Continuity of the value and optimal strategies when common priors change*

Ezra Einy†, Ori Haimanko‡, Biligbaatar Tumendemberel§

March 2010

Abstract

We show that the value of a zero-sum Bayesian game is a Lipschitz continuous function of the players' common prior belief, with respect to the total variation metric on beliefs. This is unlike the case of general Bayesian games, where lower semi-continuity of Bayesian equilibrium (BE) payoffs rests on the "almost uniform" convergence of *conditional* beliefs. We also show upper semi-continuity (USC) and approximate lower semi-continuity (ALSC) of the optimal strategy correspondence, and discuss ALSC of the BE correspondence in the context of zero-sum games. In particular, the interim BE correspondence is shown to be ALSC for some classes of information structures with highly non-uniform convergence of beliefs, that would not give rise to ALSC of BE in non-zero-sum games.

JEL Classification Number: C72.

Keywords: Zero-Sum Bayesian Games, Common Prior, Value, Optimal Strategies, Interim, Ex ante, Bayesian Equilibrium, Upper Semi-Continuity, Lower Approximate Semi-Continuity.

---

*Previous version: July 2009

†Department of Economics, Ben-Gurion University of the Negev, Beer Sheva 84105, Israel; and Graduate School of Economics, Hitotsubashi University, Naka 2-1, Kunitachi, Tokyo 186-8601, Japan. E-mail: einy@bgu.ac.il

‡Department of Economics, Ben-Gurion University of the Negev, Beer Sheva 84105, Israel. E-mail: orih@bgu.ac.il. This author acknowledges the support of JSPS fellowship during the initial stage of this work.

§Graduate School of Economics, Hitotsubashi University, Naka 2-1, Kunitachi, Tokyo 186-8601, Japan. E-mail: em071007@g.hit-u.ac.jp

1 Introduction

Bayesian games describe situations where there is uncertainty about the players' payoffs, and the players may have different private information about the realized state of nature that affects the payoffs. Being a natural framework for modelling numerous real world issues, it has been a subject of extensive investigation in the literature. In particular, the question of continuity of Bayesian equilibria (BE) with respect to changes in players' information endowments received notable attention.

Rubinstein's (1989) example of electronic mail game had demonstrated that seemingly "small" deviations from common knowledge of the payoffs in a game may produce big differences in BE payoffs and strategies. However, these deviations do not lead to *approximate* common knowledge of the payoffs in the sense of Monderer and Samet (1989), who show that approximate common knowledge of the payoffs with high ex-ante probability, guarantees existence of BE similar to the BE in the limit case of commonly known payoffs.

Among other repercussions, these developments reinvigorated the strand of research (to which the present paper also belongs) that studies the effects of small changes in *players' common prior belief*[1] on BE. In an earlier work, Milgrom and Weber (1985) showed *upper semi-continuity (USC)* of the BE correspondence with respect to the common prior, under a very general condition requiring that the common prior be sufficiently "spread-out" on the product of players' types. This condition is satisfied trivially in the important case where each player has at most *countably many* types, which is equivalent to assuming that his private information is given by a *countable* partition of the space of states of nature.[2] In this latter framework, the works of Kajii and Morris (1994, 1998)) and Engl (1995) are particularly noteworthy.

Engl (1995) investigated *(approximate) lower semi-continuity ((A)LSC)* of the BE expected payoff correspondence, under the uniform setwise convergence topology on priors. The ALSC means that for any BE in a game and any $\varepsilon > 0$, there is an ε-equilibrium with close expected payoffs in the same

[1] Even with small changes in the common prior belief, common knowledge of the payoffs may be lost, even in the approximate sense. It was shown by Kajii and Morris (1994); their example is an elaboration on Rubinstein's (1989) electronic mail game, and it is repeated in Example 3 here.

[2] This is the set-up in both Rubinstein (1989) and Monderer and Samet (1989).

game, for any close enough common prior.[3] Engl (1995) showed that the BE expected payoff correspondence is ALSC, assuming that the approximating ε-BE are with respect to *ex-ante* expected utilities. However, if players evaluate the consequences of their strategic choices at the interim stage, following the receipt of private information, they are in fact concerned with their *interim* expected utility, that takes into account their private information and is based on the correspondingly updated prior belief. But while ex-ante and interim BE are the same, this is not true for the approximate, ε-BE, since an ex-ante ε-best response may be hugely suboptimal for some realizations of the player's private information, albeit with small probability. Kajii and Morris (1994, 1998) showed that, if the approximate ε-BE are taken in the interim sense, ALSC of the BE expected payoff and strategy correspondences may fail if priors are converging only setwise. They showed that to obtain ALSC of the interim BE expected payoff correspondence, uniform across bounded games, it is necessary (and sufficient) to additionally assume *almost uniform convergence* of beliefs *conditional* on players' private information (i.e., that the closeness of conditional beliefs becomes approximate common knowledge with high ex-ante probability).

In this work we consider *zero-sum* Bayesian games. These games recently came into spotlight, particularly in the context of characterizing the value-of-information function (see, e.g., Lehrer and Rosenberg (2006, 2007)). We start by showing that the value of a zero-sum game is a *Lipschitz continuous* function of players' common prior belief, with respect to the total variation metric on the set of priors; see Theorem 1. (This metric induces the setwise convergence topology on priors.)

Although being in line with Engl's (1994) result on the ALSC of the ex-ante BE expected payoff correspondence, Theorem 1 implies a previously unnoticed fact. Since pairs of optimal strategies are both interim and ex-ante BE in a zero-sum Bayesian game, and the value ($\equiv$the expected BE payoff) is a continuous function of the common prior, the *interim* BE expected payoff correspondence is in fact LSC (and in particular ALSC) when restricted to zero-sum games. Thus, the assumptions of Kajii and Morris (1998) on the convergence of conditional beliefs, which are necessary for ALSC in the non-zero-sum setting, are not needed in the context of zero-sum Bayesian games.[4]

[3]The stronger notion of *lower semi-continuity (LSC)* requires that the BE expected payoff be approximable by expected payoffs of exact BE in games with close enough common priors.

[4]Bayesian zero-sum games were explicitly mentioned in Kajii and Morris (1997) in the

Theorem 1 bears semblance to another uniform continuity result for the value of zero-sum games, in Einy et al (2008), which was established in a different setting: the common prior of the players was fixed, but their information fields were variable (as in Monderer and Samet (1996)), and the set of fields was endowed with the Boylan pseudo-metric. This result does not imply Theorem 1, however, as the latter deals with variable common priors.

We further show that the *optimal strategy* correspondence is both USC (Proposition 1) and ALSC (Proposition 2) with respect to the total variation metric on priors. Since optimal strategies are both ex-ante and interim BE strategies in zero-sum Bayesian games, Proposition 1 implies that the ex-ante and the interim BE correspondences are USC.[5] However, the notion of ALSC uses ε-optimal strategies to approximate the given optimal strategy, and ε-optimal strategies are defined with respect to the *ex-ante* expected payoffs in the game. Thus Proposition 2 implies that the ex-ante BE correspondence is ALSC in zero-sum Bayesian games, but it sheds no light on the ALSC of the *interim* BE correspondence in these games.

As was mentioned, the interim BE correspondence may not be ALSC with respect to the total variation metric on priors, in cases where conditional beliefs do not converge almost uniformly (see Kajii and Morris (1994)). Our last two results show that, at least in some circumstances, the interim BE correspondence is ALSC in zero-sum Bayesian games without any assumptions on the convergence of conditional beliefs. Proposition 3 identifies one such instance in games where each player has an infinite information partition; the main assumption is that the knowledge of the player's own type allows him to guess the type of the other player while making a bounded error. To contrast this with the non-zero-sum case, we recall in Example 3 a non-zero-sum Bayesian game constructed in Kajii and Morris (1994), with information partitions of the type admitted by Proposition 3, and even with

context of robustness to incomplete information (of the equilibrium payoff in the *complete* information game), under the assumption that each player has a unique optimal strategy; see Proposition 3.2 and the paragraph following it in that paper. However, the aim here is to highlight the LSC of the interim BE expected payoff correspondence for all zero-sum games with *incomplete* information.

[5] The BE correspondence is, in fact, USC for general Bayesian games in our setting, although Proposition 1 is stated only for zero-sum games (being concerned with USC of *optimal strategies*). This is suggested, but not implied, by the general USC result of Milgrom and Weber (1985), which is established in a somewhat different framework.

common knowledge of payoffs in the limit of a converging sequence of common priors, for which the interim BE correspondence is not ALSC. Finally, when at least one of the players has a finite information partition, ALSC of the interim BE correspondence obtains without further assumptions, see Proposition 4.

The paper is organized as follows. The setup is described in section 2 and our results are stated and proved in section 3.

2 Preliminaries

2.1 Zero-Sum Bayesian Games

We consider zero-sum games with two players, $i = 1, 2$. Games are played in an uncertain environment, which affects payoff functions of the players. The underlying uncertainty is described by a probability space $(\Omega, \mathcal{F}, \mu)$, where Ω is a set of states of nature, $\mathcal{F}$ is a σ-field of events in Ω, and μ is a countably additive probability measure on $(\Omega, \mathcal{F})$ that represents the *common prior belief* of the players about the distribution of the realized state of nature. The *information endowment* of player i is given by an (at most) countable and $\mathcal{F}$-measurable partition Π^i of Ω. Given $\omega \in \Omega$, denote by $\Pi^i(\omega)$ the element of the partition Π^i that contains ω. If ω was realized, player i only knows that the realized state of nature belongs to $\Pi^i(\omega)$.

Each player $i = 1, 2$ has a set S^i of *strategies*, which is a convex and compact subset of a Euclidean space[6]. Additionally, there is a measurable[7] real valued *payoff function* $u : \Omega \times S^1 \times S^2 \to \mathbb{R}$. At every state of nature $\omega \in \Omega$, $u(\omega, s^1, s^2)$ is the payoff received by player 1, and $-u(\omega, s^1, s^2)$ is the payoff of player 2, when each player i chooses to play s^i. We assume that, at every $\omega \in \Omega$, each player's payoff is continuous and concave in his own strategy; that is, $u(\omega, \cdot, s^2)$ is continuous and concave for a fixed $s^2 \in S^2$, and $u(\omega, s^1, \cdot)$ is continuous and convex for a fixed $s^1 \in S^1$. We further assume

[6] All our results, with the exception of Proposition 3, use only the fact that S^i is a compact and metrizable subset of a topological vector space. For Proposition 3, the assumption of S^i being a Banach space, not necessarily of finite dimension, would have sufficed. We, however, confine ourselves to the finite-dimension framework, so as to avoid unnecessary generality.

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

that $|u|$ is bounded on $\Omega \times S^1 \times S^2$ by some $M > 0$ (Remark 1 below explains the necessity of this assumption).

The probability space $(\Omega, \mathcal{F}, \mu)$, information endowments Π^1 and Π^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 the common prior, we keep all the attributes of the game fixed henceforth, with the exception of $\mu \in \Delta(\Omega, \mathcal{F}) \equiv$the set of all countably additive probability measures on $(\Omega, \mathcal{F})$. For any $\mu \in \Delta(\Omega, \mathcal{F})$, the associated zero-sum Bayesian game will be denoted by $G(\mu)$.

A *Bayesian strategy* of player i is a Π^i-measurable function $x^i : \Omega \to S^i$, i.e., x^i is constant on any $\pi^i \in \Pi^i$. The set of all Bayesian strategies of player i will be denoted by X^i. Clearly, X^i can be identified with the function space $(S^i)^{\Pi^i}$, which is convex and compact in the product topology, and also metrizable in it since Π^i is at most countable. Given $\mu \in \Delta(\Omega, \mathcal{F})$, the *expected payoff* of player 1 (and the expected loss of player 2) when $x^i \in X^i$ is chosen by $i = 1, 2$ is

$$U_\mu(x^1, x^2) \equiv \int_\Omega u\left(\omega, x^1(w), x^2(w)\right) d\mu(\omega).$$

Remark 1. In order for the expected payoff function U_μ to be well defined for *a given* $\mu \in \Delta(\Omega, \mathcal{F})$, μ-integrability of an $\mathcal{F}$-measurable

$$f(\omega) \equiv \sup_{(s^1,s^2)\in S^1\times S^2} \left|u\left(\omega, s^1, s^2\right)\right|$$

would have sufficed, without the need to assume uniform boundedness of u as we did earlier. However, since our interest lies in changing common priors in the game with a fixed utility function, f needs to be integrable with respect to *all* $\mu \in \Delta(\Omega, \mathcal{F})$. This, in fact, implies the existence of $M = \sup_{\omega\in\Omega} f(\omega) < \infty$.

With our assumptions on u, the expected payoff function U_μ is continuous and concave in $x^1 \in X^1$ for a fixed $x^2 \in X^2$, and continuous and convex in $x^2 \in X^2$ for a fixed $x^1 \in X^1$ (the continuity is implied by the bounded convergence theorem). Thus, Sion minimax theorem (see, e.g., Theorem A.7

in Sorin (2002)) guarantees existence of the *value* $v(\mu)$ in each game $G(\mu)$: the following inequality holds,

$$\min_{x^2 \in X^2} \max_{x^1 \in X^1} U_\mu(x^1, x^2) = \max_{x^1 \in X^1} \min_{x^2 \in X^2} U_\mu(x^1, x^2), \tag{1}$$

and $v(\mu)$ is defined as the common value of the two expressions in (1).

Given $\varepsilon \geq 0$, $\widehat{x}^1 \in X^1$ is called *ε-optimal* for player 1 in $G(\mu)$ if

$$U_\mu(\widehat{x}^1, x^2) \geq v(\mu) - \varepsilon$$

for any $x^2 \in X^2$. Similarly, $\widehat{x}^2 \in X^2$ is called *ε-optimal* for player 2 in $G(\mu)$ if

$$U_\mu(x^1, \widehat{x}^2) \leq v(\mu) + \varepsilon$$

for any $x^1 \in X^1$. If a strategy x^i is 0-optimal for player i, it is called *optimal* for i. The set of ε-optimal strategies of player i in $G(\mu)$ will be denoted by $\mathcal{O}^i_\varepsilon(\mu)$. It is convex and compact. The notation for $\mathcal{O}^i_0(\mu)$, the set of optimal strategies, will be simplified to $\mathcal{O}^i(\mu)$. Since the value exists, $\mathcal{O}^i_\varepsilon(\mu)$ is a non-empty set for any $\varepsilon \geq 0$.

Optimality of a strategy is closely related to the concept of equilibrium. A pair $(\widehat{x}^1, \widehat{x}^2) \in X^1 \times X^2$ is called an *ex-ante Bayesian ε-equilibrium* (henceforth *ε-EBE* for short) if

$$U_\mu(\widehat{x}^1, \widehat{x}^2) \geq U_\mu(x^1, \widehat{x}^2) - \varepsilon \tag{2}$$

for any $x^1 \in X^1$, and

$$U_\mu(\widehat{x}^1, \widehat{x}^2) \leq U_\mu(\widehat{x}^1, x^2) + \varepsilon \tag{3}$$

for any $x^2 \in X^2$. Denote by $EBE_\varepsilon(\mu)$ the set of all ε-EBE in $G(\mu)$, and simplify $EBE_0(\mu)$ to $EBE(\mu)$. If $(\widehat{x}^1, \widehat{x}^2) \in EBE(\mu)$, we will call it an *ex-ante Bayesian equilibrium* (*EBE* for short).

Remark 2. Note that, for every $\varepsilon \geq 0$,

$$\mathcal{O}^1_\varepsilon(\mu) \times \mathcal{O}^2_\varepsilon(\mu) \subset EBE_{2\varepsilon}(\mu)$$

and

$$EBE_\varepsilon(\mu) \subset \mathcal{O}^1_{2\varepsilon}(\mu) \times \mathcal{O}^2_{2\varepsilon}(\mu).$$

In particular,

$$\mathcal{O}^1(\mu) \times \mathcal{O}^2(\mu) = EBE(\mu),$$

and the value $v(\mu)$ is the unique EBE payoff (to player 1) in the game $G(\mu)$.

Example 1 (Matrix Bayesian Game). Assume that each player i has n_i pure strategies, and S^i is the $(n_i - 1)$-dimensional simplex of i's mixed strategies. Assume further that in each $\omega \in \Omega$, the payoff function is given by

$$u\left(\omega, s^1, s^2\right) = s^1 A(\omega) s^2, \tag{4}$$

where strategy $s^1 \in S^1$ is regarded as a row vector, $s^2 \in S^2$ – as a column vector, and $A(\omega)$ is an $n_1 \times n_2$-matrix, with $A(\omega)_{j,k}$ being the payoff of player 1 when he chooses pure strategy j and 2 – pure strategy k, which is uniformly bounded across Ω. Then the strategy sets of players and the payoff function satisfy all the conditions listed above, and the associated zero-sum Bayesian game is amenable to our analysis.

2.2 Interim Expected Payoffs

The notions of the value of a game, and of the optimality of strategies, are defined with respect to players' *ex-ante* expected payoffs. In other words, players are assumed to evaluate their utilities before any private information is revealed. However, they may conceivably want to evaluate the consequences of their strategic choices at the interim stage, following the receipt of private information. In other words, players may be concerned with their *interim* expected payoff, that takes into account their private information and is based on the appropriately updated prior belief.

To formalize the discussion, let $\mu \in \Delta(\Omega, \mathcal{F})$. For any $\omega \in \Omega$ and $i = 1, 2$, denote by $\Pi^i(\omega)$ the element of partition Π^i that contains ω. If $\mu(\Pi^i(\omega)) > 0$, denote by $\mu_{\Pi^i(\omega)} \in \Delta(\Omega, \mathcal{F})$ the conditional belief of player i, given his information at ω, i.e., for any $A \in \mathcal{F}$,

$$\mu_{\Pi^i(\omega)}(A) = \mu\left(A \mid \Pi^i(\omega)\right) = \frac{\mu\left(A \cap \Pi^i(\omega)\right)}{\mu\left(\Pi^i(\omega)\right)}. \tag{5}$$

The function $U_{\mu_{\Pi^i(\omega)}}(\cdot, \cdot)$ will be referred to as the *interim expected payoff* given $\Pi^i(\omega)$.

For $\varepsilon \geq 0$, a pair $(\widehat{x}^1, \widehat{x}^2) \in X^1 \times X^2$ is called an *interim Bayesian ε-equilibrium* (henceforth, *ε-IBE* for short) in $G(\mu)$ if

$$U_{\mu_{\Pi^1(\omega)}}(\widehat{x}^1, \widehat{x}^2) \geq U_{\mu_{\Pi^1(\omega)}}(x^1, \widehat{x}^2) - \varepsilon \tag{6}$$

for every $x^1 \in X^1$ and every $\omega \in \Omega$ with $\mu\left(\Pi^1\left(\omega\right)\right) > 0$, and

$$U_{\mu_{\Pi^2(\omega)}}(\widehat{x}^1, \widehat{x}^2) \leq U_{\mu_{\Pi^2(\omega)}}(\widehat{x}^1, x^2) + \varepsilon \tag{7}$$

for every $x^2 \in X^2$ and every $\omega \in \Omega$ with $\mu\left(\Pi^2\left(\omega\right)\right) > 0$. Denote by $IBE_\varepsilon\left(\mu\right)$ the set of all ε-IBE in $G(\mu)$, and simplify $IBE_0\left(\mu\right)$ to $IBE\left(\mu\right)$. If $(\widehat{x}^1, \widehat{x}^2) \in IBE\left(\mu\right)$, we will call it an *interim Bayesian equilibrium* (*IBE* for short).

Remark 3. When $\varepsilon = 0$, there is no distinction between IBE and EBE. Definitions embodied in (2), (3) and (6), (7) are equivalent, as are indeed the notions of IBE and EBE in general, non-zero-sum, games. Accordingly, neither the value of a zero-sum Bayesian game (viewed as the ex-ante payoff in an IBE) nor the optimal strategies (viewed as IBE strategies) need not be redefined in the interim expected payoffs setting.

Remark 4. When $\varepsilon > 0$, the definition of ε-IBE is significantly more demanding than that of ε-EBE. Although any ε-IBE is in particular an ε-EBE, i.e., $IBE_\varepsilon\left(\mu\right) \subset EBE_\varepsilon\left(\mu\right)$, as follows from integrating both sides in (6) and (7) over Ω, the opposite is not true. In terms of the interim expected payoffs $U_{\mu_{\Pi^i(\omega)}}(\cdot, \cdot)$, the definition of $(\widehat{x}^1, \widehat{x}^2) \in EBE_\varepsilon\left(\mu\right)$ implies that

$$U_{\mu_{\Pi^1(\omega)}}(\widehat{x}^1, \widehat{x}^2) \geq U_{\mu_{\Pi^1(\omega)}}(x^1, \widehat{x}^2) - \frac{\varepsilon}{\mu\left(\Pi^1\left(\omega\right)\right)} \tag{8}$$

for every $x^1 \in X^1$ and every $\omega \in \Omega$ with $\mu\left(\Pi^1\left(\omega\right)\right) > 0$, and

$$U_{\mu_{\Pi^2(\omega)}}(\widehat{x}^1, \widehat{x}^2) \leq U_{\mu_{\Pi^2(\omega)}}(\widehat{x}^1, x^2) + \frac{\varepsilon}{\mu\left(\Pi^2\left(\omega\right)\right)} \tag{9}$$

for every $x^2 \in X^2$ and every $\omega \in \Omega$ with $\mu\left(\Pi^2\left(\omega\right)\right) > 0$. This indicates that although an ε-EBE strategy $\widehat{x}^i$ is ex-ante an ε-best response against $\widehat{x}^j$, it may be hugely *interim*-suboptimal in states of nature ω with low probability $\mu\left(\Pi^i\left(\omega\right)\right)$, thereby failing to be an ε'-IBE strategy for all sufficiently small ε'. (See, e.g., Example 2 in section 3.3.)

2.3 Topology on Common Priors

Consider the *total variation metric* d on $\Delta(\Omega, \mathcal{F})$, given by

$$d(\mu, \mu') = \sup_{E \in \mathcal{F}} |\mu(E) - \mu'(E)| \tag{10}$$

for any $\mu, \mu' \in \Delta(\Omega, \mathcal{F})$.

The following lemma shows that the expected payoff U_μ is a Lipschitz function of μ with respect to d, for a fixed $(x^1, x^2) \in X^1 \times X^2$.

Lemma 1. *For any $(x^1, x^2) \in X^1 \times X^2$ and $\mu, \mu' \in \Delta(\Omega, \mathcal{F})$,*

$$\left|U_\mu(x^1, x^2) - U_{\mu'}(x^1, x^2)\right| \leq 2Md(\mu, \mu')$$

Proof. For any $\mu, \mu' \in \Delta(\Omega, \mathcal{F})$,

$$\sup \int_\Omega f(\omega)\, d(\mu - \mu')(\omega) = 2d(\mu, \mu'), \tag{11}$$

where the supremum is taken over all $\mathcal{F}$-measurable functions $f : \Omega \to [-1, 1]$ (see, e.g., Lemma 1 on p. 360 in Shiryaev (1996)). Given any $(x^1, x^2) \in X^1 \times X^2$, note that, by the boundedness of u and (11),

$$\begin{aligned} \left|U_\mu(x^1, x^2) - U_{\mu'}(x^1, x^2)\right| &= \left|\int_\Omega u\left(\omega, x^1(\omega), x^2(\omega)\right) d\mu(\omega) \right. \\ &\quad \left. - \int_\Omega u\left(\omega, x^1(\omega), x^2(\omega)\right) d\mu'(\omega)\right| \\ &= \left|\int_\Omega u\left(\omega, x^1(\omega), x^2(\omega)\right) d(\mu - \mu')(\omega)\right| \\ &\leq 2Md(\mu, \mu'). \end{aligned}$$

□

3 Results

3.1 Continuity of Value

Our main result establishes Lipschitz continuity of the value:

Theorem 1. *The value $v(\mu)$ is a Lipschitz continuous function of μ with respect to the metric d: for any $\mu, \mu' \in \Delta(\Omega, \mathcal{F})$,*

$$|v(\mu) - v(\mu')| \leq 2Md(\mu, \mu'). \tag{12}$$

Proof. Let $\hat{x}^1 \in \mathcal{O}^1(\mu)$ and $x^2 \in X^2$. The optimality of $\hat{x}^1$ in $G(\mu)$ and Lemma 1 imply that

$$\begin{aligned} U_{\mu'}(\hat{x}^1, x^2) &\geq U_\mu(\hat{x}^1, x^2) - 2Md(\mu, \mu') \\ &\geq v(\mu) - 2Md(\mu, \mu'). \end{aligned}$$

This holds for every $x^2 \in X^2$, and hence it follows that

$$\begin{aligned} v(\mu') &= \max_{x^1 \in X^1} \min_{x^2 \in X^2} U_{\mu'}(x^1, x^2) && (13a) \\ &\geq \min_{x^2 \in X^2} U_{\mu'}(\hat{x}^1, x^2) \\ &\geq v(\mu) - 2Md(\mu, \mu'). && (13b) \end{aligned}$$

Similarly, starting with $\hat{x}^2 \in \mathcal{O}^2(\mu)$ we obtain

$$\begin{aligned} v(\mu') &= \min_{x^2 \in X^2} \max_{x^1 \in X^1} U_{\mu'}(x^1, x^2) && (14a) \\ &\leq \max_{x^1 \in X^1} U_{\mu'}(x^1, \hat{x}^2) \\ &\leq v(\mu) + 2Md(\mu, \mu') && (14b) \end{aligned}$$

The combination of (13) and (14) yields (12). □

Theorem 1 implies, in particular, that the single-valued EBE (or, equivalently via Remark 3, IBE) *expected payoff* correspondence is both lower and upper semi-continuous with respect to the total variation metric on the common prior, when restricted to zero-sum Bayesian games. This stands in contrast to the general, non-zero-sum case. As was shown by Kajii and Morris (1998), IBE payoffs in a non-zero-sum Bayesian game $G(\mu)$ may be quite far from ε-IBE payoffs in $G(\mu')$ for all small enough $\varepsilon > 0$, even if $d(\mu, \mu')$ is arbitrarily small, when the beliefs *conditional* on each player's private information do not converge uniformly. But, in the zero-sum case, the IBE payoff in $G(\mu)$ is approximated by the *exact* IBE payoff in $G(\mu')$ (not merely an ε-IBE payoff) when $d(\mu, \mu') \to 0$.

3.2 Upper Semi-continuity of Optimal Strategies

Otimal strategies too have strong continuity properties with respect to the total variation metric on the common prior, as we will see in this and the next subsection. Given a sequence $\{\mu_n\}_{n=0}^{\infty} \subset \Delta(\Omega, \mathcal{F})$ such that $\lim_{n\to\infty} \mu_n = \mu_0$, we say that the optimal strategy correspondence is *upper semi-continuous* (USC) along $\{\mu_n\}_{n=1}^{\infty}$ if the following holds: for any player i and any sequence $\{\hat{x}_n^i\}_{n=1}^{\infty} \subset X^i$ such that $\lim_{n\to\infty} \hat{x}_n^i = \hat{x}_0^i$ and $\hat{x}_n^i \in \mathcal{O}^i(\mu_n)$ for each $n \geq 1$, $\hat{x}_0^i \in \mathcal{O}^i(\mu_0)$.

Proposition 1. *The optimal strategy correspondence is USC along any convergent sequence* $\{\mu_n\}_{n=1}^{\infty} \subset \Delta(\Omega, \mathcal{F})$.

Proof. Let $\{\mu_n\}_{n=1}^{\infty} \subset \Delta(\Omega, \mathcal{F})$ and $\{\hat{x}_n^i\}_{n=1}^{\infty} \subset X^i$ be such that $\lim_{n\to\infty} \mu_n = \mu_0$, $\lim_{n\to\infty} \hat{x}_n^i = \hat{x}_0^i$, and $\hat{x}_n^i \in O^i(\mu_n)$ for each $n \geq 1$. We will assume $i = 1$, the case of $i = 2$ being analogous. Take any $x^2 \in X^2$. By assumption, for any $n \geq 1$

$$U_{\mu_n}(\hat{x}_n^1, x^2) \geq v(\mu_n). \tag{15}$$

Using Lemma 1, we obtain

$$\begin{aligned} & \left|U_{\mu_n}(\hat{x}_n^1, x^2) - U_{\mu_0}(\hat{x}_0^1, x^2)\right| \\ \leq\ & \left|U_{\mu_n}(\hat{x}_n^1, x^2) - U_{\mu_0}(\hat{x}_n^1, x^2)\right| + \left|U_{\mu_0}(\hat{x}_n^1, x^2) - U_{\mu_0}(\hat{x}_0^1, x^2)\right| \\ \leq\ & 2Md(\mu_n, \mu_0) + \left|U_{\mu_0}(\hat{x}_n^1, x^2) - U_{\mu_0}(\hat{x}_0^1, x^2)\right|, \end{aligned}$$

and thus

$$\lim_{n\to\infty} U_{\mu_n}(\hat{x}_n^1, x^2) = \lim_{n\to\infty} U_{\mu_0}(\hat{x}_n^1, x^2) = U_{\mu_0}(\hat{x}_0^1, x^2)$$

by continuity of U_{μ_0} in the first variable. Now, taking the limits of both sides in (15) and using Theorem 1 yields

$$U_{\mu_0}(\hat{x}_0^1, x^2) \geq v(\mu_0).$$

Since this holds for every $x^2 \in X^2$, $\hat{x}_0^1 \in \mathcal{O}^i(\mu_0)$. $\square$

Since $\mathcal{O}^1(\mu) \times \mathcal{O}^2(\mu) = EBE(\mu)$ for every $\mu \in \Delta(\Omega, \mathcal{F})$, as was mentioned in Remark 2, Proposition 1 shows that the EBE correspondence (and, equivalently by Remark 3, the IBE correspondence) is USC in zero-sum Bayesian games. In fact, this is true for *general* Bayesian games in our set-up, as can be easily seen.[8]

3.3 Approximate Lower Semi-continuity of Optimal Strategies, EBE, and IBE

Defining lower semi-continuity of the optimal strategy correspondence requires some care. Its straightforward version will not work: given $\{\mu_n\}_{n=0}^{\infty} \subset \Delta(\Omega, \mathcal{F})$ with $\lim_{n\to\infty} \mu_n = \mu_0$ and $\hat{x}_0^i \in \mathcal{O}^i(\mu_0)$, we may not be able to find a sequence $\{\hat{x}_n^i\}_{n=1}^{\infty} \subset X^i$ such that $\lim_{n\to\infty} \hat{x}_n^i = \hat{x}_0^i$ and $\hat{x}_n^i$ is *optimal* in $G(\mu_n)$ for each $n \geq 1$. Indeed, even in a simple decision problem (i.e., a one-player game) not all payoff maximizers may be approximable by maximizers in nearby problems. Thus, the appropriate notion of lower semi-continuity of the optimal strategy correspondence with respect to the common prior is the following. We will say that the optimal strategy correspondence is *approximately lower semi-continuous (ALSC)* along a sequence $\{\mu_n\}_{n=1}^{\infty}$ with $\lim_{n\to\infty} \mu_n = \mu_0$ if the following holds: Given any $\hat{x}_0^i \in \mathcal{O}^i(\mu_0)$ for some player i, and any $\varepsilon > 0$, there exists a sequence $\{\hat{x}_n^i\}_{n=1}^{\infty} \subset X^i$ with $\lim_{n\to\infty} \hat{x}_n^i = \hat{x}_0^i$, such that $\hat{x}_n^i$ is (merely) ε-*optimal* in $G(\mu_n)$ for every $n \geq 1$.

Similarly, the EBE (respectively, IBE) correspondence is defined to be *ALSC* along a sequence $\{\mu_n\}_{n=1}^{\infty}$ with $\lim_{n\to\infty} \mu_n = \mu_0$ by the requirement that, given any $(\hat{x}^1, \hat{x}^2) \in EBE(\mu_0)$ (respectively, $IBE(\mu_0)$) and any $\varepsilon > 0$, there exists a sequence $\{(\hat{x}_n^1, \hat{x}_n^2)\}_{n=1}^{\infty} \subset X^1 \times X^2$ with $\lim_{n\to\infty} (\hat{x}_n^1, \hat{x}_n^2) =$

[8] Milgrom and Weber (1985) have a very general result on USC of the BE correspondence for general Bayesian games, but in a somewhat different framework.

$(\hat{x}_0^1, \hat{x}_0^2)$, such that $(\hat{x}_n^1, \hat{x}_n^2) \in EBE_\varepsilon(\mu_n)$ (respectively, $IBE_\varepsilon(\mu_n)$) for every $n \geq 1$.

It follows from the next proposition that the optimal strategy correspondence is ALSC, with the sequence $\{\hat{x}_n^i\}_{n=1}^\infty$ being the strategy $\hat{x}_0^i$ itself:

Proposition 2. *For every $\mu, \mu' \in \Delta(\Omega, \mathcal{F})$ and every $i = 1, 2$,*

$$\mathcal{O}^i(\mu) \subset \mathcal{O}^i_{4Md(\mu,\mu')}(\mu').$$

Proof. As in the proof of Proposition 1, we will only consider the case of $i = 1$. Fix any $\mu \in \Delta(\Omega, \mathcal{F})$ and let $\hat{x}^1 \in \mathcal{O}^1(\mu)$. By Lemma 1, optimality of $\hat{x}^1$, and Theorem 1, for any $x^2 \in X^2$ and any $\mu' \in \Delta(\Omega, \mathcal{F})$,

$$\begin{aligned} U_{\mu'}(\hat{x}^1, x^2) &\geq U_\mu(\hat{x}^1, x^2) - 2Md(\mu, \mu') \\ &\geq v(\mu) - 2Md(\mu, \mu') \\ &\geq v(\mu') - 4Md(\mu, \mu'). \end{aligned}$$

This shows that $\hat{x}^1$ is indeed $4Md(\mu, \mu')$-optimal for player 1 in $G(\mu')$. □

According to Proposition 2, if $\lim_{n\to\infty} \mu_n = \mu_0$ then, for a given $\varepsilon > 0$, any $\hat{x}_0^i$ which is optimal in $G(\mu_0)$ is also ε-optimal in $G(\mu_n)$ for all sufficiently large n. However, while optimality of a strategy has an interpretation in terms of both ex-ante and interim expected payoffs (since $\mathcal{O}^1(\mu) \times \mathcal{O}^2(\mu) = EBE(\mu) = IBE(\mu)$, by Remarks 2 and 3), this is no longer so with ε-optimality which is a purely ex-ante concept (as expounded in Remark 4). Thus, although Proposition 2 trivially implies that the EBE correspondence is ALSC[9] along any converging sequence $\{\mu_n\}_{n=1}^\infty$ (since $EBE(\mu_0) = \mathcal{O}^1(\mu_0) \times \mathcal{O}^2(\mu_0)$ and $EBE_\varepsilon(\mu_n) \supset \mathcal{O}^1_{\frac{\varepsilon}{2}}(\mu_n) \times \mathcal{O}^2_{\frac{\varepsilon}{2}}(\mu_n)$ by Remark 2), it remains mute on IBE. And indeed, a pair $(\hat{x}^1, \hat{x}^2) \in \mathcal{O}^1(\mu_0) \times \mathcal{O}^2(\mu_0) = IBE(\mu_0)$ may fail to be in $IBE_\varepsilon(\mu_n)$ for all $n \geq 1$ and all small enough ε:

Example 2. Let $\Omega = \mathbb{Z}_+$ (the set of non-negative integers), $S^1 = [0, 1]$, $S^2 = \{0\}$, $\Pi^1 = \Pi^2 \equiv \{\{2n, 2n+1\} : n \in \mathbb{Z}_+\}$, and, finally,

$$u(\omega, s^1, s^2) \equiv \begin{cases} -(s^1)^2, & \text{if } \omega \text{ is odd,} \\ 0, & \text{is } \omega \text{ is even.} \end{cases}.$$

[9] This is, in fact, well known even for *general* Bayesian games, due to Engl (1995).

If μ_0 is a probability measure with the full support on $\{2n : n \in \mathbb{Z}_+\}$, then, clearly, $\hat{x}_0^1(\cdot) \equiv 1$ is an optimal strategy of player 1 in $G(\mu_0)$. Consider, however, a sequence $\{\mu_n\}_{n=1}^{\infty}$ of probability measure, such that, for every n, μ_n is identical to μ_0 on all subsets of $\Omega \backslash \{2n, 2n+1\}$, but $\mu_n(\{2n\}) = \mu_n(\{2n+1\}) = \frac{1}{2}\mu_0(\{2n\})$. Then clearly $\lim_{n\to\infty} \mu_n = \mu_0$, but $(\hat{x}_0^1, 0) \notin IBE_\varepsilon(\mu_n)$ for every $n \geq 1$ and all $\varepsilon \in [0, \frac{1}{2})$.

However, the failure of some $(\hat{x}_0^1, \hat{x}_0^2) \in IBE(\mu_0)$ to be in $IBE_\varepsilon(\mu_n)$ for all small enough $\varepsilon > 0$, does not rule out that the IBE correspondence is ALSC: it does not preclude the possibility that $(\hat{x}_0^1, \hat{x}_0^2)$ is *approximable* by a sequence $\{(\hat{x}_n^1, \hat{x}_n^2)\}_{n=1}^{\infty}$ with $(\hat{x}_n^1, \hat{x}_n^2) \in IBE_\varepsilon(\mu_n)$ for each $n \geq 1$. Indeed, the IBE $(\hat{x}_0^1, 0)$ in Example 2 is the limit of the sequence $\{(\hat{x}_n^1, 0)\}_{n=1}^{\infty}$, where $\hat{x}_n^1(\{2k, 2k+1\}) \equiv \begin{cases} 1, & \text{if } k \neq n, \\ 0, & \text{if } k = n \end{cases}$ for all $k \in \mathbb{Z}_+$, and $\{(\hat{x}_n^1, 0)\}_{n=1}^{\infty}$ are ε-IBE (and even IBE) in $\{G(\mu_n)\}_{n=1}^{\infty}$. But is the IBE correspondence ALSC in general zero-sum Bayesian games?

This is, at present, an open problem. However, we will show that the the IBE correspondence is ALSC in zero-sum games, with information structures belonging to a certain interesting class, for which the ALSC property would have failed had the games been *non-zero-sum.* We start by recalling the following example, due to Kajii and Morris (1994), showing that the IBE correspondence in non-zero-sum Bayesian games is not ALSC with respect to the total variation metric on $\Delta(\Omega, \mathcal{F})$:

Example 3 (Kajii and Morris (1994), Section 4.2). This is an elaboration of the coordinated attack problem in the computer science literature, and an electronic mail game of Rubinstein (1989); for a more methodical and detailed presentation of the example the reader is referred to Kajii and Morris (1994). Let $\Omega = \mathbb{Z}_+ \times \mathbb{Z}_+$, $S^1 = S^2 = [0, 1]$, and assume that each player i can discern only the i^{th} coordinate in each state $(t^1, t^2) \in \Omega$, i.e., that $\Pi^1((t^1, t^2)) = \{t^1\} \times \mathbb{Z}_+$ and $\Pi^1((t^1, t^2)) = \mathbb{Z}_+ \times \{t^2\}$. Furthermore, let $\mu_n(\{(0, 0)\}) = \frac{1}{2}$, $\mu_n(\{(0, k)\}) = \left(\frac{1}{2}\right)^{k+1} \alpha^n$, $\mu_n(\{(k, k)\}) = \left(\frac{1}{2}\right)^{2k}(1 - \alpha^n)$, and $\mu_n(\{(k, k+1)\}) = \left(\frac{1}{2}\right)^{2k+1}(1 - \alpha^n)$ for all $k \geq 1$, where $\alpha^n \to_{n\to\infty} 0$; and let $\mu_n(\{(t^1, t^2)\}) = 0$ for all other $(t^1, t^2) \in \Omega$. The limit measure, $\mu_0 = \lim_{n\to\infty} \mu_n$, is thus given by $\mu_0(\{(0, 0)\}) = \frac{1}{2}$, $\mu_0(\{(k, k)\}) = \left(\frac{1}{2}\right)^{2k}$, $\mu_0(\{(k, k+1)\}) = \left(\frac{1}{2}\right)^{2k+1}$ for all $k \geq 1$, and $\mu_0(\{(t^1, t^2)\}) = 0$ otherwise.

In each state of nature (t^1, t^2), strategy $s^i \in [0,1]$ is the probability of choosing pure strategy $Safe$ from the binary set of pure strategies $\{Dangerous, Safe\}$. In each $(t^1, t^2) \in \Omega$ where $t^1 = 0$, the pure strategy payoffs are given by

	$Dangerous$	$Safe$
$Dangerous$	$(-10,-10)$	$(-10,1)$
$Safe$	$(1,-10)$	$(1,1)$

and, when $t^1 \neq 0$, by

	$Dangerous$	$Safe$
$Dangerous$	$(2,2)$	$(-10,1)$
$Safe$	$(1,-10)$	$(1,1)$

.

Kajii and Morris (1994) show that an IBE in the non-zero-sum Bayesian game $G(\mu_0)$, where both players play $Safe$ when it is (commonly) known that $t^1 = 0$ and $Dangerous$ otherwise, is far from *all* ε-IBE in $\{G(\mu_n)\}_{n=1}^{\infty}$, for all sufficiently low values of ε, and in all non-zero states of nature (in fact, they show that even IBE *payoffs* are far apart). What is at fault in that example is a *non-uniform* (across Ω) convergence of prior beliefs conditional on players' private information (i.e., of measures $\left((\mu_n)_{\Pi^i(\omega)}\right)_{i=1,2;\omega\in\Omega}$), which occurs despite that $\lim_{n\to\infty}\mu_n = \mu_0$ in the total variation metric (see also Example 2 above). More precisely, the problem lies in the lack of *almost uniform convergence* of conditional beliefs, which is defined, roughly speaking, by the requirement that the closeness of conditional beliefs becomes approximate common knowledge with high ex-ante probability.[10]

We will show in our next Proposition 3 that the (almost) uniform convergence of conditional beliefs mentioned in Example 3 is by no means necessary for the IBE correspondence to be ALSC in *zero-sum* Bayesian games. The proposition makes certain assumptions on the support of $\{\mu_n\}_{n=0}^{\infty}$ (that are satisfied, e.g., by the information structure in Example 3). But, contrary to

[10]The notion of almost uniform convergence of conditional beliefs is defined and expounded upon in Kajii and Morris (1994), where they also show (on p. 19) that it is not satisfied in this example. We do not attempt to give a formal definition here, since this would require a lengthy digression, and this notion's relevance is primarily for the issue of IBE *expected payoff* continuity in general (non-zero-sum) games, which is not our focus.

the case in Example 3, in zero-sum Bayesian games these assumptions guarantee ALSC behavior along $\{\mu_n\}_{n=1}^{\infty}$, without any additional requirement on the convergence of conditional beliefs.

Before we state Proposition 3, the following convention is in order. For each $i = 1, 2$, we will write the elements of the (at most countable) partition Π^i as an indexed sequence $\{\pi_j^i\}_{j=0}^{T^i}$, where $T^i \in \{\infty\} \cup \mathbb{Z}_+$. When i's private information is given by π_j^i (i.e., when the realized state of nature ω is such that $\Pi^i(\omega) = \pi_j^i$), the index j of π_j^i may be referred to as the *type* of player i.

Proposition 3. *Let $\{\mu_n\}_{n=0}^{\infty} \subset \Delta(\Omega, \mathcal{F})$ be such that $\lim_{n\to\infty} \mu_n = \mu_0$, and assume that in zero-sum Bayesian games $\{G(\mu_n)\}_{n=0}^{\infty}$:*

(I) *there exists $C > 0$ such that, for every $\omega \in \Omega$, the payoff function $u(\omega, \cdot, \cdot)$ is Lipschitz continuous with a constant C with respect to the Euclidean norm on $S^1 \times S^2$;*

(II) *$T^1 = T^2 = \infty$, and there exist functions $t^1, t^2 : \mathbb{Z}_+ \to \mathbb{Z}_+$ with $\lim_{n\to\infty} \min(t^1, t^2)(n) = \infty$, and integers $K, L \geq 0$ such that, for every $i = 1, 2$:*

(a) $\mu_0(\pi_j^i) > 0$ for every $j \geq 0$;

(b) if $j \leq L$, the measure $(\mu_0)_{\pi_j^i}$ has a support on the set $\bigcup_{k=0}^{L} \pi_k^{-i}$, where $-i$ denotes the rival of player i.

and

(c) for every $n \geq 0$, if $j > L$ and $\mu_n(\pi_j^i) > 0$ then the measure $(\mu_n)_{\pi_j^i}$ has a support on the set $\bigcup_{k=0}^{L} \pi_k^{-i} \cup \bigcup_{k=t^i(j)}^{t^i(j)+K} \pi_k^{-i}$.

Then the IBE correspondence is ALSC along $\{\mu_n\}_{n=1}^{\infty}$.

Note that assumption (I) is implied by the uniform boundedness of u if $\{G(\mu_n)\}_{n=0}^{\infty}$ are Bayesian *matrix* games (see Example 1). Assumptions (IIb,c) mean that although player i may be unsure of the type of his rival$-i$, given the knowledge of his own type he knows that $-i$'s type is either very low (not exceeding L) or can be estimated via the function t^i with an error of at most K. This assumption is satisfied by the information structure in Example 3 above.

Proof of Proposition 3. Let $\{\mu_n\}_{n=0}^{\infty} \subset \Delta(\Omega, \mathcal{F})$ be a sequence satisfying all the assumptions. Fix $(\hat{x}_0^1, \hat{x}_0^2) \in IBE(\mu_0)$ and $\varepsilon > 0$.We begin with the following notation. For any $n \geq 0$, denote by $G'(\mu_n)$ an auxilliary game that is identical to $G(\mu_n)$, in all but the following aspect: the strategy-set $X^{i\prime}$ of player i in $G'(\mu_n)$ is the subset of X^i consisting of all strategies that coinside with $\hat{x}_0^i$ on $\bigcup_{k=0}^{L} \pi_k^i$. The notions of ($\varepsilon$-)optimal strategies of player i, (ε-)EBE, and (ε-)IBE are defined for $G'(\mu_n)$ in the same way as for $G(\mu_n)$, bearing in mind the constraint on strategies, and their (non-empty, by our assumptions on the game) sets will be denoted by $\mathcal{O}_\varepsilon^{i\prime}(\mu_n)$, $IBE'_\varepsilon(\mu_n)$, and $IBE'_\varepsilon(\mu_n)$, respectively. The subindex will be droped if $\varepsilon = 0$. Note that Remarks 2, 3, and 4 hold in their enterity if the sets $\mathcal{O}_\varepsilon^{i}(\mu_n)$, $IBE_\varepsilon(\mu_n)$, and $IBE_\varepsilon(\mu_n)$ are replaced by their "tagged" counterparts, $\mathcal{O}_\varepsilon^{i\prime}(\mu_n)$, $IBE'_\varepsilon(\mu_n)$, and $IBE'_\varepsilon(\mu_n)$, respectively; we will refer to the "tagged" versions of the remarks as **Remarks 2'**, **3'**, and **4'**.

Consider now a sequence $\{(\widehat{y}_n^1, \widehat{y}_n^2)\}_{n=1}^{\infty} \subset X^{1\prime} \times X^{2\prime}$ such that $(\widehat{y}_n^1, \widehat{y}_n^2) \in \mathcal{O}^{1\prime}(\mu_n) \times \mathcal{O}^{2\prime}(\mu_n)$ for each $n \geq 1$. Let

$$0 < \delta \leq \delta_0 \equiv \frac{\varepsilon}{8CK^2\left(\max_{s^1 \in S^1} \|s^1\| + \max_{s^2 \in S^2} \|s^2\|\right)}, \tag{16}$$

where $\|\cdot\|$ denotes the Euclidean norm on both S^1 and S^2. Define a sequence $\left\{\left(\overline{x}_{n,\delta}^1, \overline{x}_{n,\delta}^2\right)\right\}_{n=1}^{\infty} \subset X^{1\prime} \times X^{2\prime}$ as follows: for every $n \geq 1$, $i = 1, 2$, and $j > L$, set

$$\overline{x}_{n,\delta}^i\left(\pi_j^i\right) \equiv \max\left(1 - \delta(j-L), 0\right) \cdot \widehat{x}_0^i\left(\pi_j^i\right) + \min\left(\delta(j-L), 1\right) \cdot \widehat{y}_n^i\left(\pi_j^i\right) \tag{17}$$

(recall also that each $\overline{x}_{n,\delta}^i \in X^{i\prime}$ coinsides with $\hat{x}_0^i$ on $\bigcup_{k=1}^{L} \pi_k^i$, and hence, when $j \leq L$, $\overline{x}_{n,\delta}^i\left(\pi_j^i\right) = \hat{x}_0^i\left(\pi_j^i\right)$). Thus $\left(\overline{x}_{n,\delta}^1, \overline{x}_{n,\delta}^2\right)$ is a state-dependent *mixture* of $(\hat{x}_0^1, \hat{x}_0^2)$ with $(\widehat{y}_n^1, \widehat{y}_n^2)$, with the property that $\overline{x}_{n,\delta}^i\left(\pi_j^i\right)$ is gradually transformed from $\widehat{x}_0^i\left(\pi_j^i\right)$ into $\widehat{y}_n^i\left(\pi_j^i\right)$ when j moves from L to[11] $L+\left[\frac{1}{\delta}\right]+1$.

We claim that $\left(\overline{x}_{n,\delta}^1, \overline{x}_{n,\delta}^2\right) \in IBE_\varepsilon(\mu_n)$ for each sufficiently large n. We will only show that there exists $N(\delta) > 0$ such that

$$U_{(\mu_n)_{\pi_j^1}}(\overline{x}_{n,\delta}^1, \overline{x}_{n,\delta}^2) \geq U_{(\mu_n)_{\pi_j^1}}(x^1, \overline{x}_{n,\delta}^2) - \varepsilon \tag{18}$$

for every $x^1 \in X^1$, every $j \geq 0$, and every $n \geq N(\delta)$. It could then be established similarly that

$$U_{(\mu_n)_{\pi_j^2}}(\overline{x}_{n,\delta}^1, \overline{x}_{n,\delta}^2) \leq U_{(\mu_n)_{\pi_j^2}}(\overline{x}_{n,\delta}^1, x^2) + \varepsilon \tag{19}$$

for every $x^2 \in X^2$, every $j \geq 0$, and every $n \geq N(\delta)$, and thus (6) and (7) in the definition of ε-IBE will follow.

Since $\lim_{n\to\infty} \min(t^1, t^2)(n) = \infty$, we can find an integer $J(\delta) \geq L + \left[\frac{1}{\delta}\right] + 1$ such that, for all $j \geq J(\delta)$ and $i = 1, 2$, $t^i(j) \geq L + \left[\frac{1}{\delta}\right] + 1$. Given $x^1 \in X^1$, consider the following three cases.

Case 1: $j \geq J(\delta)$. Note that, by (17) and the choice of $J(\delta)$, $\overline{x}_{n,\delta}^1(\omega) = \widehat{y}_n^1(\omega)$ for every $\omega \in \pi_j^1$ and $\overline{x}_{n,\delta}^2(\omega) = \widehat{y}_n^2(\omega)$ for every $\omega \in \bigcup_{k=t^1(j)}^{t^1(j)+K} \pi_k^2$. Thus, by assumption (IIc),

$$U_{(\mu_n)_{\pi_j^1}}(\overline{x}_{n,\delta}^1, \overline{x}_{n,\delta}^2) = U_{(\mu_n)_{\pi_j^1}}(\widehat{y}_n^1, \widehat{y}_n^2) \tag{20}$$

and

$$U_{(\mu_n)_{\pi_j^1}}(x^1, \overline{x}_{n,\delta}^2) = U_{(\mu_n)_{\pi_j^1}}(x^1, \widehat{y}_n^2). \tag{21}$$

But $(\widehat{y}_n^1, \widehat{y}_n^2) \in \mathcal{O}^{1\prime}(\mu_n) \times \mathcal{O}^{2\prime}(\mu_n) = IBE'(\mu_n)$ (see Remarks 2' and 3'). This fact and (20), (21) imply (18) for $x^1 \in X^{1\prime}$. But since any $x^1 \in X^1$ is identical to some $x^{1\prime} \in X^{1\prime}$ when restricted to π_j^1, (18) in fact holds for all $x^1 \in X^1$, and every n.

Case 2: $L < j < J(\delta)$. Denote

$$\overline{\overline{x}}_{n,\delta}^1 \equiv \max(1 - \delta(j-L), 0) \cdot \widehat{x}_0^1 + \min(\delta(j-L), 1) \cdot \widehat{y}_n^1$$

[11] Henceforth, $\left[\frac{1}{\delta}\right]$ will stand for the integer part of $\frac{1}{\delta}$.

and

$$\overline{\overline{x}}_{n,\delta}^{2} \equiv \max\left(1-\delta\left(t^{1}\left(j\right)-L\right),0\right)\cdot\widehat{x}_{0}^{2}+\min\left(\delta\left(t^{1}\left(j\right)-L\right),1\right)\cdot\widehat{y}_{n}^{2}$$

(note a subtle but important difference between $\overline{x}_{n,\delta}^{i}$ and $\overline{\overline{x}}_{n,\delta}^{i}$ – in the former the coefficients in the convex combination are state-dependent, while in the latter they are not). It follows from the definition of $\overline{x}_{n,\delta}^{i}$ and $\overline{\overline{x}}_{n,\delta}^{i}$ and assumptions (I) and (II) that

$$\left|U_{(\mu_n)_{\pi_j^1}}(\overline{x}_{n,\delta}^{1},\overline{x}_{n,\delta}^{2})-U_{(\mu_n)_{\pi_j^1}}(\overline{\overline{x}}_{n,\delta}^{1},\overline{\overline{x}}_{n,\delta}^{2})\right|$$

(by assumption (IIc))

$$\begin{aligned}\leq\ &\sum_{k=0}^{L}\int_{\pi_j^1\cap\pi_k^2}\left|u\left(\omega,\overline{x}_{n,\delta}^{1}(w),\overline{x}_{n,\delta}^{2}\left(w\right)\right)-u\left(\omega,\overline{\overline{x}}_{n,\delta}^{1}(w),\overline{\overline{x}}_{n,\delta}^{2}\left(w\right)\right)\right|d\left(\mu_n\right)_{\pi_j^1}(\omega)\\ &+\sum_{k=t^1(j)}^{t^1(j)+K}\int_{\pi_j^1\cap\pi_k^2}\left|u\left(\omega,\overline{x}_{n,\delta}^{1}(w),\overline{x}_{n,\delta}^{2}\left(w\right)\right)-u\left(\omega,\overline{\overline{x}}_{n,\delta}^{1}(w),\overline{\overline{x}}_{n,\delta}^{2}\left(w\right)\right)\right|d\left(\mu_n\right)_{\pi_j^1}(\omega)\end{aligned}$$

(by the definition of $\overline{x}_{n,\delta}^{1}$ and $\overline{\overline{x}}_{n,\delta}^{1}$)

$$\begin{aligned}=\ &\sum_{k=0}^{L}\int_{\pi_j^1\cap\pi_k^2}\left|u\left(\omega,\overline{x}_{n,\delta}^{1}(w),\hat{x}_{0}^{2}\left(w\right)\right)-u\left(\omega,\overline{x}_{n,\delta}^{1}(w),\hat{x}_{0}^{2}\left(w\right)\right)\right|d\left(\mu_n\right)_{\pi_j^1}(\omega)\\ &+\sum_{k=t^1(j)}^{t^1(j)+K}\int_{\pi_j^1\cap\pi_k^2}\left|u\left(\omega,\overline{x}_{n,\delta}^{1}(w),\overline{x}_{n,\delta}^{2}\left(w\right)\right)-u\left(\omega,\overline{x}_{n,\delta}^{1}(w),\overline{\overline{x}}_{n,\delta}^{2}\left(w\right)\right)\right|d\left(\mu_n\right)_{\pi_j^1}(\omega)\end{aligned}$$

(by assumption (I) and the definition of $\overline{x}_{n,\delta}^{2}$ and $\overline{\overline{x}}_{n,\delta}^{2}$)

$$\leq\sum_{k=t^1(j)}^{t^1(j)+K}\int_{\pi_j^1\cap\pi_k^2}C\delta\left(k-t^1\left(j\right)\right)\left(\left\|\widehat{x}_{0}^{2}\left(w\right)\right\|+\left\|\widehat{y}_{n}^{2}\left(w\right)\right\|\right)d\left(\mu_n\right)_{\pi_j^1}(\omega)$$

$$\leq 2C\delta K^{2}\max_{s^2\in S^2}\left\|s^{2}\right\|.$$

Similarly, it can be shown that

$$\left|U_{(\mu_n)_{\pi_j^1}}(x^{1},\overline{x}_{n,\delta}^{2})-U_{(\mu_n)_{\pi_j^1}}(x^{1},\overline{\overline{x}}_{n,\delta}^{2})\right|\leq 2C\delta K^{2}\max_{s^2\in S^2}\left\|s^{2}\right\|.$$

Thus,

$$U_{(\mu_n)_{\pi^1_j}}(x^1, \overline{x}^2_{n,\delta}) - U_{(\mu_n)_{\pi^1_j}}(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta}) \tag{22a}$$

$$\leq \ U_{(\mu_n)_{\pi^1_j}}(x^1, \overline{\overline{x}}^2_{n,\delta}) - U_{(\mu_n)_{\pi^1_j}}(\overline{\overline{x}}^1_{n,\delta}, \overline{\overline{x}}^2_{n,\delta}) + 4C\delta K^2 \max_{s^2 \in S^2} \left\| s^2 \right\|. \tag{22b}$$

Since $\lim_{n\to\infty} \mu_n = \mu_0$, it follows from assumption (IIa) on μ_0 that there exists $N'(\delta) > 0$ such that

$$\varepsilon' \equiv \frac{\varepsilon}{2} \inf \left\{ \mu_n\left(\pi^i_j\right) \mid i \in \{1,2\}; L \leq j \leq J(\delta); N'(\delta) \leq n \right\} > 0. \tag{23}$$

By Proposition 2, the assumption that $\lim_{n\to\infty} \mu_n = \mu_0$, and the fact that $(\hat{x}^1_0, \hat{x}^2_0) \in \mathcal{O}^1(\mu_0) \times \mathcal{O}^2(\mu_0)$ (implied by the choice of $(\hat{x}^1_0, \hat{x}^2_0)$ and Remark 3), there exists $N''(\delta) \geq N'(\delta)$ such that for every $n \geq N''(\delta)$, $(\hat{x}^1_0, \hat{x}^2_0) \in \mathcal{O}^1_{\varepsilon'}(\mu_n) \times \mathcal{O}^2_{\varepsilon'}(\mu_n)$. Thus $(\hat{x}^1_0, \hat{x}^2_0) \in \mathcal{O}^{1\prime}_{\varepsilon'}(\mu_n) \times \mathcal{O}^{2\prime}_{\varepsilon'}(\mu_n)$ as well for every $n \geq N''(\delta)$. But then clearly

$$\overline{\overline{x}}^i_{n,\delta} \in \mathcal{O}^{i\prime}_{\varepsilon'}(\mu_n) \tag{24}$$

for $i = 1, 2$ and every $n \geq N''(\delta)$, being a convex combination of ε'-optimal strategies. Therefore, by (8) in Remark 4', for every $x^1 \in X^{1\prime}$ and $n \geq N''(\delta)$

$$U_{(\mu_n)_{\pi^1_j}}(x^1, \overline{\overline{x}}^2_{n,\delta}) - U_{(\mu_n)_{\pi^1_j}}(\overline{\overline{x}}^1_{n,\delta}, \overline{\overline{x}}^2_{n,\delta}) \leq \frac{\varepsilon'}{\mu_n\left(\pi^1_j\right)}. \tag{25}$$

From (22) and (25), for every $x^1 \in X^{1\prime}$ and $n \geq N''(\delta)$

$$U_{(\mu_n)_{\pi^1_j}}(x^1, \overline{x}^2_{n,\delta}) - U_{(\mu_n)_{\pi^1_j}}(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta}) \leq \frac{\varepsilon'}{\mu_n\left(\pi^1_j\right)} + 4C\delta K^2 \max_{s^2 \in S^2} \left\| s^2 \right\| \leq \varepsilon, \tag{26}$$

where the last inequality is implied by the definition of ε' and the choice of δ (see (16), (23)). As in Case 1, this is true for any $x^1 \in X^1$, not just $x^1 \in X^{1\prime}$.

Case 3: $0 \leq j \leq L$. Since $(\hat{x}^1_0, \hat{x}^2_0) \in IBE(\mu_0)$ and $(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta})$ is identical to $(\hat{x}^1_0, \hat{x}^2_0)$ on some support of $(\mu_0)_{\pi^1_j}$ (by the definition of $(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta})$ and assumption (IIb)), we have

$$U_{(\mu_0)_{\pi^1_j}}(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta}) \geq U_{(\mu_0)_{\pi^1_j}}(x^1, \overline{x}^2_{n,\delta}) \tag{27}$$

for every $x^1 \in X^1$. However, as

$$\lim_{n\to\infty} (\mu_n)_{\pi_j^1} \left(\pi_k^2\right) = (\mu_0)_{\pi_j^1} \left(\pi_k^2\right)$$

for $k = 0, ..., L$ and $(\mu_0)_{\pi_j^1} \left(\bigcup_{k=0}^{L} \pi_k^2 \right) = 1$ (by (IIb)), and as u is uniformly bounded, it is obvious that

$$\lim_{n\to\infty} \left[U_{(\mu_n)_{\pi_j^1}}(x^1, x^2) - U_{(\mu_0)_{\pi_j^1}}(x^1, x^2) \right] = 0 \tag{28}$$

uniformly in (x^1, x^2). It follows from (27) and (28) that there exists $N' > 0$ such that (18) holds for every $x^1 \in X^1$, every $0 \le j \le L$, and every $n \ge N'$.

We conclude that (18) (and, similarly, (19)) hold, for every $x^1 \in X^1$ and $x^2 \in X^2$, every $j \ge 0$, and every $n \ge N(\delta)$, where $N(\delta) = \max\{N''(\delta), N'\}$. Thus, indeed, $\left(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta}\right) \in IBE_\varepsilon(\mu_n)$ for each $n \ge N(\delta)$.

It can be assumed w.l.o.g. that $\left\{N\left(\frac{\delta_0}{k}\right)\right\}_{k=1}^{\infty}$ is a strictly increasing sequence of positive integers (recall the definition of δ_0 in (16)). Consider a sequence $\{(\hat{x}^1_n, \hat{x}^2_n)\}_{n=1}^{\infty} \subset X^1 \times X^2$ defined by $(\hat{x}^1_n, \hat{x}^2_n) \equiv \left(\overline{x}^1_{n,\frac{\delta_0}{k}}, \overline{x}^2_{n,\frac{\delta_0}{k}}\right)$ if $N(\frac{\delta_0}{k}) \le n < N(\frac{\delta_0}{k+1})$ for $k \ge 1$.[12] It follows from (17) that $\lim_{n\to\infty}(\hat{x}^1_n, \hat{x}^2_n) = (\hat{x}^1_0, \hat{x}^2_0)$. Furthermore, since for every $0 < \delta \le \delta_0$ and every $n \ge N(\delta)$ it was shown that $\left(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta}\right) \in IBE_\varepsilon(\mu_n)$, it follows that $(\hat{x}^1_n, \hat{x}^2_n) \in IBE_\varepsilon(\mu_n)$ for every $n \ge 1$. $\square$

Remark 5. It is worthwhile to stress the role played in the proof by the zero-sum assumption on the game (without which, as we know, the ALSC of the IBE correspondence does not obtain). Recall that the approximating equilibrium $\left(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta}\right)$ in $G(\mu_n)$, defined in (17), is a state-dependent mixture of the given IBE $(\hat{x}^1_0, \hat{x}^2_0)$ in $G(\mu_0)$ with an IBE $(\widehat{y}^1_n, \widehat{y}^2_n)$ in $G'(\mu_n)$, in which the weight of $(\hat{x}^1_0, \hat{x}^2_0)$ is moving gradually from 1 to 0 as the players' types increase. The proof uses the fact that mixtures of (ε-)optimal strategies remain (ε-)optimal (see the argument establishing (24)), to deduce that $\left(\overline{x}^1_{n,\delta}, \overline{x}^2_{n,\delta}\right)$ is an ε-IBE in $G(\mu_n)$ for the types considered in Case 2. Note

[12] If $n < N(\delta_0)$, let $\left(\hat{x}^1_n, \hat{x}^2_n\right)$ be an arbitrarilty chosen element of $IBE_\varepsilon(\mu_n)$.

that this step could not have been done for non-zero-sum games, where the sets of (ε-)IBE strategies are typically non-convex, and thus mixtures may be far from (ε-)IBE strategies.

Remark 6. Proposition 3 remains valid when instead of (IIb) and (IIc), it is assumed that for every $n \geq 0$ and $j \geq 0$, $(\mu_n)_{\pi^i_j}$ is supported only on $\bigcup_{k=t^i(j)}^{t^i(j)+K} \pi^{-i}_k$. The proof follows the same line as the current one, but is simpler.

When at least one player has a finite number of types, ALSC obtains without assuming either (I) or (II) of Proposition 3. Convergence of common priors implies, in this case, almost uniform convergence of conditional beliefs, and thus the IBE *expected payoff* correspondence is ALSC (for general, not just the zero-sum, games), according to the main result of Kajii and Morris (1994, 1998). The ALSC of the IBE *strategy* correspondence transpires from the proof of that result. For the sake of completeness, however, we state and prove the following proposition, for zero-sum Bayesian games with a finite number of types for at least one player.

Proposition 4. *Assume* $\min(T^1, T^2) < \infty$. *Let* $\{\mu_n\}_{n=0}^{\infty} \subset \Delta(\Omega, \mathcal{F})$ *be such that* $\lim_{n\to\infty} \mu_n = \mu_0$, *and* $(\mu_0)\left(\pi^i_j\right) > 0$ *for every* $i = 1, 2$ *and every* $0 \leq j < \min(T^i+1, \infty)$. *Then the IBE correspondence is ALSC along* $\{\mu_n\}_{n=1}^{\infty}$.

Proof. Assume w.l.o.g. that $T^1 < \infty$. Let $\{\mu_n\}_{n=0}^{\infty} \subset \Delta(\Omega, \mathcal{F})$ be a sequence satisfying all the assumptions, and fix $(\hat{x}^1_0, \hat{x}^2_0) \in IBE(\mu_0)$, $\varepsilon > 0$. Also consider a sequence $\{y^2_n\}_{n=1}^{\infty} \subset X^2$ such that y^2_n is an (ex-ante, or, equivalently, interim) best response of player 2 to the strategy $\hat{x}^1_0$ of player 1 in the game $G(\mu_n)$, for each $n \geq 1$.

Let

$$0 < \delta \leq \delta_0 \equiv \frac{\varepsilon}{16M} \min_{0 \leq j \leq T^1} \mu_0\left(\pi^1_j\right), \tag{29}$$

and let an integer $0 \leq J(\delta) < \min(T^2+1, \infty)$ be such that $\mu(\bigcup_{j=0}^{J(\delta)} \pi^2_j) > 1 - \delta$.

Define a sequence $\left\{\left(\overline{x}_{n,\delta}^{2}\right)\right\}_{n=1}^{\infty} \subset X^2$ as follows: for every $n \geq 1$,

$$\overline{x}_{n,\delta}^{2}(\omega) \equiv \begin{cases} \hat{x}_0^2(\omega), & \text{if } \omega \in \bigcup_{j=0}^{J(\delta)} \pi_j^2; \\ y_n^2(\omega), & \text{otherwise.} \end{cases} \tag{30}$$

The strategy $\hat{x}_0^1$, being in $\mathcal{O}^1(\mu_0)$ by Remarks 2 and 3, is also in $\mathcal{O}^1_{4Md(\mu_n,\mu_0)}(\mu_n)$ by Proposition 2, and thus, for every $0 \leq j \leq T^1$ with $\mu_n\left(\pi_j^1\right) > 0$ and every $x^1 \in X^1$,

$$U_{(\mu_n)_{\pi_j^1}}(\hat{x}_0^1, \hat{x}_0^2) \geq U_{(\mu_n)_{\pi_j^1}}(x^1, \hat{x}_0^2) - \frac{4Md(\mu_n, \mu_0)}{\mu_n\left(\pi_j^1\right)} \tag{31}$$

by (8) in Remark 4. Since clearly, for every such j and x^1

$$\left| U_{(\mu_n)_{\pi_j^1}}(x^1, \hat{x}_0^2) - U_{(\mu_n)_{\pi_j^1}}(x^1, \overline{x}_{n,\delta}^2) \right|$$

$$\leq 2M\left(\mu_n\right)_{\pi_j^1}\left(\Omega \backslash \bigcup_{j=0}^{J(\delta)} \pi_j^2\right) \leq \frac{2M\delta}{\mu_n\left(\pi_j^1\right)},$$

(31) implies that

$$U_{(\mu_n)_{\pi_j^1}}(\hat{x}_0^1, \overline{x}_{n,\delta}^2) \geq U_{(\mu_n)_{\pi_j^1}}(x^1, \overline{x}_{n,\delta}^2) - \frac{4M(d(\mu_n, \mu_0) + \delta)}{\mu_n\left(\pi_j^1\right)}. \tag{32}$$

By our assumptions on $\{\mu_n\}_{n=0}^{\infty}$, there exists $N_1(\delta) > 0$ such that $d(\mu_n, \mu_0) \leq \delta_0$ and $\min_{0 \leq j \leq T^1} \mu_n\left(\pi_j^1\right) \geq \frac{1}{2} \min_{0 \leq j \leq T^1} \mu_0\left(\pi_j^1\right)$ for every $n \geq N_1(\delta)$. The choice of δ in (29) and the inequality (32) guarantee that

$$U_{(\mu_n)_{\pi_j^1}}(\hat{x}_0^1, \overline{x}_{n,\delta}^2) \geq U_{(\mu_n)_{\pi_j^1}}(x^1, \overline{x}_{n,\delta}^2) - \varepsilon \tag{33}$$

for every $0 \leq j \leq T^1$, every $x^1 \in X^1$, and every $n \geq N_1(\delta)$.

Just as in (31), for every $0 \leq j \leq J(\delta)$ with $\mu_n\left(\pi_j^2\right) > 0$ and $x^2 \in X^2$

$$U_{(\mu_n)_{\pi_j^2}}(\hat{x}_0^1, \hat{x}_0^2) \leq U_{(\mu_n)_{\pi_j^2}}(\hat{x}_0^1, x^2) + \frac{4Md(\mu_n, \mu_0)}{\mu_n\left(\pi_j^2\right)}. \tag{34}$$

By our assumptions on $\{\mu_n\}_{n=1}^{\infty}$, there exists $N_2(\delta) > 0$ such that $\min_{0 \le j \le J(\delta)} \mu_n\left(\pi_j^2\right) > 0$ and $\max_{0 \le j \le J(\delta)} \frac{4Md(\mu_n,\mu_0)}{\mu_n\left(\pi_j^2\right)} \le \varepsilon$ for every $n \ge N_2(\delta)$, and therefore for every $0 \le j \le J(\delta)$, every $x^2 \in X^2$, and every $n \ge N_2(\delta)$

$$U_{(\mu_n)_{\pi_j^2}}(\hat{x}_0^1, \hat{x}_0^2) \le U_{(\mu_n)_{\pi_j^2}}(\hat{x}_0^1, x^2) + \varepsilon.$$

But by (30), $\hat{x}_0^1(\omega) = \overline{x}_{n,\delta}^2(\omega)$ for all $\omega \in \pi_j^2$, and thus in fact the following inequality holds:

$$U_{(\mu_n)_{\pi_j^2}}(\hat{x}_0^1, \overline{x}_{n,\delta}^2) \le U_{(\mu_n)_{\pi_j^2}}(\hat{x}_0^1, x^2) + \varepsilon. \tag{35}$$

When $j > J(\delta)$, $\overline{x}_{n,\delta}^2(\omega) = y_n^2(\omega)$ for all $\omega \in \pi_j^2$, and by the definition of $\{y_n^2\}_{n=1}^{\infty}$ as the sequence of best responses to $\hat{x}_0^1$ in the games $\{G(\mu_n)\}_{n=1}^{\infty}$,

$$U_{(\mu_n)_{\pi_j^2}}(\hat{x}_0^1, \overline{x}_{n,\delta}^2) \le U_{(\mu_n)_{\pi_j^2}}(\hat{x}_0^1, x^2)$$

for every $J(\delta) \le j < \min(T^2+1, \infty)$, every $x^2 \in X^2$ and every $n \ge N_2(\delta)$. Thus, (35) in fact holds for every $0 \le j < \min(T^2+1, \infty)$, every $x^2 \in X^2$, and every $n \ge N_2(\delta)$. This fact, coupled with (33), shows that $(\hat{x}_0^1, \overline{x}_{n,\delta}^2) \in IBE_{\varepsilon}(\mu_n)$ for each $n \ge N(\delta) \equiv \max(N_1(\delta), N_2(\delta))$.

Since $\lim_{\delta \to 0} \overline{x}_{n,\delta}^2 = \hat{x}_0^2$, the construction in the last paragraph of the proof of proposition 3 can be repeated to create out of $\left\{\overline{x}_{n,\delta}^2\right\}_{n=1}^{\infty}$ a sequence $\{\hat{x}_n^2\}_{n=1}^{\infty} \subset X^2$ such that $\lim_{n \to \infty}(\hat{x}_0^1, \hat{x}_n^2) = (\hat{x}_0^1, \hat{x}_0^2)$, and $(\hat{x}_0^1, \hat{x}_n^2) \in IBE_{\varepsilon}(\mu_n)$ for every $n \ge 1$. □

References

1. Einy E., O. Haimanko, D. Moreno, and B. Shitovitz. (2008). "Uniform Continuity of the Value of Zero-Sum Games with Differential Information," *Mathematics of Operations Research* **33**, pp. 552-560.

2. Engl, G. (1995). "Lower Hemicontinuity of the Nash Equilibrium Correspondence", *Games and Economic Behavior* **9**, pp. 151-160.

3. Kajii, A., and S. Morris (1994). "Payoff Continuity in Incomplete Information Games and Almost Uniform Convergence of Beliefs," CARESS Working paper # 94-17.

4. Kajii, A., and S. Morris (1997). "The Robustness of Equilibria to Incomplete Information," *Econometrica* **65**, pp. 1283-1309.

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

6. Lehrer, E. and D. Rosenberg (2006). "What restrictions do Bayesian games impose on the value of information?" *Journal of Mathematical Economics* **42**, pp. 343-357.

7. Lehrer, E. and D. Rosenberg (2007). "Evaluating information in zero-sum games with incomplete information on both sides," mimeo.

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

9. Monderer, D. and D. Samet (1989). "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior* **1**, pp. 170-190.

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

11. Rubinstein, A., (1989). "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge", *American Economic Review* **79**, pp. 385-391.

12. Shiryaev, A. N. (1996). *Probability*, Second Edition. Graduate Texts in Mathematics, Springer-Verlag.

13. Sion, M. (1958). "On General Minimax Theorem," *Pacific Journal of Mathematics* **8**, pp. 171-176.

14. Sorin, S. (2002). *A First Course on Zero-Sum Repeated Games.* Mathematiques and Applications, Vol. 37. Springer-Verlag.