Available online at www.sciencedirect.com
SCIENCE DIRECT®

Journal of Economic Theory 120 (2005) 206–238

JOURNAL OF Economic Theory

http://www.elsevier.com/locate/jet

Continuous-time games of timing

Rida Laraki,[a] Eilon Solan,[b,c] and Nicolas Vieille[d,*]

[a] CNRS and Laboratoire d'Econométrie de l'Ecole Polytechnique, 1, rue Descartes, 75005 Paris, France
[b] MEDS Department, Kellogg School of Management, Northwestern University, USA
[c] School of Mathematical Sciences, Tel Aviv University, Israel
[d] Département Finance et Economie, HEC School of Management (GREQHEC), 1, rue de la Libération, 78 351 Jouy-en-Josas, France

Received 1 June 2003; final version received 5 February 2004
Available online 10 June 2004

Abstract

We address the question of existence of equilibrium in general timing games with complete information. Under weak assumptions, any two-player timing game has a Markov subgame perfect ε-equilibrium, for each $\varepsilon > 0$. This result is tight. For some classes of games (symmetric games, games with cumulative payoffs), stronger existence results are established.

JEL classification: C72; C73

Keywords: Timing games; Continuous-time games; Games of timing; War of attrition; Preemption games

0. Introduction

Many economic and political interactions revolve around timing. A well-known example is the class of war of attrition games, in which the decision of each player is when to quit, and the game ends in the victory of the player who held on longer. These games were introduced by Maynard Smith [19], and later analyzed by a number of authors. Hendricks et al. [15] provide a characterization of equilibrium payoffs for complete information, continuous time wars of attrition played over a compact time interval. Several models that resemble wars of attrition were studied in the literature. Ghemawat and Nalebuff [14] analyze the exit decision of two

*Corresponding author.
E-mail addresses: laraki@poly.polytechnique.fr (R. Laraki), solan@kellogg.northwestern.edu, eilons@post.tau.ac.il (E. Solan), vieille@hec.fr (N. Vieille).

0022-0531/$ - see front matter
doi:10.1016/j.jet.2004.02.001

competing firms in a declining market, and assume that the market will eventually not be profitable if none of the two firms ever drops from the market (see also [7]). Fudenberg and Tirole [12] look at an incomplete information setup, in which there is a small probability that either firm will find it dominant to stay in forever. More recently, Bilodeua and Slivinski [4] studied a model where a volunteer for a public service is needed, while Bulow and Klemperer [5] consider multi-player auctions as generalized wars of attrition.

Another important class of timing games are preemption games, in which each player prefers to stop first. The analysis is then sensitive to the specification of the payoff, were the two players to stop simultaneously, see [11,13, pp. 126–128].

Yet another class of timing games consists of duel games. These are two-player zero-sum games. In the simplest version, both players are endowed with one bullet, and have to choose when to fire. As time goes, the two players get closer and the accuracy of their shooting improves. These games are similar to preemption games in that a player who decides to act may be viewed as preempting her opponent. However, as opposed to preemption games, in duel games a player has no guarantee that firing first would result in a victory. We refer the reader to Karlin [16] for a detailed presentation of duel games, and to Radzik and Raghavan [24] for an updated survey.

There are many timing games that do not fall neatly into any of these known categories. Consider for instance the standard case of a declining market, with two initially present firms. If the monopoly profits in that market are not decreasing—e.g. if the market has a cyclical component—or if the monopoly profits remain consistently above the outside option, the game fails to be a war of attrition (see [13, p. 122]). In another setup, when two firms compete on the patenting or the introduction of new technology, their interaction has the flavor of a preemption game. But each such firm also has an incentive to wait, since the probability of higher payoffs increases with time (and, presumably, with product quality). LaCasse et al. [17] studied a model where volunteers for several jobs are needed. When only one volunteer is needed, the model reduces to a standard war of attrition, but when there are several jobs, the strategic considerations are more complex.

The present paper addresses the question of existence of equilibrium in general timing games. It provides a framework that includes all timing games discussed in the literature, together with many other, and a unified analysis of all these games.

For our purposes, a continuous-time game of timing is described by a set I of players, and, for each non-empty subset of players $S \subseteq I$, a function $u_S : [0, \infty) \to \mathbf{R}^I$, with the interpretation that $u_S(t)$ is the payoff vector if the players in S—called the *leaders*—are the first to act, and they do so at time t. In addition, player i's time preferences are described by a discount rate δ_i.[1]

Our first result is a general existence result for *two-player* games: assuming $u_S(\cdot)$ is continuous and bounded for each S, a Markov subgame-perfect ϵ-equilibrium exists, for each $\epsilon > 0$.

[1] The model that is described here is of a game with complete information. We shall argue that some of our results extend to games with symmetric incomplete information.

This general existence result does not hold with more than two players. Nevertheless, an existence result can still be established for most cases of economic interest. As an illustration, we consider two such classes of games.

For *symmetric* games, our existence result is valid irrespective of the number of players. Moreover, the corresponding strategy profile is *pure*. Hence, any symmetric timing game has a Markov subgame perfect ε-equilibrium in pure strategies. However, a symmetric ϵ-equilibrium need not exist.

In most cases of interest, the payoff of a player who acts at time t can be written as the sum of a payoff incurred up to t and of an outside option. As a consequence, the payoff to such a leader is *independent* of the identity of the other leaders. We call these games *games with cumulative payoffs*. For such games, our existence result is valid for any number of players.

A point of interpretation is worth stressing here. In some applications once a player drops from the game, the interaction continues among the remaining players. Our existence result allows for the analysis of such games, using backward induction and applying the existence result inductively. Specifically, the payoff $u_S^i(t)$, for $i \notin S$, should rather be interpreted as the sum of the payoff accrued to i up to time t, and of an equilibrium payoff to player i in the continuation game—the timing game that starts at time t and with set of players $I \backslash S$. Our technique can be used to study interactions in which each player can act a bounded number of times, and the payoff depends on past and current behavior of the players.

Most of our proofs are constructive. In addition, our existence results are tight. Indeed, we exhibit a two-player zero-sum game with no exact Nash equilibrium, and a three-player zero-sum game with no Nash ϵ-equilibrium, provided ϵ is sufficiently small. In these two examples, payoffs are constant over time.

Finally, we provide a restrictive condition under which existence of a Nash ϵ-equilibrium for every $\epsilon > 0$ implies the existence of an exact equilibrium. The condition is that the function u_S is constant for each $S \subseteq I$, and that players are not discounting payoffs (but we do not impose any restriction on the number of players). Incidentally, this establishes the existence of a Nash equilibrium for the corresponding class of two-player timing games, a class of games for which none of the known sufficient conditions for equilibrium existence hold (see, e.g., [25]).

We conclude the introduction by discussing a few conceptual issues. We adhere to the classic view of continuous-time models as idealized versions of discrete-time models, which allows for the use of powerful tools of mathematical analysis within a simple framework.

In this respect, the use of continuous-time *repeated* games has been controversial. In such games, a "naive" definition of a strategy profile need not yield a well-defined outcome. This difficulty has been discussed at length, e.g., in [1,2,28,31]. All these authors provide various cumbersome restrictions on strategies, at the cost of losing the conceptual elegance of the continuous-time framework. Perry and Reny [22,23] adopt another approach. They assume that players have a waiting time: once changing the current action, a player needs to wait a pre-specified period until he can change his action again. We emphasize that such a difficulty in the notion of a strategy does not arise in timing games. The reason is that at any time t, there is only one relevant history of play—the history in which no one stopped. As a consequence, the difficulty of adapting one's own behavior to past behavior of the other players disappears.

A second issue is more directly related to timing games. It turns out that the set of equilibrium outcomes in continuous-time timing games need not coincide with the limit set of equilibrium outcomes of discrete time versions of the game, when the length of the time periods shrinks to zero. This issue was first raised in [11], through the *grab-the-dollar* game.[2] In this game each of two players (with the same discount rate) can grab a dollar that lies between them, at any time. The game terminates once at least one of the players grabs the dollar. If at that time only one player grabbed the dollar, he receives 1, and his opponent receives 0. If both grabbed it, both lose 1. In the discrete-time version of this game, the players are only allowed to act at exogenously given times (t_n), where the sequence (t_n) is increasing. The unique symmetric equilibrium has both players grab the dollar with probability $1/2$ at every time t_n (if the game still goes on at that stage), yielding a payoff of zero to both players. When the stage length decreases to zero, the symmetric equilibrium strategies do not converge to any strategy profile of the continuous-time version, since such a limit strategy would have to stop with probability $1/2$ at any time.

In our view, the problem which arises in the grab-the-dollar game is best seen as a lack of upper semi-continuity, as the time period decreases to zero. However, as was also pointed out in [10] in a different context, some kind of lower semi-continuity always holds: given any $\epsilon' > \epsilon > 0$, any ϵ-equilibrium profile for the continuous-time model is still, when discretized, an ϵ'-equilibrium in the discrete-time versions of the game, provided the time period is short enough.

To summarize, our analysis of a given continuous-time timing game yields a subgame perfect ε-equilibrium profile of the continuous-time game. Moreover, this profile yields an approximate equilibrium of *all* discrete-time versions of the game, provided time periods are short. However, not all ε'-equilibria of the discrete-time versions will be obtained this way.

The paper is organized as follows. In Section 1 we state our assumptions and our results. Section 2 contains the proof of the general existence result for two-player games while the discussion of specific classes of games is postponed to Section 3. The proofs of the assertions relative to Markov equilibrium are given in Section 4. All examples are collected in Section 5. Section 6 contains the proof of a result on the existence of exact equilibrium. Finally, Section 7 concludes with further discussion and few extensions.

1. The model and the main results

The set of non-negative reals $[0, \infty)$ is also denoted by $\mathbf{R}^+$, and for every $t \in \mathbf{R}^+$ we identify $[t, \infty] = [t, \infty) \cup \{\infty\}$.

[2] Our discussion follows closely the discussion in [11].

1.1. The model

A *game of timing* Γ is given by:

- A finite set of players I, and a discount rate $\delta_i \in \mathbf{R}^+$ for each player $i \in I$.
- For every non-empty subset $\emptyset \subset S \subseteq I$, a continuous and bounded function $u_S : [0, \infty) \to \mathbf{R}^I$.

To be consistent with the terminology of games in extensive form (see, e.g., [21, p. 103], we define a *plan of action* (or *plan* in short) of player i to be simply a time t_i to act, namely an element of $[0, \infty]$, where the alternative $t_i = \infty$ corresponds to never acting. Such a time does not define a strategy in the usual sense, since it does not prescribe what to do, if the game were to start after t_i.

Given a pure plan profile $(t_i)_{i \in I}$, we let $\theta := \min_{i \in I} t_i$ denote the terminal time, and $S_* := \{i \in I \mid t_i = \theta\}$ be the coalition of leaders. The payoff $g^i((t_j)_j)$ to player i is $e^{-\delta_i \theta} u^i_{S_*}(\theta)$ if $\theta < \infty$—i.e., if the game terminates in finite time—and 0 otherwise.

In most timing games of economic interest, the players incur costs, or receive profits prior to the end of the game, and the discounted sum of profits/costs up to t is bounded as a function of t. This case reduces to the case under study here by deducting/adding the total cost/profit up to time t from the discounted $u_S(t)$. Hence, our standing assumption that $g^i = 0$ if $\theta = \infty$ is a normalization convention, and entails no loss of generality.

1.2. Strategies and payoffs

A *mixed plan* for player i is a probability distribution σ^i over the set $[0, \infty]$. The *expected payoff* given a plan profile $\sigma = (\sigma^i)_{i \in I}$ is

$$\gamma^i_0(\sigma) = \mathbf{E}_{\otimes_{i \in I} \sigma^i}[g^i(t_1, \ldots, t_I)]. \tag{1}$$

The subscript reminds that payoffs are discounted back to time zero. We denote by $\gamma^i_t(\sigma) = e^{\delta_i t} \gamma^i_0(\sigma)$, the expected payoff discounted to time t.

In finite extensive form games, the notions of pure and mixed plans do not suffice when studying subgame-perfect equilibria. This is still the case here. Indeed, pure and mixed plans indicate when the player acts for the *first* time. However, they do not indicate how the player plays if the game starts at some time $t > 0$ which is beyond his acting time.

For every $t \geqslant 0$, the subgame that starts at time t is the game of timing Γ_t with player set I, where the payoff function when coalition S terminates is $u'_S(s) = u_S(t + s)$. Thus, payoffs are evaluated at time t.

Definition 1.1. A *strategy* of player 1 is a function $\widehat{\sigma}^i : t \mapsto \sigma^i_t$ that assigns to each $t \geqslant 0$ a mixed plan σ^i_t that satisfies

- Properness: σ^i_t assigns probability one to $[t, \infty]$.

- Consistency: For every $0 \leqslant t < s$ and every Borel set $A \subseteq [s, \infty]$, one has

$$\sigma_t^i(A) = (1 - \sigma_t^i([t, s))\sigma_s^i(A).$$

The properness condition asserts that σ_t^i is a mixed plan in the subgame that starts at time t: the probability that player i acts before time t is 0. The consistency condition asserts that as long as a plan does not act with probability 1, later strategies can be calculated by Bayes' rule.

The consistency requirement is closely related to the consistency property of conditional probability systems, see [20]. This is no coincidence, since a strategy (σ_t^i) can be interpreted as a conditional probability system over $[0, +\infty]$.

Given a strategy profile $\widehat{\sigma} = (\widehat{\sigma}^i)$, a player $i \in I$ and a time $t \in \mathbf{R}^+$, we denote by $\gamma_t^i(\widehat{\sigma}) := \gamma_t^i(\sigma_t)$ the payoff induced by $\widehat{\sigma}$ in the subgame starting at time t.

A Markov strategy is a strategy that depends only on payoff relevant past events, see [18]. In the context of timing games, this requirement is expressed as follows. A real number $T \in \mathbf{R}^+$ is a *period* of the game if $u_S(t + T) = u_S(t)$, for every $t \in \mathbf{R}^+$ and every $S \subseteq I$. A strategy profile σ is Markov if, for every $t \in \mathbf{R}^+$ and every $i \in I$, the mixed plan σ_{t+T}^i is obtained from σ_t^i by translation: for each Borel set $A \subseteq \mathbf{R}^+$, one has $\sigma_t^i(A) = \sigma_{t+T}^i(A + T)$.

1.3. Main results

Let $\epsilon > 0$ be given. A profile of mixed plans is a *Nash ϵ-equilibrium* if no player can profit more than ϵ by deviating to any other mixed strategy. Equivalently, no player can profit more than ϵ by deviating to a pure plan.

A profile of strategies $\widehat{\sigma} = (\sigma_t)_{t \geqslant 0}$ is a *subgame-perfect ϵ-equilibrium* if for every $t \geqslant 0$, the profile σ_t is a Nash ϵ-equilibrium in the subgame that starts at time t (when payoffs are discounted to time t).

We now state a general existence result for two-player games. The proof appears in Section 2.

Theorem 1.2. *Every two-player discounted game of timing in continuous time admits a Markov subgame-perfect ϵ-equilibrium, for every $\epsilon > 0$. If $\delta_i = 0$ for some i, the game admits a Nash ϵ-equilibrium, for each $\epsilon > 0$.*

The proof is essentially constructive. In many cases of interest, a *pure* subgame-perfect ϵ-equilibrium exists.

Section 3 deals with some classes of timing games of specific interest. We first analyze *games with cumulative payoffs*, defined by the property that for $i \in S$, the payoff $u_S^i(t)$ does not depend on which other player(s) happen to act at time t. Formally, $u_S^i(t) = u_{\{i\}}^i(t)$ for every player i and every subset S that contains i. This class includes games in which each player receives a stream of payoffs until he exits from the game (and the game proceeds with the remaining players). In particular, it includes models of shrinking markets, (see, e.g., [12,14]). It can also accommodate the case in which there is a collection $\mathcal{S}$ of winning coalitions, and the game terminates at the first time t in which the coalition of remaining players S_t is a winning coalition. One model of this sort is the model of multi-object auctions studied in [5].

Theorem 1.3. *Every game with cumulative payoffs has a subgame-perfect ϵ-equilibrium, for each $\epsilon > 0$. Moreover,*

- *there is such a profile in which symmetric players play the same strategy;*[3]
- *there is a Markov subgame-perfect ϵ-equilibrium, provided not all functions $u_S(\cdot)$, $S \subseteq I$, are constant.*

In many cases of economic interest, the players enjoy symmetric roles, in the sense that the payoff $u^i_S(t)$ to player i if S acts depends only on t, on the size of S, and on whether i belongs to S or not. Formally, a symmetric I-player game of timing is described by functions $\alpha_k : \mathbf{R}^+ \to \mathbf{R}$, $\beta_k : \mathbf{R}^+ \to \mathbf{R}$, $k \in \{1, \ldots, |I|\}$, with the interpretation that, for $|S| = k$, one has $u^i_S(t) = \alpha_k(t)$ if $i \in S$, and $u^i_S(t) = \beta_k(t)$ if $i \notin S$. For symmetric games, our existence result is surprisingly strong.

Theorem 1.4. *Every symmetric discounted game of timing admits a pure Markov subgame-perfect ϵ-equilibrium, for each $\epsilon > 0$.*

The *grab-the-dollar* game is an example of a symmetric game that does not have a symmetric ϵ-equilibrium, provided ϵ is sufficiently small.

Finally, in Section 1.5, we prove that under somewhat restrictive assumptions, the existence of an ϵ-equilibrium implies the existence of an equilibrium.

Theorem 1.5. *Let I be a finite set of players, let $u_S(\cdot)$ be a constant function for each $\emptyset \neq S \subseteq I$, and let $\delta_i = 0$ for each $i \in I$. If the game of timing $(I, (u_S)_S)$ has an ϵ-equilibrium for each $\epsilon > 0$, then it also has a zero equilibrium.*

In particular, combined with Theorem 1.2, Theorem 1.5 implies that every two-player, constant-payoff, undiscounted game of timing has a (mixed) Nash equilibrium. This equilibrium existence result is not standard. It is worth noting that it does *not* follow from the most general existence result due to Reny [25]. Indeed, Theorem 3.1 in Reny assumes that both strategy spaces are compact Hausdorff spaces, and that the game is so-called better-reply secure. In the context of timing games, one is tempted to endow the mixed strategy spaces with the topology

[3] Players i and j are symmetric if (i) $u^i_S = u^j_S$, for every S that either contains both i and j, or none of them, (ii) $u^i_{S\cup\{i\}} = u^j_{S\cup\{j\}}$ for every S that contains neither i nor j, and (iii) $\delta_i = \delta_j$.

of weak convergence.[4] Consider the constant-payoff timing game defined by $u_{\{1\}} = (3,1)$, $u_{\{2\}} = (0,0)$ and $u_{\{1,2\}} = (2,3/2)$, and any strategy profile σ where player 1 acts at time zero, but player 2 does not act at time zero: $\sigma^1(\{0\}) = 1$ and $\sigma^2(\{0\}) = 0$. Plainly, σ yields the payoff $(3,1)$ but is not an equilibrium. Since 3 is the highest payoff player 1 may possibly get in the game, player 1 cannot secure at σ a higher payoff, in the sense of Reny. On the other hand, any strategy $\tilde{\sigma}^2$ of player 2 that secures at σ a payoff strictly above one must act with some positive probability η at time zero. Let now σ^1_n be a sequence of strategies that weakly converges to σ^1 and with no atom at time zero. Plainly $\lim_{n\to\infty}\gamma^2(\sigma^1_n, \tilde{\sigma}^2) = (1-\eta) < 1$, hence Reny's condition does not hold.

2. Subgame-perfect equilibria in two-player games

This section is devoted to the proof of Theorem 1.2. The proof combines a backward induction argument with a compactness, or diagonal extraction, principle. We provide here a brief outline.

We start with few definitions, that will be in use throughout the section. We let a two-player game of timing $(u_S(\cdot))_{\emptyset\neq S\subseteq\{1,2\}}$ be given, together with the discount rates $\delta_1, \delta_2 > 0$ of the two players. For ease of presentation, we denote by $a(\cdot)$, $b(\cdot)$ and $c(\cdot)$ the three functions $u_{\{1\}}(\cdot)$, $u_{\{2\}}(\cdot)$ and $u_{\{1,2\}}(\cdot)$, respectively.

Note that for every continuous function $f : \mathbf{R}^+ \to \mathbf{R}^N$, and every $\eta, \delta > 0$, there is a strictly increasing sequence $(t_k)_k$, with limit ∞, such that for every k and every $t_k \leqslant s < t \leqslant t_{k+1}$, $||e^{-\delta(s-t)}f(s) - f(t)|| < \eta$.

Given $\epsilon > 0$, we let $\eta > 0$ be small enough. We apply the previous paragraph to the $\mathbf{R}^6$-valued function $f = (a,b,c)$, to η and to $\delta = \min\{\delta_1, \delta_2\}$, and obtain a sequence $(t_k)_k$ that strictly increases to ∞.

The proof is divided into two parts. Given $n \in \mathbf{N}$, we consider the version of the timing game that terminates at time t_n with a payoff of zero if no player acted before. In this game with finite horizon, we define inductively, for $0 \leqslant k < n$, a strategy profile $\widehat{\sigma}_k(n)$ over the time interval $[t_k, t_{k+1})$. We prove that the profile obtained by concatenating the profiles $\widehat{\sigma}_k(n)$ is a subgame-perfect ϵ-equilibrium in the game with finite horizon.

Next, we let n go to ∞. We observe that, for fixed k, the sequence $(\widehat{\sigma}_k(n))_n$ takes only finitely many values, so that by a diagonal extraction argument a limit $\widehat{\sigma}$ of $\widehat{\sigma}(n)$ exists. This limit is our candidate for a subgame-perfect ε-equilibrium.

2.1. Induction games

The induction step mentioned above takes as given a timing game played between times t_k and t_{k+1} and with a terminal payoff that may differ from zero. We deal in this section with such games.

[4] As in the two-player zero-sum timing game of Example 5.1 in Reny.

Given $0 \leqslant \tau < \theta < \infty$ and $v \in \mathbf{R}^2$, we define the *induction game* $G([\tau,\theta);v)$ to be the game that starts at time τ, and ends at time θ with a payoff of v if no player acted in between. In this game, each player is allowed to act at any time in $[\tau,\theta)$, and the payoff is v if no one ever acts. Since the interval $[\tau,\theta)$ is homeomorphic to $\mathbf{R}^+$, the induction game is formally equivalent to a game of timing, as introduced in Section 1, except that the terminal payoff may differ from zero, and that discounting is not exponential. The definitions of pure plans, mixed plans and strategies, as well as of a subgame-perfect ϵ-equilibrium, are analogous to those given for infinite horizon games. Hence, a *pure* plan in the induction game is an element in $[\tau,\theta) \cup \{\infty\}$, while a strategy of player i is a map $\widehat{\sigma}^i$ that assigns to each $t \in [\tau,\theta)$ a probability distribution over $[\tau,\theta) \cup \{\infty\}$, and satisfies the analogs of the Properness and Consistency requirements of Definition 1.1.

We shall later obtain strategy profiles in the infinite-horizon game by concatenating profiles of successive induction games. For clarity, we use the letter g for the payoff function in $G([\tau,\theta);v)$: given a strategy profile $\widehat{\sigma}$ in $G([\tau,\theta);v)$ and $t \in [\tau,\theta)$, $g_t(\widehat{\sigma})$ is the payoff induced by $\widehat{\sigma}$ in the subgame starting from t, and evaluated at time t.

2.1.1. Classification

We will say that the induction game $G([\tau,\theta);v)$ is of:

Type C if $c^1(\tau) \geqslant b^1(\tau)$ and $c^2(\tau) \geqslant a^2(\tau)$.

Type V if $e^{-\delta_1(\theta-\tau)}v^1 + \eta \geqslant a^1(\tau)$ and $e^{-\delta_2(\theta-\tau)}v^2 + \eta \geqslant b^2(\tau)$.

Type A1 if $a^1(\tau) \geqslant e^{-\delta_1(\theta-\tau)}v^1 + \eta$ and $a^2(\tau) \geqslant c^2(\tau)$.

Type B1 if $b^2(\tau) \geqslant e^{-\delta_2(\theta-\tau)}v^2 + \eta$ and $b^1(\tau) \geqslant c^1(\tau)$.

Type A2 if $a^1(\tau) \geqslant e^{-\delta_1(\theta-\tau)}v^1 + \eta$ and $a^2(\tau) \geqslant b^2(\tau)$.

Type B2 if $b^2(\tau) \geqslant e^{-\delta_2(\theta-\tau)}v^2 + \eta$ and $b^1(\tau) \geqslant a^1(\tau)$.

Type A3 if $a^1(\tau) \geqslant b^1(\tau)$ and $a^2(\tau) \geqslant c^2(\tau)$.

Type B3 if $b^2(\tau) \geqslant a^2(\tau)$ and $b^1(\tau) \geqslant c^1(\tau)$.

Each of these types may easily be interpreted. In a game of type C, the players will agree to act simultaneously. In a game of type V, the players will agree not to act on $[\tau,\theta)$.

Each induction game has at least one type, and possibly several. Indeed, assume that $G([\tau,\theta);v)$ has no type. If $a^1(\tau) \geqslant e^{-\delta_1(\theta-\tau)}v^1 + \eta$, one must have $a^2(\tau) < b^2(\tau)$ by A2, $b^1(\tau) < c^1(\tau)$ by B3, $a^2(\tau) > c^2(\tau)$ by C and $a^1(\tau) < e^{-\delta_1(\theta-\tau)}v^1 + \eta$ by A1—a contradiction. If $a^1(\tau) < e^{-\delta_1(\theta-\tau)}v^1 + \eta$ then one must have $b^2(\tau) \geqslant e^{-\delta_2(\theta-\tau)}v^2 + \eta$ by V, so that by the previous chain of implications, applied to player 2, one reaches a contradiction.

Plainly if (v_n) is a convergent sequence in $\mathbf{R}^2$, with limit v, and if the induction game $G([\tau,\theta);v_n)$ is of type T for every n, then $G([\tau,\theta);v)$ is also of type T.

2.1.2. Definition of the strategy profile

We next proceed to define a strategy profile $\widehat{\sigma}$ in the game $G([\tau,\theta);v)$. The payoff that will correspond to (σ_t^1,σ_t^2) is $c(t)$ (resp. v discounted to time t) if the type is C (resp. V), and is approximately $a(t)$ (resp. $b(t)$) if the type is A1, A2 or A3 (resp. B1, B2 or B3).

If the game is of

- type C, we let σ_t^i act with probability one at time t, for each $t\in[t,\theta)$, and $i=1,2$; hence $\gamma_t(\sigma_t)=c(t)$;
- type V, we let σ_t^i act with probability zero over the time interval $[t,\theta)$, for each t and $i=1,2$; hence $\gamma_t^i(\sigma_t)=e^{-\delta_i(\theta-t)}v^i$;
- type A1, we let σ_t^1 act with probability one at time t, and σ_t^2 assign probability zero to $[t,t_{k+1})$; hence $\gamma_t(\sigma)=a(t)$;
- type A2, we let σ_t^1 be the uniform distribution over $[t,\theta)$, and σ_t^2 act with probability zero over the time interval $[t,\theta)$; hence $\gamma_t(\sigma_t)\approx a(t)$ provided the maximal variation of a over the interval $[\tau,\theta)$ is small;
- type A3, we let σ_t^1 act with probability one at time t, and σ_t^2 be the uniform distribution over $[t,\theta)$; hence $\gamma_t(\sigma_t)=a(t)$.

Finally, types B1–B3 correspond, respectively, to types A1–A3, when exchanging the roles of the two players, and the definition of σ_t^i for those types is to be deduced from the definitions for their symmetric counterpart.

It is clear that $\widehat{\sigma}$ satisfies the properness requirement, and one can verify that it also satisfies the consistency requirement.

As explained earlier, the inductive proof will apply this construction to time intervals $[\tau,\theta)$ over which the maximal variation of $u_S(\cdot)$ is close to zero, for each $S\subseteq I$. We now prove that, under such assumptions, the profile $\widehat{\sigma}$ is a subgame-perfect ϵ-equilibrium of the game $G([\tau,\theta);v)$.

Proposition 2.1. *Let $\tau,\theta\in\mathbf{R}^+$ and $v\in\mathbf{R}^2$ be given. Assume that, for every $f\in\{a,b,c\}$, and for $\delta=\min\{\delta_1,\delta_2\}$, and $\tau\leqslant s<t<\theta$ one has $\|e^{-\delta(s-t)}f(s)-f(t)\|<\eta$ and moreover that $(1-e^{-\delta(\theta-\tau)})\|v\|<\eta$. Then, for each $t\in[\tau,\theta)$, the profile (σ_t^1,σ_t^2) is a 4η-equilibrium of the game $G([t,\theta);v)$. Moreover, if σ_t^2 assigns probability one to ∞, then player 1 does not profit by not acting, and the same holds when exchanging the roles of the two players.*

Proof. Let $t\in[\tau,\theta)$ be arbitrary. We prove that no pure plan of player 1 improves upon σ_t by more than 4η. The argument for player 2 is symmetric.

Assume that under σ_t^2 player 2 does not act in the interval $[t,\theta)$ (types V, A1, A2). Any deviation of player 1 yields at most

$$\max\{e^{-\delta_1(\theta-t)}v^1,\sup_{s\in[t,\theta]}e^{-\delta_1(s-t)}a^1(s)\}\leqslant\max\{e^{-\delta_1(\theta-t)}v^1,a^1(t)\}+\eta, \tag{2}$$

whereas the payoff to player 1 under (σ_t^1,σ_t^2) is $e^{-\delta_1(\theta-t)}v^1$ if the type is V, $a^1(t)$ if the type is A1, and at least $\inf_{s\in[t,\theta]}e^{-\delta_1(s-t)}a^1(s)\geqslant a^1(t)-\eta$ if the type is A2. In each case, by the definition of the types, this payoff is higher than the quantity in (2) minus 2η.

Observe that by not acting player 1 receives $e^{-\delta_1(\theta-t)}v^1$ which is at most what he receives in each of these cases. This establishes the second assertion of the Proposition.

Assume next that under σ_t^2 player 2 acts at time t (types C, B1, B3). Any pure deviation of player 1 yields either $b^1(t)$ or $c^1(t)$. However, the payoff to player 1 under (σ_t^1, σ_t^2) is $c^1(t)$ (resp. $b^1(t)$) if the type is C (resp. B1 or B3), which, by the definition of the types, is equal in both cases to $\max\{b^1(t), c^1(t)\}$.

Assume finally that σ_t^2 is the uniform distribution over $[t, \theta)$ (types A3, B2). Any deviation of player 1 yields at most $\max\{a^1(t), b^1(t)\} + \eta$. However, the payoff to player 1 under (σ_t^1, σ_t^2) is at least $a^1(t) - \eta$ (resp. $b^1(t) - \eta$) if the type is A3 (resp. B2), which, by the definition of the types, is equal in both cases to $\max\{a^1(t), b^1(t)\} - \eta$. In particular, player 1 cannot gain more than 2η by deviating. $\square$

2.2. *Proof of Theorem 1.2*

We here explicit the induction and the limit argument that were sketched in the introduction to this section.

Given $n \in \mathbf{N}$, we associate to each $k \in \{0, \dots, n\}$ a payoff $v_k(n) \in \mathbf{R}^2$ and a type $j_k(n)$, as follows:

- we set $v_n(n) := (0, 0)$;
- for $k < n$, we let $j_k(n)$ be a type of the induction game $G([t_k, t_{k+1}); v_{k+1}(n))$, and we let $v_k(n)$ be the payoff induced by the 4η-equilibrium that was defined in Section 2.1.1: $v_k(n) = g_{t_k}(\sigma_{t_k})$.

We now let n go to infinity. Since there are finitely many types, and since payoffs are bounded, a diagonal extraction argument implies that there is an increasing sequence of indices $(n_m)_{m\in\mathbf{N}}$ such that the sequences $(v_k(n_m))_{m\in\mathbf{N}}$ and $(j_k(n_m))_{m\in\mathbf{N}}$ converge for every $k \geqslant 0$. Denote for every $k \geqslant 0$ $v_k = \lim_{m\to\infty} v_k(n_m)$ and $j_k = \lim_{m\to\infty} j_k(n_m)$. By the remark at the end of Section 2.1.2, j_k is a type of $G([t_k, t_{k+1}); v_{k+1})$.

We next proceed to the definition of a strategy profile $(\widehat{\sigma}^1, \widehat{\sigma}^2)$ in the timing game (with infinite horizon). Given $k \in \mathbf{N}$, we denote by $(\widehat{\sigma}^{1,k}, \widehat{\sigma}^{2,k})$ the strategy profile in the game $G([t_k, t_{k+1}); v_{k+1})$ corresponding to type j_k, as defined in Section 2.1.2. Note that, for $i \in I$ and $t \in [t_k, t_{k+1})$, $\sigma_t^{i,k}$ is a probability distribution over $[0, \infty]$ which gives probability 1 to $[t, t_{k+1}) \cup \{\infty\}$.

By Proposition 2.1, for each $t \in [t_k, t_{k+1})$, the profile $(\sigma_t^{1,k}, \sigma_t^{2,k})$ is a 4η-equilibrium of the game $G([t, t_{k+1}); v_{k+1})$.

Intuitively, we shall define $\widehat{\sigma}_t^i$, $t \in \mathbf{R}^+$, as the concatenation of the different strategies $(\widehat{\sigma}^{i,k})_{k\in\mathbf{N}}$. Formally, this is achieved via the following construction.

Given a mixed plan σ^i in an induction game $G([t, t'); v)$ and a mixed plan σ'^i in an induction game $G([t', t''); v')$, we define their concatenation $\sigma^i \circ \sigma'^i$ to be the distribution in $G([t,t'');v')$ that assigns probability $\sigma^i(A)$ to every Borel set $A \subseteq [t,t')$, and probability $(1-\sigma^i([t,t'))\sigma'^i(A)$ to every Borel set $A \subseteq [t',t'') \cup \{\infty\}$. For every k and every $t \in [t_k, t_{k+1})$ define

$$\sigma_t^i = \sigma_t^{i,k} \circ \sigma_{t_{k+1}}^{i,k+1} \circ \sigma_{t_{k+2}}^{i,k+2} \circ \cdots .$$

One can verify that $\widehat{\sigma}^i = (\sigma_t^i)_{t \in \mathbf{R}^+}$ satisfies both the Properness and the Consistency requirement in Definition 1.1. We omit this verification.

Proposition 2.2. *The strategy profile $\widehat{\sigma}$ is a subgame-perfect ϵ-equilibrium of the timing game.*

Proof. We first claim that $\gamma_{t_k}(\widehat{\sigma}) = v_k$ for each $k \in \mathbf{N}$. Indeed, since $\widehat{\sigma}$ is defined as the concatenation of the profiles $\widehat{\sigma}^k$, the equation that links $\gamma_{t_k}(\widehat{\sigma})$ to $\gamma_{t_{k+1}}(\widehat{\sigma})$ is the same as the relation between v_k and v_{k+1}: if at least one player acts with probability one on the interval $[t_k, t_{k+1})$, both v_k and $\gamma_{t_k}(\widehat{\sigma})$ coincide with the corresponding payoff. On the other hand, if both players act with probability zero on $[t_k, t_{k+1})$, then $\gamma_{t_k}^i(\widehat{\sigma}) = e^{-\delta_i(t_{k+1}-t_k)}\gamma_{t_{k+1}}^i(\widehat{\sigma})$ and $v_k^i = e^{-\delta_i(t_{k+1}-t_k)}v_{k+1}^i$. Therefore, for a given k, either (i) there is $k_* > k$ such that at least one player acts with probability one on the interval $[t_{k_*}, t_{k_*+1})$, in which case reasoning backwards from k_* yields $\gamma_{t_k}(\widehat{\sigma}) = v_k$, or (ii) no such k_* exists, in which case the equality $v_k^i = e^{-\delta_i(t_l - t_k)}v_l^i$ holds for each $l > k$. Since payoffs are bounded, by letting l go to infinity we obtain $v_k = 0$ for each k, so that as above $\gamma_{t_k}(\widehat{\sigma}) = 0$.

Let $k \in \mathbf{N}$ and $t \in [t_k, t_{k+1})$ be given. We shall prove that, for each pure plan $\sigma_t'^1$ in the timing game starting at t, one has

$$\gamma_t^1(\sigma_t'^1, \sigma_t^2) \leqslant \gamma_t^1(\sigma_t^1, \sigma_t^2) + \epsilon. \tag{3}$$

Since the roles of the two players are symmetric, this will imply that (σ_t^1, σ_t^2) is an ϵ-equilibrium of the game starting at time t. Since t is arbitrary, the Proposition will follow.

Since it is a pure plan, $\sigma_t'^1$ assigns probability one to some element $t_* \in [t, \infty) \cup \{\infty\}$. We first deal with the case $t_* < \infty$.

Let $k_* \in \mathbf{N}$ be the unique integer such that $t_* \in [t_{k_*}, t_{k_*+1})$. Let $k_{**} \geqslant k$ be the first integer such that the type of the game $G([t_{k_{**}}, t_{k_{**}+1}); v_{k_{**}+1})$ is either C, B1, B2, A3 or B3 (with $k_{**} = \infty$ if no such integer exists). By the definition of the plan of player 2, the game terminates before time $t_{k_{**}+1}$ with probability one, whatever player 1 plays. Set $\widehat{k} = \min\{k_*, k_{**}\}$.

We prove that for every $k < k' \leqslant \widehat{k}$, the expected payoff of player 1 if player 2 follows $\sigma_{t_{k'}}^2$ and player 1 acts at time t_*, discounted to $t_{k'}$, is at most $v_{k'}^1 + 4\eta$.

For $k' = \widehat{k}$ this follows since (σ_t^1, σ_t^2) is a 4η-equilibrium of the induction game $G([t_{\hat{k}}, t_{\hat{k}+1}); v_{\hat{k}+1})$.[5] Assume we proved the claim for $k'+1$. Since player 2 does not

[5] Strictly speaking, σ_t^i need not be an admissible plan in $G([t_{\hat{k}}, t_{\hat{k}+1}); v_{\hat{k}+1})$, but it induces one when collapsing $[t_{\hat{k}+1}, \infty]$ to ∞.

act before time $t_{k'+1}$, the type $j_{k'}$ of the game $G([t_{k'}, t_{k'+1}); v_{k'+1})$ must be V, A1 or A2. By the induction hypothesis, the expected payoff of player 1 if player 2 follows $\sigma^2_{t_{k'}}$ and player 1 acts at time t_*, discounted to $t_{k'}$, is at most $e^{-\delta_1(t_{k'+1}-t_{k'})}(v^1_{k'+1} + 4\eta) \leqslant e^{-\delta_1(t_{k'+1}-t_{k'})}v^1_{k'+1} + 4\eta$. By the second assertion of Proposition 2.1 this last quantity is at most $v^1_{k'} + 4\eta$, as desired. The same argument, applied to the induction game $G([t, t_{k+1}); v_{k+1})$, delivers now (3).

For every $t \in [0, \infty]$ denote by $\delta(t)$ the pure plan that acts at time t with probability 1.

If $t_* = \infty$, then, since $\delta_1 > 0$ and by the first part,

$$\gamma^1_t(\delta(\infty), \sigma^2_t) = \lim_{\tilde{t} \to \infty} \gamma^1_t(\delta(\tilde{t}), \sigma^2_t) \leqslant \gamma^1_t(\sigma_t) + 4\eta. \qquad \square \tag{4}$$

Comment. We now argue that if $\delta_1 = 0$ (or $\delta_2 = 0$), that is, if at least one of the players does not discount, then a Nash ε-equilibrium exists.

For every n and k, let $(\widehat{\sigma}^{1,k}(n), \widehat{\sigma}^{2,k}(n))$ be the strategies defined in Section 2.1 for type $j_k(n)$ in the game $G([t_k, t_{k+1}); v_k(n))$. Denote $\sigma^i_0(n) = \sigma^{i,1}_{t_1}(n) \circ \sigma^{i,2}_{t_2}(n) \circ \cdots \circ \sigma^{i,n-1}_{t_{n-1}}(n)$. If under $(\sigma^1_0(n), \sigma^2_0(n))$ both players act with probability 1 before time t_n, the arguments we presented in the proof of Proposition 2.2 imply that $(\sigma^1_0(n), \sigma^2_0(n))$ is an ε-equilibrium.

Assume, then, that under $\sigma^2_0(n)$ player 2 never acts, for every n. Then $j_k(n)$ is V, A1 or A2 for every k and every n. The construction in Section 2.1.2 implies that $v^1_k(n) \geqslant 0$ for every k and every n. In particular, the plan $\delta(\infty)$ that never acts cannot be a profitable deviation of player 1. Let n be sufficiently large such that for some $t < t_n$ one has $a^1(t) \geqslant \sup_{s \in [0,\infty)} a^1(s) - \eta$ and for some $t' < t_n$ one has $b^2(t') \geqslant \sup_{s \in [0,\infty)} b^2(s) - \eta$. In words, the best payoff by acting alone occurs before time t_n. One can verify that $(\sigma^1_0(n), \sigma^2_0(n))$ is a 5η-equilibrium.

Corollary 2.3. *Assume that, for every t one has either* (i) $b^1(t) \geqslant c^1(t)$ *and* $a^2(t) \geqslant c^2(t)$, *or* (ii) $b^1(t) \leqslant c^1(t)$ *and* $a^2(t) \leqslant c^2(t)$. *Then for every* $\varepsilon > 0$,

- *if* $\min\{\delta_1, \delta_2\} > 0$, *there exists a pure subgame-perfect* ϵ*-equilibrium.*
- *if* $\min\{\delta_1, \delta_2\} = 0$, *there exists a pure* ϵ*-equilibrium.*

Observe that in wars of attrition, condition (i) holds for every t.

Proof. It suffices to show that all the induction games $G([t_k, t_{k+1}), v_{k+1}(n))$ that appear in the proof are of types C, V, A1 or B1. This is a matter of straightforward verification. $\square$

3. Proofs of Theorems 1.3 and 1.4

Many proofs in this section are minor variations upon the proof of Theorem 1.2. Hence few details will be omitted. Again, the proofs of the assertions relative to Markov equilibrium are postponed to Section 4.

3.1. Games with cumulative payoff

We here prove Theorem 1.3. Let Γ be a game with cumulative payoffs. Fix a strictly increasing sequence (s_n) with $s_0 = 0$ and $\lim_{n\to\infty} s_n = \infty$ such that $\sup_n \sup_{s_n \leqslant s < t \leqslant s_{n+1}} |e^{-\delta(s-s_n)} u_S^i(s) - u_S^i(t)| < \epsilon$ for every non-empty subset $S \subseteq I$ and every player i. Define an auxiliary game Γ^* in which players can act *only* at times $\{s_n, n \geqslant 0\}$ and *must* continue in all other times. The auxiliary game Γ^* is equivalent to a discounted[6] game Γ^{**} in discrete time with countably many states s_n. The stochastic game Γ^{**} has quite a specific structure: at state s_n, each player can either act or not. If at least one player acts, the game reaches an absorbing state. If no one acts, the game moves to state s_{n+1}.

Every strategy profile τ_{**} in the game Γ^{**} naturally induces a strategy profile in the game Γ^*, and therefore it induces a strategy profile $\widehat{\tau}$ in the game Γ. Observe that for every n, the expected payoff under τ_{**} starting from state s_n is equal to the expected payoff induced by $\widehat{\tau}$ in Γ, starting from time s_n.

By Fink, [8] the discounted stochastic game Γ^{**} has a subgame-perfect 0-equilibrium $\tau_{**} = (\tau_{**}^i)_{i \in I}$. Moreover, there is such a subgame-perfect 0-equilibrium in which symmetric players play the same strategy.

Denote by $\widehat{\sigma}$ the profile of strategies in Γ induced by τ_{**}. Then $\widehat{\sigma}^i = \widehat{\sigma}^j$ for every pair of symmetric players $i \neq j$. Moreover, under $\widehat{\sigma}$ players act only at times $(s_n)_{n\geqslant 0}$, that is, the probability distribution σ_t^i gives weight one to the set $\{s_n, n \geqslant 0\}$, for each $t \in \mathbf{R}^+$.

We will prove that $\widehat{\sigma}$ is a subgame-perfect ϵ-equilibrium. Let $t \in \mathbf{R}^+$ be given, and let τ^i be a pure plan of player i in the subgame starting at time t, which acts at time $t_i \in [t, \infty]$.

We denote by $\tilde{\tau}^i$ the auxiliary pure plan that acts at time s_k, where $k \in \mathbf{N} \cup \{\infty\}$ is the minimal integer such $s_k \geqslant t_i$. By construction, under both (σ_t^{-i}, τ^i) and $(\sigma_t^{-i}, \tilde{\tau}^i)$ no player in $S \setminus \{i\}$ acts in the time interval (t_i, s_k). Therefore,

$$|\gamma_t^i(\sigma_t^{-i}, \tilde{\tau}^i) - \gamma_t^i(\sigma_t^{-i}, \tau^i)| < |e^{-\delta_i(s_k - t_i)} u_{\{i\}}^i(t_i) - u_{\{i\}}^i(s_k)| \leqslant \epsilon. \tag{5}$$

The pure plan $\tilde{\tau}^i$ is a valid plan in Γ^*, and therefore naturally induces a pure strategy $\tilde{\tau}_{**}^i$ in Γ^{**}. Since τ_{**} is a subgame-perfect 0-equilibrium, the payoff induced by $(\tilde{\tau}_{**}^i, \tau_{**}^{-i})$ in the stochastic game Γ^{**}, starting from state s_k, does not improve upon the payoff induced by τ_{**} in that game. Since these payoffs coincide with $\gamma_t^i(\sigma_t^{-i}, \tilde{\tau}^i)$

[6] with a state-dependent discount factor.

and $\gamma_t^i(\sigma_t)$ respectively, and by (5), one gets

$$\gamma_t^i(\sigma_t^{-i}, \tau^i) \leqslant \gamma_t^i(\sigma_t) + \epsilon,$$

as desired.

3.2. Symmetric games

We here prove Theorem 1.4. Let an I-player symmetric timing game be given. We set

$$T_I = \{t \in [0, \infty) \mid \alpha_I(t) \geqslant \beta_{I-1}(t)\}$$

and

$$T_k = \{t \in [0, \infty) \mid \alpha_k(t) \geqslant \beta_{k-1}(t) \quad \text{and} \quad \alpha_{k+1}(t) \leqslant \beta_k(t)\}, \text{ for } k = 2, 3, \ldots, I-1.$$

If $t \in T_I$ then the plan profile in which all players act at time t is a 0-equilibrium in Γ_t. Indeed, under this profile the payoff for all players is $\alpha_I(t)$, while any deviator who will not act at time t will receive $\beta_{I-1}(t) \leqslant \alpha_I(t)$.

Similarly, if $t \in T_k$, for $k = 2, \ldots, I-1$, any plan profile in which exactly k players act at time t is a 0-equilibrium in the game starting from time t. Indeed, any one of the k players who acts at time t receives $\alpha_k(t)$, while if such a player deviates and does not act at time t he will receive $\beta_{k-1}(t) \leqslant \alpha_k(t)$. Any one of the $I-k$ players who does not act at time t receives $\beta_k(t)$, while if such a player deviates and acts at time t he will receive $\alpha_{k+1}(t) \leqslant \beta_k(t)$.

For $k = 2, 3, \ldots, I$, we let T_k^* be the closure of the interior of T_k. Then each T_k^* is the union of at most countably many disjoint closed intervals: $T_k^* = \bigcup_{n=1}^{\infty} [c_n^k, d_n^k]$. Set $\widehat{T}_k = \bigcup_{n=1}^{\infty} [c_n^k, d_n^k)$.

We set $T_0 = [0, \infty) \setminus \bigcup_{k=2}^{I} \widehat{T}_k$. Observe that $T_0 = \bigcup_{n=1}^{\infty} [c_n^0, d_n^0)$ is a union of disjoint half-closed half-open intervals.

Given $t \in \mathbf{R}^+$, one has $t \in \bigcup_{k \geqslant 2} T_k$ as soon as $\alpha_2(t) \geqslant \beta_1(t)$. Therefore, $\alpha_2(t) \leqslant \beta_1(t)$ for every $t \in T_0$.

We already defined a pure 0-equilibrium for initial times $t \in \bigcup_k \widehat{T}_k$. To complete the proof, it is now sufficient to prove that a subgame-perfect ϵ-equilibrium exists in each game $G([c_n^0, d_n^0); v)$, where v is the equilibrium payoff we defined starting from time d_n^0. If $d_n^0 = \infty$, we set this terminal payoff to zero. To prove this claim, we shall mimic the proof of Theorem 1.2. We shall only sketch the main steps of the proof. We let the game $G([c_n^0, d_n^0); v)$ and $\epsilon > 0$ be given. Choose $\eta > 0$ to be very small. Consider an increasing sequence $(t_k)_k$ that converges to d_n^0 and such that $\sup_{s,t \in [t_k, t_{k+1}]} |e^{-\delta(s-t_k)} \alpha_1(s) - \alpha_1(t)| < \eta$. If $d_n^0 < \infty$, we define the sequence so that it contains only finitely many terms $(t_k)_{k \leqslant K}$, with $t_K = d_n^0$. In that case, the profile is constructed by backward induction, starting with the game $G([t_{K-1}, d_n^0); v)$. If

$d_n^0 = \infty$, the sequence (t_k) contains infinitely many terms, and the induction proceeds as in the proof of Theorem 1.2, as explained below.

Fix $k \in \mathbf{N}$, and look at the game $G([t_k, t_{k+1}); v_k(n))$ that appears in the induction step. We use the symmetry of the game to simplify the classification into types. Specifically, we say that $G([t_k, t_{k+1}); v_k(n))$ is of

Type V if $e^{-\delta(t_{k+1}-t_k)}\min_{i \in I} v_k^i(n) + \eta \geqslant \alpha_1(t_k)$.

Type $1i$ if $e^{-\delta(t_{k+1}-t_k)} v_k^i(n) + \eta < \alpha_1(t_k)$.

Following the proof of Theorem 1.2, we define a *pure* strategy profile in the game $G([t_k, t_{k+1}); v_k(n))$, depending on the type of that game. If it is of type V, we let σ_t^i act with probability zero on the time interval $[t, t_{k+1})$, for each $t \in [t_k, t_{k+1})$. If it is of type $1i$ for some i, we let σ_t^i act with probability one at t, and σ_j^i act with probability zero on the time interval $[t, t_{k+1})$, for each $j \neq i$ and $t \in [t_k, t_{k+1})$. The rest of the proof follows the proof of Theorem 1.2.

4. Markov equilibrium

We here collect all proofs that relate to Markov strategies. It will be convenient to describe the set of Markov strategies, when payoffs are constant. Let $\widehat{\sigma}^i$ be a Markov strategy of player i. If $\sigma_0^i(0) = 1$ then $\sigma_t^i(0) = 1$ for every $t \in \mathbf{R}^+$: under $\widehat{\sigma}^i$ the player acts at every time t.

If $\sigma_0^i(0) < 1$ then $\sigma_0^i(\eta) < 1$ for some $\eta > 0$ sufficiently small. By the Markov requirement, this implies that $\sigma_0^i(s) < 1$ for every $s \in \mathbf{R}^+$; indeed, by induction over k, $\sigma_0^i((k+1)\eta) = \sigma_0^i(k\eta) + (1 - \sigma_0^i(k\eta))\sigma_0^i(\eta) < 1$. Moreover, the Markov requirement implies that $(1 - \sigma_0^i(t))(1 - \sigma_0^i(s)) = 1 - \sigma_0^i(t+s)$, so that by the characterization of the exponential distribution (see, e.g., [3, p. 189]) σ_0 is an exponential distribution over $\mathbf{R}^+$, and for $t > 0$ σ_t is obtained by translation. To summarize, if a strategy $\widehat{\sigma}$ is Markov, then σ_t is obtained from σ_0 by translation. Moreover, σ_0 is either a unit mass located at 0 or ∞, or is an exponential distribution over $[0, \infty)$. Conversely, any such strategy has the Markov property.

Proposition 4.1. *Every two player game has a Markov subgame-perfect ϵ-equilibrium, for each $\epsilon > 0$.*

Proof. We shall use the notations of Section 2. We first assume that $a(\cdot), b(\cdot)$ and $c(\cdot)$ are constant, and we adapt the proof of Theorem 1.2. Since payoffs are constant, it is sufficient for our proof to consider only one induction game $G([0, \infty); \vec{0})$. In most cases (i.e., C, V, A1, B1, A2 and B3 for player 2, A3 and B2 for player 1) the strategies we defined are either never to act, or always to act, which are Markov. In the other four cases replace the current definition of σ_t^i by an exponential distribution over $[t, \infty)$ with sufficiently high parameter α. Given $\epsilon > 0$, if α is sufficiently high, then under the new definition the game terminates before time $t + \epsilon$ with probability at least $1 - \epsilon$; since the payoff functions are constant this implies that no player can profit in discounted terms more than 3ϵ by deviating, provided ϵ is sufficiently small.

Next, we assume that the functions $a(\cdot), b(\cdot)$ and $c(\cdot)$ have a common period $T<\infty$. We shall discuss two cases. Up to symmetries, these cases exhaust all possible cases.

Case 1: $a^1(t) \leqslant b^1(t)$ *and* $a^2(t) \geqslant b^2(t)$ *for each* $t \in \mathbf{R}^+$. In a sense, each player would rather see his opponent stop. We adapt the proof of Theorem 1.3, see Section 3.1. We shall only sketch the proof, without providing all the details. Given $\varepsilon>0$, we let $\eta>0$ be small enough, and let $0=t_0<t_1<\cdots<t_n=T$ be a finite subdivision of $[0, T]$, such that a, b and c do not vary by more than η on each subinterval $[t_k, t_{k+1}]$, $k=0,1,\ldots,n-1$.

Consider the stochastic game Γ^{**} with finitely many states labelled $t_0, \ldots, t_{n-1}$ where (i) the game moves cyclically from one state to the next one in the sequence (and from t_{n-1} to t_0) as long as no player ever acts, (ii) player 1 (resp. player 2) can only act in states with *odd* index (resp. with *even* index), and (iii) the payoff by acting at state t_k is $a(t_k)$ or $b(t_k)$ depending on k. The game Γ^{**} has a subgame-perfect equilibrium $\widehat{\sigma}$ in stationary strategies—strategies that depend only on the current state. When reverting to the interpretation of t_k as a time rather than a state, this profile corresponds to a periodic profile—still denoted $\widehat{\sigma}$—in the timing game. We derive a modified, periodic strategy profile $\widehat{\tau}$ as follows. Loosely, if player i stops with probability p at time t_k under $\widehat{\sigma}$, we will have him act under $\widehat{\tau}$ with probability p over the whole time-interval $[t_k, t_{k+1})$. Specifically, for $k<n$, the mixed plan $\tau^i_{t_k}$ has no atoms, assigns to the interval $[t_k, t_{k+1})$ the probability $\sigma^i_{t_k}(\{t_k\})$ with which $\sigma^i_{t_k}$ acts at time t_k, and can be calculated using Bayes' rule from $\tau^i_{t_{k+1}}$ on the interval $[t_{k+1}, \infty]$. For $t \neq t_k$, τ_t is defined via Bayes rule. Note that, for each $t \in \mathbf{R}^+$, the payoffs $\gamma_t(\widehat{\sigma})$ and $\gamma_t(\widehat{\tau})$ differ by at most η.

We claim that $\widehat{\tau}$ is a subgame-perfect ε-equilibrium, provided η is small enough. Plainly, it is enough to prove that player 1 cannot deviate profitably in the game that starts at time 0. This claim is supported by the following arguments.

Let $\tilde{\tau}^i_0$ be a pure plan of player 1 in the timing game. If it never acts, it is payoff equivalent—up to η—to the plan in Γ^{**} that never acts.[7] If it acts at time $t \in [t_k, t_{k+1})$ for some odd k, it is payoff-equivalent to the plan in Γ^{**} that acts at state t_k. Finally, if it acts at time $t \in [t_k, t_{k+1})$ for some even k, it yields a *lower* payoff than the plan that acts at time t_{k+1}, by the assumption on payoffs.

Case 2: $a^2(t_*)<b^2(t_*)$ for some $t_* \in \mathbf{R}^+$. We start with a simple observation. Assume that, for some $t \in \mathbf{R}^+$ and $\eta>0$, there is a profile $\widehat{\sigma}$ such that (i) $\widehat{\sigma}$ is a subgame-perfect ε-equilibrium in $G([t, t+\eta); v)$, irrespective of v and (ii) for each $s \in [t, t+\eta)$, under σ_s, at least one player will act before $t+\eta$. Then there is a Markov ε-equilibrium.

Indeed, by translation we can assume that $t \geqslant T$. By the backward-induction argument presented in Section 2.2 we construct a pure ε-equilibrium in the period

[7] To be precise: faced with τ_0^{-i} in the timing game, it yields approximately the same payoff as the plan *never act* in Γ^{**}, faced with $\widehat{\sigma}^{-i}$.

$[t+\eta-T, t+\eta]$. By (2), the strategy profile in the original game that is defined by repeating periodically this ε-equilibrium is a subgame-perfect ε-equilibrium in the original game.

Given this fact, we shall mimic the proof of Theorem 1.2, see Section 2.2, where we choose the sequence (t_k) so that $t_* = t_{k_*}$ for some $k_* \in \mathbf{N}$. If, for some $n \in \mathbf{N}$, the induction game $G([t_{k_*}, t_{k_*+1}); v_{k_*}(n))$ is either of type A3, B3 or C, we may apply the above observation with $[t, t+\eta) = [t_{k_*}, t_{k_*+1})$ and the result follows. Otherwise, it must be that $a^1(t_*) < b^1(t_*)$. Indeed, since $a^2(t_*) < b^2(t_*)$, one first has $b^1(t_*) < c^1(t_*)$ by B3, next $a^2(t_*) > c^2(t_*)$ by C, and finally $a^1(t_*) < b^1(t_*)$ by A3.

To conclude, we let $[t, t+\eta) = [t_{k_*}, t_{k_*+1})$, and define a profile $\widehat{\sigma}$ in $G([t, t+\eta); v)$ by having both players acting time be distributed according to an exponential distribution[8] over $[t, t+\eta)$. The parameter of player 2's distribution is chosen to be much larger than the parameter of player 1's distribution. We then apply the basic observation. □

Next, we show that in symmetric games and in games with non-constant cumulative payoff a Markov ϵ-equilibrium always exists, irrespective of the number of players.

Proposition 4.2. *Every multi-player symmetric game of timing has a pure Markov subgame-perfect ε-equilibrium.*

Proof. We modify the proof given in Section 3.2. If payoffs are constant, the proof is similar to the proof of Proposition 4.1.

Assume now that the payoffs are periodic with period $T > 0$. We shall use the observation made in Case 2 of the previous proof. Observe that if $t \in T_k$ for some $k = 2, \ldots, K$, and if $\eta > 0$ is small enough, then the profile that requires k players to act and $I - k$ players to continue satisfies the two requirements of that observation. Therefore, we can assume w.l.o.g. that $T_0 = [0, \infty)$.

If $\sup \alpha_1(\cdot) \leqslant 0$, there is a subgame-perfect equilibrium in which no player ever acts. Thus, we may assume that $\sup \alpha_1 > 0$. We divide the proof in three cases. Since α_1 and β_1 are continuous, these exhaust all possible cases.

Case 1: $\alpha_1(t) = \beta_1(t)$ *for some* t. We let η be small enough, and let (t_n) be an increasing sequence with limit $t+\eta$ and such that $t_0 = t$. We define σ as follows: player 1 (resp. player 2) acts at each time $s \in [t_n, t_{n+1})$ for even n (resp. for odd n). Players $3, 4, \ldots, I$ never act. We then use the first observation.

Case 2: $\alpha_1(t) > \beta_1(t)$ *for each* $t \in \mathbf{R}^+$. We divide the time interval $[0, T]$ into a large, finite, even number of intervals, and define a periodic profile $\widehat{\sigma}$ as follows: player 1 (resp. player 2) acts at each time $s \in [t_n, t_{n+1})$ for even n (resp. for odd n). Players $3, 4, \ldots, I$ never act. It is straightforward to check that $\widehat{\sigma}$ is a subgame-perfect ε-equilibrium, provided the partition of $[0, T]$ is fine enough.

[8] To be precise, it is the image of an exponential distribution over $\mathbf{R}^+$ under an increasing homeomorphism that maps $\mathbf{R}^+$ to $[t, t+\eta)$.

Case 3: $\alpha_1(t) < \beta_1(t)$ *for each* $t \in \mathbf{R}^+$. Choose $t_* \geqslant T$ such that $\alpha_1(t_*) = \sup_{t \in \mathbf{R}^+} \alpha_1(t)$, and let $\eta > 0$ be small enough. We divide the period $[t_* - T + \varepsilon, t_* + \varepsilon)$ into finitely many small intervals $[t_k, t_{k+1})$, $k = 0, \ldots, k_*$ and apply the backward construction that appears in the proof of Theorem 1.4. We initialize the induction with player 1 acting at each $s \in [t_{k_*}, t_{k_*+1})$, while players $2, \ldots, I$ do not act on $[t_{k_*}, t_{k_*+1})$. Hence $v^1_{k_*} = \alpha_1(t_*)$, while $v^i_{k_*} = \beta_1(t_*)$ for each $i = 2, \ldots, I$. One can check inductively that $0 < v^1_k < v^i_k$ for each $k = 1, \ldots, k_*$ and $i = 2, \ldots, I$—so that each induction game is either of type 1–1 or C, while the last one, $G([t_0, t_1); v_1)$ is of type 1–1. Therefore, this construction generates a periodic profile. $\square$

Proposition 4.3. *In every multi-player game with non-constant cumulative payoffs a Markov subgame-perfect* ε*-equilibrium exists. Moreover, there is a Markov equilibrium where symmetric players play the same strategy.*

Proof. The proof is essentially the same as the proof of Theorem 1.2. All one should note is that since payoffs are periodic, one can construct the stochastic game Γ^{**} in discrete time to have finitely many states, that correspond to one period of the game in continuous time. $\square$

5. Examples

In the present section we study several examples, which show that the results we present in the paper are sharp. We first exhibit a two-player zero-sum game with no Nash (exact) equilibrium. Next, we analyze a three-player zero-sum game with no ε-equilibrium, provided ε is sufficiently small. As mentioned in the Introduction, the grab-the-dollar game is a symmetric game with no symmetric ε-equilibrium, but it does admit a pure (non-symmetric) equilibrium. Our third example is one of a two-player symmetric game with no pure equilibrium. We conclude with a three-player game with cumulative payoffs, that has no Markov ε-equilibrium.

5.1. A two-player zero-sum game with no equilibrium

Consider the two-player zero-sum game defined by $u^1_S(t) = 1$ if $|S| = 1$ and $u^1_{\{1,2\}}(t) = 0$, with $\delta_1 > 0$.

We first argue that player 1 can guarantee a payoff $1 - \epsilon$, for every $\epsilon > 0$. Indeed, consider the mixed plan σ^1 that acts at a random time in the interval $[0, \eta]$, where $\eta > 0$ satisfies $e^{-\delta_1 \eta} \geqslant 1 - \epsilon$. Formally, the corresponding c.d.f. F^1 is defined by $F^1_t = \min\{t/\eta, 1\}$. Since player 1 acts at a random time, the probability that both players act simultaneously is 0, whatever be the plan used by player 2. Since the game terminates by time η, player 1's payoff is 1 with probability 1, and taking the discount rate into account, his expected payoff is at least $e^{-\delta_1 \eta} \geqslant 1 - \epsilon$. Since the highest payoff in the game is 1, this means that the value of the game exists, and is equal to 1.

Don't Act

	Don't Act	Act
Don't Act		$-1, 1, 0$
Act	$1, 0, -1$	$0, -1, 1$

Act

	Don't Act	Act
Don't Act	$0, -1, 1$	$1, 0, -1$
Act	$-1, 1, 0$	$0, 0, 0$

Fig. 1.

We now claim that player 1 has no optimal strategy. Indeed, the discounted payoff of player 1 equals 1 only if, with probability one, the game terminates at time 0, and only one player acts at that time. This can happen only if one player acts with probability one at time 0, while the other does not act. However, if player 1 acts with probability 1 at time 0, it is optimal for player 2 to act at time 0 as well, whereas if player 1 does not act at time 0, it is optimal for player 2 not to act at time 0 as well.

5.2. *A three-player zero-sum game with no ϵ-equilibrium*

We here analyze the three-player zero-sum game of timing with constant payoffs that is defined by[9] $u^i_{\{i\}}(t) = 1$, $u^{i+1}_{\{i\}}(t) = 0$, $u^{i+2}_{\{i\}}(t) = -1$, $u^i_{\{i,i+1\}}(t) = 0$, $u^{i+1}_{\{i,i+1\}}(t) = -1$, $u^{i+2}_{\{i,i+1\}}(t) = 1$ and $u^i_{\{1,2,3\}}(t) = 0$ for every $i \in I$ and every $t \in \mathbf{R}^+$. The game is described by the matrix given in Fig. 1 in which players 1, 2 and 3 choose, respectively, a row, a column and a matrix. We assume that the three players have the same discount rate $\delta \geqslant 0$. The value of δ plays no role in the analysis. In particular, we allow for the possibility that $\delta = 0$, allowing in effect for the case of an un-discounted game.

We prove that this game has no ϵ-equilibrium, provided $\epsilon > 0$ is small enough. It is interesting to recall that three-player game of timing in *discrete* time do have a subgame-perfect equilibrium (see [8,29]). Thus, this example stands in sharp contrast with known results in discrete time.

We first verify that this game has no (exact) equilibrium. Let σ be a plan profile. If σ is an equilibrium, the probability that the game terminates at time 0 is below one. Otherwise, at least one player, say player 1, would act with probability one at time 0. By the equilibrium condition, player 2 would act with probability 0: given that player 1 acts, *act* is a strictly dominated action for player 2. Hence, player 3 would act with probability one at time 0, and player 1 would find it optimal not to act at time 0—a contradiction. Next, given that the game does not terminate at time 0, each player i can get a payoff arbitrarily close to one, by acting immediately after time 0, that is, by acting at time $t > 0$, where t is sufficiently small so that the probability that σ^{i+1} or σ^{i+2} act in the time interval $(0, t]$ is arbitrarily small. Thus, the continuation equilibrium payoff of each player must be at least one—a contradiction to the zero-sum property. Hence σ is not an equilibrium.

[9] Here addition is understood modulo 3.

We now prove that the game has no ϵ-equilibrium. For every $w \in [-1,1]^3$ let $G(w)$ be the one-shot game with payoff matrix as in Fig. 1, where the payoff if no player acts is w. The result of the previous paragraph can be rephrased as follows: for every $w \in [-1,1]^3$ with $\sum_{i=1}^{3} w^i = 0$, the probability that the game terminates at time 0, under any Nash equilibrium in $G(w)$, is strictly less than 1. Since the correspondence that assigns to each $w \in [-1,1]^3$ and every $\epsilon > 0$ the set of ϵ-equilibria of the game $G(w)$ has a closed graph, there is $\epsilon > 0$ such that for every $w \in [-1,1]^3$ with $\sum_{i=1}^{3} w^i = 0$, the probability that the game terminates at time 0, under any ϵ-equilibrium in $G(w)$, is strictly less than $1 - 2\epsilon$.

Let σ be an ϵ-equilibrium of the timing game. In particular, the probabilities $\sigma^i(\{0\})$ assigned to *act* at time zero form an ϵ-equilibrium of the game $G(w)$, taking for w the continuation payoff vector in the game. Since the game is zero-sum, the continuation payoff at time 0 of at least one player is non-positive. As argued above, by acting right after time 0, this player can improve his payoff by almost 1 if the game is not terminated at time 0. By the previous paragraph, this event has probability at least 2ϵ, hence the deviation improves by more than ϵ—a contradiction.

5.3. A symmetric game with no pure equilibrium

We here provide a symmetric two-player game with no pure equilibrium. It is defined by

$$\alpha_2(t) = 0 \text{ for every } t: \text{ if both players act simultaneously, no-one gets anything,}$$

$$\alpha_1(t) = t\mathbf{1}_{\{t \leqslant 1\}} + (2-t)\mathbf{1}_{\{1<t<5/2\}} - \tfrac{1}{2}\mathbf{1}_{\{t \geqslant 5/2\}}: \text{ if only one player is to act, he will do it at time 1,}$$

$$\beta_1(t) = \tfrac{1}{4}\mathbf{1}_{\{t \leqslant 1/4\}} + (\tfrac{1}{2} - t)\mathbf{1}_{\{1/4<t<3/2\}} - \mathbf{1}_{\{t \geqslant 3/2\}}.$$

Graphically, the payoff functions look as follows.

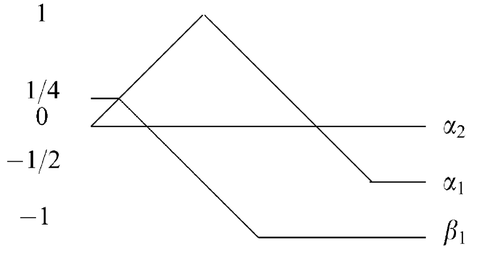

We assume $\delta_1 = \delta_2 = 0$, but our arguments remain valid as long as the discount rates are sufficiently small.

Observe that the plan profile in which both players act at a random time uniformly chosen from the interval $[1/4, 1/4+\epsilon]$ is a symmetric ϵ-equilibrium. Indeed, the corresponding payoff to both players is $1/4$, whereas the best payoff a player can get by deviating is at most $1/4+\epsilon$. It is also easy to verify that the plan profile in which player 1 acts at time $1/4$ and player 2 acts at time $1/4+\varepsilon$ is an ε-equilibrium.

Assume that there is a pure Nash equilibrium. If both players act simultaneously at time $t_* \in \mathbf{R}^+ \cup \{\infty\}$, the equilibrium payoff is 0. Since $\beta_1(t) > \alpha_2(t)$ for $t < 1/2$, we must have $t_* \geqslant 1/2$. Each player would then rather act alone at some time $0 < t < \min\{2, t_*\}$.

By symmetry, it is now sufficient to assume that players 1 acts at time t_*, and player 2 acts at time $t_{**} > t_*$, possibly infinity. Since $\alpha_2(t) > \beta_1(t)$ for $t > 1/2$, we must have $t_* \leqslant 1/2$. Since the function α_1 increases until $t=1$, player 1 is better off by acting at any time $t \in (t_*, \min\{t_{**}, 1\})$.

5.4. A game with cumulative payoffs and no Markov equilibrium

Consider the following three-player game with constant cumulative payoffs.

	Don't Act		Act	
	Don't Act	Act	Don't Act	Act
Don't Act		$2,1,-\frac{5}{2}$	$-\frac{5}{2},2,1$	$2,1,1$
Act	$1,-\frac{5}{2},2$	$1,1,2$	$1,2,1$	$1,1,1$

As shown in Section 4, the only Markov strategies $\widehat{\sigma}^i = (\sigma^i_t)_t$ are either (i) σ^i_t acts at time t, for every t, or (ii) σ^i_t assigns probability 1 to ∞, for each t or (iii) σ^i_t is an exponential distribution over $[t, +\infty)$.

Suppose to the contrary that we are given a Markov equilibrium in this game.

If some player, say player 3, follows the strategy of type (i), the best reply of player 2 is to follow the strategy of type (i) as well, so that the best reply of player 1 is to follow a strategy of type (ii) or (iii), so that the best reply of player 3 is to follow a strategy of type (ii) or (iii) as well, a contradiction.

Otherwise, all players play strategies of type (ii) or (iii), so that either the game never terminates, or it terminates by a single player. If the game terminates by a single player the sum of payoffs to the three players is $1/2$. In particular, in this case the expected payoff of at least one player is below $1/2$, but that player can receive 1 by acting at time 0.

Consequently, the game admits no Markov ε-equilibrium, provided ε is sufficiently small.

Observe that the strategy profile in which each player acts with probability $1/2$ whenever t is an integer, and does not act otherwise, is a non-Markov Nash equilibrium.

6. An equilibrium existence result

We here prove Theorem 1.5. It will be convenient to describe a mixed plan σ^i of player i by its cumulative distribution function (c.d.f.), i.e., by the function $F^i : \mathbf{R}^+ \to [0,1]$ defined by $F^i_t = \sigma^i([0,t])$. Plainly, F^i is right-continuous and non-decreasing. Note also that $1 - \lim_{t \nearrow \infty} F_t^i$ is the probability under σ^i that player i never acts, and that F_0^i is the probability that player i acts immediately. We let $\mathscr{F}$ denote the set of all such functions F^i.

Given $F \in \mathscr{F}$ and $t \in [0, \infty]$, we let $F_{t-} = \lim_{s \nearrow t} F_s$ denote the left-limit of F at t (with $F_{0-} := 0$ and $F_{\infty -} = \lim_{t \to \infty} F_t$) and we denote by $\Delta F_t := F_t - F_{t-}$ the jump of F at t.

When expressed in terms of c.d.f's, formula (1) reduces to

$$\gamma_0^l(F^1, \ldots, F^I) = \sum_{i \in I} \int_{[0,\infty)} e^{-\delta_l t} u_{\{i\}}^l(t) \prod_{j \neq i} (1 - F_t^j)\, dF_t^i + \sum_{S \subseteq I, |S| \geqslant 2} \sum_{t=0}^{\infty} u_S^l(t) \prod_{i \in S} \Delta F_t^i \prod_{i \notin S} (1 - F_t^i),$$

where the integral is a Stieltjes integral w.r.t. F^i (the notation $\int_{[0,\infty)}$ stresses that the jump of F^i at zero is explicitly taken into account in the value of the integral).

The proof of Theorem 1.5 relies on a compactness principle. We shall exhibit a compact set $\mathscr{G}$ of profiles that satisfies:

(a) if there is an ε-equilibrium, then there is an ε-equilibrium in $\mathscr{G}$, and
(b) the payoff function $\gamma(\cdot)$ is continuous on $\mathscr{G}$.

The second property will imply that any accumulation point of ε-equilibria in $\mathscr{G}$, as ε goes to 0, is an equilibrium, while the first property, together with the compactness of $\mathscr{G}$, will imply that under the assumptions of Theorem 1.5 such an accumulation point exists.

The set $\mathscr{F}^I$ of all profiles, endowed with the weak topology, does not satisfy the second property, since the payoff function is not continuous over $\mathscr{F}^I$. Discontinuities may arise for two reasons. First, in the weak topology, several atoms may merge to a single atom at the limit. Second, a sequence of non-atomic distributions may weakly converge to an atomic distribution.

We illustrate these two phenomena with two examples. Both examples involve two players. We let $F = (F^1, F^2)$ be the profile in which both players act with probability 1 at time 0: $F_t^i = 1$ for every $t \in \mathbf{R}^+$.

Example 1. Player 1 acts with probability 1 at time 0, while player 2 acts with probability 1 at time $1/n$. Formally, for every $n \in \mathbf{N}$, $F^1(n) = F^1$ whereas $F_t^2(n) = \mathbf{1}_{t \geqslant 1/n}$. Plainly the sequence $(F(n))$ weakly converges to F, but $\gamma(F(n)) = u_{\{1\}}$ while $\gamma(F) = u_{\{1,2\}}$.

Example 2. Both players act uniformly in the interval $[0, 1/n]$. Formally, $F_t^1(n) = F_t^2(n) = \min\{1, nt\}$. The sequence $(F(n))$ weakly converges to F. Since for every $n \in \mathbf{N}$ the probability that under $F(n)$ both players act simultaneously is 0, $\gamma(F(n)) = \frac{1}{2} u_{\{1\}} + \frac{1}{2} u_{\{2\}}$, while $\gamma(F) = u_{\{1,2\}}$.

Roughly speaking, the auxiliary space $\mathcal{G}$ contains all profiles $G=(G^1,\dots,G^I)$ that satisfy (A) if G^i has a jump of ΔG^i_t at t, then all G^j's are constant in the interval $(t-\Delta G^i_t, t)$, and (B) the slope of $\frac{1}{n}\sum_i G^i$ is 1 whenever this function is continuous.

The first requirement implies that as one goes to the limit, it cannot be that two atoms merge. Indeed, if for each $n\in\mathbf{N}$ $G^i(n)$ and $G^j(n)$ have discontinuities at t_n and s_n, respectively, with $t_n<s_n$, then $\Delta G^j_{s_n}(n)$ is bounded by s_n-t_n. Therefore, if $\lim s_n = \lim\, t_n$ then the atom of $G^j(n)$ at s_n vanishes at the limit.

The second requirement implies that a sequence of non-atomic distributions in $\mathcal{G}$ cannot converge to an atomic distribution, since the slope of $G^i(n)$ is uniformly bounded by I.

We now turn to the formal presentation. Recall that $\mathcal{F}$ is the space of all functions $F:\ \mathbf{R}^+\to[0,1]$ that are non-decreasing and right-continuous. It is in bijection with the set of probability measures μ over $[0,+\infty]$. We denote by λ the Lebesgue measure over $[0,+\infty)$. The set of atoms of μ^i (or equivalently, of discontinuities of F^i) is denoted by A_{μ^i}. Let $\mathcal{G}\subset(\mathcal{F})^I$ be the space of all $\mu=(\mu^1,\dots,\mu^n)$ that satisfy the following conditions.

(0) The support of each μ^i is an interval $[0,T_i]$, with $T_i\leqslant I$.

(A) For each $i\in I$ and $t\in A_{\mu^i}$, one has $\mu^j_{[t-\mu^i_t,t)}=0$ for every $j\in I$. Set $T_\mu := \mathbf{R}^+\backslash(\bigcup_i \bigcup_{t\in A_{\mu^i}}[t-\mu^i_t,t])$.

(B) One has $\frac{1}{I}\sum_i\mu^i_A=\frac{1}{I}\sum_i\lambda_{A\cap[0,T_i]}$, for every $A\subseteq T_\mu$.

By Helly's Theorem [3, Theorem 25.9] and Theorem 25.10 in [3], the set $\mathcal{G}$ is compact for the topology of weak convergence.

Theorem 1.5 follows immediately from Lemmas 6.1 and 6.2, using the compactness of $\mathcal{G}$.

Lemma 6.1. *Let $\varepsilon>0$ be given. If the game has an ε-equilibrium, then it has an ε-equilibrium in $\mathcal{G}$.*

The proof of this lemma appears in Section 6.1.

We denote by Δ^i the set of pure plans of player i.

Lemma 6.2. *The payoff function γ is continuous over $\mathcal{G}$. Moreover, let $(G(n))_{n\in\mathbf{N}}$ be a convergent sequence in $\mathcal{G}$, with limit G, and let $\tilde{G}^i\in\Delta^i$, for some $i\in I$. Then there exists a sequence $\tilde{G}^i(n)\in\Delta^i$, such that*

$$\lim_{n\to+\infty}\gamma^i(\tilde{G}^i(n),G^{-i}(n))=\gamma^i(\tilde{G}^i,G^{-i}).$$

The proof of this lemma appears in Section 6.2.

6.1. Time-changes

Our goal in this section is to prove Lemma 6.1. A *time-change* is a non-decreasing, right-continuous function defined over some interval of $\mathbf{R}^+$, with values in $\mathbf{R}^+$. Given an ε-equilibrium $(F^1, \ldots, F^I)$, we shall construct a time-change u such that the profile $(G^1, \ldots, G^I)$ defined by $G^i_t = F^i_{u(t)}$ is in $\mathscr{G}$, and is an ε-equilibrium.

For $s \in \mathbf{R}^+$, we define the *s-level set* of F to be the interval $F^{-1}(\{s\})$.

6.1.1. Straightening F

We here define a first time-change, relative to a given *continuous* function $F \in \mathscr{F}$. In effect, the clock will be adjusted in such a way that: (i) the duration of the level sets of F will not be affected and (ii) the increasing portions of F will be transformed into affine portions with slope one.

We first introduce a usual time-change (see, e.g., [26, Chapter 0]):

$$C_s = \inf\{t \geqslant 0 \mid F_t > s\}, \text{ for } s \in [0, F_{\infty -}).$$

The function C is defined on $[0, F_\infty)$, with values in $\mathbf{R}^+$. It is increasing (since F is continuous) and right-continuous. Moreover, the s-level set of F coincides with the interval $[C_{s-}, C_s)$.

Plainly, the function $s \mapsto F_{C_s}$ increases linearly from 0 to $F_{\infty -}$, at unit speed. We now proceed to introduce the non-trivial level sets of F. More precisely, we will let the value of F at time t be reached, under the time-change, at a time which is the sum of two components, the time F_{t-} that is needed to reach the level F_{t-} at unit speed, and the cumulative length of all level sets up to time t.

As mentioned above, the length of the $F_{t'}$-level set is $\Delta C_{F_{t'}}$. Therefore, the cumulative length of all level sets up to time t is

$$\sum_{t'<t} \Delta C_{F_{t'}} + t - C_{F_{t-}},$$

the first summation is the total length of all level sets lying entirely to the left of t, while $t - C_{F_{t-}}$ is the time elapsed since the current level set was initiated.

This leads us to introduce the function v_1 defined by

$$v_1(t) := F_t + \sum_{t'<t} \Delta C_{F_{t'}} + t - C_{F_{t-}}.$$

The next lemma lists few easy properties of v_1. The proof is omitted.

Lemma 6.3. *The function v_1 is continuous and increasing. In addition, $v_1(0) = 0$, and*[10] *$v_1(\infty -)$ is infinite or finite depending on whether F is eventually constant or not.*

[10] Recall that $f(\infty -) = \lim_{t \to \infty} f(t)$.

6.1.2. Playing with level sets

We here define a second time-change, relative to an arbitrary $F \in \mathcal{F}$. In effect, we shall adjust the length of level sets of F to the size of nearby discontinuities. Formally, the value of F at time t will be reached, according to the new clock, at time s, which is obtained from t by subtracting the cumulative length of all level sets prior to t, and by adding the cumulative sum of jumps prior to time t. That is, we set

$$v_2(t) = t + \sum_{t'<t} \Delta F_{t'} - \left(\sum_{t'<t} \Delta C_{F_{t'}} + t - C_{F_t-} \right)$$

$$= C_{F_t-} + \sum_{t'<t} \Delta F_{t'} - \sum_{t'<t} \Delta C_{F_{t'}}.$$

The proof of the following basic properties of v_2 is left to the reader.

Lemma 6.4. *The function v_2 is non-decreasing and right-continuous.*

6.1.3. Time changes and the equilibrium property

We let here $\varepsilon > 0$ and an ε-equilibrium $(F^i)_{i \in I}$ be given. Loosely speaking, our goal is to show that applying the above time changes to the profile $(F^i)_{i \in I}$ does not affect the ε-equilibrium property.

We will make extensive use of the following change-of-variable formula for Stieltjes integrals, which is a minor variation upon Proposition 4.10 in [26, Chapter 0].

Lemma 6.5. *Let $u : [a, b] \to \mathbf{R}^+$ be a right-continuous, non-decreasing map. Let $F \in \mathcal{F}$ and g be a bounded, Borel measurable map. Assume that $F_{u(t)-} = F_{u(t)}$ whenever $\Delta u(t) > 0$. Then*

$$\int_{[u(a), u(b)]} g(s)\, dF_s = \int_{[a,b]} g(u(t))\, dF_{u(t)}.$$

For $i \in I$, we let $\tilde{F}^i$ denote the continuous part of F^i: $\tilde{F}^i_t = F^i_t - \sum_{t'<t} \Delta F^i_t$ for $t \in \mathbf{R}^+$. Next, we set $\tilde{F} = \frac{1}{I} \sum_{i \in I} \tilde{F}^i$ and consider the function v_1 relative to $\tilde{F}$, as defined in Section 6.1.1. Let u_1 be the inverse map of v_1.

For $i \in I$, we define G^i to be the image of F^i under the time-change u_1: $G^i_s = F^i_{u_1(s)}$ for $s < v_1(\infty-)$ and $G^i_s = F^i_{\infty-}$ for $s \geqslant v_1(\infty-)$. Plainly, $G^i \in \mathcal{F}$ for each $i \in I$.

Lemma 6.6. *The profile $(G^i)_{i \in N}$ is an ε-equilibrium.*

Proof. We fix $i \in I$, and prove that player i has no pure deviation that increases his payoff by more than ε. Let $\tilde{G}^i$ be a pure plan.

Case 1: $\tilde{G}^i_s = 0$ *for every* $s \in \mathbf{R}^+$ (*player i never acts*). Since $(F^1, \dots, F^I)$ is an ε-equilibrium,

$$\gamma^i(\tilde{G}^i, G^{-i}) = \gamma^i(\tilde{G}^i, F^{-i}) \leqslant \gamma^i(F^i, F^{-i}) + \varepsilon = \gamma^i(G^i, G^{-i}) + \varepsilon,$$

where the equalities follow by the change-of-variable formula.

Case 2: $\tilde{G}^i_s = 1_{s \geqslant s_0}$ *for some* $s_0 \in \mathbf{R}^+$ (*player i acts at time* s_0). If $s_0 < v_1(\infty -)$, we set $t_0 = u_1(s_0)$ and we define $\tilde{F}^i_t = 1_{t \geqslant t_0}$.

Since $(F^1, \dots, F^I)$ is an ε-equilibrium,

$$\gamma^i(\tilde{G}^i, G^{-i}) = \gamma^i(\tilde{F}^i, F^{-i}) \leqslant \gamma^i(F^i, F^{-i}) + \varepsilon = \gamma^i(G^i, G^{-i}) + \varepsilon,$$

where the equalities follow by the change-of-variable formula.

Assume now that $s_0 \geqslant v_1(\infty -)$. In particular, $v_1(\infty -) < \infty$. For $\bar{s} \leqslant v_1(\infty -)$, define $1^{\bar{s}} \in \mathscr{F}$ by $1^{\bar{s}}_s = 1_{s \geqslant \bar{s}}$.

Plainly,

$$\begin{aligned}\gamma^i(\tilde{G}^i, G^{-i}) &= \gamma^i(1^{v^1(\infty -)}, G^{-i}) \\ &= \lim_{\bar{s} \nearrow v^1(\infty -)} \gamma^i(1^{\bar{s}}, G^{-i}) \leqslant \gamma^i(G^i, G^{-i}) + \varepsilon,\end{aligned}$$

where the last inequality follows by the analysis of the case $\bar{s} < v^1(\infty -)$. □

We now analyze the impact of the second time-change on $(G^i)_{i \in I}$. We let v_2 be the time-change relative to $\frac{1}{I}\sum_{i \in I} G^i$, as defined in Section 6.1.2. We let u_2 be the generalized inverse of v_2: $u_2(s) = \inf\{t : v_2(t) > s\}$. The function u_2 is defined over $[0, v_2(\infty -))$, is right-continuous and non-decreasing. Note that a level set of u_2 with positive length corresponds to a jump in v_2. Also, a jump in u_2 corresponds to a non-trivial level set of v_2. For $i \in I$, we let $H^i_s = G^i_{u_2(s)}$ for $s < v_2(\infty -)$ and $H^i_s = G^i(\infty -)$ for $s \geqslant v_2(\infty -)$.

Lemma 6.7. *The profile* $(H^i)_{i \in I}$ *is an* ε*-equilibrium in* $\mathscr{G}$.

Proof. We prove that player i has no pure profitable deviation. Let $\tilde{H}^i \in \Delta^i$ be arbitrary. The case $\tilde{H}^i = 0$ can be dealt with as in the previous proof. Assume now that $\tilde{H}^i = 1_{s \geqslant s_0}$ for some $s_0 \in \mathbf{R}^+$. As observed at the end of the previous proof, it is enough to deal with the case $s_0 < v_2(\infty -)$. Set $t_0 = u_2(s_0)$. If u_2 is continuous at s_0, the inequality $\gamma^i(\tilde{H}^i, H^{-i}) \leqslant \gamma^i(H^i, H^{-i}) + \varepsilon$ follows by the change-of-variable formula.

If u_2 is not continuous at s_0, then the change-of-variable cannot be applied (at least for the integral w.r.t. $\tilde{H}^i$). In that case, we let (s^n) be a increasing sequence of continuity points of u_2, that converges to s_0, and we let $\tilde{H}^{i,n}_s = 1_{s \geqslant s^n}$. It is not difficult to check that $\lim_{n \to \infty} \gamma^i(\tilde{H}^{i,n}, H^{-i}) = \gamma^i(\tilde{H}^i, H^{-i})$. Hence, by the previous paragraph, $\gamma^i(\tilde{H}^i, H^{-i}) \leqslant \gamma^i(H^i, H^{-i}) + \varepsilon$. Therefore, $(H^i)_{i \in I}$ is an ε-equilibrium. □

6.2. Proof of Lemma 6.2

We shall only prove the first assertion of Lemma 6.2. The second one can be established using similar ideas.

Let $(F(n))$ be a sequence in $\mathcal{G}$ that weakly converges to $F \in \mathcal{G}$.

For every non-empty subset S of I we let π_S be the probability that under F the game terminates, and the terminating coalition is S. For $n \in \mathbf{N}$, we denote by $\pi_S(n)$ the analogous probability under $F(n)$.

Since $\gamma(F) = \sum_S \pi_S u_S$ and $\gamma(F(n)) = \sum_S \pi_S(n) u_S$, it is enough to prove that $\lim_{n\to\infty} \pi_S(n) = \pi_S$ for every S.

Note first that $F^i_t = \lim_{n\to\infty} F^i_t(n)$ for each $i \in I$ and for every continuity point t of F^i. In particular, the equality holds for λ-a.e. $t \in \mathbf{R}^+$, which implies

$$\lim_{n\to\infty} F^i_{t-}(n) = F^i_{t-}, \quad \text{for every } t \in \mathbf{R}^+ \text{ and every } i \in I. \tag{6}$$

Step 1: *Relating atoms.* Let t be an atom of F^i, for some $i \in I$. Let $S^* = \{i \in I, \Delta F^i_t > 0\}$ be the set of i's such that t is an atom of F^i.

We show that for every n there is $\tau(t;n) \in \mathbf{R}^+$ such that

(A.i) $\lim_{n\to\infty} \tau(t;n) = t$,
(A.ii) $\lim_{n\to\infty} \Delta F^i_{\tau(t;n)}(n) = \Delta F^i_t$ for each $i \in I$, and
(A.iii) $\lim_{n\to\infty} F^i_{\tau(t;n)}(n) = F^i_t$ for each $i \in I$.

Let $\varepsilon \in (0,t)$ satisfy $\Delta F^i_t > (2I+5)\varepsilon$ for every $i \in S^*$.[11] In addition, we assume that both $t+\varepsilon$ and $t-\varepsilon$ are continuity points of F^i.

For n large enough, $F^i_{t+\varepsilon}(n) - F^i_{t-\varepsilon}(n) \geqslant F^i_{t+\varepsilon} - F^i_{t-\varepsilon} - \varepsilon \geqslant \Delta F^i_t - \varepsilon$. Let $\tau^i(t;n)$ be the infimum over all discontinuities of $F^i(n)$ in the interval $[t-\varepsilon, t+\varepsilon]$, and set $\tau(t;n) = \min_{i\in S^*} \tau^i(t;n)$. Since $F(n) \in \mathcal{G}$, one has

$$\sum_{s\in[t-\varepsilon,t+\varepsilon]} \Delta F^i_s(n) \geqslant F^i_{t+\varepsilon}(n) - F^i_{t-\varepsilon}(n) - 2I\varepsilon, \text{ and } \sum_{s\in(\tau(t;n),t+\varepsilon]} \Delta F^i_s(n) \leqslant 2\varepsilon. \tag{7}$$

Eq. (7) implies that $\Delta F^i_{\tau(t;n)} \geqslant F^i_{t+\varepsilon}(n) - F^i_{t-\varepsilon}(n) - 2(I+1)\varepsilon \geqslant \Delta F^i_t - (2I+3)\varepsilon$. Therefore, for $i \in S^*$, $\Delta F^i_{\tau(t;n)} > 0$, so that $\tau^i(t;n) = \tau(t;n)$, and moreover $\Delta F^i_{\tau(t;n)}(n) \geqslant \Delta F^i_t - 5\varepsilon$.[12] Therefore,

$$\liminf_n \Delta F^i_{\tau(t;n)}(n) \geqslant \Delta F^i_t - 5\varepsilon. \tag{8}$$

This implies that $\lim_{n\to\infty} \tau(t;n) = t$, so that (A.i) holds. Indeed, otherwise there would be a subsequence of $(\tau(t;n))_n$—still denoted $(\tau(t;n))_n$—such that

[11] If $t = 0$, the condition $\varepsilon < t$ is omitted, and in the sequel $t - \varepsilon$ is replaced by t.

[12] For further use, we note the following additional consequence. Strictly speaking, the sequence $(\tau(t;n))_n$ depends on ε, and should rather be denoted by $(\tau^\varepsilon(t;n))_n$. For $\varepsilon' < \varepsilon$, one has $\tau^\varepsilon(t;n) \leqslant \tau^{\varepsilon'}(t;n)$ whenever the two sides are well-defined. The last inequality in the text implies that $\tau^\varepsilon(t;n) = \tau^{\varepsilon'}(t;n)$ for n large enough. In that sense, the sequence $(\tau^\varepsilon(t;n))_n$ is (asymptotically) independent of ε.

$\lim_{n\to+\infty}\tau(t;n) = t' \neq t$. By repeating the above argument with $\varepsilon' \in (0,\varepsilon)$ small enough so that $t' \notin [t-\varepsilon', t+\varepsilon']$, we would construct another sequence $(\tau'(t;n))_n$ such that $\lim_{n\to+\infty}\Delta F^i_{\tau'(t;n)}(n) = \Delta F^i_t$, for each $i \in I$—a contradiction to the second inequality in (7). By weak convergence, (A.i) implies that (A.ii) holds whenever $\Delta F^i_t = 0$, or, equivalently, whenever $i \notin S^*$.

We now prove that (A.ii) holds for $i \in S^*$ as well. Since $F^i_{t+\varepsilon} - F^i_{t-\varepsilon} \leqslant \Delta F^i_t + I\varepsilon$, one has $\Delta F^i_{\tau(t;n)} \leqslant F^i_{t+\varepsilon}(n) - F^i_{t-\varepsilon}(n) \leqslant \Delta F^i_t + (I+1)\varepsilon$, provided n is large enough. Therefore, $\limsup_n \Delta F^i_{\tau(t;n)}(n) \leqslant \Delta F^i_t + 2\eta$, which, together with (8), and since ε is arbitrary, yields

$$\lim_{n\to+\infty} \Delta F^i_{\tau(t;n)-}(n) = \Delta F^i_t, \quad \text{for each } i \in I, \tag{9}$$

so that (A.ii) holds.

Finally, we show that $\lim_{n\to\infty} F^i_{\tau(t;n)-}(n) = F^i_{t-}$, for each $i \in I$, which, together with (A.ii), implies that (A.iii) holds. W.l.o.g., we may assume that the sequence $(\tau(t;n))_n$ is monotonic. Assume first that it is non-decreasing, and let $\varepsilon > 0$ be given. Choose $t' < t$ such that $F^i_{t'-} \geqslant F^i_{t-} - \varepsilon$. Then, for n large enough, one has by (6)

$$F^i_{t'-} - \varepsilon \leqslant F^i_{t'-}(n) \leqslant F^i_{\tau(t;n)}(n) \leqslant F^i_{t-}(n) \leqslant F^i_{t-} + \varepsilon.$$

If the sequence $(\tau(t;n))_n$ is non-increasing, then $F^i_{\tau(t;n)-}(n) = F^i_{t-}(n)$ for n large, hence by (6) the claim still holds.

Step 2: $\lim_{n\to\infty}\pi_S(n) = \pi_S$ *whenever* $|S| \geqslant 2$. Suppose $S \subseteq I$ with $|S| \geqslant 2$. For the sake of clarity, we set $g^S_t := \prod_{j\notin S}(1-F^j_t)$, and $h^S_t := \prod_{i\in S}\Delta F^i_t$ for $S \subset I$ and $t \in \mathbf{R}^+$. Then

$$\pi_S = \sum_{t\in\mathbf{R}^+} \prod_{j\notin S}(1-F^j_t)\prod_{i\in S}\Delta F^i_t = \sum_{t\in\mathbf{R}^+} g^S_t h^S_t,$$

and a similar expression holds for $\pi_S(n)$.

Fix $i \in S$, and let $\varepsilon > 0$ be arbitrary. Let $A \subset \mathbf{R}^+$ be a finite set of atoms that almost exhausts the atoms of F^i: $\sum_{t\in A}\Delta F^i_t \geqslant \sum_{t\in\mathbf{R}^+}\Delta F^i_t - \varepsilon$.

By (A.ii) and (A.iii), $\lim_{n\to+\infty} g^S_{\tau(t;n)}(n)h^S_{\tau(t;n)}(n) = g^S_t h^S_t$ for every $t \in \mathbf{R}^+$. In particular, since A is a finite set,

$$\lim_{n\to\infty} \sum_{t\in A} g^S_{\tau(t;n)}(n)h^S_{\tau(t;n)}(n) = \sum_{t\in A} g^S_t h^S_t. \tag{10}$$

Moreover,

$$\sum_{t\notin A} g^S_t h^S_t \leqslant \sum_{t\notin S} \Delta F^i_t < \varepsilon. \tag{11}$$

For $n \in \mathbf{N}$ set $A_n := \{\tau(t;n) : t \in A\}$. Our goal is to prove that

$$\lim_{n\to\infty} \sum_{t\notin A_n} g^S_t h^S_t = 0, \tag{12}$$

which, together with (10) and (11) implies that $\lim_{n\to\infty}\pi_S(n) = \pi_S$, provided $|S|\geqslant 2$.

Let $\delta_n := \sup\{\Delta F^i_s(n) : s\notin A_n, i\in I\}$ (with $\sup\emptyset = 0$) be the maximal size of the remaining discontinuities, and let t_n achieve the supremum, up to $1/n$. We claim that $\lim_{n\to\infty}\delta_n = 0$. Indeed, since the support of F^i is included in $[0, I]$, the sequence (t_n) converges, up to a subsequence, to some $t\in\mathbf{R}^+$. If $\Delta F^i_t > 0$ for some $i\in I$, then $\lim_{n\to\infty}\Delta F^j_{t_n}(n) = 0$ since $t_n \neq \tau(t;n)$ for each n. If $\Delta F^i_t = 0$ then by weak convergence $\lim_{n\to\infty}\Delta F^i_t(n) = 0$. Therefore, $\lim_{n\to\infty}\delta_n = 0$.

For every two sequences $(x_k, y_k)_{k=1}^{\infty}$ such that $0\leqslant x_k, y_k\leqslant\delta<1$ and $\sum_k x_k, \sum_k y_k\leqslant 1$ one has $\sum_k x_k y_k\leqslant\delta$. Since $|S|\geqslant 2$, and since $g^S_t(n)h^S_t(n)$ is non-zero on at most a countable set of t's, (12) holds.

Step 3: *$\lim_{n\to\infty}\pi_S(n) = \pi_S$ whenever $S = \{i\}$ is a singleton.* Let $\varepsilon>0$ be arbitrary. We prove that $\pi_{\{i\}} - 3\varepsilon\leqslant\liminf_{n\to\infty}\pi_{\{i\}}(n)$ and $\limsup_{n\to\infty}\pi_{\{i\}}(n)\leqslant\pi_{\{i\}} + 3\varepsilon$.

As in step 2, let $A\subset\mathbf{R}^+$ be a finite set such that $\sum_{t\in A}\Delta F^i_t\geqslant\sum_{t\in\mathbf{R}^+}\Delta F^i_t - \varepsilon$. We assume that A contains 0 if $\Delta F^i_0>0$.

Since A is finite, we may assume w.l.o.g. that for every n, the finite set $\{\tau(t;n), t\in A\}$ contains $|A|$ different elements.

Denote $\widehat{F}^i_t = F^i_t - \sum_{s<t, s\in A}\Delta F^i_s$ and $\widehat{F}^i_t(n) = F^i_t(n) - \sum_{s<t, s\in A}\Delta F^i_{\tau(s;n)}$. This is the part of F^i (resp. $F^i(n)$) without the atoms in A. Then $(\widehat{F}^i(n))$ weakly converges to $\widehat{F}^i$.

Choose a finite sequence $0<t_1<\ldots<t_K = I+1$ such that

(B.i) $\widehat{F}^i_{t_{k+1}} - \widehat{F}^i_{t_k}<\varepsilon$ for each $k = 0, \ldots, K-1$ (with $\widehat{F}^i_{t_0} = 0$).
(B.ii) $t_1, \ldots, t_K$ are continuity points of F^j, for every $j\in I$.

We now modify the distributions F^i and $(F^i(n))_{n\in\mathbf{N}}$, and construct completely atomic distributions $\bar{F}^i$, $\underline{F}^i$, $(\bar{F}^i(n))_{n\in\mathbf{N}}$, and $(\underline{F}^i(n))_{n\in\mathbf{N}}$ as follows:

- $\bar{F}^i$: every $t\in A$ is an atom of $\bar{F}^i$ with size ΔF^i_t. In addition, each $(t_k)_{k=1}^{K-1}$ is an atom; the weight of this atom is equal to $\widehat{F}^i_{t_{k+1}} - \widehat{F}^i_{t_k}$.
- $\underline{F}^i$: every $t\in A$ is an atom of $\bar{F}^i$ with size ΔF^i_t. In addition, each $(t_k)_{k=2}^{K}$ is an atom; the weight of this atom is equal to $\widehat{F}^i_{t_k} - \widehat{F}^i_{t_{k-1}}$.
- $\bar{F}^i(n)$ and $\underline{F}^i(n)$ are defined analogously w.r.t. $F^i(n)$.[13]

Thus, under $\bar{F}^i$ player i acts earlier than under F^i, whereas under $\underline{F}^i$ he acts later.

Observe that in this definition, we ignored the part of $\widehat{F}^i$ prior to time t_1, but by (B.i) this part has small weight. Let $\bar{\pi}_{\{i\}}$, $\underline{\pi}_{\{i\}}$, $\bar{\pi}_{\{i\}}(n)$ and $\underline{\pi}_{\{i\}}(n)$ be analogous to $\pi_{\{i\}}$ under $(\bar{F}^i, F^{-i})$, $(\underline{F}^i, F^{-i})$, $(\bar{F}^i(n), F^{-i}(n))$ and $(\underline{F}^i(n), F^{-i}(n))$, respectively.

[13] Note that, for n large enough, the two sets $\{\tau(t;n), t\in A\}$ and $\{t_k, k = 1, \ldots, K\}$ are disjoint.

By (B.i) we have

$$\bar{\pi}_{\{i\}} + \varepsilon \geqslant \pi_{\{i\}} \geqslant \underline{\pi}_{\{i\}}, \quad \text{and} \quad \bar{\pi}_{\{i\}}(n) + \varepsilon \geqslant \pi_{\{i\}}(n) \geqslant \underline{\pi}_{\{i\}}(n) \ \ \forall n \in \mathbf{N}. \tag{13}$$

Moreover,

$$\bar{\pi}_{\{i\}} - \underline{\pi}_{\{i\}} < 2\varepsilon. \tag{14}$$

Since $\bar{F}^i$ is completely atomic, we can derive an explicit formula for $\bar{\pi}_{\{i\}}$:

$$\bar{\pi}_{\{i\}} = \sum_{k=1}^{K-1} \prod_{j \neq i} (1 - F^j_{t_k}) \Delta \bar{F}^i_{t_k} + \sum_{t \in A} \prod_{j \neq i} (1 - F^i_j) \Delta \bar{F}^i_t. \tag{15}$$

One has a similar expression for $\underline{\pi}_{\{i\}}$. For $\bar{\pi}_{\{i\}}(n)$ one has

$$\bar{\pi}_{\{i\}}(n) = \sum_{k=1}^{K-1} \prod_{j \neq i} (1 - F^j_{t_k}(n)) \Delta \bar{F}^i_{t_k}(n) + \sum_{t \in A} \prod_{j \neq i} (1 - F^j_{\tau(t;n)}(n)) \Delta \bar{F}^i_{\tau(t;n)}(n). \tag{16}$$

By (A.ii) and (A.iii), since $(F^i(n))$ weakly converges to F^i, and since (t_k) are continuity points of F^i, $\lim_{n\to\infty} \Delta \bar{F}^i_{t_k}(n) = \Delta \bar{F}^i_{t_k}(n)$. Since the (t_k) are continuity points of $(F^j)_{j \neq i}$, $\lim_{n\to\infty} \bar{F}^j_{t_k}(n) = \bar{F}^j_{t_k}(n)$. Therefore, again using (A.ii) and (A.iii), we obtain $\lim_{n\to\infty} \bar{\pi}_{\{i\}}(n) = \bar{\pi}_{\{i\}}$. Similarly, one obtains $\lim_{n\to\infty} \underline{\pi}_{\{i\}}(n) = \underline{\pi}_{\{i\}}$. These two inequalities, together with (13) and (14), deliver the claim.

7. Comments and extensions

In this paper we analyzed continuous-time games of timing with complete information. Even though in general ε-equilibria may fail to exist, in several classes of economic interest we proved the existence of a subgame-perfect ε-equilibrium.

One feature of games in continuous time that was critically used is that in these games a player can mask the time in which he acts. That is, a player who wishes to act at a certain time t_0, can instead act at a random time t close to t_0, thereby concealing the exact moment in which he acts. This way the player can guarantee that no other player will act at the very same moment he does. Since payoffs are continuous, this concealment is not too costly. This idea was used to construct explicit ε-equilibria.

The failure of existence of equilibrium in games with more than two players seems to be related to the non-convergence of the fictitious play dynamics in 3×3 two-player games (see [27]) and to the existence of cyclic equilibrium in undiscounted stochastic games (see [9]). In our model, as in [9], there are three players that are ordered on a circle; each player prefers the player to his right to act, while he does not want the player to his left to act. Since time is continuous, this structure leads to the problem of "choosing the smallest positive real number": If the game is not terminated at time 0, each player would rather act before his opponents act. In [9], time is discrete, so that even though there are no symmetric equilibria, there is a cyclic one. In [27] game the action i of each player is a best reply to action $i+1 \bmod 3$ of his opponent, and this structure leads to the non-convergence of the fictitious play dynamics.

We conclude by discussing which insights can be gained for the analysis of discrete time games with short time periods, and some extensions of our results.

Let $\widehat{\sigma}$ be a subgame-perfect ε-equilibrium of a continuous-time game of timing. Consider a discrete-time version of the game, in which the players are allowed to stop only at times t_n, $n \in \mathbf{N}$, where $(t_n)_n$ is a strictly increasing sequence in $\mathbf{R}^+$. We denote by $\widehat{\tau}$ the discretized version of $\widehat{\sigma}$, defined as follows: at time t_n, assuming no player acted before, player i acts with probability $\sigma^i_{t_{n-1}}((t_{n-1}, t_n])$ (and acts with probability $\sigma^i_0(\{0\})$ at time zero, if $t_0 = 0$). In words, at t_n, player i assigns to *act* the probability with which he would have acted between t_{n-1} and t_n, had he been allowed to act at any time. Assuming all functions u_S are continuous, it is easy to check that $\widehat{\tau}$ is, say, a subgame-perfect 2ε-equilibrium of the game in discrete time, provided $\sup_n |t_n - t_{n-1}|$ is small enough. Moreover, this result does not rely on the sequence (t_n) being known *ex ante*. Specifically, assume that the sequence (t_n) is a random sequence that increases a.s. to ∞, and assume that players get to know the value of t_n at time t_n only.[14] Since the probability to act at time t_n is computed *ex post*, as a function of the interval $(t_{n-1}, t_n]$, the profile $\widehat{\tau}$ is well-defined. Moreover, it is a subgame-perfect 2ε-equilibrium provided that, with high probability, $\sup_n |t_n - t_{n-1}|$ is small enough. Thus, our analysis of the continuous-time game gives an easy scheme for constructing approximate equilibria in a large class of discrete time scenarios.

Finally, we discuss weakenings of the complete information assumption. Our approach does not extend to games with asymmetric information. Nevertheless, it yields partial results in the case of games with symmetric incomplete information. In these games, u_S is a stochastic process, for every $S \subseteq I$, whose law is publicly known. At any time, all the players have the same information on the realization of the payoff processes.[15] These games were first introduced by Dynkin [6] in a two-player zero-sum discrete-time setting. Since then, they have come to be known as Dynkin games in the theory of stochastic processes, and a very extensive literature has been devoted to the zero-sum case, see, e.g., [30] and the references therein.

Acknowledgments

We thank Geir Asheim, Elchanan Ben-Porath, Itzhak Gilboa and an anonymous referee for useful comments.

References

[1] J. Bergin, A model of strategic behavior in repeated games, J. Math. Econ. 21 (1992) 113–153.
[2] J. Bergin, W.B. MacLeod, Continuous time repeated games, Int. Econ. Rev. 34 (1993) 21–37.

[14] We need not assume that the players have any prior information on the law of the sequence (t_n).

[15] For example, they may know past and present values of u_S, $S \subseteq I$, and therefore learn the paths $u_S(\cdot)$, for $\emptyset \neq S \subseteq I$, as time unfolds.

[3] P. Billingsley, Probability and Measure, Wiley, New York, 1995.
[4] M. Bilodeau, A. Slivinski, Toilet cleaning and department chairing: volunteering a public service, J. Public Econ. 59 (1996) 299–308.
[5] J. Bulow, P. Klemperer, The generalized war of attrition, Amer. Econ. Rev. 89 (1999) 175–189.
[6] E.B. Dynkin, Game variant of a problem on optimal stopping, Soviet Math. Dokl. 10 (1969) 270–274.
[7] C.H. Fine, L. Li, Equilibrium exit in stochastically declining industries, Games Econ. Behav. 1 (1989) 40–59.
[8] A.M. Fink, Equilibrium in a stochastic n-person game, J. Sci. Hiroshima Univ. 28 (1964) 89–93.
[9] J. Flesch, F. Thuijsman, K. Vrieze, Cyclic Markov equilibria in stochastic games, Int. J. Game Theory 26 (1997) 303–314.
[10] D. Fudenberg, D.K. Levine, Limit games and limit equilibria, J. Econ. Theory 38 (1986) 261–279.
[11] D. Fudenberg, J. Tirole, Preemption and rent equalization in the adoption of new technology, Rev. Econ. Stud. LII (1985) 383–401.
[12] D. Fudenberg, J. Tirole, A theory of exit in duopoly, Econometrica 54 (1986) 943–960.
[13] D. Fudenberg, J. Tirole, Game Theory, The MIT Press, Cambridge, 1991.
[14] P. Ghemawat, B. Nalebuff, Exit, RAND J. Econ. 16 (1985) 184–194.
[15] K. Hendricks, A. Weiss, C. Wilson, The war of attrition in continuous time with complete information, Int. Econ. Rev. 29 (1988) 663–680.
[16] S. Karlin, Mathematical Methods and Theory in Games, Programming and Economics, Vol. 2, Addison-Wesley, Reading, MA, 1959.
[17] C. LaCasse, C. Ponsatí, V. Barham, Chores, Games Econ. Behav. 39 (2002) 237–281.
[18] E. Maskin, J. Tirole, Markov perfect equilibrium I. Observable actions, J. Econ. Theory 100 (2001) 191–219.
[19] J. Maynard-Smith, The theory of games and the evolution of animal conflicts, J. Theory Biol. 47 (1974) 209–221.
[20] R.B. Myerson, Multistage games with communication, Econometrica 54 (1986) 323–358.
[21] M.J. Osborne, A. Rubinstein, A Course in Game Theory, MIT Press, Cambridge, 1999.
[22] M. Perry, P.J. Reny, A non-cooperative bargaining model with strategically timed offers, J. Econ. Theory 59 (1993) 50–77.
[23] M. Perry P.J. Reny, A non-cooperative view of coalition formation and the core, Econometrica 62 (1994) 795–817.
[24] T. Radzik, T.E.S. Raghavan, Duels, in: R.J. Aumann, S. Hart (Eds.), Handbook of Game Theory with Economic Applications, Vol. 2, 1994, pp. 761–768.
[25] P.J. Reny, On the existence of pure and mixed strategy Nash equilibria in discontinuous games, Econometrica 67 (1999) 1029–1056.
[26] D. Revuz, M. Yor, Continuous Martingales and Brownian Motion, Springer, Berlin, 2000.
[27] L.S. Shapley, Some topics in two-person games, in: Advances in Game Theory, Annals of Mathematical Studies, Princeton University Press, Princeton, NJ, 5 (1964), 1–28.
[28] L.K. Simon, M.B. Stinchcombe, Extensive form games in continuous time: pure strategies, Econometrica 57 (1989) 1171–1214.
[29] E. Solan, Three-person absorbing games, Math. Oper. Res. 24 (1999) 669–698.
[30] E. Solan, N. Vieille, Stopping games—recent results, Ann. ISDG, to appear.
[31] M.B. Stinchcombe, Maximal strategy sets for continuous-time game theory, J. Econ. Theory 56 (1992) 235–265.