Perfect correlated equilibria in stopping games

Heller, Yuval

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 define a new solution concept for an undiscounted dynamic game - a perfect uniform normal-form constant-expectation correlated approximate equilibrium with a canonical and universal correlation device. This equilibrium has the following appealing properties: (1) "Trembling-hand" perfectness - players do not use non-credible threats; (2) Uniformness - it is an approximate equilibrium in any long enough finite-horizon game and in any discounted game with a high enough discount factor; (3) Normal-form correlation - The strategy of a player depends on a private signal he receives before the game starts (which can be induced by "cheap-talk" among the players); (4) Constant expectation - The expected payoff of each player almost does not change when he receives his signal; (5) Universal correlation device - the device does not depend on the specific parameters of the game. (6) Canonical - each signal is equivalent to a strategy. We demonstrate the use of this equilibrium by proving its existence in every undiscounted multi-player stopping game.

Key words: Keywords: stochastic games, stopping games, perfect correlated equilibrium, distribution equilibrium, Ramsey Theorem. JEL classification: C73

1 Introduction

Consider the following example of strategic interaction in the financial markets:

[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.

Example 1 *The Bureau of Labor Statistics publishes each month a news release on U.S. employment situation (ES). This news release is announced in the middle of the trading day in the European stock markets.*[2] *The ES announcement has strong impact on these markets (see Nikkinen et al. [25] and the references within). Empirical studies (see for example, Christie-David, Chaudhry and Khan [6]) show that a few dozen minutes elapse before financial instruments adjust to such announcements. This gap of time (the "adjustment period") may provide an opportunity for substantial profit by quick trading ("news-playing"). Consider the strategic interaction between a few traders of a financial institution that coordinate their trading actions in the adjustment period. Each trader is responsible for some financial instruments, and can make buy and sell orders for these instruments during the adjustment period. The traders share a common objective - maximizing the profit of the institution. In addition to this, each trader has also a private objective - maximizing the profit that is made in financial instruments that are under his responsibility (which influence his bonuses and prestige).*

Three natural questions arise when modeling the strategic interaction among the traders in this example: (1) Which kind of game should be used? (2) Which solution concept should be chosen? (3) Is it possible to prove the existence of this solution in this game? which simplifying assumptions are needed?

We begin by dealing with the first question. The adjustment period is relatively short in absolute terms - a few dozen minutes. Nevertheless, the traders have many opportunities to act, as they can make different orders in each fraction of a second. In addition, the point in time where the markets have fully adjusted may not be known to the players in real-time. Thus, it seems more appropriate to model this situation as a stochastic (dynamic) game with *infinite-horizon*, rather then modeling it as a game with a fixed finite large number of stages (see Rubinstein [30, Sect. 5], for a discussion why even short strategic interactions may be better analyzed as infinite-horizon games).

Each buy order induces a single-stage profit (or loss) that depends on the difference between the value of the bought amount at the buying time and at the end of the adjustment period, and similarly for sell orders. The total payoff of each trader depends on his own profit (the sum of his single-stage profits) and on the profit of all the traders in the company. As all profits and losses are accumulated in the same day, it is natural to assume that these sums are *undiscounted*: there is no difference between earning a dollar at early or late stages of the game.

Due to the above arguments we model the interaction in example 1 as an undiscounted infinite-horizon stochastic game. We now deal with the second

[2] It is published on the first Friday of each month at 13:30 London time (8:30 Eastern Time).

question: which solution concept is appropriate for this game. In order to avoid the use of equilibria that rely on "non-credible" threats of punishment, we require the solution concept to satisfy the standard requirement of "trembling-hand" perfection (Selten, [31]). [3]

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 a private signal to each player before the start of play. The strategy of each player can then depend on the private signal that he received. [4] It is well known that a correlation device can be induced by pre-play non-biding communication among the players ("cheap-talk"). [5]

For sequential games, two main versions of correlated equilibrium have been studied (see e.g., Forges [10]): normal-form correlated equilibrium, where each player receives a private signal only before the game starts, and extensive-form correlated equilibrium, where each player receives a private signal at each stage of the game. In example 1, the traders can freely communicate and coordinate their future actions before the play starts (that is, before the adjustment period begins). On the other hand, communication and coordination along the play are very costly: the adjustment period is short (a few dozen minutes), and each moment that is spent on communication may slow down the traders and limit their potential profits. Thus, the smaller set of normal-form correlated equilibria is more appropriate to example 1. [6]

The above arguments limit the plausible outcomes of the game to the set of perfect normal-form correlated equilibria. A few papers defined and studied the properties of perfect correlated equilibria in finite games, see e.g., Myerson ([20,21]) and Dhillon and Mertens ([7]). As infinite undiscounted games may only admit approximate equilibria, we generalize the definition of the last

[3] In the rest of this paper we use the shorter term "perfection" to denote "trembling-hand" perfection.

[4] Correlated equilibria in finite games 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.

[5] Ben-Porath ([4]) shows that if there are at least three players and if there are two distinct Nash equilibrium payoffs for each player, then each correlated equilibrium distribution can be implemented as a perfect Nash equilibrium of an extended game with pre-play "cheap-talk". Urbano and Villa ([38]) demonstrate the use of cryptographic methods to achieve similar implementation when there are only two players. Heller ([14]) shows that the implementation can be done in way that is also resistant to joint deviations of coalitions.

[6] Note that every normal-form correlated equilibrium is an extensive-form correlated equilibrium, but the converse is not true.

paper, and define a *perfect correlated* (δ, ϵ)*-equilibrium*, as a strategy profile where with probability of at-least $1-\delta$, no player can earn more than ϵ by deviating at any stage of the game and after any history of play (as formally defined in Section 2).

We further impose four more requirements from the solution concept:

- *Uniformness* - An equilibrium of an undiscounted game is uniform if it is an approximate equilibrium in any long enough finite-horizon stopping game and in any discounted stopping game with a high enough discount factor. [7]
- *Canonical correlation device* - each signal is equivalent to a strategy.
- *Universal correlation device* - The correlation device does not depend on the specific payoff functions of the players (it only depends on the number of players and on ϵ). This property allows the traders to use the same correlation device for all news-playing interactions, rather than devising a new correlation device each time.
- *Approximate constant-expectation* equilibrium - Sorin ([37]) defines a distribution equilibrium in a normal-form finite game, as a correlated equilibrium where the expected payoff of each player is independent of his signal. Without this property the implementation of a correlated equilibrium by pre-play communication is much more complex: a trader who receives a bad signal that induces a low payoff, may not cooperate in the rest of the communication, and this may interfere with the construction of the correlated profile. We generalize Sorin's definition and define a (δ, ϵ)*-constant-expectation correlated equilibrium,* as a correlated equilibrium where the expected payoff of a player almost does not change when he receives his signal.

The first contribution of this paper is the presentation of this new solution concept for undiscounted dynamic games: a perfect uniform normal-form constant-expectation correlated approximate equilibrium with a canonical and universal correlation device. Our second contribution is demonstrating its use in a specific family of dynamic games (stopping games, described below). We hope that this concept will be useful in future study of other dynamic games. [8]

We now deal with the third question: proving the existence of this equilibrium. We prove it under two simplifying assumptions on the strategic interaction. Each trader in example 1 may act several times during the adjustment phase. Specifically, for each financial instrument under his responsibility, the trader chooses a stage to buy and a stage to sell. Our first simplifying assumption assumes that each trader only acts once: he chooses a single stage in the entire game where he "stops" - makes a single buy or sell order. The second

[7] See Aumann and Mashcler ([3]) for arguments in favor of this notion.

[8] Our definition is analog to the notion of sub-game perfect (δ, ϵ)-equilibrium presented in Mashiah-Yaakovi ([17]), where it is proven that such equilibrium exists in multi-player stopping games where at any stage a single player is allowed to stop.

simplifying assumption requires that throughout the game, the traders have symmetric information on the financial markets (such as past prices of the different markets).

Our model of the strategic interaction and the solution concept, are also appropriate in other situations, as demonstrated in the following examples.

Example 2 *A few countries plan to ally together in a war against another country. The allying countries share a common objective - maximizing the military success against the common enemy. In addition to this, each country has private objectives, such as maximizing the territories and resources it occupies during the war, and minimizing its losses.*

Example 3 *A few male animals compete over the relative positions they shall occupy in the social hierarchy or pack order. This competition is often settled by "a war of attrition" (Maynard Smith, [18]). In such a contest the animals use "ritualized" fighting and do not seriously injure the opponents. The winner is the contestant who continues the war for the longest time. Excessive persistence has the disadvantages of waste of time and energy in the contest.*

Examples 2 and 3 share similar properties to example 1. These similarities make our model and solution concept appropriate to these examples as well:

- In both examples, the war is relatively short in absolute time (a modern war typically lasts a couple of weeks; a war of attrition usually lasts a few hours or days). Nevertheless its length is not bounded, and it consists a large unknown number of stages. Thus, it seems appropriate to model these situations as undiscounted infinite-horizon stochastic games.
- Normal-form correlation is appropriate to both examples. The country Representatives in example 2 can communicate and coordinate their future actions before the war begins. On the other hand, secure communication and coordination during the war is costly and noisy. Shmida and Peleg ([33]) discuss how a normal-form (but not extensive-form) correlation device can be induced in nature by phenotypic conditional behavior. [9]
- Each country in example 2 does many actions in the battlefield, but usually only a few of them are crucial to the outcome of the war. A simplifying model may concentrate on the most important action of each country, such as the timing of the main military attack. The only choice of strategy of each animal in example 3, is the maximal period for which he is prepared to continue in the contest (and this period may depend on the set of animals

[9] They present an example ([33, Section 5]) of butterflies who compete for sunspot clearings in a forest in order to fertilize females. When two butterflies meet in a sunspot, they engage in a war of attrition. The period of time each butterfly was in the spot before the fighting, is used as a normal-form correlation device: a "senior" butterfly stays for a long time in the war, while a "new" butterfly gives up early.

who are still competing).
- The constant-expectation requirement defined earlier is specially appealing in a biological setup as in example 3. As discussed in [37], constant-expectation is a necessary requirement for the stability of the population in evolutionary setups. [10]

Under the two simplifying assumptions mentioned earlier, the strategic interaction is a *stopping game*. An undiscounted discrete stopping 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 actions while he has been active and on the state variable.

Stopping games have been introduced by Dynkin ([8]) 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 [11] and Mamer [16]), struggle of survival among firms in a declining market (see e.g., Fudenberg and Tirole [12], Ghemawat and Nalebuff [13], and Fine and Li [9]), auctions (see e.g., Krishna and Morgan [15]), lobbying (see e.g., Bulow and Klemperer [5]), and conflict among animals (see e.g., Nalebuff and Riley [22]).

Much work has been devoted to the study of undiscounted 2-player stopping games. This problem, when the payoffs have a special structure, was studied, among others, by Neveu ([24]), Mamer ([16]), Morimoto ([19]), Ohtsubo ([27]), Nowak and Szajowski ([26]), Rosenberg, Solan and Vieille ([29]), and Neumann, Ramsey and Szajowski ([23]). Those authors provided various sufficient conditions under which (Nash) ϵ-equilibria exist. Recently, Shmaya and Solan ([32]) 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.

Our main result shows that for every $\delta, \epsilon > 0$, a multi-player stopping game admits a perfect normal-form constant-expectation correlated (δ, ϵ)-equilibrium with a canonical and universal correlation device. The proof relies on using stochastic variation of Ramsey's theorem ([32]) to reduce the problem to that

[10] This is demonstrated in [37, example 1]. Consider a symmetric two-player game where the payoff (fitness) is 1 if both players play A, 2 if both play B and 0 otherwise. Consider a correlated equilibrium in some population: half of the population are type A - they always play against other A-s and they play action A; the other half are type B- they always play against other B-s and they play action B. This equilibrium does not satisfy the constane-expectation property, and it is not atable in an evolutionary setup: type B gets a higher fitness and they would invade the whole population.

of studying the properties of correlated ϵ-equilibria in multi-player absorbing games.[11] The study uses the result of Solan and Vohra [36] 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 induced games "deep enough" in the tree. 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. The formal model in Sections 2-7 deals only with stopping games that terminate as soon as any of the players stop. In Section 8 we discuss how to apply our result for more general stopping games.

2 Model and Main Result

Definition 4 *A* stopping 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 stopping game is played as follows. At each stage n, each player is informed which elements of $\mathcal{F}_n$ include ω (the state of the world), 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.

Remark 5 *Definition 4 describes games that end as soon as one of the players stop. In Section 8 we discuss how to extend our results to more generalized stopping games where a player that stops gets his payoff and becomes passive for the rest of the game, and the game continues with the other players.*

Definition 6 defines a correlation device:

Definition 6 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, and *(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.

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

Given a correlation device $\mathcal{D}$, we define an extended game $G(\mathcal{D})$. The game $G(\mathcal{D})$ is played exactly as G, except that before the game starts, a signal combination $m = (m^i)_{i \in I}$ is drawn according to μ, and each player i is privately 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 $x^i_n : (\Omega \times M^i) \to [0,1]$. The interpretation is that $x^i_n(\omega, m^i)$ is the probability by which player i stops at stage 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, each player has a positive probability to stop and a positive probability to continue. Formally: for each $i \in I$, $m^i \in M^i$, $n \in \mathbf{N}$: $0 < x^i_n(\omega, m^i) < 1$.

Consider a function that assigns a correlation device to each stopping game, given some positive values of δ and ϵ. We say that the assigned correlation device is *universal* if it depends only on the number of players and ϵ.

Definition 7 Let f be a function that assigns to each stopping 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 assigned device a universal correlation device.

A correlation device $\mathcal{D} = (M, \mu)$ is *canonical* if each signal $m^i \in M^i$ is equivalent to a strategy of player i.

Definition 8 Let G be a stopping game. A correlation device $\mathcal{D} = (M, \mu)$ is *canonical* given the strategy profile x in $G(\mathcal{D})$ if , for each player i there is an injection between M^i and his set of strategies in G. That is $x(m^i) \neq x(m'^i)$ for each $m^i \neq m'^i$. [12]

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 $S_{\theta_i} \subseteq I$ be the coalition that stops at stage θ_i, and if $\theta_i = \infty$ let $S_{\theta_i} \subseteq I$ be the coalition of players who never stops in the game. The expected payoff of player i under the strategy profile $x = (x^i)_{i \in I}$ is given by $\gamma^i(x) = \mathbf{E}_x \left(R^i_{S_{\theta_i}, \theta_i} \right)$ 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 E: $\gamma^i(x|E) = \mathbf{E}_x \left(R^i_{S, \theta_i} | E \right)$, and let it arbitrarily equal to 0 when $p(E) = 0$.

[12] The standard definition of canonical correlation device for finite games, is that the set of signals is equal to the set of strategy profiles. Our definition is somewhat different because the set of signals is finite, while the set of strategies is infinite.

Given a correlation device $\mathcal{D}$ and $\delta, \epsilon > 0$, we say that a profile x has (δ, ϵ)-constant-expectation if with high probability the expected payoff of a player almost does not change when he obtains his signal.[13]

Definition 9 Let G be a stopping game, $\epsilon, \delta > 0$, $\mathcal{D} = (M, \mu)$ a correlation device. The strategy profile x in $G(\mathcal{D})$ is (δ, ϵ)-*constant-expectation* if there is a set $M' \subseteq M$ satisfying $\mu(M') > 1 - \delta$, such that for every player $i \in I$ and every signal $m^i \in (M')^i$: $|\gamma^i(x|m^i) - \gamma^i(x)| \leq \epsilon$, where $\gamma^i(x|m^i)$ is the expected payoff of player i where all players follow x, conditioned on that player i received a signal m^i.

The strategy x^i is ϵ-*best reply* for player i when all his opponents follow x^{-i} if for every strategy y^i of player i: $\gamma^i(x) \geq \gamma^i(x^{-i}, y^i) - \epsilon$. Similarly, x^i is ϵ-best reply conditioned on E if $\gamma^i(x|E) \geq \gamma^{\mathbf{i}}(x^{-i}, y^i|E) - \epsilon$.

Given $\omega \in \Omega$ let $H_n(\omega) \subseteq \mathcal{F}_n$ be the elements $\mathcal{F}_n$ that include ω: $H_n(\omega) = \{F_n \in \mathcal{F}_n | \omega \in F_n\}$, and let $\mathcal{H}_n$ the set of all such sets: $\mathcal{H}_n = \{H_n(\omega) | \omega \in \Omega\}$. Let $G(H_n, \mathcal{D})$ be the induced stopping game that begins at stage n, when the players are informed of $H_n \in \mathcal{H}_n$ (i.e, they are informed that the elements of $\mathcal{F}_n$ that include ω are the elements of 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 $G(H_n, \mathcal{D})$. We now define a few auxiliary definitions that are used to define a perfect correlated (δ, ϵ)-equilibrium. The definition extends [7]'s definition of perfect correlated equilibrium in normal-form finite games.

Definition 10 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 E and given M', if* there exists a sequence $(y_k)_{k \in \mathbf{N}} = (y_k^i)_{k \in \mathbf{N}, i \in I}$ 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 $i \in I$, $n \in \mathbf{N}$, $H_n \in \mathcal{H}_n$, x^i is ϵ-best reply for player $i \in I$ in the induced game $G(H_n, \mathcal{D})$ when all his opponents $j \in -i$ use $(1 - \epsilon_k) x^j + \epsilon_k y_k^j$, conditioned on E and given that the signal profile is included in M'.

That is, x *is a* perfect ϵ-equilibrium conditioned on E and given M', if it is a limit of completely mixed profiles y_k, such that for each player i, x^i is *is ϵ-best reply for y_k^{-i} whenever the state ω is in E and the signal profile is in M'.*

Remark 11 *In the setup of stopping games, the history up to stage n only includes the symmetric information on ω, which is given by $H_n \in \mathcal{H}_n$. In a more general stochastic game, Def. 10 would remain the same, except that $\mathcal{H}_n$ should be modified to denote the set of all possible histories of length n.*

[13] This generalizes [37]'s definition of distribution equilibrium for finite normal-form games, which was discussed in Section 1.

A profile is *a* perfect (δ, ϵ)-equilibrium if it is an ϵ-equilibrium conditioned on E and given M', where E and M' have probabilities of at least $1-\delta$.

Definition 12 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 $p(E) > 1-\delta$, $\mu(M') > 1-\delta$, and x is a perfect ϵ-equilibrium of $G(\mathcal{D})$ conditioned on E and given M'.

Finally, we define a perfect correlated (δ, ϵ)-equilibrium.

Definition 13 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 14 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 perfect normal-form (δ, ϵ)-constant-expectation correlated (δ, ϵ)-equilibrium with a canonical and universal correlation device.

Remark 15 The perfect correlated (δ, ϵ)-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. 29) or by applying a general observation made by [34, Prop. 2.13].

3 Sketch of the Proof

In this section we provide the main ideas of the proof. Let G be a stopping game. To simplify the presentation, assume that $\mathcal{F}_n$ is trivial for every n, so that the payoff process is deterministic, 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 stages corresponds to a single stage of the absorbing game.[14] 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 ([36, Prop. 4.10]) proved a classification result for absorbing

[14] Recall that an absorbing game is a stochastic game with a single non-absorbing state.

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. This distribution has special properties that allow to construct a correlated ϵ-equilibrium.

Assign to each pair of non-negative integers $k < l$ an element from a finite set of *colors* $c(k,l)$; the color is a couple where the first element, which is 1, 2 or 3, denotes which case of the classification result holds in $G(k,l)$, and the second element is a vector in a dense subset of $[-1,1]^n$ that approximates the equilibrium payoff in $G(k,l)$. A consequence of Ramsey's theorem ([28]) 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, then between stages k_j and k_{j+1} the players follow a periodic 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, 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 2, then always continuing is an equilibrium. If the case indicated by $c(k_1,k_2)$ is 3, then we adapt the procedure presented by Solan and Vohra for the construction of a correlated ϵ-equilibrium in a quitting game ([35, Section 4.2]). The adaptation is required to allow the construction of a perfect (δ,ϵ)-equilibrium, despite the use of punishments in the procedure.

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 induced 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 3ϵ-equilibrium.

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 [32] for two-player stopping games, and by using their stochastic version of Ramsey's theorem.

4 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 stopping game $G=(I,\Omega,\mathcal{A},p,\mathcal{F},R)$ is a $(\delta,3\epsilon)$-equilibrium in the stopping 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 stopping game G. Let $R_n(\omega)$ be the payoff matrix at stage n.

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 $E=\bigcup_{d\in D}\left(\left\{\omega\in\bar{A}_d|\omega\notin A_d\right\}\bigcup\left\{\omega\in\bar{B}_d|\omega\notin B_{d,N_0}\right\}\right)$. 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 $\alpha^i_F=\max\left(d^i_{\{i\}}|d\in D_F\right)$. That is D_F is the sets of payoff matrices that repeat infinitely often in F, and α^i_F is the maximal payoff a player can get by stopping alone in these matrices.

The following standard lemma shows that it is enough to show that every induced game $G\left(H_n,\mathcal{D}\right)$ deep enough in the tree $(n>N_0)$ has an approximate constant-expectation perfect correlated equilibrium.

Lemma 16 Let G be a stopping game, $\delta,\epsilon>0$, $\mathcal{D}=(M,\mu)$ a correlation device, $M'\subseteq M$ a set of signals such that $\mu(M')>1-\delta$, τ *a* bounded stopping time, and $E\subseteq\Omega$ an event such that $p(E)>1-\delta$. Assume that for every $\omega\in\Omega$ and for every $H=H_{\tau(\omega)}\in\mathcal{H}_{\tau(\omega)}$, there is a constant-expectation perfect (δ,ϵ)-equilibrium x_H of $G(H,\mathcal{D})$ conditioned on E and given M'. Then $G\left(\mathcal{D}\right)$ admits a constant-expectation perfect $(2\delta,2\epsilon)$-equilibrium.

PROOF. It is well known that any finite-stage game admits a 0-equilibrium (see, e.g., [29, Prop. 3.1]). Since τ is bounded, $p(E)\geq 1-\delta$ and $\mu(M')\geq 1-\delta$, the following strategy profile x is a $(2\delta,2\epsilon)$-equilibrium in $G(\mathcal{D})$:

- Until stage τ, play an equilibrium in the game that terminates at τ, if no player stops before that stage, with a terminal payoff $\gamma^i(x_H)$.

- If the game has not terminated by stage τ, play from that stage on the profile x_H in $G(H, \mathcal{D})$.

5 stopping Games on Finite trees

An important building block in our analysis is stopping 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 stopping game on a finite tree* (or simply a *game on a tree*) is a tuple $T = \left(I, V, V_{leaf}, r, V_{stop}, (C_v, p_v, R_v)_{v \in V \setminus V_{leaf}}\right)$, where:

- I is a finite non-empty set of players.
- $\left(V, V_{leaf}, r, (C_v)_{v \in V \setminus V_{leaf}}\right)$ 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 $v \in V \setminus V_{leaf}$, $C_v \subseteq V \setminus \{r\}$ is the nonempty set of children of v. We denote by $V_0 = V \setminus V_{leaf}$ 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 stopping 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 $v \in V \setminus V_{stop}$ 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. A game on a 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 $\gamma_T^i(x) = \gamma^i(x)$ 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). 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 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.

5.2 Representative Finite Approximations

In the following subsections we are going to use finite games on trees to represent periodic stopping games. Since the state space Ω is arbitrary, while games on trees only represent games with a finite state space, we need to approximate $\mathcal{F}$ by representative finite partitions. This can be done by using the method presented in Shmaya and Solan ([32, Sect. 6]). For each number $n \geq 0$ and bounded stopping time σ we define a representative finite partition $\mathcal{G}_{n,\sigma}$ of Ω such that: (1) $\mathcal{G}_{n,\sigma}$ refines $\mathcal{G}_{k,\tau}$ whenever $k \leq n$ and $\tau \leq \sigma$. (2) $\mathcal{G}_{n,\sigma}$ is $\mathcal{F}_n$-measurable. (3) $\mathcal{G}_{n,\sigma}$ contains all the information relevant to the players until σ is reached. Given $k \geq 0$, $\omega \in \Omega$, and $\tau \geq k$, let $F_{k,\tau}(\omega)$ be the element of $\mathcal{G}_{k,\tau}$ that includes ω.

Let $n < \sigma$ be a bounded stopping time, and $F \in \mathcal{G}_{n,\sigma}$. We define the game on a tree $T_{n,\sigma}(F)$ as follows: The game begins at stage n, when the state $\omega \in F \subseteq \mathcal{G}_{n,\sigma}$ 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 $\omega \in F \subseteq \mathcal{G}_{n,\sigma}$ is randomly chosen). Players are only allowed to stop in nodes where the matrix payoff is in D_F (repeats infinitely often in the infinite stopping game). Formally:

Definition 18 Let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a stopping game , $n \geq 0$ a number, $n < \tau$ a bounded stopping time, $(\mathcal{G}_{k,\tau})_{\tau \geq k \geq n}$ representative finite approximating partitions of $\mathcal{F}$, and $F \in \mathcal{G}_{n,\tau}$. The game on the finite tree $T_{n,\tau}(F)$ is $\left(I, V, V_{leaf}, r, V_{stop}, (C_v, p_v, R_v)_{v \in V \setminus V_{leaf}}\right)$ where:

- $V = \bigcup_{\substack{\omega \in F \\ n \leq k \leq \tau(\omega)}} \{F_{k,\tau}(\omega)\}$, $V_{leaf} = \bigcup_{\omega \in F} \left\{F_{\tau(\omega),\tau}(\omega)\right\}$, $r = F$, $V_{stop} = \{v \in V | d_v \in D_F\}$
- R_v, C_v, p_v are defined by induction. Assume that $v \in V \setminus V_{leaf}$ and $v \in \mathcal{G}_{k,\tau}$

for some $n \leq k$, then: $R_v = R_n(v)$, $C_v = \{F_{k+1,\tau} \in \mathcal{G}_{k+1,\tau} | F_{k+1,\tau} \subseteq v\}$, and $p_v(F_{k+1,\tau}) = p(F_{k+1,\tau}|v)$.

5.3 *Equivalence with Absorbing Games*

A stopping 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 (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 non-absorbing action profile.

Adapting [36]'s Prop. 4.10 to the two special properties gives the following:

Definition 19 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. [15]

Proposition 20 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 in case 1.

Lemma 21 In case 1 of prop. *20,* T admits a stationary absorbing perfect ϵ-equilibrium $x \neq c$.

[15] Originally part 3 of Prop. 20 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 $v \in V_{stop}$.

PROOF. Let T_ϵ be a perturbed version of the game on a tree T: In T_ϵ when a non-empty coalition wishes to stop at some node, there is a probability ϵ^2 that the "stopping request is ignored", and the game continues to the next stage. 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. 20 applies, then T_ϵ admits a perfect stationary equilibrium x_ϵ, and x_ϵ is a perfect stationary absorbing ϵ-equilibrium in T.

5.4 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. 20), by adapting to the multi-player case the methods presented in [32, 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 [32], and the proof is omitted as the changes are minor.

Lemma 22 *Let G be a stopping game, $n > 0$ a number, $\sigma > n$ a bounded stopping time, $F \in \mathcal{G}_{n,\sigma}$, 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.

We now define a subgame of a game on a tree.

Definition 23 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 elements of $\mathcal{G}_{n,\tau}$ into two kinds: *simple* and *complicated*. This division will be used in the following sections. The simple sets have at least one of the following properties: (1) There is a player that receives a negative payoff whenever he stops alone. (2) There is a distribution over the set of action profiles in which a single player stops. Moreover, each player receives payoff α_i^F when he stops, and approximately this is also his average payoff when another player stops.

Definition 24 Let G be a stopping game, $\epsilon > 0$, and $N_0 \leq n$ a number, and $\tau > n$ a bounded stopping time. The set $F \in \mathcal{G}_{n,\tau}$ is *ϵ-simple* if one of the following holds:

(1) For every $i \in I$: $\alpha_F^i < 0$. 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.

The next proposition analyzes stationary $\epsilon-$equilibria that yield high payoffs to all the players. The proposition is an adaptation of Prop. 5.5 in [32, Sec. 8]. The proof is omitted as the changes compared with [32] are minor.

Proposition 25 Let G be a stopping game, $N_0 \leq n$ a number, $\sigma > n$ a bounded stopping time, $F \in \mathcal{G}_{n,\sigma}$ 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 ([28,32]), to disassemble an infinite stopping game into games on finite trees with special properties. We begin by defining an $\mathcal{F}$-consistent C-valued NT-function.

Definition 26 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 27 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 ([32, Theorem 4.3]):

Proposition 28 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 $\sigma > n \geq 0$ and $F \in \mathcal{G}_{n,\sigma}$, such that $c_{n,\sigma}(F)$ is a C-valued $\mathcal{F}$-consistent NT-function.

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 W be a finite covering of $[-1,1]^{|I|}$ with (not necessarily disjoint) rectangles $([a^i, a^i + \epsilon])_{i\in I}$, 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 stopping game, $n \geq 0$, $\sigma > n$ a bounded stopping time, and $F \in \mathcal{G}_{n,\sigma}$. 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 such 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:[16]

- Set $T^{(0)} = T_{n,\sigma}(F)$.
- For $1 \leq j \leq J$ apply Prop. 20 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: $\gamma^i(x^{-i}_j, y^i) \leq a^i_j + 8\epsilon$.

[16] The procedure is an adaptation of the 2-player procedure described in [32, Sec. 5]

(c) $\pi(x_j) \geq \epsilon^2 \times p_{T^{(j)},T^{(j-1)}}$.

- If $T^{(J)}$ is trivial (the only node is the root), set $c_{n,\sigma}(F) = allbad$; otherwise due to Prop. 20 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 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. 20. 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 $g(T^{(J)})$.

Each strategy profile x_j, as given by Prop. 20, is a profile in $T^{(j-1)}$. We consider is as a profile in T by letting it 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. 28 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 $F \in \mathcal{G}_{\sigma_1,\sigma_2,}$, 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: $A_{\infty,j} = \left\{w \in \Omega | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}}\left(F_{\sigma_k(\omega)}\right) = \infty\right\}$, $A_{\epsilon,j} = \left\{w \in \Omega | \sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}}\left(F_{\sigma_k(\omega)}\right) \leq \frac{\epsilon}{|J|}\right\}$. 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) $p\left(A_{\epsilon,j} | \bar{A}_{\epsilon,j}\right) \geq 1 - \frac{\delta}{6\cdot|J|}$. (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 $p(E') \geq 1 - \delta$):

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

7 Approximate Constant-Expectation Perfect Correlated Equilibrium

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

Proposition 29 *Let G be a stopping game, $\delta, \epsilon > 0$, $E' \subseteq \Omega$, σ_1 and σ_2 be defined as in the previous subsection, and $F \in \mathcal{G}_{\sigma_1,\sigma_2}$. 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 E' and given M', where $\mu(M') > 1 - \delta$.*

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 adapt the methods of [32, Sec.7].

7.1 There exists $j \in J$ such that $F \in \bar{A}_{\infty,j}$

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 $F \in \bigcap_{j \in J} \bar{A}_{\epsilon,j}$ 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 E'.

7.3 $F \in \bigcap_{j \in J} \bar{A}_{\epsilon,j}$ 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 $i \in I$ such that: $a^i_w \leq \alpha^i_F - 2\epsilon$. 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 22 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 $q^i_{\sigma_k,\sigma_{k+1}} = p(\exists \sigma_k \leq n < \sigma_{k+1},\ R^i_{i,n} =$

α^i_F, $R^i_{i,n} \in D_F$). On E', $R^i_{i,n} = \alpha^i_F$ infinitely often and $\sum_{j=1..J}\sum_{k=1..\infty} \lambda_{j,\sigma_k,\sigma_{k+1}} < \epsilon$. 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 E'.

7.4 $F \in \bigcap_{j\in J} \bar{A}_{\epsilon,j}$ *and* $c_F = (1, w_1, w_2) \in (1 \times W \times W)$

The construction in this case is as an adaptation of the procedure of [35], which deals with quitting games (stationary stopping 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. 20. The definition of α^i_F implies that $\alpha^i_F = g^i(T_{\sigma_1,\sigma_2}) \in w^i_2$. 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 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 $\hat{l}$ 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 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 $M_{\delta,D_F} \subseteq M_{D_F}$ be the 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 $\mu(M_{\delta,D_F}) \leq \frac{\delta}{2^D}$. 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: $n = (m_i) \bmod \hat{T} + T + 1$,[17] and continues in all other stages. Let the universal correlation device $\mathcal{D} = (M, \mu)$ be the Cartesian multiplication: $\mathcal{D} =$

[17] 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.

$\prod_{D_F \subseteq D} \mathcal{D}_{D_F}$. Similarly let $M' = M \setminus \prod_{D_F \subseteq D} M_{\delta, D_F}$. Observe that $\mu(M') \geq 1 - \delta$.

If the players follow the strategy profile x_F then the game is absorbed with probability 1 conditioned on E' and the expected payoff satisfies $\alpha_F^i \leq \gamma_F^i(x) \in w_1^i$. Moreover, if $\hat{T} >> T$, then immediately after receiving his signal m_i (assuming $m \in M'$) no player can infer from his signal whether 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 $\omega \in E'$ and given $m \in M'$. 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 $\tau_{\hat{l}}^i$, the stage the quitter stops . Moreover, if T is sufficiently large, then, with high probability, player j does not know when he recieves his signal 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 not 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 $m \in M'$). 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 $\omega \in E'$ and given $m \in M'$.

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 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) $\theta(d, i) > 0 \Rightarrow R_{\{i\},d}^i = \alpha_F^i$. (2) $\alpha_F^i + \epsilon \geq \sum_{j \in I, d \in D_F} \theta(d, j) \cdot R_{\{j\},d}^i \geq \alpha_F^i - \epsilon$. 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 $\omega \in E'$ and given $m \in M'$.

8 Generalized Stopping Games

In the previous sections we only dealt with *simple* stopping games, which end as soon as any player stops. In this section we show how to extend our result to more generalized stopping games, where the game terminates only after all the players stopped.

A *generalized* stopping games is played as follows. 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 actions while he has been active and on the state variable. Theorem 14 shows that every simple stopping game admits a perfect normal-form constant-expectation correlated approximate equilibrium with a canonical and universal correlation device. For brevity, we will relate to such an equilibrium in the rest of this section as a "good" approximate equilibrium. We now sketch the outline of the proof that every generalized stopping game admits a "good" approximate equilibrium.

Assume by induction that any m-player stopping game admits a "good" approximate equilibrium. Given a generalized stopping game G' with $m+1$ players, we construct an auxiliary simple stopping game G with the following payoff process:

- When $i \in S$: $R^i_{S,n}$ is equal to the payoff of player i in the generalized game G' when coalition S stops at stage n, while no other player stopped before.
- When $i \notin S$: $R^i_{S,n}$ is the payoff of player i in a "good" approximate equilibrium of the induced generalized $m+1-|S|$player stopping game that begins at stage $n+1$ with the players $I \setminus S$. Such an equilibrium exists due to the induction hypothesis.

Due to Theorem 14, the simple game G admits a "good" approximate equilibrium x. x induces a "good" approximate equilibrium x' in the generalized game G' in a natural way:

- The players follow x as long as all the players continue.
- As soon as some of the players stop, the remaining active players play the "good" approximate equilibrium of the induced generalized stopping game with fewer players.

References

[1] R.J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 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] E. Ben-Porath. Communication without mediation: expending the set of equilibrium outcomes by "cheap" pre-play procedures. *Journal of Economic Theory* 80 (1998), 108-122.

[5] J. Bulow, P. Klemperer. The generalized war of attrition. *American Economic Review* 89 (2001), 175-189.

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

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

[8] E. B. Dynkin, Game variant of a problem on optimal stopping. *Soviet mathematics - Doklady.* 10 (1969), 270-274.

[9] C.H. Fine, L. Li. Equilibrium exit in stochastically declining industries. *Games and Economic Behavior* 1 (1989), 40-59.

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

[11] D. Fudenberg, J.Tirole. Preemption and Rent Equalization in the Adoption of New Technology. *Review of Economic Studies* LII (1985), 383-401.

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

[13] P. Ghemawat, B. Nalebuff. Exit. *RAND Journal of Economics* 16 (1985),184-194.

[14] Y. Heller. Minority-proof cheap-talk protocol, manuscript (2009), http://www.tau.ac.il/~helleryu/minority.pdf

[15] V. Krishna, J. Morgan. An Analysis of the War of Attrition and the All-Pay Auction. *Journal of Economic Theory* 72 (1997), 343-362.

[16] J.W. Mamer. Monotone stopping games. *Journal of Applied Probability* 24 (1987), 386-401.

[17] A. Mashiah-Yaakovi. Subgame perfect equilibria in stopping games. mimeo (2009), http://www.math.tau.ac.il/~ayalam/publication.files/general stopping game 3.pdf

[18] J. Maynard Smith. The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology* 47 (1974), 209-221.

[19] H. Morimoto. Non-zero-sum discrete parameter stochastic games with stopping times. *Probability Theory and Related Fields* 72 (1986), 155-160.

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

[21] Myerson R (1986b), Acceptable and predominant correlated equilibria. *International Journal of Game Theory* 15 (3), 133-154.

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

[23] P. Neumann, D. Ramsey, and K. Szajowski. Randomized stopping times in Dynkin games. *Journal of Applied Mathematics and Mechanics* 82 (2002), 811-819.

[24] J. Neveu, Discrete-Parameters Martingales. Borth-Holland, Amsterdam (1975).

[25] J. Nikkinen, M. Omran, P. Sahlstrom, J Aijo. Global stock market reactions to scheduled U.S. macroeconomic news announcements. *Global Finance Journal* 17, 92-104 (2006).

[26] 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). Birkhauser, Boston.

[27] Y. Ohtsubo. On a discrete-time non-zero-sum Dynkin problem with monotonicity. *Journal of Applied Probability*, 28 (1991), 466-472.

[28] F. Ramsey, On a problem of formal logic. *Proceedings of the London Mathematical Society* 30 (1930), 264-286.

[29] D. Rosenberg, E. Solan, N. Vieille, Stopping games with randomized strategies. *Probability Theory and Related Fields* 119 (2001), 433-451.

[30] A. Rubinstein, Comments on the Interpretation of Game Theory. *Econometrica* 59 (1991), 909-924.

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

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

[33] 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. Guth, P. Hammerstein, B. Moldovanu, E. van Damme. Springer-Verlag (1997), 149-170.

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

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

[36] E. Solan, V. Vohra, Correlated equilibrium payoffs and public signaling in absorbing games, *International Journal of Game Theory* 31 (2002), 91-121.

[37] S. Sorin, Distribution equilibrium I: definition and equilibrium. manuscript. 1998.

[38] Urbano, A., Vila J.E, Computational complexity and communication: coordination in two-player games. *Econometrica* 70 (2002), 1893-1927.