



HAL
open science

Nash equilibria with partial monitoring; Computation and Lemke-Howson algorithm

Vianney Perchet

► **To cite this version:**

Vianney Perchet. Nash equilibria with partial monitoring; Computation and Lemke-Howson algorithm. 2013. hal-00773217

HAL Id: hal-00773217

<https://hal.science/hal-00773217>

Preprint submitted on 12 Jan 2013

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.

Nash equilibria with partial monitoring; Computation and Lemke-Howson algorithm.

Vianney Perchet *

January 12, 2013

Abstract

In two player bi-matrix games with partial monitoring, actions played are not observed, only some messages are received. Those games satisfy a crucial property of usual bi-matrix games: there are only a finite number of required (mixed) best replies. This is very helpful while investigating sets of Nash equilibria: for instance, in some cases, it allows to relate it to the set of equilibria of some auxiliary game with full monitoring.

In the general case, the Lemke-Howson algorithm is extended and, under some genericity assumption, its output are Nash equilibria of the original game. As a by product, we obtain an oddness property on their number.

Introduction

In finite games, proving the existence of Nash equilibria [10, 11] is not very challenging, as they are fixed points of some correspondence. On the other hand, computing the whole set of Nash equilibria (or exhibiting some of its topological properties) is quite hard [12]. Similar statements can be made in games where actions chosen or actual payoff mappings are (partially) unknown. These games are getting increasing interest and have been referred as *robust* [1], *ambiguous* [2], with *uncertainty* [6], *partially specified* [7], and so on. Indeed, Nash equilibria are defined similarly as fixed points of some complicated – yet regular – correspondence; existence is then ensured, almost always using the very same argument of Nash [10], Kakutani’s fixed point theorem. So the focus shall not be existence, but characterizations and computation of these equilibria.

In full generality, and as expected as it is a more complex set-up, this turns out to be a very challenging problem [1, Section 5]. We therefore consider here the class of bi-matrix games with *partial monitoring*, see e.g., [9], which contains all two-player finite games. In this framework, players might not observe perfectly their opponent’s actions (yet we always assume that one knows his own choice); they only receive messages. Depending on the game, actions and messages can in fact be correlated as well as independent; we

*Laboratoire de Probabilités et de Modèles Aléatoires, Université Paris 7, 175 rue du Chevaleret, 75013 Paris. vianney.perchet@normalesup.org

could even assume that the latter is random, but up to some lifting, this can be reduced to the deterministic case, see [14]. These games are therefore described by two pair of matrices: a first pair for payoffs and a second pair for messages received.

Players, facing uncertainties upon their payoffs, cannot directly maximize them. As it is usual now [5, 3], we assume that they optimize their behavior with respect to the worst possible scenario, leading to *maxmin expected utility*.

Using topological properties of linear mappings and projection, we recover surprisingly the following fundamental property of finite bi-matrix games with full monitoring (when actions are observed). There exists a fixed finite subset of (mixed) actions containing best-replies to any action of the opponent. While obvious with full monitoring by considering whole set of pure actions, this result is not immediate with partial monitoring (and actually incorrect in another class of games than the one considered here).

In the subclass of games called with *semi-standard* information structure, developed in Section 2, this allows the construction of an auxiliary game with full monitoring such that its Nash equilibria are (in some sense) also equilibria of the original game. So any property with full monitoring holds for this type of games.

In the general case, this direct reduction is incorrect. Yet we prove in Section 4 that Nash equilibria satisfy again another usual properties of full monitoring, see [18]. Using this, sets of Nash equilibria are characterized and some of them can be computed using the Lemke-Howson algorithm [8], recalled briefly in Section 3. These computations are illustrated in Section 5; other claims are also, as often as possible, accompanied by examples. Interestingly, since Nash equilibria – even with partial monitoring – are endpoints of a special instance of the Lemke-Howson algorithm, some oddness property of their set is preserved (as soon as some genericity assumption is satisfied).

1 Two players game with partial monitoring

Consider a finite two players game Γ where action of player 1 (resp. player 2) is by \mathcal{A} (resp. \mathcal{B}) and his payoff mapping is $u : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ (resp. $v : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$), extended multi-linearly to $\mathcal{X} \times \mathcal{Y}$. We denote by $\mathcal{X} = \Delta(\mathcal{A})$ and $\mathcal{Y} = \Delta(\mathcal{B})$ mixed action sets of both players. We also assume that they have partial monitoring: they do not observe actions of their opponent but receive messages instead, see [9]. Formally, there exist two convex compact sets of messages \mathcal{H} and \mathcal{M} and two signaling mappings H and M from $\mathcal{A} \times \mathcal{B}$ into \mathcal{H} or \mathcal{M} (also extended multi-linearly) such that if players choose $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, player 1 gets a payoff of $u(x, y)$ but he only observes the message $H(x, y) \in \mathcal{H}$. On his side, player 2 gets a payoff of $v(x, y)$ and he observes $M(x, y) \in \mathcal{M}$.

No matter his choice of actions, player 1 cannot distinguish between y and $y' \in \mathcal{Y}$ satisfying $H(a, y) = H(a, y')$ for every $a \in \mathcal{A}$. We thus define the *maximal informative mapping* $\mathbf{H} : \mathcal{Y} \rightarrow \mathcal{H}^{\mathcal{A}}$ (A stands for the cardinality of \mathcal{A}) by:

$$\forall y \in \mathcal{Y}, \mathbf{H}(y) = \left[H(a, y) \right]_{a \in \mathcal{A}} \in \mathcal{H}^{\mathcal{A}}.$$

Similarly, the maximal informative payoff of player 2, $\mathbf{M} : \mathcal{X} \rightarrow \mathcal{M}^B$, is defined by

$$\forall x \in \mathcal{X}, \mathbf{M}(x) = \left[M(x, b) \right]_{b \in \mathcal{B}} \in \mathcal{M}^B.$$

These linear mappings induce *uncertainty correspondences* $\Phi : \mathcal{Y} \rightrightarrows \mathbb{R}^A$ and $\Psi : \mathcal{X} \rightrightarrows \mathbb{R}^B$ defined by:

$$\Phi(y) = \{u(\cdot, y') \in \mathbb{R}^A; \mathbf{H}(y') = \mathbf{H}(y)\} \text{ and } \Psi(x) = \{v(x', \cdot) \in \mathbb{R}^B; \mathbf{M}(x') = \mathbf{M}(x)\}.$$

Informally, if player 2 chooses $y \in \mathcal{Y}$, then player 1 cannot distinguish it from any other y' that have the same image under \mathbf{H} ; thus, if he plays $x \in \mathcal{X}$, he cannot compute his actual payoff as he only infer that it will be on the form $\langle x, U \rangle$ for some unknown U that must belong to $\Phi(y)$ (which is also equal to $\Phi(y')$).

When dealing with uncertainties, best replies are extended, following [5, 1], into

$$\text{BR}_1 : \mathcal{P}(\mathbb{R}^A) \rightrightarrows \mathcal{X} \text{ with } \text{BR}_1(\mathcal{U}) = \arg \max_{x \in \mathcal{X}} \inf_{U \in \mathcal{U}} \langle x, U \rangle ,$$

where $\mathcal{P}(\mathbb{R}^A)$ is the family of subsets of \mathbb{R}^A . This is well-defined since $x \mapsto \inf_{U \in \mathcal{U}} \langle x, U \rangle$ is concave and upper semi-continuous hence maxima are attained. $\text{BR}_2 : \mathcal{P}(\mathbb{R}^B) \rightrightarrows \mathcal{Y}$ is defined in a similar way. Definition 1 below of Nash equilibria with partial monitoring (see also [14] for more details and explanations) follows naturally.

Definition 1 $(x^*, y^*) \in \mathcal{X} \times \mathcal{Y}$ is a Nash equilibrium of a game with partial monitoring iff $x^* \in \text{BR}_1(\Phi(y^*))$ and $y^* \in \text{BR}_2(\Psi(x^*))$, i.e., iff

$$x^* \in \arg \max_{x \in \mathcal{X}} \inf_{U \in \Phi(y^*)} \langle x, U \rangle \text{ and } y^* \in \arg \max_{y \in \mathcal{Y}} \inf_{V \in \Psi(x^*)} \langle x, V \rangle .$$

2 A warm-up: semi-standard structure

We first consider an easy case: games with a *semi-standard* information structure. Informally, it implies that action sets are partitioned into subsets of undistinguishable actions (but it is always possible to distinguish between these subsets).

Definition 2 The information of player 1 (and similarly for player 2) is semi-standard if there exists a partition $\{\mathcal{B}_i; i \in \mathcal{I}\}$ of \mathcal{B} such that

- i) If b and b' belong to the same cell \mathcal{B}_i then $\mathbf{H}(b) = \mathbf{H}(b') = \mathbf{H}_i$ and
- ii) The family $\{\mathbf{H}_i; i \in \mathcal{I}\}$ is linearly independent, i.e. if $\sum_{i \in \mathcal{I}} \lambda_i \mathbf{H}_i = \sum_{i \in \mathcal{I}} \gamma_i \mathbf{H}_i$ then λ_i and γ_i must be equal, for every $i \in \mathcal{I}$.

A game has a semi-standard structure if both \mathbf{H} and \mathbf{M} satisfy these properties.

In particular, this means that, for every $y \in \mathcal{Y}$, given $\mathbf{H}(y) \in \mathcal{H}^A$, player 1 can only infer $\{y_i; i \in \mathcal{I}\}$ where $y_i = \sum_{b \in \mathcal{B}_i} y[b]$ is the probability (accordingly to y) of choosing an action in \mathcal{B}_i .

Example 1 If $\mathcal{H} = [0, 1]^d$ and, no matter $b \in \mathcal{B}$, $H(a, b) = H(a', b) = e_b$ where e_b is a vector with only one non-zero coordinate which is 1, then player 1 has a semi-standard information structure. However, if we do not assume that $H(a, b) = H(a', b)$, then this is no longer true.

Indeed, let $\mathcal{A} = \{a, a'\}$, $\mathcal{B} = \{b_1, b_2, b_3, b_4\}$, $\mathcal{H} = [0, 1]^2$ and H be represented as

$$H: \begin{array}{c} a \\ a' \end{array} \begin{array}{|c|c|c|c|} \hline b_1 & b_2 & b_3 & b_4 \\ \hline e_1 & e_2 & e_1 & e_2 \\ \hline e_1 & e_2 & e_2 & e_1 \\ \hline \end{array} \quad \text{with } e_1 = (1, 0) \text{ and } e_2 = (0, 1).$$

The decomposition of point *i*) of Definition 2 must be $\mathbf{H}_1 = (e_1, e_1)$, $\mathbf{H}_2 = (e_2, e_2)$ and so on. However, point *ii*) of the same definition is not satisfied since

$$\mathbf{H} \left(\frac{b_1 + b_2}{2} \right) = \left(\frac{e_1 + e_2}{2}, \frac{e_1 + e_2}{2} \right) \neq \mathbf{H} \left(\frac{b_3 + b_4}{2} \right).$$

In this framework, following Lemma 1 allows an easy reduction from partial to full monitoring. But we need to recall first the general concept of *polytopial complex* (a polytope is the convex hull of a finite number of points¹) on which our results rely:

Definition 3 A finite set $\{P_k; k \in \mathcal{K}\}$ is a polytopial complex of a polytope $P \subset \mathbb{R}^d$ with non-empty interior if:

- i*) For every $k \in \mathcal{K}$, $P_k \subset P$ is a polytope with non empty interior;
- ii*) The union $\bigcup_{k \in \mathcal{K}} P_k$ is equal to P ;
- iii*) Every intersection of two different polytopes $P_k \cap P_{k'}$ has an empty interior.

The following Lemma 1 is an adaptation of an argument stated in [13, Theorem 34].

Lemma 1 There exists a finite subset $\{x_\ell; \ell \in \mathcal{L}\}$ of \mathcal{X} that contains, for every $y \in \mathcal{Y}$, a maximizer of the program $\max_{x \in \mathcal{X}} \min_{U \in \Phi(y)} \langle x, U \rangle$ and such that its convex hull contains the whole set of maximizers. Moreover, there exists a polytopial complex $\{\mathcal{Y}_\ell; \ell \in \mathcal{L}\}$ of \mathcal{Y} such that, for every $\ell \in \mathcal{L}$, x_ℓ is a maximizer on \mathcal{Y}_ℓ .

Similarly, we denote by $\{y_k; k \in \mathcal{K}\}$ the set defined in a dual way for player 2.

Proof: Define, for every $i \in \mathcal{I}$, the set of compatible outcomes with \mathbf{H}_i by:

$$\mathcal{U}_i = \{u(\cdot, y) \in \mathbb{R}^A; y \text{ s.t. } \mathbf{H}(y) = \mathbf{H}_i\} = \text{co} \{u(\cdot, b); b \in \mathcal{B}_i\},$$

where co stands for the convex hull; in particular, $\mathcal{U}_i = \Phi(b)$, for all $b \in \mathcal{B}_i$ and it is a polytope. So the mapping Φ is linear on \mathcal{Y} since² it is defined, for every $y \in \mathcal{Y}$, by

$$\Phi(y) = \sum_{i \in \mathcal{I}} y_i \mathcal{U}_i = \sum_{i \in \mathcal{I}} \sum_{b \in \mathcal{B}_i} y[b] \mathcal{U}_i = \sum_{b \in \mathcal{B}} y[b] \Phi(b).$$

¹A polytope can also be defined, in a totally equivalent way, as a compact and non-empty intersection of a finite number of half-planes

²Actually, the semi-standard structure could also be defined through the linearity of Φ .

Given $x^* \in \mathcal{X}$ and $y \in \mathcal{Y}$, if $U^\# \in \Phi(y)$ is a minimizer of $\min_{U \in \Phi(y)} \langle x^*, U \rangle$ then it can be assumed that $U^\#$ is a vertex of $\Phi(y)$, because a linear program is always minimized on a vertex of the admissible polytope. And necessarily $-x^*$ must belong to the normal cone to $\Phi(y)$ at $U^\#$ [16, Theorem 27.4, page 270]. As a consequence, $-x^*$ must belong to the intersection of $-\mathcal{X}$ and a normal cone; more precisely, since $\langle x, U^\# \rangle$ is linear, $-x^*$ must be one of the vertices (or a convex combination of them) of this intersection.

However, $\Phi(\cdot)$ is linear on \mathcal{Y} , so normal cones at vertices – their set is called normal fan – are constant, see [22, Example 7.3, page 193] and [4, page 530]. As a consequence, there exists a finite number of intersection between $-\mathcal{X}$ and normal cones and they all have a finite number of vertices. The set of every possible vertices is denoted by $-\{x_\ell; \ell \in \mathcal{L}\}$ and it always contains a maximizer (and any maximizer must belong to its convex hull).

Since Φ is linear, $y \mapsto \min_{U \in \Phi(y)} \langle x_\ell, U \rangle$ is also linear, for every $\ell \in \mathcal{L}$; so x_ℓ is a maximizer on a polytopial subset of \mathcal{Y} . \square

Remark 1 *Lemma 1 might be surprising to reader familiar with linear programming. Indeed, it is quite clear that if $u_1(\cdot, y)$ is linear then it is always maximized at one of the vertices of \mathcal{X} . However, in our case, $\min_{U \in \Phi(y)} \langle \cdot, U \rangle$ is not linear but only concave. So it can be maximized anywhere in \mathcal{X} , even in its interior.*

So without some regularity of Φ , the result would obviously be wrong. The key point of the proof is that, in our framework, Φ is itself induced by the minimization of another linear mapping. Lemma 1 holds because $\min_{U \in \Phi(y)} \langle \cdot, U \rangle$ is not just any concave mapping, but it has this extra specific property.

We now introduce an auxiliary game $\tilde{\Gamma}$, with full monitoring, such that its Nash equilibria somehow coincide with Nash equilibria of Γ , the original game. Respective action sets of player 1 and 2 are \mathcal{L} and \mathcal{K} and payoff mappings

$$\tilde{u}(\ell, k) = \min_{U \in \Phi(y_k)} \langle x_\ell, U \rangle \quad \text{and} \quad \tilde{v}(\ell, k) = \min_{V \in \Psi(x_\ell)} \langle y_k, V \rangle.$$

Any pair of mixed actions $(\mathbf{x}, \mathbf{y}) \in \Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ induces a pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ defined by $x = \mathbb{E}_{\mathbf{x}}[x_\ell] \in \mathcal{X}$. This means that, for every $a \in \mathcal{A}$, the weight put by x on a is $x[a] := \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] x_\ell[a]$; similarly, y is defined by $y = \mathbb{E}_{\mathbf{y}}[y_k] \in \mathcal{Y}$.

Theorem 2 *Every Nash equilibrium of $\tilde{\Gamma}$ induces a Nash equilibrium of Γ and, reciprocally, every Nash equilibrium of Γ is induced by a Nash equilibrium of $\tilde{\Gamma}$.*

Proof: Let (\mathbf{x}, \mathbf{y}) be a Nash equilibrium of $\tilde{\Gamma}$ and (x, y) the induced mixed actions. By linearity of Φ , one has $\sum_{k \in \mathcal{K}} \mathbf{y}[k] \Phi(y_k) = \Phi(y)$ thus

$$\begin{aligned} \tilde{u}(\mathbf{x}, \mathbf{y}) &= \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] \sum_{k \in \mathcal{K}} \mathbf{y}[k] \tilde{u}(\ell, k) = \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] \sum_{k \in \mathcal{K}} \mathbf{y}[k] \min_{U_k \in \Phi(y_k)} \langle x_\ell, U_k \rangle \\ &= \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] \min_{U \in \sum_{k \in \mathcal{K}} \mathbf{y}[k] \Phi(y_k)} \langle x_\ell, U \rangle = \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] \min_{U \in \Phi(y)} \langle x_\ell, U \rangle \\ &\leq \min_{U \in \Phi(y)} \left\langle \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] x_\ell, U \right\rangle = \min_{U \in \Phi(y)} \langle x, U \rangle. \end{aligned}$$

Therefore with, respectively, the fact that (\mathbf{x}, \mathbf{y}) is a Nash equilibrium, the linearity of Φ_1 and Lemma 1, this implies that

$$\min_{U \in \Phi(y)} \langle x, U \rangle \geq \tilde{u}(\mathbf{x}, \mathbf{y}) \geq \max_{\ell \in \mathcal{L}} \tilde{u}(\ell, \mathbf{y}) = \max_{\ell \in \mathcal{L}} \min_{U \in \Phi(y)} \langle x_\ell, U \rangle = \max_{x' \in \mathcal{X}} \min_{U \in \Phi(\mathbf{y})} \langle x', U \rangle.$$

Hence we have proved that $x \in \text{BR}_1(\Phi(y))$; similarly $y \in \text{BR}_2(\Psi(x))$, so (x, y) is a Nash equilibrium of Γ .

Reciprocally, let (x, y) be a Nash equilibrium of Γ . Lemma 1 implies that x is a convex combinations of mixed actions in $\{x_\ell; \ell \in \mathcal{L}\}$ that maximize $\min_{U \in \Phi(y)} \langle x_\ell, U \rangle$. Denote by $\mathbf{x} \in \Delta(\mathcal{L})$ this convex combination and define \mathbf{y} in a dual way.

Since $\mathbf{y} \in \Delta(\mathcal{K})$ induces y , then one has, for every $\ell' \in \mathcal{L}$:

$$\tilde{u}(\ell', \mathbf{y}) \leq \max_{x' \in \mathcal{X}} \min_{U \in \Phi(y)} \langle x', U \rangle = \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] \min_{U \in \Phi(y)} \langle x_\ell, U \rangle = \tilde{u}(\mathbf{x}, \mathbf{y}),$$

where we used respectively the linearity of Φ , the fact that $\mathbf{x}_\ell > 0$ if x_ℓ is a maximizer and again the linearity of Φ . Therefore \mathbf{x} is a best reply to \mathbf{y} and the converse is true by symmetry: (\mathbf{x}, \mathbf{y}) is a Nash equilibrium of $\tilde{\Gamma}$. \square

Theorem 2 implies that one just has to compute the set of Nash Equilibria of $\tilde{\Gamma}$ in order to describe the set of Nash equilibria of Γ . For example, one might consider the Lemke-Howson algorithm [8] – or LH-algorithm for short – recalled briefly in the following section.

If $\tilde{\Gamma}$ satisfies some *non-degeneracy* assumption, the LH-algorithm outputs a subset of Nash equilibria of both $\tilde{\Gamma}$ and Γ . The specific assumption and how to modify and apply this algorithm to any game are detailed in [19].

3 Quick reminder on Lemke-Howson algorithm

The Lemke-Howson algorithm of [8] is designed to compute Nash equilibria of a two-player finite game with full monitoring. It is based on the decomposition of \mathcal{X} and \mathcal{Y} into best-replies areas. Recall that $\mathcal{Y}_a := \left\{ y \in \mathcal{Y} \text{ s.t. } a \in \arg\max_{a' \in \mathcal{A}} u(a', y) \right\} \subset \mathcal{Y}$, for any $a \in \mathcal{A}$, is the a -th best-reply area of player 1. The *genericity* assumption required by the LH-algorithm is the following:

Assumption 1 $\{\mathcal{Y}_a; a \in \mathcal{A}\}$ forms a polytopial complex of \mathcal{Y} and any $y \in \mathcal{Y}$ belongs to at most m_y best reply areas \mathcal{Y}_a , where m_y is the size of the support of y . The similar condition holds for $\{\mathcal{X}_B; b \in \mathcal{B}\}$.

Stated otherwise, Assumption 1 means that every $y \in \mathcal{Y}$ has at most m_y best replies.

Each \mathcal{Y}_a is a polytope, so denote by V_2 and E_2 the set of all vertices and edges of these sets (necessarily $\mathcal{B} \subset V_2$). For technical purpose, we also assume that V_2 contains another (abstract) point 0_2 such that $(0_2, b)$ belongs to E_2 for every $b \in \mathcal{B}$. This defines a graph $\mathcal{G}_2 = (V_2, E_2)$ over \mathcal{Y} and similarly a graph $\mathcal{G}_1 = (V_1, E_1)$ over \mathcal{X} . To each vertex $v_2 \in V_2$ (and to each $v_1 \in \mathcal{B}_1$) is associated the following set of labels:

$$L(v_2) := \left\{ a \in \mathcal{A} \text{ s.t. } v_2 \in \mathcal{Y}_a \right\} \cup \left\{ b \in \mathcal{B} \text{ s.t. } v_2[b] = 0 \right\} \subset \mathcal{A} \cup \mathcal{B},$$

i.e., its best replies and pure actions on which it does not put any weight. Label sets of abstract points 0_1 and 0_2 are $L(0_1) = \mathcal{A}$ and $L(0_2) = \mathcal{B}$.

This induces a product labelled graph $\mathcal{G}_0 = (V_0, E_0)$ over $\mathcal{X} \times \mathcal{Y}$, whose set of vertices is the cartesian product $V_0 = V_1 \times V_2$ and such that there exists an edge in E_0 between (v_1, v_2) and (v'_1, v'_2) if and only if $v_1 = v'_1$ and $(v_2, v'_2) \in E_2$ or $v_2 = v'_2$ and $(v_1, v'_1) \in E_1$. The set of labels of (v_1, v_2) is $L(v_1, v_2) = L(v_1) \cup L(v_2)$.

Nash equilibria are exactly *fully labeled* pairs (v_1, v_2) , i.e., if $L(v_1, v_2) = \mathcal{A}_1 \cup \mathcal{A}_2$; indeed, this means that an action a is either not played (if $v_1[a] = 0$) or a best reply to v_2 (if $v_2 \in \mathcal{Y}_a$). The LH-algorithm walks along edges of \mathcal{G}_0 , from vertices to vertices, and stops at a one of those points. We describe quickly in the remaining of this section how it works generically (i.e. for almost all games); for more details we refer to [18, 20] and references therein.

Starting at $v_0 = (0_1, 0_2)$ (which is fully labeled), one label ℓ in $\mathcal{A} \cup \mathcal{B}$ is chosen arbitrarily. The LH algorithm visits sequentially *almost fully labeled* vertices $(v_t)_{t \in \mathbb{N}}$ of \mathcal{G}_0 , i.e., points such that $L(v_t) \supset \mathcal{A} \cup \mathcal{B} \setminus \{\ell\}$ and (v_t, v_{t+1}) is an edge in E_0 . Generically, at any v_t there exists at most one point (apart from v_{t-1}) satisfying both properties, and any end point must be fully labeled.

As a consequence, when starting from any almost fully labeled point vertex, LH algorithm follows either a cycle (and stops when returning to a previously visited point) or a path whose endpoints are necessarily Nash equilibria (or $(0_1, 0_2)$). This property can be used, for example, to prove that the number of Nash equilibria is generically odd.

4 Characterization and computation of Nash equilibria

Without the semi-standard structure, Lemma 1 and Theorem 2 might not hold since Φ is not linear (this is illustrated in Example 2). However, we will show that, in the general case, we still have a similar property: Φ is piece-wise linear. This means that Φ is linear on a polytopial complex of \mathcal{Y} (see the following Lemma 3). Using this, it will be easy to show (in Lemma 4 below) that best-replies areas forms a polytopial complex,

allowing the generalization of LH-algorithm. Such decompositions have been recently used in related frameworks, see e.g. [21].

Example 2 Assume that $\mathcal{A} = \{T; B\}$, $\mathcal{B} = \{L, C, R\}$ and $\mathcal{H} = [0, 1]$. Payoffs and player 1's message matrices (player 2 has full monitoring) are given respectively by:

$$u: \begin{array}{c|ccc} & L & C & R \\ \hline T & (1, 1) & (0, 0) & (0, 0) \\ \hline B & (0, 0) & (1, 2) & (1, 0) \end{array} \quad H: \begin{array}{c|ccc} & L & C & R \\ \hline T & 0 & 1 & 1/3 \\ \hline B & 0 & 1 & 1/3 \end{array}$$

Player 1 cannot distinguish between the mixed action $2/3L + 1/3C$ and the pure action R .

Following notations of Lemma 1, one has $\{x_\ell; \ell \in \mathcal{L}\} = \{T, B, M\}$ where $M = 1/2T + 1/2B$ and³ $\{y_k; k \in \mathcal{K}\} = \{L, C, R\}$. Thus $\tilde{\Gamma}$ is defined by the following matrix:

$$\begin{array}{c|ccc} & L & C & R \\ \hline T & (1, 1) & (0, 0) & (0, 0) \\ \hline B & (0, 0) & (1, 2) & (1/3, 0) \\ \hline M & (1/2, 1/2) & (1/2, 1) & (1/2, 0) \end{array}$$

This game has three Nash Equilibria: (T, L) , (B, C) and $(2/3T + 1/3B, 1/2L + 1/2C)$. Although the first two are indeed Nash equilibria of Γ , this is not true for the last one. Indeed, $\Phi(1/2L + 1/2C) = \{(\lambda/2; 1 - \lambda/2); \lambda \in [0, 1]\}$ and its best response is $\{T\}$.

Actually, and as we shall see in Example 5, Γ has three Nash equilibria which are (T, L) , (B, C) and $(1/3T + 2/3M, 3/4L + 1/4C) = (2/3T + 1/3B, 3/4L + 1/4C)$

Lemma 3 The correspondence Φ is piecewise linear on \mathcal{Y} .

Proof: Since \mathbf{H} is linear from \mathcal{Y} into \mathcal{H}^A , then $\mu \mapsto \mathbf{H}^{-1}(\mu)$ is piecewise linear on \mathcal{H}^A , see [4, page 530] and [15, Proposition 2.4, page 221]. Therefore, by composition, $y \mapsto \mathbf{H}^{-1}(\mathbf{H}(y))$ is piecewise linear on \mathcal{Y} and $y \mapsto u\left(\cdot, \mathbf{H}^{-1}(\mathbf{H}(y))\right)$ – which is by definition Φ – is also piecewise linear on \mathcal{Y} . \square

So Lemma 1 can be rephrased as follows.

Lemma 4 There exists a finite subset $\{x_\ell; \ell \in \mathcal{L}\}$ of \mathcal{X} that contains, for every $y \in \mathcal{Y}$, a maximizer of the program $\max_{x \in \mathcal{X}} \min_{U \in \Phi(y)} \langle x, U \rangle$ and such that its convex hull contains the set of maximizers.

Moreover, for every $\ell \in \mathcal{L}$, x_ℓ is a maximizer on \mathcal{Y}_ℓ which is a finite union of polytopes. Similarly, we denote by $\{y_k; k \in \mathcal{K}\}$ and $\{\mathcal{X}_k; k \in \mathcal{K}\}$ the finite sets for player 2.

Proof: One just has to consider the polytopial complex $\{P_i; i \in \mathcal{I}\}$ with respect to which Φ and Ψ are piecewise linear and apply Lemma 1 on each P_i . \square

Our main result is the following characterization of Nash equilibria in a general game with partial monitoring. We recall that $\mathbf{x} \in \Delta(\mathcal{L})$ induces the mixed action $x = \mathbb{E}_{\mathbf{x}}[x_\ell]$ where $x[a]$, the weight put by x on $a \in \mathcal{A}$, is equal to $\sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell]x_\ell[a]$.

³To be extremely rigorous, the pure action R should be removed since it is never a best response.

Theorem 5 *Nash equilibria of G_H are induced by points in $\Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ that are fully labelled with respect to the two decompositions $\{\mathbf{Y}_\ell; \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k; k \in \mathcal{K}\}$ (and to the label set $\mathcal{L} \cup \mathcal{K}$) defined by*

$$\mathbf{Y}_\ell = \left\{ \mathbf{y} \in \Delta(\mathcal{K}) \text{ s.t. } \mathbb{E}_{\mathbf{y}}[y_k] \in Y_\ell \right\} = \left\{ \mathbf{y} \in \Delta(\mathcal{K}) \text{ s.t. } x_\ell \in \arg \max_{\ell' \in \mathcal{L}} \inf_{U \in \Phi(\mathbb{E}_{\mathbf{y}}[y_k])} \langle x_{\ell'}, U \rangle \right\}$$

and similarly for \mathbf{X}_k .

Proof: Consider any fully labelled point $(\mathbf{x}, \mathbf{y}) \in \Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ and the induced mixed actions $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. By definition (see Section 3), for every $\ell \in \mathcal{L}$ and $k \in \mathcal{K}$, either $\mathbf{x}[\ell] = 0$ or \mathbf{y} belongs to \mathbf{Y}_ℓ (and similarly either $\mathbf{y}[k] = 0$ or $\mathbf{x} \in \mathbf{X}_k$).

As a consequence, x is a best reply to y (and reciprocally) since:

$$\min_{U \in \Phi(y)} \langle x, U \rangle = \min_{U \in \Phi(y)} \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] \langle x_\ell, U \rangle \geq \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] \min_{U \in \Phi(y)} \langle x_\ell, U \rangle \geq \max_{x' \in \mathcal{X}} \min_{U \in \Phi(y)} \langle x', U \rangle.$$

Therefore, any fully labelled point induces a Nash equilibrium of Γ .

Reciprocally, if (x, y) is a Nash equilibrium of Γ then Lemma 4 implies that x and y belong to the convex hull of $\{x_\ell; \ell \in \mathcal{L}\}$ and $\{y_k; k \in \mathcal{K}\}$. More precisely, x is a convex combination of the maximizers of $\min_{U \in \Phi(y)} \langle x_\ell, U \rangle$ (i.e. those x_ℓ such that $y \in Y_\ell$). If we denote this convex combination as $x = \sum_{\ell \in \mathcal{L}} \mathbf{x}[\ell] x_\ell$, then necessarily either $\mathbf{x}[\ell] = 0$ or y belongs to Y_ℓ (and $\mathbf{y} \in \mathbf{Y}_\ell$). Therefore (\mathbf{x}, \mathbf{y}) is fully labeled. \square

It remains to describe why the LH algorithm can be used in this framework. First, recall that every set Y_ℓ or X_k provided by Lemma 4 is a finite union of polytopes. So, up to an arbitrary subdivision of these non-convex unions (associated with maybe a duplication of some mixed actions, see Example 3 below), we can assume that $\{Y_\ell; \ell \in \mathcal{L}\}$ and $\{X_k; k \in \mathcal{K}\}$ are finite families of polytopes.

Lemma 6 *Any element of the families $\{\mathbf{Y}_l; l \in \mathcal{L}\}$ and $\{\mathbf{X}_k; k \in \mathcal{K}\}$ is a polytope.*

Proof: Since, by definition,

$$Y_\ell = \left\{ y \in \mathcal{Y} \text{ s.t. } x_\ell \in \arg \max_{\ell' \in \mathcal{L}} \inf_{U \in \Phi(y)} \langle x_{\ell'}, U \rangle \right\}$$

is a polytope of \mathbb{R}^B , there exists a finite family $\{b_t \in \mathbb{R}^B, c_t \in \mathbb{R}; t \in \mathcal{T}_\ell\}$ such that

$$Y_\ell = \bigcap_{t \in \mathcal{T}_\ell} \left\{ y \in \mathcal{Y} \text{ s.t. } \langle y, b_t \rangle \leq c_t \right\}.$$

Therefore, \mathbf{Y}_ℓ is also a polytope of \mathbb{R}^K as it can be written as

$$\mathbf{Y}_\ell = \bigcap_{t \in \mathcal{T}_\ell} \left\{ \mathbf{y} \in \Delta(\mathcal{K}) \text{ s.t. } \langle \mathbb{E}_{\mathbf{y}}[y], b_t \rangle \leq c_t \right\} = \bigcap_{t \in \mathcal{T}_\ell} \left\{ \mathbf{y} \in \Delta(\mathcal{K}) \text{ s.t. } \langle \mathbf{y}, (\langle y_k, b_t \rangle)_{k \in \mathcal{K}} \rangle \leq c_t \right\}.$$

Similar arguments hold for $\{\mathbf{X}_k; k \in \mathcal{K}\}$. \square

Using this important property, we can generalize the LH-algorithm to games with uncertainties satisfying some non-degeneracy assumptions.

Theorem 7 *If $\{\mathbf{Y}_\ell; \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k; k \in \mathcal{K}\}$ satisfy Assumption 1, then any end-point of Lemke-Howson algorithm induces a Nash equilibrium of Γ .*

Proof: If $\{\mathbf{Y}_\ell; \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k; k \in \mathcal{K}\}$ satisfy the non-degeneracy Assumption 1, any end point of the LH-algorithm is fully labelled, hence a Nash equilibrium of Γ . \square

Remark 2 *It is not compulsory to use the induced polytopial complexes of $\Delta(\mathcal{L})$ and $\Delta(\mathcal{K})$. One can work directly in \mathcal{X} and \mathcal{Y} by considering the projection of the skeleton of the complexes $\{\mathbf{Y}_\ell; \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k; k \in \mathcal{K}\}$ onto them. However, the graphs generated might not be planar and there are, at first glance, no guarantee that the LH-algorithm will work. In the proof of Theorem 5, it is a lifting of the problem that ensures that graphs are planar.*

The fact that there was an odd number of Nash equilibria in the game of Example 2 (continued in Section 5 below) is therefore not surprising; with full monitoring and the non-degeneracy assumption, this can be proved using the LH-algorithm. Therefore, as soon as $\{\mathbf{Y}_\ell; \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k; k \in \mathcal{K}\}$ satisfy this assumption, there will exist an odd number of fully labelled points in $\Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ inducing Nash equilibria.

In some cases, the main argument of the proof of Theorem 5 can be rephrased as follows. The game Γ is, in fact, equivalent to a game $\widehat{\Gamma}$ with full monitoring, with action spaces \mathcal{L} and \mathcal{K} and with payoffs defined in a arbitrary way so that the polytopial complexes induced by the best-replies areas coincide with $\{\mathbf{Y}_\ell; \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k; k \in \mathcal{K}\}$. However, the existence of such abstracts payoffs might not be ensured in general (or it can depend on the duplication of the mixed actions chosen, see Example 3). Anyway, whenever it is possible, it is again almost instantaneous to understand that Nash equilibria of Γ and $\widehat{\Gamma}$ coincide.

Example 3 *Consider the game defined by, respectively, the following payoffs and signal (in \mathbb{R}^2) matrices for the row player:*

	L	M	C	R
T	4	4	4	0
B	3	3	3	3

	L	M	C	R
T	(0,0)	(0,1)	(1,0)	(1,1)
B	(0,0)	(0,1)	(1,0)	(1,1)

Given the signal $(\alpha, \beta) \in [0, 1]^2$, the best response is B if α and β are both bigger than 0.25 and the best response is T is either α or β is smaller than 0.25. Therefore Y_B is convex but Y_T is not (but it is the union of two polytopes).

Assume that the column player has a full monitoring and that his four action might be best responses, then \mathbf{Y}_T is not convex and the decomposition $\{\mathbf{Y}_B, \mathbf{Y}_T\}$ cannot be induced by some equivalent game with full monitoring.

On the other hand, one can find a decomposition of Y_T into two polytopes, namely $Y_{T_1} = H^{-1}\left(\{(\alpha, \beta) \in [0, 1]^2 \text{ s.t. } \alpha \leq \min(0.25, \beta)\}\right)$ and similarly $Y_{T_2} = H^{-1}\left(\{(\alpha, \beta) \in [0, 1]^2 \text{ s.t. } \beta \leq \min(0.25, \alpha)\}\right)$. It is easy to see that $\{\mathbf{Y}_{T_1}, \mathbf{Y}_{T_2}, \mathbf{Y}_B\}$ can be induced by some completely auxiliary game with full monitoring – this decomposition is said to be regular, see [22, Definition 5.3 and page 132]. And with respect to this decomposition, $\mathbf{x} \in \Delta(\{T_1, T_2, B\})$ induces the mixed action $x \in \Delta(\{T, B\})$ defined by $x[T] = \mathbf{x}[T_1] + \mathbf{x}[T_2]$.

5 Examples with partial monitoring or in robust games

Consider again Example 2. The polytopial complexes $\{Y_\ell; \ell \in \mathcal{L}\} := \{Y_B, Y_M, Y_T\}$ and $\{X_k; k \in \mathcal{K}\} := \{X_C, X_L\}$ are represented in the following figure 1.




Figure 1: On the left \mathcal{X} and on the right \mathcal{Y} and $\Delta(\mathcal{K})$ with their complexes.

In order to describe how the LH-algorithm works, we will denote a vertex of the product graph by the cartesian product of its labels (in this example the set of labels is $\{T, B, M, L, R, C\}$); for example the vertex represented with a black dot in figure 1 is denoted by $\{R, T, M\} \times \{B, C, L\}$.

The first step in the LH-algorithm is to drop one label arbitrarily; If the label M is dropped then the first vertex visited by the algorithm is $\{L, R, C\} \times \{B, T, C\}$. The label C appears twice, so in order to get rid of one of them, the algorithm chooses at the next step the vertex $\{L, R, B\} \times \{B, T, C\}$ and the following vertex is $\{L, R, B\} \times \{M, T, C\}$. It is fully labelled, thus an end point of the algorithm, hence $(B, C) \in \mathcal{X} \times \mathcal{Y}$ is a pure Nash equilibrium of Γ .

Similarly, If T is dropped at the first stage, then the first vertex is $\{L, R, C\} \times \{B, M, L\}$ and the second $\{R, C, T\} \times \{B, M, L\}$. So (T, L) is also a pure Nash equilibrium of Γ .

Starting again from this point and dropping the label C makes the LH-algorithm visit $\{R, T, M\} \times \{B, M, L\}$, and then $\{R, T, M\} \times \{B, L, C\}$ which is also a Nash equilibrium. It corresponds to $(\mathbf{x}, \mathbf{y}) = (1/3T + 2/3M, 3/4L + 1/4C) \in \Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ which induces $(x, y) = (2/3T + 1/3B, 3/4L + 1/4C)$ which is a (mixed) Nash equilibrium of Γ .

One can check the remaining vertices of the product graph to be convinced that there does not exist any more equilibria.

We now quickly treat the case of *robust games* where players observe their opponents actions but their payoff mapping is unknown; the only information is that u belongs to some polytope \mathbf{U} (and v to some \mathbf{V}). Then under those assumptions uncertainties correspondence Φ and Ψ might not be piece-wise linear.

Example 4 Assume that the payoff matrix of player 1 belongs to the convex hull of the following two matrices, i.e., $\mathbf{U} = \{\lambda u_1 + (1 - \lambda)u_2; \lambda \in [0, 1]\}$ with

$$\begin{array}{c}
 \begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\
 \begin{array}{c} T_1 \\ u_1 = T_2 \\ B_1 \\ B_2 \end{array} & \begin{array}{|c|c|} \hline 1 & 0 \\ \hline 0 & 0 \\ \hline 0 & 1 \\ \hline 0 & 0 \\ \hline \end{array}
 \end{array}
 \quad \text{and} \quad
 \begin{array}{c}
 \begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\
 \begin{array}{c} T_1 \\ u_2 = T_2 \\ B_1 \\ B_2 \end{array} & \begin{array}{|c|c|} \hline 0 & 0 \\ \hline 1 & 0 \\ \hline 0 & 0 \\ \hline 0 & 1 \\ \hline \end{array}
 \end{array}
 \end{array}$$

Then for any $y \in [0, 1]$,

$$\Phi(yL + (1 - y)R) = \left\{ \left(\begin{array}{c} y\lambda \\ y(1 - \lambda) \\ (1 - y)\lambda \\ (1 - y)(1 - \lambda) \end{array} \right) ; \lambda \in [0, 1] \right\}$$

which is not piece-wise linear in y . Indeed $\Phi(yL + (1 - y)R)$ can be seen as the set of product probability distributions over $\{T, B\} \times \{1, 2\}$ with first marginal $yT + (1 - y)B$.

As a consequence, Lemma 4 might not hold so Lemke-Howson algorithm can not, in general, be extended (see e.g., [1, Section 5] for alternative technics). On the other hand, if both players have only 2 actions, then it is not difficult to see that Φ and Ψ are piecewise linear (as they cannot *turn* as in higher dimensions); so in that specific case, our results extend.

Some open questions

Important questions remains open. We have shown that under some regularity (or non-degeneracy) assumption on the decomposition into best reply areas, Nash equilibria are induced by an odd number of points. The characterization of such games (maybe as a large semi-algebraic class or such that a game chosen uniformly in some open ball satisfy it with probability one) appears to be a real challenging problem here. With full monitoring, one just has to check that vectors $u(\cdot, b)$ and $v(a, \cdot)$ are in some generic position. With partial monitoring, one must first control the fact that x_ℓ and y_k are themselves in generic position, then that \mathbf{Y}_ℓ and \mathbf{X}_k also satisfy regularity conditions; moreover, genericity can be described with respect to the mappings u and v (as in full monitoring) or to H and M , or to simultaneously all of them. Answering this question will most probably require a deeper understanding of how normal cones evolve with u, v, H and M .

Other questions concern wether index and stability of these equilibria can be defined and studied, see [17]: for instance, we can wonder which equilibria remains in a neighborhood of a given game. The complexity of computing these equilibria, and wether it is in the same class than with full monitoring [12], must also be addressed.

Acknowledgments: I am grateful to S. Sorin for his – as always – useful comments and to F. Riedel, B. von Stengel and G. Vigeral for their wise remarks.

References

- [1] Aghassi, M. & Bertsimas, D. (2006). Robust Game Theory *Math. Program., Ser. B* 107, 231–273
- [2] Bade, S. (2010). *Ambiguous act equilibria. Games. Econ. Behav.*, 71, 246–260.
- [3] Ben-Tal, A., El Ghaoui, L. & Nemirovski, A. (2009). *Robust Optimization*, Princeton University Press
- [4] Billera, L.J & Sturmfels, B. (1992) Fiber polytopes. *Ann. Math.*, 135, 527–549
- [5] Gilboa, I. & Schmeidler, D. (1989) Maxmin expected utility with a non-unique prior. *Journal. of Math. Econom.*, 61, 141–153
- [6] Klibanoff, P. (1996) Uncertainty, decision and normal form games *Manuscript*
- [7] Lehrer, E. (2007) Partially specified probabilities: decisions and games. *mimeo*.
- [8] Lemke, C.E. & Howson, J.T. (1964) Equilibrium points of bimatrix games *J. SIAM*, 12, 413–423
- [9] Mertens, J.-F., Sorin, S. & Zamir, S. (1994) *Repeated Games*. CORE discussion paper, 9420–9422.
- [10] Nash, J.F. (1950). Equilibrium points in N -person games *Proc. Nat. Acad. Sci. USA.*, 36, 48–49.
- [11] Nash, J.F. (1951). Non-cooperative games *Ann. Math.*, 54, 286–295.
- [12] Papadimitriou, C. (2007). The complexity of finding Nash Equilibria In: *Algorithmic Game Theory*, eds Nisan, N., Roughgarden, T., Tardos, E. and Vazirani, V, Cambridge University Press, Cambridge, 29–52
- [13] Perchet, V. (2011). Internal Regret with Partial Monitoring Calibration-Based Optimal Algorithms, *Journal of Machine Learning Research*, 12, 1893–1921
- [14] Perchet, V. (2012). A note on robust Nash equilibria in games with uncertainties *manuscript*
- [15] Rambau, K. & Ziegler, G.M. (1996) Projections of polytopes and the generalized Baues conjecture. *Discrete Comput. Geom.*, 16, 215–237
- [16] Rockafellar, R.T. (1970) *Convex Analysis*. Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J.
- [17] von Schemde, A. (2005). *Index and Stability in Bimatrix Games*. Springer, Berlin
- [18] Shapley, L.S. (1974) A note on the Lemke-Howson algorithm *Mathematical Programming Study 1: Pivoting and Extensions*, 4, 22–55

- [19] von Stengel, B. (2002). Computing equilibria for two-person games. In: *Handbook of Game Theory with Economic Applications*, eds R.J. Aumann and S.Hart, Elsevier, Amsterdam, 3, 1723–1759
- [20] von Stengel, B. (2007). Equilibrium computation for two-player games in strategic and extensive form In: *Algorithmic Game Theory*, eds Nisan, N., Roughgarden, T., Tardos, E. and Vazirani, V, Cambridge University Press, Cambridge, 53–78
- [21] von Stengel, B & Zamir, S. (2010). Leadership games with convex strategy sets *Games Econom. Behav.*, 69, 446–457
- [22] Ziegler, G.M. (1995). *Lectures on Polytopes*. Springer, New York