

Repeated games with public uncertain duration process

Abraham Neyman · Sylvain Sorin

Accepted: 27 October 2009 / Published online: 13 November 2009
© Springer-Verlag 2009

Abstract We consider repeated games where the number of repetitions θ is unknown. The information about the uncertain duration can change during the play of the game. This is described by an uncertain duration process Θ that defines the probability law of the signals that players receive at each stage about the duration. To each repeated game Γ and uncertain duration process Θ is associated the Θ -repeated game Γ_Θ . A public uncertain duration process is one where the uncertainty about the duration is the same for all players. We establish a recursive formula for the value V_Θ of a repeated two-person zero-sum game Γ_Θ with a public uncertain duration process Θ . We study asymptotic properties of the normalized value $v_\Theta = V_\Theta/E(\theta)$ as the expected duration $E(\theta)$ goes to infinity. We extend and unify several asymptotic results on the existence of $\lim v_n$ and $\lim v_\lambda$ and their equality to $\lim v_\Theta$. This analysis applies in particular to stochastic games and repeated games of incomplete information.

Keywords Repeated games · Uncertain duration · Recursive formula · Asymptotic analysis · Stochastic games · Incomplete information

A. Neyman
Institute of Mathematics, Center for the Study of Rationality,
The Hebrew University of Jerusalem, Givat Ram, 91904 Jerusalem, Israel
e-mail: aneyman@huji.ac.il

S. Sorin (✉)
Equipe Combinatoire and Optimisation, CNRS FRE 3232, Faculté de Mathématiques,
Université P. et M. Curie-Paris 6, 175 rue du Chevaleret, 75013 Paris, France
e-mail: sorin@math.jussieu.fr

S. Sorin
Laboratoire d'Econométrie, Ecole Polytechnique, Paris, France

1 Introduction

We consider repeated games with an uncertain number of stages. For simplicity, we describe in detail the case of two-person repeated games with symmetric uncertainty about the number of stages. (The extension to general n -person games and/or asymmetric uncertainty about the number of stages is straightforward.) The model consists of two basic components:

(a) First, a repeated game Γ is given and described as follows. M is a state space on which a family of normal form two-person games is defined by move spaces I and J for Player 1 and Player 2 respectively, and a payoff function $g = (g^1, g^2)$ from $M \times I \times J \times \mathbb{R}^2$. To simplify the description of the model, we assume here that the sets I , J , and M are finite. (Suitable topological and measurability conditions are needed in the infinite case.)

The initial state m_1 is chosen at random and the players receive some information about it, say a_1 (resp. b_1) for Player 1 (resp. Player 2). The choice of the triple (m_1, a_1, b_1) is performed according to some probability π on $M \times A \times B$, where A and B are signal sets. In addition, after each stage the players obtain some further information about the previous choice of moves and about both the previous and the new state. This is represented by a map Q from $M \times I \times J$ to probabilities on $M \times A \times B$. At stage t , given the state m_t and the moves (i_t, j_t) , a triple $(m_{t+1}, a_{t+1}, b_{t+1})$ is chosen at random according to the distribution $Q(m_t, i_t, j_t)$. The new state is m_{t+1} , and the signal a_{t+1} (resp. b_{t+1}) is transmitted to Player 1 (resp. Player 2). A play of the game is thus a sequence $m_1, a_1, b_1, i_1, j_1, m_2, a_2, b_2, i_2, j_2, \dots$, while the information of Player 1 before his play at stage t is a private history of the form $(a_1, i_1, a_2, i_2, \dots, a_t)$, and similarly for Player 2. The associated sequence of stage payoffs is g_1, g_2, \dots , with $g_t = g(m_t, i_t, j_t) \in \mathbb{R}^2$. Note that a play of the game consists of a sequence of states, signals, and moves. The repeated game Γ is thus represented by the tuple $\langle M, I, J, g, \pi, Q, A, B \rangle$ and this description is public knowledge.

A behavioral (resp. pure) strategy σ for Player 1 in Γ is a map from private histories in Γ to probabilities on the set I of moves, denoted $\Delta(I)$ (resp. to I). A strategy τ of Player 2 is defined similarly. The corresponding sets of behavioral strategies are denoted Σ and \mathcal{T} .

The space of all private histories $h = (a_1, i_1, \dots, i_{t-1}, a_t)$ of Player 1 before the play at stage t is denoted H_t^1 , and $H^1 = \cup_{t \geq 1} H_t^1$ is the set of all private histories of Player 1. The set of private histories H^2 of Player 2 is defined similarly. For finite sets M, I, J, A, B , the sets of finite histories H^1 and H^2 are countable and the sets of pure (resp. behavioral) strategies of Player 1 and of Player 2 are the compact (in the product topology) Cartesian product spaces I^{H^1} (resp. $(\Delta(I))^{H^1}$) and J^{H^2} (resp. $(\Delta(J))^{H^2}$).

Nature's choices in the repeated game Γ (with finite sets M, I, J, A, B) are represented by a list X of independent random variables: X_0 is an $M \times A \times B$ -valued random variable with distribution π (that selects the initial triple (m_1, a_1, b_1)), and, for each $t \geq 1$ and $(m, i, j) \in M \times I \times J$, $X_{t,m,i,j}$ is an $M \times A \times B$ -valued random variable with distribution $Q(m, i, j)$ (that selects the triple $m_{t+1}, a_{t+1}, b_{t+1}$ when $m_t = m$, $i_t = i$, and $j_t = j$). Nature's choices X can be viewed as a random variable with values in the infinite Cartesian product $(Q \times A \times B) \times \prod_{t \geq 1, m \in M, i \in I, j \in J} (M \times A \times B)$.

A pair of pure strategies, σ of Player 1 and τ of Player 2, together with the realization X of Nature's choices, define inductively the play of the repeated games Γ as follows: $(m_1, a_1, b_1) = X_0, i_t = \sigma(a_1, i_1, \dots, a_t), j_t = \tau(b_1, j_1, \dots, b_t)$, and $(m_{t+1}, a_{t+1}, b_{t+1}) = X_{t, m_t, i_t, j_t}$.

Two special cases of repeated games that have been extensively considered are stochastic games and repeated games with incomplete information.

We describe a stochastic game in the simplest framework of standard signaling (or perfect monitoring). The initial signal to the players is the state, namely, $a_1 = b_1 = m_1$, and at each subsequent stage the signal to both players is the previous pair of moves and the new state: $a_{t+1} = b_{t+1} = \{i_t, j_t, m_{t+1}\}$. Hence, formally $A = B = M \cup (I \times J \times M)$. It follows that a play can be identified with a sequence $m_1, i_1, j_1, m_2, i_2, j_2, \dots$, and the information of each player before his play at stage t is the sequence of states and moves $m_1, i_1, j_1, \dots, i_{t-1}, j_{t-1}, m_t$. In addition, since the initial state is publicly known, the analysis is usually done conditional on m_1 ; hence the initial probability π is replaced by a Dirac mass at some point in M .

As for the game with lack of information on both sides (in the so-called dependent case) the traditional description is as follows. To each m in M corresponds a two-person $I \times J$ game G^m . Nature chooses $m \in M$ according to a publicly known probability distribution p on M . Each player gets partial information regarding the actual state $m \in M$; Player 1 (resp. Player 2) observes the realization of a deterministic signal $\ell^1(m)$ (resp. $\ell^2(m)$). Equivalently, $M^1 = \{M_1^1, \dots, M_C^1\}$ and $M^2 = \{M_1^2, \dots, M_D^2\}$ are two partitions of the set M , and following the choice of $m \in M$, Player 1 is informed of c and Player 2 is informed of d where $m \in M_c^1 \cap M_d^2$. The state m is chosen once and for all according to p , and the game G^m is played repeatedly, with standard signaling (perfect monitoring). Using the general presentation above, the probability π is the one induced by p on $M \times C \times D$ and $g(m, i, j) = G_{ij}^m$. Finally, $Q(m_t, i_t, j_t)$ is the unit mass on (m_t, i_t, j_t) and $a_{t+1} = b_{t+1} = (i_t, j_t)$; hence formally $A = C \cup (I \times J)$ and $B = D \cup (I \times J)$.

One basic distinction between the two classes is that in stochastic games the information of the players on the state space is at each stage identical while it is asymmetric in repeated games with incomplete information. More generally, we refer to any repeated game with identical information about the state as a stochastic game.

(b) The second component of the model is the number of repetitions θ , unknown to the players. θ is an integer-valued random variable defined on a probability space $(\Omega, \mathcal{B}, \mu)$ with finite expectation $E(\theta)$. The players receive partial information about the value of θ via a sequence of public signals $s_0, s_1, \dots, s_t, \dots$. Each signal s_t is a measurable function defined on the probability space $(\Omega, \mathcal{B}, \mu)$ and with finite range S . The random variable (s_0, s_1, \dots) is independent of the random variable X that defines Nature's choices in Γ . This defines an uncertain duration process $\Theta = \langle (\Omega, \mathcal{B}, \mu), (s_t)_{t \geq 0}, \theta \rangle$.

To each repeated game Γ and uncertain duration process Θ is associated an extended repeated game Γ_Θ , called the Θ -repeated game. The Θ -repeated game is played essentially like the original game Γ , but, in addition, following the play at stage t , namely (i_t, j_t) , and before the next play at stage $t + 1$, the players receive the public signal $s_t(\omega)$. Formally, a play is now an infinite sequence $s_0, m_1, a_1, b_1, i_1, j_1, s_1, m_2, a_2, b_2, i_2, j_2, s_2, \dots$, while the information of Player 1

corresponds to finite private histories like $(s_0, a_1, i_1, s_1, a_2, i_2, s_2, \dots, a_t)$. Note that the sequence of payoffs g_1, g_2, \dots , is a function of the sequence of states and stage actions, and is thus unaffected (directly) by the process Θ ; but the total payoff in Γ_Θ , $\sum_{t=1}^\infty g_t I(\theta \geq t)(\omega) = \sum_{t=1}^{\theta(\omega)} g_t$, is affected by the process Θ . Note, however, that the payoff on a play is a function of (only) the sequence of states and moves and of the duration. Since the play after stage t on the event $\theta \leq t$ is irrelevant, one can enlarge the signal space so that the signal s_t conveys also the information whether $\theta \leq t$ or not. Thus, we can assume, whenever convenient, that θ is a stopping time w.r.t. the filtration generated by the signals, namely, the increasing sequence of fields $\mathcal{F}_t = \sigma(s_0, \dots, s_t), t = 0, 1, \dots$

A behavioral (resp. pure) strategy σ for Player 1 in Γ_Θ is a map from private histories in Γ_Θ to $\Delta(I)$ (resp. I). A strategy τ is defined similarly for Player 2. The associated sets of behavioral (or mixed) strategies are Σ_Θ and \mathcal{T}_Θ . If M, I, J, A, B are finite sets and each s_t has a countable range (so that measurability requirements are not needed), then the strategy sets Σ_Θ and \mathcal{T}_Θ are compact. A strategy $\sigma \in \Sigma$ is identified with the strategy $\tilde{\sigma} \in \Sigma_\Theta$ that ignores the public signals s_0, s_1, s_2, \dots , i.e., $\tilde{\sigma}(s_0, a_1, i_1, s_1, \dots, a_t) = \sigma(a_1, i_1, \dots, a_t)$. Thus Σ and \mathcal{T} are identified with subsets of Σ_Θ and \mathcal{T}_Θ respectively.

Given a repeated game $\Gamma = \langle M, I, J, g, \pi, Q, A, B \rangle$ and an uncertain duration process Θ , a pair of strategies (σ, τ) in $\Sigma_\Theta \times \mathcal{T}_\Theta$ induces a distribution $P_{\sigma, \tau, \Theta}$ (or $P_{\sigma, \tau, \mu}$ for simplicity) on plays; hence on the sequence of payoffs (g_t) . The (un-normalized) payoff in Γ_Θ is $G_\Theta(\sigma, \tau) = E_{\sigma, \tau, \mu}(\sum_{t=1}^\infty g_t I(\theta \geq t)) \in \mathbb{R}^2$, where $E_{\sigma, \tau, \mu}$ is the expectation w.r.t. distribution $P_{\sigma, \tau, \mu}$.

In the zero-sum case, the value $V_\Theta(\Gamma)$ of the (un-normalized) Θ -repeated game Γ_Θ , with I, J , and M finite (or infinite with the suitable measurability and continuity assumptions) is

$$\begin{aligned} V_\Theta(\Gamma) &= \max_{\sigma \in \Sigma_\Theta} \min_{\tau \in \mathcal{T}_\Theta} G_\Theta(\sigma, \tau) \\ &= \min_{\tau \in \mathcal{T}_\Theta} \max_{\sigma \in \Sigma_\Theta} G_\Theta(\sigma, \tau). \end{aligned}$$

The existence of $V_\Theta(\Gamma)$ follows from the usual minmax theorem: considering mixed strategies, the payoff is bilinear and (even jointly) continuous and both strategy spaces are convex and compact.

The value $v_\Theta(\Gamma)$ of the normalized Θ -repeated game is $v_\Theta(\Gamma) = \frac{V_\Theta(\Gamma)}{E(\theta)}$.

We are interested in the asymptotic behavior of $v_\Theta(\Gamma)$ as the expected duration $E(\theta)$ goes to ∞ .

Remark (1) The above model of public uncertain duration extends naturally to a model of asymmetric uncertain duration with n players. The asymmetric uncertainty is modeled by private signals; i.e., the signal s_t is a profile $s_t = (s_t^1, \dots, s_t^n)$ of signals for each player and the information to, say, Player 1 before the play at stage t is thus the private history $(s_0^1, a_1, i_1, s_1^1, a_2, i_2, s_2^1, \dots, a_t)$. Then a behavioral (resp. pure) strategy σ of Player 1 is a map from such histories to the set of mixed moves $\Delta(I)$ (resp. pure moves I). An additional possible extension is

to the model where the distribution of the duration signal s_t depends also on the play up to stage t .

- (2) The previous extension of a game via a random duration process applies as well to any multistage game with an associated sequence of stage payoffs.
- (3) The case of an asymmetric uncertain duration, as distinct from an asymmetric uncertain duration process, corresponds to the signaling structure where signals about the duration occur only before the “start of the game”, namely, $(s(\omega), I(\theta > t)) = s_t(\omega)$ for every $t \geq 1$; alternatively, if we do not assume that, for every Player i , θ is a stopping time w.r.t. the filtration generated by the duration signals to Player i , it corresponds to $s_t(\omega)$ being independent of t , namely, $s_t(\omega) = s_0(\omega)$.
- (4) The case of uncertain duration with no signals about the duration corresponds to $s_t(\omega) = 1$ if $\theta(\omega) = t$ and $s_t(\omega) = 0$ if $\theta(\omega) \neq t$; alternatively, if we do not assume that θ is a stopping time w.r.t. the filtration generated by the duration signals, it corresponds to $s_t(\omega)$ being a constant that is independent of t and ω . In this case the strategy sets Σ_Θ and \mathcal{T}_Θ are equal to Σ and \mathcal{T} and the evaluation of the payoffs can be written as $\sum_t \rho_t g_t$ with $\rho_t = \mu(\theta \geq t) / E_\mu(\theta)$. The case where θ is deterministic ($\theta = n$) is the classical n -stage repeated game. The payoff is $E_{\sigma, \tau, \mu}(\frac{1}{n} \sum_{t=1}^n g_t)$ and we write v_n for v_Θ . The λ -discounted game corresponds to $\mu(\theta \geq t) = (1 - \lambda)^{t-1}$. Since $E(\theta) = 1/\lambda$ the payoff is $E_{\sigma, \tau, \mu}(\sum_{t=1}^\infty \lambda(1 - \lambda)^{t-1} g_t)$ and we use the notation v_λ for v_Θ .
- (5) Repeated games with asymmetric uncertain duration were studied in [Neyman \(1999\)](#) and [Neyman \(2009b\)](#), demonstrating that asymmetric information about the duration leads to results that differ significantly from those in the public uncertain duration case.

2 Initial results

The next property confirms the robustness of the uniform value.

We first recall the definition. Following [Mertens et al. \(1994, Chap. IV, Sect. 1\)](#) we say that Player 1 can guarantee w in Γ if: for any $\varepsilon > 0$, there exists a strategy σ of Player 1 in Γ and a number of stages T such that, for any strategy τ of Player 2 in Γ and any $t \geq T$,

$$E_{\sigma, \tau} \left(\sum_{\ell=1}^t g_\ell \right) \geq t(w - \varepsilon).$$

Similarly, Player 2 can guarantee w in Γ if: for any $\varepsilon > 0$, there exists a strategy τ of Player 2 in Γ and a number of stages T such that for any strategy σ of Player 1 in Γ and any $t \geq T$,

$$E_{\sigma, \tau} \left(\sum_{\ell=1}^t g_\ell \right) \leq t(w + \varepsilon).$$

The uniform value v_∞ exists if both players can guarantee it.

Theorem 1 *If Player 1 can guarantee w in Γ , then $\liminf_{E(\theta) \rightarrow \infty} v_\Theta(\Gamma) \geq w$. Moreover, $\forall \varepsilon > 0 \exists (\sigma, T) \in \Sigma \times \mathbb{R}$ such that for every uncertain duration process Θ with $E(\theta) \geq T$ and $\forall \tau \in \mathcal{T}_\Theta$ we have $g_\Theta(\sigma, \tau) \geq w - \varepsilon$. In particular, if the infinite game Γ has a uniform value v_∞ , then*

$$\lim_{E(\theta) \rightarrow \infty} v_\Theta(\Gamma) = v_\infty$$

and, moreover, $\forall \varepsilon > 0 \exists (\sigma^*, \tau^*, T) \in \Sigma \times \mathcal{T} \times \mathbb{R}$ s.t. $\forall (\sigma, \tau) \in \Sigma_\Theta \times \mathcal{T}_\Theta$ and $\forall \Theta$ with $E(\theta) > T$ we have $g_\Theta(\sigma, \tau^*) - \varepsilon \leq v_\infty \leq g_\Theta(\sigma^*, \tau) + \varepsilon$.

Proof Assume that Player 1 can guarantee w in Γ . Given $\varepsilon > 0$, let σ be a strategy of Player 1 in Γ and let T be a positive integer so that for every strategy τ of Player 2 in Γ and every $t \geq T$,

$$E_{\sigma, \tau} \left(\sum_{\ell=1}^t g_\ell \right) \geq t(w - \varepsilon).$$

Given a pure strategy τ of Player 2 in Γ_Θ and ω in Ω we denote by τ_ω the strategy of Player 2 in Γ given by

$$\tau_\omega(b_1, j_1, b_2, \dots, b_t) = \tau(s_0(\omega), b_1, j_1, s_1(\omega), b_2, j_2, s_2(\omega), \dots, b_t).$$

It follows that if τ is a pure strategy of Player 2 in Γ_Θ and if ω in Ω satisfies $\theta(\omega) \geq T$, then

$$E_{\sigma, \tau_\omega} \left(\sum_{\ell=1}^{\theta(\omega)} g_\ell \right) \geq \theta(\omega)(w - \varepsilon)$$

and therefore for every $\omega \in \Omega$

$$E_{\sigma, \tau_\omega} \left(\sum_{\ell=1}^{\theta(\omega)} g_\ell \right) \geq \theta(\omega)(w - \varepsilon) - \|g\|T$$

where $\|g\| = \sup_{M \times I \times J} |g(i, j, m)|$. As $E_{\sigma, \tau, \mu} \left(\sum_{\ell=1}^{\theta} g_\ell \right) = E_\mu \left(E_{\sigma, \tau_\omega} \left(\sum_{\ell=1}^{\theta(\omega)} g_\ell \right) \right)$, we deduce that for any strategy τ of Player 2 in Γ_Θ

$$E_{\sigma, \tau, \mu} \left(\sum_{t=1}^{\theta} g_t \right) \geq E(\theta)(w - \varepsilon) - \|g\|T.$$

□

Remark (1) Note that the strategy σ of the pair (σ, T) in the “moreover” part of the theorem is in Σ . Therefore, the theorem holds also for asymmetric uncertain duration processes.

- (2) The proof of Theorem 1 applies also to n -person games with either symmetric or asymmetric uncertain duration processes. The applications are both for the uniform maxmin (and the uniform minmax) and for uniform equilibrium.

We start with the comments applicable to the uniform maxmin and to the uniform minmax. In an n -person repeated game Γ Player i can guarantee w^i if: for every $\varepsilon > 0$, there exists a strategy σ^i of Player i in Γ and a number of stages T such that for any strategy profile σ^{-i} of the other players in Γ and any $t \geq T$, $E_{\sigma^i, \sigma^{-i}} \sum_{\ell=1}^t g_\ell^i \geq t(w^i - \varepsilon)$. The uniform maxmin of Player i in the repeated game Γ exists and equals w^i if Player i can guarantee w^i in Γ and for every strategy σ^i of Player i in Γ there is a strategy profile σ^{-i} of the other players in Γ and a positive integer T such that for any $t \geq T$, $E_{\sigma^i, \sigma^{-i}} (\sum_{\ell=1}^t g_\ell) \leq t(w^i + \varepsilon)$. The maxmin of Player i in Γ_Θ is $v_\Theta^i(\Gamma) := \max_{\sigma^i \in \Sigma_\Theta^i} \min_{\sigma^{-i} \in \times_{j \neq i} \Sigma_\Theta^j} \frac{1}{E(\theta)} E_{\sigma^i, \sigma^{-i}} \left(\sum_{\ell=1}^\theta g_\ell^i \right)$.

Similarly, Player i can protect w^i , if for every strategy profile σ^{-i} of the other players in Γ and every $\varepsilon > 0$ there is a strategy σ^i of Player i and a positive integer T such that for any $t \geq T$, $E_{\sigma^i, \sigma^{-i}} \sum_{\ell=1}^t g_\ell^i \geq t(w^i - \varepsilon)$. The uniform minmax of Player i in the repeated game Γ exists and equals w^i if Player i can protect w^i in Γ and for every $\varepsilon > 0$ there is a strategy profile σ^{-i} of the other players in Γ and a positive integer T such that for any $t \geq T$ and every strategy σ^i of Player i , $E_{\sigma^i, \sigma^{-i}} (\sum_{\ell=1}^t g_\ell) \leq t(w^i + \varepsilon)$. The minmax of Player i in Γ_Θ is $\bar{v}_\Theta^i(\Gamma) := \min_{\sigma^{-i} \in \times_{j \neq i} \Sigma_\Theta^j} \max_{\sigma^i \in \Sigma_\Theta^i} \frac{1}{E(\theta)} E_{\sigma^i, \sigma^{-i}} \left(\sum_{\ell=1}^\theta g_\ell^i \right)$.

The proof of Theorem 1 implies that if Player i can guarantee (resp. protect) w^i in Γ , then $\liminf_{E(\theta) \rightarrow \infty} v_\Theta^i(\Gamma) \geq w^i$ (resp. $\liminf_{E(\theta) \rightarrow \infty} \bar{v}_\Theta^i(\Gamma) \geq w^i$), and if the infinite game Γ has a uniform maxmin (resp. minmax) v_∞^i (resp. \bar{v}_∞^i), then $\lim_{E(\theta) \rightarrow \infty} v_\Theta^i(\Gamma) = v_\infty^i$ (resp. $\lim_{E(\theta) \rightarrow \infty} \bar{v}_\Theta^i(\Gamma) = \bar{v}_\infty^i$).

We follow with comments regarding uniform equilibrium. A strategy profile $\sigma = (\sigma^1, \sigma^2, \dots)$ is an ε -uniform equilibrium in the repeated game Γ if there is a positive integer T such that for every $t \geq T$, every Player i , and every strategy τ^i of Player i , we have $E_\sigma (\sum_{\ell=1}^t g_\ell^i) + t\varepsilon \geq E_{\sigma^{-i}, \tau^i} (\sum_{\ell=1}^t g_\ell^i)$. An ε -equilibrium of Γ_Θ is a strategy profile σ such that for every Player i and every strategy τ^i of Player i we have $E_\sigma \left(\sum_{\ell=1}^\theta g_\ell^i \right) + E(\theta)\varepsilon \geq E_{\sigma^{-i}, \tau^i} \left(\sum_{\ell=1}^\theta g_\ell^i \right)$.

The proof of Theorem 1 implies that if σ is an ε -uniform equilibrium in the repeated game Γ , then for every $\varepsilon' > \varepsilon$ there is T such that σ is an ε' -equilibrium of Γ_Θ for every uncertain duration process Θ with $E(\theta) \geq T$.

- (3) Two-person zero-sum stochastic games (resp. n -person stochastic games) with finitely many states and actions and with perfect monitoring, and absorbing games with compact action spaces and perfect monitoring, have a value (Mertens and Neyman 1981; Mertens et al. 2009) (resp. a maxmin and minmax (Neyman 2003a)) and thus, in particular, a uniform value (resp. a uniform maxmin and minmax). However, without perfect monitoring these classes of two-person zero-sum stochastic games do not have a uniform value. For example, the big match does not have a uniform value when players do not observe the actions of the other players. Therefore Theorem 1 does not imply directly that the limit of $v_\Theta(\Gamma)$, as

$E(\theta) \rightarrow \infty$, exists in these repeated games without perfect monitoring and is independent of the signals of players' moves. This limiting result holds and it will follow from Theorem 1 in conjunction with the recursive formula developed in the next section.

3 The extended recursive structure

3.1 Recursive structure

Shapley (1953) associates to a two-person zero-sum stochastic game the Shapley operator Ψ that maps real-valued functions defined on the state space M to themselves:

$$\Psi(f)[m] = \sup_{x \in \Delta(I)} \inf_{y \in \Delta(J)} \{g(m, x, y) + E_{m,x,y} [f(m')]\} \quad (1)$$

where $g(m, x, y)$ is the bilinear extension of $g(m, \cdot, \cdot)$ to $\Delta(I) \times \Delta(J)$ and the expectation $E_{m,x,y}$ is with respect to the law of m' given by $Q(m, x, y)$: the bilinear extension to $\Delta(I) \times \Delta(J)$ of the transition $Q(m, \cdot, \cdot)$. $\Psi(f)[m]$ is interpreted as the value of a one-stage game played as the one-shot stochastic game with the payoff function being the sum of the stage payoff of the stochastic game and the value of the function f at the new state. The iterates of the operator Ψ evaluated at 0 express the values of the finitely repeated stochastic game:

$$\Psi^n(0) = V_n. \quad (2)$$

A similar operator can be introduced in the zero-sum case, for any repeated game considered in Sect. 1 (Mertens et al. 1994, Chap. IV, Sect. 3). In contrast to the case of stochastic games where the operator acts on functions defined on the state space of the game, the general case involves operators on functions defined on an auxiliary enlarged state space M' . In fact, to each repeated game $\Gamma = \langle M, I, J, g, \pi, Q, A, B \rangle$, one associates an auxiliary stochastic game $\Gamma' = \langle M', I', J', g', \pi', Q' \rangle$ having the same n -stage (and λ -discounted) value. A formula like (1) defines the corresponding Shapley operator and the n -stage value is expressed by the n -th iterate of the operator evaluated at 0.

The purpose of this section is to extend formulas (1) and (2) for repeated games (hence in particular also for stochastic games) with an uncertain duration process. Given an uncertain duration process Θ , equality of the type $V_\Theta(\Gamma) = \Psi^\Theta(0)$ will be obtained where Ψ is the extended Shapley operator of the auxiliary game. The generalized iterate Ψ^Θ will in fact be defined (in Sect. 3.2, following Neyman (2003b)) for any nonexpansive mapping Ψ .

Next we illustrate extended Shapley operators Ψ associated with several specific classes of repeated games:

(1) Repeated games with a publicly known state.

This is the class where, at every stage t , each private signal a_t or b_t contains at least the current state m_t . Since in this framework the state is public knowledge,

the Shapley operator (1) introduced in the perfect monitoring case and acting on functions defined on M itself is the auxiliary Shapley operator.

- (2) Repeated games with publicly known probability distribution on the state space. This is the class with public signals and perfect monitoring, i.e., where at every stage t , a_t equals b_t and contains at least (i_t, j_t) , namely, the pair of moves. The auxiliary stochastic game Γ' can be chosen as follows: the state space $M' = \Delta(M)$ is the space of probability distributions on M ; $I' = I, J' = J, g'(\cdot, i, j)$ is the linear extension of $g(\cdot, i, j)$ to M' . Finally, for each $\eta \in M', Q'(\eta, i, j)$ is the probability on M' defined as follows: let $\rho(\eta, i, j)(a) = E_\eta(Q(m, i, j)(a))$ be the probability of signal a induced in Γ by (η, i, j) . Each signal a determines, given (η, i, j) , a conditional probability distribution $\eta(a)$ on the state space M , hence a point in M' . Then $Q'(\eta, i, j)[\eta'] = \sum_{a:\eta(a)=\eta'} \rho(\eta, i, j)(a)$. The resulting auxiliary Shapley operator is:

$$\Psi(f)[\eta] = \sup_{x \in \Delta(I)} \inf_{y \in \Delta(J)} \{g'(\eta, x, y) + E_{\eta,x,y}(f(\eta'))\}$$

where $g'(\eta, x, y) := \sum_{i,j} x(i)y(j)g'(\eta, i, j)$ is the multi-linear extension of $g'(\eta, i, j)$ and $E_{x,y,\eta}(f(\eta')) := \sum_{i,j,\eta'} x(i)y(j)Q'(\eta, i, j)[\eta'](f(\eta'))$.

- (3) Repeated games with incomplete information: the independent case with perfect monitoring (Aumann and Maschler 1995).

M is a product space $K \times L, \pi$ is a product probability $p \otimes q$ with $p \in \Delta(K), q \in \Delta(L)$, and, in addition, $a_1 = k$ and $b_1 = \ell$. The state $m = (k, \ell)$ corresponds to the type of the players and each player knows his own type and holds a prior on the other player's type. From stage 1 on, the state is fixed and the information of the players is $a_{t+1} = b_{t+1} = \{i_t, j_t\}$.

The auxiliary stochastic game Γ' can be chosen as follows: the state space M' is $\Delta(K) \times \Delta(L)$ and is interpreted as the space of beliefs on the true state; $I' = \Delta(I)^K$ and $J' = \Delta(J)^L$ correspond to type-dependent mixed moves of the players; g' is defined on $M' \times I' \times J'$ by $g'(p, q, x, y) = \sum_{k,\ell} p^k q^\ell g(k, \ell, x^k, y^\ell)$. The transition Q' is introduced next: given (p, q, x, y) , let $x(i) = \sum_k x_i^k p^k$ and $p(i)$ be the conditional probability on K given the move i , that is, $p^k(i) = \frac{p^k x_i^k}{x(i)}$ (and similarly for y and q). Then $Q'(p, q, x, y)(p', q') = \sum_{i,j:(p(i),q(j))=(p',q')} x(i)y(j)$.

The resulting auxiliary Shapley operator is:

$$\begin{aligned} \Psi(f)[p, q] &= \sup_{x \in I'} \inf_{y \in J'} \left\{ \sum_{k,\ell} p^k q^\ell g(k, \ell, x^k, y^\ell) + \sum_{i,j} x(i)y(j)f(p(i), q(j)) \right\} \\ &= \sup_{x \in I'} \inf_{y \in J'} \{g'(p, q, x, y) + E_{p,q,x,y}[f(p', q')]\}. \end{aligned}$$

- (4) Repeated games where Player 1 knows more than Player 2. This is the case where a_t determines b_t and j_t and w.l.o.g. determines also i_t (we always assume perfect recall). The auxiliary stochastic game Γ' can be chosen as follows: the state space M' is $\Delta(\Delta(M))$. The set $S = \Delta(M)$ describes the

possible beliefs of Player 1 on the state space M and the set $M' = \Delta(S)$ stands for the beliefs of Player 2 about Player 1's beliefs. $I' = \Delta(I)^S$, $J' = \Delta(J)$, and for $\eta \in M'$, $x \in I'$ and $y \in J'$, the payoff function is $g'(\eta, x, y) = \int_S [\sum_j y_j \int_M \sum_i g(m, i, j) x_i^s s(dm)] \eta(ds)$. It remains to specify $Q'(\eta, x, y)$. A signal a of Player 1 defines (through Q) a posterior probability on M , namely, a point in S . Given a signal b and a move j of Player 2, Player 2 can compute the conditional distribution on the signals a of Player 1, and hence a point in M' . Finally, (η, x, y) and Q determine the law of the signals b . The resulting auxiliary Shapley operator is

$$\Psi(f)[\eta] = \sup_{x \in \Delta(I)^S} \inf_{y \in \Delta(J)} \{g'(\eta, x, y) + E_{\eta, x, y}(f(\eta'))\}.$$

Note that the above procedure aims at building from a repeated game Γ , another repeated game Γ' , which is in fact a stochastic game, such that the n -stage and λ -discounted values of both games coincide. However, there is in general no direct relation between strategies in both games; in particular optimal strategies in Γ' need not have counterpart optimal strategies in the original repeated game Γ . In examples 1 and 2 above, the state space in the auxiliary stochastic game Γ' is public knowledge in the original game Γ , which enables a natural map from strategies in Γ' to strategies in Γ that preserves optimality. In example 3, the auxiliary state space corresponds to beliefs of the players, which can be computed on the basis of their type-dependent mixed moves, but which are not observable during the play of the original game Γ . Finally, in example 4, the auxiliary state space that models the probabilistic information of Player 2 relies again on the knowledge of Player 1's strategy that Player 2 does not observe in the play of Γ . However, the better-informed Player 1 can compute it, and hence optimal strategies for Player 1 in Γ' translate also to optimal strategies in Γ .

We turn to the construction of the auxiliary Shapley operator for an arbitrary repeated game Γ as defined in Sect. 1. It relies on the recursive structure based on the representation of an information scheme, namely, a probability distribution on $M \times A \times B$, where A and B are arbitrary signal sets for Players 1 and 2 respectively. (We follow [Mertens et al. \(1994, Chap. IV, Sect. 3\)](#).)

Given M , there exists a space W with the following properties ([Mertens and Zamir 1985](#)):

- (1) There exists a homeomorphism φ from W to $\Delta(M \times W)$.
- (2) If W^i denotes a copy of W , any information scheme has a canonical representation as a probability P on $\mathcal{U} = M \times W^1 \times W^2$ such that, for $\{i, j\} = \{1, 2\}$, the conditional probability $P(\cdot | w^i)$ on $M \times W^j$ coincides with $\varphi(w^i)$, P almost surely.

The set of such probabilities (called *consistent*) is denoted \mathcal{P} . The set W^i is called the *type set* of player i and \mathcal{U} is called the *universal belief space*. Given a consistent probability P , a type w^i in W^i , which corresponds to the signal to Player i , is thus identified with its image $\varphi(w^i)$, which is a distribution on $M \times W^j$ ($j \neq i$), namely, on states and types of the opponent. It coincides with the beliefs he computes, given his type.

The game $\Gamma = \langle M, I, J, g, \pi, Q, A, B \rangle$ is value-equivalent to $\Gamma' = \langle M', I', J', g', \pi', Q' \rangle$ where $M' = \mathcal{P}$, the signal space to player i would be W^i , and the corresponding distribution on \mathcal{U} would be $P_1 = P$. Similarly, given σ and τ , strategies in Γ , the distribution on plays up to moves at stage t defines a consistent probability P_t on \mathcal{U} , which corresponds to the state of knowledge on M at stage t . The game Γ is value-equivalent to a new game where σ and τ are played for $t - 1$ stages and where the game restarts at stage t with P_t (as a function of the play and σ and τ) as initial distribution on state and signals. The (behavioral) strategies used at stage t in Γ allow us to construct P_{t+1} . They can be replaced, for each player i , by a function of his type w_t^i at that stage that induces the same payoff (Mertens et al. 1994, Chap. III, Proposition 4.5). Hence the play at time t is specified by P_t and maps α_t and β_t from type sets to mixed actions. This triple determines the stage payoff via the marginal distribution on $M \times I \times J$ and P_{t+1} as a function $H(P_t, \alpha_t, \beta_t)$ in $\Delta(M \times W^1 \times W^2)$.

The extension of the Shapley operator (1) is the following operator Ψ defined on bounded functions on \mathcal{P} :

$$\Psi(f)[P] = \sup_{\alpha} \inf_{\beta} \{g'(P, \alpha, \beta) + f[H(P, \alpha, \beta)]\} \tag{3}$$

where α (resp. β) is a map from W^1 to $\Delta(I)$ (resp. W^2 to $\Delta(J)$) and $g'(P, \alpha, \beta)$ denotes the expectation of $g(m, i, j)$ with respect to the probability induced by α, β and P on $M \times I \times J$. Hence the auxiliary repeated game is actually a stochastic game with a deterministic transition defined by the function H .

The previous explicit examples used this property with a reduction of both spaces \mathcal{U} and \mathcal{P} to subspaces \mathcal{U}_0 and \mathcal{P}_0 where the support of each probability P in \mathcal{P}_0 is included in \mathcal{U}_0 .

In example 1 as well as in standard stochastic games satisfying (1), rather than dealing with probability distributions on states, the public knowledge of the state in Γ allows us to choose $M' = M$ in Γ' , but the transition is no longer deterministic. A similar remark applies to example 2, where one works with $\Delta(M)$ rather than $\Delta(\Delta(M))$. Also, in the last two examples public knowledge of the moves or comparison of the information enables us to reduce the level of iterations needed when working with H .

3.2 Uncertain duration process and extended orbits of a nonexpansive operator

Let Ψ be a nonexpansive map from a Banach space X to itself. Following (Neyman 2003b), we generalize here the iterates Ψ^n to operators Ψ^Θ that act on X and are defined for an arbitrary uncertain duration process $\Theta = \langle (\Omega, \mathcal{B}, \mu), (s_t)_{t \geq 0}, \theta \rangle$. Ψ^Θ captures the idea of a “generalized” random number of iterations of the nonexpansive map Ψ . Moreover, when θ is a stopping time w.r.t. the increasing sequence of algebras $\mathcal{F}_t = \sigma(s_0, \dots, s_t)$, the domain of the operator Ψ^Θ extends from X to all \mathcal{F}_θ -measurable functions $x : \Omega \rightarrow X$.

To an uncertain duration process Θ , where θ is a stopping time, we associate a probability tree as follows. The terminal nodes $T = T_\Theta$ are all finite sequences of signals $v = (s_0, \dots, s_t)$ with positive μ probability and $t = \theta(s_0, \dots, s_t)$. The set of nodes $N = N_\Theta$ is the set of all initial segments $(s_0, \dots, s_r), r \leq t$, of a terminal node

(s_0, \dots, s_t) . For every node $v = (s_0, \dots, s_r)$, let $k(v) = \max\{\theta(v') - r; v' \text{ terminal node with initial segment } s_0, \dots, s_r\}$. The root of the tree is the node of the empty string of signals $v = \emptyset$. The probability measure μ on (Ω, \mathcal{B}) induces a probability measure μ_T on the countable set of terminal nodes T . Given a node $v = (s_0, \dots, s_r)$ and an integrable function $f : \Omega \rightarrow \mathbb{R}$ we denote by $E(f \mid v)$ the value of the conditional expectation $E(f \mid \mathcal{F}_r)$ at v .

The next theorem asserts that given an integrable function x on the terminal nodes and a nonexpansive map Ψ there is a uniquely determined Ψ iterate \bar{x} defined on all non-terminal nodes.

Theorem 2 *Given a function $x : T \rightarrow X$ (which is identified with a function $x : \Omega \rightarrow X$, measurable w.r.t. \mathcal{F}_θ) with finite expectation $E_{\mu_T}(\|x\|) = E_\mu(\|x\|) < \infty$, and a nonexpansive map $\Psi : X \rightarrow X$, there is a unique extension of the function x to a function \bar{x} defined on all nodes N such that*

$$\bar{x}(\emptyset) = E(\bar{x}(s_0)), \tag{4}$$

and for every non-terminal node $v = (s_0, \dots, s_r) \neq \emptyset$,

$$\bar{x}(v) = \Psi(E(\bar{x}(v, s_{r+1}) \mid v)) \tag{5}$$

$$\|\bar{x}(v)\| \leq E(\theta - r \mid v)\|\Psi(0)\| + E(\|x\| \mid v). \tag{6}$$

In addition, given two functions $x, y : T \rightarrow X$ with finite expectation, the following inequality holds:

$$\|\bar{x}(v) - \bar{y}(v)\| \leq E(\|x - y\| \mid v). \tag{7}$$

Proof If $\theta = 0$, Eq. 4 defines uniquely $\bar{x}(\emptyset)$ and the inequalities (6) and (7) hold. Assume that $E(\theta) > 0$. We first prove the lemma for the case where θ is bounded, by induction on the number of nodes in N_θ . Let $v \neq \emptyset$ be a node with $k(v) = 1$; equivalently, v is a maximal non-terminal node (i.e., all subsequent nodes are terminal nodes). Given two functions x, y from T_θ to X , Eq. 5 defines $\bar{x}(v)$ and $\bar{y}(v)$ and $\|\bar{x}(v) - \bar{y}(v)\| \leq E(\|y - x\| \mid v)$ by nonexpansiveness. Therefore, by the induction hypothesis, \bar{x} and \bar{y} , which are uniquely defined on all nodes by backward induction using (5), will satisfy for any other non-terminal node v' , $\|\bar{x}(v') - \bar{y}(v')\| \leq E(\|x - y\| \mid v')$, i.e., (7) holds. (6) holds as well by backward induction. In fact, fix a node v with $k(v) = 1$; note that $\bar{x}(v)$ is defined by (5), and replace it by a terminal node. This defines a stopping time θ' with associated process Θ' . Set $y : T_{\Theta'} \rightarrow X$ by $y(v) = \bar{x}(v)$ and $y = x$ at all other terminal nodes in $T_{\Theta'}$. Note that $\|y(v)\| = \|\bar{x}(v)\| \leq \|\Psi(0)\| + E(\|x\| \mid v)$. At any node $v' = (s_0, \dots, s_r)$ in $N_{\Theta'} \setminus T_{\Theta'}$, one has $\bar{x}(v') = \bar{y}(v')$ so that $\|\bar{x}(v')\| = \|\bar{y}(v')\| \leq E(\theta' - r \mid v')\|\Psi(0)\| + E(\|y\| \mid v')$ hence $\|\bar{x}(v')\| \leq E(\theta' - r \mid v')\|\Psi(0)\| + E(\|x\| \mid v') + \text{Prob}(v \mid v')\|\Psi(0)\| = E(\theta - r \mid v')\|\Psi(0)\| + E(\|x\| \mid v')$.

Assume finally that θ is unbounded with finite expectation. Fix a node $v = (s_0, \dots, s_r)$ and a sufficiently large $n \geq r$. Define the stopping time $\theta \wedge n$ and let $\Theta \wedge n = ((\Omega, \mathcal{B}, \mu), (s_t)_{t \geq 0}, \theta \wedge n)$ be the associated uncertain duration process.

Consider y and z , two functions on $T_{\Theta \wedge n}$ that coincide with x on $T_{\Theta \wedge n} \cap T_{\Theta}$ and such that for every $v' \in T_{\Theta \wedge n}$,

$$\|\bar{y}(v')\| + \|\bar{z}(v')\| \leq 2E((\theta - n)^+ | v')\|\Psi(0)\| + 2E(\|x\| | v').$$

It follows that

$$\|\bar{y}(v) - \bar{z}(v)\| \leq E(\|y - z\| | v) \leq 2E((\theta - n)^+ | v)\|\Psi(0)\| + 2E(\|x\| I(\theta > n) | v)$$

and the upper bound goes to 0 as $n \rightarrow \infty$. In particular, if y_n coincides with x on $T_{\Theta \wedge n} \cap T_{\Theta}$ and equals 0 on $T_{\Theta \wedge n} \setminus T_{\Theta}$, $\bar{y}_n(v)$ converges to a limit denoted by $\bar{x}(v)$. The last argument proves existence of the extension and the previous one shows uniqueness. \square

Definition The Θ -iterate of Ψ is defined on the set of \mathcal{F}_{θ} -measurable functions $x : \Omega \rightarrow X$ by

$$\Psi^{\Theta}(x) = \bar{x}(\emptyset).$$

Comments Equations 5, 6, and 7 are in fact true in a more general setting. Let θ' be another \mathcal{F}_t -stopping time and write Θ' for the associated uncertain duration process. Assume $\theta' \leq \theta$ so that $T_{\Theta'} \subset N_{\Theta}$. Define a $\mathcal{F}_{\theta'}$ -measurable function y by $y(v) = \bar{x}(v)$ for v in $T_{\Theta'}$. Then

$$\Psi^{\Theta}(x) = \Psi^{\Theta'}(y). \tag{8}$$

If θ and θ' are two stopping times (w.r.t. $(\mathcal{F}_t)_{t \geq 0}$) with finite expectations, $\Psi^{\Theta}(0) = \Psi^{\Theta \wedge \Theta'}(y)$ where $E(\|y\|) \leq E(\theta - (\theta \wedge \theta'))\|\Psi(0)\|$ and thus $\|\Psi^{\Theta}(0) - \Psi^{\Theta \wedge \Theta'}(0)\| \leq E(\theta - (\theta \wedge \theta'))\|\Psi(0)\|$. Similarly, $\|\Psi^{\Theta'}(0) - \Psi^{\Theta \wedge \Theta'}(0)\| \leq E(\theta' - (\theta \wedge \theta'))\|\Psi(0)\|$. As $|\theta - \theta'| = \theta - (\theta \wedge \theta') + \theta' - (\theta \wedge \theta')$, we have,

$$\|\Psi^{\Theta'}(0) - \Psi^{\Theta}(0)\| \leq E(|\theta' - \theta|)\|\Psi(0)\|.$$

If $v = (s_0, \dots, s_r)$ is a non-terminal node of the uncertain duration process Θ , we denote by $\Theta(v)$ the remaining uncertain duration process after v (in particular, $\theta(v) = \theta - r$). The associated probability tree is thus the sub-tree with root v , endowed with the corresponding conditional probability. If v is a terminal node we identify the identity operator with $\Psi^{\Theta(v)}$. With this notation one has, for every $v \in N_{\Theta}$,

$$\bar{x}(v) = \Psi^{\Theta(v)}(x),$$

and thus, in particular, for any non-terminal node $v = (s_0, \dots, s_r) \neq \emptyset$,

$$\Psi^{\Theta(v)}(\cdot) = \Psi(E(\Psi^{\Theta(v, s_{r+1})}(\cdot)) | v). \tag{9}$$

Note that Eq. 6 is needed for uniqueness only in the case where θ is unbounded. However, uniqueness follows also with a weaker requirement: $\|\bar{x}(v)\| \leq KE(\theta - r |$

$v) \|\Psi(0)\| + KE(\|x\| \mid v)$ for some constant $K \geq 1$. Nevertheless, some bound is needed. Indeed, if θ is unbounded and (s_0, s_1, \dots) is an infinite sequence of signals so that $\mu(s_0, \dots, s_t) > 0$ for every t , for any $z \in X$ we can find a function $\bar{y} : N_\Theta \rightarrow X$ that coincides with our defined \bar{x} on all nodes v that are not initial segments of the infinite sequence (s_0, s_1, \dots) and so that $\bar{y}(\emptyset) = z$ and \bar{y} obeys (5). Indeed, define inductively $\bar{y}(s_0, \dots, s_t)$ so that (5) holds also for $\bar{y}(s_0, \dots, s_{t-1})$.

3.3 Uncertain duration process and extended recursive formula

Here we establish an extended recursive formula for the value of the Θ -repeated game Γ_Θ .

Since the law of θ is independent of the moves and states, one obtains that the value of the game Γ extended by the uncertain duration process Θ satisfies the following extension of (2):

Theorem 3

$$V_\Theta(\Gamma) = \Psi^\Theta(0) \tag{10}$$

where Ψ is given by (3).

Proof As the duration signals are public, following, e.g., Mertens et al. (1994, Chap. IV, Sect. 3), the equality holds for any bounded uncertain duration process Θ . Moreover, as $V_{\Theta \wedge n}$ and $\Psi^{\Theta \wedge n}(0)$ converge to V_Θ and $\Psi^\Theta(0)$ respectively the equalities $V_{\Theta \wedge n} = \Psi^{\Theta \wedge n}(0)$ imply that $V_\Theta = \Psi^\Theta(0)$. \square

Equation 9 implies that for any non-terminal node $v = (s_0, \dots, s_r) \neq \emptyset$,

$$V_{\Theta(v)}(\Gamma) = \Psi(E(V_{\Theta(v, s_{r+1})}(\Gamma) \mid v)) \tag{11}$$

and from (4) one has:

$$V_\Theta(\Gamma) = E(V_{\Theta(s_0)}(\Gamma)).$$

In particular, this recursive formula implies the following property on optimal strategies.

Theorem 4 *Assume that the state variable $P_t \in \mathcal{P}$ is public knowledge. Then each player has an optimal strategy that at each stage t is only a function of P_t and $\Theta(v)$, $v \in \mathcal{F}_t$.*

To be more specific, consider a stochastic game with a publicly known state, as previously defined in Sect. 3.1, example 1. The above result implies that both players have optimal strategies that depend only upon the remaining uncertain duration process $\Theta(v)$ and the current state m_t . Hence the value v_Θ is the same whatever the additional information on moves may be. However, in the case of full monitoring or at least when the signals a_t and b_t allow both players to compute g_{t-1} , Mertens and Neyman (1981) proved the existence of a uniform value. Theorem 1 thus implies:

Corollary 1 *In a stochastic game with a publicly known state, $\lim_{E(\theta) \rightarrow \infty} v_{\Theta}$ exists and is independent of the signals on moves.*

Similarly, [Mertens et al. \(2009\)](#) proved the existence of a uniform value in absorbing games with compact action spaces: these are stochastic games where only one state, say m , is nonabsorbing. (Recall that an absorbing state is a state that once reached cannot be left.) The action spaces are compact sets X and Y . Given (x, y) in $X \times Y$ the game starting from m remains in stage m with probability $q(x, y)$ and the payoff in this event is $g(x, y)$. Otherwise an absorbing state is reached and one can assume that there is an absorbing payoff $\rho(x, y)$. These three functions are continuous on $X \times Y$. Theorem 1 thus implies:

Corollary 2 *In an absorbing game with compact action spaces and with a publicly known state, $\lim_{E(\theta) \rightarrow \infty} v_{\Theta}$ exists and is independent of the signals on moves.*

Note that these results hold for any signals on moves, and hence also in cases where the uniform value may not exist. They rely on the previous recursive formula for v_{Θ} that implies, as mentioned above, that both players have optimal strategies that depend only upon the remaining uncertain duration process $\Theta(v)$ and the current state m_t .

4 Operator approach

In this section we extend previously known inequalities of the values of the n -stage game v_n and the λ -discounted game v_{λ} to corresponding inequalities of the values v_{Θ} of the repeated games with uncertain duration processes.

4.1 Variational bounds

The nonexpansive operators arising in repeated games act on spaces of real-valued functions endowed with the uniform norm and in addition these operators are monotonic. Let Ψ be a monotonic nonexpansive mapping on a convex cone \mathcal{F} of bounded real functions that contains the constants. We first recall a definition from [Sorin \(2004\)](#) that extends [Rosenberg and Sorin \(2001\)](#).

Definition 1 \mathcal{L}^+ is the set of functions f in \mathcal{F} for which there exists $L \in \mathbb{R}$ such that

$$\Psi(Kf) \leq (K + 1)f, \quad \forall K \geq L.$$

Such a function yields an upper bound for the iterates, $\Psi^n(0) \leq nf + 2L\|f\|$, which implies that $\limsup_{n \rightarrow \infty} \frac{\Psi^n(0)}{n} \leq f$ for any $f \in \mathcal{L}^+$; see [Rosenberg and Sorin \(2001\)](#) and [Sorin \(2004\)](#). The next result generalizes this inequality to any uncertain duration process.

Theorem 5 *Assume $f \in \mathcal{L}^+$. For any uncertain duration process Θ ,*

$$\Psi^{\Theta}(0) \leq E(\theta)f + 2L\|f\|.$$

Proof Assume $f \in \mathcal{L}^+$ and let L be the constant associated to f . We prove the stronger property that for any uncertain duration process Θ and every function $K : T_\Theta \rightarrow \mathbb{R}$ with $K \geq L$ and $E(K) < \infty$,

$$\Psi^\Theta(Kf) \leq (E(K) + E(\theta))f. \tag{12}$$

Obviously, (12) holds when $E(\theta) = 0$. The proof of (12) for bounded uncertain duration processes (with $E(\theta) > 0$) is by induction on N_Θ , that is, on the number of nodes of the probability tree associated with Θ . Let Θ' be an uncertain duration process (defined on the same duration signal space $(\Omega, \mu, (s_t)_{t \geq 0})$) with $\theta - 1 \leq \theta' < \theta$. Thus $|N_{\Theta'}| < |N_\Theta|$. Define the function $K' : T_{\Theta'} \rightarrow \mathbb{R}$ by $K'(v) = E(\theta - r \mid v) + E(K \mid v)$ for a terminal node $v = (s_0, \dots, s_r)$ of Θ' . As $f \in \mathcal{L}^+$, we have $\Psi^{\Theta(v)}(Kf) \leq (E(\theta - r \mid v) + E(K \mid v))f = K'(v)f$ by induction. As $K'(v) = E(\theta - r \mid v) + E(K \mid v) \geq L$, the induction hypothesis implies that $\Psi^{\Theta'}(K'f) \leq (E(\theta') + E(K'))f$. Note that $E(\theta') + E(K') = E(\theta) + E(K)$ and thus $\Psi^{\Theta'}(K'f) \leq (E(\theta) + E(K))f$. As $\Psi^\Theta(Kf) = \Psi^{\Theta'}(K'f)$ we conclude that $\Psi^\Theta(Kf) \leq (E(\theta) + E(K))f$. We prove (12) for an unbounded duration process by truncation: define the function $K \wedge n$ on the terminal nodes of $\Theta \wedge n$ by $(K \wedge n)(v) = E(K \mid v)$. As $\Psi^{\Theta \wedge n}((K \wedge n)f) \rightarrow_{n \rightarrow \infty} \Psi^\Theta(Kf)$ and $E(\theta \wedge n) + E(K \wedge n) \rightarrow_{n \rightarrow \infty} E(\theta) + E(K)$, (12) holds for any uncertain duration process. \square

A function f belongs to \mathcal{C}^+ if it satisfies the following: For all $\delta > 0$, there exists L_δ such that

$$\Psi(Kf) \leq (K + 1)f + \delta, \quad \forall K \geq L_\delta.$$

Similarly a function f belongs to \mathcal{C}^- if it satisfies the following: For all $\delta > 0$, there exists L_δ such that

$$\Psi(Kf) \geq (K + 1)f - \delta, \quad \forall K \geq L_\delta.$$

If a function f belongs to \mathcal{C}^+ , then

$$\Psi(K(f + \delta)) \leq (K + 1)(f + \delta), \quad \forall K \geq L_\delta,$$

and thus $f + \delta$ belongs to \mathcal{L}^+ for all $\delta > 0$; one obtains then an upper bound:

Corollary 3 *Let $f \in \mathcal{C}^+$; then*

$$\limsup_{E(\theta) \rightarrow \infty} \frac{\Psi^\Theta(0)}{E(\theta)} \leq f.$$

We now apply this property to continuous absorbing games. Proposition 7 and Corollary 8 in Rosenberg and Sorin (2001) prove that in this case the intersection of the closure of \mathcal{C}^+ and of \mathcal{C}^- is nonempty; hence reduced to one point. This provides an alternative proof of Corollary 2.

In the same spirit let us consider continuous recursive games. These are stochastic games where the set M_0 of nonabsorbing states is finite and the payoff is 0 at each of these states. The action spaces are compact sets X and Y . The absorbing payoffs as well as the transitions are continuous on $X \times Y$. From Proposition 16 in [Sorin \(2003\)](#), which proves that in this case also the intersection of the closure of C^+ and of C^- is nonempty, one deduces the following result:

Corollary 4 *If Γ is a continuous recursive game, there exists w in \mathbb{R}^{M_0} such that:*

$$\lim_{E(\theta) \rightarrow \infty} v_{\Theta}(\Gamma) = w.$$

4.2 Bounded variation of v_{λ}

We first recall a result from [Neyman \(2003b\)](#) dealing with a nonexpansive mapping Ψ on a Banach space X . Define $v_n = \frac{\Psi^n(0)}{n}$ and for $0 < \lambda \leq 1$ let $V_{\lambda} = \frac{v_{\lambda}}{\lambda}$ be the (unique) fixed point of the mapping $x \mapsto \Psi((1 - \lambda)x)$.

Definition 2 The function $\lambda \mapsto v_{\lambda}$ is of bounded variation (over $(0, 1]$) if there exists a constant C such that for any decreasing sequence λ_i with $0 < \lambda_{i+1} \leq \lambda_i \leq 1$,

$$\sum_i \|v_{\lambda_{i+1}} - v_{\lambda_i}\| \leq C.$$

If the function $\lambda \mapsto v_{\lambda}$ has bounded variation, then v_{λ} converges to a limit w as $\lambda \rightarrow 0^+$ and it is shown further in [Neyman \(2003b\)](#) that it implies the convergence of v_n to the same limit.

We will establish here a similar property under an additional monotonicity hypothesis on the uncertain duration process.

Definition 3 The uncertain duration process is monotonic if for every terminal node $v = (s_0, \dots, s_r)$, the conditional expectations $E(\theta - t \mid s_0, \dots, s_t)$ decrease in t , $0 \leq t \leq r$.

The interpretation is that the expected remaining duration decreases over time, implying that the relative weight of the present increases as the process evolves. Typical examples include finite length (where the ratio is $1/n$ if the remaining duration is n) and discounted factor uncertain duration (where the ratio is the constant λ).

We follow the proof in [Neyman \(2003b\)](#) to obtain

Theorem 6 *Assume v_{λ} is of bounded variation. Set $w = \lim_{\lambda \rightarrow 0^+} v_{\lambda}$. For every $\varepsilon > 0$, there exists N such that for any monotonic duration process satisfying $E(\theta) > N$,*

$$\|v_{\Theta} - w\| \leq \varepsilon.$$

Proof On the event $\{t < \theta\}$ we set $\rho_t = E(\theta - t \mid s_0, \dots, s_t)$, $\lambda_t = 1/\rho_t$, $w_t = v_{\lambda_t}$. On the event $\{\theta \leq t\}$ we set $\rho_t = 0$ and $w_t = 0$. We will prove that

$$\|\Psi^\Theta(0) - V_{1/E(\theta)}\| \leq E \left(\sum_{t \geq 0} \rho_{t+1} \|w_{t+1} - w_t\| \right). \tag{13}$$

For every $t \geq 0$ we define the random variables U_t and W_t as follows: on the event $\{t < \theta\}$ we define $U_t = \Psi^\Theta(s_0, \dots, s_t)(0)$ and $W_t = V_{\lambda_t} = v_{\lambda_t}/\lambda_t$; on the event $\{t \geq \theta\}$ we set $U_t = W_t = 0$. As θ is a stopping time the X -valued random variables U_t and W_t are measurable w.r.t. \mathcal{F}_t .

On the event $\{t < \theta\}$, $U_t = \Psi(E(U_{t+1} \mid \mathcal{F}_t))$, $W_t = \Psi((1 - \lambda_t)W_t)$ and $E(\rho_{t+1} \mid \mathcal{F}_t) = \rho_t - 1$. Therefore using the nonexpansiveness of Ψ followed by the triangle inequality and thereafter the equality $E(\rho_{t+1} \mid \mathcal{F}_t) = \rho_t - 1$, we have on $\theta > t$:

$$\begin{aligned} \|U_t - W_t\| &\leq \|E(U_{t+1} \mid \mathcal{F}_t) - ((1 - \lambda_t)W_t)\| \\ &\leq \|E(U_{t+1} - W_{t+1} \mid \mathcal{F}_t)\| + \|E(W_{t+1} \mid \mathcal{F}_t) - (1 - \lambda_t)W_t\| \\ &\leq E(\|U_{t+1} - W_{t+1}\| \mid \mathcal{F}_t) + \|E(\rho_{t+1}w_{t+1} - (\rho_t - 1)w_t \mid \mathcal{F}_t)\| \\ &\leq E(\|U_{t+1} - W_{t+1}\| \mid \mathcal{F}_t) + E(\rho_{t+1}\|w_{t+1} - w_t\| \mid \mathcal{F}_t). \end{aligned}$$

On the event $\theta \leq t$, $\|U_t - W_t\| = 0$. Therefore we have everywhere

$$\|U_t - W_t\| \leq E(\|U_{t+1} - W_{t+1}\| \mid \mathcal{F}_t) + E(\rho_{t+1}\|w_{t+1} - w_t\| \mid \mathcal{F}_t).$$

Summing the expectations (conditional on \mathcal{F}_0) of the above inequalities over $T \geq t \geq 0$ we deduce that

$$\|U_0 - W_0\| \leq E(\|U_T - W_T\| \mid \mathcal{F}_0) + E \left(\sum_{0 \leq t < T} \rho_{t+1} \|w_{t+1} - w_t\| \mid \mathcal{F}_0 \right).$$

As $\|U_T\| + \|W_T\| \leq 2\rho_T \|\Psi(0)\|$ (by non expansiveness), $E(\|U_T - W_T\| \mid \mathcal{F}_0) \leq 2E(\rho_T \mid \mathcal{F}_0) \|\Psi(0)\|$, which converges to 0 as $T \rightarrow \infty$, and therefore

$$\|U_0 - W_0\| \leq E \left(\sum_{t \geq 0} \rho_{t+1} \|w_{t+1} - w_t\| \mid \mathcal{F}_0 \right).$$

Observe that $\rho_0 = E(\theta)$ and therefore $\Psi^\Theta(0) = E(U_0)$ and $V_{1/E(\theta)} = E(W_0)$, which proves (13).

Fix $\varepsilon > 0$ and let K be sufficiently large so that the variation of v_λ over the interval $(0, 1/K)$ is less than ε . Assume that $E(\theta) > K$ and let θ' be the smallest r so that $E(\theta - r \mid (s_0, \dots, s_r)) < K$. We have

$$\|\Psi^\Theta(0) - V_{1/E(\theta)}\| \leq E \left(\sum_{0 \leq t < \theta'} \rho_{t+1} \|v_{\lambda_{t+1}} - v_{\lambda_t}\| \right) + E \left(\sum_{t \geq \theta'} \rho_{t+1} \|v_{\lambda_{t+1}} - v_{\lambda_t}\| \right).$$

The monotonicity of the uncertain duration process Θ implies that the sequence λ_t is monotonic. Therefore, for any $t \geq 0$, $\rho_t \leq \rho_0$, and for every $t \geq \theta'$, $\rho_{t+1} \leq K$. Hence:

$$\begin{aligned} \|\Psi^\Theta(0) - V_{1/E(\theta)}\| &\leq E \left(\rho_0 \sum_{0 \leq t < \theta'} \|v_{\lambda_{t+1}} - v_{\lambda_t}\| \right) + E \left(K \sum_{t \geq \theta'} \|v_{\lambda_{t+1}} - v_{\lambda_t}\| \right) \\ &\leq \rho_0 \varepsilon + KC, \end{aligned}$$

where C bounds the variation of the function v_λ .

Thus if $E(\theta) > KC/\varepsilon$, we deduce that $\|v_\Theta - v_{1/E(\theta)}\| < 2\varepsilon$, which imply that $\|v_\Theta - w\| < 3\varepsilon$. □

The inequality (13) has an alternative formulation using the probability tree associated with the uncertain duration process Θ . For every terminal node $v = (s_0, \dots, s_r)$, define

$$f(v) = \sum_{t=1}^r (r-t) \|w(s_0, \dots, s_t) - w(s_0, \dots, s_{t-1})\|,$$

where $\rho(s_0, \dots, s_t) = E(\theta - t \mid s_0, \dots, s_t)$ and $w(s_0, \dots, s_t) = v_{1/\rho(s_0, \dots, s_t)}$ if (s_0, \dots, s_t) is a non-terminal node; and it = 0 if (s_0, \dots, s_t) is a terminal node. Recall that μ_T is the probability induced on the terminal nodes. The alternative formulation is:

$$\|\Psi^\Theta(0) - V_{1/E(\theta)}\| \leq E_{\mu_T}(f(v)).$$

5 Game with lack of information on both sides

We consider here games with incomplete information (with finitely many states and finitely many actions; namely, M , I , and J are finite sets) as defined in Sect. 1.

In the case of finitely repeated games, it is proved in Mertens and Zamir (1971) that $v(p) = \lim_{n \rightarrow \infty} v_n(p)$ exists. Moreover, the error term $\|v_n - v\|$ ($:= \max_p |v_n(p) - v(p)|$) is bounded by a constant times $\frac{1}{\sqrt{n}}$. For the λ -discounted game it is also proved in Mertens and Zamir (1971) that $\|v_\lambda - v\|$ is bounded by a term of the order of $\sqrt{\lambda}$.

The purpose of this section is to extend these results of Mertens and Zamir (1971) to general public uncertain duration processes.

Theorem 7 *The limit of $v_\Theta(p)$ as $E(\theta) \rightarrow \infty$ exists, equals $v(p)$, and*

$$\|v_\Theta - v\| \leq O \left(\frac{1}{\sqrt{E(\theta)}} \right).$$

Proof We use the minmax theorem. It is thus sufficient to prove that there is a constant R such that for every uncertain duration process Θ and every strategy τ of Player 2, there is a strategy σ of Player 1 such that

$$E_{p,\sigma,\tau,\mu} \left(\sum_{t \geq 1} g_t I(\theta \geq t) \right) \geq v(p) E_\mu(\theta) - R \sqrt{E_\mu(\theta)} \tag{14}$$

where $E_{p,\sigma,\tau,\mu}$ is the expectation w.r.t. the probability defined on plays of the game Γ_Θ by the initial probability p on the state space, the strategy pair σ and τ , and by the probability μ , and E_μ is the expectation with respect to the probability μ .

We denote by \mathcal{H}_t the σ -algebra generated by the sequence of moves $i_1, j_1, \dots, i_{t-1}, j_{t-1}$ and by the sequence s_0, s_1, \dots, s_t of public signals prior to the play at stage t . W.l.o.g. we assume that the event $\theta \geq t$ is measurable w.r.t. \mathcal{H}_t . The expectation of a random variable is the expectation of its conditional expectation. Therefore

$$E \left(\sum_{t \geq 1} I(\theta \geq t) g_t \right) = E \left(\sum_{t \geq 1} I(\theta \geq t) E(g_t \mid \mathcal{H}_t) \right) \tag{15}$$

where E stands for the more explicit $E_{p,\sigma,\tau,\mu}$.

For every state $m \in M$ let G^m denote the $I \times J$ matrix $G^m_{i,j} = g(m, i, j)$. Set $\|G^m\| = \max_{i,j} |G^m_{i,j}|$, and let $\|G\| = \max_{m \in M} \|G^m\|$.

Let Θ and τ be given. From Theorem 4.4 and the proof of Proposition 4.3 in [Mertens and Zamir \(1971\)](#), there exists a strategy σ of Player 1 and a martingale $p_1, \tilde{p}_1, \dots, p_t, \tilde{p}_t, \dots$, with values in $\Delta(M)$ and $p_1 = p$, and where p_t is measurable w.r.t. \mathcal{H}_t , such that

$$v(p_t) \text{ is a submartingale,} \tag{16}$$

$$E_{p,\sigma,\tau,\mu}(g_t \mid \mathcal{H}_t) \geq v(p_t) - \|G\| E_{p,\sigma,\tau,\mu}(\|p_{t+1} - \tilde{p}_t\|_1 \mid \mathcal{H}_t), \tag{17}$$

and for any $\ell \geq t$, conditional on \mathcal{H}_t , the random variables

$$I(\theta \geq \ell) \text{ and } p_{t+1} \text{ are independent.} \tag{18}$$

By (15) and (17),

$$\begin{aligned} E \left(\sum_{t \geq 1} I(\theta \geq t) g_t \right) &\geq \sum_{t \geq 1} E(I(\theta \geq t) v(p_t)) \\ &\quad - \|G\| \sum_{t \geq 1} E(I(\theta \geq t) E(\|p_{t+1} - \tilde{p}_t\|_1 \mid \mathcal{H}_t)). \end{aligned} \tag{19}$$

□

Note that

$$E \left(\sum_{t \geq 1} I(\theta \geq t) E(\|p_{t+1} - \tilde{p}_t\|_1 \mid \mathcal{H}_t) \right) = E \left(\sum_{m \in M} \sum_{t \geq 1} I(\theta \geq t) |p_{t+1}(m) - \tilde{p}_t(m)| \right).$$

But we have:

Lemma 1 For any martingale $\{q_t\}$ with values in $[0, 1]$ and expectation q :

$$\begin{aligned} \sum_{t \geq 1} E \left(|q_{t+1} - q_t|^2 \right) &\leq q(1 - q) \\ E \left(\sum_{t \geq 1} I(\theta \geq t) |q_{t+1} - q_t| \right) &\leq \sqrt{E(\theta)} \sqrt{q(1 - q)}. \end{aligned}$$

Proof The first inequality follows from the fact that the differences $(q_{t+1} - q_t)$ of the martingale $\{q_t\}$ are uncorrelated. Hence

$$\sum_{t \geq 1}^T E \left(|q_{t+1} - q_t|^2 \right) = E \left(\left(\sum_{t \geq 1}^T q_{t+1} - q_t \right)^2 \right) \leq E \left((q_{T+1} - q)^2 \right)$$

and the variance of q_{T+1} is at most $q(1 - q)$. For the second property one has

$$\begin{aligned} E \left(\sum_{t \geq 1} I(\theta \geq t) |q_{t+1} - q_t| \right) &= \sum_{t \geq 1} E(I(\theta \geq t) |q_{t+1} - q_t|) \\ &\leq \sum_{t \geq 1} \sqrt{E(I(\theta \geq t)) E(|q_{t+1} - q_t|^2)} \\ &\leq \sqrt{\sum_{t \geq 1} E(I(\theta \geq t))} \sqrt{\sum_{t \geq 1} E(|q_{t+1} - q_t|^2)} \\ &\leq \sqrt{E(\theta)} \sqrt{q(1 - q)} \end{aligned}$$

where the first equality follows from the additivity of the expectation, the first inequality follows from the Cauchy Schwarz inequality applied to the functions $I(\theta \geq t)$ and $|q_{t+1} - q_t|$, the second inequality follows from the Cauchy Schwarz inequality applied to the vectors $(\sqrt{E(I(\theta \geq t))})_{t \geq 1}$ and $(\sqrt{E(|q_{t+1} - q_t|^2)})_{t \geq 1}$, and the last inequality follows from the first inequality of the lemma. □

Hence

$$\begin{aligned} E \left(\sum_{t \geq 1} I(\theta \geq t) E(\|p_{t+1} - \tilde{p}_t\|_1 \mid \mathcal{H}_t) \right) &\leq \sum_m (E(\theta))^{1/2} \sqrt{p^m(1-p^m)} \\ &\leq \sqrt{|M| - 1} (E(\theta))^{1/2} \end{aligned} \quad (20)$$

where the last inequality follows from Jensen's inequality and the concavity of the square root, and $|M|$ denotes the number of elements of the set M .

Next we show that

$$E \left(\sum_{t \geq 1} I(\theta \geq t) v(p_t) \right) \geq E(\theta) v(p_1). \quad (21)$$

Fix $\ell \geq t \geq 1$. As $I(\theta \geq \ell)$ and p_t are, by (18), independent conditionally on \mathcal{H}_{t-1} , and as, by (16), $v(p_t)$ is a submartingale,

$$E(I(\theta \geq \ell) v(p_t) \mid \mathcal{H}_{t-1}) \geq E(I(\theta \geq \ell) v(p_{t-1}) \mid \mathcal{H}_{t-1}).$$

Therefore,

$$E(I(\theta \geq t) v(p_t)) \geq E(I(\theta \geq t) v(p_1)). \quad (22)$$

As $E(\sum_{t \geq 1} I(\theta \geq t)) = E(\theta)$, summing inequality (22) over $t \geq 1$ proves inequality (21).

Inequalities (19), (21), and (20) yield formula (14) with the constant $R = \|G\| \sqrt{|M| - 1}$. \square

Comments

- (1) In the framework of games with state-independent signaling matrices (Mertens 1971), the proof of Theorem 7 implies, using Mertens et al. (1994, Lemma 4.6, p. 355), that

$$\|v_\Theta - v\| \leq O(E(\theta)^{-\frac{1}{3}}).$$

- (2) The conclusion of the theorem does not hold in the case of non-public uncertain duration (Neyman 2009b). In the non-public uncertain duration, Neyman (2009b) demonstrates the following result. If $\underline{v}(p)$ and $\bar{v}(p)$ denote the maxmin and the minmax of the repeated game, then $\liminf v_\Theta = \underline{v}(p)$ and $\limsup v_\Theta = \bar{v}(p)$ as $E(\theta) \rightarrow \infty$ (and Θ ranges over all asymmetric uncertain durations).
- (3) Theorem 7 holds (with the same proof) also for the model of an uncertain duration process where the distribution of the duration signal s_t is also a function of the sequence of moves $i_1, j_1, \dots, i_{t-1}, j_{t-1}$. If the distribution of the duration

signal s_t depends also on the state m we can no longer apply the proof. In fact, the limit of $v_{\Theta}(p)$ as $E(\theta)$ goes to infinity need not exist in this case. The essential difference is that as there is asymmetry in the information on the state space, the dependence of the duration signal, say s_1 , on the state m , results in asymmetric information regarding the duration.

- (4) Following [Neyman \(2009a,c\)](#), one can improve on the constant R in formula (14). For example, $\|G\|$ can be replaced with $E_p \|G^m\|$, and the term $\sqrt{|M| - 1}$ can be replaced by a constant C that is independent of the size of M (but is dependent on the size of I and J).

6 Open problems

Among other classes where the existence of the limit of the value of the discounted game has been established are recursive games with incomplete information on one side ([Rosenberg and Vieille 2000](#)) and absorbing games with lack of information on one side ([Rosenberg 1999](#)). In both cases one has in addition $\lim_{\lambda \rightarrow 0} v_{\lambda} = \lim_{n \rightarrow \infty} v_n$. The asymptotic behavior of v_{Θ} in such games deserves further research.

In the framework of general dynamic programming, [Lehrer and Sorin \(1992\)](#) proves the equivalence between the uniform convergence of the functions v_n and the uniform convergence of the functions v_{λ} . Whether this property has an extension to some family v_{Θ} is unknown. However, in the special class of finite action and state spaces with signals, the result of [Rosenberg et al. \(2002\)](#) on the existence of a uniform value proves in particular together with our Theorem 1 that v_{Θ} converges to a limit as $E(\theta) \rightarrow \infty$. Partial results extending this equivalence property to models with uncertain duration are in [Monderer and Sorin \(1993\)](#).

Acknowledgements Research supported in part by the Israel Science Foundation grants 382/98 and 1123/06, the Center for the Study of Rationality of the Hebrew University of Jerusalem (Israel), and by a grant ANR-08-BLAN-0294-01 (France).

References

- Aumann RJ, Maschler M (1995) Repeated games with incomplete information, with the collaboration of R. Stearns. MIT Press, Cambridge
- Lehrer E, Sorin S (1992) A uniform Tauberian theorem in dynamic programming. *Math Oper Res* 17:303–307
- Mertens J-F (1971) The value of two-person zero-sum repeated games: the extensive case. *Int J Game Theory* 1:217–227
- Mertens J-F, Neyman A (1981) Stochastic games. *Int J Game Theory* 10:53–66
- Mertens J-F, Zamir S (1971) The value of two-person zero-sum repeated games with lack of information on both sides. *Int J Game Theory* 1:39–64
- Mertens J-F, Zamir S (1985) Formulation of Bayesian analysis for games with incomplete information. *Int J Game Theory* 14:1–29
- Mertens J-F, Sorin S, Zamir S (1994) Repeated games. C.O.R.E. D.P. 9420, 9421, 9422
- Mertens J-F, Neyman A, Rosenberg D (2009) Absorbing games with compact action spaces. *Math Oper Res* 34:257–262
- Monderer D, Sorin S (1993) Asymptotic properties in dynamic programming. *Int J Game Theory* 22:1–11
- Neyman A (1999) Cooperation in repeated games when the number of stages is not commonly known. *Econometrica* 67:45–64

- Neyman A (2003a) Stochastic games: existence of the minmax. In: Neyman A, Sorin S (eds) Stochastic games and applications. NATO ASI series. Kluwer Academic Publishers, Dordrecht, pp 173–193
- Neyman A (2003b) Stochastic games and nonexpansive maps. In: Neyman A, Sorin S (eds) Stochastic games and applications. NATO ASI series. Kluwer Academic Publishers, Dordrecht, pp 397–415
- Neyman A (2009a) The maximal variation of martingales of probabilities and repeated games with incomplete information. DP 510, Center for the Study of Rationality, Hebrew University
- Neyman A (2009b) The value of two-person zero-sum repeated games with incomplete information and uncertain duration. DP 512, Center for the Study of Rationality, Hebrew University
- Neyman A (2009c) The error term in repeated games with incomplete information. DP 522, Center for the Study of Rationality, Hebrew University
- Rosenberg D (1999) Zero-sum absorbing games with incomplete information on one side: asymptotic analysis. *SIAM J Control Optim* 39:208–225
- Rosenberg D, Sorin S (2001) An operator approach to zero-sum repeated games. *Isr J Math* 121:221–246
- Rosenberg D, Vieille N (2000) The maxmin of recursive games with lack of information on one side. *Math Oper Res* 25:23–35
- Rosenberg D, Solan E, Vieille N (2002) Blackwell optimality in Markov decision processes with partial observation. *Ann Stat* 30:1178–1193
- Shapley LS (1953) Stochastic games. *Proc Natl Acad Sci USA* 39:1095–1100
- Sorin S (2003) Operator approach to stochastic games. In: Neyman A, Sorin S (eds) Stochastic games and applications. NATO ASI series. Kluwer Academic Publishers, Dordrecht, pp 417–426
- Sorin S (2004) Asymptotic properties of monotonic nonexpansive mappings. *Discret Events Dyn Syst* 14:109–122