisfying $p((\Omega\backslash E)|F_n) > 0$, x^i is ϵ-best reply for player $i \in I$ in the induced game $G(h_n, F_n, \mathcal{D})$ 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 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 $\sup_{n \in (\mathbf{N} \bigcup \infty)} \|R_n\|_\infty \in L^1(p)$ (*integrable payoffs*). Then for every $\delta, \epsilon > 0$, G has a prefect correlated (δ, ϵ)-equilibrium.

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 [30, 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 Any multi-player stopping game with integrable payoffs admits a 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 game terminates. 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, the terminal payoff at stage n in $G(k, l)$ 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, it has two special properties: It is recursive (payoff in the non-absorbing state is 0), and there is a unique action profile with a 0 absorbing probability. Solan and Vohra ([32, Prop. 4.10]) proved a classification result for absorbing games. Applying it to the two special properties yields that $G(k, l)$ has one of the following: (1) A stationary absorbing equilibrium. (2) A 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 to construct a correlated ϵ-equilibrium.

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 an ϵ-approximation of the equilibrium payoff. A consequence of RamseyâĂŹs theorem ([24]) is that there is an increasing sequence of integers $0 \leq k_1 < k_2 < ...$ such that $c(k_1,k_2) = c(k_j,k_{j+1})$ for every j.

Assume first that $k_1 = 0$. A perfect correlated 3ϵ-equilibrium is constructed as follows. The construction 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 be a perfect 3ϵ-equilibrium in G in case 1, it is needed to verify that: (1) The equilibrium in each $G(k,l)$ is ϵ-perfect. (2) 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. If the case indicated by $c(k_1,k_2)$ is 3, then we adept the procedure presented by Solan and Vohra for the construction of a correlated ϵ-equilibrium in a quitting game ([31, Section 4.2]). As part of the adaptation we require the correlation device to be universal and (δ,ϵ)-constant-expectation.

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 deterministic 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 existence of (δ,ϵ)-constant-expectation perfect correlated (δ,ϵ)-equilibrium in G' (with a universal correlation device) implies naturally a similar 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 [28] for 2-player stopping games, and by using their stochastic version of Ramsey's theorem.

4 Reductions

In this section we make 3 reductions to the problem of existence of perfect correlated (δ, ϵ)-equilibrium in stopping games: (1) We reduce it to the problem of existence of such equilibrium in terminating games. (2) We further reduce it to the problem of existence of such an equilibrium in tree-like terminating games with a finite-range payoff process. (3) Finally, we reduce it to the problem of existence of such equilibrium in an induced game deep enough in the tree, where with high probability each payoff occurs infinitely often or does not occur at all. Thus, in the following sections we deal only with terminating games with a finite-range payoff process deep enough in the tree.

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 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 $x^i_n : (\Omega \times M^i) \to [0, 1]$. A perfect correlated (δ, ϵ)-equilibrium and a perfect correlated payoff vector are defined in an analog way to Sec. 2. A profile x in $G(\mathcal{D})$ is (δ, ϵ)-constant-expectation if with high probability the expected payoff of a player almost does not change when he obtains his signal.

Definition 11 Let G be a terminating game, $\epsilon > 0$, $\mathcal{D} = (M, \mu)$ a correlation device, and x a profile in $G(\mathcal{D})$. The strategy profile x is (δ, ϵ)-*constant-expectation* if there is a set $M' \subseteq M$ satisfying $\mu(M') \leq \delta$, 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 $\gamma^i(x|m^i)$ is the expected payoff of player i where all players follow x, conditioned on receiving a message m^i.

A (δ,ϵ)-*constant-expectation* profile is similarly defined for stopping games. Consider a function that assigns a correlation device to each stopping game (given ϵ and δ). We say that the assigned correlation device is universal if it depends only on the number of players and ϵ.

Definition 12 Let f be a function that assigns to each stopping (or terminating) game G and to each $\epsilon,\delta > 0$ a correlation device $f(G,\epsilon,\delta) = \mathcal{D}(G,\epsilon,\delta)$. The function f is universal if the assigned correlation device depends only on the number of players and ϵ: $\mathcal{D}(G,\epsilon,\delta) = \mathcal{D}(|I|,\epsilon)$. Given such a function, we call the the assigned device a *universal (correlation) device.*

The following proposition reduces the problem of existence of (δ,ϵ)-constant-expectation perfect correlated (δ,ϵ)-equilibrium (with a universal device) in stopping games to the problem of existence of such equilibrium in terminating games.

Proposition 13 Assume that each terminating game with integrable payoffs admits a (δ,ϵ)-constant-expectation perfect correlated (δ,ϵ)-equilibrium for every $\delta,\epsilon > 0$ with a universal correlation device. Then any stopping game G with integrable payoffs admits such an equilibrium for every $\delta,\epsilon > 0$.

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

universal correlation device.□

4.2 Tree-like stopping game

Definition 14 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 ([28, Sec. 6]) that any 2-player terminating game can be approximated by tree-like terminating games. With minor changes, the proof can be adapted for multi-player terminating games, and for normal-form perfect correlated equilibria. This implies the following lemma (the proof is omitted):

Lemma 15 Assume that each tree-like terminating game with integrable payoffs admits a perfect correlated (δ, ϵ)-equilibrium for every $\delta, \epsilon > 0$. Then any terminating game with integrable payoffs admits such an equilibrium $\forall \delta, \epsilon > 0$.

4.3 The Induced Game $G(F, \mathcal{D})$

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

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) $\forall d \in D$, $p\left(A_d | \bar{A}_d\right) \geq 1 - \frac{\delta}{3\cdot|D|}$

(3) $\forall d \in D$, $p\left(B_{d,N_0} | \bar{B}_d\right) \geq 1 - \frac{\delta}{3\cdot|D|}$

Let $E = \bigcup_{d \in D} \left(\left\{ \omega \in \bar{A}_d | \omega \notin A_d \right\} \cup \left\{ \omega \in \bar{B}_d | \omega \notin B_{d,N_0} \right\} \right)$. Observe that $p(E) < \frac{\delta}{3}$. For any $F \in \mathcal{F}$ let $D_F = \left\{ d \in D | F \in \bar{A}_d \right\}$, and let $\alpha^i_F = \max \left(d^i_{\{i\}} | d \in D_F \right)$.

Let $G(F_n, \mathcal{D})$ be the induced terminating game that begins at stage n when the players are informed that $\omega \in F_n$. The following lemma is standard.

Lemma 16 Let G be a tree-like terminating game, $\delta, \epsilon > 0$, $\mathcal{D} = (M, \mu)$ a correlation device, $M' \subseteq M$ a set of signals s.t. $\mu(M') \le \delta$, τ *a* bounded stopping time, and $E \subseteq \Omega$ an event with $p(E) < \delta$. Assume $\forall F \in \mathcal{F}_\tau$ satisfying $p\left((\Omega \backslash E) \, | F_n \right) > 0$, there is a (δ, ϵ)-constant-expectation perfect correlated (δ, ϵ)-equilibrium x_F of $G(F, \mathcal{D})$ conditioned on $\Omega \backslash E$ and given $M \backslash M'$. Then G admits a perfect correlated $(2 \cdot \delta, 3 \cdot \epsilon)$-equilibrium with a universal device.

PROOF. It is well known that any finite-stage game admits a 0-equilibrium (see, e.g., [25, Prop. 3.1]). Since τ is bounded, $p(E) \le \delta$ and $\mu(M') \le \delta$, the following strategy profile x is a $(2 \cdot \delta, 3 \cdot \epsilon)$-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 $F = F_{\tau(\omega)}(\omega) \in \hat{\mathcal{F}}_\tau$.
- If the game has not terminated by stage τ, play from that stage on the profile x_F in $G(F, \mathcal{D})$.

5 Terminating Games on Finite trees

An important building block in our analysis is terminating games that are played on finite trees. In this section 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 coalition S 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 decide 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 $x^i : V_{Stop} \to [0,1]$; $x^i(v)$ is the probability that player 1 stops at v. Let c^i be the strategy of player i that never stops, and let $c = (c^i)_{i \in I}$. Given a stationary strategy profile $x = (x^i)_{i \in I}$, Let $\gamma^i_T(x) = \gamma^i(x)$ be the expected payoff under x, and let $\pi_T(x) = \pi(x)$ the probability 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, $\gamma^i\left(x\right) > \gamma^i\left(x^{-i}, y^i\right) - \epsilon$.

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 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 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 game on the finite tree $T_{n,\tau}(F)$ is $\left(I, V, V_{leaf}, r, V_{stop}, (C_v, p_v, R_v)_{v \in V \backslash V_{leaf}}\right)$ where:

- $V = \bigcup_{\substack{\omega \in F \\ n \leq k \leq \tau(\omega)}} \{F_k(\omega)\}$, $V_{leaf} = \bigcup_{\omega \in F} \{F_\tau(\omega)\}$, $r = F$, $V_{stop} = \{v \in V | d_v \in D_F\}$
- 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 $p_v(F_{k+1}) = p\left(F_{k+1}|v\right)$.

5.2 Equivalence with Absorbing Games

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. As an absorbing game, the game T has two special properties: (1) It is a recursive game: the payoff in the non-absorbing state is zero; (2) There is a unique action profile that is non-absorbing.

Adapting [32]'s Prop. 4.10 to the two special properties gives the following:

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 stopping 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. [8]

Proposition 21 Let T be a game on a finite tree. 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 perfect 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 $g^i_T, \tilde{v}^i_T, \eta_T, x_T$. The equilibrium in case 1 may not be perfect, as players may use non-credible threats after of-equilibrium path. The following lemma asserts that a perfect ϵ-equilibrium exists.

[8] 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 perturbations on the payoffs ($o(\epsilon)$), such that $R^i_{i,\tilde{v}^i}$ is strictly larger than any other payoff $R^i_{i,v}$ where $v \in V_{stop}$.

Lemma 22 In case 1 of prop. *21,* T 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 a stationary equilibrium. If case 1 of prop. 21 applies, then T_ϵ admits a perfect stationary 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 [28, 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 [28], and the proof is omitted as the changes are minor.

Lemma 23 *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 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 $N_0 \leq n \in \mathbf{N}$. 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) $\alpha^i_F + \epsilon \geq \sum_{j\in I,\, d\in D_F} \theta(d,j) \cdot R^i_{\{j\},d} \geq \alpha^i_F - \epsilon$.

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 (where $n > N_0$ and $F_n \subseteq F_{N_0}$).

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 [28, Sec. 8]. The proof is omitted as the changes compared with [28] are minor.

Proposition 26 Let G be a terminating game, $N_0 \leq n \in \mathbf{N}$, $\sigma > n$ a bounded stopping time, $F \in \hat{\mathcal{F}}_n$ a complicated set (w.r.t. ϵ_0), $\epsilon << \frac{\epsilon_0}{|I|\cdot|D|}$, and for each $i \in I$ let $a^i \geq \alpha^i_F - \epsilon$. 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: $a^i - \epsilon \leq \gamma^i(x)$, $\gamma^i(x^{-i}, y^i) \leq a^i + 8\epsilon$, and $\pi(x) \geq \epsilon^2 \cdot p_{T_U,T}$.

6 The Use of Ramsey Theorem

In this section we use a stochastic variation of Ramsey theorem ([24,28]), 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.

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.

Where A holds on B ($A, B \in \mathcal{F}$) iff $p(A^c \bigcap B) = 0$.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 random variable. Shmaya and Solan proved the following proposition ([28, 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.

A (hyper)-rectangle $\left([a^i, a^i+\epsilon]\right)_{i\in I}$ 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 W be a finite covering of $[-1,1]^{|I|}$ with (not necessarily disjoint) rectangles $\left([a^i, a^i+\epsilon]\right)_{i\in I}$, all of which are either good or bad. Let $B = \{b_1, b_2, ..., b_J\}$ be the set of bad rectangles in W and let $O = \{o_1, o_2, ..., o_K\}$ the set of good rectangles.

Set $C = (simple \bigcup allbad \bigcup \{1 \times O\} \bigcup \{2\} \bigcup \{3 \times W \times W\})$. 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 simple we let $c_{n,\sigma}(F) = simple$. Otherwise, F is *complicated* w.r.t. to some $\epsilon_0(F)$. In that case we assume that from now we fix ε on that $0 < \epsilon << \min_{F\in\hat{\mathcal{F}}_{N_0}} \frac{\epsilon_0(F)}{|I|\cdot|D|}$. The color $c_{n,\sigma}(F)$ is determined as follows: [9]

- Set $T^{(0)} = T_{n,\sigma}(F)$.
- For $1 \leq j \leq J$ apply Prop. 21 to $T^{(j-1)}$ and the bad rectangle $h_j = \prod_{i\in I}\left[a^i_j, a^i_j+\epsilon\right]$ 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 $i \in I$, $a^i_j - \epsilon \leq \gamma^i(x_j)$.
  (b) For every $i \in I$ and every strategy y^i: $\gamma^i(x^{-i}_j, y^i) \leq a^i_j + 8\epsilon$.
  (c) $\pi\left(x_j\right) \geq \epsilon^2 \times p_{T^{(j)},T^{(j-1)}}$.

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

(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

[9] The procedure is an adaptation of the 2-player procedure described in [28, Sec. 5]

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 $c_{n,\sigma}(F) = (3, w_1, w_2)$ where w_1 is the hyper-rectangle that includes $\gamma_{T^{(J)}}(\eta)$, and w_2 is the hyper-rectangle that includes $g(T^{(J)})$.

The strategy profiles 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 define, for every $j \in J$, $\lambda_{j,n,\sigma}(F) = p_{T^{(j)},T^{(j-1)}}$.

By Prop. 29 there exists an increasing sequence of bounded stopping times $0 < \sigma_1 < \sigma_2 < \sigma_3 < ...$ such that: $p\left(c_{\sigma_1,\sigma_2} = c_{\sigma_2,\sigma_3} = ...\right) > 1 - \frac{\delta}{3}$. For every $F \in \hat{\mathcal{F}}_{\sigma_1}$, let $c_F = c_{\sigma_1,\sigma_2}(F)$.

Let $\left(A_{\epsilon,j}, A_{\infty,j}\right)_{j\in J} \in \bigvee_{n=1..\infty} \mathcal{F}_n$ be: $A_{\infty,j} = \left\{w \in \Omega | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}}\left(F_{\sigma_k(\omega)}\right) = \infty\right\}$, $A_{\epsilon,j} = \left\{w \in \Omega | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}}\left(F_{\sigma_k(\omega)}\right) \leq \frac{\epsilon}{|J|}\right\}$. As $\left(A_{\epsilon,j}, A_{\infty,j}\right)_{j\in J} \in \bigvee_{n=1..\infty} \mathcal{F}_n$, there is large enough $N_1 \geq N_0$ and sets $\left(\bar{A}_{\epsilon,j}, \bar{A}_{\infty,j}\right)_{j\in J} \in \mathcal{F}_{N_1}$ s.t.: (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) $p\left(A_{\epsilon,j} | \bar{A}_{\epsilon,j}\right) \geq 1 - \frac{\delta}{6\cdot|J|}$. (3) $p\left(A_{\infty,j} | \bar{A}_{\infty,j}\right) \geq 1 - \frac{\delta}{6\cdot|J|}$. From now on, we assume w.l.o.g. that $\sigma_1 \geq N_1$. Let E' be defined as follows (Observe that $p(E') \leq \delta$):

$$
\begin{aligned}
E' = E \bigcup_{j\in J} & \left\{\omega \in \bar{A}_{\epsilon,j} | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}}\left(F_{\sigma_k(\omega)}\right) > \frac{\epsilon}{|J|}\right\} \\
& \left\{\omega \in \bar{A}_{\infty,j} | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}}\left(F_{\sigma_k(\omega)}\right) < \infty\right\} \\
& \bigcup \left\{\omega \in \Omega | \exists n \, s.t. \, c_{\sigma_n,\sigma_{n+1}}(\omega) \neq c_{1,2}(\omega)\right\}
\end{aligned}
$$

7 Approximate Constant-Expectation Perfect Correlated Equilibrium

We finish the proof of the main theorem by the following proposition:

Proposition 30 *Let G be a tree-like terminating game, let $\epsilon, \delta > 0$, let the event $E' \subseteq \Omega$ and σ_1 be defined as in the previous subsection, and let $F \in \hat{\mathcal{F}}_{\sigma_1}$. Then there is a universal correlation device $\mathcal{D} = (M, \mu)$ and a strategy profile x_F in the game $G(F, \mathcal{D})$, such that x_F is a perfect (δ, ϵ)-constant-expectation ϵ-equilibrium in the game $G(F, \mathcal{D})$ 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 whether $F \in \bar{A}_{\infty,j}$. The first 3 cases adapts the methods of [28, 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 $G(F, \mathcal{D})$: 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 $T_{\sigma_k,\sigma_{k+1}}$: (1) $\gamma^i_{\sigma_k,\sigma_{k+1}}(x^{-i}_j, y^i) \leq a^i_j + 8\epsilon$. (2) $\pi_{\sigma_k,\sigma_{k+1}}(x_j) \geq \epsilon^2 \times \lambda_j(T_{\sigma_k,\sigma_{k+1}})$

Those facts that outside E' the game is absorbed with probability 1, and that x_F is a 11ϵ-equilibrium conditioned on $\Omega \backslash E'$. 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 at any stage, conditioned on $\Omega \backslash E'$.

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

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: $\prod_{i \in I} [a^i_w, a^i_w + \epsilon]$. 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 $\sum\limits_{j=1..J} \sum\limits_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}} < \epsilon$. 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 $F \in \bigcap_{j\in J} \bar{A}_{\epsilon,j}$ *and* $c_F = (1, m_w, m_{w'}) \in (1 \times W \times W)$

The construction in this case is as an adaptation of the procedure of [31], which deals with quitting games (stationary terminating games where payoff is the same at all 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 $\alpha^i_F = g^i(T_{\sigma_1,\sigma_2}) \in w^i_{k'}$. 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) $\theta(d,i) > 0 \Rightarrow R^i_{i,d} = \alpha^i_F$, $\forall d' \neq d \in D_F$, $\theta(d',i) = 0$. Let $d(i) \in D_F$ be the payoff satisfying $\theta(d_i,i) > 0$. 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 $m > \max_{i\in I}(\tau^i_n)$ such that $R_m = d(i_0)$. Observe that in $\Omega\backslash E'$ each $\tau^i_n < \infty$. We now describe the correlation device $\mathcal{D}_{D_F} = (M_{D_F}, \mu_{D_F})$. Let $M^i_{D_F} = \{1, ..., \hat{T} + T + 1\}$, where $T \in \mathbf{N}$ is sufficiently large, and $\hat{T} >> T$. Let μ_{D_F} 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 chosen according to ζ. Player i receives signal $\hat{l}$.
(3) A number $l \in \mathbf{N}$ is chosen uniformly over $\left\{\hat{l}+1, \hat{l}+T\right\}$
(4) Player j_i, the punisher of player i, receives the signal l.
(5) Each other player $\tilde{i} \neq i, j$ receives the signal $l + 1$.

Let $M'_{D_F} \subseteq M_{D_F}$ 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 $\mu(M'_{D_F}) \leq \frac{\delta}{2^D}$. 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 at stages τ_n that satisfy: $n = (m_i) \bmod \hat{T} + T + 1$,[10] and continues in all other stages. Let the universal correlation device $\mathcal{D} = (M, \mu)$ be the Cartesian multiplication: $\mathcal{D} = \prod_{D_F\subseteq D} \mathcal{D}_{D_F}$. Similarly let $M' = \prod_{D_F\subseteq D} M'_{D_F}$. Observe that $\mu(M') \leq \delta$.

If the players follow the strategy profile x_F then the game is absorbed with probability 1 conditioned on $\Omega\backslash E'$ 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

[10] On equilibrium path the player stops at stage τ_n. The requirement to stop at later stages where $n = (m_i) \bmod \hat{T}+T+1$ is needed to satisfy the perfection requirement.

or not he is the quitter, thus x_F is (δ, ϵ)-constant-expectation.

We now verify that if $T, \hat{T}$ are sufficiently large, no player can gain too much by deviating at any stage of the game 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 quitter, punisher or a "regular" player, and 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. When player j has to stop, he believes that he is the quitter (assuming $m \in M \backslash M'$). This implies that the players ϵ-best-respond at all stages including while (unknowingly) punishing, and that x_F is a perfect ϵ-equilibrium in $G(F, \mathcal{D})$ conditioned on $\omega \in \Omega \backslash E'$ and given $m \in M \backslash M'$.

7.5 $c_F = simple$

If for every $i \in I$: $\alpha^i_F \leq 0$, 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 $i \in I$: (1) $\theta(d, i) > 0 \Rightarrow R^i_{\{i\},d} = \alpha^i_F$. (2) $\alpha^i_F + \epsilon \geq \sum_{j \in I,\, d \in D_F} \theta(d, j) \cdot R^i_{\{j\},d} \geq \alpha^i_F - \epsilon$. In this case, one can use a procedure similar to the one described in the previous subsection, to construct a perfect ϵ-equilibrium in $G(F, \mathcal{D})$ conditioned on $\omega \in \Omega \backslash E'$ and given $m \in M \backslash M'$.

8 Equilibrium's Special Properties in Specific Cases

In this Section we present a few examples of specific kinds of stopping game, with applicative interest, and we shortly discuss the special properties of the perfect correlated (δ, ϵ)-equilibrium in those cases.

(1) *Symmetric* stopping games: Stopping games where the payoff process is the same for all players. That is:$\forall i, j \in I$, $S \subseteq I$, $n \in \mathbf{N}$, $h^S_n \in H^S_n$, $R^i_{S,n,h^S_n} = R^j_{S_{i \leftrightarrow j},n,h_n^{S_{i \leftrightarrow j}}}$, where $S_{i \leftrightarrow j}$ is the coalition derived from S by substituting players i and j). As can be seen from the construction in Sect. 7, such games admit a symmetric perfect correlated (δ, ϵ)-equilibrium.

(2) *Eventual continuation* Games: Stopping games where late enough in the

game each player would rather continue forever than stop alone. That is, $\exists N_0 \in \mathbf{N}$, such that $\forall i \in I$, $i \in S \subseteq I$, $n > N_0$, $h_n^i \in H_n^i$, $h_\infty^S \in H_\infty^S$ $R^i_{i,n,h^i_n} \leq R^i_{S,\infty,h^S_\infty}$. Such games admit a perfect Nash-equilibrium: Players play the perfect Nash equilibrium of the finite game that ends at N_0 with an absorbing payoff of $R^i_{S,\infty,h^S_\infty}$ (where S is the set of players who have not stopped until stage N_0), and continue forever afterwords.

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] J. Bulow, P. Klemperer, The Generalized War of Attrition, Amer. Econ. Rev. 89:1(. 2001), 175-189.

[5] R. Christie-David, M. Chaudhry and W. Khan, News releases, market integration, and market leadership, The Journal of Financial Research XXV(2) (2002), 223-245.

[6] A. Dhillon, J.F. Mertens, Perfect correlated equilibria. Journal of Economic Theory 68 (1996), 279-302.

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

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

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

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

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

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

[13] V. Krishna, J. Morgan, An Analysis of the War of Attrition and the All-Pay Auction. Journal of Economic Theory, 72 (1997), 343Âś62.

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

[15] A. Mashiah-Yaakovi. Subgame Perfect Equilibria in Stopping Games. mimeo (2008).

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

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

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

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

[20] P. Neumann, D. Ramsey, and K. Szajowski. Randomized stopping times in Dynkin games. Z. Angew. Math. Mech. 82 (2002), 811âĂŞ819.

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

[22] 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.

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

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

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

[26] R. Selten, Spieltheoretische behandlung eines oligopolmodells mit nachfragetr a gheit. Zeitschrift fur die gesamte Staatswissenschaft 121(1965), 301-324.

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

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

[29] 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, E. van Damme. Springer-Verlag (1997), 149-170.

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

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

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