UNIVERSIDAD CARLOS III DE MADRID
working papers

Working Paper 04-16
Economics Series 03
March 2004

Departamento de Economía
Universidad Carlos III de Madrid
Calle Madrid, 126
28903 Getafe (Spain)
Fax (34) 91 624 98 75

UNIFORM CONTINUITY OF THE VALUE OF ZERO-SUM GAMES WITH DIFFERENTIAL INFORMATION*

Ezra Einy[1], Ori Haimanko[2], Diego Moreno[3] and Benyamin Shitovitz[4]

Abstract

We establish uniform continuity of the value for zero-sum games with differential information, when the distance between changing information fields of each player is measured by the Boylan (1971) pseudo-metric. We also show that the optimal strategy correspondence is upper semi-continuous when the information fields of players change, even with the weak topology on players' strategy sets.

Keywords: Zero-Sum Games, Differential Information, Value, Optimal Strategies, Uniform Continuity.

[1] Einy, Departament of Economics, Ben Gurion University of the Negev, Beer Sheva, Israel 84105. E-mail: einy@bgumail.bgu.ac.il
[2] Haimanko, Departament of Economics, Ben Gurion University of the Negev, Beer Sheva, Israel 84105. E-mail: orih@bgumail.bgu.ac.il
[3] Moreno, Departamento de Economía, Universidad Carlos III de Madrid, 28903 Getafe, Spain. E-mail: dmoreno@eco.uc3m.es;
[4] Shitovitz, Department of Economics, University of Haifa (Haifa, Israel 31905). E-mail: binya@econ.haifa.ac.il

*Moreno acknowledges the support of the Ministerio de Ciencia y Tecnología, grant BEC2001-0973.

1 Introduction

Bayesian games, or games with incomplete information, describe situations in which there is uncertainty about players' payoffs, and different players have (typically) different private information about the realized state of nature ω that determines the payoffs. The private information of a player i can often be represented by a partition of the space Ω of all states of nature (in which case i knows to which element of the partition the realized ω belongs), or more generally, by a σ-field z^i of measurable sets in Ω (in which case i knows, given a set in z^i, whether the realized ω is located in this set). If the attention is confined to two-person zero-sum games with incomplete information, each player has an optimal strategy and the value of a game is well defined, under quite general conditions on the expected payoff function (see Sion (1958)). This work concerns continuity of the value of a game, as a function of players' information endowments (fields), when the closeness of fields is measured by means of the Boylan (1971) pseudo-metric.

It turns out that the value has strong continuity properties. We find that, when the payoff function is Lipschitz-continuous in strategies at each state of nature,[1] the value is a uniformly continuous function of players' information fields (see Theorem 1).[2] If, in addition, the state-dependent Lispchitz constant of the payoff function is bounded, then the value is in fact a Lipchitz-continuous function of the information fields (see Corollary 1). Moreover, the correspondence describing players' optimal strategies as a function of information is upper semi-continuous, even with respect to the weak convergence topology on each player's set of strategies (see Theorem 3).

These continuity properties of the value (and optimal strategies) in zero-sum games stand somewhat in contrast to the well-known discontinuity of the Bayesian Nash equilibrium (NE) correspondence[3] in general (non zero-sum) games with incomplete information. The NE correspondence is not lower semi-continuous – that is, NE strategies may not be approachable by NE strategies in games with slightly modified information endowments – as was established by, e.g., Monderer and Samet (1996)[4] in a setting identical to ours. However, it also may not be upper semi-continuous as we show here (see Remark 2, where we consider a simple coordination game), because of the weak mode of strategy convergence that we assume. While this mode of convergence suffices for upper semi-continuity of the optimal strategy correspondence in zero-sum

[1] This requirement is satisfied, for instance, by games which have the matrix-game form in all states of nature.

[2] This result requires a mild assumption of q-integrability on the state-dependent Lipschitz constant. When this constant is merely integrable, the value is also continuous (see Theorem 2), but not uniformly.

[3] The NE payoffs correspondence is also discontinuous.

[4] In fact, Monderer and Samet (1996), as well as Kajii and Morris (1998) in a fixed-types model of incomplete information, are concerned precisely with the question of what topology on information endowments (or information structure) would lead to lower semi-continuity of NE.

games, it fails to do similar work in general games[5]. This di¤erence emphasizes the important role played by the zero-sum assumption when the continuity of equilibrium strategies is considered.

The continuity of the NE correspondence with respect to changes in information has been studied by other authors. In this paper, we use the basic set-up of Monderer and Samet (1996), who work with information …elds to describe players' varying private information, with the common prior distribution of the states of nature (prior belief) …xed at all times. This follows a certain tradition of modelling information in economic theory (see, e.g., Allen (1983), Cotter (1986), Stinchcombe (1990), VanZandt (2002), and Einy et al (2003)). However, there is another approach to continuity of NE correspondences, which is with respect to players' prior beliefs (see, e.g., Milgrom and Weber (1985), Kajii and Morris (1998))[6]. In this approach, prior beliefs are variable, but the rest of the information structure (in which the space of states of nature is assumed to be the cross product of the sets of players' types[7], and each player's private information is given by the knowledge of his type) is …xed throughout. Perturbing the common prior belief in‡uences the expected payo¤s of all agents, but does not a¤ect the players' strategy sets. However, our setting emphasizes di¤erences in information, allowing the information structure of a game to be perturbed in a way that directly a¤ects only one individual player, or in a way that a¤ects all players di¤erently. Indeed, a change in the private information of both players induces (typically di¤erent) changes in players' strategy sets, due to the constraint of the strategy's measurability with respect to the player's information …eld. While the impact of these information changes on the structure of the game might appear to be signi…cant, our theorems show that the value and the optimal strategies in zero-sum games are nevertheless well behaved with respect to these changes.

Our paper is organized as follows. The set-up is described in Section 2. Our results (Theorems 1, 2, 3 and Corollaries 1, 2) are stated and proved in Section 3; Remarks 1 and 2 appear at the end of this section. The Appendix contains the proof of a technical Lemma 1.

2 Preliminaries

We consider zero-sum games with two players, $i = 1, 2$. Games are played in an uncertain environment, which a¤ects payo¤ functions of the players. The underlying uncertainty is described by a probability space $(- , z , \mu)$, where - is the space of states of nature, z is a σ-…eld of subsets of -, and μ is a countably additive probability measure on $(- , z)$, which represents the common prior

[5] Even the payo¤ functions in a typical game would be discontinuous in the weak topology on strategies.

[6] In this context, Milgrom and Weber (1986) established upper semi-continuity of the NE correspondence under certain conditions on the information structure. The objective of Kajii and Morris (1998), as we already mentioned, is to …nd ways to obtain lower semi-continuity of the NE payo¤s correspondence.

[7] A set representing other uncertainties (not type-related) is also taken in the cross product.

of the players regarding the realized state of nature. The initial information endowment of player i is given by a σ-subfield $\mathcal{F}^i$ of $\mathcal{F}$.

Each player $i = 1, 2$ has a set S^i of strategies, which is a convex and compact subset of a Euclidean space R^{n_i}. We will assume, without loss of generality, that $\max_{s \in S^1 \cup S^2} \|s\| \leq 1$, where $\|\cdot\|$ stands for the Euclidean norm in R^{n_1} or R^{n_2}. There is, in addition, a measurable[8] real valued payoff function $u : \Omega \times S^1 \times S^2 \to R$, such that $u\left(\cdot, s^1, s^2\right)$ is integrable for every $\left(s^1, s^2\right) \in S^1 \times S^2$. For every state of nature $\omega \in \Omega$, $u_\omega\left(s^1, s^2\right) = u\left(\omega, s^1, s^2\right)$ represents the payoff received by player 1 (and the loss incurred by player 2) when each player i chooses to play s^i. We assume that each u_ω is a Lipschitz function with constant $K(\omega)$, that is,

$$\left| u_\omega\left(s^1, s^2\right) - u_\omega\left(t^1, t^2\right) \right| \leq K(\omega)\left(\left\| s^1 - t^1 \right\| + \left\| s^2 - t^2 \right\| \right). \tag{1}$$

We also assume that the function $K(\cdot)$ is $\mathcal{F}$-measurable, and that there exists $q > 1$ such that it is q-integrable[9]:

$$\int_\Omega (K(\omega))^q \, d\mu(\omega) < \infty. \tag{2}$$

The probability space $(\Omega, \mathcal{F}, \mu)$, information endowments $\mathcal{F}^1$ and $\mathcal{F}^2$, strategy sets S^1 and S^2, and the payoff function u fully describe a zero-sum Bayesian game. To concentrate on the effects of changes in information endowments, we keep all the attributes of the game fixed, with the exception of $\mathcal{F}^1$ and $\mathcal{F}^2$ that are variable. Thus, we denote the game by $G(\mathcal{F}^1, \mathcal{F}^2)$, to emphasize its changeable characteristics.

A Bayesian strategy of player i is an $\mathcal{F}^i$-measurable function $x^i : \Omega \to S^i$. The set of all Bayesian strategies of player i will be denoted by $X^i\left(\mathcal{F}^i\right)$.

For $p \geq 1$, denote by $L_p^n(\Omega, \mathcal{F}, \mu)$ the Banach space of all $\mathcal{F}$-measurable functions[10] $x : \Omega \to R^n$ such that

$$\|x\|_p \equiv \left(\int_\Omega \|x(\omega)\|^p \, d\mu(\omega) \right)^{\frac{1}{p}} < \infty \tag{3}$$

(recall that $\|\cdot\|$ stands for the Euclidean norm on R^n). We will confine most of our attention to a particular $p > 1$, given by $p = \frac{q}{q-1}$ for q used in (2). The weak topology on $L_p^n(\Omega, \mathcal{F}, \mu)$ is the one in which the linear functional $\varphi_y(x) \equiv \int_\Omega x(\omega) \cdot y(\omega) \, d\mu(\omega)$ is continuous for any given $y \in L_q^n(\Omega, \mathcal{F}, \mu)$. Note that $X^i\left(\mathcal{F}^i\right)$ is a weakly closed subset of the unit ball in $L_p^{n_i}\left(\Omega, \mathcal{F}^i, \mu\right)$ (metrizable and compact in the weak topology). In fact, since the limit of every $L_p^{n_i}(\Omega, \mathcal{F}, \mu)$-weakly converging sequence in $X^i(\mathcal{F}_i)$ is $\mathcal{F}_i$-measurable by Lemma 1 in the Appendix, $X^i(\mathcal{F}_i)$ is also a weakly compact subset of the unit ball in $L_p^{n_i}(\Omega, \mathcal{F}, \mu)$.

[8] The measurability is with respect to $\mathcal{F}$ in the first coordinate, and with respect to the Borel σ-fields in the second and third coordinates.

[9] This condition is relaxed in Theorem 2.

[10] Or, to be precise, their equivalence classes, where any two functions which are equal μ-almost everywhere are identified.

The expected payoff of player 1 (and the expected loss of player 2) when $x^i \in X^i\left(\mathcal{F}^i\right)$ is chosen by i is[11]

$$U(x^1, x^2) \equiv E\left(u_{\cdot}\left(x^1(\cdot), x^2(\cdot)\right)\right) = \int_{\Omega} u_{\omega}\left(x^1(w), x^2(w)\right) d\mu(\omega).$$

This also defines U for all $(x^1, x^2) \in X^1(\mathcal{F}) \times X^2(\mathcal{F})$.

If $\min_{x^2 \in X^2(\mathcal{F}^2)} \max_{x^1 \in X^1(\mathcal{F}^1)} U(x^1, x^2)$ and $\max_{x^1 \in X^1(\mathcal{F}^1)} \min_{x^2 \in X^2(\mathcal{F}^2)} U(x^1, x^2)$ are well defined, and

$$\min_{x^2 \in X^2(\mathcal{F}^2)} \max_{x^1 \in X^1(\mathcal{F}^1)} U(x^1, x^2) = \max_{x^1 \in X^1(\mathcal{F}^1)} \min_{x^2 \in X^2(\mathcal{F}^2)} U(x^1, x^2), \tag{4}$$

then the common value $v = v(\mathcal{F}^1, \mathcal{F}^2)$ of the two expressions in (4) is called the value of the zero-sum Bayesian game $G(\mathcal{F}^1, \mathcal{F}^2)$. Note that v is the value of $G(\mathcal{F}^1, \mathcal{F}^2)$ if and only if there exists a pair of Bayesian strategies $(x^1, x^2) \in X^1(\mathcal{F}_1) \times X^2(\mathcal{F}_2)$ such that for every $(y^1, y^2) \in X^1(\mathcal{F}_1) \times X^2(\mathcal{F}_2)$

$$U(x^1, y^2) \geq U(x^1, x^2) = v \geq U(y^1, x^2). \tag{5}$$

Strategy x^i is called optimal for player i. Any pair (x^1, x^2) of optimal strategies satisfies (5).

The value exists under quite general conditions on the expected payoff function U in the game. We shall assume that U is weakly continuous[12] on $X^1(\mathcal{F}) \times X^2(\mathcal{F})$ separately in every variable, and that it is quasi-concave in x^1 and quasi-convex in x^2. This implies existence of the value by Sion (1958) theorem.

The most prevalent form of a payoff function that gives rise to such U is the usual matrix game. In a matrix game,

$$u_{\omega}\left(s^1, s^2\right) = s^1 A(\omega) s^2, \tag{6}$$

where $A(\omega)$ is an $n_1 \times n_2$ matrix, with $A(\omega)_{i,j}$ being the payoff of player 1 when he chooses pure strategy i and 2 – pure strategy j. Accordingly, s^1, s^2 should be thought of as mixed strategies of players 1 and 2, with each S^i being the n_i-dimensional simplex. Weak continuity in each variable of the corresponding U, as well as condition (1), are guaranteed if, for instance, $a(\omega) = \max_{i,j} |A_{i,j}(\omega)|$ is q-integrable.

Finally, we define convergence of players' information endowments by means of the following pseudo-metric (introduced in Boylan (1971)) on the family $\mathcal{F}^*$ of σ-subfields of $\mathcal{F}$:

$$d(\mathcal{F}_1, \mathcal{F}_2) = \sup_{A \in \mathcal{F}_1} \inf_{B \in \mathcal{F}_2} \mu(A \triangle B) + \sup_{B \in \mathcal{F}_2} \inf_{A \in \mathcal{F}_1} \mu(A \triangle B),$$

[11] The integral below is well defined, due to integrability of each $u\left(\cdot, s^1, s^2\right)$, assumption (1), and integrability of $K(\cdot)$ (which follows from its q-integrability).

[12] Since information endowments $\mathcal{F}_1$ and $\mathcal{F}_2$ of the players may vary from game to game (while the payoff function is fixed), the weak continuity of U is assumed on the set $X^1(\mathcal{F}) \times X^2(\mathcal{F})$, and not on its proper subset of players' strategy profiles $X^1(\mathcal{F}_1) \times X^2(\mathcal{F}_2)$ in the game $G(\mathcal{F}_1, \mathcal{F}_2)$. Clearly, weak continuity of U on $X^1(\mathcal{F}) \times X^2(\mathcal{F})$ induces its weak continuity on each $X^1(\mathcal{F}_1) \times X^2(\mathcal{F}_2)$.

where $A \triangle B = (A \setminus B) \cup (B \setminus A)$ is the "symmetric difference" of A and B. If $x^i \in X^i(\mathcal{F})$ and $\mathcal{F}' \in \mathcal{F}^*$, denote by $E(x^i \mid \mathcal{F}') \in X^i(\mathcal{F}')$ the conditional expectation of x^i with respect to the field $\mathcal{F}'$. If $n_i = 1$ (that is, if $S^i \subset [-1, 1]$), it is known – see, e.g., Van Zandt (1993)[13] – that for any two $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{F}^*$,

$$\left\| E(x^i \mid \mathcal{F}_1) - E(x^i \mid \mathcal{F}_2) \right\|_1 \leq 16 d(\mathcal{F}_1, \mathcal{F}_2).$$

When $n_i > 1$,

$$\begin{aligned}
\left\| E(x^i \mid \mathcal{F}_1) - E(x^i \mid \mathcal{F}_2) \right\|_1 &= \int_\Omega \left\| E(x^i \mid \mathcal{F}_1) - E(x^i \mid \mathcal{F}_2) \right\| d\mu(\omega) \\
&\leq \int_\Omega \sqrt{n_i} \sum_{j=1}^{n_i} \left| E(x^i_j \mid \mathcal{F}_1) - E(x^i_j \mid \mathcal{F}_2) \right| d\mu(\omega) \\
&\leq \sqrt{n_i} \sum_{j=1}^{n_i} \left\| E(x^i_j \mid \mathcal{F}_1) - E(x^i_j \mid \mathcal{F}_2) \right\|_1 \\
&\leq 16 n_i^{\frac{3}{2}} d(\mathcal{F}_1, \mathcal{F}_2).
\end{aligned}$$

Consequently,

$$\left\| E(x^i \mid \mathcal{F}_1) - E(x^i \mid \mathcal{F}_2) \right\|_1 \leq 16 n_i^{\frac{3}{2}} d(\mathcal{F}_1, \mathcal{F}_2). \tag{7}$$

3 Results

Given two pairs of fields in $\mathcal{F}^*$, $(\mathcal{F}^1_1, \mathcal{F}^2_1)$ and $(\mathcal{F}^1_2, \mathcal{F}^2_2)$ (where $\mathcal{F}^i_j$ is the information endowment of player $i = 1, 2$ in pair $j = 1, 2$), the distance between them will be measured by the following pseudo-metric:

$$\bar{d}\left((\mathcal{F}^1_1, \mathcal{F}^2_1), (\mathcal{F}^1_2, \mathcal{F}^2_2)\right) \equiv \max[d(\mathcal{F}^1_1, \mathcal{F}^1_2), d(\mathcal{F}^2_1, \mathcal{F}^2_2)].$$

Theorem 1. The value $v(\mathcal{F}^1, \mathcal{F}^2)$ is a uniformly continuous function of $(\mathcal{F}^1, \mathcal{F}^2) \in \mathcal{F}^* \times \mathcal{F}^*$, with respect to the pseudo-metric $\bar{d}$. Moreover, for any two $(\mathcal{F}^1_1, \mathcal{F}^2_1), (\mathcal{F}^1_2, \mathcal{F}^2_2) \in \mathcal{F}^* \times \mathcal{F}^*$,

$$\left| v(\mathcal{F}^1_1, \mathcal{F}^2_1) - v(\mathcal{F}^1_2, \mathcal{F}^2_2) \right| \leq C \left[\bar{d}\left((\mathcal{F}^1_1, \mathcal{F}^2_1), (\mathcal{F}^1_2, \mathcal{F}^2_2)\right) \right]^{\frac{1}{p}}, \tag{8}$$

[13] Van Zandt (1993) quotes Rogge (1974) and Landers and Rogge (1986), where it is shown that $\|E(f \mid \mathcal{F}_1) - E(f \mid \mathcal{F}_2)\|_1 \leq 8d(\mathcal{F}_1, \mathcal{F}_2)$ for all $\mathcal{F}$-measurable functions f with values in $[0, 1]$.

where $C > 0$ is a constant given by

$$C \equiv 4\left(4 \max\left(n_1, n_2\right)\right)^{\frac{3}{2p}} \|K\|_q . \tag{9}$$

Proof. We will establish inequality (8), which obviously implies the first part of the theorem. For any two given $(\xi_1^1, \xi_1^2), (\xi_2^1, \xi_2^2) \in \xi^* \times \xi^*$, let $x^1 \in X^1\left(\xi_1^1\right)$ be an optimal strategy of player 1 in the game $G(\xi_1^1, \xi_1^2)$, and pick $y^2 \in X^2\left(\xi_2^2\right)$. Now denote $x_2^1 \equiv E(x^1 \mid \xi_2^1) \in X^1\left(\xi_2^1\right)$ and $y_1^2 \equiv E(y^2 \mid \xi_1^2) \in X^2\left(\xi_1^2\right)$. The optimality of x^1 in $G(\xi_1^1, \xi_1^2)$ implies

$$U(x^1, y_1^2) \geq v(\xi_1^1, \xi_1^2). \tag{10}$$

Note that
$$\left|U(x^1, y_1^2) - U\left(x_2^1, y^2\right)\right|$$
(by (1))
$$\leq \int_\Omega K(\omega)\left|x^1(\omega) - x_2^1(\omega)\right| d\mu(\omega) + \int_\Omega K(\omega)\left|y_1^2(\omega) - y^2(\omega)\right| d\mu(\omega)$$
(by the Hölder inequality)
$$\leq \|K\|_q \left(\left\|x^1 - x_2^1\right\|_p + \left\|y_1^2 - y^2\right\|_p\right)$$
(since $\left|x^1(\omega) - x_2^1(\omega)\right|, \left|y_1^2(\omega) - y^2(\omega)\right| \leq 2$ for μ-almost every $\omega \in \Omega$)
$$\leq 2^{\frac{p-1}{p}} \|K\|_q \left(\left(\int_\Omega \left|x^1(\omega) - x_2^1(\omega)\right| d\mu(\omega)\right)^{\frac{1}{p}} + \left(\int_\Omega \left|y_1^2(\omega) - y^2(\omega)\right| d\mu(\omega)\right)^{\frac{1}{p}}\right)$$
$$= 2^{\frac{p-1}{p}} \|K\|_q \left(\left\|x^1 - x_2^1\right\|_1^{\frac{1}{p}} + \left\|y_1^2 - y^2\right\|_1^{\frac{1}{p}}\right)$$
$$= 2^{\frac{p-1}{p}} \|K\|_q \left(\left\|E(x^1 \mid \xi_1^1) - E(x^1 \mid \xi_2^1)\right\|_1^{\frac{1}{p}} + \left\|E(y^2 \mid \xi_1^2) - E(y^2 \mid \xi_2^2)\right\|_1^{\frac{1}{p}}\right)$$
(by (7))
$$\leq 2^{\frac{p-1}{p}} \left(16 \max\left(n_1^{\frac{3}{2}}, n_2^{\frac{3}{2}}\right)\right)^{\frac{1}{p}} \|K\|_q \left(\left[d(\xi_1^1, \xi_2^1)\right]^{\frac{1}{p}} + \left[d(\xi_1^2, \xi_2^2)\right]^{\frac{1}{p}}\right).$$
$$\leq 4\left(8 \max\left(n_1^{\frac{3}{2}}, n_2^{\frac{3}{2}}\right)\right)^{\frac{1}{p}} \|K\|_q \left[\bar{d}\left((\xi_1^1, \xi_1^2), (\xi_2^1, \xi_2^2)\right)\right]^{\frac{1}{p}}.$$
$$= 4\left(4 \max\left(n_1, n_2\right)\right)^{\frac{3}{2p}} \|K\|_q \left[\bar{d}\left((\xi_1^1, \xi_1^2), (\xi_2^1, \xi_2^2)\right)\right]^{\frac{1}{p}}.$$
To summarize, we have shown that
$$\left|U(x^1, y_1^2) - U\left(x_2^1, y^2\right)\right| \leq C\left[\bar{d}\left((\xi_1^1, \xi_1^2), (\xi_2^1, \xi_2^2)\right)\right]^{\frac{1}{p}}. \tag{11}$$

Together with (10), (11) implies that

$$U\left(x_2^1, y^2\right) \ge v(\zeta_1^1, \zeta_1^2) - C\left[\bar{d}\left((\zeta_1^1, \zeta_1^2), (\zeta_2^1, \zeta_2^2)\right)\right]^{\frac{1}{p}}.$$

This holds for every $y^2 \in X^2\left(\zeta_2^2\right)$, and hence it follows that

$$v(\zeta_2^1, \zeta_2^2) = \max_{y^1 \in X^1(\zeta_2^1)} \min_{y^2 \in X^2(\zeta_2^2)} U(y^1, y^2) \tag{12}$$

$$\ge \min_{y^2 \in X^2(\zeta_2^2)} U(x_2^1, y^2) \ge v(\zeta_1^1, \zeta_1^2) - C\left[\bar{d}\left((\zeta_1^1, \zeta_1^2), (\zeta_2^1, \zeta_2^2)\right)\right]^{\frac{1}{p}}. \tag{13}$$

Using similar arguments (when we start from an optimal strategy $x^2 \in X^2\left(\zeta_1^2\right)$ of player 2 in the game $G(\zeta_1^1, \zeta_1^2)$) we can show that, for $x_2^2 = E(x^2 \mid \zeta_2^2) \in X^2\left(\zeta_2^2\right)$, the following inequality

$$U\left(y^1, x_2^2\right) \le v(\zeta_1^1, \zeta_1^2) + C\left[\bar{d}\left((\zeta_1^1, \zeta_1^2), (\zeta_2^1, \zeta_2^2)\right)\right]^{\frac{1}{p}}$$

holds for every $y^1 \in X^1\left(\zeta_2^1\right)$. This leads to

$$v(\zeta_2^1, \zeta_2^2) = \min_{y^2 \in X^2(\zeta_2^2)} \max_{y^1 \in X^1(\zeta_2^1)} U(y^1, y^2) \tag{14}$$

$$\le \max_{y^1 \in X^1(\zeta_2^1)} U(y^1, x_2^2) \le v(\zeta_1^1, \zeta_1^2) + C\left[\bar{d}\left((\zeta_1^1, \zeta_1^2), (\zeta_2^1, \zeta_2^2)\right)\right]^{\frac{1}{p}}. \tag{15}$$

The combination of (12)-(13) and (14)-(15) now implies (8). ∎

The continuity of the value as a function of (ζ^1, ζ^2) is, of course, an immediate implication of Theorem 1:

Corollary 1. Suppose that $\left\{\zeta_k^i\right\}_{k=1}^{\infty} \subset F^*$ is a sequence such that $\lim_{k \to \infty} \zeta_k^i = \zeta^i$ in the Boylan pseudo-metric, for $i = 1,2$. Then $\lim_{k \to \infty} v(\zeta_k^1, \zeta_k^2) = v(\zeta^1, \zeta^2)$.

If $K(\cdot)$ is a bounded function, it is obvious that (2) holds for every $q > 1$, and thus $p = \frac{q}{q-1}$ can be chosen to be arbitrarily close to 1. The constant $C = C(p)$, defined in (9), converges to the limit

$$32 \max\left(n_1^{\frac{3}{2}}, n_2^{\frac{3}{2}}\right) \|K\|_{\infty}$$

when p approaches 1 ($\|K\|_{\infty}$ stands for the essential supremum of K). Inequality (8) of Theorem 1 thus provides us with the following corollary:

Corollary 2. If $K(\cdot)$ is a bounded function, the value $v(\zeta^1, \zeta^2)$ is a Lipschitz function of $(\zeta^1, \zeta^2) \in \zeta^* \times \zeta^*$, with respect to the pseudo-metric $\bar{d}$.

It is natural to ask whether the value is continuous when $K(\cdot)$ is only integrable (that is, in $L_1^1(\Omega,\mathcal{F},\mu)$), and not q-integrable for some $q>1$ as is assumed in (2). Our next theorem shows that the continuity holds even under this more general assumption. However, it does not follow from Theorem 1 (since we do not have uniform continuity in this case) and has to be established directly (using similar techniques). The continuity of U (separately in each variable) is now assumed with respect to the $L_p^{n_i}(\Omega,\mathcal{F},\mu)$-weak topology on each coordinate,[14] for some $p>1$.

Theorem 2. The statement of Corollary 1 remains valid even if $K(\cdot)$ is only integrable. That is, if $\left\{\mathcal{F}_k^i\right\}_{k=1}^{\infty}\subset F^*$ is a sequence such that $\lim_{k\to\infty}\mathcal{F}_k^i=\mathcal{F}^i$ in the Boylan pseudo-metric, for $i=1,2$, then $\lim_{k\to\infty}v(\mathcal{F}_k^1,\mathcal{F}_k^2)=v(\mathcal{F}^1,\mathcal{F}^2)$.

Proof. Suppose by the way of contradiction that the (bounded) sequence $\left\{v(\mathcal{F}_k^1,\mathcal{F}_k^2)\right\}_{k=1}^{\infty}$ has a subsequence that converges to $v^0\neq v(\mathcal{F}^1,\mathcal{F}^2)$; without loss of generality, let this subsequence be $\left\{v(\mathcal{F}_k^1,\mathcal{F}_k^2)\right\}_{k=1}^{\infty}$ itself. Also let x_k^1 be an optimal strategy of player 1 in the game $G(\mathcal{F}_k^1,\mathcal{F}_k^2)$, for every $k=1,2,3,\ldots$ As was mentioned, $X^1(\mathcal{F})$ is a weakly compact subset of the (metrizable) unit ball in $L_p^{n_1}(\Omega,\mathcal{F},\mu)$, and therefore there is a subsequence $\left\{x_{k_l}^1\right\}_{l=1}^{\infty}$ which converges weakly to some $x^1\in X^1(\mathcal{F})$. By Lemma 1 in the Appendix x^1 is $\mathcal{F}^1$-measurable, which implies that $x^1\in X^1\left(\mathcal{F}^1\right)$.

Now fix $y^2\in X^2\left(\mathcal{F}^2\right)$, and, for every $k=1,2,3,\ldots$, let $y_{k_l}^2\equiv E(y^2\mid\mathcal{F}_{k_l}^2)\in X^2\left(\mathcal{F}_{k_l}^2\right)$. Since $x_{k_l}^1$ is an optimal strategy of 1 in $G(\mathcal{F}_{k_l}^1,\mathcal{F}_{k_l}^2)$,

$$U(x_{k_l}^1,y_{k_l}^2)\geq v(\mathcal{F}_{k_l}^1,\mathcal{F}_{k_l}^2). \tag{16}$$

Since $y_{k_l}^2\to_{l\to\infty}y^2$ in $L_1^{n_2}(\Omega,\mathcal{F},\mu)$ by (7), there is a subsequence of $\left\{y_{k_l}^2\right\}_{l=1}^{\infty}$ that converges pointwise to y^2 μ-almost everywhere; without loss of generality, the sequence itself converges pointwise. Note that

$$\left|U(x_{k_l}^1,y_{k_l}^2)-U\left(x^1,y^2\right)\right|\leq\left|U(x_{k_l}^1,y_{k_l}^2)-U\left(x_{k_l}^1,y^2\right)\right|+\left|U(x_{k_l}^1,y^2)-U\left(x^1,y^2\right)\right|$$

(by (1))

$$\leq\int_\Omega K(\omega)\left\|y_{k_l}^2(\omega)-y^2(\omega)\right\|d\mu(\omega)+\left|U(x_{k_l}^1,y^2)-U\left(x^1,y^2\right)\right|.$$

The first term in the above expression converges to zero as $l\to\infty$ by the bounded convergence theorem, and the second terms also converges to zero since

[14] This assumption is satisfied quite often. For instance, when the matrix game (as in (6)) is considered, and $\max_{i,j}|A_{i,j}(\omega)|$ is only integrable, the expected payoff function $U(x^1,x^2)=\int_\Omega x^1(w)A(\omega)x^2(w)\,d\mu(\omega)$ is L_p-weakly continuous in each coordinate separately, for every $p>1$. This is so because strategies of both players are uniformly bounded, and $A(\omega)$ can be approximated in the L_1-norm by bounded matrices.

U is weakly continuous in each variable separately. Thus, $\lim_{l\to\infty} U(x^1_{k_l}, y^2_{k_l}) = U\left(x^1, y^2\right)$, and together with (16) this implies

$$U(x^1,y^2) \geq \lim_{l\to\infty} v(\zeta^1_{k_l}, \zeta^2_{k_l}) = v^0; \tag{17}$$

this inequality holds for every $y^2 \in X^2\left(\zeta^2\right)$. Thus,

$$v(\zeta^1,\zeta^2) = \max_{y^1\in X^1(\zeta^1)} \min_{y^2\in X^2(\zeta^2)} U(y^1,y^2) \tag{18}$$

$$\geq \min_{y^2\in X^2(\zeta^2)} U(x^1,y^2) \geq v^0. \tag{19}$$

Using similar arguments (when we start from finding a limit point x^2 of a sequence $\left\{x^2_k\right\}_{k=1}^{\infty}$ of optimal strategies of player 2 in games $G(\zeta^1_k, \zeta^2_k)$) we can show that

$$U(y^1,x^2) \leq \lim_{l\to\infty} v(\zeta^1_{k_l}, \zeta^2_{k_l}) = v^0 \tag{20}$$

for every $y^1 \in X^1\left(\zeta^1\right)$. This leads to

$$v(\zeta^1,\zeta^2) = \min_{y^2\in X^2(\zeta^2)} \max_{y^1\in X^1(\zeta^1)} U(y^1,y^2) \tag{21}$$

$$\leq \max_{y^1\in X^1(\zeta^1)} U(y^1,x^2) \leq v^0. \tag{22}$$

The combination of (18)-(19) and (21)-(22) now imply $v^0 = v(\zeta^1,\zeta^2)$, contradicting the initial assumption. This contradiction establishes $\lim_{k\to\infty} v(\zeta^1_k,\zeta^2_k) = v(\zeta^1,\zeta^2)$. ∎

The following theorem follows quite easily from the proof of Theorem 2.

Theorem 3. The optimal strategy correspondence is upper semi-continuous for both players. That is, if $\left\{\zeta^i_k\right\}_{k=1}^{\infty} \subset F^*$ are such that $\lim_{k\to\infty} \zeta^i_k = \zeta^i$ in the Boylan pseudo-metric for every $i = 1,2$, and $\left\{(x^1_k,x^2_k)\right\}_{k=1}^{\infty}$ is such that (x^1_k,x^2_k) is a pair of optimal strategies in $G(\zeta^1_k,\zeta^2_k)$ and $\lim_{k\to\infty}(x^1_k,x^2_k) = (x^1,x^2)$ weakly in both coordinates, then (x^1,x^2) is a pair of optimal strategies in $G(\zeta^1,\zeta^2)$.

Proof. As was said, this uses the proof of Theorem 2. The first part of that proof (leading to (17)) can be utilized to show that $U(x^1,y^2) \geq \lim_{k\to\infty} v(\zeta^1_k,\zeta^2_k)$ for every $y^2 \in X^2\left(\zeta^2\right)$. However, by Theorem 2, $\lim_{k\to\infty} v(\zeta^1_k,\zeta^2_k) = v(\zeta^1,\zeta^2)$, and so x^1 is indeed an optimal strategy of 1 in $G(\zeta^1,\zeta^2)$. Similarly, the second part of the proof can be used to show that x^2 is an optimal strategy of 2. ∎

Remark 1. The optimal strategy correspondence is not lower semi-continuous in general.[15] That is, it may be the case that $\lim_{k\to\infty}\Sigma^i_k = \Sigma^i$ in the Boylan pseudo-metric and (x^1,x^2) is pair of optimal strategies in $G(\Sigma^1,\Sigma^2)$, but there is no sequence $\left((x^1_k,x^2_k)\right)_{k=1}^{\infty}$ of optimal strategies in $G(\Sigma^1_k,\Sigma^2_k)$ that converges to (x^1,x^2) weakly in both coordinates. Indeed, consider the situation where $\Omega = [-1,1]$, Σ is the σ-field of Borel sets in Ω, μ is the normalized Lebesgue measure on Ω, $S^1 = [0,1]$, $S^2 = \{0\}$, and $u\left(\omega,s^1,s^2\right) = \omega s^1$. Now let $\Sigma^1_k = \Sigma^2_k$ be the σ-field which is generated by all Borel subsets of $\left[-1,-1+\frac{1}{k}\right]$ and the set (an "atom") $(-1+\frac{1}{k},1]$, for all $k = 1,2,3,...$, and $\Sigma^1 = \Sigma^2 = \{\emptyset,\Omega\}$. Then clearly $\lim_{k\to\infty}\Sigma^i_k = \Sigma^i$ for $i = 1,2$. However, consider a pair $(x^1,x^2)\equiv(0,0)$ of optimal strategies in the game $G(\Sigma^1,\Sigma^2)$. Since any optimal strategy x^1_k of 1 in the game $G(\Sigma^1_k,\Sigma^2_k)$ satisfies $x^1_k(\omega) = 1$ for every $\omega \in (-1+\frac{1}{k},1]$, there exists no sequence of optimal strategies of 1 in $\left(G(\Sigma^1_k,\Sigma^2_k)\right)_{k=1}^{\infty}$ that converges to x^1. ■

Remark 2. Given Theorem 3 on upper semi-continuity of the optimal strategy correspondence for zero-sum games, it is natural to ask whether its counterpart for non-zero-sum games, the Bayesian Nash equilibrium (NE) correspondence, is upper semi-continuous in the same way. (It is clearly not lower semi-continuous, since even the optimal strategy correspondence in zero-sum games is not.) The answer to the above question is negative. The discontinuous behavior of the NE correspondence in our setting is due to a markedly weak requirement on convergence of NE strategies: they only need to converge in the weak topology. While this weak mode of converges suffices to obtain optimal strategies in the limit for zero-sum games (and adds strength to Theorem 3), the situation is different for NE in non-zero-sum games. The pitfall that the weak topology brings with it is the typical discontinuity of the expected payoff function in all strategies simultaneously[16]; in zero-sum games continuity in both variables separately did the job, but not so in general games.

To construct an example of discontinuous NE, consider a non-zero-sum Bayesian game with two players, $i = 1,2$, in which $\Omega = S^1 = S^2 = [0,1]$ (each player has two pure strategies, 0 and 1, and the open interval $(0,1)$ constitutes the set of completely mixed strategies), Σ is the σ-field of Borel sets in Ω, and μ is the normalized Lebesgue measure on Ω. Both players play the same coordination game in all states of nature: the matrix which defines players'

[15] As was already mentioned, Monderer and Samet (1996) show that the Nash equilibrium (NE) correspondence is not lower semi-continuous. The example that we present here shows the lack of lower semi-continuity of NE in zero-sum games (and even in one-person decision problems).

[16] In the fixed-types set-up of Milgrom and Weber (1986), the expected payoff functions were simultaneously continuous in all players' strategies; in fact, this was shown to imply upper semi-continuity of the NE correspondence. However, this continuity of payoff funcions was partly the result of a sufficient spread of the common prior distribution of players' types (that the assumptions of Milgrom and Weber imply in the case where each type is, say, an interval). This feature would make their analysis inapplicable in the complete information case, which is precsely what we consider in our example in the next paragraph.

payoffs for pure strategy profiles is

$$\begin{array}{c|cc} & s^2=0 & s^2=1 \\ s^1=0 & (2,2) & (0,0) \\ s^1=1 & (0,0) & (1,1) \end{array}.$$

Thus, $u^1(\omega,s^1,s^2)=u^2(\omega,s^1,s^2)\equiv s^1s^2+2(1-s^1)(1-s^2)$. Also let $\mathcal{F}_k^1=\mathcal{F}_k^2=\mathcal{F}$.

For every $k=1,2,3,\ldots$ partition $\Omega=[0,1]$ into 2^k consecutive intervals of equal length, $I_1(k),\ldots,I_{2^k}(k)$. Now consider a sequence $\{(x_k^1,x_k^2)\}_{k=1}^\infty$ of symmetric NE strategies in $G(\mathcal{F},\mathcal{F})$, given by

$$x_k^1(\omega)=x_k^2(\omega)\equiv x_k(\omega)=\begin{cases}1, & \text{if } \omega\in I_n(k) \text{ for even } n;\\ 0, & \text{if } \omega\in I_n(k) \text{ for odd } n.\end{cases}$$

It is known that $\{x_k\}_{k=1}^\infty$ converges weakly[17,18] to the constant function $x\equiv\frac{1}{2}$. However, (x,x) is clearly not an NE in $G(\mathcal{F},\mathcal{F})$. $\blacksquare$

4 Appendix

Lemma 1. Let $\{\mathcal{F}_k\}_{k=1}^\infty\subset\mathcal{F}^*$ be a sequence such that $\lim_{k\to\infty}\mathcal{F}_k=\mathcal{F}_0$ in the Boylan pseudo-metric. If $\{x_k\}_{k=1}^\infty\subset\prod_{k=1}^\infty X^i(\mathcal{F}_k)$ is a sequence of functions that converges weakly to $x\in X^i(\mathcal{F})$, then x is $\mathcal{F}_0$-measurable (that is, $x\in X^i(\mathcal{F}_0)$).

Proof. Without loss of generality, we assume that

$$\sum_{k=1}^{\infty} d(\mathcal{F}_k,\mathcal{F}_0)<\infty \tag{23}$$

(otherwise consider instead some subsequence $\{\mathcal{F}_{k_l}\}_{l=1}^\infty$ with $\sum_{l=1}^\infty d(\mathcal{F}_{k_l},\mathcal{F}_0)<\infty$). For every k denote by $\mathcal{G}_k$ the σ-field $\bigvee_{n=k}^\infty\mathcal{F}_k$, that is, the minimal σ-subfield of $\mathcal{F}$ which contains each one of $\{\mathcal{F}_n\}_{n=k}^\infty$. It follows from (23) by Corollary 2 of Van Zandt (1993) that $\lim_{k\to\infty}\mathcal{G}_k=\mathcal{F}_0$.

Since $\{x_k\}_{k=1}^\infty$ converges weakly to x, by Banach-Saks theorem there exists a sequence $\{\overline{x}_k\}_{k=1}^\infty$ that converges to x strongly (that is, in the $\|\cdot\|_p$ norm), and each $\overline{x}_k$ is a convex combination of $\{x_n\}_{n=k}^\infty$. Thus, $\overline{x}_k\in X^i(\mathcal{G}_k)$ for every $k=1,2,3,\ldots$ By Lemma 1 in Einy et al (2003), the strong limit of $\{\overline{x}_k\}_{k=1}^\infty$ is measurable with respect to $\lim_{k\to\infty}\mathcal{G}_k=\mathcal{F}_0$. We conclude that $x\in X^i(\mathcal{F}_0)$. $\blacksquare$

[17] Assuming that $p=q=2$.

[18] Indeed, $\{2(x_k-\frac{1}{2})\}_{k=1}^\infty$ is the sequence of Rademacher functions that converges weakly to zero.

References

1. Allen, B. (1983). "Neighboring Information and Distribution of Agent Characteristics under Uncertainty," Journal of Mathematical Economics 12, pp. 63-101.

2. Allen, B. and N.C. Yannelis (2001). "Di¤erential Information Economies: Introduction," Economic Theory 18, pp. 263-274.

3. Boylan, E. (1971) "Equiconvergence of Martingales," Annals of Mathematical Statistics 42, pp. 552-559.

4. Cotter, K.D. (1986). "Similarity of Information and Behavior with Pointwise Convergence Topology," Journal of Mathematical Economics 15, pp. 25-38.

5. Cotter, K.D. (1987). "Convergence of Information, Random Variable and Noise," Journal of Mathematical Economics 16, pp. 39-51.

6. Einy E., O. Haimanko, D. Moreno and B. Shitovitz (2003). "On the Continuity of Equilibrium and Core Correspondences in Economies with Di¤erential Information," DP # 02-03 of Monaster Center for Economic Research, Ben-Gurion University, and DP # 03-27 of Department of Economics, Universidad Carlos III de Madrid.

7. Kajii, A. and S. Morris (1998). "Payo¤ Continuity in Incomplete Information Games," Journal of Economic Theory 82, pp. 267-276.

8. Landers, D. and L. Rogge (1986). "An Inequality for the Hausdor¤-metric of σ-…elds," The Annals of Probability 14, pp. 724-730.

9. Milgrom, P. and R.J. Weber (1985). "Distributional Strategies for Games with Incomplete Information," Mathematics of Operations Research 10, pp. 619-632.

10. Monderer, D. and D. Samet (1996). "Proximity of Information in Games with Common Beliefs," Mathematics of Operations Research 21, pp. 707-725.

11. Rogge, L. (1974). "Uniform Inequalities for Conditional Expectations," The Annals of Probability 2, pp. 486-489.

12. Sion, M. (1958). "On General Minimax Theorem," Paci…c Journal of Mathematics 8, pp. 171-176.

13. Stinchcombe, M. (1990). "Bayesian Information Topologies," Journal of Mathematical Economics 19, pp. 233-253.

14. Van Zandt, T. (1993). "The Hausdor¤ Metric of σ-…elds and the Value of Information," The Annals of Probability 21, pp. 161-167.

15. Van Zandt, T. (2002). "Information, Measurability, and Continuous Behavior," Journal of Mathematical Economics 38, pp. 293-309.