



Munich Personal RePEc Archive

Perfect correlated equilibria in stopping games

Heller, Yuval

School of Mathematical Sciences, Tel-Aviv University

10 June 2009

Online at <https://mpra.ub.uni-muenchen.de/25891/>
MPRA Paper No. 25891, posted 18 Oct 2010 15:14 UTC

Sequential Correlated Equilibria in Stopping Games

Yuval Heller

*School of Mathematical Sciences, Tel-Aviv University, P.O. box: 39040, Tel-Aviv
69978, Israel. Phone: 972-3-640-5386. Fax: 972-3-640-9357.
Email: helleryu@post.tau.ac.il*

Abstract

In many situations, such as trade in stock exchanges, agents have many instances to act even though the duration of interactions take a relatively short time. The agents in such situations can often coordinate their actions in advance, but coordination during the game consumes too much time. An equilibrium in such situations has to be sequential in order to handle mistakes made by players. In this paper, we present a new solution concept for infinite-horizon dynamic games, which is appropriate for such situations: a sequential constant-expectation normal-form correlated approximate equilibrium. Under additional assumptions, we show that every such game admits this kind of equilibrium.

Subject classifications: games/group decisions: stochastic. financial institutions: trading.

Area of review: decision analysis.

History: received December 2009, revised: April 2010, September 2010.

1 Introduction

In the modern world there are many situations that last a relatively short time but in which agents have many instances to act, such as on-line auctions and trade in stock exchanges. In many cases different agents share similar, though not identical, goals. Such is the case when the agents work in the same financial institution, and they can coordinate their actions in order to maximize the institution's profit, as well as the contribution of each agent to this profit. In this paper we present a game theoretic model for such interactions, and we propose a new solution concept for these games, that is suitable for situations where players' utilities share a common part as well as an individual part.

To motivate the study, consider the following situation. Each month the Bureau of Labor Statistics publishes a news release on the U.S. employment situation (ES). This news

release is announced in the middle of the trading day in the European stock markets (on the first Friday of each month at 13:30 London time). The ES announcement has a strong impact on these markets (see Nikkinen et al., 2006 and the references within). Empirical studies (see for example, Christie-David, Chaudhry and Khan, 2002) show that a few tens of minutes elapse before financial instruments adjust to such announcements. This gap of time (the “adjustment period”) may provide an opportunity for substantial profit to be made by quick trading (“news-playing”). Consider the strategic interaction between a few traders in a financial institution who coordinate in advance their actions in the adjustment period. Each trader can make buy and sell orders for some financial instruments that he is responsible for. The traders share a common objective - maximizing the profit of the institution. In addition to this, each trader also has a private objective - maximizing the profit that is made in financial instruments that he is responsible for (which influences his bonuses and prestige). The traders can freely communicate before the ES announcement, but communication during the adjustment period is costly: each moment that is spent on communication may slow down the traders and limit their potential profits.

The family of strategic interactions that we study has the following properties: (1) the interaction lasts a relatively short time but agents have many instances to act; (2) different agents share similar, though not identical, goals; (3) each agent chooses his action autonomously; (4) agents can freely communicate before the game starts, but communication during play is costly or not feasible; (5) agents may occasionally make mistakes, and not execute the action that they planned to take. Three natural questions arise when modeling such strategic interactions: (1) Which kind of game should be used? (2) Which solution concept should be chosen? (3) Does a solution exist, and can we find one?

We begin by dealing with the first question. As each agent chooses his actions autonomously, we model this interaction as a noncooperative game (and not as a coalitional game; see Osborne and Rubinstein, 1994, Section IV, for discussing these two modeling approaches). Next we discuss the length of the game. The interaction is relatively short in absolute terms. Nevertheless, the agents have many opportunities to act (in the leading example, trade orders can be made in each fraction of a second). In addition, the point in time where the game ends may not be known to the players in real-time. Thus, it seems appropriate to model this situation as a stochastic (dynamic) game with an *infinite-horizon*, rather than modeling it as a game with a fixed finite large number of stages. See Rubinstein (1991) and Aumann and Maschler (1995, pages 131-137) for discussions why even short strategic interactions may be better analyzed as infinite-horizon games. Infinite horizon games have been used in a wide range of applications, such as: bargaining (Chatterjee and Samuelson, 1988), inventory control system (Bouakiz and Sobel, 1992), oligopolistic competition (Bernstein and Federgruen, 2004), and supply chain relationships (Taylor and Plambeck, 2007).

The issue raised in the second question - which solution concept is appropriate - has

several aspects. First, we discuss how each agent evaluates payoffs at different stages of the infinite-horizon game. As the interaction is short in absolute time, it is natural to assume that payoffs are evaluated without discounting. Because, in undiscounted games, payoffs that are obtained in the first T stages do not affect the total payoff, for every T ; yet the interaction in our example is finite, the solution concept should satisfy uniformity: it should be an approximate equilibrium in any long enough finite-horizon game. See Aumann and Maschler (1995, pages 138-142) for arguments in favor of this notion.

The agents in the family of games that we study, can freely communicate before the game starts, and coordinate their strategies. Aumann (1974) defined normal-form correlated equilibrium in a finite 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. Forges (1986) extended this notion to dynamic games. Under relatively mild conditions, pre-play non-binding communication among the players can implement a normal-form correlated equilibrium (see, e.g., Ben-Porath, 1998), and thus this solution concept is natural in our setup. Forges (1986) also presented the alternative notion of extensive-form correlated equilibrium, which requires communication at each stage of the game. This alternative notion is less appropriate to our family of games, because communication during play is costly or not feasible.

As players may make mistakes, or forget what they were supposed to do in the equilibrium, the behavior of the players should also be rational off the equilibrium path. That is, players should also use their best response after one player makes a mistake and deviates from the equilibrium strategy profile. This is satisfied by requiring the equilibrium to be sequential (Kreps and Wilson, 1982).

The above reasoning limits the plausible outcomes of the game to the set of sequential normal-form correlated equilibria. See Myerson (1986a, 1986b) and Dhillon and Mertens (1996) who study related notions. As infinite undiscounted games may only admit approximate equilibria, we define a *sequential normal-form correlated* (δ, ϵ) -*equilibrium*, as a strategy profile where with probability 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).

Another desirable property in our setup is that the expected payoff of each player be independent of the pre-play communication. This facilitates the implementation of the coordination among the players, as none of them may feel discriminated by the coordination process. Sorin (1998) 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. We generalize Sorin's definition, and define an approximate *constant-expectation correlated equilibrium*, as a correlated equilibrium where the expected payoff of each player hardly changes when he receives his signal. Observe that every Nash equilibrium is a

constant-expectation correlated equilibrium. In Section 6 we discuss the rationale for this notion and its basic properties.

The first contribution of this paper is the presentation of a new solution concept for undiscounted dynamic games: a sequential uniform constant-expectation normal-form correlated approximate equilibrium.

We now deal with the third question: proving the existence of this equilibrium. In this paper we prove existence under the simplifying assumption that, throughout the game, the agents have symmetric information. This assumption is reasonable in many situations. For example, in the leading example, each trader can electronically access the data on all the prices of the different markets. Although in reality each trader may actually focus only on the information that is more relevant for the financial instruments that he is responsible for, he may obtain the relevant information of other players, when necessary.

A second simplifying assumption is that each player has a finite number of actions. In the leading example, each trader has a finite set of financial instruments that he is responsible for, and for each such instrument he chooses a time to buy or a time to sell. Thus, it can be assumed that a trader's strategy is a vector of buy and sell times, one for each financial instrument that he is responsible for.

The model we study also applies to situations of a different nature, for example:

- Several countries plan to ally in a war against another country. The allying countries share a common objective - maximizing their military success against the common enemy. In addition, each country has private objectives, such as maximizing the territories and resources it occupies during the war, and minimizing its losses. This situation has similar properties to the leading example: (1) The war is relatively short in absolute time (a modern war typically lasts a couple of weeks), but it consists of an unknown large number of stages. (2) The leaders of each country can communicate and coordinate their future actions before the war begins. On the other hand, secure communication and coordination during the war may be costly and noisy. (3) Finally, usually only a few of the battlefield actions of each country are crucial to the outcome of the war (such as the timing of the main military attack).
- 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, 1974). In most cases, the animals use "ritualized" fighting and do not seriously injure their opponents. The winner is the contestant who continues the war for the longest time. Excessive persistence has the disadvantage of waste of time and energy in the contest. This situation also shares similar properties with the leading example: (1) The war of attrition is short in absolute time (usually a few hours or days), but consists of an unknown large number of stages. (2) Shmida and Peleg (1997) discuss how a normal-form correlation device can be induced in biological setups by phenotypic

conditional behavior, and Sorin (1998) discusses why the constant-expectation requirement is necessary for the stability of the population in evolutionary setups (see Section 6). (3) Finally, each animal in the war of attrition acts only once, by choosing when to quit the contest.

Under the assumptions discussed earlier, all these strategic interactions are modeled as follows. There is an unknown state variable on which players receive symmetric partial information during play. For each player i (from a finite set of players), there is a finite number, T_i , that limits the number of actions he may take during the game. At stage 1 all the players are active. At every stage n , each active player declares, independently of the others, whether he takes one of a finite number of actions or “does nothing”. A player who acted T_i times, becomes passive for the rest of the game and must “do nothing” in all subsequent stages. The payoff of a player depends on the history of actions and on the state variable. By induction one can show that the problem of equilibrium existence reduces to the case when $T_i = 1$ for every player i . Moreover, one can show that the problem further reduces to the case where each player has a single “stopping” action, and that the game ends as soon as any player stops (see Section 5).

Such a game is called a (discrete undiscounted) *stopping game*. The literature includes two variants for the definition of stopping games. Some papers (see, e.g., Shmaya and Solan, 2004) assume that the game ends as soon as any player stops. Other papers (see, e.g., Ramsey, 2007) assume that after one player stops, the other players continue to play. In this paper, we formally follow the first definition, and we show in Section 5 how our result can be applied to the second variant as well.

Stopping games were introduced by Dynkin (1969), 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, 1985; Mamer, 1987), struggle of survival among firms in a declining market (see e.g., Fudenberg and Tirole, 1986), auctions (see e.g., Krishna and Morgan, 1997), lobbying (see e.g., Bulow and Klemperer, 2001), conflict among animals (see e.g., Nalebuff and Riley, 1985), and duels (see, e.g., Karlin, 1959). Stopping games where players are allowed to stop more than once ($T_i > 1$) are investigated, among others, in Szajowski (2002), Yasuda and Szajowski (2002) and Laraki and Solan (2005).

Much work has been devoted to the study of undiscounted two-player stopping games. This problem, when the payoffs have a special structure, was studied by Neveu (1975), Mamer (1987), Morimoto (1986), Ohtsubo (1991), Nowak and Szajowski (1999), Rosenberg, Solan and Vieille (2001), Neumann, Ramsey and Szajowski (2002), and Shmaya and Solan (2004), among others. Those authors provided various sufficient conditions under which (Nash) approximate equilibria exist.

Undiscounted multi-player stopping games have mostly been modeled in the existing

literature as cooperative (coalitional) games. Assaf and Samuel-Cahn (1998a, 1998b) and Glickman (2004) have studied a model where players can only stop by an unanimous decision, and that the group’s stopping rule maximizes a specific function of the expected payoff of each player. Other papers have investigated the use of cooperative solution concepts in this setup: the core (Ohtsubo, 1996), Pareto-optima (Ohtsubo, 1995, 1998) and Shapley value (Ramsey and Cierpial, 2009). Another model, which is more related to our noncooperative framework, is a stopping game with a voting procedure. In such games, each player votes at each stage whether or not he wishes to stop the game, and there is some monotonic rule (for example, majority rule) that determines if the set of players who voted to terminate, has the power to stop the game (see Section 5 for discussing the extension of our model to include a voting procedure). This model has been studied, among others, in Kurano, Yasuda and Nakagami (1980), Yasuda, Nakagami and Kurano (1982), Szajowski and Yasuda (1997), and Ferguson (2002). All these papers assume a simplifying assumption, which is not assumed in our model, that the payoffs to the players only depend on the stage in which the game stops, but not on the identity of the stopping players. In contrast with the two-player case, there is no existence result for approximate equilibria in multi-player stopping games without this assumption.

Our main result states that for every $\delta, \epsilon > 0$, a multi-player stopping game admits a sequential uniform constant-expectation normal-form correlated (δ, ϵ) -equilibrium. We further show that the equilibrium’s correlation device has two appealing properties: (1) it is canonical - each signal is equivalent to a strategy; and (2) it does not depend on the specific parameters of the game. The proof relies on a stochastic variation of Ramsey’s theorem (Shmaya and Solan, 2004) that reduces the problem to that of studying the properties of correlated ϵ -equilibria in multi-player absorbing games (stochastic games with a single non-absorbing state). The study uses the result of Solan and Vohra (2002) that any multi-player absorbing game admits a correlated ϵ -equilibrium.

Another interesting question is characterizing the properties of the set of equilibrium payoffs, and present methods for selecting a specific equilibrium with corresponding payoff that satisfies some appealing properties, like Pareto-efficiency, maximizing the sum of payoffs (utilitarianism, efficiency), or maximizing the minimal payoff (egalitarianism). Such methods are important for the use of the model in applications, such as the leading example. Our proof is not constructive, and this question, with general payoff structure, remains open for future research. The reader is referred to Ramsey and Szajowski (2008), and the references therein, who study this problem in a two-player stopping game.

The paper is arranged as follows. Section 2 presents the model and the result. A sketch of the proof appears in Section 3. Section 4 contains the proof. In Section 5 we discuss how to apply our result, which formally deals only with “simple” stopping games, to more general situations, such as the leading example. Section 6 discusses the rationale of the notion of constant-expectation correlated equilibrium.

2 Model and Main Result

In the introduction, we presented an example of the strategic interaction among traders when some macroeconomic news is published (the leading example), and discussed how to model it by a stopping game. In this section we present the formal definitions, and state our main result.

A stopping game is defined as follows:

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

- I is a finite set of players;
- (Ω, \mathcal{A}, p) is a probability space;
- $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ is a filtration over (Ω, \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 $R_{S,n}^i$ where $i \in I$ and $\emptyset \neq S \subseteq I$.

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 set of players $S \subseteq I$, the game terminates, and the payoff to player i is $R_{S,n}^i$. If no player ever stops, the payoff to everyone is zero.

Remark 2 According to Definition 1, a stopping game ends as soon as one of the players stops. As discussed earlier, the literature also includes another definition (see, e.g., Ramsey, 2007), according to which, when one player stops, the others continue to play. In Section 5 we discuss how to apply our result to the alternative definition, and to a more complicated strategic interaction, as in the leading example, in which players have a more than one action, and may act more than once during the game.

We model the pre-play communication possibilities of the players by a correlation device:

Definition 3 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) $\mu \in \Delta(M)$ is the probability distribution according to which the device sends the signals to the players before the stopping game starts.

As discussed earlier, cheap talk communication among the players can be used to “mimic” a correlation device. Specifically, when there are at least three players, under mild conditions on the set of Nash equilibrium payoffs, any correlated equilibrium can be implemented as a sequential equilibrium of an extended game with pre-play cheap talk (Ben-Porath, 1998; see also Heller, 2010a for an implementation that is resistant to coalitional deviations). This is also true for two players, under additional cryptographic assumptions (Urbano and Vila, 2002).

Throughout the paper we denote the signal profile that the players receive from the correlation device by m . Given a normal-form correlation device \mathcal{D} , we define an extended game $G(\mathcal{D})$. The game $G(\mathcal{D})$ is played exactly as G , except that, at the outset of the game, a signal profile $m = (m^i)_{i \in I}$ is drawn according to μ , and each player i is privately informed of m^i . Then, each player may base his strategy on the signal he received.

As mentioned earlier, Shmida and Peleg (1997, Section 5) discuss how a normal-form correlation device can be induced in nature by phenotypic conditional behavior. Specifically, they present an example 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 quickly.

For simplicity of notation, let the singleton set $\{i\}$ be denoted as i , and let $-i = I \setminus \{i\}$ denote the set of all players besides player i . A (behavior) *strategy* for player i in $G(\mathcal{D})$ is an \mathcal{F} -adapted process $x^i = (x_n^i)_{n \geq 0}$, where $x_n^i : (\Omega \times M^i) \rightarrow [0, 1]$. The interpretation is that $x_n^i(\omega, m^i)$ is the probability by which player i stops at stage n when he received a signal m^i .

Let θ be the first stage in which at least one player stops, and let $\theta = \infty$ if no player ever stops. If $\theta < \infty$ let $S_\theta \subseteq I$ be the set of players that stops at stage θ . The expected payoff of player i under the strategy profile $x = (x^i)_{i \in I}$ is given by $\gamma^i(x) = \mathbf{E}_x(\mathbf{1}_{\theta < \infty} \cdot R_{S_\theta, \theta}^i)$ 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 $E \subseteq \Omega$ and a set of signal profiles $M' \subseteq M$, let $\gamma^i(x|E, M')$ be the expected payoff of player i conditioned on E and on the signal profile being in M' . Given $m' \in M'$, let $\gamma^i(x|E, M', m'^i)$ denote the expected payoff of player i conditioned on E , on the signal profile being in M' , and on the signal of player i being equal to m'^i .

The strategy x^i is an 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 and M' if $\gamma^i(x|E, M') \geq \gamma^i(x^{-i}, y^i|E, M') - \epsilon$.

We say that a profile x in $G(\mathcal{D})$ is ϵ -constant-expectation conditioned on E and M' , if whenever the state is in $E \subseteq \Omega$ and the signal profile is in M' , the expected payoff of each player changes by at most ϵ when he obtains his signal. We say that x is a (δ, ϵ) -constant-expectation if this holds for some E and M' with probability at least $1 - \delta$.

Definition 4 Let $G(\mathcal{D})$ be an extended stopping game (where $\mathcal{D} = (M, \mu)$), $M' \subseteq M$ and $E \subseteq \Omega$. The strategy profile x in $G(\mathcal{D})$ is a (δ, ϵ) -constant-expectation (where $\epsilon, \delta \geq 0$) if there is a set $M' \subseteq M$ and an event E such that $\mu(M') \geq 1 - \delta$, $p(E) \geq 1 - \delta$, for every $i \in I$ and $m' \in M'$: $|\gamma^i(x|E, M', m'^i) - \gamma^i(x|E, M')| \leq \epsilon$.

The definition of an approximate constant-expectation correlated equilibrium generalizes

Sorin (1998)'s definition of distribution equilibrium for finite normal-form games. As discussed earlier, constant-expectation correlated equilibria are more easily implemented in economic setups such as the leading example, as none of the players may feel discriminated against by the coordination process. In Section 6 we discuss the rationale of this notion and its basic properties.

Given $\omega \in \Omega$, let $H_n(\omega) \subseteq \mathcal{F}_n$ be the collection of all events in \mathcal{F}_n that include ω : $H_n(\omega) = \{F_n \in \mathcal{F}_n | \omega \in F_n\}$. $H_n(\omega)$ denotes the public history of play up to stage n , when the true state is ω . Let \mathcal{H}_n be the collection of all such histories of length n : $\mathcal{H}_n = \{H_n(\omega) | \omega \in \Omega\}$, and let $\mathcal{H} = \bigcup_{n=1.. \infty} \mathcal{H}_n$ be the set of all histories. Let $G(H_n, \mathcal{D}, m)$ be the induced stopping game that begins at stage n , when each player i has received the private signal $m^i \in M^i$, and the public history is $H_n \in \mathcal{H}_n$. For simplicity of notation, we use the same notation for a strategy profile in $G(\mathcal{D})$ and for the induced strategy profile in $G(H_n, \mathcal{D}, m)$.

As discussed earlier, we require players to also be rational off the equilibrium path. This is satisfied by requiring the equilibrium to be sequential (Kreps and Wilson, 1982). In what follows we adapt the definition of sequential equilibrium in a finite extensive-form game, to our framework of infinite extended stopping games. The adaptation includes two parts: (1) Simplifying the belief system because the only source for imperfect information on past events is due to the private signals the players received from the correlation device before the game starts. (2) Defining an approximate variation of sequential equilibrium due to the infiniteness of stopping games. Observe that we adopt the notation of Osborne and Rubinstein (1994, Chapters 6 and 12), and do not consider simultaneous moves as a source of imperfect information.

We begin by defining a *belief system* in an extended stopping game $G(\mathcal{D})$ as a profile of functions $(q^i)_{i \in I}$. Each function $q^i : \mathcal{H} \times M^i \rightarrow \Delta(M^{-i})$ assigns a distribution over the signals of the other players. The distribution is interpreted as follows: after receiving a signal m^i and observing a public history H , player i assigns probability $q^i(H, m^i)(m^{-i})$ to the signal profile of the other players being m^{-i} . Given $M' \subseteq M$, let $q^i(H, m^i | M')$ be the belief of player i over the signal profile, conditional on the signal profile being in M' .

An *assessment* in an extended stopping game $G(\mathcal{D})$ is a pair (x, q) where x is a strategy profile and q is a belief system. An assessment is ϵ -*sequentially rational*, conditioned on an event E and on M' , if every player ϵ -best replies whenever the signal profile is in M' and the state is in E . When $\epsilon = 0$ it coincides with the standard definition of sequential rationality (Kreps and Wilson, 1982). Formally:

Definition 5 Let $G(\mathcal{D})$ be an extended stopping game (where $\mathcal{D} = (M, \mu)$), $\epsilon \geq 0$, $M' \subseteq M$, and $E \subseteq \Omega$. An assessment (x, q) is ϵ -*sequentially rational in $G(\mathcal{D})$ conditioned on E and M'* , if for every $i \in I$, $\omega \in E$, $n \in \mathbb{N}$, and signal profile $m \in M'$, x^i is an ϵ -best reply for player i conditioned on E and on M' in the induced game $G(H_n(\omega), \mathcal{D}, m)$,

when his opponents play x^{-i} , and his beliefs over the signal profile are $q^i(H_n(\omega), m^i | M')$.

A strategy profile is *completely mixed* if each player assigns positive probability to every action (stop or continue) after every history. An assessment (x, q) is consistent if it is the limit of a sequence of assessments $((x_n, q_n))_{n=1}^{\infty}$ with the following properties: (1) each strategy profile x_n is completely mixed; (2) each belief system q_n is derived from x_n using Bayes' rule. An assessment is a sequential ϵ -equilibrium conditioned on E and M' , if it is ϵ -sequentially rational (conditioned on E and M') and consistent. Formally:

Definition 6 Let $G(\mathcal{D})$ be an extended stopping game (where $\mathcal{D} = (M, \mu)$), $\epsilon \geq 0$, $M' \subseteq M$, and $E \subseteq \Omega$. An assessment (x, q) is *sequential ϵ -equilibrium* in $G(\mathcal{D})$ conditioned on E and M' , if it is both ϵ -sequentially rational conditioned on E and M' and consistent.

Definition 6 extends the standard definition of sequential equilibrium. That is, when $\epsilon = 0$, $M = M'$ and $E = \Omega$, it is equivalent to the standard definition of sequential equilibrium (Kreps and Wilson, 1982).

An assessment is a sequential (δ, ϵ) -equilibrium if it is a sequential ϵ -equilibrium conditioned on E and M' , where E and M' have probabilities of at least $1 - \delta$. Formally:

Definition 7 Let $G(\mathcal{D})$ be an extended stopping game and let $\delta, \epsilon \geq 0$. An assessment (x, q) is a *sequential (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 $p(E) \geq 1 - \delta$, $\mu(M') \geq 1 - \delta$, and x is a sequential ϵ -equilibrium of $G(\mathcal{D})$ conditioned on E and M' .

Abusing notation, we say that a strategy profile x is a sequential (δ, ϵ) -equilibrium of $G(\mathcal{D})$ if there is a belief system q , such that the assessment (x, q) is a sequential (δ, ϵ) -equilibrium in $G(\mathcal{D})$. Observe that when the correlation device is trivial ($|M| = 1$) sequentiality is equivalent to subgame perfectness (Selten, 1965, 1975). Specifically, when $|M| = 1$, the definition of a (δ, ϵ) -sequential equilibrium is equivalent to the definition of a (δ, ϵ) -subgame-perfect equilibrium in Mashiah-Yaakovi (2009). Without the limitation $|M| = 1$, every (δ, ϵ) -sequential equilibrium is a (δ, ϵ) -subgame-perfect equilibrium, but the converse is not true.

We now define a sequential correlated (δ, ϵ) -equilibrium.

Definition 8 Let G be a stopping game and let $\delta, \epsilon > 0$. A *sequential correlated (δ, ϵ) -equilibrium* is a pair (\mathcal{D}, x) , where \mathcal{D} is a correlation device and x is a sequential (δ, ϵ) -equilibrium in $G(\mathcal{D})$.

We end this subsection by defining another appealing property of a correlation device: canonicity. A correlation device $\mathcal{D} = (M, \mu)$ is *canonical* if each signal is equivalent to a strategy.

Definition 9 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 $m^i \neq m'^i$.

The standard definition of a canonical correlation device for finite games (Forges, 1986) is that the set of signals is equal to the set of strategy profiles. Definition 9 is different because the set of signals is finite, while the set of strategies is infinite.

Our main result is the following:

Theorem 10 Let $G = (I, \Omega, \mathcal{A}, p, \mathcal{F}, R)$ be a multi-player stopping game with integrable payoffs ($\sup_{n \in (\mathbf{N} \cup \infty)} \|R_n\|_\infty \in L^1(p)$). Then for every $\delta, \epsilon > 0$, G has a sequential (δ, ϵ) -constant-expectation normal-form correlated (δ, ϵ) -equilibrium with a canonical correlation device. Moreover, the correlation device only depends on the number of players and ϵ , and is independent of the payoff process.

The fact that the correlation device is independent of the payoff process allows the players to use the same correlation device in every stopping game (assuming the number of players and ϵ are fixed), and avoid the difficulties of constructing a new device for each stopping game. Thus, the traders in the leading example can construct, once and for all, a correlation device, and then use it for all future strategic interactions (regardless of the specific implications of the macroeconomic news that is going to be released).

Remark 11 The (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 (Proposition 17) or by applying a general observation made by Solan and Vieille (2001).

3 Sketch of the Proof

We begin our sketch by focusing on a simple kind of stopping games - periodic stopping games on finite trees. These are stopping games with a finite filtration, where after a finite number of stages, if not stopped earlier, the game restarts at the first stage. Such games are a special kind of absorbing games (stochastic games with a single non-absorbing state, see Sorin, 2002, 5.5). Solan and Vohra (2002) studied these games and proved that they admit a correlated ϵ -equilibrium. Adapting their result to our framework, implies that every periodic stopping game has either (1) a stationary equilibrium; or (2) a set of nodes in the tree $(\tilde{v}^i)_{i \in I}$, and a distribution ζ over the players, such that if player i is chosen according to ζ , and is asked to stop in node \tilde{v}^i , then this procedure induces a correlated equilibrium (each player has an incentive to stop only when being asked to).

We strengthen the result of Solan and Vohra (2002) if case (1) holds, by showing that there is a stationary sequential ϵ -equilibrium (by “perturbing” the game to continue

with positive probability at each stage). If case (2) holds, we modify the procedure in which players are being asked to stop in two ways. First, we ask each player to stop with probability $1 - \epsilon$ (and not with probability 1 as in Solan and Vohra, 2002), to prevent players from being able to deduce that they are off the equilibrium path (even when other players deviate). This allows us to obtain sequentiality. Second, we make sure that with high probability, when a player receives his signal, he cannot deduce, which player has been asked to stop. This modification guarantees constant-expectation (which trivially holds in the first case). Finally, we adapt the methods of Shmaya and Solan (2004), to extend our result to periodic games with infinite filtrations, and to show that each such game admits a correlated (δ, ϵ) -equilibrium.

The last step in the proof adapts the methods of Shmaya and Solan (2004) to extend the equilibrium existence result to infinite non-periodic stopping games. Shmaya and Solan proved a stochastic variation of Ramsey's Theorem (1930) that allows us to divide an infinite non-periodic stopping game into an infinite sequence of periodic stopping games, and to concatenate the correlated (δ, ϵ) -equilibrium in each periodic game, into a correlated (δ, ϵ) -equilibrium in the original infinite non-periodic game. We verify that the sequentiality and constant-expectation of each equilibrium in the periodic games imply the same properties for the equilibrium in the infinite game. Moreover, we show that the concatenated correlated equilibrium uses a correlation device which is normal-form, and only depends on the number of players and ϵ .

4 Proof

This section includes five parts. Subsection 4.1 includes some notation that is used later in the proof, and shows that one can focus on proving equilibrium existence in an induced game that begins after some bounded stopping time is reached. Subsection 4.2 presents a special form of stopping games - stopping games on finite trees, and shows that such games can approximate periodic stopping games with infinite filtrations. Subsection 4.3 adapts the result of Solan and Vohra (2002) and shows that every stopping games on a finite tree admits a sequential correlated equilibrium. Subsection 4.4 presents a stochastic variation of Ramsey's theorem, which is adapted from Solan and Shmaya (2004). Finally, Subsection 4.5 uses all the previous results to prove that every (infinite and non-periodic) stopping game admits a sequential correlated equilibrium with the properties required in Theorem 10.

4.1 Preliminaries

If with probability at least $1 - \delta$, the difference between the payoffs of two stopping games G and \tilde{G} is at most ϵ , then any sequential (δ, ϵ) -equilibrium in G is a sequential

$(3\delta, 3\epsilon)$ -equilibrium in \tilde{G} . Hence now fix a stopping game G and assume without loss of generality (w.l.o.g.) that the payoff process R is uniformly bounded and that its range is finite. In fact, we assume that for some $K \in \mathbf{N}$, $R_{S,n}^i \in \left\{0, \pm\frac{1}{K}, \pm\frac{2}{K}, \dots, \pm\frac{K}{K}\right\}$ for every $n \in \mathbf{N}$. Let $D = \prod_{i \in I, \emptyset \neq S \subseteq I} \left\{0, \pm\frac{1}{K}, \pm\frac{2}{K}, \dots, \pm\frac{K}{K}\right\}$ 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 .

We now fix $\epsilon, \delta > 0$. 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 (i.o.): $A_d = \{\omega \in \Omega \mid \text{i.o. } R_n(\omega) = d\}$, and let $B_{d,k} \subseteq \bigvee_{n \in \mathbf{N}} \mathcal{F}_n$ be the event that d never occurs after stage k : $B_{d,k} = \{\omega \in \Omega \mid \forall n \geq k, R_n(\omega) \neq d\}$. 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 \mathcal{F}_{N_0} -measurable sets $(\bar{A}_d, \bar{B}_d)_{d \in D} \in \mathcal{F}_{N_0}$ that approximate A_d and B_{d,N_0} . That is: (1) For each $d \in D$: $\bar{A}_d \cap \bar{B}_d = \emptyset$ and $(\bar{A}_d \cup \bar{B}_d) = \Omega$. (2) $\forall d \in D, p(A_d \mid \bar{A}_d) \geq 1 - \frac{\delta}{3 \cdot |D|}$. (3) $\forall d \in D, p(B_{d,N_0} \mid \bar{B}_d) \geq 1 - \frac{\delta}{3 \cdot |D|}$.

Let $\Phi = \bigcup_{d \in D} \left(\{\omega \in \bar{A}_d \mid \omega \notin A_d\} \cup \{\omega \in \bar{B}_d \mid \omega \notin B_{d,N_0}\} \right)$ be the event that includes all the approximation's "errors". That is, Φ includes all states where a payoff matrix d does not repeat infinitely often even though $\omega \in \bar{A}_d$, and all states where a payoff matrix d occurs after N_0 even though $\omega \in \bar{B}_d$. Observe that $p(\Phi) < \frac{\delta}{3}$. For any $H \in \mathcal{H}$ let $D(H) = \{d \in D \mid \exists F \in H, \text{ s.t. } F \subseteq \bar{A}_d\}$ be the set of payoff matrices that repeat infinitely often after history H (outside Φ). For each player $i \in I$, let $\alpha_H^i = \max(d_{\{i\}}^i \mid d \in D(H))$ be the maximal payoff a player can get by stopping alone in one of the matrices in $D(H)$.

Consider an induced game that begins after some bounded stopping time τ is reached. The following standard lemma shows that in order to prove Theorem 10, it is enough to show that each such game has an approximate constant-expectation sequential correlated equilibrium with a canonical correlation device that depends only on $|I|$ and ϵ .

Lemma 12 Let $\mathcal{D} = (M, \mu)$ a canonical correlation device that depends only on $|I|$ and ϵ , $M' \subseteq M$ a set satisfying $\mu(M') > 1 - \delta$, $E \subseteq \Omega$ an event such that $p(E) > 1 - \delta$, and τ a bounded stopping time. Assume that for every $\omega \in E$, $m \in M'$, and $H = H_{\tau(\omega)} \in \mathcal{H}_{\tau(\omega)}$, there is a constant-expectation sequential ϵ -equilibrium, x_H , in $G(H, \mathcal{D}, m)$ conditioned on E and M' . Then $G(\mathcal{D})$ admits a (δ, ϵ) -constant-expectation sequential (δ, ϵ) -equilibrium. This implies that G admits a sequential (δ, ϵ) -constant-expectation normal-form correlated (δ, ϵ) -equilibrium with a canonical device, which depends only on $|I|$ and ϵ .

PROOF. It is well known that any finite-stage game admits a sequential 0-equilibrium. Since τ is bounded, $p(E) \geq 1 - \delta$ and $\mu(M') \geq 1 - \delta$, the following strategy profile x is a (δ, ϵ) -constant-expectation sequential (δ, ϵ) -equilibrium:

- Until stage τ , play a sequential equilibrium, which is trivially a constant-expectation equilibrium, in the finite stopping 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 , from that stage on, play the profile x_H in $G(H, \mathcal{D}, m)$.

Observe that for the concatenated profile x to be a normal-form correlated equilibrium, it is necessary that each induced game's equilibrium would be constant-expectation. Otherwise, the signal a player receives before the game starts may change his expected payoffs in the induced games, and this may create profitable deviations from x . It is also necessary that all the correlated equilibria in the induced games use the same correlation device M . Also observe that the sequentiality and constant-expectation of each equilibrium in the induced games, imply that x has these two properties.

4.2 Periodic Stopping Games on Finite Trees

Generally, a stopping game is non-periodic, has an infinite length and has an infinite filtration. We now consider a special kind of stopping game, which is periodic (with finite length) and has a finite filtration. Such a game can be modeled by a game on a finite tree. The game starts at the root and is played in stages. Each node in the tree has a matrix payoff (in case players stop at that node), and a distribution over its offspring nodes, which determines the probability that the game would continue to each of these nodes, if no player stops. Given the current node, and the sequence of nodes already visited, the players decide, simultaneously and independently, whether to stop or to continue. Let S be the set of players that decides to stop. If $S \neq \emptyset$, the play ends and the terminal payoff to each player i is determined by the node's payoff matrix. If $S = \emptyset$, a new node is chosen according to the node's distribution over its offspring. The process now repeats itself, with the offspring node being the current node. When the players reach a leaf, the new current node is the root. A game on a tree is essentially played in rounds, where each round starts at the root and ends once it reaches a leaf. Formally:

Definition 13 A stopping game on a finite tree (or simply a game on a tree) is a tuple $T = (I, V, V_{leaf}, r, (C_v, p_v, R_v)_{v \in V \setminus V_{leaf}})$, where:

- I is a finite non-empty set of players;
- $(V, r, (C_v)_{v \in V \setminus V_{leaf}})$ 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 a nonempty set of offspring of v . We denote by $V_0 = V \setminus V_{leaf}$ the set of nodes which are not leaves;

and for every $v \in V_0$:

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

Given a bounded stopping time $n < \sigma$ and history $H_n \subseteq \mathcal{H}_n$, let $G_{n,\sigma}(H_n)$ be the induced stopping game that begins at stage n , when the players are informed of H_n , and the game restarts at stage n (where a new $\omega \in H_n$ is randomly chosen), if no player stopped before reaching stage $\sigma(\omega)$. A simple adaptation of the methods of Shmaya and Solan (2004, Sections 5-6) shows that $G_{n,\sigma}(H_n)$ can be approximated by a game on a tree, $T_{n,\sigma}(H_n)$, such that every ϵ -equilibrium in $T_{n,\sigma}(H_n)$ is a 3ϵ -equilibrium in $G_{n,\sigma}(H_n)$. In the following paragraph we sketch the main idea behind this approximation. The reader is referred to Shmaya and Solan (2004) for the formal details.

For simplicity of presentation let σ be constant: $\sigma = m > n$. All that matters to the players at stage m , is the payoff matrix at this stage (because if no player stops, the game restarts at stage n with a new random $\omega \in H_n$, which is independent of the information the players have on the current ω). Thus we can cluster together the \mathcal{F}_m -measurable sets according to their payoff matrices, and have at most $|D|$ leaves in the finite tree. At stage $m - 1$, players care about both the current payoff matrix and the distribution of the payoff matrices at the next stage. Using a finite approximation to this distribution (rounding each probability up to $\epsilon/2^m$), enables clustering of \mathcal{F}_{m-1} -measurable sets into a finite number of vertices as well. Similarly, one can show by a recursive procedure that the entire game $G_{n,\sigma}(H_n)$ can be approximated by a stopping game on a finite tree.

Assuming that $n > N_0$ we perturb the game on a tree $T_{n,\sigma}(H_n)$ by not allowing players to stop in any node \bar{v} with a payoff matrix $R_{\bar{v}}$ is in \bar{B}_d . That is, in such nodes, players must continue and the game goes on to one of \bar{v} 's offspring.

4.3 Equivalence of Periodic Games and Absorbing Games

A stopping game on a finite tree $T = T_{n,\sigma}(H_n)$ is equivalent to an absorbing game (Solan and Vohra, 2002; Sorin, 2002, 5.5), where each round of T corresponds to a single stage of the absorbing game. As an absorbing 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.

Given a game on a tree T , let g^i be the maximal payoff player i can get by stopping alone. Let \tilde{v}^i be a node that gives player i his maximal payoff g^i . Adapting Proposition 4.10 in Solan and Vohra (2002) to the two special properties gives the following:

Proposition 14 Let T be a game on a finite tree. One of the following holds:

- (1) There is a stationary absorbing sequential ϵ -equilibrium x .
- (2) There is a stationary non-absorbing sequential equilibrium where all the players always continue.
- (3) There is a distribution $\zeta \in \Delta(I)$ over the players such that:

- (a) For each player $j \in I$, $\mathbf{E}_{\zeta'}(R_{\{i\}, \tilde{v}^i}^j) = \sum_{i \in I} \zeta(i) \cdot R_{\{i\}, \tilde{v}^i}^j \geq g^j$, where ζ' denote the distribution over payoff vectors $\{R_{\{i\}, \tilde{v}^i}\}_{i \in I}$ that is induced from ζ as follows: player i is chosen according to ζ , and \tilde{v}^i is the node defined above. That is, we require that the expected payoff of each player j from the induced distribution ζ' is as high as his maximal payoff when stopping alone.
- (b) Let the players in the support of ζ ($\zeta(i) > 0$) be denoted as the *stopping players*. For every stopping player i there exists a player $j_i \neq i$, the *punisher* of i , such that: $g^i \geq R_{\{j_i\}, \tilde{v}^{j_i}}^i$. That is, each stopping player prefers to stop alone at \tilde{v}^i rather than having his punisher j_i stopping alone at \tilde{v}^{j_i} .

These two properties of ζ are used in Subsection 4.5 to construct a correlated equilibrium with payoffs that are induced by ζ' . The first property prevents players from deviating by stopping when they are not asked to stop, and the second property prevent players from deviating by continuing when they are asked to stop.

Remark 15 Solan and Vohra (2002) do not guarantee that the stationary absorbing equilibrium in case (1) is sequential. Specifically, players may play irrationally after some player i is supposed to stop with probability 1 according to x^i . To prevent this, we perturb the game T . Let T_ϵ be a game similar to T , except that when a non-empty set of players wishes to stop at some node, there is a probability ϵ that the “stopping request is ignored”, and the game continues to the next stage. T_ϵ is also equivalent to an absorbing game, and Solan and Vohra (2002)’s proposition can be applied. In T_ϵ no node is ever off the equilibrium path, and thus any Nash equilibrium in T_ϵ is subgame perfect, which is equivalent to being sequential, as the correlation device is trivial (as discussed after Definition 7). Any such stationary sequential equilibrium in T_ϵ naturally defines a strategy profile in T . One can see that this profile is a stationary sequential ϵ -equilibrium in T .

4.4 A Stochastic Variation of Ramsey’s Theorem

Solan and Shmaya (2004) present a stochastic variation of Ramsey’s theorem (Ramsey, 1930), and a method to use it to disassemble an infinite (non-periodic) stopping game into games on finite trees with special properties. In this subsection we sketch the main ideas of this method, while leaving some of the formal details to the appendix.

Let C be a finite set of “colors”. An \mathcal{F} -consistent C -valued NT-function (or simply an NT-function) is a function that attaches a color $c_{n,\sigma}(\omega) = c_{n,\sigma}(H_n(\omega))$ to every induced stopping game $G_{n,\sigma}(H_n(\omega))$. Given an NT-function and two bounded stopping times $\tau_1 < \tau_2$, let $c_{\tau_1, \tau_2}(\omega) = c_{\tau_1(\omega), \tau_2(\omega)}$. Thus c_{τ_1, τ_2} is an \mathcal{F}_n -measurable random variable. Shmaya and Solan (2004, Theorem 4.3) proved the following proposition :

Proposition 16 For every finite set C , every C -valued \mathcal{F} -consistent NT-function c , and every $\epsilon > 0$, there exists an increasing sequence of bounded stopping times $0 < \sigma_1 < \sigma_2 <$

$\sigma_3 < \dots$ such that: $p(c_{\sigma_1, \sigma_2} = c_{\sigma_2, \sigma_3} = \dots) > 1 - \epsilon$.

We now present a somewhat simplified version of the NT-function that would be used to prove Theorem 10; the exact function is described in the appendix.

Let $W = \prod_{i \in I} \{0, \pm \frac{1}{K}, \dots, \pm \frac{K}{K}\}$ be a finite $1/K$ -approximation of $[-1, 1]^{|I|}$. Let $C = \{\{1, 2, 3\} \times W \times W\}$ be a set of colors, where the first component denotes which case of Proposition 14 holds in $T_{n, \sigma}(H_n(\omega))$; the second component denotes the approximate equilibrium payoff, and the third component denotes the payoff of each player when he stops alone in case 3. That is, $c_{n, \sigma}(\omega) = (case, w, g)$ is defined as follows:

- $case = 1$ if there is a stationary absorbing equilibrium in $T_{n, \sigma}(H_n(\omega))$ (that is, case (1) of Proposition 14 holds). Otherwise, $case = 2$ if there is a sequential non-absorbing equilibrium in $T_{n, \sigma}(H_n(\omega))$. Otherwise, $case = 3$ and then case (3) of Prop. 14 holds.
- w is the equilibrium payoff in cases (1) and (2), and it is the payoff that is induced from the distribution η' in case (3): $w = \mathbf{E}_{\zeta'}(R_{\{i\}, \tilde{v}^i}^j) = \sum_{i \in I} \zeta(i) \cdot R_{\{i\}, \tilde{v}^i}^j$ (where \tilde{v}^i is a node that maximizes player i 's reward when stopping alone).
- g is the maximal payoff each player can get by stopping alone in $T_{n, \sigma}(H_n(\omega))$ in case (3), and it is arbitrarily set to 0 in cases (1) and (2).

By Proposition 16 there exists an increasing sequence of bounded stopping times $0 < \sigma_1 < \sigma_2 < \sigma_3 < \dots$ such that: $p(c_{\sigma_1, \sigma_2} = c_{\sigma_2, \sigma_3} = \dots) > 1 - \frac{\delta}{3}$. We assume w.l.o.g. that $\sigma_1 > N_0$. Let $E = \Omega \setminus (\Phi \cup \{\omega \in \Omega \mid \exists n \text{ s.t. } c_{\sigma_n, \sigma_{n+1}}(\omega) \neq c_{1,2}(\omega)\})$ be the event where there are no approximation errors (as defined in Subsection 4.1) and the color of all finite trees after σ_1 is the same. Observe that $P(E) > 1 - \frac{2}{3}\delta > 1 - \delta$.

4.5 Constant-Expectation Sequential Correlated Equilibrium

We conclude this section by proving Theorem 10: showing that every (non-periodic) stopping game admits a sequential (δ, ϵ) -constant-expectation normal-form correlated (δ, ϵ) -equilibrium with a canonical correlation device. By Lemma 12, Theorem 10 is implied by the following proposition:

Proposition 17 *Let E and σ_1 be defined as in the previous subsection. There is a canonical correlation device $\mathcal{D} = (M, \mu)$, and a subset $M' \subseteq M$ satisfying $\mu(M') > 1 - \delta$, such that for every $m \in M'$ and every $\omega \in E$, there is a sequential ϵ -constant-expectation ϵ -equilibrium conditioned on E and M' , x_H , in the game $G(H, \mathcal{D}, m)$, where $H = H_{\sigma_1(\omega)}(\omega)$.*

PROOF. Let $c = c_{\sigma_1, \sigma_2}(\omega) = (case, w, g)$ be the color of the game $G_{\sigma_1(\omega), \sigma_2}(H)$. Solan and Shmaya (2004) investigated 2-player stopping games, when $case$ is equal either to 1 or 2 ($case$ 3 is only relevant to games with more than two players). They show that one can concatenate the sequential stationary Nash $\epsilon/11$ -equilibria of each approximating game on

a tree $T_{\sigma_k(\omega), \sigma_{k+1}}(H_{\sigma_k(\omega)}(\omega))$ into a sequential ϵ -equilibrium (conditioned on E), x_H , in the induced game without pre-play correlation $G(H)$. The profile x_H naturally induces a sequential ϵ -constant-expectation ϵ -equilibrium conditioned on E and M' in $G(H, \mathcal{D}, m)$, given any correlation device \mathcal{D} and any signal profile m .

For this concatenation to work when *case* = 1, Solan and Shmaya (2004) provided appropriate minimal bounds to the probability of termination in the first round of the stationary approximate equilibrium of each game on a tree $T_{\sigma_k(\omega), \sigma_{k+1}}(H_{\sigma_k(\omega)}(\omega))$, that guarantee that the concatenated profile, x_H , is absorbed with probability 1. With minor adaptations, Shmaya and Solan (2004)'s method works also in multi-player stopping games, as described in the appendix.

Thus, we only have to deal with the third case (*case* = 3). The construction in this case is an adaptation of the procedure of Solan and Vohra (2002), which deals with quitting games (stationary stopping games where the payoff is the same at all stages). Changes with respect to the original procedure are needed to guarantee constant-expectation and sequentiality (which are not satisfied in Solan and Vohra, 2002).

For each player $i \in I$, let \tilde{v}^i be a node in the tree $T_{\sigma_1, \sigma_2}(H_{\sigma_1(\omega)}(\omega))$ that gives player i his maximal reward when stopping alone - g^i . The definition of $D(H)$ (the set of payoff matrices that repeats infinitely often in H) and $\alpha^i(H)$ (the maximal single-stopper payoff in $D(H)$ - see Subsection 4.1), implies that $g^i = \alpha^i(H)$, and that $R_{\tilde{v}^i} \in D(H)$ (the payoff matrix of each node \tilde{v}^i repeats infinitely often in the non-periodic infinite stopping game, assuming that $\omega \in E$). Let ζ be the distribution over the players that satisfies (Proposition 14): 3-a) $\sum_{i \in I} \zeta(i) \cdot R_{\{\tilde{v}^i\}}^i \geq g^i$, and 3-b) for each stopping player i there is a punisher - a stopping player j_i such that $g^i \geq R_{\{\tilde{v}^{j_i}\}}^{j_i}$.

Let $(\tau_n^i)_{i \in I, n \geq 1}$ be an increasing sequence of stopping times defined by induction: τ_1^i is the first stage m in which payoff matrix $R_{\tilde{v}^i}$ is reached - $R_m(\omega) = R_{\tilde{v}^i}$; and τ_{n+1}^i is the first stage $m > \max_{j \in I}(\tau_n^j)$ such that $R_m(\omega) = R_{\tilde{v}^i}$. Observe that in E each τ_n^i is bounded (because all the payoff matrices $(R_{\tilde{v}^i})_{i \in I}$ repeat infinitely often). Let $\tau_n = \max_{i \in I}(\tau_n^i)$. Intuitively, the stopping times $(\tau_n)_{n \geq 1}$ divide the infinite (non-periodic) stopping game into rounds. In each such round (assuming $\omega \in E$), the game passes at least once through each of the payoff matrices $(d(i))_{i \in I}$ that maximize the reward of a single stopper.

We now describe an auxiliary correlation device $\mathcal{D}_{D(H)}$. The device chooses a player to stop (the *stopper*) according to the distribution ζ . Let $T \in \mathbf{N}$ be chosen sufficiently large, and let $\hat{T} \in \mathbf{N}$ be chosen to be much greater than T . The alphabet of the correlation device includes $\hat{T} + T + 1$ integers: $\forall i \in I, M_{D(H)}^i = \{1, \dots, \hat{T} + T + 1\}$.

The signal sent to each player i is interpreted as the round in which that player should stop with probability $1 - \epsilon$ when reaching payoff matrix $R_{\tilde{v}^i}$ for the first time in that round.

The stopper receives a signal \hat{l} from the uniform distribution on the integers between 1 and \hat{T} . The punisher receives signal l from the uniform distribution on the integers from $\hat{l} + 1$ to $\hat{l} + T$. Finally, all other player receive the signal $l + 1$. If the game has passed through $\hat{T} + T + 1$ rounds, then the game returns to round 1. Formally, each player i with signal $m^i \in \{1, \dots, \hat{T} + T + 1\}$ stops with probability $1 - \epsilon$ at the first time that payoff matrix $d(i)$ is reached at each round n that satisfies $n = (m_i) \bmod (\hat{T} + T + 1)$.

This mechanism ensures that upon receiving the signal, with a large probability any player's estimate of the probability that he has been chosen as the stopper (the Bayesian posterior probability) is "virtually unchanged" from the prior probability. Formally, we require that with probability $1 - \frac{\delta}{2^{|D|}}$ the posterior probabilities of all players are changed by at most ϵ . Also, if the stopper deviates, the probability of him correctly predicting the moment of punishment is very small. Hence, given the others follow their signals, the stopper has no incentive to deviate. If the game is not stopped by the stopper, then at the time at which the punisher is supposed to stop, he believes with high probability that he is the stopper and so should stop according to the argument above.

Remark 18 *In our construction players are asked to stop with probability $1 - \epsilon$. This implies that no history is ever off the equilibrium path, and thus every equilibrium is sequential. It is possible to construct a similar equilibrium in which players are asked to stop with probability 1, by carefully defining players' beliefs off the equilibrium path.*

Let the canonical correlation device $\mathcal{D} = (M, \mu)$, which only depends on $|I|$ and ϵ , be the Cartesian multiplication of the correlation devices $\mathcal{D}_{D(H)}$ for each possible set of infinitely-repeated payoff matrices: $\mathcal{D} = \prod_{D(H) \subseteq D} \mathcal{D}_{D(H)}$. Let $M' \subseteq M$ be the set of signals such that for every player the posterior probability of being chosen as the stopper by the devices $(\mathcal{D}_{D(H)})_{D(H) \subseteq D}$ are changed by at most ϵ . Observe, that the above arguments imply that $\mu(M') > 1 - \delta$, and that the obedient strategy is a sequential ϵ -constant-expectation ϵ -equilibrium in the game $G(H, \mathcal{D}, m)$ conditioned on E and M' .

We conclude by demonstrating the use of our procedure in a simple example.

Example 19 *Consider the following periodic stopping game with 3 players. At stages $3k + 1$ (resp., $3k + 2$, $3k + 3$) If player 1 (resp., player 2, player 3) stops alone the payoff vector is $(1, 0, 5)$ (resp., $(5, 1, 0)$, $(0, 5, 1)$). If players 1 and 2 (resp., players 2 and 3, players 3 and 1) stop together, the payoff is $(0, 2, 0)$ (resp., $(0, 0, 2)$, $(2, 0, 0)$). If any other non-empty set of players stop, the payoff vector is $(0, 0, 0)$. That is, at each stage $3k + i$ player i gets 1 if he stops alone, and this yields 0 for player $(i + 1) \bmod 3$ and 5 for player $(i + 2) \bmod 3$. If Player $(i + 1) \bmod 3$ stops as well, he gets 2, while the other players get 0. Observe, that each player can get a maximal payoff of 1 by stopping alone ($g = (1, 1, 1)$), and that each player i has a punisher $j_i = (i + 2) \bmod 3$.*

In what follows we demonstrate how our procedure induces the payoff $(2, 2, 2) = \frac{1}{3}(1, 0, 5) + \frac{1}{3}(5, 1, 0) + \frac{1}{3}(0, 5, 1)$ as an approximate constant-expectation sequential equilib-

rium. In this example, the sequence of stopping times (τ_n^i) is as follows: $\tau_n^i = 3 \cdot (n - 1) + i$. This sequence divides the game into rounds of length 3: round 1 includes stages 1-3, round 2 includes stages 4-6, etc.

Say, for example, that the device chose player 1 as the stopper. Then player 1 receives signal $m^1 = \hat{l}$, his punisher, player 3, receives signal $m^3 = \hat{l} + l$, and player 2 receives signal $m^2 = \hat{l} + l + 1$. Assuming that the players follow their signals, player 1 stops with probability $1 - \epsilon$ when his optimal payoff (as a single stopper) is realized in the m^1 -th round (that is, at stage $\tau_{m^1}^1 = 3 \cdot (m^1 - 1) + 1$); player 3 (resp., player 2) stops with probability $1 - \epsilon$ when his optimal payoff is realized in round m^3 -th (resp., round m^2 -th) round (if the game has not terminated earlier); player 1 stops with probability $1 - \epsilon$ when his optimal payoff is realized in the $m^1 + (\hat{T} + T + 1)$ -th round, etc.

5 Extensions

Our formal model only dealt with “simple” stopping games, which end as soon as any player stops. We now discuss how to extend our result to more generalized strategic interactions, such as the leading example.

A *generalized stopping game* is played as follows. There is an unknown state variable, on which players receive symmetric partial information during play. For each player i , there is a finite number, T_i , that limits the number of actions he may take during the game. At each stage, each player i has a finite set of “stopping” actions A_i . At stage 1 all the players are active. At every stage n , each active player declares, independently of the others, whether he takes one of the “stopping” actions in A_i or continues. A player that has stopped T_i times, becomes passive for the rest of the game and must choose “continue” in all subsequent stages. The payoff of a player depends on the history of actions and on the state variable.

A generalized stopping game is different from a “simple” stopping game in three aspects: (1) if no player ever stops the payoff is not necessarily zero; (2) each player has a few different “stopping” actions ($|A_i| > 1$); (3) each player may act a finite number of times ($T_i > 1$) until he becomes passive, and when he becomes passive, the game continues with the other players.

Proposition 14 also holds when each player has a finite number of different “stopping” actions, and when the payoff if no player ever stops is different from zero. Thus, with minor adaptations, our proof is extended to cases (1) and (2).

The third case, where each player may act a finite number of times, is handled by using backward induction. The details are standard, and we only sketch here the main idea. Let $m = \sum_i T_i$ be the total number of times the players are allowed to stop. Assume

by induction on m , that any generalized stopping game where players can stop at most n times, admits an equilibrium of our type (sequential normal-form correlated approximate equilibrium with a canonical correlation device). Given a generalized stopping game G' with m “stops”, we construct an auxiliary stopping game G with the following payoff process: $R_{S,n}^i$ is equal to the payoff of player i in an equilibrium of our type of induced generalized stopping game with total number of stops $n - |S|$ that begins at stage $n + 1$, where the T_i of each player i in S is reduced by one. Such an equilibrium exists due to the induction hypothesis. By Theorem 10, the auxiliary game G admits an equilibrium of our type x . x induces an equilibrium of our type x' in the original game G' in a natural way: players follow x as long as all the players continue; as soon as some of the players stop, the remaining active players play the equilibrium of the induced stopping game with fewer “stops”.

Our result can also be extended to stopping games with voting procedures (see, e.g., Kurano, Yasuda and Nakagami, 1980, Yasuda, Nakagami and Kurano, 1982, and Szajowski and Yasuda, 1997). In such games, each player votes at each stage whether or not he wishes to stop the game, and there is some monotonic rule (for example, a majority rule) that determines whether the game stops or continues. Observe that unlike the above existing literature, we allow the payoff process to depend on the identity of the stopping players. The adaptation of our proof to this more general setup involves a single (non-minor) change: the absorbing game that is equivalent to a stopping game on a finite tree (Subsection 4.3) does no longer have a unique non-absorbing action profile. Nevertheless, Proposition 4.10 of Solan and Vohra (2002) can still be used (but in a more generalized way than Proposition 14, which assumes a unique non-absorbing profile), and an adaptation of the public signaling methods of Solan and Vohra allows to extend our result, and prove the existence of a correlated equilibrium of our type.

6 Distribution and Constant-Expectation Correlated Equilibrium

Sorin (1998) presented the notion of distribution equilibrium for finite normal-form games as a correlated equilibrium in which the expected payoff of each agent is independent of his signal. In Section 2 we generalized this notion for dynamic games with normal-form correlation, and called it constant-expectation correlated equilibrium. In this section we present basic properties of these notions, and discuss their rationales. The first two subsections are mostly based on Sorin (1998), and are given for completeness (Sorin (1998) is an unpublished manuscript, which is not readily available).

6.1 Properties and Examples

We briefly discuss some of the properties of distribution equilibrium in normal-form games. First, every Nash equilibrium is a distribution equilibrium. Second, unlike the set of correlated equilibria, the set of distribution equilibria is not convex, as demonstrated in the “battle of the sexes” game illustrated in Table 1: both (T, R) and (B, L) are distribution equilibria, but $[0.5(T, R), 0.5(B, L)]$ is not (the payoff of a player is either 1 or 2, depending on his signal).

Table 1
“Battle of the Sexes” - a Normal-Form Two-Player Game

	L	R
T	(0, 0)	(2, 1)
B	(1, 2)	(0, 0)

The next example (Table 2, adapted from Moulin and Vial, 1978) demonstrates that distribution equilibrium can induce payoffs that dominate the payoffs of Nash equilibria. The left table describes the payoff matrix. In this example, there is a unique Nash equilibrium in which each player plays $(1/3, 1/3, 1/3)$ with payoff $4/3$. The symmetric distribution equilibrium, which is described in the right table, induces payoff 2, and it dominates the Nash equilibrium payoff.

Finally, Table 3 (left table) presents a variant of the “Chicken” game (see, Aumann, 1974). There are three pure equilibria, (D, P) , (P, D) and (P, P) with payoffs $(0, 3)$, $(3, 0)$ and $(0, 0)$, respectively. In every distribution equilibrium, the payoff to both players is at most $4\frac{2}{3}$. Indeed, if, e.g., the row player does not play C , then at least one of the players must play P with probability 1 (otherwise the column player would deviate and play P with probability 1) and the payoff to both players is at most 3. If the row player plays C with positive probability, then the probability of playing (C, D) must be at least half the probability of playing (C, C) (otherwise the row player would deviate and play D instead of C), and the payoff to the row player is at most $4\frac{2}{3}$.

The middle table presents the best distribution equilibrium in this game - a symmetric

Table 2
Two-player Game with a Nash-Dominating Distribution Equilibrium

2-Player Game			Distribution Equilibrium			
	A	B	C	A	B	C
A	(0, 0)	(1, 3)	(3, 1)	0	1/6	1/6
B	(3, 1)	(0, 0)	(1, 3)	1/6	0	1/6
C	(1, 3)	(3, 1)	(0, 0)	1/6	1/6	0

Table 3

Variant of “Chicken” Game: Best Distribution and Correlated Equilibria

“Chicken” Game			Best Distribution Eq.			Correlated Eq.					
	<i>C</i>	<i>D</i>	<i>P</i>	<i>C</i>	<i>D</i>	<i>P</i>	<i>C</i>	<i>D</i>	<i>P</i>		
<i>C</i>	(6, 6)	(2, 7)	(0, 0)	<i>C</i>	4/9	2/9	0	<i>C</i>	1/2	1/4	0
<i>D</i>	(7, 2)	(0, 0)	(0, 3)	<i>D</i>	2/9	1/9	0	<i>D</i>	1/4	0	0
<i>P</i>	(0, 0)	(3, 0)	(0, 0)	<i>P</i>	0	0	0	<i>P</i>	0	0	0

Nash equilibrium that induces payoff $4\frac{2}{3}$ to both players; each player plays *C* with probability $\frac{2}{3}$ and *D* with probability $\frac{1}{3}$. The right table presents a symmetric non-distribution correlated equilibrium that induces payoff $5\frac{1}{4}$ to both players, which strictly dominates all distribution equilibria. Because this latter correlated equilibrium is preferred by both players over all distribution equilibria, one may wonder what is the advantage of a distribution equilibrium in this example. We answer this question in the following two sections.

6.2 Population Games

A common interpretation of Nash equilibrium is a description of equilibrium behavior of populations of agents who are randomly matched to play that game (see, e.g., Aumann, 1997). If each agent faces the same pattern of matching opponents, then an equilibrium in which each agent chooses a best reply corresponds to a Nash equilibrium of the underlying game. Maliath, Samuelson and Shaked (1997) relax the assumption of uniform matching pattern. They allow different types in the population to be matched to different opponents. In such a setup, an equilibrium in which each agent chooses a best reply (given his type’s pattern of matching opponents) is a correlated equilibrium of the underlying game.

Sorin (1998) changes the framework of Maliath, Samuelson and Shaked by allowing a deviating agent to “imitate” the matching pattern of another type: agent of type *i* is allowed to join the sub-population of type *j* and to follow their matching behavior. In such a setup, an equilibrium in which each agent chooses a best pattern among the existing matching patterns, and a best reply given this pattern, is a distribution equilibrium.

Non-distribution correlated equilibria are not stable in Sorin’s setup. Consider, for example, the best symmetric correlated equilibrium in the “Chicken” game (Table 3). The population includes two types: a “*d*” type ($\frac{1}{4}$ of the population) who is matched only to “*c*” opponents and always plays *D*, and a “*c*” type ($\frac{3}{4}$ of the population) who is matched to “*c*” opponents with probability $\frac{2}{3}$ and is matched to “*d*” opponents with probability $\frac{1}{3}$, and always plays *C*. If agents of one type are allowed to imitate the matching behavior of another type, then agents of type “*c*” (with payoff $4\frac{2}{3}$) would deviate and “imitate” the matching and playing behavior of type “*d*”, which has payoff 7.

In addition, non-distribution correlated equilibria are not stable in an evolutionary setup in which the type is determined at birth, and the payoff describes the fitness of each type. In such a setup, a type that has higher expected payoff will have higher number of offspring, and therefore his share of the population will increase. For example, in the “Chicken” game, the population’s share of type “ d ” would become larger than $\frac{1}{4}$ in the following generations.

6.3 Weak Mediators

One of the interpretations of a correlation device is a mediator. A mediator is a trusted third party that chooses an action profile according to a known (correlated) probability distribution, and privately informs each player of his part of the profile (a recommended action). The probability distribution is a correlated equilibrium if it is best-reply for each player to follow his recommended action, given that all other players follow their recommended actions.

In some situations, mediators are *weak* in the sense that a player who receives a “bad” recommended action (which induces a low expected payoff) has the ability to restart the mediation process. Some examples for such situations are:

- A married couple (say, Alice and Bob) goes to a marriage counselor. If Alice is discontent from the recommendations the counselor gave her, she may ask Bob to go to another counselor. It is plausible that Bob would agree to this request, which restarts the mediation process.
- Two countries in dispute ask a powerful third country to suggest an outline for a peace conference. Such an outline may include confidential parts, such as a monetary aid given to one side for his agreement to participate in the conference. The third country confidentially informs each disputing country on its part of the outline. Each disputing country can refuse the suggested outline. In that case, the outline is canceled and the disputing countries go back to the starting position, and they may restart the peace initiative with a new mediator.

In such situations, distribution equilibria have an important advantage: they can be implemented by weak mediators without having any player wishing to restart the mediation process. On the other hand, the implementation of non-distribution correlated equilibrium is limited by players’ ability to restart the mediation. The concept of weak mediators, and its relation with pre-play communication, is more thoroughly discussed in Heller (2010b).

6.4 Dynamic Games with Normal-Form Correlation

The above rationales, presented for distribution equilibria in normal-form games, are also appropriate to our notion of constant-expectation correlated equilibria in dynamic games with normal-form correlation. In the spirit of these rationales, our definition requires that the payoff of each player is independent of the signal *before the game starts*, when it is still possible to restart the pre-play process that induces the correlated profile. Observe, that we allow that later in the game, after some signals are received (e.g., the realization of the payoff matrices in a stopping games), a player may find out that his expected continuation payoff has changed, and is different than his original expected payoff.

Appendix

In Section 4 we presented a simplified version of the coloring scheme that is used in the construction of the concatenated equilibrium. In this appendix we present the exact coloring scheme, and show how to adapt Solan and Shmaya (2004)'s methods to give appropriate lower bounds for the termination probabilities in case (1) of Proposition 14.

A.1 Limits on Per-Round Probability of Termination

In this subsection we bound the probability of termination in a single round of a game on a tree when an absorbing stationary equilibrium x exists (case (1) of Prop. 14), by adapting the methods presented in Shmaya and Solan (2004, Section 5) for two players.

A stationary strategy of player i in a game on a tree T is a function $x^i : V_0 \rightarrow [0, 1]$ (recall that $V_0 = V \setminus V_{leaf}$ is the set of nodes that are not leaves; $x^i(v)$ 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 $\gamma^i(x) = \gamma_T^i(x)$ be the expected payoff under x , and let $\pi(x) = \pi_T(x)$ be the probability that the game is stopped at the first round (before returning to the root). 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 paths 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 path reached passes through \hat{V} . For each $v \in V$, we denote by F_v the event that the path reached passes through v .

The following lemma bounds 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 Shmaya and Solan (2004), and the proof is omitted as the changes are minor.

Lemma 20 *Let G be a stopping game, $n > 0$, $\sigma > n$ a bounded stopping time, $H \in \mathcal{H}_n$*

a history, and x an absorbing stationary $\frac{\epsilon}{2}$ -equilibrium in $T_{n,\sigma}(H_n)$ such that there exists a player i with a low payoff: $\gamma^i(x) \leq \alpha_H^i - \epsilon$. Then $\pi(c^i, x^{-i}) \geq \frac{\epsilon}{6} \cdot q^i$, where $q^i = q_T^i = p\left(\bigcup_{v \in V_{stop}} \{F_v | R_{\{i\},v}^i = \alpha_H^i\}\right)$ is the probability that if no player ever stop, the game visits a node $v \in V_0$ with $R_{\{i\},v}^i = \alpha_H^i$ in the first round.

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

Definition 21 Let $T = (I, V, V_{leaf}, r, (C_v, p_v, R_v)_{v \in V \setminus V_{leaf}})$ and let $T' = (I, V', V'_{leaf}, r', (C'_v, p'_v, R'_v)_{v \in V'_0})$ be two games on trees. We say that T' is a *subgame* of T if: $V' \subseteq V$, $r' = r$, and for every $v \in V'_0$, $C'_v = C_v$, $p'_v = p_v$ and $R'_v = R_v$.

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 histories \mathcal{H}_n into two kinds: *simple* and *complicated*. A simple history has at least one of the following properties: (1) Every player 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, such that each player receives payoff α_H^i when he stops, and approximately this is also his average payoff when other players stop.

Definition 22 Let G be a stopping game, $\epsilon > 0$, $N_0 \leq n$, and $\tau > n$ a bounded stopping time. The history $H \in \mathcal{H}_n$ is ϵ -*simple* if one of the following holds:

- (1) For every $i \in I$: $\alpha_H^i < 0$. or
- (2) There is a distribution $\theta \in \Delta(D_H \times I)$ such that for each player $i \in I$:
 - (a) $\theta(d, i) > 0 \Rightarrow R_{\{i\},d}^i = \alpha_H^i$. and
 - (b) $\alpha_H^i + \epsilon \geq \sum_{j \in I, d \in D_H} \theta(d, j) \cdot R_{\{j\},d}^i \geq \alpha_H^i - \epsilon$.

H is *simple* if it is ϵ -*simple* for every $\epsilon > 0$. H is *complicated* if it is not simple, i.e.: $\exists \epsilon_0 > 0$ such that H is not ϵ_0 -simple. In that case we say that H is complicated w.r.t. ϵ_0 .

The next proposition analyzes stationary ϵ -equilibria that yield high payoffs to all the players. The proposition is an adaptation of Proposition 5.5 in Shmaya and Solan (2004). The proof is omitted as the changes are minor.

Proposition 23 Let G be a stopping game, $N_0 \leq n$ a number, $\sigma > n$ a bounded stopping time, $H \in \mathcal{H}_n$ a complicated history w.r.t. ϵ_0 , $\epsilon \ll \frac{\epsilon_0}{|I| \cdot |D|}$, and for each $i \in I$ let $a^i \geq \alpha_F^i - \epsilon$. Then there exists a set $U \subseteq V_0$ and a profile x in $T = T_{n,\sigma}(F)$ such that:

- (1) No subgame of T_U has a Nash 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 Nash 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 $\pi(x) \geq \epsilon^2 \cdot p_{T_U, T}$.

A.2 Detailed Description of The Coloring Scheme

In Subsection 4.4 we presented a simplified version of the coloring scheme that is used in the proof of Proposition 17. In this subsection, we present the details of the exact coloring scheme, which adapts the coloring scheme for two-player games in Shmaya-Solan (2004). Specifically, we provide an algorithm that attaches a color $c_{n,\sigma}(H)$ and several numbers $(\lambda_{j,n,\sigma}(H))_j$ for every $\sigma > n \geq 0$ and $H \in \mathcal{H}_n$, such that $c_{n,\sigma}(H)$ 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 $i \in I$, $\alpha_H^i - \epsilon \leq a^i$. It is *good* if there exists a player $i \in I$ such that $a^i + \epsilon \leq \alpha_H^i - \epsilon$. 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 $B = \{b_1, b_2, \dots, b_J\}$ be the set of J bad rectangles in W and let $O = \{o_1, o_2, \dots, o_K\}$ be the set of good rectangles.

Set $C = (\text{simple} \cup \text{allbad} \cup \{1 \times O\} \cup \{2\} \cup \{3 \times W \times W\})$. Let G be a stopping game, $n \geq 0$, $\sigma > n$ a bounded stopping time, and $H \in \mathcal{H}_n$. If H is simple we let $c_{n,\sigma}(H) = \text{simple}$. Otherwise, H is *complicated* w.r.t. to some $\epsilon_0(H)$. In that case we assume w.l.o.g. that $\epsilon \ll \frac{\epsilon_0(H)}{|I| \cdot |D|}$. The color $c_{n,\sigma}(H)$ is determined by the following procedure:

- Set $T^{(0)} = T_{n,\sigma}(H)$.
- For $1 \leq j \leq J$ apply Proposition 14 to $T^{(j-1)}$ and the bad rectangle $h_j = \prod_{i \in I} [a_j^i, a_j^i + \epsilon]$ to obtain a subgame $T^{(j)}$ of $T^{(j-1)}$ and strategy profile x_j in $T^{(j)}$ such that:
 - (1) No subgame of $T^{(j)}$ has a stationary ϵ -equilibrium with a corresponding payoff in h_j .
 - (2) Either $T^{(j)} = T^{(j-1)}$ or the following three conditions hold:
 - (a) For every $i \in I$, $a_j^i - \epsilon \leq \gamma^i(x_j)$.
 - (b) For every $i \in I$ and every strategy y^i : $\gamma^i(x_j^{-i}, y^i) \leq a_j^i + 8\epsilon$.
 - (c) $\pi(x_j) \geq \epsilon^2 \times p_{T^{(j)}, T^{(j-1)}}$.
- If $T^{(J)}$ is trivial (the only node is the root), set $c_{n,\sigma}(H) = \text{allbad}$; otherwise due to Proposition 14 and our procedure one of the following holds:
 - (1) $T^{(J)}$ has a sequential stationary absorbing ϵ -equilibrium x , with a payoff $\gamma(x)$ in one of the good hyper-rectangles. Let $c_{n,\sigma}(H) = (1, o_l)$, where o_l is the good rectangle that includes γ_x .
 - (2) $T^{(J)}$ has a sequential stationary non-absorbing equilibrium c , with a payoff 0. Let $c_{n,\sigma}(H) = (2)$.
 - (3) There is a correlated strategy profile $\eta \in \Delta(A)$ in $T^{(J)}$ that satisfies 3(a)+3(b)+3(c)

in Proposition 14. Let $c_{n,\sigma}(H) = (3, w_1, w_2)$ 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 Proposition 14, is a profile in $T^{(j-1)}$. We consider it 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 Proposition 16 there exists an increasing sequence of bounded stopping times $0 < \sigma_1 < \sigma_2 < \sigma_3 < \dots$ such that $p(c_{\sigma_1, \sigma_2} = c_{\sigma_2, \sigma_3} = \dots) > 1 - \frac{\delta}{3}$. For every $\omega \in \Omega$ and $H = H(\omega) \in \mathcal{H}_{\sigma_1(\omega)}$, let $c_H = c_{\sigma_1, \sigma_2}(H)$.

Let $(A_{\epsilon,j}, A_{\infty,j})_{j \in J} \in \bigvee_{n=1.. \infty} \mathcal{F}_n$ be defined as follows: $A_{\infty,j} = \left\{ \omega \in \Omega \mid \sum_{k=1.. \infty} \lambda_{j, \sigma_k, \sigma_{k+1}}(H_{\sigma_k(\omega)}(\omega)) = \infty \right\}$ is the event where the sum of the λ -s is infinite, and $A_{\epsilon,j} = \left\{ \omega \in \Omega \mid \sum_{k=1.. \infty} \lambda_{j, \sigma_k, \sigma_{k+1}}(F_{\sigma_k(\omega)}) \leq \frac{\epsilon}{|J|} \right\}$ is the event where the sum is very small. 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 $(\bar{A}_{\epsilon,j}, \bar{A}_{\infty,j})_{j \in J} \in \mathcal{F}_{N_1}$ that approximate $A_{\infty,j}$ and $A_{\epsilon,j}$: (1) For each $j \in J$, $\bar{A}_{\epsilon,j} \cap \bar{A}_{\infty,j} = \emptyset$ and $(\bar{A}_{\epsilon,j} \cup \bar{A}_{\infty,j}) = \Omega$. (2) $p(A_{\epsilon,j} \mid \bar{A}_{\epsilon,j}) \geq 1 - \frac{\delta}{6 \cdot |J|}$. (3) $p(A_{\infty,j} \mid \bar{A}_{\infty,j}) \geq 1 - \frac{\delta}{6 \cdot |J|}$. From now on, we assume w.l.o.g. that $\sigma_1 \geq N_1$. Let E' be defined as follows (Observe that $p(E') \geq 1 - \delta$):

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

That is, E' is equal to E (defined in Subsection 4.4), except that we subtract the errors in the approximations of $(A_{\epsilon,j}, A_{\infty,j})_{j \in J}$ by $(\bar{A}_{\epsilon,j} \cup \bar{A}_{\infty,j})_{j \in J}$.

A.3 Detailed Proof of Cases 1 and 2 of Proposition 17

In Subsection 4.5 we gave the details of the proof of Proposition 17 only when *case* = 3. In this subsection we give the details of the proof for the other cases, which are adaptations of the proof for the two-player case in Shmaya and Solan (2004). The proof is divided to 5 exhaustive cases according to the color of c_H and whether $H \cap \bar{A}_{\infty,j} \neq \emptyset$.

A.3.1 There exists $j \in J$ and $F \in H$ such that $F \subseteq \bar{A}_{\infty,j}$

Let $1 \leq j \leq J$ be the smallest index such that $F \subseteq \bar{A}_{\infty,j}$. Let $x_{j, \sigma_k, \sigma_{k+1}}$ be the j^{th} profile in the procedure described earlier, when applied to $T_{\sigma_k, \sigma_{k+1}}(H)$. Let x_H be the following strategy profile in $G(H, \mathcal{D}, m)$: between σ_k and σ_{k+1} play according to $x_{j, \sigma_k, \sigma_{k+1}}$. The procedure of the previous subsection 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: $a_j^i - \epsilon \leq \gamma_{\sigma_k,\sigma_{k+1}}^i(x_j) \leq a_j^i + 8\epsilon$.
- For each player $i \in I$ and for each strategy y^i in $T_{\sigma_k,\sigma_{k+1}}$: (1) $\gamma_{\sigma_k,\sigma_{k+1}}^i(x_j^{-i}, y^i) \leq a_j^i + 8\epsilon$.
(2) $\pi_{\sigma_k,\sigma_{k+1}}(x_j) \geq \epsilon^2 \times \lambda_j(T_{\sigma_k,\sigma_{k+1}})$

These facts imply that the game is absorbed with probability 1 in E' , and that x_F is a 11ϵ -equilibrium conditioned on E' . Observe that $c_H = \text{allbed}$ implies that there exists $j \in J$ and $F \in H$ such that $F \in \bar{A}_{\infty,j}$.

A.3.2 There exists $F \in H$ such that $F \subseteq \left(\bigcap_{j \in J} \bar{A}_{\epsilon,j}\right)$ and $c_H = 2$

Let x_H 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' .

A.3.3 There exists $F \in H$ such that $F \subseteq \left(\bigcap_{j \in J} \bar{A}_{\epsilon,j}\right)$ and $c_H = (1, o_k) \in (1 \times O)$

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 $o_w: \prod_{i \in I} [a_w^i, a_w^i + \epsilon]$. As o_w is good, there is a player $i \in I$ such that: $a_w^i \leq \alpha_H^i - 2\epsilon$. Let x_H be the following strategy profile in G_H : between σ_k and σ_{k+1} play according to $x_{\sigma_k,\sigma_{k+1}}$. Lemma 20 implies that $\pi(c^i, x_{\sigma_k,\sigma_{k+1}}^{-i}) \geq \frac{\epsilon}{6} \cdot q_{\sigma_k,\sigma_{k+1}}^i$, where $q_{\sigma_k,\sigma_{k+1}}^i = p(\exists \sigma_k \leq n < \sigma_{k+1}, R_{i,n}^i = \alpha_F^i, R_{i,n}^i \in D_F)$. In E' , $R_{i,n}^i = \alpha_F^i$ infinitely often and $\sum_{j=1 \dots J} \sum_{k=1 \dots \infty} \lambda_{j,\sigma_k,\sigma_{k+1}} < \epsilon$. This implies that under x_H the game is absorbed with probability 1, and that x_H is a 4ϵ -equilibrium in G , conditioned on E' .

A.3.4 There exists $F \in H$ such that $F \subseteq \left(\bigcap_{j \in J} \bar{A}_{\epsilon,j}\right)$ and $c_H = (3, w_1, w_2) \in (1 \times W \times W)$

This case was thoroughly presented in Subsection 4.5.

A.3.5. $c_H = \text{simple}$

If for every $i \in I$: $\alpha_H^i \leq 0$, then the profile in which all the players always continue is an equilibrium in E' . Otherwise, the fact that $c_H = \text{simple}$ implies that there is a distribution $\theta \in \Delta(D_H \times I)$ such that for each $i \in I$: (1) $\theta(d, i) > 0 \Rightarrow R_{\{i\},d}^i = \alpha_H^i$. (2) $\alpha_H^i + \epsilon \geq \sum_{j \in I, d \in D_F} \theta(d, j) \cdot R_{\{j\},d}^i \geq \alpha_H^i - \epsilon$. In this case, one can use a procedure similar to the one described in Subsection 4.5, to construct a sequential ϵ -equilibrium in $G(H, \mathcal{D}, m)$ conditioned on E' and M' .

Acknowledgements

This research is in partial fulfillment of the requirements for the Ph.D. in Mathematics at Tel-Aviv University, and it was supported by the Israel Science Foundation (grant number 212/09). I would like to express my deep gratitude to Eilon Solan for his careful supervision, for the continuous help he offered, and for many insightful discussions. I would like also to thank Ayala Mashiah-Yaakovi, the associate editor, and the anonymous referees for many useful comments.

References

- Assaf, D., E. Samuel-Cahn, 1998. Optimal cooperative stopping rules for maximization of the product of the expected stopped values. *Stat. Probab. Lett.* **38** 89-99.
- Assaf, D., E. Samuel-Cahn, 1998. Optimal multivariate stopping rules. *J. Appl. Probab.* **35** 693-706..
- Aumann, R.J. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* **1** 67-96.
- Aumann, R.J., M. Maschler. 1995. Repeated games with incomplete information. The MIT press.
- Aumann, R.J. 1997. Rationality and Bounded Rationality. *Games and Economic Behavior* **21** 2-14.
- Ben-Porath, E. 1998. Communication without mediation: expanding the set of equilibrium outcomes by “cheap” pre-play procedures. *Journal of Economic Theory* **80** 108-122.
- Bernstein, F., A. Federgruen, 2004. A general equilibrium model for industries with price and service competition. *Operations Research* **52** 868-886.
- Bouakiz, M., M.J. Sobel, 1992. Inventory control with an exponential utility criterion. *Operations Research* **40** 603-608.
- Bulow, J., P. Klemperer. 2001. The generalized war of attrition. *American Economic Review* **89** 175-189.
- Chatterjee K., L. Samuelson, 1988. Bargaining under two-Sided incomplete information: the unrestricted offers case. *Operations Research* **36** 605-618.
- Christie-David, R., M. Chaudhry, W. Khan. 2002. News releases, market integration, and market leadership. *Journal of Financial Research* **XXV** 223-245.

- Dhillon, A., J.F. Mertens. 1996. Perfect correlated equilibria. *Journal of Economic Theory* **68** 279-302.
- Dynkin, E.B. 1969. Game variant of a problem on optimal stopping. *Soviet mathematics - Doklady*. **10** 270-274.
- Ferguson, T.S. 2005. Selection by Committee. A.S. Nowak, K. Szajowski, eds. *Advances in Dynamic Games: Applications to Economics, Finance, Optimization, and Stochastic Control, Annals of the International Society of Dynamic Games*, Vol. 7, Part III, Birkhäuser Boston, 203-209.
- Forges, F. 1986. An Approach to Communication Equilibria. *Econometrica* **54** 1375-1385.
- Fudenberg, D., J.Tirole. 1985. Preemption and Rent Equalization in the Adoption of New Technology. *Review of Economic Studies* **LII** 383-401.
- Fudenberg, D., J.Tirole. 1986. A theory of exit in duopoly. *Econometrica* **54** 943-960.
- Glickman, H., 2004. Cooperative stopping rules in multivariate problems. *Sequential Anal.* **23** 427-449.
- Heller, Y. 2010a. Minority-proof cheap-talk protocol. *Games and Economic Behavior* **69**(2) 394-400.
- Heller, Y. 2010b. Comment on Distribution Equilibria.
<http://www.tau.ac.il/~helleryu/distribution.pdf>
- Karlin, S. 1959. *Mathematical Methods and Theory in Games, Programming and Economics*, Vol. 2, Addison-Wesley, Reading, MA.
- Kreps, M.D., R. Wilson. 1982. Sequential Equilibria. *Econometrica* **50** 863-894.
- Krishna, V., J. Morgan. 1997. An Analysis of the War of Attrition and the All-Pay Auction. *Journal of Economic Theory* **72** 343-362.
- Kurano, M., M. Yasuda, J. Nakagami, 1980. Multi-variate stopping problem with a majority rule, *J. Oper. Res. Soc. Jap.* **23** 205-223.
- Laraki, R. E. Solan, 2005. The value of zero-sum stopping games in continuous time. *SIAM J. Control Optim.* **43** 1913-1922.
- Maliath, J.G., Larry Samuelson, Avner Shaked, 1997. Correlated equilibria and local interactions. *Economic Theory* **9** 551-556.
- Mamer, J.W. 1987. Monotone stopping games. *Journal of Applied Probability* **24** 386-401.
- Mashiah-Yaakovi, A. 2009. Subgame perfect equilibria in stopping games. working paper,

School of Mathematical Sciences, Tel-Aviv University.

[http://www.math.tau.ac.il/~ayalam/publication.files/general stopping game 3.pdf](http://www.math.tau.ac.il/~ayalam/publication.files/general%20stopping%20game%203.pdf)

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

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

Moulin, H., J. P. Vial. 1978. Strategically Zero-Sum Games: The Class of Games Whose Completely Mixed Equilibria Cannot be Improved Upon, *International Journal of Game Theory* **7** 201-221.

Myerson, R. B. 1986a. Multistage Games with Communication. *Econometrica* **54** 323-358.

Myerson, R.B. 1986b. Acceptable and predominant correlated equilibria. *International Journal of Game Theory* **15** 133-154.

Nalebuff, B., J.G. Riley. 1985. Asymmetric Equilibria in the War of Attrition. *Journal of Theoretical Biology* **113** 517-27.

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

Neveu, J. 1975. *Discrete-Parameters Martingales*. Borth-Holland, Amsterdam.

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

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

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

Ohtsubo, Y. 1995. Pareto optimum in a cooperative Dynkin's stopping problem. *Nihonkai Math. J.* **6** 135-151.

Ohtsubo, Y. 1996. Core in a cooperative Dynkin's stopping problem. *RIMS Kokyuroku* **947** 13-21.

Ohtsubo, Y. 1998. Pareto optima in a multi-person cooperative stopping problem. *RIMS Kokyuroku* **1043** 184-191.

Osborne, M. J., A. Rubinstein. 1994. *A Course in Game Theory*. The MIT Press.

- Porteus, E.L. 1975. On the optimality of structured policies in countable stage decision processes. *Management Sci.* **22** 148-157.
- Ramsey, D.M. 2007. Correlated equilibria in n-player stopping games. *Scientiae Mathematicae Japonicae*, 66(1): 149-164.
- Ramsey, D., D. Cierpial, 2009. Cooperative strategies in stopping games. Vladimir Gaitsgory Pierre Bernhard and Odile Pourtallier, eds. *Advances in dynamic games: applications to economics, finance, optimization, and stochastic control, Annals of the International Society of Dynamic Games*, Vol. 10, Birkhuser, Boston, 415-430.
- Ramsey, D., K. Szajowski, 2008. Selection of a correlated equilibrium in Markov stopping games., *Eur. J. Oper. Res.* **184** 185-206.
- Ramsey, F. 1930. On a problem of formal logic. *Proceedings of the London Mathematical Society* **30** 264-286.
- Rosenberg, D., E. Solan, N. Vieille. 2001. Stopping games with randomized strategies. *Probability Theory and Related Fields* **119** 433-451.
- Rubinstein, A. 1991. Comments on the Interpretation of Game Theory. *Econometrica* **59** 909-924.
- Szajowski, K. 2002, On stopping games when more than one stop is possible. V.F.Kolchin, V.Y. Kozlov, V.V. Mazalov, Y.L. Pavlov and Y.V. Prokhorov, eds. *Probability Methods in Discrete Mathematics, Proceedings of the Fifth International Petrozavodsk Conference*, May 2000. International Science Publishers, Leiden, Netherlands, 57-72.
- Szajowski, K., M. Yasuda, 1996. Voting procedure on stopping games of Markov chain. A.H. Christer, S. Osaki, and L.C. Thomas, eds. *UK-Japanese Research Workshop on Stochastic Modeling in Innovative Manufacturing, July 21-22, 1995, Moller Center, Churchill College, University of Cambridge, UK, Lecture Notes in Economics and Mathematical Systems*, Vol. 445, Springer, 68-80.
- Selten, R. 1965. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft* **12** 301-324 and 667-689.
- Selten, R. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* **4**, 25-55.
- Shmaya, E., E. Solan. 2004. Two-player nonzero-sum stopping games in discrete time. *Annals of Probability* **32** 2733-2764.
- Shmida, A., B. Peleg. 1997. 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, 149-170.

Solan, E., N. Vieille. 2001. Quitting games, *Mathematics of Operations Research* **26** 265-285.

Solan, E., R.V. Vohra. 2001. Correlated Equilibrium in Quitting Games, *Mathematics of Operations Research* **26** 601-610.

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

Sorin, S. 1998. Distribution equilibrium I: definition and equilibrium. Papers 9835, Paris X - Nanterre, U.F.R. de Sc. Ec. Gest. Maths Infor.

Sorin, S. 2002. A first course on zero-sum repeated games, Series: Mathematics and Applications, Springer, Paris.

Taylor, T.E., E.L. Plambeck, 2007. Supply chain relationships and contracts: the impact of repeated interaction on capacity investment and procurement, *Management Science* **53** 1577-1593.

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

Yasuda, M., J. Nakagami, M. Kurano, 1982. Multi-variate stopping problem with a monotone rule, *J. Oper. Res. Soc. Jap.* **25** 334-350.

Yasuda, M., K. Szajowski, 2002. Dynkin games and its extension to a multiple stopping model, Bulletin of the Japan Society for Industrial Mathematics **12** (3) 17-28 (Japanese).