The value of Markov chain games with lack of information on one side

Jérôme Renault

CEREMADE, Université Paris Dauphine, Place du Maréchal de Lattre, 75776 Paris Cedex 16, France.

email: renault@ceremade.dauphine.fr   http://www.ceremade.dauphine.fr/~renault/

We consider a two-player zero-sum game, given by a Markov chain over a finite set of states and a family of matrix games indexed by states. The sequence of states follows the Markov chain. At the beginning of each stage, only player 1 is informed of the current state, then the corresponding matrix game is played and the actions chosen are observed by both players before proceeding to the next stage. We call such a game a Markov chain game with lack of information on one side. This model generalizes the model of Aumann and Maschler of zero-sum repeated games with lack of information on one side (which corresponds to the case where the transition matrix of the Markov chain is the identity matrix). We generalize the proof of Aumann and Maschler and, from the definition and the study of appropriate "non revealing" auxiliary games with infinitely many stages, show the existence of the uniform value. An important difference with Aumann and Maschler's model is that here, the notions for player 1 of using the information and revealing a relevant information are distinct.

1. Introduction. Repeated games with incomplete information were introduced in the sixties by Aumann and Maschler [3] (for a reedition of their work). Their simplest model is the one of two-player zero-sum repeated game with lack of information on one side and perfect observation: at the beginning of the play, a state of nature k is chosen according to an initial probability p over a finite set of states K, and announced to player 1 only. This state of nature determines a finite zero-sum game G^k which is then repeatedly played over and over, and after each stage the actions played are observed by both players. In this setup, Aumann and Maschler showed the existence of the value, and gave a famous characterization for it. They also proved the existence of optimal strategies for the players. This pioneering work has been widely extended since the sixties, and a lot of important works are dedicated to this model, or close extensions of it. For example, let us mention Aumann *et al.* [3] or Mertens and Zamir [10] for the case of incomplete information on both sides, Kohlberg [7] for an explicit construction of an optimal strategy for player 2, or Sorin [17] Hart [5] and Simon *et al.* [15] for the non-zero sum case.

We only consider here two player zero-sum games and generalize the above model to the case where the state of nature is no longer fixed once and for all at the beginning of the game, but evolves according to a given Markov chain, and at the beginning of each stage is observed by player 1 only. We call such games Markov chain games with lack of information on one side, since here also only player 2's information is incomplete. Note that this model is a special case of stochastic game with incomplete information, which may fail to have a value (Sorin [18], see also Rosenberg and Vieille [14]). Recently, Rosenberg *et al.* [13] proved the existence of the value in a particular class of games generalizing Aumann and Maschler's model. These games are defined via a collection of stochastic games : initially nature chooses a stochastic game, and player 1 only is informed of the selected game. Only one player controls the transition, and these games are neither more nor less general than the ones studied here.

In this paper, we generalize the proof of Aumann and Maschler and show the existence of the value in Markov chain games with lack of information on one side. However, several problems remain. Our expression of the value cannot be easily computed from the basic data of the game, and we present a very simple example where we cannot compute it. About optimal strategies, we obtain the existence of an optimal strategy for player 2, but we still do not know if there exists such a strategy for player 1.

A. Neyman and an anonymous associate editor suggested the possibility that there exists an alternative proof of the existence of the value based on the application of the result of Aumann and Maschler to our framework.

Section 2 of this paper contains the model. A few important examples, and the main ideas of the proof, are presented in section 3. In section 4 we formally introduce a projection matrix which will be used to quantify a notion of "relevant information" for player 2. In section 5 we study the values of the N-stage

games and show that player 2 can guarantee their limit v^*. In section 6 we define and study a notion of "non revealing" strategies for player 1. In section 7 we define auxiliary games called non revealing games, where player 1 is restricted to play a non revealing strategy. The values of the non revealing games are used in section 8 to show that player 1 can guarantee some quantity w^* in the original game. In section 9 we show that $v^* = w^*$, and this implies the existence of the value. We conclude with several remarks on the model and possible extensions. A few proofs can be found in the Appendix.

2. The model. If $\mathcal{S}$ is a finite set, $|\mathcal{S}|$ denotes its cardinality and $\Delta(\mathcal{S})$ the set of probability distributions over $\mathcal{S}$. $\Delta(\mathcal{S})$ is viewed as a subset of $\mathbb{R}^{\mathcal{S}}$. For $q = (q^s)_{s \in \mathcal{S}}$ in $\mathbb{R}^{\mathcal{S}}$, we use $\|q\| = \sum_{s \in \mathcal{S}} |q^s|$.

We denote by $K = \{1, ..., |K|\}$ the set of states, by I the set of actions of player 1 and by J the set of actions of player 2. K, I and J are assumed to be finite and non empty. We have a family $(G^k)_{k \in K}$ of $|I| \times |J|$ payoff matrices for player 1, and a Markov chain on K, given by an initial probability p on K and a transition matrix $M = (M_{kk'})_{(k,k') \in K \times K}$. All elements of M are non negative and for each k in K, $\sum_{k' \in K} M_{kk'} = 1$.

An element $q = (q^k)_{k \in K}$ in $\Delta(K)$ will also be represented by a row vector $q = (q^1, ..., q^{|K|})$, with $q^k \geq 0$ for each k and $\sum_{k \in K} q^k = 1$. If the law of the state at some stage is q, the law of the state at the next stage is the product qM. For each k in K, δ_k denotes the Dirac measure on k.

The play of the zero-sum game is as follows:

-at stage 1, k_1 is chosen according to p, and told to player 1 only. Players 1 and 2 independently choose an action in their own set of actions, $i_1 \in I$ and $j_1 \in J$ respectively. The stage payoff for player 1 is $G^{k_1}(i_1, j_1)$, (i_1, j_1) is publicly announced, and the play proceeds to stage 2.

- at stage $n \geq 2$, k_n is chosen according to $\delta_{k_{n-1}}M$, and told to player 1 only. The players independently choose an action in their own set of actions. If $i_n \in I$ and $j_n \in J$ are selected, the stage payoff for player 1 is $G^{k_n}(i_n, j_n)$. (i_n, j_n) is publicly announced, and the play proceeds to the next stage.

Note that the payoffs are not announced after each stage. Players are assumed to have perfect recall, and the whole description of the game is public knowledge. It will be convenient to keep the initial probability on K as a parameter, so we denote by $\Gamma_\infty(p)$ the game just defined.

A behavior strategy for player 1 is an element $\sigma = (\sigma_n)_{n \geq 1}$ where for each n σ_n is a mapping from the cartesian product $(K \times I \times J)^{n-1} \times K$ to $\Delta(I)$ giving the mixed action played by player 1 at stage n depending on past and current states and past actions played. Since player 2 does not observe the states, a behavior strategy for him is an element $\tau = (\tau_n)_{n \geq 1}$, where for each n τ_n is a mapping from $(I \times J)^{n-1}$ to $\Delta(J)$. Denote by Σ and $\mathcal{T}$, respectively, the set of behavior strategies of player 1 and player 2. A strategy profile (σ, τ) induces a probability distribution over the set of plays $\Omega = (K \times I \times J)^\infty$, and we denote, for each positive N, the average expected payoff for player 1 induced by (σ, τ) at the first N stages by:

$$\gamma_N^p(\sigma, \tau) = \mathbb{E}_{p,\sigma,\tau}\left(\frac{1}{N}\sum_{n=1}^{N} G^{k_n}(i_n, j_n)\right)$$

where k_n, i_n, j_n respectively denote the state, action for player 1 and action for player 2 at stage n.

We will also consider explicitly strategies in games with finitely many stages. For $N \geq 1$, Σ_N will denote the set of N-stage behavior strategies of player 1. An element σ in Σ_N is just an element $(\sigma_n)_{n \in \{1,...,N\}}$, where for each n, σ_n is a mapping from $(K \times I \times J)^{n-1} \times K$ to $\Delta(I)$. Similarly, $\mathcal{T}_N$ will denote the set of N-stages strategies for player 2. The definition of the payoff $\gamma_N^p(\sigma, \tau)$ trivially extends to the case where σ and τ are strategies with finitely many stages. The N-stage game $\Gamma_N(p)$ is defined as the zero-sum game with strategy spaces Σ_N and $\mathcal{T}_N$, and payoff function γ_N^p. By Kuhn's theorem about the equivalence between mixed and behavioral strategies, $\Gamma_N(p)$ can be viewed as a mixed extension of a finite game, and hence it has a value denoted by $v_N(p)$, and both players have optimal strategies. Note that $v_N(p)$ is also the value of the zero-sum game with strategy spaces Σ and $\mathcal{T}$, and payoff function γ_N^p.

Concerning the infinite game $\Gamma_\infty(p)$, we will use the following standard notion of uniform value.

Definition 2.1 *Let v be a real number.*

- Player 1 can guarantee v in $\Gamma_\infty(p)$ if:

$\forall \varepsilon > 0 \ \exists \sigma \in \Sigma \ \exists N_0, \ \forall N \geq N_0 \ \forall \tau \in \mathcal{T}, \quad \gamma_N^p(\sigma, \tau) \geq v - \varepsilon.$

- Player 2 can guarantee v in $\Gamma_\infty(p)$ if:

$\forall \varepsilon > 0 \ \exists \tau \in \mathcal{T} \ \exists N_0, \ \forall N \geq N_0 \ \forall \sigma \in \Sigma, \quad \gamma_N^p(\sigma, \tau) \leq v + \varepsilon.$

- v is the value of $\Gamma_\infty(p)$ if both players can guarantee v in $\Gamma_\infty(p)$.

Notice that if player 1 (resp. player 2) can guarantee v, then $\liminf_N v_N(p) \geq v$ (resp. $\limsup_N v_N(p) \leq v$). It is then easy to see that if the value v of $\Gamma_\infty(p)$ exists, it is necessarily unique and equal to $\lim_N v_N(p)$ (one can show that it is also the limit of the value of the discounted games as the discount factor goes to zero, see e.g. Sorin [19], lemma 3.1 p.27).

Definition 2.2 *Assume that the value v of $\Gamma_\infty(p)$ exists.*

A strategy σ of player 1 is said to be optimal if:
$$\forall \varepsilon > 0 \ \exists N_0, \ \forall N \geq N_0 \ \forall \tau \in \mathcal{T}, \quad \gamma_N^p(\sigma, \tau) \geq v - \varepsilon.$$

Similarly, a strategy τ of player 2 is optimal if:
$$\forall \varepsilon > 0 \ \exists N_0, \ \forall N \geq N_0 \ \forall \sigma \in \Sigma, \quad \gamma_N^p(\sigma, \tau) \leq v + \varepsilon.$$

In this paper we prove the following result.

Theorem 2.3 $\Gamma_\infty(p)$ *has a value, and player 2 has an optimal strategy.*

We will use the following notations and definitions.

If $x = (x(i))_{i \in I} \in \Delta(I)$ and $y = (y(j))_{j \in J} \in \Delta(J)$, for each state k we denote by $G^k(x,y) = \sum_{i \in I, j \in J} x(i)y(j)G^k(i,j)$ the expected payoff for player 1 in state k if i in I and j in J are selected according to the probabilities x and y, respectively. For each p in $\Delta(K)$, $u(p)$ will denote the value of the matrix game $\sum_{k \in K} p^k G^k$, i.e. :

$$u(p) = \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} \sum_{k \in K} p^k G^k(x,y).$$

It will turn out that player 1 never needs to take into account the actions played by player 2, so for simplicity we will often consider the following type of strategies. A strategy $\sigma = (\sigma_n)_n$ of player 1 (in Σ or in Σ_N) will be called independent of the actions of player 2 if for each n and $(k_1, i_1, j_1, ..., k_{n-1}, i_{n-1}, j_{n-1}, k_n)$ in $(K \times I \times J)^{n-1} \times K$, $\sigma_n(k_1, i_1, j_1, ..., k_{n-1}, i_{n-1}, j_{n-1}, k_n)$ does not depend on $(j_1, ..., j_{n-1})$. We denote by Σ^+ (or Σ_N^+ for N-stages strategies) the set of such strategies of player 1.

Finally, we denote by C an upper bound for all absolute values of payoffs, i.e. C satisfies: $|G^k(i,j)| \leq C$ for all k, i, j. And if $f : \Delta(K) \longrightarrow I\!\!R$ is bounded from above, cav$f$ denotes the pointwise smallest concave function g on $\Delta(K)$ satisfying $g(p) \geq f(p)$ for each p in $\Delta(K)$.

3. Examples and overview of the proof. We first notice the following important fact. Assume that player 1 does not use his information on the states, i.e. plays some strategy which is independent of the sequence of states. Then the belief of player 2 on the current state of stage $n+1$ (i.e. the conditional probability on the state of stage $n+1$ given the initial probability p, the strategy used by player 1 and the actions played up to stage n) is the unconditional probability over the state of stage $n+1$, so it is just pM^n. We start with examples.

Example A: A Markov chain without memory

Consider the following transition matrix (there are two states): $M = \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}$.

At each stage $n \geq 2$, the new state is chosen according to the probability $p^* = (1/2, 1/2)$ on $\Delta(K)$, independently of the previous state. Player 1 can use at each stage his knowledge of the current state, and player 2 will never learn anything relevant about future states. For any initial probability p, each stage $n \geq 2$ of the repeated game $\Gamma_\infty(p)$ is similar to the 1-stage game $\Gamma_1(p^*)$, hence $\Gamma_\infty(p)$ has a value which is $v_1(p^*)$.

A more complex irreducible aperiodic Markov chain is now presented.

Example B: An irreducible aperiodic Markov chain

The transition matrix is now: $M = \begin{pmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{pmatrix}$.

This Markov chain has a unique recurrence class which is aperiodic, and has $p^* = (1/2, 1/2)$ as unique invariant measure. $(M^n)_n$ converges to $\begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}$, hence for any p in $\Delta(K)$, $(pM^n)_n$ converges to p^*. This implies that if player 1 plays independently of the states, the belief of player 2 on the current state converges to p^*.

For any N player 1 can play in $\Gamma_\infty(p)$ as follows: (1) play for N stages an optimal strategy in $\Gamma_N(p^*)$, then (2) play independently of the states for a fixed number of stages (this number becoming much smaller than N when N is large), and come back to (1). This shows that player 1 can guarantee $\limsup_N v_N(p^*)$. Concerning player 2, recall that $pM^n \longrightarrow_{n \to \infty} p^*$. For any N, he can play as follows in $\Gamma_\infty(p)$: play for N stages an optimal strategy in $\Gamma_N(p^*)$, then forget everything and redo from start. This shows that player 2 can guarantee $\inf_N v_N(p^*)$ in $\Gamma_\infty(p)$. And we can conclude that $\Gamma_\infty(p)$ has a value which is $\inf_N v_N(p^*)$, hence independent of p.

Notice that this proof can be generalized to the case of any irreducible aperiodic Markov chain.

Example C: M is the identity matrix.

This is the standard case of repeated game with lack of information on one side. In this case, the initial state is selected at the beginning of the game and remains constant. It is well known (see Aumann Maschler [3]) that the value $v(p)$ of $\Gamma_\infty(p)$ exists and satisfies:

$$v(p) = \mathrm{cav} u(p).$$

The proof of Aumann and Maschler can be sketched as follows.

1) For each N player 2 can play as follows in $\Gamma_\infty(p)$: play for N stages an optimal strategy in $\Gamma_N(p)$, then forget everything and redo from start. We obtain that player 2 can guarantee $\inf_N v_N(p)$, so $\lim_N v_N(p)$ exists and can be guaranteed by player 2.

2) Aumann and Maschler defined, for each initial probability $p = (p^k)_{k \in K}$, a "non revealing" game corresponding to the game where player 1 does not use his information on the selected state, i.e. where he plays independently of the state. This game can be analyzed as the repetition of the average matrix game $\sum_{k \in K} p^k G^k$. Consequently, its value is just the value $u(p)$ of the average game and player 1 can guarantee this quantity by playing independently of the states. By a property, which is now standard (see for example [19, prop. 2.2. p.16]), of games with lack of information on one side, one can then deduce that player 1 can guarantee $\mathrm{cav} u(p)$ in the original game.

3) The last point is to show that $\lim_N v_N(p) \leq \mathrm{cav} u(p)$. This is done by considering, for a fixed strategy of player 1, the a posteriori of player 2 on the selected state of nature (i.e. after each stage, the belief on the state depending on the previous actions played). The sequence of a posteriori forms a martingale, and a classical bound on its L_1 variation gives the above inequality.

We insist on the fact that here, the non revealing games correspond to the game where player 1 does not use his information on the states. The relevant information for player 2 is in this case the initial state, or equivalently the recurrence class of the Markov chain. We will generalize this proof to the general case.

Remark 3.1 *Aumann and Maschler [3] also generalize their result to the case where each player observes after each stage a signal which is function of the actions played. In this setup, they already notice the*

difference betwween using and revealing information and they define a non revealing strategy of player 1 as a strategy that keeps unchanged the belief of player 2 on the state.

Example D: A periodic chain

Let K be $\{a, b\}$, and M be $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$.

The Markov chain has a unique recurrence class, which is periodic with period 2. Here the relevant information for player 2 is not the recurrence class, but if the sequence of states will be $(a, b, a, b, a, b, ...)$, or $(b, a, b, a, b, a, ...)$. By considering the play of stages by blocks of length 2, we can reduce the problem here to that of example C. The point is that M^2 is the identity matrix.

This idea of playing stages by blocks of fixed length could be used in the general case to assume w.l.o.g. that the chain is aperiodic. We will however not assume aperiodicity in the following proofs, because it would simplify almost nothing.

A last aspect of finite Markov chains is the possible existence of transient states. This does not play an important role here, because player 1 can always wait during the first stages until the current state is a recurrence one. We however present now an example with a transient state. This example, and our last example below, will be used later to illustrate several definitions.

Example E: $K = \{a, b, c\}$, and $M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1/2 & 1/4 & 1/4 \end{pmatrix}$.

Example F: $K = \{a, b, c\}$, and $M = \begin{pmatrix} 2/3 & 1/3 & 0 \\ 1/3 & 2/3 & 0 \\ 0 & 0 & 1 \end{pmatrix}$.

To prove the existence of the value in case of example F is more difficult, and we do not know how to proceed without following the general proof presented in this paper. Notice that we have here two kinds of information: a long-term information (corresponding to the recurrence class $\{a, b\}$ or $\{c\}$, as in the Aumann-Maschler case), and a short-term information (corresponding to the state a or b if we are in the class $\{a, b\}$). We now come back to the general case and describe the different steps of our proof.

Overview of the proof:

One can show, as in example B, that player 2 can guarantee the limit of $(v_N(p))_N$. This will be done in section 5 and not discussed now.

The main conceptual difficulty is to generalize point 2) of the proof of Aumann and Maschler. We will define an appropriate version of the non revealing games. In the case of example A, at each stage player 1 can use his knowledge of the current state without revealing anything important. The use of information is "free" for player 1, since information will be generated anew. On the contrary, in the Aumann and Maschler's case, as soon as player 1 uses some piece of information, the martingale of a posteriori moves and a relevant piece of information is revealed to player 2. In the general case, we will define and quantify the information of player 1 which is relevant to player 2 *in the long run*. The basic and main idea is that this information for player 2 can be described by the asymptotic behavior of the sequence of states.

Recall that if player 1 does not use his information on the states, the belief of player 2 on the current state of stage $n + 1$ is pM^n. An important property of a finite Markov chain is that the limit behavior of the sequence $(pM^n)_n$ can always be approximated by a periodic sequence $(pB_0, pB_0M, ..., pB_0M^{L-1}, pB_0, pB_0M, ...)$, L being a common multiple of the periods of the recurrence classes of the Markov chain, and B_0 being the limit, as n goes to infinity, of $(M^{nL})_n$. B_0 is unique, it is a stochastic matrix, and also a projection matrix ($B_0^2 = B_0$). This matrix takes into account the recurrence classes of M, but also the periodic aspect of these classes, and its existence is fundamental for our proof. In particular, if player 1 never uses his information the relevant information for player 2 can be summarized by $\lim_n pM^{nL} = pB_0$, since this characterizes the asymptotic periodic sequence. And

if player 1 does not use his information after some stage n, the relevant information for player 2 about the limit behavior of the sequence of states is $q_n = p_n B_{-n}$, where p_n is player 2's belief, computed after stage n, on the state of stage $n+1$, and $B_{-n} = \lim_t M^{tL-n}$.

Fix now any strategy of player 1. $(q_n)_n$ will turn out to be a martingale representing player 2's current relevant information over the limit behavior of the sequence of states (see definition 6.1 in section 6 for a formal definition of q_n). q_n will simply be called player 2's relevant information after stage n, and we define a non revealing strategy for player 1 at p as a strategy such that this martingale is almost surely constant (and, for convenience, which plays independently of the actions of player 2). So a non revealing strategy of player 1 does not necessarily play independently of the states, as in Aumann and Maschler's case. This is especially the case when the Markov chain has a unique recurrence class which is aperiodic: if player 1 chooses his actions depending on the current state for a certain number of stages, and then play independently of the states, all the information obtained by player 2 will vanish since his belief on the current state will still converge to the unique invariant probability measure. So a non revealing strategy of player 1 may give some information about the states, *in the short term*, to player 2. But it will give no relevant information in the long term to player 2. For player 1, the two notions of not using the information and not revealing a relevant information are different here.

We define the non revealing game at p, denoted by $\hat{\Gamma}_\infty(p)$, as the auxiliary game where player 1 is restricted to play a non revealing strategy at p. This auxiliary game has infinitely many stages, and cannot be analyzed as a repeated matrix game as in Aumann and Maschler's case. Indeed, in the case of a recurrent aperiodic Markov chain, all strategies of player 1 are non revealing and the original game $\Gamma_\infty(p)$ and the non revealing game $\hat{\Gamma}_\infty(p)$ coincide. Section 6 is devoted to the definition and a careful study of non revealing strategies. Section 7 deals with non revealing games, and first consists of a technical generalization of section 5 to auxiliary games where the set of strategies of player 1 is restricted in some way (which can be: no restriction at all). We notably prove a version, adapted to non revealing games, of Aumann and Maschler's "splitting lemma", and a recursive formula for the values $\hat{v}_N(p)$ of the N-stages non revealing games. At the end of section 7, we prove, using the definition of non revealing strategies, that player 1 can guarantee some quantity in the non revealing game, and then obtain that $\hat{\Gamma}_\infty(p)$ has a value $w(p)$ which satisfies: $w(p) = \inf_{N\geq 1}\hat{v}_{NL}(pB_0)$. Since player 1 can guarantee $w(p)$ in the non revealing game, *a fortiori* he can guarantee $w(p)$ in the original game $\Gamma_\infty(p)$, and we will show in section 8 that player 1 can guarantee $\text{cav}w(pB_0)$ in $\Gamma_\infty(p)$.

It will remain in section 9 to generalize point 3) of the proof of Aumann and Maschler. Indeed, our definitions of relevant information and non revealing games can only be justified if we link the values of the original N-stage games to the ones of the non revealing games. Let's consider a LNT-stage (original) game, with N fixed and T large, and fix any strategy of player 1. For each block $t = 1, ..., T$ of LN stages, we consider an approximation of player 1's strategy on this block as a non revealing one, and then we consider a best reply of player 2 against this approximation in a corresponding non revealing game with LN stages. One can then show from the bound on the L^1- variation of the martingale $(q_n)_n$ that $\lim_n v_n(p) \leq \text{cav}\hat{v}_{NL}(pB_0)$. The rest is mainly technical. We show that $\text{cav}w(pB_0) = \inf_{N\geq 1}\text{cav}\hat{v}_{NL}(pB_0)$, so $\text{cav}w(pB_0) \geq \lim_N v_N(p)$. Since in the original game with initial probability p, player 2 can guarantee $\lim_N v_N(p)$ and player 1 can guarantee $\text{cav}w(pB_0)$, we are done and the value of $\Gamma_\infty(p)$ is $\text{cav}w(pB_0) = \lim_N v_N(p)$.

4. Limit beliefs over the current state of an unobserved Markov chain. An important property of finite Markov chains is that the sequence of unconditional beliefs $(p, pM, pM^2, ..., pM^n, ...)$ can be approximated by a periodic sequence $(pB_0, pB_1, ..., pB_{L-1}, pB_0, pB_1, ...)$ that we define now.

Proposition 4.1 *There exists a positive integer L and stochastic matrices B_0, B_1,..., B_{L-1} such that:*

$$\forall l \in \{0, ..., L-1\}, \quad \lim_{n\to\infty} M^{nL+l} = B_l.$$

The proof is omitted and can be easily deduced, e.g., from Gordon [4], or Norris [11, Theorem 1.8.5, p.44]. The only thing to show is the existence of some L such that $(M^{nL})_n$ converges. From a linear algebra viewpoint, this is due to the fact that if z is a complex eigenvalue of M with $|z| \geq 1$, then there exists a positive integer n such that $z^n = 1$. L corresponds to a common multiple of all such n, or equivalently, to a common multiple of the periods of the recurrence classes of M.

As $(M^{nL})^2 \longrightarrow_{n \to \infty} B_0^2$, B_0 is a projection matrix, that is $B_0^2 = B_0$. Notice that B_0 is necessarily unique. The set of invariant probability measures of B_0 will play an important role, hence we introduce the following notations.

Notations:

$$Q = \{q \in \Delta(K), qB_0 = q\} = \{pB_0, p \in \Delta(K)\}$$
$$\text{For each } q \text{ in } Q, \quad A(q) = \{p \in \Delta(K), pB_0 = q\}.$$

If $q = qB_0$, then $qM^L = qB_0M^L = qB_0 = q$, hence it is easy to see that Q also is the set of invariant probability measures of the matrix M^L. Notice also that Q and each set $A(q)$ are polytopes.

It is convenient to define B_n for each integer n. We put, for each (possibly negative) integer n, $B_n = B_l$, where $l \in \{0, ..., L-1\}$ and $n - l$ is a multiple of L. We then have for each integer n, $M^{tL+n} \longrightarrow_{t \to \infty} B_n$, and:

$$\forall p \in \mathbb{N}, \quad M^p B_n = B_n M^p = B_{n+p}.$$

5. N-stage games. For each positive integer N and each p in $\Delta(K)$, $\Gamma_N(p)$ is the N-stage game with initial probability p, and has a value $v_N(p)$. The payoff function γ_N^p satisfies:

$$\forall \sigma \in \Sigma, \ \forall \tau \in \mathcal{T}, \quad \gamma_N^p(\sigma, \tau) = \sum_{k \in K} p^k \gamma_N^{\delta_k}(\sigma, \tau).$$

Consequently, v_N is, as a function of p, Lipschitz with constant C. Moreover it is concave, since $\Gamma_N(p)$ is a game with incomplete information on one side with player 1 as the informed player (see for example Zamir [20]). In order to prove a recursive formula for v_N, we first need some notations.

Let p be in $\Delta(K)$ representing player 2's belief on the current state at some stage $n \geq 1$, this belief being computed before stage n is played. Assume that player 1's action at this stage will be chosen according to some $x = (x^k)_{k \in K} \in \Delta(I)^K$, i.e. that player 1 plays, for each k, according to $x^k = (x^k(i))_{i \in I}$ if the current state is k. The probability that player 1 plays at stage n some action i in I will be denoted by:

$$x(p)(i) = \sum_{k \in K} p^k x^k(i).$$

For each i in I, the conditional probability on the state of stage n given that player 1 has played i at this stage is denoted by $\hat{p}(x, i) \in \Delta(K)$. We have:

$$\hat{p}(x, i) = \left(\frac{p^k x^k(i)}{x(p)(i)} \right)_{k \in K}$$

(if $x(p)(i) = 0$, $\hat{p}(x, i)$ is defined arbitrarily in $\Delta(K)$.) Obviously, $\sum_{i \in I} x(p)(i)\hat{p}(x, i) = p$. And the expected stage payoff for player 1 is, if player 2 plays according to y in $\Delta(J)$:

$$G(p, x, y) = \sum_{k \in K} p^k G^k(x^k, y).$$

We can now state the recursive formula, where v_0 is defined arbitrarily. The proof is very standard and can be found in the Appendix. It will be generalized in section 7 to non revealing games (for a general recursive formula in repeated games, see thm 3.2 p.187 in Mertens *et al.*).

Proposition 5.1 *For each $n \geq 1$ and p in $\Delta(K)$,*

$$v_n(p) = \frac{1}{n} \max_{x \in \Delta(I)^K} \min_{y \in \Delta(J)} \left(G(p, x, y) + (n-1) \sum_{i \in I} x(p)(i) \, v_{n-1}(\hat{p}(x, i)M) \right)$$

$$= \frac{1}{n} \min_{y \in \Delta(J)} \max_{x \in \Delta(I)^K} \left(G(p, x, y) + (n-1) \sum_{i \in I} x(p)(i) \, v_{n-1}(\hat{p}(x, i)M) \right)$$

In $\Gamma_n(p)$, player 1 has a strategy $\sigma = (\sigma_n)_{n \geq 1}$ such that at each stage n', $\sigma_{n'}$ only depends on player 1's past actions $i_1, ..., i_{n'-1}$ and on the current state $k_{n'}$.

Remark 5.2 $v_n(p)$ *is also the value of the n-stage stochastic game where the set of states is $\Delta(K)$, player 1's set of actions is $\Delta(I)^K$, player 2's set of actions is $\Delta(J)$, the stage payoff is given by $((p, x, y) \longmapsto G(p, x, y))$, and the transition between states is controlled by player 1 only: if the state is p, and player 1 plays x, the new state is $\hat{p}(x, i)M$ with probability $x(p)(i)$. In this stochastic game, player 1 has an optimal strategy which is a Markovian strategy: the action played at each stage only depends on the current state and on the stage number.*

Corollary 5.3 $\forall p \in \Delta(K)$, $v_n(p) - v_{n-1}(pM) \longrightarrow_{n \to \infty} 0$.

Proof: Since v_{n-1} is concave, we have for each p: $v_n(p) \leq \frac{1}{n}(v_1(p) + (n-1)v_{n-1}(pM))$. On the other hand, player 1 may play independently of the state. By taking $x \in \Delta(I)^K$ such that $x^k = x^{k'}$ for all k and k', we get: $v_n(p) \geq \frac{1}{n}(u(p) + (n-1)v_{n-1}(pM))$. Since payoffs are uniformly bounded, $v_n(p) - v_{n-1}(pM) \longrightarrow_{n \to \infty} 0$. □

We now define a function v^* on $\Delta(K)$ which will turn out to be the value of $\Gamma_\infty(p)$.

Definition 5.4 $\forall p \in \Delta(K)$, $v^*(p) = \inf_{N \geq 1} v_{NL}(pB_0)$.

v^* is concave as an infimum of concave functions. For each p, $v^*(p) = v^*(pB_0)$, hence for each q in Q the restriction of v^* to $A(q)$ is constant. The next proposition is inspired by example B.

Proposition 5.5 *For each p in $\Delta(K)$, player 2 can guarantee $v^*(p)$ in $\Gamma_\infty(p)$.*

Proof: Fix p in $\Delta(K)$, $N \geq 1$, and let τ_{NL} be an optimal strategy for player 2 in $\Gamma_{NL}(pB_0)$.

We divide the set of stages $\{1, 2, ..., n, ...\}$ into consecutive blocks B^1, B^2, ..., B^m,... of length NL. Define the strategy τ of player 2 in $\Gamma_\infty(p)$ as follows: at each block B^m, play according to τ_{NL} (and forget everything that has happened at previous blocks). For each m, B^m begins at stage $(m-1)NL+1$, and the (unconditional) probability on the state at this stage is $pM^{(m-1)NL}$.

Since $|\gamma_{NL}^{pB_0}(\sigma', \tau') - \gamma_{NL}^{pM^{(m-1)NL}}(\sigma', \tau')| \leq \|pB_0 - pM^{(m-1)NL}\|C$ for each strategy pair (σ', τ'), we have that for any strategy σ in Σ,
$$\mathbb{E}_{p,\sigma,\tau}\left(\frac{1}{NL}\sum_{n \in B^m} G^{k_n}(i_n, j_n)\right) \leq v_{NL}(pB_0) + C\|pB_0 - pM^{(m-1)NL}\|.$$
As $(M^{nL})_n$ converges to B_0, we obtain that: $\forall \varepsilon > 0 \; \exists N_0 \; \forall N_1 \geq N_0, \forall \sigma \in \Sigma \;\; \gamma_{N_1}^p(\sigma, \tau) \leq v_{NL}(pB_0) + \varepsilon$. Hence player 2 can guarantee $v_{NL}(pB_0)$ in $\Gamma_\infty(p)$. □

Corollary 5.6 $(v_n)_{n \geq 1}$ *uniformly converges to v^* on $\Delta(K)$. v^* is Lipschitz with constant C and satisfies $v^*(p) = v^*(pM)$ for each p in $\Delta(K)$.*

Proof: From the definition of guaranteeing, proposition 5.5 implies that for all p in $\Delta(K)$:
$$\limsup_n v_n(p) \leq \inf_{n \geq 1} v_{nL}(pB_0) \leq \liminf_n v_{nL}(pB_0).$$
Since $B_0^2 = B_0$, we obtain for all p in $\Delta(K)$:
$$\limsup_n v_n(pB_0) \leq \inf_{n \geq 1} v_{nL}(pB_0) \leq \liminf_n v_{nL}(pB_0).$$
So $(v_{nL}(pB_0))_n$ converges to $\inf_{n \geq 1} v_{nL}(pB_0) = v^*(p)$. Since it is clear from the definition of γ_n^p that $v_n(p) - v_{n+1}(p) \longrightarrow_{n \to \infty} 0$, $(v_n(pB_0))_n$ converges to $v^*(p)$.

Put $\underline{v}(p) = \liminf_n v_n(p)$ and $\bar{v}(p) = \limsup_n v_n(p)$. By corollary 5.3, $\underline{v}(p) = \underline{v}(pM)$ and $\bar{v}(p) = \bar{v}(pM)$. Moreover $\underline{v}$ and $\bar{v}$ are Lipschitz with constant C, hence continuous. Since $(M^{nL})_n$ converges to B_0, we have $\underline{v}(p) = \underline{v}(pB_0) = v^*(p) = \bar{v}(pB_0) = \bar{v}(p)$. And $(v_n(p))_n$ converges to $v^*(p)$.

Since $(v_n)_n$ is a sequence of continuous concave functions (pointwise) converging to v^* which is continuous and concave, and since $\Delta(K)$ is a polytope, $(v_n)_n$ uniformly converges to v^* (see for example Mertens *et al* , part A, p.46, ex. 15). □

Remark 5.7 *p being fixed, proposition 5.5 can be stated as: $\forall \varepsilon > 0 \ \exists \tau \in \mathcal{T} \ \exists N_0, \ \forall N \geq N_0 \ \forall \sigma \in \Sigma$, $\gamma_N^p(\sigma, \tau) \leq v^*(p) + \varepsilon$. It can be strengthened in order to get the existence of a single strategy τ of player 2 such that : $\forall \varepsilon > 0 \ \exists N_0, \ \forall N \geq N_0 \ \forall \sigma \in \Sigma, \quad \gamma_N^p(\sigma, \tau) \leq v^*(p) + \varepsilon$. To construct such τ, modify slightly the strategy presented in proposition 5.5 as follows. Divide the set of stages into consecutive blocks $B^1, ..., B^m, ...,$ where for each m, B^m has cardinality mL. At each block B^m, play according to an optimal strategy for player 2 in $\Gamma_{mL}(pB_0)$. As $v_{mL}(pB_0) \longrightarrow_{m \to \infty} v^*(p)$, τ has the required property. Since we will finally show that $v^*(p)$ is the value of $\Gamma_\infty(p)$, this will imply that τ is an optimal strategy of player 2 in $\Gamma_\infty(p)$.*

Remark 5.8 *It is known (see e.g. [19]) that in Aumann and Maschler's case, $(v_N(p))_{N \geq 1}$ is non-increasing. This property is no longer satisfied here, and in point 4 of the last section we give an example where $(v_{NL}(p))_{N \geq 1}$ is not non-increasing.*

The rest of the paper is devoted to the proof that player 1 can guarantee $v^*(p)$ in $\Gamma_\infty(p)$.

6. Non revealing strategies. The basic idea is that in $\Gamma_\infty(p)$, the relevant information is only the asymptotic behavior of the sequences of states, hence the projection pB_0. This is due to the fact that both players can "wait" for a large number of stages and place themselves approximately in $\Gamma_\infty(pB_0)$.

Fix some strategy σ of player 1 and some stage n, and consider the point of view of player 2 after stage n has been played (player 1 using σ). Player 2 can compute his belief p_n on the next state, i.e. on the state of stage $n + 1$. He can estimate the asymptotic behavior of this sequence of states as follows: for any $t \geq n + 1$, his belief on the state of stage t is $q_{n,t} = p_n M^{t-(n+1)}$. This sequence of beliefs can be asymptotically approximated by a periodic sequence of period L. This sequence is characterized, for example, by $\lim_{t \to \infty} q_{n,tL+1}$, which is $p_n B_{-n}$. Hence the relevant information for player 2 after stage n has been played can be summarized by $p_n B_{-n}$.

Consequently we will define non revealing strategies for player 1 as strategies such that the sequence $(p_n B_{-n})_n$ is (almost surely) constant. For convenience, and because player 1 does not need to use the actions played by his opponent, we will also require that a non revealing strategy for player 1 plays independently of the actions played by player 2. For any n in $\mathbb{N}$, we denote by $\hat{H}_n$ the set of possible actions of player 1 up to stage n. $\hat{H}_n = \{(i_1, ..., i_n) \in I^n\} = I^n$ ($\hat{H}_0$ standing for a singleton $\{h_0\}$).

Fix p in $\Delta(K)$, and assume that player 1 plays in $\Gamma_\infty(p)$ some strategy σ in Σ^+. (p, σ) induces a stochastic process $(k_1, i_1, k_2, i_2, ..., k_n, i_n, ...)$ over $(K \times I)^\infty$. For n in $\mathbb{N}$ and $h_n = (i_1, ..., i_n) \in \hat{H}_n$, we denote by $p_n(p, \sigma)(h_n) \in \Delta(K)$ the belief of player 2 on the state k_{n+1} of stage $n+1$ knowing that player 1 plays σ and h_n has occurred (i.e. player 1 played i_1 at stage 1,..., i_n at stage n) . Since $\sigma \in \Sigma^+$, this belief is independent of the actions played by player 2. $p_n(p, \sigma)(h_n)$ is just a conditional probability, indeed we simply have for each state k:

$$p_n^k(p, \sigma)(h_n) = \mathbb{P}_{p,\sigma}(k_{n+1} = k | h_n).$$

$p_n(p, \sigma)(h_n)$ is defined arbitrarily in $\Delta(K)$ if $\mathbb{P}_{p,\sigma}(h_n) = 0$ (we will proceed similarly for all further conditional probabilities).

Definition 6.1 *Fix the initial probability p in $\Delta(K)$, and the strategy σ of player 1 in Σ^+. For n in $\mathbb{N}$ and h_n in $\hat{H}_n$, the relevant information of player 2 after h_n has been played is defined as:*

$$q_n(p, \sigma)(h_n) = p_n(p, \sigma)(h_n) B_{-n} \ \in \Delta(K).$$

$p_n(p, \sigma)$ and $q_n(p, \sigma)$ are random variables defined on the measurable space $(\Omega, \hat{\mathcal{H}}_n)$, where $\hat{\mathcal{H}}_n$ is the σ-algebra generated by the projection of any play to the first n actions of player 1. $p_0(p, \sigma)$ is just p and $q_0(p, \sigma)$ is pB_0.

Lemma 6.2 *For any σ in Σ^+, $(q_n(p, \sigma))_{n \geq 0}$ is a $(\hat{\mathcal{H}}_n)_{n \geq 0}$ martingale with respect to $\mathbb{P}_{p,\sigma}$.*

Proof: Fix n in $\mathbb{N}$ and $h_n = (i_1, ..., i_n) \in \hat{H}_n$. Let i_{n+1} be in I and put $h_{n+1} = (i_1, ..., i_{n+1})$. We denote the conditional probability on the state k_{n+1} on stage $n + 1$ knowing that player 1 has played h_{n+1} by $r_n(p, \sigma)(h_{n+1})$. We have $\mathbb{E}_{p,\sigma}(r_n(p, \sigma) | h_n) = p_n(p, \sigma)(h_n)$ and $p_{n+1}(p, \sigma)(h_{n+1}) = r_n(p, \sigma)(h_{n+1})M$.

Consequently,

$$
\begin{aligned}
\mathbb{E}_{p,\sigma}(q_{n+1}(p,\sigma)|h_n) &= \mathbb{E}_{p,\sigma}(p_{n+1}(p,\sigma)B_{-1-n}|h_n)\\
&= \mathbb{E}_{p,\sigma}(r_n(p,\sigma)MB_{-1-n}|h_n)\\
&= p_n(p,\sigma)(h_n)B_{-n}\\
&= q_n(p,\sigma)(h_n).
\end{aligned}
$$

$\square$

Definition 6.3 *A strategy σ in Σ^+ is called non revealing at p if for each n in $\mathbb{N}$, $q_{n+1}(p,\sigma) = q_n(p,\sigma)$ $\mathbb{P}_{p,\sigma}$ a.s.*

We can already notice the following point. Let σ be non revealing at p, n be a multiple of L and h_n be in $\hat{H}_n$ such that $\mathbb{P}_{p,\sigma}(h_n) > 0$. Then $p_n(p,\sigma)(h_n)B_{-n} = pB_0$, hence $p_n(p,\sigma)(h_n)B_0 = p_n(p,\sigma)(h_n)B_{-n}M^n$ $= pB_0M^n = pB_0$. Thus $p_n(p,\sigma)(h_n) \in A(pB_0)$. Whenever player 1 uses a non revealing strategy at p, the belief of player 2 on the current state essentially remains in the set $A(pB_0)$. This will be used to study the values of the non revealing games.

We denote by $\hat{\Sigma}(p)$ the set of strategies of player 1 that are non revealing at p. If $\sigma = (\sigma_n)_{n\geq 1}$ in Σ^+ is independent of the states (i.e. if $\forall n \geq 1$, σ_n only depends on the $(n-1)$ first actions of player 1), then $p_n(p,\sigma) = pM^n$ for each n, hence $q_n(p,\sigma) = pB_0 = q_0(p,\sigma)$. Thus σ is non revealing at p, and $\hat{\Sigma}(p)$ contains all such strategies of player 1. We now give a characterization, which can be seen as an alternative definition, of non revealing strategies.

Fix σ in Σ^+, let n in $\mathbb{N}$ and $h_n = (i_1,...,i_n)$ in $\hat{H}_n$ be such that $\mathbb{P}_{p,\sigma}(h_n) > 0$. We denote by $x(p,\sigma)(h_n) \in \Delta(I)^K$ the expectation of player 1's action at stage $n+1$ depending on the current state:

$$
\forall i \in I, \ \forall k \in K, \ x^k(p,\sigma)(h_n)(i) = \mathbb{P}_{p,\sigma}(i_{n+1} = i|k_{n+1} = k, h_n)
$$

Recall that if $p \in \Delta(K)$ is player 2's belief on the state at some stage n, computed after stage $n-1$, and player 1 plays at this stage according to some x in $\Delta(I)^K$, then for each i in I $\hat{p}(x,i)$ denotes the updated probability on the state of stage n given that player 1 has played i at this stage. Hence player 2's information on the asymptotic behavior of the sequence of states can be described by pB_0 before stage n is played, and by $\hat{p}(x,i)B_0$ after i has been played at stage n.

Definition 6.4 *For each probability p on K, we put:*

$$
NR(p) = \{x \in \Delta(I)^K, \forall i \in I \ s.t. \ x(p)(i) > 0, \ \hat{p}(x,i)B_0 = pB_0\}.
$$

Recall that for each i in I, $x(p)(i) = \sum_{k\in K} p^k x^k(i)$. It is plain that $NR(p)$ can also be written as:

$$
NR(p) = \{x \in \Delta(I)^K, \forall i \in I, (p^k x^k(i))_{k\in K}B_0 = x(p)(i)\,pB_0\}.
$$

$NR(p)$ contains all $x = (x^k)_{k\in K}$ with $x^k = x^{k'}$ for all k and k', hence is non empty. It is clearly a convex compact subset of $\Delta(I)^K$. We can now characterize non revealing strategies at p.

Proposition 6.5 *Let σ be in Σ^+, and p be in $\Delta(K)$.*

σ is non revealing at $p \iff \forall n \in \mathbb{N}, \forall h_n \in \hat{H}_n \ s.t. \ \mathbb{P}_{p,\sigma}(h_n) > 0, x(p,\sigma)(h_n) \in NR(p_n(p,\sigma)(h_n)).$

Proof: Fix p and σ. For simplicity of notations, we omit to mention (p,σ) when writing $x(p,\sigma)(h_n)$, $p_n(p,\sigma)(h_n)$ and $q_n(p,\sigma)(h_n)$.

Let n be in $\mathbb{N}$, $h_n = (i_1,...,i_n)$ be in $\hat{H}_n$ s.t. $\mathbb{P}_{p,\sigma}(h_n) > 0$. Put for simplicity $x = x(p,\sigma)(h_n) \in \Delta(I)^K$, and $p_n = p_n(p,\sigma)(h_n) \in \Delta(K)$. For any i in I, $\mathbb{P}_{p,\sigma}(i_{n+1} = i|h_n) = \sum_{k\in K} \mathbb{P}_{p,\sigma}(k_{n+1} = k|h_n)x^k(i) = \sum_{k\in K} p_n^k x^k(i) = x(p_n)(i)$.

Fix i_{n+1} in I, and let h_{n+1} be $(i_1,...,i_n,i_{n+1})$. If $\mathbb{P}_{p,\sigma}(h_{n+1}) > 0$, we have: $(\mathbb{P}_{p,\sigma}(k_{n+1} = k|h_{n+1}))_{k\in K} = \hat{p}_n(x,i_{n+1})$. So $p_{n+1}(h_{n+1}) = \hat{p}_n(x,i_{n+1})M$, and $q_{n+1}(h_{n+1}) = \hat{p}_n(x,i_{n+1})B_{-n}$.

We now prove the $\implies$ part of the proposition. Assume that σ is non revealing at p. Then for each n in $\mathbb{N}$, h_n in $\hat{H}_n$ and i_{n+1} in I s.t. $\mathbb{P}_{p,\sigma}(h_n) > 0$ and $x(h_n)(i_{n+1}) > 0$, we have $q_{n+1}(h_n, i_{n+1}) = q_n(h_n)$. Writing p_n for $p_n(h_n)$, we have $\hat{p_n}(x(h_n), i_{n+1})B_{-n} = p_n B_{-n}$, so $\hat{p_n}(x(h_n), i_{n+1})B_0 = p_n B_0$. Hence $x(h_n) \in NR(p_n)$.

We conclude with the $\impliedby$ part. Assume that $\forall n \in \mathbb{N}, \forall h_n \in \hat{H}_n$ s.t. $\mathbb{P}_{p,\sigma}(h_n) > 0$, $x(h_n) \in NR(p_n(h_n))$.

Fix n in $\mathbb{N}$, and $h_{n+1} = (i_1, ..., i_{n+1})$ s.t. $\mathbb{P}_{p,\sigma}(h_{n+1}) > 0$. Put $h_n = (i_1, ..., i_n)$. Then $q_{n+1}(h_{n+1}) = \hat{p_n}(x(h_n), i_{n+1})B_{-n} = \hat{p_n}(x(h_n), i_{n+1})B_0 M^l$, where $n + l$ is a multiple of L. By hypothesis, $\hat{p_n}(x(h_n), i_{n+1})B_0 = p_n(h_n)B_0$, thus $q_{n+1}(h_{n+1}) = p_n(h_n)B_0 M^l = p_n(h_n)B_l = q_n(h_n)$. $(q_n)_{n\geq 0}$ is constant $\mathbb{P}_{p,\sigma}$-a.s.. $\qquad\square$

In order to illustrate the previous notions, we come back to the examples of section 3.

In examples A and B, $(M^n)_n$ converges to $\begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}$, which is B_0. For each p in $\Delta(K)$, $pB_0 = (1/2, 1/2)$ does not depend on p, hence every strategy of player 1 in Σ^+ is non revealing at p. Player 1 can use his information on the current state without revealing to player 2 any information on the limit behavior of the sequence of states. Here Q reduces to the singleton $\{(1/2, 1/2)\}$, and B_0 is the projection matrix on Q.

In examples C and D, B_0 is the identity matrix hence for each p, $NR(p) = \{x \in \Delta(I)^K, \forall (k, k')$ s.t. $p^k > 0$ and $p^{k'} > 0, x^k = x^{k'}\}$. A non revealing strategy for player 1 is, like in Aumann and Maschler case, a strategy in Σ^+ that plays at each stage independently of the current state (formally, that plays the same mixed action in all states having a positive probability). Example C shows how our definition generalizes the one of Aumann and Maschler.

These first four examples indeed only show extreme cases where the set of non revealing strategies at p is mainly independent of p. In case of example F, one obtains that $B_0 = \begin{pmatrix} 1/2 & 1/2 & 0 \\ 1/2 & 1/2 & 0 \\ 0 & 0 & 1 \end{pmatrix}$, and for each p with full support,

$$NR(p) = \left\{ x \in \Delta(I)^K, \forall i \in I, x(p)(i) = x^c(i) = \frac{p^a x^a(i) + p^b x^b(i)}{p^a + p^b} \right\}.$$

Notice that $NR(p)$ does not depend on the "weights" of p over recurrence classes, i.e. does not only depend on $p^a + p^b$, but depends on the conditional probability $p^a/(p^a + p^b)$. Here, $Q = \{(p^k)_{k\in K} \in \Delta(K), p^a = p^b\}$ and B_0 corresponds to the orthogonal projection on Q.

In Example E we have two recurrence classes, and c is a transient state. $B_0 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2/3 & 1/3 & 0 \end{pmatrix}$, and $Q = \{(p^k)_{k\in K} \in \Delta(K), p^c = 0\}$. For p in $\Delta(K)$ with full support,

$$NR(p) = \left\{ x \in \Delta(I)^K, \forall i \in I, \frac{p^a x^a(i) + 2/3p^c x^c(i)}{p^a + 2/3p^c} = \frac{p^b x^b(i) + 1/3p^c x^c(i)}{p^b + 1/3p^c} \right\}.$$

We now prove some elementary results on the structure of the set of non revealing strategies of player 1. First notice that since player 1 has perfect recall in the original game $\Gamma_\infty(p)$, by Aumann's extension [2] of Kuhn's theorem any strategy of player 1 in Σ can be viewed as a mixed strategy, i.e. as a probability distribution over the set of pure strategies of player 1 in $\Gamma_\infty(p)$, and vice-versa.

Lemma 6.6 *Fix q in Q, and S a positive integer.*

Consider $p = \sum_{s=1}^S \lambda_s p_s$, with $p_1, ..., p_S$ elements of $A(q)$, $\lambda_1, ..., \lambda_S$ non negative numbers such that $\sum_{s=1}^S \lambda_s = 1$, and for each s let σ_s be a non revealing strategy for player 1 at p_s. Then we have:

(1) If σ in Σ^+ satisfies $\mathbb{P}_{p,\sigma} = \sum_{s=1}^S \lambda_s \mathbb{P}_{p_s,\sigma_s}$, then σ is non revealing at p.

(2) There exists some strategy σ non revealing at p and such that $\mathbb{P}_{p,\sigma} = \sum_{s=1}^{S} \lambda_s \mathbb{P}_{p_s,\sigma_s}$.

Proof:

(1) Let n be in $\mathbb{N}$ and h_n be in $\hat{H}_n$ such that $\mathbb{P}_{p,\sigma}(h_n) > 0$. For each k in K,

$$
\begin{aligned}
p_n^k(p,\sigma)(h_n) &= \frac{\mathbb{P}_{p,\sigma}(h_n, k_{n+1} = k)}{\mathbb{P}_{p,\sigma}(h_n)} \\
&= \frac{\sum_{s=1}^{S} \lambda_s \mathbb{P}_{p_s,\sigma_s}(h_n, k_{n+1} = k)}{\sum_{s=1}^{S} \lambda_s \mathbb{P}_{p_s,\sigma_s}(h_n)} \\
&= \frac{\sum_{s=1}^{S} \lambda_s \mathbb{P}_{p_s,\sigma_s}(h_n) p_n^k(p_s,\sigma_s)(h_n)}{\sum_{s=1}^{S} \lambda_s \mathbb{P}_{p_s,\sigma_s}(h_n)}
\end{aligned}
$$

Hence $p_n(p,\sigma)(h_n)$ is in the convex hull of $\{p_n(p_s,\sigma_s)(h_n), s = 1, ..., S\}$. Since for each s in S, $q_n(p_s,\sigma_s)(h_n) = p_s B_0 = q$, we have $q_n(p,\sigma)(h_n) = q = pB_0$ and σ is non revealing at p.

(2) We define σ *via* the splitting procedure of Aumann and Maschler. To play according to σ, observe the first state k, then choose s in S with probability $\lambda_s p_s^k/p^k$. Finally play in the whole game according to σ_s.

σ is actually defined as a mixture of behavior strategies, and by Aumann's extension of Kuhn's theorem, σ can be viewed as an element of Σ. Moreover σ does not depend on the moves of player 2, hence can be considered as an element of Σ^+. Let now A be any measurable subset of $(K \times I)^\infty$.

$$
\begin{aligned}
\mathbb{P}_{p,\sigma}(A) &= \sum_{k \in K} \mathbb{P}_{p,\sigma}(k_1 = k) \mathbb{P}_{p,\sigma}(A|k_1 = k) \\
&= \sum_{k \in K} p^k \sum_{s \in S} \mathbb{P}_{p,\sigma}(s|k_1 = k) \mathbb{P}_{p,\sigma}(A|k_1 = k, s) \\
&= \sum_{s \in S} \lambda_s \sum_{k \in K} p_s^k \mathbb{P}_{\delta_k,\sigma_s}(A) \\
&= \sum_{s \in S} \lambda_s \mathbb{P}_{p_s,\sigma_s}(A)
\end{aligned}
$$

So $\mathbb{P}_{p,\sigma} = \sum_{s=1}^{S} \lambda_s \mathbb{P}_{p_s,\sigma_s}$, and (1) shows that σ is non revealing at p. $\qquad\square$

Remark 6.7 *In Aumann and Maschler's case, $A(q)$ is a singleton and the previous lemma is not very helpful. However, we consider lemma 6.6 as the analog, adapted to non revealing strategies, of Aumann and Maschler's splitting lemma (see e.g. Zamir [20]). This lemma may be reformulated as follows: if p is a convex combination $\sum_{s \in S} \lambda_s p_s$, and if for each s in S we are given a (revealing or non revealing) strategy σ_s of player 1 in Σ^+, then there exists a strategy σ in Σ^+ such that $\mathbb{P}_{p,\sigma} = \sum_{s=1}^{S} \lambda_s \mathbb{P}_{p_s,\sigma_s}$. The proof is exactly the same as the proof of point (2) of lemma 6.6, which is thus due to Aumann and Maschler. They used the splitting lemma to study what can be guaranteed by player 1 in the original game $\Gamma_\infty(p)$. We will proceed similarly with lemma 6.6, but to study what can be guaranteed by player 1 in an auxiliary game where he is restricted to play a non revealing strategy. This is why we concentrate here on non revealing strategies, whereas the splitting lemma was used by Aumann and Maschler to construct optimal strategies that were not non revealing.*

In the following corollary, $\hat{\Sigma}(p)$ is viewed as a subset of the set of mixed strategies of player 1.

Corollary 6.8 *The set of non revealing strategies of player 1 at p is convex.*

Proof: Let σ_1 and σ_2 be in $\hat{\Sigma}(p)$, and let λ be in $[0,1]$. Define the strategy σ of player 1 as follows: with probability λ, play σ_1 and with probability $(1 - \lambda)$, play σ_2. By point *(1)* of lemma 6.6, σ is non revealing at p. $\qquad\square$

We will also need in section 7 more elaborate properties of non revealing strategies. We now define two types of continuation strategies for player 1. Fix σ in Σ^+ and p in $\Delta(K)$, and consider that player

1 uses σ in the original game $\Gamma_\infty(p)$. Although what follows is conceptually simple, we need numerous notations to be precise.

We first define, for each n in $\mathbb{N}$ and $h_n = (i_1, ..., i_n) \in I^n$ s.t. $\mathbb{P}_{p,\sigma}(h_n) > 0$, the "expected continuation strategy" $\sigma(p, h_n)$ of player 1 after stage n and the play of h_n.

For every possible sequence of states up to stage n, $h^n = (k_1, ..., k_n)$ in K^n, we denote by $\omega(h_n, h^n)$ the finite history of states and actions of player 1 up to stage n: $\omega(h_n, h^n) = (k_1, i_1, ..., k_n, i_n)$. We denote by $\sigma(h_n, h^n)$ the behavior strategy in Σ^+ played by σ after $\omega(h_n, h^n)$. We define $\sigma(p, h_n)$ in Σ^+ as follows, similarly to the splitting procedure (the tilde now denoting random variables to avoid confusion): if $\tilde{k}_1 = k$, then choose h^n in K^n according to $\mathbb{P}_{p,\sigma}(h^n | h_n, \tilde{k}_{n+1} = k)$ and play according to $\sigma(h_n, h^n)$ (if $\tilde{k}_1 = k$ s.t. $\mathbb{P}_{p,\sigma}(\tilde{k}_{n+1} = k, h_n) = 0$, just play arbitrarily).

$\sigma(p, h_n)$ has the following interesting property. For m in $\mathbb{N}$, $(k_{n+1}, i_{n+1}, ..., k_{n+m}, i_{n+m}) \in (K \times I)^m$, define the events:
$$A = (\tilde{k}_{n+1} = k_{n+1}, \tilde{i}_{n+1} = i_{n+1}, ..., \tilde{k}_{n+m} = k_{n+m}, \tilde{i}_{n+m} = i_{n+m})$$
and
$$B = (\tilde{k}_1 = k_{n+1}, \tilde{i}_1 = i_{n+1}, ..., \tilde{k}_m = k_{n+m}, \tilde{i}_m = i_{n+m}).$$

We have:
$$
\begin{aligned}
\mathbb{P}_{p,\sigma}(A|h_n) &= \sum_{k \in K} p_n^k(p,\sigma)(h_n) \mathbb{P}_{p,\sigma}(A|h_n, \tilde{k}_{n+1} = k) \\
&= \sum_{k \in K} p_n^k(p,\sigma)(h_n) \sum_{h^n \in K^n} \mathbb{P}_{p,\sigma}(h^n|h_n, \tilde{k}_{n+1} = k) \mathbb{P}_{p,\sigma}(A|\omega(h_n, h^n), \tilde{k}_{n+1} = k) \\
&= \sum_{k \in K} p_n^k(p,\sigma)(h_n) \sum_{h^n \in K^n} \mathbb{P}_{p,\sigma}(h^n|h_n, \tilde{k}_{n+1} = k) \mathbb{P}_{\delta^k, \sigma(h_n, h^n)}(B) \\
&= \mathbb{P}_{p_n(p,\sigma)(h_n), \sigma(p, h_n)}(B)
\end{aligned}
$$

The following lemma expresses some kind of subgame property of non revealing strategies, and will be used later in the proof of proposition 7.4. Its proof can be found in the appendix.

Lemma 6.9 *Let σ be in Σ^+, and p be in $\Delta(K)$.*
$$\sigma \in \hat{\Sigma}(p) \iff \forall n \in \mathbb{N}, \forall h_n \in I^n \text{ s.t. } \mathbb{P}_{p,\sigma}(h_n) > 0, \ \sigma(p, h_n) \in \hat{\Sigma}(p_n(p,\sigma)(h_n))$$

In the same spirit, we now define for each n in $\mathbb{N}$ the expected strategy $\sigma(p, n+)$ of player 1 in $\Gamma_\infty(p)$ after stage n. If $\tilde{k}_1 = k$, choose (h_n, h^n) in $I^n \times K^n$ according to $\mathbb{P}_{p,\sigma}(\omega(h_n, h^n)|\tilde{k}_{n+1} = k)$ and play according to $\sigma(h_n, h^n)$ (if $\tilde{k}_1 = k$ such that $\mathbb{P}_{p,\sigma}(\tilde{k}_{n+1} = k) = 0$, play arbitrarily).

Note that $\sigma(p, 0+)$ is just σ (up to events with probability zero). For m in $\mathbb{N}$ and $(k_{n+1}, i_{n+1}, ..., k_{n+m}, i_{n+m}) \in (K \times I)^m$, define as before:
$$A = (\tilde{k}_{n+1} = k_{n+1}, \tilde{i}_{n+1} = i_{n+1}, ..., \tilde{k}_{n+m} = k_{n+m}, \tilde{i}_{n+m} = i_{n+m})$$
and
$$B = (\tilde{k}_1 = k_{n+1}, \tilde{i}_1 = i_{n+1}, ..., \tilde{k}_m = k_{n+m}, \tilde{i}_m = i_{n+m}).$$

One has for each k in K:
$$\mathbb{P}_{p,\sigma}(A|\tilde{k}_{n+1} = k) = \sum_{(h_n, h^n) \in I^n \times K^n} \mathbb{P}_{p,\sigma}(\omega(h_n, h^n)|\tilde{k}_{n+1} = k) \mathbb{P}_{\delta^k, \sigma(h_n, h^n)}(B)$$

Hence,
$$
\begin{aligned}
\mathbb{P}_{p,\sigma}(A) &= \sum_{k \in K} \mathbb{P}_{p,\sigma}(\tilde{k}_{n+1} = k) \sum_{(h_n, h^n) \in I^n \times K^n} \mathbb{P}_{p,\sigma}(\omega(h_n, h^n)|\tilde{k}_{n+1} = k) \mathbb{P}_{\delta^k, \sigma(h_n, h^n)}(B) \\
&= \mathbb{P}_{pM^n, \sigma(p, n+)}(B)
\end{aligned}
$$

We then obtain that:

$$
\begin{aligned}
\mathbb{P}_{pM^n,\sigma(p,n+)}(B) \;=\; \mathbb{P}_{p,\sigma}(A) &= \sum_{h_n \in I^n} \mathbb{P}_{p,\sigma}(h_n)\mathbb{P}_{p,\sigma}(A|h_n) \\
&= \sum_{h_n \in I^n} \mathbb{P}_{p,\sigma}(h_n)\mathbb{P}_{p_n(p,\sigma)(h_n),\sigma(p,h_n)}(B)
\end{aligned}
$$

Consequently, $\mathbb{P}_{pM^n,\sigma(p,n+)} = \sum_{h_n \in I^n} \mathbb{P}_{p,\sigma}(h_n)\mathbb{P}_{p_n(p,\sigma)(h_n),\sigma(p,h_n)}$.

We now have an analog of lemma 6.9 for this type of continuation strategy, which will be used later in the proof of proposition 7.7.

Lemma 6.10 *Let σ be in Σ^+ and p be in $\Delta(K)$.*

$$
\sigma \in \hat{\Sigma}(p) \Longleftrightarrow \forall n \in \mathbb{N}, \; \sigma(p,n+) \in \hat{\Sigma}(pM^n)
$$

The proof can be found in the appendix. Finally we also need to define explicitly N-stage non revealing strategies.

Definition 6.11 *If $\sigma = (\sigma_n)_{n \geq 1}$ is in $\hat{\Sigma}(p)$, the restriction of σ to the first N-stages, i.e. the strategy $(\sigma_n)_{n=1,...,N}$, is called a N-stage non revealing strategy for player 1 at p. We denote by $\hat{\Sigma}_N(p)$ the set of such strategies.*

7. Non revealing games.

Definition 7.1 *For p in $\Delta(K)$, the non revealing game at p, denoted by $\hat{\Gamma}_\infty(p)$, is the game obtained from $\Gamma_\infty(p)$ by restricting player 1 to play a strategy in $\hat{\Sigma}(p)$.*

The notions of guaranteeing and value in $\hat{\Gamma}_\infty(p)$ are defined as in definition 2.1: one just has to replace everywhere $\Gamma_\infty(p)$ by $\hat{\Gamma}_\infty(p)$ and Σ by $\hat{\Sigma}(p)$. The N-stages non revealing games are defined as follows:

Definition 7.2 *For $N \geq 1$, the N-stage non revealing game at p is defined as the zero-sum game $\hat{\Gamma}_N(p)$ with strategy spaces $\hat{\Sigma}_N(p)$ for player 1, $\mathcal{T}_N$ for player 2 and with payoff function γ_N^p for player 1.*

We first study the value of the N-stage games.

Proposition 7.3 *For each positive N, and probability p on K:*

(1) $\hat{\Gamma}_N(p)$ has a value, denoted by $\hat{v}_N(p)$ and both players have optimal strategies.

(2) $\hat{v}_N$ is an upper semi-continuous (u.s.c.) mapping from $\Delta(K)$ to $\mathbb{R}$, and $|\hat{v}_N(p)| \leq C$.

(3) For each q in Q, the restriction of $\hat{v}_N$ to $A(q)$ is concave.

Proof: $\hat{\Gamma}_N(p)$ is the zero-sum game $(\hat{\Sigma}_N(p), \mathcal{T}_N, \gamma_N^p)$. In the original N-stage game $(\Sigma_N, \mathcal{T}_N, \gamma_N^p)$, the sets of mixed strategies of both players are the convex hull of a finite number of points, hence are convex subsets of some Euclidean space. By corollary 6.8, $\hat{\Sigma}_N(p)$ is a convex subset of Σ_N (we always identify mixed and behavior strategies). Since γ_N^p is bilinear over $\Sigma_N \times \mathcal{T}_N$, if we show that $\hat{\Sigma}_N(p)$ is closed in Σ_N we can apply Sion's theorem to obtain that $\hat{\Gamma}_N(p)$ has a value and both players have optimal strategies.

Put $H = \cup_{n=0}^{N-1}\hat{H}_{n+1}$, and consider the mapping F from $\Delta(K) \times \Sigma_N^+$ to $(\mathbb{R}^K)^H$ such that for each (p,σ) in $\Delta(K) \times \Sigma_N^+$, $F(p,\sigma)$ is:

$$
\left(\mathbb{P}_{p,\sigma}(h_{n+1}) \left(q_{n+1}^k(p,\sigma)(h_{n+1}) - q_n^k(p,\sigma)(i_1,...,i_n) \right) \right)_{k \in K, n=0,...,N-1, h_{n+1}=(i_1,...,i_{n+1}) \in \hat{H}_{n+1}}.
$$

It is clear that we have $F(p,\sigma) = 0$ if and only if $\sigma \in \hat{\Sigma}_N(p)$.

For each (p,σ), each n in $\{0,...,N-1\}$ and each $h_{n+1} = (i_1,...,i_{n+1})$ in $\hat{H}_{n+1}$, we have:

$$
\mathbb{P}_{p,\sigma}(h_{n+1}) \left(q_{n+1}(p,\sigma)(h_{n+1}) - q_n(p,\sigma)(h_n) \right)
$$

$$= I\!\!P_{p,\sigma}(h_{n+1})\left(p_{n+1}(p,\sigma)(h_{n+1})B_{-1-n} - p_n(p,\sigma)(h_n)B_{-n}\right)$$

$$= (I\!\!P_{p,\sigma}(h_{n+1}, k_{n+2} = k))_k B_{-1-n} - \frac{I\!\!P_{p,\sigma}(h_{n+1})}{I\!\!P_{p,\sigma}(h_n)}(I\!\!P_{p,\sigma}(h_n, k_{n+1} = k))_k B_{-n}.$$

(with an obvious convention if $I\!\!P_{p,\sigma}(h_n) = 0$.)

The mapping $((p,\sigma) \longrightarrow I\!\!P_{p,\sigma}(h))$ being continuous for any finite history h in $(K \times I)^N$, F is continuous.

This implies that $\hat{\Sigma}_N(p)$ is closed in Σ_N^+. Since Σ_N^+ is a compact subset of Σ_N, so is $\hat{\Sigma}_N(p)$. Thus (1) is proved by Sion's theorem.

Secondly, the correspondence from $\Delta(K)$ to Σ_N^+ which associates to each p the set $\hat{\Sigma}_N(p)$ has a closed graph. By a maximum theorem (see e.g. Aubin and Cellina [1, Theorem 5 p.52]), this easily implies that $\hat{v}_N$ is upper semi-continuous. $|\hat{v}_n(p)| \leq C$ being obvious, (2) is proved.

Finally, (3) is just a consequence of our "splitting lemma": with the same notations as lemma 6.6, choose for each s σ_s an optimal strategy of player 1 in $\hat{\Sigma}_N(p_s)$. □

We can now write a recursive formula for $\hat{v}_n$ ($\hat{v}_0$ being defined arbitrarily). The proof is very similar to that of proposition 5.1 and can be found in the Appendix.

Proposition 7.4 *For each $n \geq 1$ and p in $\Delta(K)$,*

$$\hat{v}_n(p) = \frac{1}{n} \max_{x \in NR(p)} \min_{y \in \Delta(J)} \left(G(p, x, y) + (n-1) \sum_{i \in I} x(p)(i)\, \hat{v}_{n-1}(\hat{p}(x, i)M) \right)$$

$$= \frac{1}{n} \min_{y \in \Delta(J)} \max_{x \in NR(p)} \left(G(p, x, y) + (n-1) \sum_{i \in I} x(p)(i)\, \hat{v}_{n-1}(\hat{p}(x, i)M) \right)$$

As for v_n, we have for each $n \geq 1$ and p in $\Delta(K)$ that $\hat{v}_n(p) \geq \frac{1}{n}(u(p) + (n-1)\hat{v}_{n-1}(pM))$, and by concavity of $\hat{v}_{n-1}$ on $A(pMB_0)$, $\hat{v}_n(p) \leq \frac{1}{n}(\hat{v}_1(p) + (n-1)\hat{v}_{n-1}(pM))$. Hence $\hat{v}_n(p) - \hat{v}_{n-1}(pM) \longrightarrow_{n \to \infty} 0$. As $\hat{v}_n(p) = \max_{\sigma \in \hat{\Sigma}(p)} \min_{\tau \in \mathcal{T}} \gamma_n^p(\sigma, \tau)$ it is also clear that $\hat{v}_n(p) - \hat{v}_{n-1}(p) \longrightarrow_{n \to \infty} 0$. Thus $\hat{v}_n(p) - \hat{v}_n(pM) \longrightarrow_{n \to \infty} 0$. A main interest of the recursive formula for $\hat{v}_n$ is the following corollary.

Corollary 7.5 *For each n, $\hat{v}_n$ is continuous.*

Proof:

We first modify the recursive formula by changing variables. For any p in $\Delta(K)$ and x in $NR(p)$, we put $z(p,x) = (p^k x^k(i))_{k \in K, i \in I} \in (I\!\!R^K)^I$. Define Z as the correspondence from $\Delta(K)$ to $(I\!\!R^K)^I$ which associates to each probability p the set $\{z(p,x), x \in NR(p)\}$.

For $z = (z_i^k)_{k \in K, i \in I}$ in the non-negative orthant $(I\!\!R^K)_+^I$, we put for each i in I, $z(i) = \sum_{k \in K} z_i^k$ and $z_i = (z_i^k)_{k \in K} \in I\!\!R^K$. $\frac{z_i}{z(i)}$ is thus an element of $\Delta(K)$ (defined arbitrarily if $z(i) = 0$). Define finally $H(z) = \min_{y \in \Delta(J)} \sum_{k \in K} \sum_{i \in I} z_i^k G^k(i, y)$.

The recursive formula becomes: $\forall n \geq 1, \forall p \in \Delta(K)$,

$$\hat{v}_n(p) = \frac{1}{n} \max_{z \in Z(p)} \left(H(z) + (n-1) \sum_{i \in I} z(i)\, \hat{v}_{n-1}(\frac{z_i}{z(i)}M) \right).$$

We have $Z(p) = \{z \in (I\!\!R^K)_+^I, \sum_{i \in I} z_i = p \text{ and } \forall i \in i, z_i B_0 = z(i)pB_0\}$, hence $Z(p)$ is a polytope. H is continuous, and if $f : \Delta(K) \longrightarrow I\!\!R$ is continuous, so is the mapping $\left(z \longrightarrow \sum_{i \in I} z(i)f(\frac{z_i}{z(i)}M)\right)$. If we prove that the correspondence Z is continuous (i.e. lower and upper semi-continuous) then by induction we obtain that $\hat{v}_n$ is continuous for each n. As it is clear that Z has a compact graph, we just have to show that Z is lower semi-continuous (l.s.c. for short).

Fix p in $\Delta(K)$ and z in $Z(p)$. We put $D = \{(k, i) \in K \times I, z(i)p^k - z_i^k > 0\}$. In order to prove that Z is l.s.c., we will show that:

$$\forall p' \in \Delta(K), \exists z^* \in Z(p') \text{ with } \|z - z^*\| \le \|p - p'\|(1 + A/B),$$

where $A = \sum_{k \in K, i \in I} |p^k z(i) - z_i^k|$ and $B = \min_{(k,i) \in D} z(i)p^k - z_i^k$, with the convention that $A/B = 0$ if $D = \emptyset$.

First define z' in $(\mathbb{R}^K)^I$ s.t. for each i and k, $z_i'^k = z_i^k + z(i)(p'^k - p^k)$. We have $\sum_{i \in I} z_i' = p'$, for each i in I, $z'(i) = z(i)$, $(z_i' - z'(i)p')B_0 = 0$ and $\|z - z'\| = \|p - p'\|$. If $D = \emptyset$, $z' \in (\mathbb{R}^K)_+^I$ and we are done. Assume now that z' is not in $(\mathbb{R}^K)_+^I$.

Define z'' in $(\mathbb{R}^K)_+^I$ such that for each i and k, $z_i''^k = p'^k z(i)$. z'' belongs to $Z(p')$. Finally put, for every λ in $[0,1]$, $z(\lambda) = (1 - \lambda)z' + \lambda z''$. For each λ, $z(\lambda) \in Z(p')$ if and only if $z(\lambda)$ has non-negative coordinates.

Let $\lambda^* \in (0,1]$ be $\max_{(k,i) \in D} \frac{z(i)p^k - z_i^k - z(i)p'^k}{z(i)p^k - z_i^k}$. A simple computation shows that $z(\lambda^*) \in (\mathbb{R}^K)_+^I$ and that $\lambda^* \le \|p - p'\|/B$. Hence:

$$
\begin{aligned}
\|z(\lambda^*) - z\| &= \sum_{k \in K} \sum_{i \in I} |z_i^k(\lambda^*) - z_i^k| \\
&= \sum_{k \in K} \sum_{i \in I} |z(i)(p'^k - p^k) + \lambda^* p^k z(i) - \lambda^* z_i^k| \\
&\le \sum_{k \in K} \sum_{i \in I} z(i)|p'^k - p^k| + \lambda^* \sum_{k \in K} \sum_{i \in I} |p^k z(i) - z_i^k| \\
&\le \|p - p'\| + \|p - p'\|A/B
\end{aligned}
$$

This proves that Z is l.s.c., and we obtain that for each n, $\hat{v}_n$ is continuous. □

We can now study the non revealing games with infinitely many stages. We first define an auxiliary mapping from $\Delta(K)$ to $\mathbb{R}$.

Definition 7.6 $\forall p \in \Delta(K), \quad w(p) = \inf_{N \ge 1} \hat{v}_{NL}(pB_0).$

As an infimum of continuous functions, w is u.s.c. on $\Delta(K)$. For each q in Q, it is constant on $A(q)$. $w(p)$ will turn out to be the value of the non revealing game $\hat{\Gamma}_\infty(p)$. The following proposition is the analog of proposition 5.5 for non revealing games.

Proposition 7.7 *In any non revealing game $\hat{\Gamma}_\infty(p)$, player 2 can guarantee $w(p)$.*

Proof: Let N be a positive multiple of L, and define a strategy τ of player 2 as follows. Divide the set of stages into consecutive blocks $B^1, ..., B^m, ...$ of length N. For each positive m, τ plays on block B^m an optimal strategy τ_m in $\hat{\Gamma}_N(pM^{(m-1)N})$, independently on what happened at previous blocks.

Fix now σ a strategy for player 1 in $\hat{\Sigma}(p)$. At some block B^m, the expected payoff induced by p, σ and τ at B^m is:

$$\mathbb{E}_{p,\sigma,\tau}\left(\tfrac{1}{N} \sum_{n \in B^m} G^{k_n}(i_n, j_n)\right)$$

Since σ is independent of the moves of player 2, this payoff only depends on τ_m and on the stochastic process induced by p and σ over the states and actions of player 1 at this block. But since $\mathbb{P}_{pM^{(m-1)N}, \sigma(p,(m-1)N+)}(B) = \mathbb{P}_{p,\sigma}(A)$ for any events A and B as before (see before lemma 6.10), we have that:

$$\mathbb{E}_{p,\sigma,\tau}\left(\frac{1}{N} \sum_{n \in B^m} G^{k_n}(i_n, j_n)\right) = \mathbb{E}_{pM^{(m-1)N}, \sigma(p,(m-1)N+),\tau_m}\left(\frac{1}{N} \sum_{n=1}^{N} G^{k_n}(i_n, j_n)\right)$$

By lemma 6.10, $\sigma(p, (m-1)N+)$ is in $\hat{\Sigma}(pM^{(m-1)N})$, hence by definition of τ_m, this expected payoff is at most $\hat{v}_N(pM^{(m-1)N})$. Consequently, for each S:

$$\gamma_{NS}^p(\sigma, \tau) \le \frac{1}{S} \sum_{s=1}^{S} \hat{v}_N(pM^{(s-1)N}),$$

and by concavity of $\hat{v}_N$ on $A(pB_0)$, we get:

$$\gamma^p_{NS}(\sigma,\tau) \le \hat{v}_N\left(\frac{1}{S}\sum_{s=1}^{S} pM^{(s-1)N}\right).$$

$\hat{v}_N$ being continuous by corollary 7.5, the right-hand side converges to $\hat{v}_N(pB_0)$ as S goes to infinity, and this convergence is uniform in σ. We have obtained:

$$\forall \varepsilon > 0 \; \exists T_0 \; \forall T \ge T_0 \; \forall \sigma \in \hat{\Sigma}(p) \;\; \gamma^p_T(\sigma,\tau) \le \hat{v}_N(pB_0) + \varepsilon. \hspace{2cm} \square$$

Note that the previous proof also directly gives the following interesting property.

Proposition 7.8 *Let p be in $\Delta(K)$, S and N be positive integers with N a multiple of L. Then:*

$$\hat{v}_{NS}(p) \le \frac{1}{S}\sum_{s=1}^{S} \hat{v}_N(pM^{(s-1)N}).$$

For each q in Q, we have $qM^L = q$, hence the previous proposition gives : $\forall N \ge 1$, $\forall S \ge 1$, $\hat{v}_{LNS}(q) \le \hat{v}_{LN}(q)$. We obtain the following corollary, which will be used later in section 9.

Corollary 7.9 *For each q in Q, $(\hat{v}_{L2^N}(q))_{N\ge 1}$ is non-increasing.*

We can now also prove the convergence of the value of the N-stage games.

Corollary 7.10 *For each p in $\Delta(K)$, $\hat{v}_N(p) \longrightarrow_{N\to\infty} w(p)$.*

Proof:

For each p in $\Delta(K)$, by proposition 7.7 player 2 can guarantee $w(p)$ in $\hat{\Gamma}_\infty(p)$, hence: $\limsup_N \hat{v}_N(p) \le w(p)$.

Consequently, for each q in Q:

$$\limsup_N \hat{v}_N(q) \le \inf_N \hat{v}_{NL}(q).$$

so $\hat{v}_{NL}(q) \longrightarrow_{N\to\infty} w(q)$. As $\hat{v}_N(q) - \hat{v}_{N-1}(q) \longrightarrow_{n\to\infty} 0$, we obtain $\hat{v}_N(q) \longrightarrow_{N\to\infty} w(q)$.

Define now, for each p in $\Delta(K)$, $\underline{v}(p) = \liminf_N \hat{v}_N(p)$. As for each p in $\Delta(K)$, $\hat{v}_N(p) - \hat{v}_N(pM) \longrightarrow_{N\to\infty} 0$, we have for each N, $\underline{v}(p) = \underline{v}(pM) = ... = \underline{v}(pM^N)$.

Fix finally p in $\Delta(K)$, and put $q = pB_0 \in Q$. Since for each N $\hat{v}_N$ is concave on $A(q)$ so is $\underline{v}$. Since $A(q)$ is a polytope, $\underline{v}$ is then necessarily l.s.c. on $A(q)$ (a reference for this is Mertens *et al*, part A, p.46, ex. 15). As $pM^{NL} \longrightarrow_{N\to\infty} q$, we get that $\underline{v}(p) \ge \underline{v}(q) = w(q)$.

Summing, we obtain that $\limsup_N \hat{v}_N(p) \le w(p) = w(q) \le \underline{v}(p)$, and the corollary is proved. $\hspace{1cm} \square$

We can now conclude about the value of the non revealing games.

Theorem 7.11 *For each p in $\Delta(K)$, the non revealing game $\hat{\Gamma}_\infty(p)$ has a value which is $w(p)$.*

Proof: Fix p in $\Delta(K)$, and put $q = pB_0$, and $w = w(p)$. By proposition 7.7, we just have to prove that player 1 can guarantee $w(p)$ in $\hat{\Gamma}_\infty(p)$.

$A(q)$ is a polytope and for each N $\hat{v}_N$ is concave on $A(q)$, hence by the previous corollary we can obtain:

$$\forall \varepsilon > 0 \; \exists N_0 \; \forall N \ge N_0 \;\; \forall p' \in A(q) \; \hat{v}_N(p') \ge w - \varepsilon.$$

Fix now $\varepsilon > 0$, and let N_0 be as above. Divide the set of stages into consecutive blocks $B^1,...,B^m,...$ of length N_0L. We define σ by induction on the blocks as follows:

- at block B^1, σ plays an optimal strategy in $\hat{\Gamma}_{N_0L}(p)$.

- at block B^m, σ is defined as follows: if $h_{(m-1)N_0L}$ in $\hat{H}_{(m-1)N_0L}$ has been played by player 1 at previous blocks, put $p_m = p_{(m-1)N_0L}(p,\sigma)(h_{(m-1)N_0L})$. Notice that p_m only depends on the definition of σ on previous blocks. σ plays at block B^m after $h_{(m-1)N_0L}$ an optimal strategy in $\hat{\Gamma}_{N_0L}(p_m)$.

By proposition 6.5, σ is a non revealing strategy at p. And for each m and $h_{(m-1)N_0 L}$ such that $\mathbb{P}_{p,\sigma}(h_{(m-1)N_0 L}) > 0$, $p_{(m-1)N_0 L}(p, \sigma)(h_{(m-1)N_0 L})$ is in $A(q)$. Consequently, for any strategy τ of player 2, we have:

$$\forall m \geq 1, \ \mathbb{E}_{p,\sigma,\tau}\left(\frac{1}{N_0 L}\sum_{n \in B^m} G^{k_n}(i_n, j_n)\right) \geq w - \varepsilon.$$

Hence: $\exists N_1 \ \forall N \geq N_1 \forall \tau \in \mathcal{T}, \ \ \gamma_N^p(\sigma, \tau) \geq w - 2\varepsilon.$ □

8. Player 1 can guarantee cav$w(pB_0)$ in the original game. We investigate here what can be guaranteed by player 1 in the original game. By theorem 7.11, player 1 can guarantee $w(p) = \inf_{N \geq 1}\hat{v}_{NL}(pB_0)$ in the non revealing game $\hat{\Gamma}_\infty(p)$. Thus *a fortiori* player 1 can guarantee $w(p)$ in $\Gamma_\infty(p)$. We now define another mapping from $\Delta(K)$ to $\mathbb{R}$.

Definition 8.1 $\forall p \in \Delta(K), \ \ w^*(p) = \text{cav}w(pB_0).$

Notice that $(w(p) = w(pB_0) \ \forall p)$ does not necessarily imply $(\text{cav}w(p) = \text{cav}w(pB_0) \ \forall p)$. Recall that if f is an u.s.c. mapping from $\Delta(K)$ to $\mathbb{R}$, then cav$f$ is continuous and for each p in $\Delta(K)$, cav$f(p)$ is:

$$\max\{\sum_{s \in S}\lambda_s f(p_s), S \text{ finite set }, \forall s \in S \ \lambda_s \geq 0, p_s \in \Delta(K), \sum_{s \in S}\lambda_s = 1, \sum_{s \in S}\lambda_s p_s = p\}.$$

By Caratheodory's theorem, one can show (see e.g. Rockafellar [12, corollary 17.1.5 p.157]) that the finite set S in the above maximum can always be chosen with cardinality $|K|$. So we may impose $S = \{1, ..., |K|\} = K$ in the above expression.

Lemma 8.2 *Let f be an u.s.c. mapping from $\Delta(K)$ to $\mathbb{R}$ such that for each p in $\Delta(K)$, player 1 can guarantee $f(p)$ in $\Gamma_\infty(p)$. Then for each p in $\Delta(K)$, player 1 can guarantee $\text{cav}f(pB_0)$.*

Proof: By the splitting procedure (see remark 6.7, or e.g. Zamir [20]), it is standard that for all p player 1 can guarantee cav$f(p)$ in $\Gamma_\infty(p)$. As cav$f$ is continuous, to prove the lemma we show that if $g : \Delta(K) \longrightarrow \mathbb{R}$ is continuous and if for each p player 1 can guarantee $g(p)$ in $\Gamma_\infty(p)$, then for each p he can guarantee $g(pB_0)$ in $\Gamma_\infty(p)$.

Fix p in $\Delta(K)$, and $\varepsilon > 0$. For $N \geq 1$, let σ_N in Σ be a strategy of player 1 such that: $\exists T_0 \ \forall T \geq T_0$ $\forall \tau \in \mathcal{T}, \gamma_T^{pM^N}(\sigma_N, \tau) \geq g(pM^N) - \varepsilon$. Define $\sigma(N)$ as follows: play arbitrarily independently of the states up to stage N, then play according to σ_N. It is clear that $\sigma(N)$ guarantees $g(pM^N)$ - ε in $\Gamma_\infty(p)$. For N multiple of L and large enough, $\sigma(N)$ thus guarantees $g(pB_0) - 2\varepsilon$ in $\Gamma_\infty(p)$.   □

As player 1 can guarantee $w(p)$ in $\Gamma_\infty(p)$ for each p, the following corollary is immediate.

Corollary 8.3 *In $\Gamma_\infty(p)$, player 1 can guarantee $w^*(p)$.*

9. Value of the original game By proposition 5.5 and by corollary 8.3, we know that for each p in $\Delta(K)$, in the game $\Gamma_\infty(p)$:

player 2 can guarantee $v^*(p) = \lim_{n \to \infty}v_n(p) = \inf_{N \geq 1}v_{NL}(pB_0)$.

player 1 can guarantee $w^*(p) = \text{cav}w(pB_0),$

   where $w(p) = \lim_{n \to \infty}\hat{v}_n(p) = \inf_{N \geq 1}\hat{v}_{NL}(pB_0)$.

We finally show that $v^*(p) = w^*(p)$, proving that $\Gamma_\infty(p)$ has a value. This is done in two steps. First, we show that $v^*(p) = \inf_{N \geq 1}\text{cav}\hat{v}_{NL}(pB_0)$. Then we will prove that $\inf_{N \geq 1}\text{cav}\hat{v}_{NL}(pB_0) = \text{cav}w(pB_0)$.

The following proposition establishes the link between the limit of the values of the N-stage original games and the values of the N-stage non revealing games. Indeed, it justifies our definition of non revealing games.

Proposition 9.1 *For p in $\Delta(K)$ and $N \geq 1$, $v^*(p) \leq \mathrm{cav}\hat{v}_{NL}(pB_0)$.*

We first roughly explain the ideas of the proof. Fix N a multiple of L, we have to show that $v^*(p) \leq \mathrm{cav}\hat{v}_N(pB_0)$. Consider a strategy σ of player 1 which is optimal in a NT-stage game $\Gamma_{NT}(p)$, with T large. We construct a strategy τ for player 2 such that $\gamma_{NT}^p(\sigma, \tau) \leq f(N, T)$, where f satisfies $\limsup_T f(N, T) \leq \mathrm{cav}\hat{v}_N(pB_0)$. σ need not be non revealing at p, but since $(q_n(p, \sigma))_{n \geq 0}$ is a martingale, we know by the classical bound on its L_1 variation (see for example [20, p.122]), that for each S, $\frac{1}{S}\sum_{s=0}^{S-1} \mathbb{E}_{p,\sigma}(\|q_{s+1}(p, \sigma) - q_s(p, \sigma)\|) \leq \frac{|K|}{\sqrt{S}}$, hence is small when S is large. The NT stages will be viewed as T blocks of length N. At the beginning of each block m, player 2 will compute his belief p_m on the current state and the N-stage strategy σ_m to be played at this block by player 1 using σ. (p_m, σ_m) will be approximated by a non revealing pair $(\hat{p}_m, \hat{\sigma}_m)$, and player 2 will play at block B_m a best response against $\hat{\sigma}_m$ in the non revealing game $\hat{\Gamma}_N(\hat{p}_m)$. Then player 1's payoff at this block will be at most $\hat{v}_N(\hat{p}_m)$ plus an error term depending of $\|(p_m, \sigma_m) - (\hat{p}_m, \hat{\sigma}_m)\|$. By the previous bound, it will be possible to control the average error term and to show that it vanishes as T goes to infinity. And we obtain the upper bound of $\mathrm{cav}\hat{v}_N(pB_0)$ by collecting the expected average of the non revealing N-stage values. We now formally prove proposition 9.1.

Proof:

Let N be a positive multiple of L. N will remain fixed in all what follows, and we will prove that : $\forall p \in \Delta(K)$, $v^*(p) \leq \mathrm{cav}\hat{v}_N(pB_0)$.

(1) It will be convenient to consider the following subset $\bar{\Sigma}_N$ of N-stages behavior strategies of player 1. A strategy $\sigma = (\sigma_n)_{n=1,\dots,N}$ in Σ_N is in $\bar{\Sigma}_N$ if at each stage n, σ_n only depends on player 1's past actions $i_1,\dots,i_{n-1}$ and on the current state k_n. By proposition 5.1, for each p player 1 has an optimal strategy in $\Gamma_N(p)$ which belongs to $\bar{\Sigma}_N$.

We need to be precise about the strategy spaces and the norms we will use. Recall that if $x = (x^s)_{s \in \mathcal{S}}$ is an element of an Euclidean space $\mathbb{R}^{\mathcal{S}}$, we use $\|x\| = \sum_{s \in \mathcal{S}} |x^s|$. We view $\Delta(I)$ as a subset of $\mathbb{R}^I$. An element σ in $\bar{\Sigma}_N$ can then be seen as a particular mapping from $\cup_{n=1}^N (K \times I \times J)^{n-1} \times K$ to $\mathbb{R}^I$, hence as an element of the Euclidean space $(\mathbb{R}^I)^{\cup_{n=1}^N (K \times I \times J)^{n-1} \times K}$. One can find a positive constant $C_1 \geq C$ such that for all σ and σ' in $\bar{\Sigma}_N$, for each p in $\Delta(K)$ and τ in $\mathcal{T}_N$, we have:

$$|\gamma_N^p(\sigma, \tau) - \gamma_N^p(\sigma', \tau)| \leq C_1\|\sigma - \sigma'\|.$$

Notice that C_1 depends on N.

Denote by R the set $\Delta(K) \times \bar{\Sigma}_N$. If (p, σ) and (p', σ') are elements of R, we use: $\|(p, \sigma) - (p', \sigma')\| = \|p - p'\| + \|\sigma - \sigma'\|$, where $\|p - p'\| = \sum_{k \in K} |p^k - p'^k|$. R is thus compact.

As in the proof of proposition 7.3, we put $H = \cup_{n=0}^{N-1} \hat{H}_{n+1}$. We consider the following mapping F from R to $(\mathbb{R}^K)^H$ such that for each (p, σ) in $\Delta(K) \times \bar{\Sigma}_N$, $F(p, \sigma)$ is:

$$\left(\mathbb{P}_{p,\sigma}(h_{n+1})\left(q_{n+1}(p, \sigma)(h_{n+1}) - q_n(p, \sigma)(i_1, \dots, i_n)\right)\right)_{n=0,\dots,N-1, h_{n+1}=(i_1,\dots,i_{n+1}) \in \hat{H}_{n+1}}.$$

For (p, σ) in R, we have $F(p, \sigma) = 0$ if and only if $\sigma \in \hat{\Sigma}_N(p)$, and we define $\hat{R} = \{(p, \sigma) \in R, \ F(p, \sigma) = 0\}$. F is continuous as in the proof of proposition 7.3, so $\hat{R}$ is compact. Notice that for (p, σ) in R we have:

$$\|F(p, \sigma)\| = \sum_{n=0}^{N-1} \mathbb{E}_{p,\sigma}\left(\|q_{n+1}(p, \sigma) - q_n(p, \sigma)\|\right).$$

(2) Fix now $\varepsilon > 0$.

Assume that for each $t > 0$, one can find (p_t, σ_t) in R such that: $\forall(\hat{p}, \hat{\sigma}) \in \hat{R}$, $\|(p_t, \sigma_t) - (\hat{p}, \hat{\sigma})\| > \varepsilon + t\|F(p_t, \sigma_t)\|$. By compacity of R and continuity of F, one can consider a cluster point of the sequence $(p_t, \sigma_t)_t$ and obtain the existence of some element (p^*, σ^*) in R satisfying: $\forall(\hat{p}, \hat{\sigma}) \in \hat{R}, \|(p^*, \sigma^*) - (\hat{p}, \hat{\sigma})\| \geq \varepsilon > 0$, and $\|F(p^*, \sigma^*)\| = 0$. This is impossible by definition of $\hat{R}$. Consequently, one can find a positive constant C_2, depending on N and ε, such that:

$$\forall(p, \sigma) \in R, \ \exists(\hat{p}, \hat{\sigma}) \in \hat{R}, \ \|(p, \sigma) - (\hat{p}, \hat{\sigma})\| \leq \varepsilon + C_2\|F(p, \sigma)\|.$$

(3) Now comes the main part of the proof. N and ε being fixed as before, we also fix an element p in $\Delta(K)$ in all what follows. Let $C_1 = C_1(N)$ and $C_2 = C_2(N,\varepsilon)$ be defined as in points (1) and (2).

Let now T be a positive integer, and let σ in $\bar{\Sigma}_{NT}$ (i.e. σ is in Σ_{NT} and only depends on past actions and on the current state) be an optimal strategy for player 1 in $\Gamma_{NT}(p)$. $(q_n(p,\sigma))_{n \geq 1}$ is, with respect to $\mathbb{P}_{p,\sigma}$, a martingale with values in $\Delta(K)$, so we have the usual bound:

$$\frac{1}{NT} \sum_{n=0}^{NT-1} \mathbb{E}_{p,\sigma}\left(\|q_{n+1}(p,\sigma) - q_n(p,\sigma)\|\right) \leq \frac{|K|}{\sqrt{NT}}.$$

We define τ a NT-stage strategy of player 2 as follows. Divide the set of stages into consecutive blocks $B^1,..., B^T$ of length N. For $m \in \{1,...,T\}$, τ plays at block $B^m = \{(m-1)N+1,...,mN\}$ as follows: if $h_m = (i_1,...,i_{(m-1)N}) \in \hat{H}_{(m-1)N}$ has been played by player 1 at previous blocks, let $p_m(h_m) = p_{(m-1)N}(p,\sigma)(h_m)$ in $\Delta(K)$ be player 2's belief on the state of the beginning of block B^m. Consider the strategy $\sigma_m(h_m)$ played by player 1 using σ at block B^m after h_m has occurred. $\sigma_m(h_m)$ belongs to $\bar{\Sigma}_N$, so it is possible to choose some $(\hat{p}_m(h_m), \hat{\sigma}_m(h_m))$ in $\hat{R}$ such that: $\|(\hat{p}_m(h_m), \hat{\sigma}_m(h_m)) - (p_m(h_m), \sigma_m(h_m))\| \leq \varepsilon + C_2\|F(p_m(h_m), \sigma_m(h_m))\|$. At block B^m after h_m, τ plays a best response $\tau_m(h_m)$ against $\hat{\sigma}_m(h_m)$ in the non revealing game $\hat{\Gamma}_N(\hat{p}_m(h_m))$.

By definition of σ, $\gamma_{NT}^p(\sigma,\tau) \geq v_{NT}(p)$. We now compute an upper bound for $\gamma_{NT}^p(\sigma,\tau)$. Let m be in $\{1,...,T\}$ and h_m be in $\hat{H}_{(m-1)N}$.

$$\mathbb{E}_{p,\sigma,\tau}\left(\frac{1}{N} \sum_{n \in B^m} G^{k_n}(i_n, j_n)\big|h_m\right) = \gamma_N^{p_m(h_m)}(\sigma_m(h_m), \tau_m(h_m))$$

$$\leq C\|p_m(h_m) - \hat{p}_m(h_m)\| + \gamma_N^{\hat{p}_m(h_m)}(\sigma_m(h_m), \tau_m(h_m))$$

$$\leq C\|p_m(h_m) - \hat{p}_m(h_m)\| + \gamma_N^{\hat{p}_m(h_m)}(\hat{\sigma}_m(h_m), \tau_m(h_m)) + C_1\|(\sigma_m(h_m) - \hat{\sigma}(h_m)\|$$

$$\leq C_1\left(\varepsilon + C_2\|F(p_m(h_m), \sigma_m(h_m))\|\right) + \hat{v}_N(\hat{p}_m(h_m)),$$

since $\hat{\sigma}_m(h_m)$ is non revealing at $\hat{p}_m(h_m)$ and $\tau_m(h_m)$ is optimal in $\hat{\Gamma}_N(\hat{p}_m(h_m))$.

Consequently, if we condition on all possible histories h_m we obtain that at block B^m player 1's payoff $\mathbb{E}_{p,\sigma,\tau}\left(\frac{1}{N} \sum_{n \in B^m} G^{k_n}(i_n, j_n)\right)$ is at most:

$$\sum_{h_m \in \hat{H}_{(m-1)N}} \mathbb{P}_{p,\sigma}(h_m)\left(C_1\varepsilon + C_1C_2\|F(p_m(h_m), \sigma_m(h_m))\| + \hat{v}_N(\hat{p}_m(h_m))\right).$$

We have $\sum_{h_m \in \hat{H}_{(m-1)N}} \mathbb{P}_{p,\sigma}(h_m)p_m(h_m) = pM^{(m-1)N}$, and we put $\bar{p}_m =$

$\sum_{h_m \in \hat{H}_{(m-1)N}} \mathbb{P}_{p,\sigma}(h_m)\hat{p}_m(h_m)$. Since $\text{cav}\hat{v}_N$ is concave and above $\hat{v}_N$ we obtain that $\mathbb{E}_{p,\sigma,\tau}\left(\frac{1}{N} \sum_{n \in B^m} G^{k_n}(i_n, j_n)\right)$ is at most:

$$C_1\varepsilon + C_1C_2 \sum_{h_m \in \hat{H}_{(m-1)N}} \mathbb{P}_{p,\sigma}(h_m)\|F(p_m(h_m), \sigma_m(h_m))\| + \text{cav}\hat{v}_N(\bar{p}_m).$$

Summing up over blocks, we get:

$$\gamma_{NT}^p(\sigma,\tau) \leq C_1\varepsilon + C_1C_2\frac{1}{T} \sum_{m=1}^{T} \sum_{h_m \in \hat{H}_{(m-1)N}} \mathbb{P}_{p,\sigma}(h_m)\|F(p_m(h_m), \sigma_m(h_m))\| + \frac{1}{T} \sum_{m=1}^{T} \text{cav}\hat{v}_N(\bar{p}_m).$$

We now bound the second and third terms on the right-hand side.

First, for each m and h_m in $\hat{H}_{(m-1)N}$ we have:

$$\|F(p_m(h_m), \sigma_m(h_m))\| = \sum_{n=0}^{N-1} \mathbb{E}_{p_m(h_m), \sigma_m(h_m)}(\|q_{n+1}(p_m(h_m), \sigma_m(h_m)) - q_n(p_m(h_m), \sigma_m(h_m))\|)$$

$$= \sum_{n \in B^m} \mathbb{E}_{p,\sigma}(\|q_{n+1}(p,\sigma) - q_n(p,\sigma)\| \,|h_m))$$

Thus we get:

$$\frac{1}{T}\sum_{m=1}^{T}\sum_{h_m\in\hat{H}_{(m-1)N}} \mathbb{P}_{p,\sigma}(h_m)\|F(p_m(h_m),\sigma_m(h_m))\| = \frac{1}{T}\sum_{n=0}^{NT-1}\mathbb{E}_{p,\sigma}\left(\|q_{n+1}(p,\sigma)-q_n(p,\sigma)\|\right)$$

$$\leq \frac{|K|\sqrt{N}}{\sqrt{T}}.$$

Secondly, we have $\frac{1}{T}\sum_{m=1}^{T}\mathrm{cav}\hat{v}_N(\bar{p}_m)\leq \mathrm{cav}\hat{v}_N(\frac{1}{T}\sum_{m=1}^{T}\bar{p}_m)$ and :

$$\|\frac{1}{T}\sum_{m=1}^{T}\bar{p}_m - \frac{1}{T}\sum_{m=1}^{T}pM^{(m-1)N}\| = \|\frac{1}{T}\sum_{m=1}^{T}\sum_{h_m}\mathbb{P}_{p,\sigma}(h_m)(p_m(h_m)-\hat{p}_m(h_m))\|$$

$$\leq \frac{1}{T}\sum_{m=1}^{T}\sum_{h_m}\mathbb{P}_{p,\sigma}(h_m)(\varepsilon + C_2\|F(p_m(h_m),\sigma_m(h_m))\|)$$

$$\leq \varepsilon + C_2\frac{|K|\sqrt{N}}{\sqrt{T}}.$$

Summing up, we have obtained:

$$\gamma_{NT}^p(\sigma,\tau)\leq C_1\varepsilon + \frac{C_1C_2|K|\sqrt{N}}{\sqrt{T}} + \mathrm{cav}\hat{v}_N(\frac{1}{T}\sum_{m=1}^{T}\bar{p}_m)$$

where $\|\frac{1}{T}\sum_{m=1}^{T}\bar{p}_m - \frac{1}{T}\sum_{m=1}^{T}pM^{(m-1)N}\|\leq \varepsilon + C_2\frac{|K|\sqrt{N}}{\sqrt{T}}$.

$\mathrm{cav}\hat{v}_N$ being continuous, we then have that $v_{NT}(p)$ is at most:

$$C_1\varepsilon + \frac{C_1C_2|K|\sqrt{N}}{\sqrt{T}} + \max\left\{\mathrm{cav}\hat{v}_N(p'), p' \text{ s.t. } \|p'-\frac{1}{T}\sum_{m=1}^{T}pM^{(m-1)N}\|\leq \varepsilon + C_2\frac{|K|\sqrt{N}}{\sqrt{T}}\right\}.$$

Recall that N, ε, C_1 and C_2 are fixed. $\frac{1}{T}\sum_{m=1}^{T}pM^{(m-1)N}\longrightarrow_{T\to\infty}pB_0$, so one can find T_0 such that for each $T\geq T_0$:

$$v_{NT}(p)\leq C_1\varepsilon + \frac{C_1C_2|K|\sqrt{N}}{\sqrt{T}} + \max\{\mathrm{cav}\hat{v}_N(p'), p' \text{ s.t. } \|p'-pB_0\|\leq 3\varepsilon\}.$$

And we obtain:

$$v^*(p) = \limsup_{T}v_{NT}(p)\leq C_1\varepsilon + \max\{\mathrm{cav}\hat{v}_N(p'), p' \text{ s.t. } \|p'-pB_0\|\leq 3\varepsilon\}.$$

(4) We can now conclude the proof of proposition 9.1. The above inequality holds for each $\varepsilon > 0$. Since C_1 does not depend on ε and $\mathrm{cav}\hat{v}_N$ is continuous, we have $v^*(p)\leq \mathrm{cav}\hat{v}_N(pB_0)$ as wanted. $\qquad\square$

Corollary 9.2 *For each p in $\Delta(K)$,*

$$v^*(p) = \inf_{N\geq 1}\mathrm{cav}\hat{v}_{NL}(pB_0) = \lim_{N\to\infty}\mathrm{cav}\hat{v}_{NL}(pB_0).$$

proof: By the previous proposition, for each p we have $v^*(p)\leq \inf_{N\geq 1}\mathrm{cav}\hat{v}_{NL}(pB_0)$. For each N, v_N is concave and above $\hat{v}_N$, hence we obtain: $\inf_{N\geq 1}\mathrm{cav}\hat{v}_{NL}(pB_0)\geq v^*(p) = \lim_{N\to\infty}v_N(pB_0)\geq \limsup_{N\to\infty}\mathrm{cav}\hat{v}_N(pB_0)$. Thus $(\mathrm{cav}\hat{v}_{NL}(pB_0))_N$ converges to $\inf_{N\geq 1}\mathrm{cav}\hat{v}_{NL}(pB_0) = v^*(p)$. $\qquad\square$

It now just remains to show that $w^*(p) = \inf_{N\geq 1}\mathrm{cav}\hat{v}_{NL}(pB_0)$. We start with a lemma, which proof can be found in the Appendix.

Lemma 9.3 *Let $(f_n)_{n\geq 1}$ be a non-increasing sequence of u.s.c. mappings from $\Delta(K)$ to $\mathbb{R}$ pointwise converging to some mapping f from $\Delta(K)$ to $\mathbb{R}$. Then the sequence $(\mathrm{cav}f_n)_{n\geq 1}$ uniformly converges to $\mathrm{cav}f$.*

Proposition 9.4 *For each p in $\Delta(K)$,*

$$w^*(p) = \inf_{N \geq 1} \mathrm{cav} \hat{v}_{NL}(pB_0).$$

Proof: As $w(p) = \inf_{N \geq 1} \hat{v}_{NL}(pB_0)$, it is clear that $w^*(p) \leq \inf_{N \geq 1} \mathrm{cav} \hat{v}_{NL}(pB_0)$. We define, for each positive n, two mappings f_n and u_n from $\Delta(K)$ to $\mathrm{I\!R}$ such that for each p, $f_n(p) = \hat{v}_{2^n L}(pB_0)$ and $u_n(p) = \hat{v}_{2^n L}(p)$. Notice that $\mathrm{cav} f_n(p)$ may not be $\mathrm{cav}\hat{v}_{2^n L}(pB_0)$.

1. For each n, f_n is continuous and $(f_n)_n$ pointwise converges to w. By corollary 7.9, $(f_n)_n$ is non-increasing so by lemma 9.3, we first obtain that $(\mathrm{cav} f_n)_{n \geq 1}$ uniformly converges to $\mathrm{cav} f$. In particular, for each p in $\Delta(K)$ we have:

$\mathrm{cav} f_n(p) \longrightarrow_{n \to \infty} \mathrm{cav} w(p)$.

2. By proposition 7.8, we also have: $\forall p \in \Delta(K), \forall N \geq 1, \forall T \geq 1$:

$$\hat{v}_{NLT}(p) \leq \frac{1}{T} \sum_{t=1}^{T} \hat{v}_{NL}(pM^{(t-1)NL}).$$

$\hat{v}_{NL}$ being concave on $A(pB_0)$, we obtain:

$$\hat{v}_{NLT}(p) \quad \leq \quad \hat{v}_{NL}\left(\frac{1}{T} \sum_{t=1}^{T} pM^{(t-1)NL} \right).$$

Hence: $\forall N \geq 1, \forall T \geq 1$,

$$\hat{v}_{L2^N 2^T}(p) \quad \leq \quad \hat{v}_{L2^N}\left(\frac{1}{2^T} \sum_{t=1}^{2^T} pM^{(t-1)2^N L} \right),$$

$$u_{N+T}(p) \quad \leq \quad u_N\left(\frac{1}{2^T} \sum_{t=1}^{2^T} pM^{(t-1)2^N L} \right).$$

Fix $N \geq 1$ and $\varepsilon > 0$. u_N being uniformly continuous, one can find T_0 such that: $\forall T \geq T_0$, $\forall p \in \Delta(K)$, $u_{N+T}(p) \leq u_N(pB_0) + \varepsilon = f_N(p) + \varepsilon$. Thus for each p in $\Delta(K)$ and $T \geq T_0$, we have $\mathrm{cav} u_{N+T}(p) \leq \mathrm{cav} f_N(p) + \varepsilon$, and so $\mathrm{cav} u_{N+T}(pB_0) \leq \mathrm{cav} f_N(pB_0) + \varepsilon$. Hence:

$$\mathrm{cav} f_N(pB_0) + \varepsilon \quad \geq \quad \inf_T \mathrm{cav} u_{N+T}(pB_0)$$
$$\geq \quad \inf_T \mathrm{cav} \hat{v}_{TL}(pB_0).$$

We have obtained: $\forall N \geq 1$, $\forall \varepsilon > 0$, $\forall p \in \Delta(K)$, $\inf_T \mathrm{cav} \hat{v}_{TL}(pB_0) \leq \mathrm{cav} f_N(pB_0) + \varepsilon$. Thus $\inf_T \mathrm{cav} \hat{v}_{TL}(pB_0) \leq \inf_N \mathrm{cav} f_N(pB_0)$, and this last quantity is just $\mathrm{cav} w(pB_0)$ by point 1. Consequently, $w^*(p) = \inf_{N \geq 1} \mathrm{cav} \hat{v}_{NL}(pB_0)$. $\qquad\qquad\Box$

It just remains to conclude.

Proof of Theorem 2.3 By proposition 5.5, player 2 can guarantee $v^*(p)$ in $\Gamma_\infty(p)$. By corollary 8.3, player 1 can guarantee $w^*(p) = \mathrm{cav} w(pB_0)$ in $\Gamma_\infty(p)$. By corollary 9.2 and proposition 9.4, $v^*(p) = \inf_{N \geq 1} \mathrm{cav} \hat{v}_{NL}(pB_0) = w^*(p)$. Hence $\Gamma_\infty(p)$ has a value which is:

$$v^*(p) = w^*(p) = \mathrm{cav} w(pB_0) = \inf_{N \geq 1} \mathrm{cav} \hat{v}_{NL}(pB_0).$$

And player 2 has an optimal strategy by remark 5.7. $\qquad\qquad\Box$

10. Concluding Remarks.

1. Observation of player 1

A) The fact that player 1 observes after each stage the action played by player 2 plays no role. Hence if player 1 does not observe these actions, or just imperfectly observes them, the game still has the same value.

B) Assume now that player 1 observes at the beginning of stage 1 the whole sequence of states k_1, $k_2,..., k_n,....$ Then the N-stage values $v_N(p)$ still satisfy the same recursive formula given by proposition 5.1, and player 2 can still guarantee their limit as shown by proposition 5.5. Hence again this modified game has the same value (*idem* if player 1 observes the sequence of states in any manner such that, for each n, player 1 knows the state k_n before choosing his action of stage n).

2. Optimal strategy for player 1

By slightly improving the proof of theorem 7.11, one can show that player 1 has a strategy σ that guarantees $w(p)$ in $\hat{\Gamma}_\infty(p)$ (just play by consecutive blocks B^m of length mL, and at B^m play an optimal strategy in $\hat{\Gamma}_{mL}(p_m)$, if p_m is the belief of player 2 on the current state at the beginning of the block). w being u.s.c., by the splitting procedure one can construct σ that guarantees $\text{cav}w(p)$ in $\Gamma_\infty(p)$. But it is unclear whether there exists σ that guarantees the value $\text{cav}w(pB_0)$ for player 1, hence the existence of an optimal strategy for player 1 in $\Gamma_\infty(p)$ is an open question.

3. Values of the N-stage games

A) The proof of proposition 9.1 gives no useful bound on the speed of convergence of $(v_N(p))_N$ to $v^*(p)$, contrary to the Aumann-Maschler case where this convergence is according to $1/\sqrt{N}$.

B) It is unclear whether $\hat{v}_N$, even $\hat{v}_1$, is Lipschitz or not. A proof of this would simplify things and give insights for the previous points.

4. Computing the value

The value is not easy to compute, even in a simple example such as: $K = \{a, b\}$, $M = \begin{pmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{pmatrix}$, $G^a = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ and $G^b = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$.

Very recently, A. Marino [8] and Horner *et al.* [6] gave independent and different proofs for this case, obtaining that: $v_\infty(p) = 0.4$ for each p. Keeping the same payoffs, Horner *et al.* compute the value for any transition matrix $M = \begin{pmatrix} \alpha & 1-\alpha \\ 1-\alpha & \alpha \end{pmatrix}$, with $1/2 \leq \alpha \leq 2/3$. The case $\alpha = 3/4$ appears more difficult, and now seems a challenging example.

5. Generalizations

An interesting generalization is the case where player 2 observes after each stage a signal depending on the action just played by player 1 and on the current state. Is it possible, as in the Aumann Maschler's case, to generalize to this setup the definition of non revealing strategies in order to get the existence of the value ?

Another generalization, focussing on the stochastic game aspect, is the case where at each stage the current state is chosen according to a probability distribution depending on the previous state *and* on the actions just played by both players. However, it is known (Sorin [18]) that the value may fail to exist in this general case. But in the intermediate case where the transitions do not depend on player 2's actions, what about the existence of the value ?

Appendix A.

Proof of proposition 5.1 It goes by induction on n. The result holds for $n = 1$ by definition of $\Gamma_1(p)$. Fix $n \geq 2$, and assume that the proposition holds for $n - 1$ and all p. Fix p in $\Delta(K)$ and consider the game $\Gamma_n(p)$.

We define an auxiliary zero-sum game $A_n(p)$ with strategy spaces $\Delta(I)^K$ for player 1 and $\Delta(J)$ for player 2, and payoff function for player 1 defined by: $f_n^p(x, y) = \frac{1}{n}(G(p, x, y) + (n - 1)\sum_{i \in I} x(p)(i)v_{n-1}(\hat{p}(x, i)M))$ for all x in $\Delta(I)^K$ and y in $\Delta(J)$.

Consider $x = \lambda x' + (1-\lambda)x''$, with $\lambda \in [0,1]$ and x', x'' in $\Delta(I)^K$. For each i, we have $x(p)(i)\,\hat{p}(x,i) = \lambda x'(p)(i)\,\hat{p}(x',i) + (1-\lambda)x''(p)(i)\,\hat{p}(x'',i)$, hence by concavity of v_{n-1},
$$x(p)(i)\,v_{n-1}(\hat{p}(x,i)M) \geq \lambda x'(p)(i)\,v_{n-1}(\hat{p}(x',i)M) + (1-\lambda)x''(p)(i)\,v_{n-1}(\hat{p}(x'',i)M).$$

This shows that f_n^p is concave in x. As it is convex in y and continuous, by Sion's theorem [16], $A_n(p)$ has a value that we denote by $f_n(p)$.

We now formally prove that player 1 can guarantee $f_n(p)$ in $\Gamma_n(p)$. Let σ in Σ_n be as follows: - at stage 1, play some x^* optimal for player 1 in $A_n(p)$. - if i in I is the action played at stage 1, play from stage 2 to stage n an optimal strategy σ_i for player 1 in the game $\Gamma_{n-1}(\hat{p}(x,i)M)$. Let τ be in $\mathcal{T}$, and denote by y in $\Delta(J)$ the mixed action played by τ at stage 1 and for each $(i,j) \in I \times J$, by $\tau_{i,j}$ the strategy played by player 2 using τ at stages 2,...,n if (i,j) is played at stage 1. We have:
$$\gamma_n^p(\sigma,\tau) = \frac{1}{n}\left(G(p,x^*,y) + (n-1)E_{p,\sigma,\tau}\left(\frac{1}{n-1}\sum_{n'=2}^n G^{k_{n'}}(i_{n'},j_{n'})\right)\right)$$

and $E_{p,\sigma,\tau}\left(\frac{1}{n-1}\sum_{n'=2}^n G^{k_{n'}}(i_{n'},j_{n'})\right) =$

$$\sum_{i,j,k} x^*(p)(i)y(j)P_{p,\sigma,\tau}(k_2 = k|i_1 = i)E_{p,\sigma,\tau}\left(\left.\frac{1}{n-1}\sum_{n'=2}^n G^{k_{n'}}(i_{n'},j_{n'})\right| k_2 = k, i_1 = i, j_1 = j\right)$$

$= \sum_{i \in I} x^*(p)(i)\sum_{j \in J} y(j)\gamma_{n-1}^{\hat{p}(x^*,i)M}(\sigma_i,\tau_{i,j})$,

since $(P_{p,\sigma,\tau}(k_2 = k|i_1 = i))_{k \in K} = \hat{p}(x^*,i)M$ for all i.

For each i and j we have by definition of σ_i: $\gamma_{n-1}^{\hat{p}(x^*,i)M}(\sigma_i,\tau_{i,j}) \geq v_{n-1}(\hat{p}(x^*,i)M)$. Hence $\gamma_n^p(\sigma,\tau) \geq f_n^p(x^*,y) \geq f_n(p)$, and $v_n(p) \geq f_n(p)$.

It is similarly possible to show that player 2 can defend $f_n(p)$ in $\Gamma_n(p)$. Fix σ in Σ, and denote by x in $\Delta(I)^K$ the strategy induced by σ at stage 1. Define τ as follows: - at stage 1, play y^* in $\Delta(J)$ such that: $G(p,x,y^*) = \min_{y \in \Delta(J)} G(p,x,y)$, -at stages 2 to n, if i has been played at stage 1 by player 1, play τ_i optimal in the game $\Gamma_{n-1}(\hat{p}(x,i)M)$. Similar computations as before show that $\gamma_n^p(\sigma,\tau) \leq f_n(p)$. Finally $v_n(p) \leq f_n(p)$, thus $v_n(p) = f_n(p)$.

An optimal strategy for player 1 in $\Gamma_n(p)$ can be constructed as follows: at each stage n', compute the belief p' of player 2 on the current state (depending on the actions $i_1, ..., i_{n'-1}$ previously played by player 1), and play according to some x optimal in the zero-sum game $A_{n-n'+1}(p')$. This strategy only depends on his own past actions $i_1, ..., i_{n'-1}$ and on the current state $k_{n'}$. □

Proof of lemma 6.9

$\Longleftarrow$ just take $n = 0$.

$\Longrightarrow$ Assume that σ is in $\hat{\Sigma}(p)$, let n be in N, and let $h_n = (i_1, ..., i_n) \in I^n$ be such that $P_{p,\sigma}(h_n) > 0$. We have, by definition of $\hat{\Sigma}(p)$, $p_n(p,\sigma)(h_n)B_{-n} = pB_0$, and have to show that $\sigma(p,h_n) \in \hat{\Sigma}(p_n(p,\sigma)(h_n))$. We put for simplicity $\sigma' = \sigma(p,h_n)$ and $p' = p_n(p,\sigma)(h_n)$.

Fix m in N, and $h(m) = (i_{n+1}, ..., i_{n+m})$ in I^m. Put $\omega = (i_1, ..., i_n, i_{n+1}, ..., i_{n+m})$, $A = (\tilde{i}_{n+1} = i_{n+1}, ..., \tilde{i}_{n+m} = i_{n+m})$, and $B = (\tilde{i}_1 = i_{n+1}, ..., \tilde{i}_m = i_{n+m})$.

Assume that $P_{p',\sigma'}(B) > 0$. We have to show that $p_m(p',\sigma')(h(m))B_{-m} = p'B_0$. $P_{p,\sigma}(\omega) = P_{p,\sigma}(h_n)P_{p,\sigma}(A|h_n) = P_{p,\sigma}(h_n)P_{p',\sigma'}(B) > 0$. Since σ is in $\hat{\Sigma}(p)$, $p_{n+m}(p,\sigma)(\omega)B_{-n-m} = pB_0$. For each k in K,
$$\begin{aligned} P_{p,\sigma}(\omega, \tilde{k}_{n+m+1} = k) &= P_{p,\sigma}(h_n)P_{p,\sigma}(A, \tilde{k}_{n+m+1} = k|h_n) \\ &= P_{p,\sigma}(h_n)P_{p',\sigma'}(B, \tilde{k}_{m+1} = k) \end{aligned}$$

So we obtain:
$$p_{n+m}^k(p,\sigma)(\omega) = \frac{P_{p',\sigma'}(B, \tilde{k}_{m+1} = k)}{P_{p',\sigma'}(B)} = p_m^k(p',\sigma')(h(m)).$$

Hence $p_m(p',\sigma')(h(m))B_{-n-m} = p_{n+m}(p,\sigma)(\omega)B_{-n-m} = pB_0 = p'B_{-n}$. Multiplying by M^n both sides, we obtain that $p_m(p',\sigma')(h(m))B_{-m} = p'B_0$. □

Proof of lemma 6.10 $\Longleftarrow$ is clear. We prove the $\Longrightarrow$ part.

Let n be in $I\!\!N$. Put $p' = pM^n$, and $\sigma' = \sigma(p, n+)$. We have:

$$\mathbb{P}_{p',\sigma'} = \sum_{h_n \in I^n} \mathbb{P}_{p,\sigma}(h_n) \mathbb{P}_{p_n(p,\sigma)(h_n),\sigma(p,h_n)}.$$
$$p' = \sum_{h_n \in I^n} \mathbb{P}_{p,\sigma}(h_n) p_n(p,\sigma)(h_n).$$

For each h_n such that $\mathbb{P}_{p,\sigma}(h_n) > 0$, by the previous lemma $\sigma(p, h_n) \in \hat{\Sigma}(p_n(p,\sigma)(h_n))$ and since $\sigma \in \hat{\Sigma}(p)$, $p_n(p,\sigma)(h_n) B_{-n} = pB_0$. Hence $p_n(p,\sigma)(h_n) B_0 = p_n(p,\sigma)(h_n) B_{-n} M^n = pM^n B_0$, and $p_n(p,\sigma)(h_n) \in A(pM^n B_0)$. By lemma 6.6, part *(1)*, $\sigma(p, n+) \in \hat{\Sigma}(pM^n)$. $\qquad\square$

Proof of proposition 7.4 For $n = 1$, the result is clear since a strategy σ in $\hat{\Sigma}_1(p)$ can be seen as an element of $NR(p)$ by proposition 6.5. Fix $n \geq 2$, and assume that the proposition holds for $n - 1$ for all p. Fix p in $\Delta(K)$.

Define the auxiliary zero-sum game $\hat{A}_n(p)$ with strategy spaces $NR(p)$ for player 1 and $\Delta(J)$ for player 2, and payoff function for player 1 defined by: $\hat{f}_n^p(x, y) = \frac{1}{n}(G(p, x, y) + (n-1) \sum_{i \in I} x(p)(i) \hat{v}_{n-1}(\hat{p}(x, i)M))$ for all x in $NR(p)$ and y in $\Delta(J)$. We are going to apply Sion's theorem to prove that $\hat{A}_n(p)$ has a value.

x being fixed, $(y \longrightarrow \hat{f}_n^p(x, y))$ is affine hence convex and continuous. Consider $x = \lambda x' + (1 - \lambda)x''$, with $\lambda \in [0, 1]$ and x', x'' in $NR(p)$. For each i s.t. $x'(p)(i) > 0$ and $x''(p)(i) > 0$, we have $\hat{p}(x, i)M = \frac{\lambda x'(p)(i)}{x(p)(i)} \hat{p}(x', i)M + \frac{(1-\lambda)x''(p)(i)}{x(p)(i)} \hat{p}(x'', i)M$, and by definition of $NR(p)$, $\hat{p}(x', i)B_0 M = pB_0 M = \hat{p}(x'', i)B_0 M$. Since $B_0 M = MB_0$, we obtain that $\hat{p}(x', i)M$, $\hat{p}(x'', i)M$ and $\hat{p}(x, i)M$ all belong to $A(pMB_0)$. By concavity of $\hat{v}_{n-1}$ on $A(pMB_0)$,
$$x(p)(i)\,\hat{v}_{n-1}(\hat{p}(x, i)M) \geq \lambda x'(p)(i)\,\hat{v}_{n-1}(\hat{p}(x', i)M) + (1 - \lambda)x''(p)(i)\,\hat{v}_{n-1}(\hat{p}(x'', i)M).$$

This proves that $\hat{f}_n^p$ is concave in x. Consider now a sequence $(x_t)_t$ of elements in $NR(p)$ converging to some x. For each i s.t. $x(p)(i) > 0$, $\hat{p}(x_t, i) \longrightarrow_{t \to \infty} \hat{p}(x, i)$ hence, since $\hat{v}_{n-1}$ is u.s.c., $\limsup_t \sum_{i \in I} x(p_t)(i)\hat{v}_{n-1}(\hat{p}(x_t, i)M) \leq \sum_{i \in I} x(p)(i)\hat{v}_{n-1}(\hat{p}(x, i)M)$. This shows that y being fixed, $(x \longrightarrow \hat{f}_n^p(x, y))$ is upper semi-continuous. By Sion's theorem, $\hat{A}_n(p)$ has a value that we denote by $\hat{f}_n(p)$.

We now show that $\hat{f}_n(p)$ is the value of $\hat{\Gamma}_n(p)$. As in the proof of proposition 5.1, we define σ in Σ_N as follows: - at stage 1, play some x^* optimal for player 1 in $\hat{A}_n(p)$. - from stage 2 to n, if i in I is the action played at stage 1, play an optimal strategy σ_i for player 1 in the game $\hat{\Gamma}_{n-1}(\hat{p}(x, i)M)$. Using proposition 6.5, one can see that σ is non revealing at p. As in the proof of proposition 5.1, σ guarantees $\hat{f}_n(p)$ in $\hat{\Gamma}_n(p)$. Similarly, we show that player 2 can defend $\hat{f}_n(p)$ in $\hat{\Gamma}_n(p)$. Fix σ in $\hat{\Sigma}_N(p)$, and denote by x in $NR(p)$ the strategy induced by σ at stage 1. For each i in I s.t. $x(p)(i) > 0$, the strategy induced by σ at stage 2 to n if i has been played at stage 1 is in $\hat{\Sigma}_{n-1}(\hat{p}(x, i)M)$ by lemma 6.9. Consequently one can construct, as in the proof of proposition 5.1, a strategy τ for player 2 satisfying $\gamma_n^p(\sigma, \tau) \leq \hat{f}_n(p)$. $\qquad\square$

Proof of lemma 9.3 Fix p in $\Delta(K)$. It is plain that $\lim_n \mathrm{cav} f_n(p) \geq \mathrm{cav} f(p)$. One can take a fixed finite set S, e.g. $S = \{1, ..., |K|\}$, satisfying: for each $n \geq 1$, there exists $(p_s^n)_{s \in S}$, $(\lambda_s^n)_{s \in S}$ such that for each s in S, $p_s^n \in \Delta(K)$, $\lambda_s^n \geq 0$, $\sum_{s \in S} \lambda_s^n = 1$, $\sum_{s \in S} \lambda_s^n p_s^n = p$ and $\mathrm{cav} f_n(p) = \sum_{s \in S} \lambda_s^n f_n(p_s^n)$. Taking converging subsequences, one can find an increasing mapping ψ from the set of positive integers to itself, $(\lambda_s)_{s \in S}$ and $(p_s)_{s \in S}$ s.t. for each s in S, $\lambda_s^{\psi(n)} \longrightarrow_{n \to \infty} \lambda_s$ and $p_s^{\psi(n)} \longrightarrow_{n \to \infty} p_s$.

Then for each $n_0 \geq 1$ we have for each $n \geq n_0$:
$$\begin{aligned}
\mathrm{cav} f_{\psi(n)}(p) &= \sum_{s \in S} \lambda_s^{\psi(n)} f_{\psi(n)}(p_s^{\psi(n)}) \\
&\leq \sum_{s \in S} \lambda_s^{\psi(n)} f_{n_0}(p_s^{\psi(n)}),
\end{aligned}$$

since $(f_n)_{n \geq 1}$ is non-increasing. So $\inf_{n \geq 1} \mathrm{cav} f_{\psi(n)}(p) \leq \sum_{s \in S} \lambda_s f_{n_0}(p_s)$, since f_{n_0} is u.s.c.

Hence $\mathrm{cav} f(p) \leq \inf_{n \geq 1} \mathrm{cav} f_n(p) \leq \sum_{s \in S} \lambda_s f_{n_0}(p_s)$ for each $n_0 \geq 1$. Thus $\mathrm{cav} f(p) \leq \inf_{n \geq 1} \mathrm{cav} f_n(p) \leq \sum_{s \in S} \lambda_s f(p_s) \leq \mathrm{cav} f(p)$, and $\inf_n \mathrm{cav} f_n(p) = \mathrm{cav} f(p)$.

We have obtained that $(\mathrm{cav} f_n)_{n \geq 1}$ pointwise converges to $\mathrm{cav} f$. For each n, f_n is u.s.c. so $\mathrm{cav} f_n$

is continuous. f is u.s.c as an infimum of u.s.c. mappings, so cav$f$ also is continuous. $\Delta(K)$ being a polytope, the convergence of $(\text{cav} f_n)_{n \geq 1}$ to cav$f$ is uniform (see e.g. Mertens *et al* [9, part A, p.46, ex. 15]). $\square$

References

[1] J. P. Aubin and A. Cellina, *Differential Inclusions*, Springer-Verlag, 1984.

[2] R. J. Aumann, *Mixed and behavior strategies in infinite extensive games*, In Dresher, Shapley and Tucker, 1964, pp. 627–650.

[3] R. J. Aumann and M. Maschler, *Repeated games with incomplete information*, with the collaboration of R. Stearns, Cambridge, MA: MIT Press, 1995.

[4] P. Gordon,*Theory of finite Markov chains*, in French, Dunod, Paris, 1964.

[5] S. Hart, *Nonzero-sum two-person repeated games with incomplete information*, Mathematics of Operations Research **10** (1985), 117-153.

[6] J. Horner, D. Rosenberg, E. Solan, and N. Vieille, *On the value of a specific Markov game*, mimeo, 2005.

[7] E. Kohlberg, *Optimal strategies in repeated games with incomplete information*, International Journal of Game Theory **4** (1975), 7-24.

[8] A. Marino, *The value and optimal strategies of a particular Markov chain game*, preprint, University Paris I, 2005.

[9] J. F. Mertens, S. Sorin and S. Zamir, *Repeated Games*, Parts A and B, CORE DP 9420-9421, 1994.

[10] J. F. Mertens and S. Zamir, *The value of two-person zero-sum repeated games with lack of information on both sides*, International Journal of Game Theory **1** (1971-1972), 39-64.

[11] J. R. Norris, *Markov chains*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 1997.

[12] R. T. Rockafellar, *Convex analysis*, Princeton University Press, 1970.

[13] D. Rosenberg, E. Solan and N. Vieille, *Stochastic games with a single controller and incomplete information*, SIAM Journal on Control and Optimization **43** (2002), 86-110.

[14] D. Rosenberg, and N. Vieille, *The maxmin of recursive games with lack of information on one side*, Mathematics of Operations Research **25** (2000), 23-35.

[15] R. S. Simon, S. Spież and H. Toruńczyk, *The existence of equilibria in certain games, separation for families of convex functions and a theorem of Borsuk-Ulam type*, Israel Journal of Mathematics **92** (1995), 1-21.

[16] M. Sion, *On general minimax theorems*, Pacific Journal of Mathematics **8** (1995), 171-176.

[17] S. Sorin, *Some results on the existence of Nash equilibrium for non-zero sum games with incomplete information*, International Journal of Game Theory **12** (1983), 193-205.

[18] S. Sorin, *Big Match with lack of information on one side (Part I)*, International Journal of Game Theory **13** (1984), 173-204.

[19] S. Sorin, *A first course on zero-sum repeated games*, Mathématiques & Applications, Springer, 2002.

[20] S. Zamir, *Repeated Games of Incomplete Information: Zero-Sum*, Handbook of Game Theory with Economic Applications, edited by Aumann and Hart. Vol I, chapter 6.