HAL
open science

Duality and optimal strategies in the finitely repeated zero-sum games with incomplete information on both sides

Bernard de Meyer, Alexandre Marino

To cite this version:

Bernard de Meyer, Alexandre Marino. Duality and optimal strategies in the finitely repeated zero-sum games with incomplete information on both sides. 2005. halshs-00193996

HAL Id: halshs-00193996

https://shs.hal.science/halshs-00193996

Submitted on 5 Dec 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ PARIS 1 – PANTHÉON - SORBONNE –

CERMSEM

UMR CNRS 8095

Cahiers de la MSE

Duality and optimal strategies in the finitely repeated zero-sum games with incomplete information on both sides

Bernard DE MEYER, CERMSEM

Alexandre MARINO, CERMSEM

2005.27

CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

Maison des Sciences Économiques, 106-112 boulevard de L'Hôpital, 75647 Paris Cedex 13

http://mse.univ-paris1.fr/Publicat.htm

ISSN : 1624-0340

DUALITY AND OPTIMAL STRATEGIES IN THE FINITELY REPEATED ZERO-SUM GAMES WITH INCOMPLETE INFORMATION ON BOTH SIDES

BERNARD DE MEYER & ALEXANDRE MARINO

MARCH 20, 2005

Abstract. The recursive formula for the value of the zero-sum repeated games with incomplete information on both sides is known for a long time. As it is explained in the paper, the usual proof of this formula is in a sense non constructive: it just claims that the players are unable to guarantee a better payoff than the one prescribed by the formula, but it does not indicates how the players can guarantee this amount.

In this paper we aim to give a constructive approach to this formula using duality techniques. This will allow us to recursively describe the optimal strategies in those games and to apply these results to games with infinite action spaces.

1. Introduction

This paper is devoted to the analysis of the optimal strategies in the repeated zero-sum game with incomplete information on both sides in the independent case. These games were introduced by Aumann, Maschler [1] and Stearns [7]. The model is described as follows: At an initial stage, nature choses as pair of states (k, l) in $(K \times L)$ with two independent probability distributions p, q on K and L respectively. Player 1 is then informed of k but not of l while, on the contrary, player 2 is informed of l but not of k. To each pair (k, l) corresponds a matrix $A_k^l := [A_{k,i}^{l,j}]_{i,j}$ in $\mathbb{R}^{I\times J}$, where I and J are the respective action sets of player 1 and 2, and the game A_k^l is the played during n consecutive rounds: at each stage $m = 1, \ldots, n$, the players select simultaneously an action in their respective action set: $i_m \in I$ for player 1 and $j_m \in J$ for player 2. The pair (i_m, j_m) is then publicly announced before proceeding to the next stage. At

2000 *Mathematics Subject Classification.* 91A05, 91A10, 91A15, 91A20.

Key words and phrases. repeated games, dual games, incomplete information, recurrence formula.

address: CERMSEM, Université Paris 1-Panthéon Sorbonne
Maison des Sciences Economiques — 106-112, Bd de l'Hôpital, 75013 Paris .

the end of the game, player 2 pays $\sum_{m=1}^{n} A_{k,i_m}^{l,j_m}$ to player 1. The previous description is common knowledge to both players, including the probabilities p, q and the matrices A_k^l.

The game thus described is denoted $G_n(p,q)$.

Let us first consider the finite case where K, L, I, and J are finite sets. For a finite set I, we denote by $\Delta(I)$ the set of probability distribution on I. We also denote by h_m the sequence $(i_1, j_1, \ldots, i_m, j_m)$ of moves up to stage m so that $h_m \in H_m := (I \times J)^m$.

A behavior strategy σ for player 1 in $G_n(p,q)$ is then a sequence $\sigma = (\sigma_1, \ldots, \sigma_n)$ where $\sigma_m : K \times H_{m-1} \to \Delta(I)$. $\sigma_m(k, h_{m-1})$ is the probability distribution used by player 1 to select his action at round m, given his previous observations (k, h_{m-1}). Similarly, a strategy τ for player 2 is a sequence $\tau = (\tau_1, \ldots, \tau_n)$ where $\tau_m : L \times H_{m-1} \to \Delta(J)$. A pair (σ, τ) of strategies, join to the initial probabilities (p,q) on the sates of nature induces a probability $\Pi^n_{(p,q,\sigma,\tau)}$on $(K \times L \times H_n)$. The payoff of player 1 in this game is then:

$$g_n(p,q,\sigma,\tau) := E_{\Pi^n_{(p,q,\sigma,\tau)}}[\sum_{m=1}^{n} A_{k,i_m}^{l,j_m}],$$

where the expectation is taken with respect to $\Pi^n_{(p,q,\sigma,\tau)}$. We will define $\underline{V}_n(p,q)$ and $\overline{V}_n(p,q)$ as the best amounts guaranteed by player 1 and 2 respectively:

$$\underline{V}_n(p,q) = \sup_{\sigma} \inf_{\tau} g_n(p,q,\sigma,\tau) \text{ and } \overline{V}_n(p,q) = \inf_{\tau} \sup_{\sigma} g_n(p,q,\sigma,\tau)$$

The functions $\underline{V}_n$ and $\overline{V}_n$ are continuous, concave in p and convex in q. They satisfy to $\underline{V}_n(p,q) \leq \overline{V}_n(p,q)$. In the finite case, it is well known that, the game $G_n(p,q)$ has a value $V_n(p,q)$ wich means that $\underline{V}_n(p,q) = \overline{V}_n(p,q) = V_n(p,q)$. Futhermore both players have optimal behavior strategies σ^* and τ^*:

$$\underline{V}_n(p,q) = \inf_{\tau} g_n(p,q,\sigma^*,\tau) \text{ and } \overline{V}_n(p,q) = \sup_{\sigma} g_n(p,q,\sigma,\tau^*)$$

Let us now turn to the recursive structure of $G_n(p,q)$: a strategy $\sigma = (\sigma_1, \ldots, \sigma_{n+1})$ in $G_{n+1}(p,q)$ may be seen as a pair (σ_1, σ^+) where

$$\sigma^+ = (\sigma_2, \ldots, \sigma_{n+1})$$

is in fact a strategy in a game of lenght n depending on the first moves (i_1, j_1). Similarly, a strategy τ for player 2 is viewed as $\tau = (\tau_1, \tau^+)$.

Let us now consider the probability π (resp. λ) on $(K \times I)$ (resp. $(L \times J)$) induced by (p, σ_1) (resp. (q, τ_1)). Let us denote by s the marginal distribution of π on I and let p^{i_1} be the conditional probability on K given i_1. Similarly, let t the marginal distribution of λ on J and let q^{j_1} be the conditional probability on L given j_1.

The payoff $g_{n+1}(p,q,\sigma,\tau)$ may then be computed as follows: the expectation of the first stage payoff is just $g_1(p,q,\sigma_1,\tau_1)$. Conditioned on i_1, j_1, the expectation of the n following terms is just $g_n(p^{i_1}, q^{j_1}, \sigma^+(i_1,j_1), \tau^+(i_1,j_1))$. Therefore:

$$(1)\quad g_{n+1}(p,q,\sigma,\tau) = g_1(p,q,\sigma_1,\tau_1) + \sum_{i_1,j_1} s_{i_1} t_{j_1} g_n(p^{i_1}, q^{j_1}, \sigma^+(i_1,j_1), \tau^+(i_1,j_1)).$$

At a first glans, if σ, τ are optimal in $G_{n+1}(p,q)$, this formula suggests that $\sigma^+(i_1,j_1)$ and $\tau^+(i_1,j_1)$ should be optimal strategies in $G_n(p^{i_1}, q^{j_1})$, leading to the following recursive formula :

Theorem 1.1.

$$V_{n+1} = \underline{T}(V_n) = \overline{T}(V_n)$$

with the recursive operators $\underline{T}$ and $\overline{T}$ defined as follows:

$$\underline{T}(f)(p,q) = \sup_{\sigma_1} \inf_{\tau_1} \left\{ g_1(p,q,\sigma_1,\tau_1) + \sum_{i_1,j_1} s_{i_1} t_{j_1} f(p^{i_1}, q^{j_1}) \right\}$$

$$\overline{T}(f)(p,q) = \inf_{\tau_1} \sup_{\sigma_1} \left\{ g_1(p,q,\sigma_1,\tau_1) + \sum_{i_1,j_1} s_{i_1} t_{j_1} f(p^{i_1}, q^{j_1}) \right\}$$

The usual proof of this theorem is as follows: When playing a best reply to a strategy σ of player 1 in $G_{n+1}(p,q)$, player 2 is supposed to know the strategy σ_1. Since he is also aware of his own strategy τ_1, he may compute both a posteriori p^{i_1} and q^{j_1}. If he then plays $\tau^+(i_1,j_1)$ a best reply in $G_n(p^{i_1}, q^{j_1})$ against $\sigma^+(i_1,j_1)$, player 1 will get less than $\underline{V}_n(p^{i_1}, q^{j_1})$ in the n last stages of $G_{n+1}(p,q)$. Since player 2 can still minimize the procedure on τ_1, we conclude that the strategy σ of player 1 guarantees a payoff less than $\underline{T}(\underline{V}_n)(p,q)$. In other words, $\underline{V}_{n+1} \leq \underline{T}(\underline{V}_n)$. A symmetrical argument leads to $\overline{V}_{n+1} \geq \overline{T}(\overline{V}_n)$.

Next, observe that $\forall f : \overline{T}(f) \geq \underline{T}(f)$. So, using the fact that G_n has a value V_n, we get:

$$\overline{V}_{n+1} \geq \overline{T}(\overline{V}_n) = \overline{T}(V_n) \geq \underline{T}(V_n) = \underline{T}(\underline{V}_n) \geq \underline{V}_{n+1}$$

Since G_{n+1} has also a value: $V_{n+1} = \overline{V}_{n+1} = \underline{V}_{n+1}$, the theorem is proved. ∎

This proof of the recursive formula is by no way constructive: it just claims that player 1 is unable to guarantee more than $\underline{T}(\underline{V}_n)(p,q)$, but it does not provide a strategy of player 1 that guarantee this amount.

To explain this in other words, the only strategy built in the last proof is a reply τ of player 2 to a given strategy of player 1. Let us call τ° this reply of player 2 to an optimal strategy σ^* of player 1. τ° is a best reply of player 2

against σ^*, but it could fail to be an optimal strategy of player 2. Indeed, it prescribes to play from the second stage on a strategy $\tau^+(i_1, j_1)$ which is an optimal strategy in $G_n(p^{*i_1}, q^{j_1})$, where p^{*i_1} is the conditional probability on K given that player 1 has used σ_1^* to select i_1. So, if player 1 deviates from σ^*, the true a posteriori p^{i_1} induced by the deviation may differ from p^{*i_1} and player 2 will still use the strategy $\tau^+(i_1, j_1)$ which could fail to be optimal in $G_n(p^{i_1}, q^{j_1})$. So when playing against τ°, player 1 could have profitable deviations from σ^*. τ° would therefore not be an optimal strategy. An example of this kind, where player 2 has no optimal strategy based on the a posteriori p^{*i_1} is presented in exercise 4, in chapter 5 of [5].

An other problem with the previous proof is that it assumes that $G_{n+1}(p, q)$ has a value. This is always the case for finite games. For games with infinite sets of actions however, it is tempting to deduce the existence of the value of $G_{n+1}(p, q)$ from the existence of a value in G_n, using the recursive structure. This is the way we proceed in [4]. This would be impossible with the argument in previous proof: we could only deduce that $\overline{V}_{n+1} \geq \overline{T}(V_n) \geq \underline{T}(V_n) \geq \underline{V}_{n+1}$, but we could not conclude to the equality $\overline{V}_{n+1} = \underline{V}_{n+1}$!

Our aim in this paper is to provide optimal strategies in $G_{n+1}(p, q)$.

We will prove in theorem 3.2 that $\underline{V}_{n+1} \geq \underline{T}(\underline{V}_n)$ by providing a strategy of player 1 that guarantees this amount. Symmetrically, we provide a strategy of player 2 that guarantees him $\overline{T}(\overline{V}_n)$, and so $\overline{T}(\overline{V}_n) \geq \overline{V}_{n+1}$.

Since in the finite case, we know by theorem 1.1 that $\underline{T}(\underline{V}_n) = V_{n+1} = \overline{T}(\overline{V}_n)$, these strategies are optimal.

These results are also useful for games with infinite action sets: provide one can argue that $\overline{T}(V_n) = \underline{T}(V_n)$, one deduces recursively the existence of the value for $G_{n+1}(p, q)$, since

$$\overline{T}(V_n) = \overline{T}(\overline{V}_n) \geq \overline{V}_{n+1} \geq \underline{V}_{n+1} \geq \underline{T}(\underline{V}_n) = \underline{T}(V_n). \tag{2}$$

Since our aim is to prepare the last section of the paper where we analyze the infinite action space games, where no general min-max theorem applies to guarantee the existence of V_n, we will deal with the finite case as if $\overline{V}_n$ and $\underline{V}_n$ were different functions. Even more, care will be taken in our proofs for the finite case to never use a "min-max" theorem that would not applies in the infinite case.

The dual games were introduced in [2] and [3] for games with incomplete information on one side to describe recursively the optimal strategies of the uninformed player. In games with incomplete information on both sides, both players are partially uninformed. We introduce the corresponding dual games in the next section.

2. The dual games

Let us first consider the amount guaranteed by a strategy σ of player 1 in $G_n(p,q)$. With obvious notations, we get:

$$\inf_\tau g_n(p,q,\sigma,\tau) = \inf_{\tau=(\tau^1,\dots,\tau^L)} \sum_l q_l \cdot g_n(p,l,\sigma,\tau^l) = \sum_l q_l \cdot y_l(p,\sigma) = \langle q, y(p,\sigma)\rangle,$$

where $\langle\cdot,\cdot\rangle$ stands for the euclidean product in $\mathbb{R}^L$, and

$$y_l(p,\sigma) := \inf_{\tau^l} g_n(p,l,\sigma,\tau^l).$$

The definition of $\underline{V}_n(p,q)$ indicates that

$$\forall p,q : \langle q, y(p,\sigma)\rangle = \inf_\tau g_n(p,q,\sigma,\tau) \leq \underline{V}_n(p,q),$$

and the equality $\langle q, y(p,\sigma)\rangle = \underline{V}_n(p,q)$ holds if and only if σ is optimal in $G_n(p,q)$. In particular, $\langle q, y(p,\sigma)\rangle$ is then a tangent hyperplane at q of the convex function $q \to \underline{V}_n(p,q)$.

In the following $\partial\underline{V}_n(p,q)$ will denote the under-gradient at q of that function:

$$\partial\underline{V}_n(p,q) := \{y | \forall q' : \underline{V}_n(p,q') \geq \underline{V}_n(p,q) + \langle q'-q, y\rangle\}$$

Our previous discussion indicates that if σ is optimal in $G_n(p,q)$, then $y(p,\sigma) \in \partial\underline{V}_n(p,q)$.

As it will appear in the next section, the relevant question to design recursively optimal strategies is as follows: given an affine functional $f(q) = \langle y, q\rangle + \alpha$ such that

$$\forall q : f(q) \leq V_n(p,q), \tag{3}$$

is there a strategy σ such that

$$\forall q : f(q) \leq \langle y(p,\sigma), q\rangle? \tag{4}$$

To answer this question it is useful to consider the Fenchel transform in q of the convex function $q \to \underline{V}_n(p,q)$: For $y \in \mathbb{R}^L$, we set:

$$\underline{V}_n^*(p,y) := \sup_q \langle q, y\rangle - \underline{V}_n(p,q)$$

As a supremum of convex functions, the function $\underline{V}_n^*$ is then convex in (p,y) on $\Delta(K) \times \mathbb{R}^L$.

For relation (3) to hold, one must then have $\alpha \leq -\underline{V}_n^*(p,y)$, so that $\forall q : f(q) \leq \langle y, q\rangle - \underline{V}_n^*(p,y)$.

The function $\underline{V}_n^*(p,y)$ is related the following dual game $G_n^*(p,y)$: At the initial stage of this game, nature choses k with the lottery p and informs player 1. Contrarily to $G_n(p,q)$, nature does not select l, but l is chosen privately

by player 2. Then the game proceeds as in $G_n(p,q)$, so that the strategies σ for player 1 are the same as in $G_n(p,q)$. For player 2 however, a strategy in $G_n^*(p,y)$ is a pair (q,τ), with $q \in \Delta(L)$ and τ a strategy in $G_n(p,q)$. The payoff $g_n^*(p,y,\sigma,(q,\tau))$ paid by player 1 (the minimizer in $G_n^*(p,y)$) to player 2 is then

$$g_n^*(p,y,\sigma,(q,\tau)) := \langle y,q\rangle - g_n(p,q,\sigma,\tau).$$

Let us next define $\underline{W}_n(p,y) = \sup_{q,\tau}\inf_\sigma g_n^*(p,y,\sigma,(q,\tau))$ and $\overline{W}_n(p,y) = \inf_\sigma \sup_{q,\tau} g_n^*(p,y,\sigma,(q,\tau))$.

We then have the following theorem:

Theorem 2.1. *$\overline{W}_n(p,y) = \underline{V}_n^*(p,y)$ and $\underline{W}_n(p,y) = \overline{V}_n^*(p,y)$.*

Proof: The following prove is designed to work with infinite action spaces: the "min-max" theorem used here is on vector payoffs instead of on strategies σ. Let $Y(p)$ be the convex set

$$Y(p) := \{y \in \mathbb{R}^L | \exists \sigma : \forall l : y_l \leq y_l(p,\sigma)\},$$

and let $\overline{Y(p)}$ be its closure in $\mathbb{R}^L$. Then

$$\underline{V}_n(p,q) = \sup_\sigma \langle y(p,\sigma),q\rangle = \sup_{y\in Y(p)} \langle y,q\rangle = \sup_{y\in\overline{Y(p)}} \langle y,q\rangle.$$

Now

$$\overline{W}_n(p,y) = \inf_\sigma \sup_q \left\{\langle y,q\rangle - \inf_\tau g_n(p,q,\sigma,\tau)\right\} = \inf_\sigma \sup_q \langle y - y(p,\sigma),q\rangle$$

Since any $z \in Y(p)$ is dominated by some $y(p,\sigma)$, we find

$$\overline{W}_n(p,y) = \inf_{z\in Y(p)} \sup_q \langle y-z,q\rangle = \inf_{z\in\overline{Y(p)}} \sup_q \langle y-z,q\rangle$$

Next, we may apply the "min-max" theorem for a bilinear functional with two closed convex strategy strategy spaces, one of which is compact, and we get thus

$$\overline{W}_n(p,y) = \sup_q \inf_{z\in\overline{Y(p)}} \langle y-z,q\rangle = \sup_q \{\langle y,q\rangle - \underline{V}_n(p,q)\} = \underline{V}_n^*(p,y)$$

On the other hand,

$$\begin{aligned}\underline{W}_n(p,y) &= \sup_{q,\tau}\inf_\sigma \{\langle y,q\rangle - g_n(p,q,\sigma,\tau)\} \\ &= \sup_q \{\langle y,q\rangle - \inf_\tau \sup_\sigma g_n(p,q,\sigma,\tau)\} \\ &= \overline{V}_n^*(p,y)\end{aligned}$$

This concludes the proof. ∎

We are now able to answer our previous question: Let σ be an optimal strategy of player 1 in $G_n^*(p,y)$. Then, $\forall q,\tau$: $\overline{W}_n(p,y) \geq \langle y,q\rangle - g_n(p,q,\sigma,\tau)$,

therefore, $\forall q$:

$$\langle y(p,\sigma),q\rangle = \inf_{\tau} g_n(p,q,\sigma,\tau) \geq \langle y,q\rangle - \underline{V}_n^*(p,y) \geq f(q). \tag{5}$$

Let us finally remark that if, for some q, $y \in \partial \underline{V}_n(p,q)$, then Fenchel lemma indicates that $V_n(p,q) = \langle y,q\rangle - \underline{V}_n^*(p,y)$, and the above inequality indicates that σ guarantees $\underline{V}_n(p,q)$ in $G_n(p,q)$:

Theorem 2.2. *Let $y \in \partial \underline{V}_n(p,q)$, and let σ be an optimal strategy of player 1 in $G_n^*(p,y)$. Then σ is optimal in $G_n(p,q)$.*

This last result indicates how to get optimal strategies in the primal game, having optimal strategies in the dual one.

3. The primal recursive formula

Let us come back on formula (1). Suppose σ_1 is already fixed. Given an array $y_{i,j}$ of vectors in $\mathbb{R}^L$, player 1 may decide to play $\sigma^+(i_1,j_1)$ an optimal strategy in $G_n^*(p^{i_1}, y_{i_1,j_1})$. As indicates relation (5), for all strategy τ^+:

$$\begin{aligned} g_n(p^{i_1}, q^{j_1}, \sigma^+(i_1,j_1), \tau^+(i_1,j_1)) &\geq \langle y(p^{i_1}, \sigma^+(i_1,j_1)), q^{j_1}\rangle \\ &\geq \langle y_{i_1,j_1}, q^{j_1}\rangle - \underline{V}_n^*(p^{i_1}, y_{i_1,j_1}) \end{aligned}$$

and so, if $\overline{y}_j := \sum_i s_i y_{i,j}$, formula (1) gives:

$$g_{n+1}(p,q,\sigma,\tau) \geq g_1(p,q,\sigma_1,\tau_1) + \sum_{j_1} t_{j_1} \langle \overline{y}_{j_1}, q^{j_1}\rangle - \sum_{j_1} t_{j_1} \sum_{i_1} s_{i_1} \underline{V}_n^*(p^{i_1}, y_{i_1,j_1})$$

We now have to indicate how player 1 will chose the array $y_{i,j}$. He will proceed in two steps: suppose $\overline{y}_j$ is fixed, he has then advantage to pick the $y_{i,j}$ among the solutions of the following minimization problem $\Psi(p,\sigma_1,\overline{y}_j)$, where

$$\Psi(p,\sigma_1,\overline{y}) := \inf_{y_i : \overline{y} := \sum_i s_i y_i} \sum_i s_i \underline{V}_n^*(p^i, y_i)$$

Lemma 3.1. *Let f_{p,σ_1} be defined as the convex function*

$$f_{p,\sigma_1}(q) := \sum_i s_i \underline{V}_n(p^i, q).$$

Then the problem $\Psi(p,\sigma_1,\overline{y})$ has optimal solutions and

$$\Psi(p,\sigma_1,\overline{y}) = f_{p,\sigma_1}^*(\overline{y}). \tag{6}$$

Proof: First of all observe that $\forall q : \underline{V}_n^*(p^i, y_i) \geq \langle y_i, q\rangle - \underline{V}_n(p^i, q)$, and thus $\Psi(p,\sigma_1,\overline{y}) \geq \langle \overline{y}, q\rangle - f_{p,\sigma_1}(q)$. This holds for all q, so $\Psi(p,\sigma_1,\overline{y}) \geq f_{p,\sigma_1}^*(\overline{y})$.

On the other hand, let q^* be a solution of the maximization problem:

$$\sup_q \langle \overline{y}, q\rangle - f_{p,\sigma_1}(q),$$

then $\overline{y} \in \partial f_{p,\sigma_1}(q^*)$. Now, the functions $q \to \underline{V}_n(p^i, q)$ are finite on $\Delta(L)$, and we conclude with Theoreme 23.8 in [6] that

$$\partial f_{p,\sigma_1}(q^*) = \sum_i s_i \partial \underline{V}_n(p^i, q^*). \tag{7}$$

In particular, there exists $y_i^* \in \partial \underline{V}_n(p^i, q^*)$ such that $\overline{y} = \sum_i s_i y_i^*$. Now observe that:

$$\begin{array}{rcl} \Psi(p, \sigma_1, \overline{y}) & \leq & \sum_i s_i \underline{V}_n^*(p^i, y_i^*) \\ & = & \sum_i s_i \left\{ \langle y_i^*, q^* \rangle - \underline{V}_n(p^i, q^*) \right\} \\ & = & \langle \overline{y}, q^* \rangle - f_{p,\sigma_1}(q^*) \\ & = & f_{p,\sigma_1}^*(\overline{y}) \end{array}$$

So both formula (6) and the optimality of y_i^* are prooven. ∎

Suppose thus that player one picks optimal $y_{i,j}$ in the problem $\Psi(p, \sigma_1, \overline{y}_j)$. He guarantees then:

$$g_{n+1}(p, q, \sigma, \tau) \geq g_1(p, q, \sigma_1, \tau_1) + \sum_{j_1} t_{j_1} \langle \overline{y}_{j_1}, q^{j_1} \rangle - \sum_{j_1} t_{j_1} f_{p,\sigma_1}^*(\overline{y}_{j_1})$$

Next let A_{p,σ_1}^j denote the L-dimensional with l-th component equal to

$$A_{p,\sigma_1}^j := \sum_{k,i} p_k \sigma_{1,k,i} A_{k,i}^{l,j}.$$

With this definition, we get $g_1(p, q, \sigma_1, \tau_1) = \sum_{j_1} t_{j_1} \langle A_{p,\sigma_1}^{j_1}, q^{j_1} \rangle$. Therefore:

$$g_{n+1}(p, q, \sigma, \tau) \geq \sum_{j_1} t_{j_1} \langle A_{p,\sigma_1}^{j_1} + \overline{y}_{j_1}, q^{j_1} \rangle - \sum_{j_1} t_{j_1} f_{p,\sigma_1}^*(\overline{y}_j)$$

Suppose next that player 1 picks $y \in \mathbb{R}^L$, and plays $\overline{y}_{j_1} := y - A_{p,\sigma_1}^{j_1}$. Since $\sum_j t_j q^j = q$, the first sum in the last relation will then be independent of the strategy τ_1 of player 2. It follows:

$$\begin{array}{rcl} g_{n+1}(p, q, \sigma, \tau) & \geq & \langle y, q \rangle - \sum_{j_1} t_{j_1} f_{p,\sigma_1}^*(y - A_{p,\sigma_1}^{j_1}) \\ & \geq & \langle y, q \rangle - \sup_{j_1} f_{p,\sigma_1}^*(y - A_{p,\sigma_1}^{j_1}) \end{array} \tag{8}$$

We will next prove that choosing appropriate σ_1 and y, player 1 can guarantee $\underline{T}(\underline{V}_n)(p, q)$:

$$\begin{array}{rcl} g_{n+1}(p, q, \sigma, \tau) & \geq & \langle y, q \rangle - \sup_{t \in \Delta(J)} \sum_{j_1} t_{j_1} f_{p,\sigma_1}^*(y - A_{p,\sigma_1}^{j_1}) \\ & = & \langle y, q \rangle - \sup\limits_{\substack{t \in \Delta(J) \\ r^1 \ldots r^J \in \Delta(L)}} \sum_{j_1} t_{j_1} \left\{ \langle y - A_{p,\sigma_1}^{j_1}, r^{j_1} \rangle - f_{p,\sigma_1}(r^{j_1}) \right\} \end{array}$$

Let $\overline{r}$ denote $\sum_{j_1} t_{j_1} r^{j_1}$. The maximization over t, r can be split in a maximization over $\overline{r} \in \Delta(L)$ and then a maximization over t, r with the constraint $\overline{r} = \sum_{j_1} t_{j_1} r^{j_1}$. This last maximization is clearly equivalent to a maximization

over a strategy τ_1 of player 2 in $G_1(p,\overline{r})$, inducing a probability λ on $(J \times L)$, whose marginal on J is t and the conditional on L are the r^{j_1}. In this way, $\sum_{j_1} t_{j_1} \langle A^{j_1}_{p,\sigma_1}, r^{j_1} \rangle = g_1(p,\overline{r},\sigma_1,\tau_1)$, and we get:

$$g_{n+1}(p,q,\sigma,\tau) \geq \inf_{\overline{r}} \{\langle y, q-\overline{r}\rangle + H(p,\sigma_1,\overline{r})\}$$

where $H(p,\sigma_1,\overline{r}) := \inf_{\tau_1} \left(g_1(p,\overline{r},\sigma_1,\tau_1) + \sum_{j_1} t_{j_1} f_{p,\sigma_1}(r^{j_1}) \right)$. We will prove in lemma 3.4 that $H(p,\sigma_1,\overline{r})$ is a convex function of $\overline{r}$. If player 1 chooses $y \in \partial H(p,\sigma_1,q)$ then $\forall \overline{r}$: $\langle y, q-\overline{r}\rangle + H(p,\sigma_1,\overline{r}) \geq H(p,\sigma_1,q)$, and thus

$$g_{n+1}(p,q,\sigma,\tau) \geq H(p,\sigma_1,q)$$

Replacing now f_{p,σ_1} by its value, we get:

$$H(p,\sigma_1,q) = \inf_{\tau_1} \left(g_1(p,q,\sigma_1,\tau_1) + \sum_{i_1,j_1} s_{i_1} t_{j_1} \underline{V}_n(p^{i_1}, q^{j_1}) \right) \tag{9}$$

Since player 1 can still maximize over σ_1, we just have proved that player 1 can guarantee

$$\sup_{\sigma_1} H(p,\sigma_1,q) \tag{10}$$

proceeding as follows:

(1) He first selects an optimal σ_1 in (10), that is, an optimal strategy in the problem $\underline{T}(\underline{V}_n)(p,q)$.
(2) He then computes the function $r \to H(p,\sigma_1,r)$ and picks $y \in \partial H(p,\sigma_1,q)$.
(3) He next defines $\overline{y}_j$ as $\overline{y}_j = y - A^j_{p,\sigma_1}$ and finds optimal $y_{i,j}$ in the problem $\Psi(p,\sigma_1,\overline{y}_j)$ as in the proof of lemma 3.1.
(4) Finally, he selects $\sigma^+(i,j)$ an optimal strategy in $G^*_n(p,y_{i,j})$.

The next theorem is thus proved.

Theorem 3.2. *With the above described strategy, player 1 guarantees* $\underline{T}(\underline{V}_n)(p,q)$ *in* $G_{n+1}(p,q)$. *Therefore:* $\underline{V}_{n+1}(p,q) \geq \underline{T}(\underline{V}_n)(p,q)$

The first part of the proof of theorem 1.1 indicates that $\underline{V}_{n+1}(p,q) \leq \underline{T}(\underline{V}_n)(p,q)$, and this result will hold even for games with infinite action spaces: it uses no min-max argument. We may then conclude:

Corollary 3.3. $\underline{V}_{n+1}(p,q) = \underline{T}(\underline{V}_n)(p,q)$ *and the above described strategy is thus optimal in* $G_{n+1}(p,q)$.

It just remains for us to prove the following lemma:

Lemma 3.4. *The function* $H(p,\sigma_1,\overline{r})$ *is convex in* $\overline{r}$.

Proof: Let us denote $\Delta_{\overline{r}}$ the set of probabilities λ on $(J \times L)$, whose marginal $\lambda_{|L}$ on L is $\overline{r}$. As mentioned above, a strategy τ_1, joint to $\overline{r}$, induces a probability λ in $\Delta_{\overline{r}}$, and conversely, any such λ is induced by some τ_1.

Let next e_l be the l-th element of the canonical basis of $\mathbb{R}^L$. The mapping $e : l \to e_l$ is then a random vector on $(J \times L)$, and $r^{j_1} = E_\lambda[e|j_1]$. Similarly, the mapping $A_{p,\sigma_1} : (l, j_1) \to A^{l,j_1}_{p,\sigma_1}$ is a random variable and $E_\lambda[A_{p,\sigma_1}] = g_1(p, \overline{r}, \sigma_1, \tau_1)$. We get therefore

$$H(p, \sigma_1, \overline{r}) := \inf_{\lambda \in \Delta_{\overline{r}}} E_\lambda[A_{p,\sigma_1} + f_{p,\sigma_1}(E_\lambda[e|j_1])].$$

Let now $\pi_0, \pi_1 \geq 0$, with $\pi_0 + \pi_1 = 1$, let $\overline{r}_0, \overline{r}_1, \overline{r}_\pi \in \Delta(L)$, with $\overline{r}_\pi = \pi_1 \overline{r}_1 + \pi_0 \overline{r}_0$. Let $\lambda_u \in \Delta_{\overline{r}_u}$, for u in $\{0, 1\}$. Then $\pi, \lambda_1, \lambda_0$ induce a probability μ on $(\{0, 1\} \times J \times L)$: first pick u at random in $\{0, 1\}$, with probability π_1 of u being 1. Then, conditionally to u, use the lottery λ_u to select (j_1, l). The marginal λ_π of μ on $(J \times L)$ is obviously in $\Delta_{\overline{r}_\pi}$. Next observe that, due to Jensen's inequality and the convexity of f_{p,σ_1}:

$$\begin{array}{rcl} \sum_u \pi_u E_{\lambda_u}[A_{p,\sigma_1} + f_{p,\sigma_1}(E_{\lambda_u}[e|j_1])] & = & E_\mu[A_{p,\sigma_1} + f_{p,\sigma_1}(E_{\lambda_u}[e|j_1])] \\ & = & E_\mu[A_{p,\sigma_1} + f_{p,\sigma_1}(E_\mu[e|j_1, u])] \\ & \geq & E_\mu[A_{p,\sigma_1} + f_{p,\sigma_1}(E_\mu[e|j_1])] \\ & = & E_{\lambda_\pi}[A_{p,\sigma_1} + f_{p,\sigma_1}(E_{\lambda_\pi}[e|j_1])] \\ & \geq & H(p, \sigma_1, \overline{r}_\pi) \end{array}$$

Minimizing the left hand side in λ_0 and λ_1, we obtain:

$$\sum_u \pi_u H(p, \sigma_1, \overline{r}_u) \geq H(p, \sigma_1, \overline{r}_\pi)$$

and the convexity is thus prooved. ∎

4. The dual recursive structure

The construction of the optimal strategy in $G_{n+1}(p, q)$ of last section is not completely satisfactory: the procedure ends up in point 4) by selecting optimal strategies in the dual game $G^*_n(p, y_{i,j})$ but it does not explain how to construct such strategies. The purpose of this section is to construct recursively optimal strategies in the dual game. It turns out that this construction will be "self-contained" and truly recursive: finding optimal strategies in G^*_{n+1} will end up in finding optimal strategies in G^*_n.

Given σ_1, let us consider the following strategy $\sigma = (\sigma_1, \sigma^+)$ in $G^*_{n+1}(p, y)$: player 1 sets $\overline{y}_j = y - A^j_{p,\sigma_1}$ and finds optimal $y_{i,j}$ in the problem $\Psi(p, \sigma_1, \overline{y}_j)$ as in the proof of lemma 3.1. He then plays $\sigma^+(i_1, j_1)$ an optimal strategy in $G^*_n(p, y_{i_1,j_1})$. This is exactly what we prescribed for player 1 in the beginning

of last section. In particular, this strategy was not depending on q in the last section, so that inequality (8) holds for all q, τ:

$$\sup_{j_1} f^*_{p,\sigma_1}(y - A^{j_1}_{p,\sigma_1}) \geq \langle y, q\rangle - g_{n+1}(p, q, \sigma, \tau) = g^*_{n+1}(p, y, \sigma, (q, \tau))$$

So, with lemma 3.1, and the definition of Ψ.

$$\begin{aligned} g^*_{n+1}(p, y, \sigma, (q, \tau)) &\leq \sup_{j_1} f^*_{p,\sigma_1}(y - A^{j_1}_{p,\sigma_1}) \\ &= \sup_{j_1} \Psi(p, \sigma_1, y - A^{j_1}_{p,\sigma_1}) \\ &= \sup_{j_1} \inf_{y_i : \sum_i s_i y_i = y - A^{j_1}_{p,\sigma_1}} \textstyle\sum_i s_i \underline{V}^*_n(p^i, y_i) \\ &= \inf_{y_{i,j} : \sum_i s_i y_{i,j} = y - A^{j}_{p,\sigma_1}} \sup_{j_1} \textstyle\sum_i s_i \underline{V}^*_n(p^i, y_{i,j_1}) \end{aligned} \tag{11}$$

Notice that there is no "min-max" theorem needed to derive the last equation: We just allowed the variables y_i to depend on j_1: the new variables are $y_{i,j}$.

With theorem 2.1, $\underline{V}^*_n(p^i, y_{i,j_1}) = \overline{W}_n(p^i, y_{i,j_1})$. It is next convenient to define $m_{i,j} := y_{i,j} - y + A^j_{p,\sigma_1}$, so that $\sum_i s_i m_{i,j} = 0$, and to take $m_{i,j}$ as minimization variables:

$$g^*_{n+1}(p, y, \sigma, (q, \tau)) \leq \inf_{m_{i,j} : \sum_i s_i m_{i,j} = 0} \sup_{j_1} \textstyle\sum_{i_1} s_{i_1} \overline{W}_n(p^{i_1}, y - A^{j_1}_{p,\sigma_1} + m_{i_1,j_1}) \tag{12}$$

Let still player 1 minimize this procedure over σ_1. It follows:

Theorem 4.1. *The above defined strategy σ guarantees $\overline{T}^*(\overline{W}_n)(p, y)$ to player 1 in $G^*_{n+1}(p, y)$, where, for a convex function W on $(\Delta(K) \times \mathbb{R}^L)$:*

$$\overline{T}^*(W)(p, y) := \inf_{\substack{\sigma_1 \\ m_{i,j} : \sum_i s_i m_{i,j} = 0}} \sup_{j_1} \sum_{i_1} s_{i_1} W(p^{i_1}, y - A^{j_1}_{p,\sigma_1} + m_{i_1,j_1}).$$

In particular: $\overline{W}_{n+1}(p, y) \leq \overline{T}^*(\overline{W}_n)(p, y)$

We next will prove the following corollary:

Corollary 4.2. *$\overline{W}_{n+1}(p, y) = \overline{T}^*(\overline{W}_n)(p, y)$ and the strategy σ is thus optimal in $G^*_{n+1}(p, y)$.*

Proof: If player 1 uses as strategy $\sigma = (\sigma_1, \sigma^+)$ in $G^*_{n+1}(p, y)$, player 2 may reply the following strategy (q, τ), with $\tau = (\tau_1, \tau^+)$: for a given choice of q, τ_1, he computes the a posteriori p^{i_1}, q^{j_1} and plays a best reply $\tau^+(i_1, j_1)$ against $\sigma^+(i_1, j_1)$ in $G_n(p^{i_1}, q^{j_1})$. Since

$$g_n(p^{i_1}, q^{j_1}, \sigma^+(i_1, j_1), \tau^+(i_1, j_1)) \leq \underline{V}_n(p^{i_1}, q^{j_1}),$$

we get

$$\begin{aligned} g^*_n(p, y, \sigma, (q, \tau)) \geq & \langle y, q\rangle - g_1(p, q, \sigma_1, \tau_1) - \textstyle\sum_{i_1,j_1} s_{i_1} t_{j_1} \underline{V}_n(p^{i_1}, q^{j_1}) \\ = & \textstyle\sum_{j_1} t_{j_1} \left(\langle y - A^{j_1}_{p,\sigma_1}, q^{j_1}\rangle - \sum_{i_1} s_{i_1} \underline{V}_n(p^{i_1}, q^{j_1})\right) \end{aligned}$$

The reply (q,τ) of player 2 we will consider is that corresponding to the choice of q,τ_1 maximizing this last quantity. This turns out to be a maximization over the joint law λ on $(J\times L)$. In turn, it is equivalent to a maximization (t,q^{j_1}), without any constraint on $\sum_j t_j q^j$. So:

$$\begin{aligned} g_n^*(p,y,\sigma,(q,\tau)) \geq & \sup_t \sum_{j_1} t_{j_1} \sup_{q^{j_1}} \left(\langle y - A^{j_1}_{p,\sigma_1}, q^{j_1}\rangle - \sum_{i_1} s_{i_1} \underline{V}_n(p^{i_1},q^{j_1})\right) \\ = & \sup_{j_1} f^*_{p,\sigma_1}(y - A^{j_1}_{p,\sigma_1}). \end{aligned}$$

We then derive as in equations (11) and (12) that

$$\begin{aligned} \sup_{j_1} f^*_{p,\sigma_1}(y - A^{j_1}_{p,\sigma_1}) = & \inf_{m_{i,j}:\sum_i s_i m_{i,j}=0} \sup_{j_1} \sum_{i_1} s_{i_1} \overline{W}_n(p^{i_1}, y - A^{j_1}_{p,\sigma_1} + m_{i_1,j_1}) \\ \geq & \overline{T}^*(\overline{W}_n)(p,y) \end{aligned}$$

So, with no strategy, player 1 will be able to guarantee a better payoff in $G^*_{n+1}(y,p)$ than $\overline{T}^*(\overline{W}_n)(p,y)$, and the corollary is proved. ∎

We thus gave a recursive procedure to construct optimal strategies in the dual game. Now, instead of using the construction of the previous section to play optimally in $G_{n+1}(p,q)$, player 1 can use theorem 2.2: He picks $y \in \partial \underline{V}_{n+1}(p,q)$, and then plays optimally in $G^*_{n+1}(p,y)$, with the recursive procedure introduced in this section.

5. Games with infinite action spaces

In this section, we generalize the previous results to games where I and J are infinite sets. K and L are still finite sets. The sets I and J are then equipped with σ-algebras $\mathcal{I}$ and $\mathcal{J}$ respectively. We will assume that $\forall k,l$, the mapping $(i,j)\to A^{l,j}_{k,i}$ is bounded and measurable on $(\mathcal{I}\otimes\mathcal{J})$. The natural σ-algebra on the set of histories H_m is then $\mathcal{H}_m := (\mathcal{I}\otimes\mathcal{J})^{\otimes m}$. A behavior strategy σ for player 1 in $G_n(p,q)$ is then a n-uple $(\sigma_1,\dots,\sigma_n)$ of transition probabilities σ_m from $K\times H_{m-1}$ to I which means: $\sigma_m : (k,h_{m-1},A)\in(K\times H_{m-1}\times\mathcal{I})\to\sigma_m(k,h_{m-1})[A]\in[0,1]$ satifying $\forall k,h_{m-1}$: $\sigma_m(k,h_{m-1})[\cdot]$ is a probability measure on $(I,\mathcal{I})$, and $\forall k,A$, $\sigma_m(k,h_{m-1})[A]$ is $\mathcal{H}_m$ measurable. A strategy of player 2 is defined in a similar way. To each (p,q,σ,τ) corresponds a unique probability measure $\Pi^n_{(p,q,\sigma,\tau)}$ on $(K\times L\times H_n, \mathcal{P}(K)\otimes\mathcal{P}(L)\otimes\mathcal{H}_n)$. Since the payoff map $A^{l,j}_{k,i}$ is bounded and measurable, we are allowed to define $g_n(p,q,\sigma,\tau) := E_{\Pi^n_{(p,q,\sigma,\tau)}}[\sum_{m=1}^n A^{l,j_m}_{k,i_m}]$. The definitions of $\overline{V}_n$, $\underline{V}_n$, $\overline{W}_n$ and $\underline{W}_n$ are thus exactly the same as in the finite case, and the a posteriori p^{i_1} and q^{j_1} are defined as the conditional probabilities of $\Pi^1_{(p,q,\sigma_1,\tau_1)}$ on K and L given i_1 and j_1. The sums in the definition of the recursive operators $\underline{T}$ and $\overline{T}$ are to be replaced by expectations:

$$\underline{T}(f)(p,q) = \sup_{\sigma_1}\inf_{\tau_1}\left\{g_1(p,q,\sigma_1,\tau_1) + E_{\Pi^1_{(p,q,\sigma_1,\tau_1)}}[f(p^{i_1},q^{j_1})]\right\}$$

Let $\mathcal{V}$ denote the set of Lipschitz functions $f(p,q)$ on $\Delta(K)\times\Delta(L)$ that are concave in p and convex in q. The result we aim to prove in this section is the next theorem. For all $V\in\mathcal{V}$ such that $\underline{V}_n > V$, we will provide strategies of player 1 that guarantee him $\underline{T}(V)$.

Theorem 5.1. *If $\underline{V}_n \geq V$, where $V\in\mathcal{V}$, then $\underline{V}_{n+1}\geq \underline{T}(V)$.*

Proof: Since $\forall\epsilon>0$, $\underline{T}(V-\epsilon)=\underline{T}(V)-\epsilon$, it is sufficient to prove the result for $V<\underline{V}_n$. In this case, we also have $\forall p,y$: $V^*(p,y)>\underline{V}^*_n(p,y)=\overline{W}_n(p,y)$.

In the infinite games, optimal strategies may fail to exist. However, due to the strict inequality, $\forall p,y$, there must exist a strategy $\sigma^+_{p,y}$ in $G^*_n(p,y)$ that warrantees strictly less than $V^*(p,y)$ to player 1. Since the payoffs map $A^{l,j}_{k,i}$ is bounded and V^* is continuous, the set $O(p,y)$ of $(p',y')\in\Delta(K)\times\mathbb{R}^L$ such that $\sigma^+_{p,y}$ warrantees $V^*(p',y')$ in $G^*_n(p',y')$ is a neighborhood of (p,y). There exists therefore a sequence $\{(p_m,y_m)\}_{m\in\mathbb{N}}$ such that $\cup_m O(p_m,y_m)=\Delta(K)\times\mathbb{R}^L$. The map $(p,y)\to\sigma^+(p,y)$ defined as $\sigma^+(p,y):=\sigma^+_{p_{m^*},y_{m^*}}$, where m^* is the smallest integer m with $(p,y)\in O(p_m,y_m)$ satisfies then

- for all ℓ, the map $(p,y)\to\sigma^+_\ell(p,y)(k,h_{\ell-1})$ is a transition probability from $(\Delta(K)\times\mathbb{R}^L\times K\times H_{\ell-1})$ to I.
- $\forall p,y$: $\sigma^+(p,y)$ warrantees $V^*(p,y)$ to player 1 in $G^*_n(p,y)$.

The argument of section 3 can now be adapted to this setting: Given a first stage strategy σ_1 and a measurable mapping $y:(i_1,j_1)\to y_{i_1,j_1}\in\mathbb{R}^L$, player 1 may decide to play $\sigma^+(p^{i_1},y_{i_1,j_1})$ from stage 2 on in $G_{n+1}(p,q)$. Since $\sigma^+(p,y)$ warrantees $V^*(p,y)$ to player 1 in $G^*_n(p,y)$, we get

$$g_n(p^{i_1},q^{j_1},\sigma^+(i_1,j_1),\tau^+(i_1,j_1))\geq\langle y_{i_1,j_1},q^{j_1}\rangle-V^*(p^{i_1},y_{i_1,j_1}).$$

Let s and t denote the marginal distribution of i_1 and j_1 under $\Pi^1_{(p,q,\sigma_1,\tau_1)}$. In the following $E_s[\cdot]$ and $E_t[\cdot]$ are short hand writings for $\int_I\cdot ds(i_1)$ and $\int_J\cdot dt(j_1)$. If $\overline{y}_j:=E_s[y_{i,j}]$, formula (1) gives:

$$g_{n+1}(p,q,\sigma,\tau)\geq g_1(p,q,\sigma_1,\tau_1)+E_t\left[\langle\overline{y}_{j_1},q^{j_1}\rangle-E_s[V^*(p^{i_1},y_{i_1,j_1})]\right].$$

As in section 3, player 1 would have advantage to choose $i_1\to y_{i_1,j_1}$ optimal in the problem $\Psi(p,\sigma_1,\overline{y}_{j_1})$, where

$$\Psi(p,\sigma_1,\overline{y}):=\inf_{y:\overline{y}:=E_s[y_{i_1}]}E_s[V^*(p^{i_1},y_{i_1})]$$

Lemma 3.1 also holds in this setting, with $f_{p,\sigma_1}(q):=E_s[V(p^{i_1},q)]$. The only difficulty to adapt the prove of section 3 is to generalize equation (7). With the Lipschitz property of V, we prove in theorem 5.3 that there exists a measurable mapping $y:i\to\mathbb{R}^L$ satisfying $E_s[y_{i_1}]=\overline{y}$ and for s-a.e i_1: $y_{i_1}\in\partial V(p^{i_1},q^*)$. We get in this way $\Psi(p,\sigma_1,\overline{y})=f^*_{p,\sigma_1}(\overline{y})$.

We next prove that for all measurable map $\overline{y} : j_1 \to \overline{y}_{j_1}$, $\forall \epsilon > 0$, there exists a measurable array $y : (i_1, j_1) \to y_{i_1,j_1}$ such that $\forall j_1 : E_s[y_{i_1,j_1}] = \overline{y}_{j_1}$ and

$$\forall j_1 : E_s[V^*(p^{i_1}, y_{i_1,j_1})] \leq f^*_{p,\sigma_1}(\overline{y}_{j_1}) + \epsilon \tag{13}$$

The function f^*_{p,σ_1} is Lipschitz, and we may therefore consider a triangulation of $\mathbb{R}^L$ in a countable number of L-dimensional simplices with small enough diameter to insure that the linear interpolation $\overline{f^*_{p,\sigma_1}}$ of f^*_{p,σ_1} at the extreme points of a simplex S satisfies $\overline{f^*_{p,\sigma_1}} \leq f^*_{p,\sigma_1} + \epsilon$ on the interior of S. We define then $y(\overline{y}, i)$ on $S \times I$ as the linear interpolation on S of optimal solutions of $\Psi(p, \sigma_1, \overline{y})$ at the extreme points of the simplex S. Obviously $E_s[y(\overline{y}, i_1)] = \overline{y}$, and, due to the convexity of V^*, we get $E_s[V^*(p^{i_1}, y(\overline{y}, i_1))] \leq \overline{f^*_{p,\sigma_1}}(\overline{y})$. The array $y_{i_1,j_i} := y(\overline{y}_{j_1}, i_1)$ will then satisfy (13).

With such arrays y, Player 1 guarantees up to an arbitrarily small ϵ:

$$\inf_{\tau_1} g_1(p, q, \sigma_1, \tau_1) + E_t\left[\langle \overline{y}_{j_1}, q^{j_1}\rangle - f^*_{p,\sigma_1}(\overline{y}_{j_1})\right]$$

The proof next follows exactly as in section 3, replacing summations by expectations. ∎

As announced in the introduction, the last theorem has a corollary:

Corollary 5.2. *If $\forall V \in \mathcal{V} : \underline{T}(V) = \overline{T}(V) \in \mathcal{V}$, then, $\forall n, p, q$, the game $G_n(p, q)$ has a value $V_n(p, q)$, and $V_{n+1} = \underline{T}(V_n) \in \mathcal{V}$.*

Proof: The proof just consists of equation (2). ∎

It remains for us to prove the next theorem:

Theorem 5.3. *Let $(\Omega, \mathcal{A}, \mu)$ be probability space, let U be a convex subset of $\mathbb{R}^L$, let f be a function $\Omega \times U \to \mathbb{R}$ satisfying*

- *$\forall \omega$: the mapping $q \to f(\omega, q)$ is convex.*
- *$\exists M : \forall q, q', \omega : |f(\omega, q) - f(\omega, q')| \leq M|q - q'|$.*
- *$\forall q$: the mapping $\omega \to f(\omega, q)$ is in $\mathbb{L}^1(\Omega, \mathcal{A}, \mu)$.*

The function $f_\mu(q) := E_\mu[f(\omega, q)]$ is then clearly convex and M-Lipschitz in q. Let next $\overline{y} \in \partial f_\mu(q_0)$.

Then there exists a measurable map $y : \Omega \to \mathbb{R}^L$ such that
1) for μ-a.e. $\omega : y(\omega) \in \partial f(\omega, q_0)$.
2) $\overline{y} = E_\mu[y(\omega)]$

Proof: Using a translation, there is no loss of generality to assume $q_0 = 0 \in U$. Then, considering the mapping $g(\omega, q) := f(\omega, q) - f(\omega, 0) - \langle \overline{y}, q\rangle$, and the corresponding $g_\mu(q) := E_\mu[g(\omega, q)]$, we get $\forall \omega : g(\omega, 0) = 0 = g_\mu(0)$ and $\forall q : g_\mu(q) \geq 0$.

Let $\mathcal{S}$ denote the set of (α, X) where α and X are respectively $\mathbb{R}$- and $\mathbb{R}^L$-valued mappings in $\mathbb{L}^1(\Omega, \mathcal{A}, \mu)$. Let us then define

$$\mathcal{R} := \{(\alpha, X) \in \mathcal{S} | E_\mu[\alpha(\omega)] > E_\mu[g(\omega, X(\omega))]\}$$

Our hypotheses on f imply in particular that the map $\omega \to g(\omega, X(\omega))$ is $\mathcal{A}$-measurable and in $\mathbb{L}^1(\Omega, \mathcal{A}, \mu)$. Furthermore the map $X \to E_\mu[g(\omega, X(\omega))]$ is continuous for the $\mathbb{L}^1$-norm, so that $\mathcal{R}$ is an open convex subset of $\mathcal{S}$.

Let us next define the linear space $\mathcal{T}$ as:

$$\mathcal{T} := \{(\alpha, X) \in \mathcal{S} | E_\mu[\alpha(\omega)] = 0, \text{ and } \exists \overline{x} \in \mathbb{R}^L \text{ such that } \mu\text{-a.s. } X(\omega) = \overline{x}\}.$$

Now observe that $\mathcal{R} \cap \mathcal{T} = \emptyset$. Would indeed (α, X) belong to $\mathcal{R} \cap \mathcal{T}$, we would have μ-a.s. $X(\omega) = \overline{x}$, and $0 = E_\mu[\alpha(\omega)] > E_\mu[g(\omega, X(\omega))] = g_\mu(\overline{x}) \geq 0$.

There must therefore exist a linear functional ϕ on $\mathcal{S}$ such that

$$\phi(\mathcal{R}) > 0 = \phi(\mathcal{T}).$$

Since the dual of $\mathbb{L}^1$ is $\mathbb{L}^\infty$, there must exist a $\mathbb{R}$-valued λ and a $\mathbb{R}^L$-valued Z in $\mathbb{L}^\infty(\Omega, \mathcal{A}, \mu)$ such that

$$\forall (\alpha, X) \in \mathcal{S} : \phi(\alpha, X) = E_\mu[\lambda(\omega)\alpha(\omega) - \langle Z(\omega), X(\omega)\rangle].$$

From $0 = \phi(\mathcal{T})$, it is easy to derive that $E_\mu[Z(\omega)] = 0$ and that $\exists \overline{\lambda} \in \mathbb{R}$ such that μ-a.s. $\lambda(\omega) = \overline{\lambda}$.

Next, $\forall \epsilon > 0$, $\forall X \in \mathbb{L}^1(\Omega, \mathcal{A}, \mu)$, the pair (α, X) belongs to $\mathcal{R}$, where $\alpha(\omega) := g(\omega, X(\omega)) + \epsilon$. So, $\phi(\mathcal{R}) > 0$ with $X \equiv 0$, implies in particular $\overline{\lambda} > 0$, and ϕ may be normalized so as to take $\overline{\lambda} = 1$. Finally, we get $\forall \epsilon > 0$, $\forall X \in \mathbb{L}^1(\Omega, \mathcal{A}, \mu)$: $E_\mu[g(\omega, X(\omega))] + \epsilon > E_\mu[\langle Z(\omega), X(\omega)\rangle]$ and thus, $\forall X \in \mathbb{L}^1(\Omega, \mathcal{A}, \mu)$: $E_\mu[g(\omega, X(\omega))] \geq E_\mu[\langle Z(\omega), X(\omega)\rangle]$.

For $A \in \mathcal{A}$ and $x \in \mathbb{R}^L$, we may apply the last inequality to $X(\omega) := \mathbb{1}_A(\omega)x$, and we get: $E_\mu[\mathbb{1}_A g(\omega, x)] \geq E_\mu[\mathbb{1}_A \langle Z(\omega), x\rangle]$. Therefore, for all $x \in \mathbb{R}^L$: $\mu(\Omega_x) = 1$, where $\Omega_x = \{\omega \in \Omega : g(\omega, x) \geq \langle Z(\omega), x\rangle\}$. So, if $\Omega' := \cap_{x \in \mathbb{Q}^L} \Omega_x$, we get $\mu(\Omega') = 1$, since $\mathbb{Q}^L$ is a countable set, and $\forall \omega \in \Omega'$, $\forall x \in \mathbb{Q}^L$: $g(\omega, x) \geq \langle Z(\omega), x\rangle$. Due to the continuity of $g(\omega, .)$, the last inequality holds in fact for all $\forall x \in \mathbb{R}^L$, so that $\forall \omega \in \Omega' : Z(\omega) \in \partial g(\omega, 0)$.

Hence, if we define $y(\omega) := \overline{y} + Z(\omega)$, we get μ-a.s.: $y(\omega) \in \partial f(\omega, 0)$ and $E_\mu[y(\omega)] = \overline{y} + E_\mu[Z(\omega)] = \overline{y}$. This concludes the proof of the theorem. ∎

References

[1] Aumann, Robert J., and Michael B. Maschler: "Repeated games of incomplete information: the zerosum extensive case", in Mathematica, Inc. (1968), chap. III, pp. 37–116.

[2] De Meyer, Bernard:"Repeated Games and Partial Differential Equations", in Mathematics of Operations Research, Vol. 21, No1, 1996, pp. 209–236.

[3] De Meyer, Bernard: "Repeated games, Duality, and the Central Limit Theorem", in Mathematics of Operation Research, Vol 21, No 1, 1996, pp. 237-251.

[4] De Meyer, Bernard and Alexandre Marino: "Repeated market games with lack of information on both sides", Cahier de la MSE 2004/66, Université Paris 1 (Panthéon Sorbonne), France, 2004.

[5] Mertens, Jean-François; Sylvain Sorin and Shmuel Zamir: "Repeated Games" Core Discussion papers 9420, 9421, 9422, Core, Université Catholique de Louvain, Belgium, 1994.

[6] Rockafellar, R.Tyrrell: *Convex Analysis*, Princeton, New Jersey, Princeton university press, 1970.

[7] Stearns, Richard E.: "A formal information concept for games with incomplete information", in Mathematica, Inc. (1967), chap. IV, pp. 405–433.