RAIRO-Oper. Res. 48 (2014) 365–371
DOI: 10.1051/ro/2014001

RAIRO Operations Research
www.rairo-ro.org

A NOTE ON ROBUST NASH EQUILIBRIA WITH UNCERTAINTIES *

Vianney Perchet[1]

Abstract. In this short note, we investigate the framework where agents or players have some uncertainties upon their payoffs or losses, the behavior (or the type, number or any other characteristics) of other players. More specifically, we introduce an extension of the concept of Nash equilibria that generalize different solution concepts called by their authors, and depending on the context, either as robust, ambiguous, partially specified or with uncertainty aversion. We provide a simple necessary and sufficient condition that guarantees its existence and we show that it is actually a selection of conjectural (or self-confirming) equilibria. We finally conclude by how this concept can and should be defined in games with partial monitoring in order to preserve existence properties.

Keywords. Robust games, robust Nash equilibria, uncertainties, partial monitoring, conjectural equilibria.

Mathematics Subject Classification. 91A10, 91B52.

1. Introduction

Uncertainties or ambiguities have been introduced in several fields (such as optimization [5, 6], operations research [7–9, 14] or game theory [1, 4, 12]) to take into account the fact that real agents do not have a perfect knowledge of their environment, an infinite memory or rationality, *etc.* There are basically two different approaches to handle them, either to assume the existence of an underlying

Received February 20, 2013. Accepted November 6, 2013.

* *Part of this research was supported by the ANR (ANR-10-BLAN-0112).*

[1] Laboratoire de Probabilités et de Modèles Aléatoires, UMR 7599, Université Paris Diderot, 8 place FM/13, 75013 Paris, France. vianney.perchet@normalesup.org

Article published by EDP Sciences

© EDP Sciences, ROADEF, SMAI 2014

https://doi.org/10.1051/ro/2014001 Published online by Cambridge University Press

probability distribution (sometimes refereed as a *stochastic* or *Bayesian* approach) or to accept the existence of non-unique environment and to find a solution that behaves well in any of them (this could be refereed as a *minmax approach* or *maximization with a non-unique prior*, following [12]).

This note is dedicated to some theoretical consequences of uncertainties on game theory and on the concept of Nash equilibrium, the latter needing to be adapted. Indeed, its original definition strongly relies on the absence of ambiguity. Quoting Nash himself, we recall that equilibria are "N-tuple of [mixed] strategies, one for each player, [that] may be regarded as a point in the product space obtained by multiplying the N strategy spaces of the players" [20], "such that each player's mixed strategy maximizes his payoff if the strategies of the others are held fixed. Thus each player's strategy is optimal against those of the others" [21]. It implicitly assumes that each player knows not only his own reward mapping, but also his set of opponents and their strategies. However, in some cases (which can rely on empirical or intuitive results as in Ellsberg's paradox [10]), they do not have perfect knowledge of their own payoff mapping (or preferences can not be represented by such mappings) or might have only partial information upon their opponents, their action set (as in games on networks, or played on internet), *etc.*

Different, yet quite related, models have emerged to encompass uncertainties. From the minmax approach, we can cite *robustness* [1] (when payoff mappings are unknown, as a reference to robust optimization [6]) or *uncertainty aversion* [16], *ambiguity* [3], *partially specified probabilities* [17] (when strategies actually played at an equilibrium is not perfectly known to players, their only knowledge is that they must belong to some given sets). In the stochastic approach, players formulate a *conjecture* upon their opponents, and then maximize payoffs with respect to this conjecture; this is the idea behind *conjectural* and *self-confirming* equilibria [4, 11, 15].

The main objective of this note is to formalize and unify in a general framework different notions of Nash equilibria with uncertainties (we decided to keep the name of robust Nash equilibria) and to provide a simple necessary and sufficient condition guaranteeing their existence. We also show how the aforementioned two different approaches articulate, as our concept is in fact a selection of conjectural equilibria. We finally describe how equilibria should be defined in games with partial monitoring [18, 19, 24]. Results are, as often as possible, based on simple and illustrating examples.

2. Robust Nash equilibria

We consider N-player games where the action set of player $n \in \mathcal{N} := \{1, \dots, N\}$ is denoted by $\mathcal{X}_n \subset \mathbb{R}^{A_n}$ and his payoff mapping by $u_n : \mathcal{X} \to \mathbb{R}$, where $\mathcal{X} = \prod_{m\in\mathcal{N}} \mathcal{X}_m$. We assume that $\mathcal{X}_n$ is a compact and convex set and u_n is multilinear, which means that $u_n(\cdot, x_{-n})$ is linear for every $x_{-n} \in \mathcal{X}_{-n} := \prod_{m\neq n} \mathcal{X}_m$.

Let us recall basic facts on Nash equilibria. Define, for every $n \in \mathcal{N}$, the best reply correspondence BR_n from $\mathbb{R}^{A_n}$ to $\mathcal{X}_n$ by $\mathrm{BR}_n(U_n) := \arg\max_{x_n\in\mathcal{X}_n} \langle x_n, U_n\rangle$.

https://doi.org/10.1051/ro/2014001 Published online by Cambridge University Press

Then $x^* = (x_1^*, \dots, x_n^*) \in \mathcal{X}$ is a Nash equilibrium iff there exists, for every $n \in N$, $U_n \in \mathbb{R}^{A_n}$ satisfying $x_n^* \in \mathrm{BR}_n(U_n)$ and $U_n = u_n\left(\cdot, x_{-n}^*\right)$. As it is a fixed point of $x \to \prod_{n\in\mathcal{N}} \mathrm{BR}_n\left(u_n(\cdot, x_{-n})\right)$, its existence is ensured by Kakutani's theorem. This definition of equilibria highlights the fact that an equilibrium can be decomposed into two components: an *optimization* one, as every player is maximizing again U_n, and a *compatibility* one, as $U_n = u_n(\cdot, x_{-n}^*)$.

Best replies are extended with uncertainties in the minmax approach [12], into

$$\mathrm{BR}_n : \mathcal{P}(\mathbb{R}^{A_n}) \to \mathcal{P}(\mathcal{X}_n) \quad \text{with} \quad \mathrm{BR}_n(\mathcal{U}_n) = \arg \max_{x_n \in \mathcal{X}_n} \inf_{U_n \in \mathcal{U}_n} \langle x_n, U_n \rangle,$$

where, for any set E, $\mathcal{P}(E)$ is the family of its subsets. This is well-defined since $x \mapsto \inf_{U_n\in\mathcal{U}_n} \langle x, U_n\rangle$ is concave and upper semi-continuous hence maxima are attained.

Remark 2.1. Evaluation of payoff decreases with respect to uncertainties, *i.e.*, the subset $\mathcal{U}_n$. Indeed, if $\mathcal{V}_n \subset \mathcal{U}_n$, then $\inf_{V_n\in\mathcal{V}_n} \langle x_n, V_n\rangle \geq \inf_{U_n\in\mathcal{U}_n} \langle x_n, U_n\rangle$, for every $x_n \in \mathcal{X}_n$. This is referred as "uncertainty aversion" of players, see for instance [12]: the more information a player has, the more he values his payoff.

No assumptions are made on origins or structure of uncertainties; they are represented, for every $n \in \mathcal{N}$, by a given mapping $\Phi_n : \mathcal{X}_{-n} \to \mathcal{P}(\mathbb{R}^{A_n})$. The original framework, called *full monitoring case*, corresponds to $\Phi_n(x_{-n}) = \left\{u_n(\cdot, x_{-n})\right\}$.

Definition 2.2. $x^* = (x_1^*, \dots, x_n^*) \in \mathcal{X}$ is a robust Nash equilibrium iff there exists, for every $n \in \mathcal{N}$, $\mathcal{U}_n \subset \mathbb{R}^{A_n}$ satisfying $x_n^* \in \mathrm{BR}_n(\mathcal{U}_n)$ and $\mathcal{U}_n = \Phi_n\left(x_{-n}^*\right)$.

Here again, our concept of equilibria has an optimization and a compatibility component. Existence of Robust Nash equilibria is ensured under a mild regularity assumption, namely the continuity of the mappings Φ_n.

We recall that a multivalued mapping $\Psi : \mathbb{R}^k \to \mathcal{P}(R^d)$ is continuous at $x^* \in \mathbb{R}^k$ if it is upper semi-continuous (for every sequence $x_n \in \mathbb{R}^k$ converging to x^*, if $z_n \in \Phi(x_n)$ converges to some $z \in \mathbb{R}^d$, then $z \in \Phi(x^*)$) and lower semi-continuous (for any sequence $x_n \in \mathbb{R}^k$ converging to x^* and any $z^* \in \Phi(x^*)$, there exists a subsequence n_k and $z_k \in \Phi(x_{n_k})$ such that z_{n_k} converges to z^*).

Proposition 2.3. *If every Φ_n is continuous, there exist robust Nash equilibria.*

Proof. Robust Nash equilibria are fixed points of the correspondence defined, for every $x \in \mathcal{X}$, by $\mathrm{BR}\left[\Phi(x)\right] = \prod_{n\in\mathcal{N}} \mathrm{BR}_n\left[\Phi_n(x_{-n})\right]$, which is always a compact non-empty convex subset of $\mathcal{X}$. If Φ is continuous then $\mathrm{BR}\left[\Phi(\cdot)\right]$ has a closed graph, hence by Kakutani's theorem, it has fixed points that are Nash equilibria. □

https://doi.org/10.1051/ro/2014001 Published online by Cambridge University Press

An alternative proof[2] of this result is to appeal to the existence theorem of Nash−Glicksberg [13]

A key property is that existence of Nash equilibria is not implied only by either upper nor lower semi-continuity of Φ, as illustrated in the following Example 2.4.

Example 2.4. Consider the following bi-matrix game, whose unique Nash equilibrium in full monitoring is $(x^*, y^*) = (1/2T + 1/2B, 2/3L + 1/3R)$, defined by

	L	R
T	1;0	0;1
B	0;1	2;0

$\mathcal{X}_1 = \Delta(\{T, B\})$, $\mathcal{X}_2 = \Delta(\{L, R\})$, $\Phi_2(x) = \{u_2(x, \cdot)\}$, $\Phi_1(y^*) = \{u_1(\cdot, y); y \in \mathcal{X}_2\}$ and $\Phi_1(y) = \{u_1(\cdot, y)\}$ otherwise.
Φ_2 is continuous, Φ_1 upper semi-continuous, yet no robust Nash equilibrium exists.

If Φ_1 is modified into $\Phi_1'(R) = \{u_1(\cdot, R)\}$ and $\Phi_1'(y) = \{u_1(\cdot, y); y \in \mathcal{X}_2\}$ otherwise, then Φ_1' is lower semi-continuous and no robust Nash equilibrium exists.

Equilibria concepts that can be found in the literature correspond to specific structures of Φ_n: define for instance, as in [1], $\Phi_n(x_{-n}) = \{u(\cdot, x_{-n}); u \in \mathbf{U}\}$ where $\mathbf{U}$ is some given convex family of possible payoff mappings or, as in [17], $\Phi_n(x_{-n}) = \{u_n(\cdot, q_{-n}); q_{-n} \in \prod_{m \neq n} \mathcal{X}_m[x_m]\}$ where $\mathcal{X}_m[x_m] \subset \mathcal{X}_m$ is defined by a small number of linear (in x_{-m}) mappings.

Similarly to the full monitoring case, we can naturally define ε-robust equilibria.

Definition 2.5. Given $\varepsilon > 0$, $x^* = (x_1^*, \dots, x_n^*) \in \mathcal{X}$ is an ε-robust equilibrium iff there exists, for every $n \in \mathcal{N}$, $\mathcal{U}_n \subset \mathbb{R}^{A_n}$ satisfying $\mathcal{U}_n = \Phi_n\left(x_{-n}^*\right)$ and

$$\inf_{U_n \in \mathcal{U}_n} \langle x_n^*, U_n \rangle \geq \sup_{x_n \in \mathcal{X}_n} \inf_{U_n \in \mathcal{U}_n} \langle x_n, U_n \rangle - \varepsilon.$$

It is possible to relate perturbation of robust equilibria and ε-robust equilibria.

Proposition 2.6. *If Φ is a continuous with a bounded range then, for every $\varepsilon > 0$, there exists $\delta_\varepsilon > 0$ such that every δ-perturbation of a robust Nash equilibria, for $\delta \leq \delta_\varepsilon$, is an ε-robust Nash equilibria.*

If Φ is continuous but with an unbounded range then, no matter $\varepsilon > 0$, there might exist no ε-robust equilibria apart from robust equilibria.

Proof. The first part is an immediate consequence of continuity of Φ and of $x \mapsto \inf_{U \in \mathcal{U}} \langle x, U \rangle$ if $\mathcal{U}$ is bounded.

For the second part, consider $N = 2$, $\mathcal{X}_1 = \mathcal{X}_2 = \{(x, y) \in \mathbb{R}_+^2 \text{ s.t. } x + y = 1\}$ and both $\Phi_1 = \Phi_2$ are constant equal to

$$\Phi_1(x_2) = \Phi_2(x_1) := \{(\alpha, \beta) \text{ s.t. } \alpha \geq 0, \beta \leq 0\} =: \mathcal{U}^*, \quad \forall x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2.$$

There exists a unique robust equilibria $x^* = (x_1^*, x_2^*)$, with $x_1^* = x_2^* = (1, 0)$, and no other ε-Nash equilibria. Indeed $\inf_{U \in \mathcal{U}^*} \langle x_i^*, U \rangle = 0$ and, no matter $x \neq x_i^*$, $\inf_{U \in \mathcal{U}^*} \langle x, U \rangle = -\infty$. □

[2]We thank an anonymous referee for this suggestion.

https://doi.org/10.1051/ro/2014001 Published online by Cambridge University Press

3. Selection of conjectural equilibria

Conjectural, *self-confirming* or *subjective* equilibria [4, 11, 15] can be related to robust Nash equilibria. Recall that $x^* \in \mathcal{X}$ is a conjectural equilibrium of a game with uncertainties if, for every $n \in \mathcal{N}$, there exists a conjecture $V_n^\star$ on the possible set of outcomes, *i.e.*, an element of $\mathrm{co}\left(\Phi(x^*_{-n})\right)$, the convex hull of $\Phi(x^*_{-n})$, such that x_n^* is a best reply to $\{V_n^\star\}$, see *e.g.* [2]. Equivalently, $V_n^\star$ can be represented as a probability distribution over $\Phi(x^*_{-n})$.

Sets of conjectural equilibria can be very large and even equal to $\mathcal{X}$. This happens in every game (without strictly dominated strategies) such that $\Phi_n(\cdot) = \{u_n(\cdot, x_{-n})\}$. As player actually observe nothing, those games are called *in the dark* and they have generically only one robust Nash equilibrium, each player playing his maxmin strategy.

Proposition 3.1. *A robust Nash equilibrium is a conjectural equilibrium.*

Proof. Any robust Nash equilibrium x^* satisfies, by linearity of $\langle x, \cdot \rangle$,

$$x_n^* \in \arg\max_{x \in \mathcal{X}_n} \min_{U_n \in \Phi_n(x^*_{-n})} \langle x, U_n \rangle = \arg\max_{x \in \mathcal{X}_n} \min_{V_n \in \mathrm{co}\left(\Phi_n(x^*_{-n})\right)} \langle x, V_n \rangle.$$

So x_n^* is an optimal strategy in the zero-sum game with action sets $\mathcal{X}_n$, $\mathrm{co}\left(\Phi_n(x^*_{-n})\right)$ and payoff $\langle x, V_n \rangle$. It remains to let V_n^* be any optimal strategy of the second player. □

A conjecture could also be defined as a subset of possible outcomes instead of a probability distribution upon them, as in [16]: an equilibria is then a pair $(x^*, \mathcal{U})$ such that $x_n^* \in \mathcal{U}_n$ and $u_n(\cdot, x^*_{-n}) \in \mathcal{U}_n$. A robust Nash equilibria obviously still satisfies this requirement.

4. Equilibria of games with partial monitoring

An important class of games where uncertainties appear are *finite game with partial monitoring*, see [19, 24]. Here sets of *pure and mixed action* of player n are respectively $\mathcal{A}_n$ and $\mathcal{X}_n = \Delta(\mathcal{A}_n)$ and players do not observe actions of their opponents but they receive *messages*. Formally, there exist a convex compact set of messages $\mathcal{H}$ and signaling mappings H_n from $\mathcal{A} := \prod_{n \in \mathcal{N}} \mathcal{A}_n$ into $\mathcal{H}$, extended multi-linearly to $\mathcal{X}$. Given $a \in \mathcal{A}$, player n receives the message $H_n(a)$.

No matter his choice of actions, player n cannot distinguish between x_{-n} and x'_{-n} in $\mathcal{X}_{-n}$ satisfying $H_n(a, x_{-n}) = H_n\left(a, x'_{-n}\right)$ for all $a \in \mathcal{A}_n$. As usual in this setup [18, 22, 24], we define the *maximal informative mapping* $\mathbf{H}_n : \mathcal{X}_n \to \mathcal{H}^{A_n}$ by:

$$\forall\, x_{-n} \in \mathcal{X}_{-n},\ \mathbf{H}_n(x_{-n}) = \left[H_n\left(a, x_{-n}\right)\right]_{a \in \mathcal{A}_n} \in \mathcal{H}^{A_n}.$$

These mappings naturally define the correspondences $\Phi_n : \mathcal{X}_{-n} \to \mathcal{P}(\mathbb{R}^{A_n})$ by:

$$\Phi_n\left(x_{-n}\right) := \left\{u_n(\cdot, x'_{-n}) \in \mathbb{R}^{A_n};\ \mathbf{H}_n\left(x'_{-n}\right) = \mathbf{H}_n\left(x_{-n}\right)\right\}. \tag{4.1}$$

https://doi.org/10.1051/ro/2014001 Published online by Cambridge University Press

Definition 4.1. $x^* \in \mathcal{X}$ is a Nash equilibrium of a game with partial monitoring $\mathbf{H}$ iff it is a robust Nash equilibrium, with uncertainties Φ_n defined by equation (4.1).

$\mathbf{H}_n$ and u_n are continuous, so Φ is continuous and Nash equilibria always exist.

Example 4.2. Consider the game with payoffs given by the left matrix and $\mathcal{H} = \{a, b, c\}$. Player 2 has full monitoring, so H_2 is not represented, and H_1 is the matrix on the right:

$$u_1; u_2 = \begin{array}{c|ccc} & L & M & R \\ T & 2;0 & 1;0 & 1;2 \\ B & 0;1 & 2;0 & 2;2 \end{array} \qquad \text{and} \qquad H_1 = \begin{array}{c|ccc} & L & M & R \\ T & a & a & b \\ B & c & c & c \end{array}$$

Actions L and M are undistinguishable so, for every $\lambda \in [0,1]$ and $\eta \in [0, \lambda]$:

$$\begin{aligned} \Phi_1\Big(\lambda L + (1-\lambda)R\Big) &= \Phi_1\Big(\lambda M + (1-\lambda)R\Big) = \Phi_1\Big(\eta L + (\lambda-\eta)M + (1-\lambda)R\Big) \\ &= \Big\{(1+\gamma, 2-2\gamma)\,;\ \gamma \in [0,\lambda]\Big\}, \end{aligned}$$

where $(1+\gamma, 2-2\gamma)$ are respective payoffs of T and B for some γ. As a consequence:

$$\mathrm{BR}_1\left(\Phi_1\big(\lambda L + (1-\lambda)R\big)\right) = \begin{cases} \{2/3T + 1/3B\} & \text{if } \lambda < 1/3 \\ \{B\} & \text{if } \lambda > 1/3 \\ \Delta(\{T,B\}) & \text{if } \lambda = 1/3, \end{cases}$$

and, since R is a strictly dominating strategy, the only Nash equilibrium is (B, R).

A usual objection is that, given $x^* = (x_n^*, x_{-n}^*)$, there might exist $a_n \in \mathcal{A}_n$ such that $x_n^*[a_n]$, the weight put by x_n^* on a_n, is zero; stated otherwise, these a_n are not in $\mathrm{supp}(x_n^*)$, the support of x_n^*. So, player n cannot observe $H_n(a_n, x_{-n}^*)$ nor compute $\mathbf{H}_n(x_n^*)$, as in Example 4.2. So we should consider instead of Φ_n the following correspondence $\widehat{\Phi}_n : \mathcal{X} \to \mathbb{R}^{A_n}$ defined by

$$\widehat{\Phi}_n(x) = \left\{u_n(\cdot, x'_{-n}) \in \mathbb{R}^{A_n};\ H_n(a_n, x'_{-n}) = H_n(a_n, x_{-n}),\ \forall a_n \in \mathrm{supp}(x_n^*)\right\}.$$

Proposition 4.3. *With respect to $\widehat{\Phi}$, there exist games without any equilibria or such that any perturbation of equilibria is not an ε-equilibrium (even if $\widehat{\Phi}$ has a bounded range).*

Proof. In Example 4.2, if (B, R) is played, the only message received is c so $\widehat{\Phi}(B, R) = \{(1+\gamma, 2-2\gamma)\,; \gamma \in [0,1]\}$. Its best reply is T; yet, for every $\delta \in (0,1]$, best reply to $\widehat{\Phi}(\delta T + (1-\delta)B, R) = \{(1,2)\}$ is B. So this game has no equilibria.

For the second part of the proposition, consider the following two players game:

$$(u_1, u_2) = \begin{array}{c|cc} & L & R \\ T & -1;0 & 1;1 \\ B & 0;0 & 0;1 \end{array} \qquad \text{and} \qquad H_1 = \begin{array}{c|cc} & L & R \\ T & a & b \\ B & c & c \end{array}$$

https://doi.org/10.1051/ro/2014001 Published online by Cambridge University Press

(B,R) is an equilibrium since $\widehat{\Phi}(B,R)=\Big\{(\lambda,0)\,;\ \lambda\in[-1,1]\Big\}$. However, for every $\delta>0, \widehat{\Phi}\Big(\delta T+(1-\delta)B,R\Big)=\Big\{(1,0)\Big\}$, so $\delta T+(1-\delta)B$ is not an ε-equilibrium. □

Random messages can be embedded into this framework. Assume that there exists a finite set $\mathcal{S}$ and given $a\in\mathcal{A}$, player n receives the signal $s\in\mathcal{S}$ of law $s_n(a)\in\Delta(\mathcal{S})$. We define $\mathcal{H}:=\Delta(\mathcal{S})$ and $\mathbf{H}_n(a_{-n}):=\big[s_n(a,a_{-n})\big]_{a\in\mathcal{A}_n}$. Although $\mathbf{H}_n(a_{-n})$ is a vector of laws, unbiased estimators can estimate it at an arbitrarily small cost, see *e.g.* [22].

References

[1] M. Aghassi and D. Bertsimas, Robust game theory. *Math. Program., Ser. B* **107** (2006) 231–273.
[2] Y. Azrieli, On pure conjectural equilibrium with non-manipulable information. *Int. J. Game Theory* **38** (2009) 209–219.
[3] S. Bade, Ambiguous act equilibria. *Games. Econ. Behav.* **71** (2010) 246–260.
[4] P. Battigalli and D. Guaitoli, *Conjectural equilibrium.* Mimeo (1988).
[5] A. Ben-Tal and A. Nemirovski, Robust convex optimization. *Math. Oper. Res.* **23** (1998) 769–805.
[6] A. Ben-Tal, L. El Ghaoui and A. Nemirovski, *Robust optimization.* Princeton University Press (2009).
[7] A. Ben-Tal, B. Golany and S. Shtern, Robust multi echelon multi period inventory control. *Eur. J. Oper. Res.* **199** (2009) 922–935.
[8] D. Bertsimas. and A. Thiele, Robust and data-driven optimization: modern decision-making under uncertainty, in *Tutorials on Oper. Res.* (2006) Chap. 4, 122–195.
[9] A. Candia-Véjar, E. Álvarez-Miranda and N. Maculan, Minmax regret combinatorial optimization problems: an algorithmic perspective. *RAIRO – Oper. Res.* **45** (2011) 101–129.
[10] D. Ellsberg, Risk, ambiguity and the Savage axiom. *Quat. J. Econom.* **75** (1961) 643–669.
[11] D. Fudenberg and D.K. Levine, Self-confirming equilibrium. *Econometrica* **61** (1993) 523–545.
[12] I. Gilboa and D. Schmeidler, Maxmin expected utility with a non-unique prior. *J. Math. Econ.* **61** (1989) 141–153.
[13] I. Glicksberg, A further generalization of the Kakutani fixed point theorem, with applications to Nash equilibrium points. *Proc. Am. Math. Soc.* **3** (1952) 170–174.
[14] C. Hitch, Uncertainties in operations research. *Oper. Res.* **8** (1960) 437–445.
[15] E. Kalai and E. Lehrer, Subjective equilibrium in repeated games. *Econometrica* **61** (1993) 1231–1240.
[16] P. Klibanoff, Uncertainty, decision and normal form games. Mimeo (1996).
[17] E. Lehrer, Partially specified probabilities: decisions and games. *Am. Econ. J. Micro.* **4** (2012) 70–100.
[18] G. Lugosi, S. Mannor and G. Stoltz, Strategies for prediction under imperfect monitoring. *Math. Oper. Res.* **33** (2008) 513–528.
[19] J.-F. Mertens, S. Sorin and S. Zamir, *Repeated games, CORE discussion paper* (1994) 9420–9422.
[20] J.F. Nash, Equilibrium points in N-person games. *Proc. Natl. Acad. Sci.* USA **36** (1950) 48–49.
[21] J.F. Nash, Non-cooperative games. *Ann. Math.* **54** (1951) 286–295.
[22] V. Perchet, Internal regret with partial monitoring, calibration-based optimal algorithms. *J. Mach. Learn. Res.* **12** (2011) 1893–1921.
[23] F. Riedel and L. Sass, The strategic use of ambiguity. Mimeo (2011).
[24] A. Rustichini, Minimizing regret: the general case. *Games Econom. Behav.* **29** (1999) 224–243.

https://doi.org/10.1051/ro/2014001 Published online by Cambridge University Press