STOCHASTIC GAMES WITH INCOMPLETE INFORMATION

SYLVAIN SORIN
Université P. et M. Curie and École Polytechnique
Paris, France

Abstract. This chapter deals with stochastic games where the state is not publicly known.

1. Introduction

In this chapter we consider situations described as stochastic games but where the state is not known to all the players. This is in fact a special case of a much more general class of games where both informational and stochastic features are present; see [3]. We will introduce here a collection of partial results: most of the models deal with two-person zero-sum games with lack of information on one side and standard signalling. First, structural properties related to the recursive structure will be studied, and then asymptotic aspects will be considered.

The main purpose of this chapter is to show that on the one hand some tools extend easily from stochastic games or incomplete information games to games having both aspects (see also [17], [3]). On the other hand we aim at presenting some fundamental differences with classical stochastic games: nonexistence of the value, and non-algebraicity of the minmax, maxmin or $\lim v_n$.

2. Classification

We introduce in this section several models of two-player stochastic games (in decreasing order of generality) where player 1 knows the state and has more information than player 2. It follows (see the entry laws in [3]) that the natural state parameter will be the beliefs of player 2 on the state space.

2.1. LEVEL 1

The game is specified by a state space S and a map q from $S\times A\times B$ to probabilities on $S\times\Omega$ where Ω is a set of public signals. We assume that each signal includes the move of player 2; hence ω determines $b = b(\omega)$. Let p be a probability on S according to which the initial state is chosen. It is then announced to player 1, while player 2 knows only p. At stage n, the state z_n and the moves a_n and b_n determine the payoff $r_n = r(z_n, a_n, b_n)$. $q(z_n, a_n, b_n)$ is the law of (z_{n+1}, ω_{n+1}) where z_{n+1} is the new state and ω_{n+1} is the public signal to both players. Write as usual $G_n(p)$ (resp. $G_\lambda(p)$) for the n-stage (resp. λ-discounted) version of the game. Stage after stage, the strategy of player 1 and the public signal determine a posterior distribution on S; hence a recursive formula for v_n and v_λ holds and properties of optimal strategies of player 1 obtain. The natural state space is thus $\Delta(S)$. This covers the usual stochastic game when $\Omega = S\times A\times B$ and q satisfies $\sum_{z'\in S} q(z', (z', a, b)|z, a, b) = 1$, for all (z, a, b): the state is public.

The standard stochastic game with incomplete information is obtained when $\Omega = A\times B$: the signal is the pair of moves and gives no information on the state; see [7], [6].

2.2. LEVEL 2: A FAMILY OF STOCHASTIC GAMES

The framework here corresponds to a finite family G^k, $k\in K$, of stochastic games on the same state space Ξ and with the same action spaces A and B. Hence, formally one has $S = K\times\Xi$. ξ_1 in Ξ is given and known to the players. k is chosen according to some initial probability p and player 1 is informed of k but not player 2. Then the game G^k is played with at each stage a new state and a public signal determined by $q^k(\xi, a, b)$ in $\Delta(\Xi\times\Omega)$. We also assume that the signal contains the new state in Ξ and the move of player 2. Note that since the transition depends on k the signal may be revealing on K even if player 1 plays without using his information $(x^{k,\xi} = x^{k',\xi})$.

The main difference with the previous case is that the beliefs at each stage are of the form $(\tilde{p}, \tilde{\xi})$ where $\tilde{p}$ is a martingale on K. The natural state space is $\Delta(K)\times\Xi$.

Note. If player 1 is not informed and the signal is symmetric, one obtains the framework of [18].

2.3. LEVEL 3: STOCHASTIC GAMES WITH VECTOR PAYOFFS

In this subcase the law of (ξ, ω) is independent of k. Basically, the game is then a (standard) stochastic game with state space Ξ but vector payoffs in $\mathbb{R}^K$. Player 1 knows the true component while player 2 has only an initial probability on it. If Ξ is reduced to one state, one has a game with lack of information on one side.

2.4. LEVEL 4: ABSORBING GAMES WITH VECTOR PAYOFFS

This is a specific case of Level 3 where the state can change at most once and then we are in a situation of a game with incomplete information on one side.

2.5. LEVEL 5: ABSORBING PAYOFF GAMES WITH VECTOR PAYOFFS

While in a standard two-person zero-sum stochastic game one can replace the payoffs, once an absorbing state is reached, by a constant (the value of the game at that state), this is no longer true with vector payoffs. This particular subclass of Level 4 assumes that the vector payoff is a constant once an absorbing state is reached; hence the name absorbing payoff.

3. Recursive Formula

We consider Level 1 and provide a formal construction of the recursive formula. The analysis is now standard [10], [6], [11].

Given an initial probability p, a stationary Markov strategy $\{x^z\}$, with $x^z \in \Delta(A)$, and a signal ω, define the conditional probability on S given ω by

$$\tilde{p}^z(\omega) = Prob(z|\omega) = \frac{Prob(z,\omega)}{Prob(\omega)}$$

with $Prob(z,\omega) = \sum_{z',a} p^{z'} x^{z'}(a) q(z,\omega|z',a,b(\omega))$, $Prob(\omega) = \sum_z Prob(z,\omega)$. Then one has

Proposition 1

$$nv_n(p) = \max_{X^S} \min_B \{\sum\nolimits_{z,a} p^z x^z(a) r(z,a,b) + (n-1) E_{p,x,b}[v_{n-1}(\tilde{p}(\omega))]\}$$

$$v_\lambda(p) = \max_{X^S} \min_B \{\lambda \sum\nolimits_{z,a} p^z x^z(a) r(z,a,b) + (1-\lambda) E_{p,x,b}[v_\lambda(\tilde{p}(\omega))]\}$$

Proof. The proof is the same in both cases and relies on the minmax theorem, namely that v_n (for example) exists. It is clear that player 1 can obtain the amount corresponding to the right-hand side: he plays at stage 1 some x realizing the maximum and, from stage 2 on, optimally in the game $G_{n-1}(\tilde{p}(\omega))$. On the other hand, given any strategy σ of player 1, we describe a strategy τ of player 2 achieving the same amount. In fact, σ determines x used by player 1 at stage one, then player 2 chooses a one-stage best reply according to the above formula and afterwards, given the signal ω, plays optimally in the remaining game $G_{n-1}(\tilde{p}(\omega))$. ∎

We deduce from these recursive formula properties of player 1's optimal strategies analogous to those available in stochastic or incomplete information games.

Corollary 1 *Player 1 has an optimal Markov strategy in $G_n(p)$ and an optimal Markov stationary strategy in $G_\lambda(p)$, where the state space is $\Delta(S)$.*

Recall that, on the contrary, the recursive formula does not allow one to construct recursively optimal strategies of player 2 since it involves the computation of the posterior distribution, unknown to player 2 who ignores x.

Remarks. It is actually enough for player 1 to be able to compute the beliefs of player 2; he doesn't have to know player 2's move. For example, in the standard case where the signal of player 2 is (a, b), the new information of player 2, namely a, is independent of his own move; hence player 1 does not need to know b to compute player 2's posterior distribution. However, if player 1 does not know the beliefs of player 2, the recursive structure (see [3]) does not allow one to write a recursive formula on $\Delta(S)$.

A similar construction holds in the case of lack of information on both sides: $S = S_1 \times S_2$, player 1 knows the first component, player 2 the second and the public signal contains both moves.

Melolidakis [8] deals with standard stochastic games with lack of information on one side and positive stop probabilities: the payoff is the sum of the stage payoffs but every day there is a positive probability $\rho(z, a, b)$ to stop the game. Then the same kind of properties as above are established.

4. Concavity and Duality

We still consider Level 1 but the result extends to the general case of incomplete information and follows from the original (1966) proof by Aumann and Maschler [1].

Proposition 2 *$v_n(p)$ and $v_\lambda(p)$ are concave on $\Delta(S)$.*

Proof. Again, the proof is similar in both cases and we will consider $G_n(p)$. Denote by $\sigma = \{\sigma^z\}$ and τ the mixed strategies of players 1 and 2 respectively, and by $\gamma^z(\sigma^z, \tau)$ the (bilinear) normalized payoff in G_n starting from state z. The payoff in $G_n(p)$ is then $\sum_{z\in S} p^z \gamma^z(\sigma^z, \tau)$.

Let $p = \sum_{\ell \in L} \alpha_\ell p_\ell$ with $L = \{1, 2\}$, p, p_ℓ in $\Delta(S)$ and $\{\alpha_\ell\}$ in $\Delta(L)$. Choose σ_ℓ optimal for player 1 in the game $G_n(p_\ell)$. Define σ as: if the initial state is z, use σ_ℓ with probability $\alpha_\ell \frac{p_\ell^z}{p^z}$. The corresponding payoff is

$$\sum_z p^z \gamma^z(\sigma^z, \tau) = \sum_z p^z \sum_\ell \alpha_\ell \frac{p_\ell^z}{p^z} \gamma^z(\sigma_\ell^z, \tau)$$

$$= \sum_{\ell} \alpha_\ell \sum_z p_\ell^z \gamma^z(\sigma_\ell^z, \tau);$$

hence by the choice of σ_ℓ

$$\sum_z p^z \gamma^z(\sigma^z, \tau) \geq \sum_\ell \alpha_\ell v_n(p_\ell), \qquad \forall \tau$$

which achieves the proof. ∎

We now follow [4] (see also [9], [11]) in introducing the dual game of the incomplete information game $G(p)$ (that stands for $G_n(p)$ or $G_\lambda(p)$) with payoff on $\Sigma^S \times \mathcal{T}$ defined by $\sum_{z \in S} p^z \gamma^z(\sigma^z, \tau)$.

For any ζ in $\mathbb{R}^S$, the dual game $G^*(\zeta)$ is played on $\Delta(S) \times \Sigma^S$ for player 1 and on $\mathcal{T}$ for player 2 with payoff

$$\gamma^*(p, \sigma; \tau) = \sum_z p^z(\gamma^z(\sigma^z, \tau) - \zeta^z).$$

In other words, player 1 chooses the initial state z, then the game evolves as in G and there is an additional payoff of $-\zeta^z$. Denoting by $v^*(\zeta)$ the value of $G^{*(}\zeta)$ one has

Proposition 3

$$v^*(\zeta) = \max_{p \in \Delta(S)} \{v(p) - \langle p, \zeta \rangle\} \tag{1}$$

$$v(p) = \inf_{\zeta \in \mathbb{R}^S} \{v^*(\zeta) + \langle p, \zeta \rangle\}. \tag{2}$$

Proof. By definition,

$$\begin{aligned} v^*(\zeta) &= \sup_{p,\sigma} \inf_\tau \sum_z p^z(\gamma^z(\sigma^z, \tau) - \zeta^z) \\ &= \sup_p (\sup_\sigma \inf_\tau \sum_z p^z \gamma^z(\sigma^z, \tau) - \langle p, \zeta \rangle) \\ &= \sup_p (v(p) - \langle p, \zeta \rangle) \end{aligned}$$

and v is continuous (even Lipschitz); hence the sup is reached.

Since v is in addition concave, equation (2) follows from Fenchel's duality. ∎

Corollary 2 *Given p, let ζ be ε-optimal in (2) and τ be ε-optimal in $G^*(\zeta)$. Then τ is 2ε-optimal in $G(p)$.*

Proof. By the choice of τ,

$$\gamma^z(\sigma^z, \tau) - \zeta^z \leq v^*(\zeta) + \varepsilon, \qquad \forall \sigma.$$

Hence

$$\begin{aligned}\sum\nolimits_z p^z\gamma^z(\sigma^z,\tau) &\leq v^*(\zeta)+\langle p,\zeta\rangle+\varepsilon\\ &\leq v(p)+2\varepsilon, \qquad \forall\sigma\end{aligned}$$

by the choice of ζ. ■

The interest of this result is to deduce properties of optimal strategies of player 2 in $G(p)$ from similar properties in $G^*(\zeta)$.

5. Markov Strategies for the Uninformed Player

The recursive formula obtained in the primal game $G(p)$ (Proposition 1) is used here to get a similar representation for the value of the dual game $G^*(\zeta)$ and eventually to deduce properties of optimal strategies of player 2 in G^*. This approach follows [4], [11].

One starts with the duality equation (1)

$$v_n^*(\zeta)=\max_p\{v_n(p)-\langle p,\zeta\rangle\}$$

and the recursive formula

$$nv_n(p)=\max_{X^S}\min_B\{\sum\nolimits_{z,a}p^zx^z(a)r(z,a,b)+(n-1)E_{p,x,b}[v_{n-1}(\tilde p(\omega))]\}.$$

Recall that $X=\Delta(A)$ and $Y=\Delta(B)$; hence, by introducing $\pi\in\Pi=\Delta(S\times A)$, where the marginal π_S on S plays the role of p and the conditional on A given z the role of x^z, one obtains

$$\begin{aligned}nv_n^*(\zeta) &= \max_{\pi\in\Pi}\min_{y\in Y}\{\sum\nolimits_{z,a,b}\pi(z,a)y(b)r(z,a,b)\\ &\quad +(n-1)E_{\pi,y}[v_{n-1}(\tilde\pi_S(\omega))]-n\langle\zeta,\pi_S\rangle\},\end{aligned}\tag{3}$$

where $\tilde\pi_S(\omega)$ is the conditional distribution on S, given ω.

Denote by (**C**) the following condition:

(**C**) $\qquad \pi\mapsto E_{\pi,y}[v_{n-1}(\tilde\pi_S(\omega))]$ is concave.

Proposition 4 *Under (***C***) one has*

$$\begin{aligned}nv_n^*(\zeta) &= \min_{y\in Y}\inf_{\rho:\Omega\to\mathbb{R}^S}\max_{\pi\in\Pi}\{\sum\nolimits_{z,a,b}\pi(z,a)y(b)r(z,a,b)\\ &\quad +(n-1)E_{\pi,y}[v_{n-1}^*(\rho(\omega))+\langle\tilde\pi_S(\omega),\rho(\omega)\rangle]-n\langle\zeta,\pi_S\rangle\}.\end{aligned}\tag{4}$$

Proof. The payoff in (3) is concave in Π and linear in Y, both sets being convex compact; hence Sion's minmax theorem applies, so that

$$\begin{aligned} nv_n^*(\zeta) &= \min_{y\in Y}\max_{\pi\in\Pi}\{\textstyle\sum_{z,a,b}\pi(z,a)y(b)r(z,a,b) \\ &\quad +(n-1)E_{\pi,y}[v_{n-1}(\tilde{\pi}_S(\omega))] - n\langle\zeta,\pi_S\rangle\}. \end{aligned}$$

We now use the other duality equation (2) to obtain

$$\begin{aligned} nv_n^*(\zeta) &= \min_{y\in Y}\max_{\pi\in\Pi}\{\textstyle\sum_{z,a,b}\pi(z,a)y(b)r(z,a,b) \\ &\quad +(n-1)E_{\pi,y}[\inf_{\zeta'}\{v_{n-1}^*(\zeta')+\langle\tilde{\pi}_S(\omega),\zeta'\rangle\}] - n\langle\zeta,\pi_S\rangle\}. \end{aligned}$$

Finally, from (1) v^* is convex; hence one has, again using the minmax theorem,

$$\begin{aligned} \max\ & {}_{\pi\in\Pi}\{\textstyle\sum_{z,a,b}\pi(z,a)y(b)r(z,a,b) \\ & +(n-1)E_{\pi,y}[\inf_{\zeta'}\{v_{n-1}^*(\zeta')+\langle\tilde{\pi}_S(\omega),\zeta'\rangle\}] - n\langle\zeta,\pi_S\rangle\} \\ = \ & \max_{\pi\in\Pi}\inf_{\rho:\Omega\to\mathbb{R}^S}\{\textstyle\sum_{z,a,b}\pi(z,a)y(b)r(z,a,b) \\ & +(n-1)E_{\pi,y}[v_{n-1}^*(\rho(\omega))+\langle\tilde{\pi}_S(\omega)),\rho(\omega)\rangle] - n\langle\zeta,\pi_S\rangle\} \\ = \ & \inf_{\rho:\Omega\to\mathbb{R}^S}\max_{\pi\in\Pi}\{\textstyle\sum_{z,a,b}\pi(z,a)y(b)r(z,a,b) \\ & +(n-1)E_{\pi,y}[v_{n-1}^*(\rho(\omega))+\langle\tilde{\pi}_S(\omega),\rho(\omega)\rangle] - n\langle\zeta,\pi_S\rangle\}, \end{aligned}$$

which gives (4). ∎

Proposition 5 *Condition (***C***) holds for games at Level 3.*

Proof. Recall that the state space S is decomposed as $\Xi\times K$. The component ξ is known to both and is varying, and the component k is fixed and known to player 1 only. The signal ω contains the new state $\xi(\omega)$ so that one has, at ξ with $\pi\in\Delta(K\times A)$,

$$Prob(\omega) = Prob(\omega,\xi(\omega)) = \textstyle\sum_{k,a}\pi(k,a)q(w,\xi(\omega)|\xi,a,b(\omega)).$$

Since the transition on Ξ is independent of k, the law of $\tilde{\pi}_K(.,\xi')$ for a given ξ' will be a martingale. Explicitly in our framework, ξ and b being given, $E_{\pi,y}[v_{n-1}(\tilde{\pi}_S(\omega))]$ is of the form

$$\textstyle\sum_{\omega,\xi'}Prob_\pi(\omega,\xi')v_{n-1}(\tilde{\pi}_K(\omega),\xi').$$

Let $\pi = \alpha\pi_1 + (1-\alpha)\pi_2$, then for each ξ' fixed one has

$$\begin{aligned}
\tilde{\pi}^k_K(\omega) &= \frac{Prob_\pi(\omega,k)}{Prob_\pi(\omega)} = \alpha\frac{Prob_{\pi_1}(\omega,k)}{Prob_\pi(\omega)} + (1-\alpha)\frac{Prob_{\pi_2}(\omega,k)}{Prob_\pi(\omega)} \\
&= \alpha\tilde{\pi}^k_{K,1}(\omega)\frac{Prob_{\pi_1}(\omega)}{Prob_\pi(\omega)} + (1-\alpha)\tilde{\pi}^k_{K,2}(\omega)\frac{Prob_{\pi_2}(\omega,k)}{Prob_\pi(\omega)}
\end{aligned}$$

and the result follows from v_{n-1} being concave (Proposition 3). ∎

It follows that (at Level 3) an optimal strategy of player 2 in the dual game consists in choosing y and ρ optimally in the above equation (4), to play y at stage one and then, given the signal ω, optimally in the remaining game at state $\zeta = \rho(\omega)$.

By Corollary 2 an analog property holds in the primal game.

Remarks. A similar result is true even if the law of the signals depends upon k (but not the transition on the states ξ). The crucial point is the martingale property of $\tilde{\pi}(\omega,\xi')$; see [11].

In the framework of stochastic games with lack of information on both sides, there are two dual games corresponding to the information of each player. Explicitly, let the state space be of the form $\Xi\times K_1\times K_2$. k_1 and k_2 are fixed and correspond to the private information of the players. In addition, the transition on Ξ is independent of (k_1,k_2) and (ξ,a,b) are announced to both players. Then an optimal strategy of player 1 will be based on a triple $(\tilde{\xi},\tilde{p}_1,\tilde{\zeta}_1)$ corresponding to the stochastic state, the beliefs of player 2 and player 1's vector parameter ζ_1 representing his uncertainty on K_2. Player 1's "state space" is thus $\Xi\times\Delta(K_1)\times\mathbb{R}^{K_2}$; see [11].

At level 1, for $G_n(p)$, Krausz and Rieder [6] use the finite aspect to describe the value as the solution of a linear program. Optimal strategies of player 1 are obtained as optimal variables. Using the fact that $v_n(p)$ is piecewise linear they also obtain inductively optimal strategies of player 2 as the solution to the dual program.

We now turn to the study of large games: first, asymptotic analysis of v_n or v_λ in the compact case and then properties of the maxmin and minmax in the uniform approach.

6. Recursive Games with Incomplete Information

We follow here the work of Rosenberg and Vieille [13].

Consider a finite two-person zero-sum *recursive game with absorbing payoffs*: the payoff is either 0 or absorbing. We denote by S the set of non-absorbing states, by A the set of absorbing states and by p the initial probability on S according to which the state is chosen and announced to player 1. At each stage n, given the strategy of player 1, player 2 computes the posterior probability ρ_n, conditional to the fact that the payoff is not absorbing. (If only the state were absorbing, player 2 should know it in order to play optimally after absorption; here, on the contrary, in case an absorbing state is reached the future payoff is constant; hence player 2 can concentrate on the complementary event to select his strategy.)

The recursive formula is given on $\Delta(S)$ by

$$\Phi(\alpha, f)(p) = \mathtt{val}_{X^S \times Y}\{(1-\alpha)(\pi(p,x,y)E(f(\rho)) + (1-\pi(p,x,y))E(a))\}$$

where $\pi(p,x,y)$ is the probability of remaining in S and ρ the corresponding posterior probability. a stands for the absorbing payoff with $\|a\|\leq 1$. Obviously, $v_n = v_\lambda = v_\infty$ on A and we are interested in their behavior on $\Delta(S)$.

Theorem 1

$$\max\min = \lim_{n\to\infty} v_n = \lim_{\lambda\to 0} v_\lambda.$$

Proof. Let w be an accumulation point of the family $\{v_\lambda\}$ (which is uniformly Lipschitz on $\Delta(S)$). One has

$$\Phi(0, w) = w.$$

As long as w is negative, if player 1 uses an optimal strategy in the "projective game" corresponding to $\Phi(0, w)$, this will guarantee w since the current payoff is $0 \geq w$. However, the argument fails if $w(p) > 0$ and the idea is then to play optimally in a discounted game with $\|v_\lambda - w\|$ small.

The main lines of the proof are sketched below. Given $\varepsilon > 0$, let λ such that

$$\|v_\lambda - w\| \leq \varepsilon^2.$$

Write $\mathbf{x}_\lambda(p)$ (resp. $\mathbf{x}(p)$) for an optimal strategy of player 1 in $\Phi(\lambda, v_\lambda)(p)$ (resp. $\Phi(0, w)(p)$). Inductively, a strategy σ of player 1 and stopping times θ_ℓ are defined as follows. Let $\theta_1 = \min\{m : w(\rho_m) > \varepsilon\}$ and play $\mathbf{x}(\rho_n)$ at each stage n until θ_1 (excluded). Let then $\theta_2 = \min\{m \geq \theta_1; v_\lambda(\rho_m) < 0\}$ and play $\mathbf{x}_\lambda(\rho_n)$ at each stage n from θ_1 until θ_2 (excluded).

More generally, play $\mathbf{x}(\rho_n)$ from stage $\theta_{2\ell}$ to $\theta_{2\ell+1} = \min\{m : w(\rho_m) > \varepsilon\}$(excluded) and play $\mathbf{x}_\lambda(\rho_n)$ from $\theta_{2\ell+1}$ until $\theta_{2\ell+2} = \min\{m; v_\lambda(\rho_m) < 0\}$(excluded).

Define u_n to be $w(\rho_n)$ at nodes where player 1 is using $\mathbf{x}$ (i.e., playing optimally for $\Phi(0, w)$), namely for $\theta_{2\ell} \leq n < \theta_{2\ell+1}$. Let u_n be $v_\lambda(\rho_n)$ otherwise. We call the first set of nodes "increasing" and the other set "decreasing."

The first property is that u_n is essentially a submartingale. This is clear if one starts at an increasing node and stays in this set since by the choice of σ:

$$E_{\sigma,\tau}(u_{n+1}|\mathcal{H}_n) = E_{\sigma,\tau}(w|\mathcal{H}_n) \geq \Phi(0, w) = w = u_n.$$

Similarly, if the initial node is decreasing and one remains in this set, one obtains using the fact that $v_\lambda(p_n)\geq 0$ (by the choice of the stopping time):

$$E_{\sigma,\tau}((1-\lambda)u_{n+1}|\mathcal{H}_n) = E_{\sigma,\tau}((1-\lambda)v_\lambda|\mathcal{H}_n)\geq\Phi(\lambda, v_\lambda) = v_\lambda = u_n\geq 0,$$

so that

$$E_{\sigma,\tau}(u_{n+1}|\mathcal{H}_n)\geq u_n.$$

Now if one of the new nodes changes from decreasing to increasing or vice versa, the error is at most ε^2; hence in all cases

$$E_{\sigma,\tau}(u_{n+1}|\mathcal{H}_n)\geq u_n - \varepsilon^2 P(n+1\in\Theta|\mathcal{H}_n),$$

where Θ is the set of all stopping times $\{\theta_\ell\}$.

The second property is a bound on the error term using the fact that the stopping times count the upcrossing of the band $[0,\varepsilon]$ by the sequence u_n. If η^N denotes the number of stopping times θ_ℓ before stage N and $\eta = \lim \eta^N$ one has

$$E(\eta) \leq \frac{2}{\varepsilon - \varepsilon^2}$$

and one uses

$$\sum\nolimits_n P(n+1\in\Theta) = \sum\nolimits_\ell \sum\nolimits_{n+1} P(\theta_\ell = n+1)\leq E(\eta) + 1$$

to get finally

$$E(u_n) \geq u_1 - 5\varepsilon.$$

The last point is to compare u_n to the current payoff in the game. Until absorption the current payoff is 0, hence near w (or v_λ) as long as $w\leq\varepsilon$. Define $\mathcal{A}_n$ to be the set of non-absorbing nodes with $w(\rho_n) > \varepsilon$. One obtains

$$E(g_n) \geq u_1 - 7\varepsilon - 2P(\mathcal{A}_n).$$

Denoting by ξ the absorbing time, the crucial property is that $\forall\varepsilon, \lambda, \exists N$ such that

$$P(\xi \leq n + N|\mathcal{A}_n) \geq \varepsilon/2.$$

This result follows from the fact that given a node in $\mathcal{A}_n$, player 1 is using $\mathbf{x}_\lambda$ as long as $v_\lambda(\rho_m)\geq 0$. Now, before absorption, $E((1-\lambda)v_\lambda) \geq v_\lambda$. Since v_λ is bounded, positive and increasing geometrically there is a positive probability of absorption in finite time.

One then deduces that $\sum_n P(\mathcal{A}_n)$ is uniformly bounded; hence

$$E(\overline{g}_n) \geq w - 8\varepsilon$$

for n large enough. Since the strategy of player 1 is independent of the length of the game, this implies that player 1 can guarantee w.

Given any strategy σ of player 1, player 2 can compute the posterior distribution ρ_n as well and use the "dual" of the previous strategy. The same bound (independent of σ) thus implies that $\max\min = w$ and moreover $\lim v_n = \lim v_\lambda = w$. ∎

An example is given in [13] showing that max min and min max may differ.

On the other hand the previous proof shows that the crucial point is the knowledge of the beliefs parameter. Hence one obtains

Proposition 6 *Consider a recursive game with absorbing payoffs and lack of information on both sides. Then*

$$\lim_{n\to\infty} v_n = \lim_{\lambda\to 0} v_\lambda.$$

7. Absorbing Games with Incomplete Information: Level 4

We will just mention here a recent result by Rosenberg [12].

Theorem 2 *For absorbing games with vector payoffs and incomplete information on one side, both* $\lim v_n$ *and* $\lim v_\lambda$ *exist and coincide.*

The proof is very involved and uses the operator approach (see [17]) to obtain variational properties satisfied by any accumulation point of the family $\{v_\lambda\}$ and then to deduce uniqueness.

8. Absorbing Games with Incomplete Information: Level 5

This is a collection of partial results introducing new tools and ideas that may be useful in more general cases. The games under consideration have a structure similar to the Big Match of Blackwell and Ferguson [2] (see also [21]), namely, these are absorbing games where one of the players controls the absorption. However, there is some incomplete information on the state; hence the name for this class.

8.1. "BIG MATCH" WITH INCOMPLETE INFORMATION: TYPE I

We consider a family of games of the following form

$$G^k = \begin{pmatrix} a_1^{k*} & a_2^{k*} & \dots & a_m^{k*} \\ b_1^k & b_2^k & \dots & b_m^k \end{pmatrix}$$

where the first line is absorbing. k belongs to a finite set K and is selected according to p in $\Delta(K)$. Player 1 knows k while player 2 knows only p. The analysis follows [14].

8.1.1. *Asymptotic Analysis*

The use of the recursive formula allows us to deduce properties of optimal strategies. In particular, in our case the value of the game is the same if both players are restricted to strategies independent of the past: first, the information transmitted to player 2 is independent of his own moves, so one can ignore them; second, there is only one past history of moves of player 1 to take into consideration, namely Bottom up to the current stage (excluded). This suggests the construction of an asymptotic game $\mathcal{G}$ played between time 0 and 1 and described as follows. ρ^k is the law of the stopping time θ corresponding to the first stage where player 1 plays Top, if k is announced: $\rho^k(t) = Prob_{\sigma^k}(\theta \leq t)$.

f is a map from $[0,1]$ to $\Delta(Y)$, $f(t)$ being the mixed strategy used by player 2 at time t.

The payoff is given by $L(\{\rho\}, f) = \sum_k p^k L^k(\rho^k, f)$ where L^k is the payoff in game k, expressed as the integral between 0 and 1 of the "payoff at time t":

$$L^k(\rho^k, f) = \int_0^1 L_t^k(\rho^k, f)dt$$

with, letting $A^k f = \sum a_j^k f_j$ and similarly for $B^k f$, the following expression for L_t^k:

$$L_t^k(\rho^k, f) = \int_0^t A^k f(s)\rho^k(ds) + (1 - \rho^k(t))B^k f(t).$$

The first term corresponds to the absorbing component and the second term to the non-absorbing one.

Theorem 3 *1) The game $\mathcal{G}$ has a value w.*
2) $\lim_{n\to\infty} v_n = \lim_{\lambda\to 0} v_\lambda = w$.

Proof. The existence of a value follows from Sion's minmax theorem. Consider now (ε)-optimal strategies $(\rho = \{\rho^k\}, f)$ in $\mathcal{G}$. They induce natural discretizations $(\rho(n), f(n))$ or $(\rho(\lambda), f(\lambda))$ in G_n or G_λ corresponding to piecewise constant approximations on the intervals of the form $[\frac{m}{n}, \frac{m+1}{n}]$ or $[\sum_{t=0}^{m-1} \lambda\ (1-\lambda)^{t-1}, \sum_{t=0}^{m} \lambda (1-\lambda)^{t-1}]$. It is then easy to see by continuity of the payoffs that $\rho(n)$ will guarantee w up to some constant $\times[(1/n) + \varepsilon]$ in G_n and a dual result holds for $f(n)$. A similar property is obtained for G_λ. ∎

8.1.2. *Maxmin*
The construction relies on properties of the previous auxiliary game $\mathcal{G}$ and the result is the following:

Theorem 4

$$maxmin = w.$$

Proof. We first prove that player 2 can defend w. Let f be an ε-optimal strategy of player 2 in $\mathcal{G}$. Player 2 will mimic f in order to generate through the strategy σ of player 1 a family of distributions $\{\mu^k\}$ such that by playing "up to level t" the payoff will be near $L_t(\{\mu\}, f) = \sum_k L_t^k(\mu^k, f)$. Since by the choice of f, $L(\{\mu\}, f) = \int_0^1 L_t(\{\mu\}, f)$ is less than $\omega + \varepsilon$, there exists t^* with $L_{t^*}(\{\mu\}, f) \leq w + \varepsilon$. This will define the strategy τ of player 2 as: follow f up to level t^*.

Formally, we consider a discrete-valued approximation $\tilde{f}$ of f, $\tilde{f}$ being equal to f_i on $[t_i, t_{i+1}]$, with $i \in I$, finite.

Given σ, the positive measures μ^k are defined inductively as follows.

$\overline{\mu}^k(t_1) = Prob_{\sigma^k,\tau_1}(\theta < +\infty)$ where τ_1 is f_1 i.i.d.

Let N_1 be such that the above probabilities are almost achieved by that stage for all k; this defines $\mu^k(t_1) = Prob_{\sigma,\tau_1}(\theta \leq N_1)$.

$\overline{\mu}^k(t_2) = Prob_{\sigma^k,\tau_2}(\theta < +\infty)$, where τ_2 is τ_1 up to stage N_1 and then f_2 i.i.d. One introduces N_2, τ_2 as above and so on.

$\overline{\mu}^k(t_i) = Prob_{\sigma^k,\tau_i}(\theta < +\infty)$ where τ_i is τ_{i-1} up to N_{i-1} and then f_1 i.i.d.

It is then clear that the payoff induced in G_n, for n large enough, by σ and τ_i, will be of the form

$$\sum\nolimits_k p^k \{\sum\nolimits_{j \leq i} (\mu^k(t_j) - \mu^k(t_{j-1})) A^k f_j + (1 - \mu^k(t_i)) B^k f_i\},$$

hence near $L_{t_i}(\{\mu\}, \tilde{f})$. Since $\int_0^1 L_t(\mu, \tilde{f}) dt$ is at most w (up to some approximation), there exists an index i^* with $L_{t_i}(\{\mu\}, \tilde{f})$ below $w + O(\varepsilon)$. Finally, the strategy τ_{i^*} defends w.

The proof that player 1 can guarantee w is more intricate. One first shows the existence of a couple of optimal strategies ($\rho = \{\rho^k\}, f$) in $\mathcal{G}$ that are essentially equalizing, namely such that $L_t(\{\rho\}, f)$ is near w for all t. In fact, consider $\{\rho\}$ optimal for player 1 in $\mathcal{G}$ and the game $\wp$ where player 1 chooses t, player 2 chooses f and the payoff is $L_t(\{\rho\}, f)$.

Proposition 7 *The game $\wp$ has a value, w.*

Proof. The existence of a value follows again from Sion's minmax theorem. Since player 1 can choose the uniform distribution on $[0, 1]$ and so generate $L(\{\rho\}, f)$, the value w' is at least w. If $w' > w$, an optimal strategy of player 1, hence a cumulative distribution function on $[0, 1]$, α, could be used to "renormalize" the time and induce in $\mathcal{G}$ through the image of ρ by α a payoff always at least w'. ■

The idea of the proof is now to follow the "path defined by f and ρ." Basically, given k, player 1 will choose t according to the distribution ρ^k and play the strategy δ_t where δ_t is defined inductively as follows. Consider the non-absorbing payoff at time t

$$\sum\nolimits_k p^k(1-\mu^k(t))B^k f(t) = b_t(f).$$

Player 1 then uses a "Big Match" strategy blocking whenever the non-absorbing payoff evaluated through $b_t(.)$ is less than $b_t(f)$. The equalizing property of f then implies that the absorbing payoff will be at least the one corresponding to f. It follows that the total payoff is minorized by an expectation of terms of the form $L_t(\{\rho\}, f)$, hence the result. ■

8.1.3. *Minmax*

Theorem 5 *minmax* $= v_1$, *value of the one-shot game.*

Proof. It is clear that by playing i.i.d. an optimal strategy y in the one-shot game player 2 will induce an expected payoff at any stage n of the form

$$g_1(p;\alpha,y) = \sum\nolimits_k p^k(\alpha^k A^k y + (1-\alpha^k)B^k y),$$

where $\alpha^k = Prob_{\sigma^k,\tau}(\theta \leq n)$, hence less than v_1.

To prove that player 1 can defend v_1, let $\alpha = \{\alpha^k\}$ be an optimal strategy for him in $G_1(p)$. Knowing τ, player 1 evaluates the non-absorbing component of the payoff at stage n given α, namely:

$$c_n = \sum\nolimits_k p^k(1-\alpha^k)B^k \overline{y}_n,$$

where $\overline{y}_n = E(\tau(h_n)|\theta \geq n)$ is the expected mixed move of player 2 at stage n, conditional to Bottom up to that stage. Let N be such that $c_N > \sup_n c_n - \varepsilon$; then player 1 plays Bottom up to stage N excluded, then once α at stage N and always Bottom thereafter. For n larger than N, the expected payoff will be of the form

$$\sum\nolimits_k p^k(\alpha^k A^k \overline{y}_N) + c_n,$$

hence greater than $g_1(p;\alpha,\overline{y}_N) - \varepsilon$, which gives the result. ■

Example. Consider the following game with $p = p^1 = Prob(G^1)$:

$$G^1 = \begin{pmatrix} 1^* & 0^* \\ 0 & 0 \end{pmatrix} \qquad G^2 = \begin{pmatrix} 0^* & 0^* \\ 0 & 1 \end{pmatrix}.$$

Then one has

$$\overline{v}(p) = v_1(p) = \min(p, 1-p)$$

$$\underline{v}(p) = \lim_{n\to\infty} v_n(p) = \lim_{\lambda\to 0} v_\lambda(p) = (1-p)(1-exp(-\frac{p}{1-p})).$$

In particular, the uniform value does not exist, and the asymptotic value and the maxmin are transcendental functions: at $p = \frac{1}{2}$ one obtains $\underline{v}(\frac{1}{2}) = \frac{1}{2}(1-\frac{1}{e})$ while all the data are rational numbers.

8.1.4. *Extensions*

We study here the extension to Level 4. The games are of the form

$$\begin{pmatrix} A_1^{k*} & A_2^{k*} & ... & A_m^{k*} & ... \\ b_1^k & b_2^k & ... & b_m^k & ... \end{pmatrix}$$

where $A_1 = \{A_1^k\}$, ... $A_m = \{A_m^k\}$, ... are games with incomplete information corresponding to absorbing states. It follows that when player 1 plays Top the payoff is not absorbing and the strategic behavior thereafter (hence also the payoff) will be a function of the past. Let $v_m(p)$ be the uniform value of the game A_m with initial distribution p [1]. The recursive formula implies that the absorbing payoff is approximately $\sum_m v_m(p^T)y_m$ (where p^T is the conditional distribution given Top and y_m the probability of playing column m) if the number of stages that remains to be played is large enough.

Consider now the continuous time game $\mathcal{G}$. Given a profile $\rho = \{\rho^k\}$, denote by $p^T(t)$ (resp. $p^B(t)$) the conditional probability on K given $\theta = t$ (resp. $\theta > t$). The payoff is defined as

$$M(\{\rho\}, f) = \int_0^1 M_t(\{\rho\}, f)dt$$

where the payoff at time t is given by

$$M_t(\{\rho\}, f) = \int_0^t (\sum\nolimits_m v_m(p^T(s))f_m(s))d\overline{\rho}(s) + (1-\overline{\rho}(t))(\sum\nolimits_k p_k^B(t)b^k f(t)),$$

$\overline{\rho}(t) = \sum_k p^k\rho^k(t)$ being the average probability of the event $\{\theta \leq t\}$.

M is still a concave function of ρ (due to the concavity of each v_m) and Sion's theorem still applies. The analog of Theorem 3 then holds. One shows that player 1 can obtain w in large games, and using the minmax theorem, that he cannot get better. Similarly, the analysis of the maxmin follows the same lines.

Concerning the minmax one is led to introduce a family of games as follows. For each game A_m consider the set Ξ_m of vector payoffs (in $\mathbb{R}^K$) that player 2 can approach (see [1]), namely such that

$$\langle p, \xi_m \rangle \geq v_m(p) \quad \forall p \in \Delta(K).$$

Given a profile $\{\xi_m\}$ of vectors in $\prod_m \Xi_m$ we consider the game $\mathcal{A}(\xi, p)$, where each component is given by

$$\mathcal{A}^k(\xi) = \begin{pmatrix} \xi_1^{k*} & \xi_2^{k*} & \dots & \xi_m^{k*} & \dots \\ b_1^k & b_2^k & \dots & b_m^k & \dots \end{pmatrix}.$$

By construction for each such ξ player 2 can guarantee (in the original game) the minmax of $\mathcal{A}(\xi, p)$ which is the value of the one-shot version, say $\nu_1(\xi, p)$. One then has

$$\text{minmax} = \min_{\xi \in \Xi} \nu_1(\xi, p).$$

In fact, by playing optimally for the minmax in $\mathcal{A}(\xi, p)$, player 1 is anticipating the behavior of player 2, after absorption (namely, approach ξ_m if absorption occurred when playing m). The best player 2 could do then would be to choose a supporting hyperplane to v_m at the current posterior p^T. This defines a correspondence C from $\prod_m \Xi_m$ to itself. One shows that C is u.s.c. with convex values; hence it has a fixed point ξ^*. Playing optimally for the minmax against τ in $\mathcal{A}(\xi^*, p)$ will then guarantee an absorbing payoff above ξ^*, hence a total payoff above $\nu_1(\xi^*, p)$.

Note that this construction is reminiscent of the approach in [16].

8.2. "BIG MATCH" WITH INCOMPLETE INFORMATION: TYPE II

We here consider games of the form

$$G^k = \begin{pmatrix} a_1^* & b_1 \\ a_2^* & b_2 \\ \dots & \dots \\ a_m^* & b_m \\ \dots & \dots \end{pmatrix}$$

where the first column is absorbing: Player 2 controls the transition. As usual the game G^k is chosen with probability p^k and announced to player 1. We follow the approach in [15].

8.2.1. *Asymptotic Analysis and Maxmin*

The analysis is roughly similar in both cases and based on the tools developed for incomplete information games [1]. Let u be the value of the non-revealing game (where player 1 is not transmitting any information on k). A crucial property is that this value does not depend upon the length of the game and then one shows immediately that player 1 can guarantee $\texttt{Cav}u\ (p)$, where $\texttt{Cav}$ denotes the concavification operator on the simplex $\Delta(K)$. Since player 2 has a "non-absorbing" move he can (in the compact case or for the maxmin), knowing σ, observe the variation of the martingale of posterior probabilities on K. Except for a vanishing fraction of stages this variation is small; hence player 1 is almost playing non-revealing so that a best reply of player 2 gives a payoff near u at the current posterior. The result follows by averaging in time and taking expectation, using Jensen's inequality. We thus obtain

Theorem 6

$$maxmin = \lim_{n\to\infty} v_n = \lim_{\lambda\to 0} v_\lambda.$$

8.2.2. *Minmax*

The analysis in this case requires quite specific tools and is related to the question of approachability in stochastic games with vector payoffs. Rather than providing complete proofs which are quite long and painful, we will only give hints concerning the tools used on two examples.

Example 1

$$G^1 = \begin{pmatrix} 1^* & 0 \\ 0^* & 0 \end{pmatrix} \quad G^2 = \begin{pmatrix} 0^* & 0 \\ 0^* & 1 \end{pmatrix}$$

One easily has, with $p = Prob(k = 1)$, that: $u(p) = p(1-p)$; hence

$$\underline{v}(p) = \lim_{n\to\infty} v_n(p) = \lim_{\lambda\to 0} v_\lambda(p) = \texttt{Cav}\ u(p) = p(1-p).$$

However,

$$\overline{v}(p) = p(1 - exp(1 - \frac{(1-p)}{p})),$$

which is obtained as follows. Denote by β_t, $0\le t\le 1$, an ε-optimal strategy of player 2 in the game:

$$\begin{pmatrix} t^* & -t \\ -(1-t)^* & (1-t) \end{pmatrix}$$

(hence absorbing with probability near 1 if the frequency of Bottom exceeds t). Player 2 will choose t according to some distribution ρ and then play β_t. A best reply of player 1 is then to start by playing Top and to decrease slowly his frequency of Top, in order to get an absorbing payoff as high as possible. This leads to the following quantity that player 2 can guarantee:

$$\tilde{v}(p) = \inf_\rho \sup_t \{p \int_0^1 (1-s)\rho(ds) + (1-p)t(1-\rho([0,t])\}.$$

To prove that player 1 can defend $\tilde{v}$ let him construct such a measure ρ starting from the strategy τ of player 2. A discretization will be obtained by playing Bottom with frequency $\frac{\ell}{N}, \ell = 0, ..., N$, for N large. $\tilde{R}(0)$ is thus the probability of absorption given "always Top." It is almost achieved at stage N_0; this defines the quantity $R(0)$. Inductively, $\tilde{R}(\ell)$ is the probability of absorption given the previous strategy until stage $N_{\ell-1}$ and then $(1-\frac{\ell}{N}, \frac{\ell}{N})$ i.i.d. By choosing ℓ and using the associated strategy player 1 can thus achieve $\tilde{v}$.

Example 2

$$G^1 = \begin{pmatrix} 1^* & 0 \\ 0^* & 1 \end{pmatrix} \quad G^2 = \begin{pmatrix} 0^* & 3/4 \\ 1^* & 0 \end{pmatrix}$$

Given a point C in $\mathbb{R}^2$, we say that player 2 can approach C if for any ε there exists τ and N such that for any σ: $\overline{\gamma}_n^k(\sigma, \tau) \leq C^k + \varepsilon$ for $n \geq N$.

Clearly, player 2 can approach $X = (1, 3/7)$ by playing optimally in G^2. He can also approach $Y = (1/2, 3/4)$ by playing a sophisticated optimal strategy in G^1: start as in an optimal strategy in G^1 but control both the absorption probability (q) and the expected absorbing payoff (a) to satisfy $qa + (1 - q) \geq 1/2$: as soon as the opposite equality holds player 2 can play anything in G^1 and get a payoff less than $1/2$, in particular playing optimally in G^2. This allows him to approach Y.

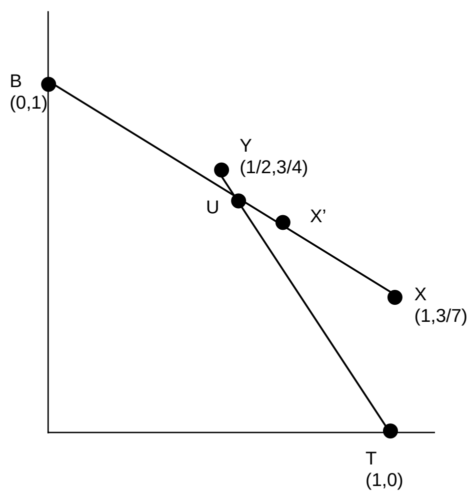

Figure 1. The approachable set

Let $T = (1, 0)$ and $B = (0, 1)$. We will show that player 2 can also approach U, which is the intersection of the segments $[BX]$ and $[TY]$. Note that $U = (1/13)T + (12/13)Y = (6/13)B + (7/13)X$. Player 2 then plays

Left with probability 1/13. If Top is played the absorbing event is $(q, a) = (1/13, T)$; hence it remains to approach Y. Otherwise, the absorbing event is $(q, a) = (1/13, B)$; hence it remains to approach X' with $U = (1/13)B + (12/13)X'$. Now choose a point U' on TY Pareto-dominated by X' and start again. An example of such a procedure is given by:

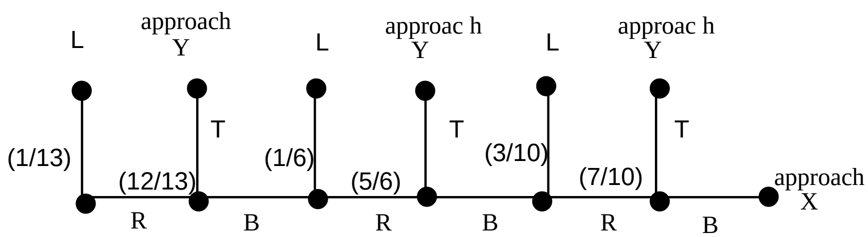

Figure 2. The approachable strategy

As for player 1, by playing Top until exhausting the probability of absorption and then eventually optimally in G^1 he forces a vector payoff minorized by a point Z of the form: $\alpha T + (1 - \alpha)Y$, hence on $[TY]$. Similarly, by playing Bottom and then eventually optimally in G^2, player 1 can "defend" the payoffs above the segment $[XB]$.

Finally, it is easy to see that the set of points that player 2 can approach is convex and that similarly player 1 can defend any convex combination of half-spaces that he can defend.

It follows that the "approachable set" is the set $\mathcal{C}$ of points C with $C^k \geq Z^k$ for some Z in the convex hull of (X, Y, U). Finally, the maxmin $\overline{v}$ is simply the support function of $\mathcal{C}$:

$$\overline{v}(p) = \min_{C \in \mathcal{C}} \langle C, p \rangle.$$

9. Comments

Let us first mention several results related to the current framework:

- Melolidakis [7] gives conditions for v_∞ to exist in games at Level 1 where the transition is independent of the moves of the players.

- Ferguson, Shapley and Weber [5] studied a game with two states where player 2 is informed of the transition on the states only in one direction (from 2 to 1). The natural state space is then the number of stages since the last announcement and the existence of a uniform value is obtained.

- More properties related to games studied in Section 8 can be found in [15], [19], [10].

Absorbing games with incomplete information were introduced as auxiliary tools to study games with incomplete information and state-dependent signalling matrices: this is the case when even by playing independently of his information a player may reveal it. An example is given by the following case, in [16]. The state is (k, ℓ). Player 1 knows k and player 2 knows ℓ. Each game $A^{k\ell}$ is 2×2 and the signalling matrices are as follows.

$$H^{11} = \begin{pmatrix} T & L \\ P & Q \end{pmatrix} \quad H^{12} = \begin{pmatrix} T & R \\ P & Q \end{pmatrix}$$

$$H^{21} = \begin{pmatrix} B & L \\ P & Q \end{pmatrix} \quad H^{22} = \begin{pmatrix} B & R \\ P & Q \end{pmatrix}$$

As soon as player 1 plays Top some game is revealed and one can assume the state absorbing.

Finally, in all the cases studied up to now where player 1 is more informed than player 2, the maxmin is equal to the asymptotic value ($\lim v_n$ and $\lim v_\lambda$), and it is conjectured that this is a general property for this class. More intuition for this to hold can be obtained using the general recursive approach (see [3]).

References

1. Aumann, R.J. and Maschler, M. (1995) *Repeated Games with Incomplete Information* (with the collaboration of R. E. Stearns), MIT Press, Cambridge, MA.
2. Blackwell, D. and Ferguson, T.S. (1968) The big match, *Annals of Mathematical Statistics* **39**, 159–163.
3. Coulomb, J.-M. (2003) Games with a recursive structure, in A. Neyman and S. Sorin (eds.), *Stochastic Games and Applications*, NATO Science Series C, Mathematical and Physical Sciences, Vol. 570, Kluwer Academic Publishers, Dordrecht, Chapter 28, pp. 427–442.
4. De Meyer, B. (1996) Repeated games and partial differential equations, *Mathematics of Operations Research* **21**, 209–236.
5. Ferguson, T., Shapley, L.S. and Weber, R. (2003) Notes on a stochastic game with information structure, *International Journal of Game Theory* **31**, 223–228.
6. Krausz, A. and Rieder, U. (1997) Markov games with incomplete information, *Mathematical Methods of Operations Research* **46**, 263–279.
7. Melolidakis, C. (1989) On stochastic games with lack of information on one side, *International Journal of Game Theory* **18**, 1–29.
8. Melolidakis, C. (1991) Stochastic games with lack of information on one side and positive stop probabilities, in T.E.S. Raghavan et al. (eds.), *Stochastic Games and Related Topics*, Kluwer Academic Publishers, Dordrecht, pp. 113–126.
9. Mertens, J.-F. (1998) The speed of convergence in repeated games with incomplete information on one side, *International Journal of Game Theory* **27**, 343–357.
10. Mertens, J.-F., Sorin, S. and Zamir, S. (1994) Repeated games, CORE Discussion Papers 9420, 9421, 9422, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.
11. Rosenberg, D. (1998) Duality and Markovian strategies, *International Journal of Game Theory* **27**, 577–597.

12. Rosenberg, D. (2000) Zero-sum absorbing games with incomplete information on one side: Asymptotic analysis, *SIAM Journal of Control and Optimization* **39**, 208–225.
13. Rosenberg, D. and Vieille, N. (2000) The maxmin of recursive games with lack of information on one side, *Mathematics of Operations Research* **25**, 23–35.
14. Sorin, S. (1984) "Big match" with lack of information on one side (Part I), *International Journal of Game Theory* **13**, 201–255.
15. Sorin, S. (1985) "Big match" with lack of information on one side (Part II), *International Journal of Game Theory* **14**, 173–204.
16. Sorin, S. (1985) On a repeated game with state-dependent signalling matrices, *International Journal of Game Theory* **14**, 249–272.
17. Sorin, S. (2003) The operator approach to zero-sum stochastic games, in A. Neyman and S. Sorin (eds.), *Stochastic Games and Applications*, NATO Science Series C, Mathematical and Physical Sciences, Vol. 570, Kluwer Academic Publishers, Dordrecht, Chapter 27, pp. 417–426.
18. Sorin, S. (2003) Symmetric incomplete information games as stochastic games, in A. Neyman and S. Sorin (eds.), *Stochastic Games and Applications*, NATO Science Series C, Mathematical and Physical Sciences, Vol. 570, Kluwer Academic Publishers, Dordrecht, Chapter 21, pp. 323–334.
19. Sorin, S. and Zamir, S. (1991) "Big match" with lack of information on one side III, in T.E.S. Raghavan et al. (eds.), *Stochastic Games and Related Topics*, Kluwer Academic Publishers, Dordrecht, pp. 101–112.
20. Thuijsman, F. (2003) Recursive games, in A. Neyman and S. Sorin (eds.), *Stochastic Games and Applications*, NATO Science Series C, Mathematical and Physical Sciences, Vol. 570, Kluwer Academic Publishers, Dordrecht, Chapter 16, pp. 253–264.
21. Thuijsman, F. (2003) The big match and the Paris match, in A. Neyman and S. Sorin (eds.), *Stochastic Games and Applications*, NATO Science Series C, Mathematical and Physical Sciences, Vol. 570, Kluwer Academic Publishers, Dordrecht, Chapter 12, pp. 195–204.