

Convergence and approximation results for non-cooperative Bayesian games: learning theorems^{*}

Leonidas C. Koutsougeras and Nicholas C. Yannellis

Department of Economics, University of Illinois at Urbana-Champaign, Champaign, Illinois 61820, USA

Received: July 19, 1992; revised version February 5, 1993

Summary. Let T denote a continuous time horizon and $\{G^t: t \in T\}$ be a net (generalized sequence) of Bayesian games. We show that: (i) if $\{x^t: t \in T\}$ is a net of Bayesian Nash Equilibrium (BNE) strategies for G^t , we can extract a subsequence which converges to a limit full information BNE strategy for a one shot limit full information Bayesian game. (ii) If $\{x^t: t \in T\}$ is a net of approximate or ε_t -BNE strategies for the game G^t we can still extract a subsequence which converges to the one shot limit full information equilibrium BNE strategy. (iii) Given a limit full information BNE strategy of a one shot limit full information Bayesian game, we can find a net of ε_t -BNE strategies $\{x^t: t \in T\}$ in $\{G^t: t \in T\}$ which converges to the limit full information BNE strategy of the one shot game.

1. Introduction

The definition of a Bayesian game and the notion of Bayesian rationality (or Bayesian Nash equilibrium) [see for instance Aumann (1987, p. 6)] are given as follows:

Let $(\Omega, \mathcal{F}, \mu)$ be a probability measure space and Y be a linear topological space. A Bayesian game $G = \{(X_i, u_i, \mathcal{F}_i, \mu): i = 1, 2, \dots, n\}$ is a set of quadruples where

- (1) $X_i: \Omega \rightarrow 2^Y$ is the *random strategy correspondence* of player¹ i ,
- (2) $u_i: \Omega \times \prod_{j=1}^n Y_j \rightarrow \mathbf{R}$ is the *random payoff function* of player i ,
- (3) \mathcal{F}_i is the *private information* of player i , which is a partition of (Ω, \mathcal{F}) , and

^{*} We wish to thank Larry Blume, Mark Feldman, Jim Jordan, Charlie Kahn, Stefan Krasa, Gregory Michalopoulos, Wayne Shafer, Bart Taub, and Anne Villamil for several useful discussions. The financial support of the University of Illinois at Urbana-Champaign Campus Research Board is gratefully acknowledged.

¹ Notice that in Aumann (1987) X_i is a fixed set and doesn't depend on the states of nature. The present framework is more general and reduces to that of Aumann simply by setting for each $\omega \in \Omega$, the set $X_i(\omega)$ to be equal to a constant subset Z_i of Y . Also note that 2^Y denotes the set of all non-empty subsets of Y .

- (4) μ is a probability measure on (Ω, \mathcal{F}) denoting the *common prior* of each player.

A *Bayesian Nash equilibrium* (BNE) for G is a function $x: \Omega \rightarrow \prod_{i=1}^n Y_i$ such that each $x_i: \Omega \rightarrow Y_i$ is \mathcal{F}_i -measurable $x_i(\omega) \in X_i(\omega)$ μ -a.e. and for all i , $\int u_i(\omega, x(\omega)) d\mu(\omega) \geq \int u_i(\omega, x_1(\omega), \dots, x_{i-1}(\omega), y_i, x_{i+1}(\omega), \dots, x_n(\omega)) d\mu(\omega)$ for any \mathcal{F}_i -measurable function $y_i: \Omega \rightarrow Y_i$, $y_i(\omega) \in X_i(\omega)$ μ -a.e.

Consider now the above game in a dynamic framework. Specifically, let T be the set $\{1, 2, \dots\}$ denoting the *time horizon*. Denote by $\sigma(u_i, X_i)$ the σ -algebra that the random payoff function and the random strategy set of player i generate. This is the initial information of player i . At any given point in time t in T the private information set of player i is defined as:

$$(1.1) \quad \mathcal{F}_i^t = \sigma(u_i, X_i, x^{t-1}, x^{t-2}, \dots)$$

where x^{t-1}, x^{t-2}, \dots are past period Bayesian Nash equilibrium strategies. In other words the private information of a player i at any given point in time consists of his/her initial information $\sigma(u_i, X_i)$ together with the information that BNE strategies generated in all previous periods, i.e., $t-1, t-2, \dots$. Note that in this setting in period $t+1$ the private information set of player i will be, $\mathcal{F}_i^{t+1} = \mathcal{F}_i^t \vee \sigma(x^t)$, (where $\sigma(x^t)$ is the information that the BNE strategy x_t generated at period t and $\mathcal{F}_i^t \vee \sigma(x^t)$ denotes the "join," i.e., the smallest σ -algebra containing \mathcal{F}_i^t and $\sigma(x^t)$). Hence, for each player i and each time period t we have that:

$$\mathcal{F}_i^t \subseteq \mathcal{F}_i^{t+1} \subseteq \mathcal{F}_i^{t+2} \subseteq \dots$$

The above expression represents a *learning process* for player i and it generates a sequence of Bayesian games $\{G^t: t \in T\}$ defined as above where the private information set of each player is given by (1.1). In other words, in period t each player's strategy is based on the initial information as well as the information that BNE strategies generated in the previous periods. In this setting agents behave myopically, i.e., they do not form expectations over the entire future horizon but only for the current period, i.e., their expected payoff is based on the current period private information. Since the private information set of each player becomes finer from period to period, the expected payoff of each player is changing from period to period as a result of the new acquired information. Note that in this scenario the learning process for a player is a direct consequence of observing the BNE strategies from period to period and refining his/her private information. In this framework, clearly the information that the equilibrium strategy generates at a given time t in T , will effect the equilibrium outcome in subsequent periods, e.g., $t+1, t+2, \dots$. Let us now denote the one shot *full information* game by $\bar{G} = \{(X_i, u_i, \bar{\mathcal{F}}_i, \mu): i = 1, 2, \dots, n\}$ where $\bar{\mathcal{F}}_i$ is the pooled information of player i over the entire horizon, i.e., $\bar{\mathcal{F}}_i = \bigvee_{t=1}^{\infty} \mathcal{F}_i^t$. Since any BNE strategy for each player i in the game \bar{G} has the property that it is $\bigvee_{t=1}^{\infty} \mathcal{F}_i^t$ -measurable, we call such a Bayesian Nash equilibrium strategy as a *full information* BNE strategy.

The basic questions that this paper addresses are the following:

(i) If $\{G^t: t = 1, 2, \dots\}$ is a sequence of Bayesian games and x^t is a sequence of BNE strategies for the game G^t , can we extract a subsequence which converges to a full information BNE strategy for the game \bar{G} ? In other words, will the learning process described above eventually lead to a full information BNE strategy?

(ii) If $\{G^t: t = 1, 2, \dots\}$ is a sequence of Bayesian games and x^t is a sequence of approximate or ε_t -BNE strategies for the game G^t can we still extract a subsequence which converges to a full information BNE strategy for the game \bar{G} ? In other words, can we obtain the counterpart of question (i) for the case of an approximate or ε_t -BNE which may be viewed respectively, as bounded rational learning will converge to the full information BNE.

(iii) Given a full information BNE strategy for the full information game \bar{G} can we find a sequence of approximate or ε_t -BNE strategies x^t in G^t which converges to the full information BNE strategy? Roughly speaking, can we approximate (or reach) a full information BNE strategy by a sequence of ε_t -BNE strategies? Alternatively, given a full information BNE strategy can it be reached by a path of plays with bounded rational players (i.e., players find “nearly” optimal responses)?

We provide a positive answer to the above questions. Note that roughly speaking (ii) and (iii) may be viewed respectively as a kind of upper semicontinuity and lower semicontinuity of the ε_t -BNE correspondence.

It should be pointed out that aspects of question (i) have already been addressed by several authors, notably Feldman (1987), and subsequently by Jordan (1991); Nyarko (1992) [see also the excellent survey of Blume and Easley (1992)], but in a different setting. In particular, we don't require each player's strategy set to be finite, we have a continuum of states, we allow for continuous time, payoff functions need not be linear and the convergence is not in probability as it is the case in the Feldman (1987), Jordan (1991) and Nyarko (1992) papers. The continuous time setting that we allow makes our results interesting to the Finance literature where continuous time models are particularly attractive. To the best of our knowledge, questions (ii) and (iii) are addressed for the first time.

A few comments on the methodology. In view of the fact that we allow for continuous time and a continuum of states, one needs to work with strategies which form a net (generalized sequence) in an infinite dimensional strategy space. The compactness and continuity arguments in this framework are not straightforward and some rather non-elementary functional analytic results seem to be required. We have collected most of the results needed for our proofs in Section 2.

The rest of the paper is organized as follows: Sections 3 and 4 contain the main results of the paper, i.e., convergence and approximation theorems for games with mixed and pure strategy Bayesian Nash equilibrium. The proofs of all our results are given in Section 5.

2. Mathematical preliminaries

2.1 Notation

\mathbb{R}^n denotes the n -fold Cartesian product of the set of real numbers \mathbb{R} .

\mathbb{R}_{++} denotes the strictly positive elements of \mathbb{R} .

$\text{con } A$ denotes the convex hull of the set A .

$\overline{\text{con}} A$ denotes the closed convex hull of the set A .

2^A denotes the set of all nonempty subsets of the set A .

\emptyset denotes the empty set.

$/$ denotes set theoretic subtraction.

If $A \subset X$, where X is a Banach space, $\text{cl } A$ denotes the norm closure of A .

2.2 Definitions

Let T and X be sets. The *graph* of the set-valued function (or correspondence), $\phi: T \rightarrow 2^X$ is denoted by $G_\phi = \{(t, y) \in T \times X : y \in \phi(t)\}$. Let now (T, τ, μ) be a complete, finite measure space, and X be a separable Banach space. The correspondence $\phi: T \rightarrow 2^X$ is said to have a *measurable graph* if $G_\phi \in \tau \otimes \beta(X)$, where $\beta(X)$ denotes the Borel σ -algebra on X and \otimes denotes product σ -algebra. The correspondence $\phi: T \rightarrow 2^X$ is said to be *lower measurable* if for every open subset V of X , the set $\{t \in T : \phi(t) \cap V \neq \emptyset\}$ is an element of τ . Recall [see Debreu (1966), p. 359] that if $\phi: T \rightarrow 2^X$ has a measurable graph, then ϕ is lower measurable. Furthermore, if $\phi(\cdot)$ is closed valued and lower measurable then $\phi: T \rightarrow 2^X$ has a measurable graph. A result of Aumann says that if (T, τ, μ) is a complete, finite measure space, X is a separable metric space and $\phi: T \rightarrow 2^X$ is a nonempty valued correspondence having a measurable graph, then $\phi(\cdot)$ admits a *measurable selection*, i.e., there exists a measurable function $f: T \rightarrow X$ such that $f(t) \in \phi(t)$ μ -a.e.

We now define the notion of a Bochner integrable function. We will follow closely Diestel-Uhl (1977). Let (T, τ, μ) be a finite measure space and X be a Banach space. A function $f: T \rightarrow X$ is called *simple* if there exist x_1, x_2, \dots, x_n in X and $\alpha_1, \alpha_2, \dots, \alpha_n$ in τ such that $f = \sum_{i=1}^n x_i \chi_{\alpha_i}$, where $\chi_{\alpha_i}(t) = 1$ if $t \in \alpha_i$ and $\chi_{\alpha_i}(t) = 0$ if $t \notin \alpha_i$. A function $f: T \rightarrow X$ is said to be μ -*measurable* if there exists a sequence of simple functions $f_n: T \rightarrow X$ such that $\lim_{n \rightarrow \infty} \|f_n(t) - f(t)\| = 0$ for almost all $t \in T$. A μ -measurable function $f: T \rightarrow X$ is said to be *Bochner integrable* if there exists a sequence of simple functions $\{f_n: n = 1, 2, \dots\}$ such that

$$\lim_{n \rightarrow \infty} \int_T \|f_n(t) - f(t)\| d\mu(t) = 0.$$

In this case we define for each $E \in \tau$ the integral to be $\int_E f(t) d\mu(t) = \lim_{n \rightarrow \infty} \int_E f_n(t) d\mu(t)$.

It can be shown [see Diestel-Uhl (1977), Theorem 2, p. 45] that, if $\phi: T \rightarrow X$ is a μ -measurable function then f is Bochner integrable if and only if $\int_T \|f(t)\| d\mu(t) < \infty$. It is important to note that the *Dominated Convergence Theorem* holds for Bochner integrable functions, in particular, if $f_n: T \rightarrow X$ ($n = 1, 2, \dots$) is a sequence of Bochner integrable functions such that $\lim_{n \rightarrow \infty} f_n(t) = f(t)$ μ -a.e., and $\|f_n(t)\| \leq g(t)$ μ -a.e., where

$g \in L_1(\mu, \mathbb{R})$, then f is Bochner integrable and $\lim_{n \rightarrow \infty} \int_T \|f_n(t) - f(t)\| d\mu(t) = 0$.

We denote by $L_1(\mu, X)$ the space of equivalence classes of X -valued Bochner integrable functions $x: T \rightarrow X$ normed by

$$\|x\| = \int_T \|x(t)\| d\mu(t).$$

It is a standard result that normed by the functional $\|\cdot\|$ above, $L_1(\mu, X)$ becomes a Banach space [see Diestel-Uhl (1977), p. 50]. We denote by S_ϕ the set of all selections from $\phi: T \rightarrow 2^X$ that belong to the space $L_1(\mu, X)$, i.e.,

$$S_\phi = \{x \in L_1(\mu, X) : x(t) \in \phi(t) \text{ } \mu\text{-a.e.}\},$$

i.e., S_ϕ is the set of all Bochner integrable selections from $\phi(\cdot)$. Using the above set and following Aumann (1965) we can define the integral of the correspondence $\phi: T \rightarrow 2^X$ as follows:

$$\int_T \phi(t) d\mu(t) = \{\int_T x(t) d\mu(t) : x \in S_\phi\}.$$

We will denote the above integral by $\int \phi$. Recall that the correspondence $\phi: T \rightarrow 2^X$ is said to be *integrably bounded* if there exists a map $h \in L_1(\mu, \mathbb{R})$ such that $\sup\{\|x\| : x \in \phi(t)\} \leq h(t)$ μ -a.e. Moreover, note that if T is a complete measure space, X is a separable Banach space and $\phi: T \rightarrow 2^X$ is an integrably bounded, nonempty valued correspondence having a measurable graph, then by the Aumann measurable selection theorem we can conclude that S_ϕ is nonempty and therefore $\int_T \phi(t) d\mu(t)$ is nonempty as well. If in addition to the fact that $\phi: T \rightarrow 2^X$ is integrably bounded and nonempty it is also weakly compact and convex valued then by *Diestel's Theorem* [see for instance Yannelis (1991), Theorem 3.1] we can conclude that S_ϕ is weakly compact in $L_1(\mu, X)$.

We close this section by defining the notion of a martingale and stating the martingale convergence theorem. Let I be a directed set and let $\{\mathcal{F}_i : i \in I\}$ be a monotone increasing net of sub- σ -fields of τ (i.e., $\mathcal{F}_{i_1} \subseteq \mathcal{F}_{i_2}$ for $i_1 \leq i_2$, i_1, i_2 in I). A net $\{x_i : i \in I\}$ in $L_1(\mu, X)$ is a *martingale* if

$$E(x_i | \mathcal{F}_{i_1}) = x_{i_1} \quad \text{for all } i \geq i_1.$$

We will denote the above martingale by $\{x_i, \mathcal{F}_i\}_{i \in I}$. The proof of the following *martingale convergence theorem* can be found in Diestel-Uhl (1977, p. 126). A martingale $\{x_i, \mathcal{F}_i\}_{i \in I}$ in $L_1(\mu, X)$ converges in the $L_1(\mu, X)$ -norm if and only if there exists x in $L_1(\mu, X)$ such that $E(x | \mathcal{F}_i) = x_i$ for all $i \in I$. Recall [see for instance Diestel-Uhl (1977, p. 129)] that if the martingale $\{x_i, \mathcal{F}_i\}_{i \in I}$ converges in the $L_1(\mu, X)$ -norm to $x \in L_1(\mu, X)$, it also converges almost everywhere, i.e., $\lim_{i \rightarrow \infty} x_i = x$ almost everywhere.

3. Convergence and approximation theorems for mixed strategy Bayesian Nash equilibria

3.1 Bayesian games and Bayesian Nash Equilibria

Let $(\Omega, \mathcal{F}, \mu)$ be a complete, probability measure space, and Y be a separable Banach space. As previously, a *Bayesian game* $G = \{(X_i, u_i, \mathcal{F}_i, \mu) : i = 1, 2, \dots, n\}$ is a set of quadruples where:

- (1) $X_i: \Omega \rightarrow 2^Y$ is the *random strategy correspondence* of player i ,
- (2) $u_i: \Omega \times \prod_{j=1}^n Y_j \rightarrow \mathbb{R}$ is the *random payoff function* of player i ,

- (3) \mathcal{F}_i is the *private information* set of player i , where \mathcal{F}_i is a (finite, measurable) partition of (Ω, \mathcal{F}) ,
 (4) μ is a probability measure on (Ω, \mathcal{F}) denoting the *common prior* of each player.²

Denote by L_{X_i} the set of all Bochner integrable and \mathcal{F}_i -measurable selections from the set-valued function $X_i: \Omega \rightarrow 2^Y$, i.e., $L_{X_i} = \{x_i \in L_1(\mu, Y): x_i: \Omega \rightarrow Y \text{ is } \mathcal{F}_i\text{-measurable and } x_i(\omega) \in X_i(\omega) \text{ } \mu\text{-a.e.}\}$. Let $L_x = \prod_{i=1}^n L_{X_i}$, and $L_{\tilde{x}_i} = \prod_{j \neq i} L_{X_j}$. Denote the elements of $L_{\tilde{x}_i}$ by \tilde{x}_i . The *expected payoff* of player i is a function $v_i: L_x \rightarrow \mathbb{R}$ defined by

$$v_i(x) = \int u_i(\omega, x(\omega)) d\mu(\omega).$$

The strategy $x^* \in L_x$ is said to be a *Bayesian Nash equilibrium* for the Bayesian game $G = \{(X_i, u_i, \mathcal{F}_i, \mu): i = 1, 2, \dots, n\}$ if for all i , $(i = 1, 2, \dots, n)$

$$v_i(x^*) = \max_{x_i \in L_{X_i}} v_i(x_i, \tilde{x}_i^*).$$

Note that $x^* \in L_x$ implies that each x_i^* is \mathcal{F}_i -measurable and therefore the vector $x^* = (x_1^*, x_2^*, \dots, x_n^*)$ is $\bigvee_{i=1}^n \mathcal{F}_i$ -measurable.

Suppose that $G = \{(X_i, u_i, \mathcal{F}_i, \mu): i = 1, 2, \dots, n\}$ satisfies the following assumptions for all i , $(i = 1, 2, \dots, n)$.

(a.3.1) $X_i: \Omega \rightarrow 2^Y$ is integrably bounded, weakly compact, convex, nonempty valued and \mathcal{F}_i -lower measurable correspondence,

(a.3.2) for each $\omega \in \Omega$, $u_i(\omega, \cdot)$ is weakly continuous on $\prod_{j=1}^n Y_j$, and for each

$$x \in \prod_{j=1}^n Y_j, u_i(\cdot, x) \text{ is } \mathcal{F}_i\text{-measurable,}$$

(a.3.3) for each $\omega \in \Omega$, and each $\tilde{x}_i \in \tilde{Y}_i = \prod_{j \neq i} Y_j$, $u_i(\omega, x_i, \tilde{x}_i)$ is a concave function of x_i on Y_i ,

(a.3.4) u_i is integrably bounded.

It was shown in Yannelis-Rustichini (1991) that under the assumptions (a.3.1)–(a.3.4)

² One may allow for different priors as follows: Let $q_i: \Omega \rightarrow \mathbb{R}_{++}$ be a Radon-Nikodym derivative (density function) denoting the prior of agent i . For each $i = 1, \dots, n$, denote by $E_i(\omega)$ the event in \mathcal{F}_i containing the realized state of nature $\omega \in \Omega$ and suppose that $\int_{t \in E_i(\omega)} q_i(t) d\mu(t) > 0$. Given $E_i(\omega) \in \mathcal{F}_i$ define the *conditional expected utility* of agent i as follows:

$$\int_{t \in E_i(\omega)} u_i(t, x_i(t)) q_i(t | E_i(\omega)) d\mu(t),$$

where

$$q_i(t | E_i(\omega)) = \begin{cases} 0 & \text{if } t \notin E_i(\omega) \\ \frac{q_i(t)}{\int_{t \in E_i(\omega)} q_i(t) d\mu(t)} & \text{if } t \in E_i(\omega). \end{cases}$$

All the results of the paper remain valid if we use the above conditional expected utility formulation. However, for the simplicity of the exposition we do not do so.

the game $G = \{(X_i, u_i, \mathcal{F}_i, \mu): i = 1, 2, \dots, n\}$ has a (mixed strategy) Bayesian Nash equilibrium.

3.2 Learning

Let T be any directed set (countable or uncountable) denoting the *time horizon*. Denote by $\sigma(u_i, X_i)$ the σ -algebra that the random payoff function and random strategy set of player i generate. This is the initial information of agent i . However, the private information set of player i at time $t \in T$, is not only $\sigma(u_i, X_i)$ but also the information that past period Bayesian Nash equilibrium strategies (denoted by x^{t^1} for $t^1 < t$, t, t^1 in T) have generated. Hence the private information set of player i at time t is defined as:

$$(3.1) \quad \mathcal{F}_i^t = \sigma(u_i, X_i, \{x^{t^1}: t^1 < t\}).$$

The private information set of player i in period $t^0 > t$ will be $\mathcal{F}_i^{t^0} = \mathcal{F}_i^t \vee \sigma(\{x^{t'}: t \leq t' < t^0\})$, and consequently, for each player i and each time period t' we have that:

$$(3.2) \quad \mathcal{F}_i^{t'} \subseteq \mathcal{F}_i^{t^0} \quad \text{for } t' \leq t^0, t', t^0 \text{ in } T.$$

The expression (3.2) represents a *learning process* for player i .

A learning process generates a net of Bayesian games $\{G^t: t \in T\}$, where $G^t = \{(X_i, u_i, \mathcal{F}_i^t, \mu): i = 1, 2, \dots, n\}$. As previously,

- (1) $X_i: \Omega \rightarrow 2^Y$ is the *random strategy set* of player i ,
- (2) $u_i: \Omega \times \prod_{j=1}^n Y_j \rightarrow \mathbb{R}$ is the *random payoff function* of player i ,
- (3) \mathcal{F}_i^t is the *private information set* of player i at time t , given by the expression (3.1), and
- (4) μ is the *common prior* of each player.

Let $L_{X_i^t} = \{x_i \in L_1(\mu, Y): x_i \text{ is } \mathcal{F}_i^t\text{-measurable and } x_i(\omega) \in X_i(\omega) \text{ } \mu\text{-a.e.}\}$. Set $L_{X^t} = \prod_{i=1}^n L_{X_i^t}$. Define the *expected utility* of player i , $v_i: L_{X^t} \rightarrow \mathbb{R}$ by

$$v_i(x) = \int u_i(\omega, x(\omega)) d\mu(\omega).$$

The interpretation of the above dynamic game is as follows: In period t each player's strategy is based on the initial information (i.e., $\sigma(u_i, X_i)$) and the information that all Bayesian Nash equilibrium strategies have generated in the previous periods. Note that each player doesn't form expectations for future periods but only for the current period t . (Recall that the expected payoff is based on the current period private information.) However, since the private information set of each player is increasing it follows that his/her expected payoff is changing from period to period as a result of the increased information. Denote by BNE (G^t) the *set of all Bayesian Nash equilibrium strategies* for G^t at time t .

3.3 The full information Bayesian game

Let $\bar{\mathcal{F}}_i$ be the *pooled information set* of player i over the entire time horizon T , i.e., $\bar{\mathcal{F}}_i = \bigvee_{t \in T} \mathcal{F}_i^t$. (Note that $\bigvee_{t \in T} \mathcal{F}_i^t$ denotes the "join" of the \mathcal{F}_i^t , i.e., the minimal

σ -algebra containing all \mathcal{F}_i^t .) The Bayesian game $\bar{G} = \{(\bar{X}_i, u_i, \bar{\mathcal{F}}_i, \mu): i = 1, 2, \dots, n\}$ where u_i, μ are defined as previously but \bar{X}_i is now $\bar{\mathcal{F}}_i$ -measurable, is called the limit full information Bayesian game. Notice that the special case of symmetric information, i.e., $\bar{\mathcal{F}}_i = \bar{\mathcal{F}}_j$ for $i \neq j$ ($i, j = 1, 2, \dots, n$) follows from our setting. However, we don't need to insist on symmetric information in the Bayesian game \bar{G} . Denote by $\text{BNE}(\bar{G})$ the set of all Bayesian Nash equilibrium strategies for the game \bar{G} , i.e., $x^* \in \text{BNE}(\bar{G})$ implies that $x^* \in L_{\bar{X}} = \prod_{i=1}^n L_{\bar{X}_i}$ and³ for all i ,

$$v_i(x^*) = \max_{x_i \in L_{X_i}} v_i(x_i, \tilde{x}_i^*).$$

3.4 Remarks

3.4.1: As we discussed above, the expression (3.2) represents a learning process. However, this is not the only way to generate a learning process. We now follow some ideas of McKelvey-Page (1986) and define an alternative learning process as follows:

Let $f: \prod_{i=1}^n L_{X_i} \rightarrow R$ be interpreted as a publicly observable statistic. For each $t \in T$ if $x^t \in \text{BNE}(G^t)$ then the statistic f conveys information specified as $\sigma(f(x^t))$, i.e., the smallest σ -field with respect to which $f(x^t)$ is measurable. This information refines the information already available to each player at time t , i.e., up to time $t \in T$, the information available to each player i , is given by

$$(3.4.1) \quad \mathcal{F}_i^t = \sigma(u_i, X_i, \{f(x^{t^1}): t^1 < t\}),$$

where $\sigma(\{f(x^{t^1}): t^1 < t\})$ is the publicly known information up to (but not including) period t . Note that $\sigma(f(x^t))$ is not contained in \mathcal{F}_i^t , but for $t^0 > t$, $\mathcal{F}_i^{t^0} = \mathcal{F}_i^t \vee \sigma(\{f(x^{t'}): t \leq t' < t^0\})$ and $\sigma(f(x^t))$ is contained in $\mathcal{F}_i^{t^0}$. Hence we have the following learning process, for each player i ,

$$(3.4.2) \quad \mathcal{F}_i^{t'} \subseteq \mathcal{F}_i^{t^0} \quad \text{for } t' \leq t^0, \quad t', t^0 \text{ in } T.$$

It should be noted that for each $i \in I$, $\mathcal{F}_i^{t^1} = \mathcal{F}_i^t$ for all $t \geq t^1$ if and only if for at least one $s \in T$, $s > t^1$, $\sigma(\{f(x^s): s > t^1\}) \subseteq \sigma(\{f(x^r): r < t^1\})$. In this case the publicly observed statistic does not convey any new information after period t^1 and the learning process of each player stops. As a result $x^{t^1} \in \text{BNE}(G^t)$ for all periods $t \geq t^1$, since the information structure does not change after period t^1 .

In other words, the evolution of the information partitions [i.e., the learning process represented by the expression (3.4.2)] stops if and only if the information conveyed by the publicly observable statistic is common knowledge to all players. Now if we define $\bar{\mathcal{F}}_i = \bigvee_{t \in T} \mathcal{F}_i^t$ where each \mathcal{F}_i^t is as in (3.4.1), then for each i ,

$$(3.4.3) \quad \bar{\mathcal{F}}_i = \sigma(u_i, X_i, \{f(x^t): t \in T\}).$$

Thus, in the limit the information conveyed by the statistic f is common knowledge.

³ Obviously, $L_{\bar{X}_i}$ is the set $\{x_i \in L_1(\mu, Y): \text{each } x_i \text{ is } \bar{\mathcal{F}}_i\text{-measurable and } x_i(\omega) \in \bar{X}_i(\omega) \mu\text{-a.e.}\}$.

That is, players have enough information to predict the statistic. In this respect, the strategy vector for the limiting game has the flavor of a Rational Expectations Equilibrium [see also McKelvey-Page (1986, p. 122)].

Notice that even in the full information Bayesian game, agents may have different information partitions. However, the asymmetry in information partitions expressed by (3.4.3) persists as long as the statistic f does not convey any new information. Hence, the term full information Bayesian game refers to the fact that players have learned all the information that can be conveyed by the statistic f which aggregates the information available to each player. For example, if

$\sigma(f(x^t)) = \bigvee_{i=1}^n \mathcal{F}_i^t$ then x^t is a *fully revealing* BNE for the game G^t , and in this case

for $t' \geq t$ the information partition of each agent will be the same and therefore in the full information Bayesian game \bar{G} players will have symmetric information.

3.4.2: If in the full information game \bar{G} , we assume that Ω is a finite set and that each player has a finite set of strategies X_i , then the concept of BNE or Bayesian rationality coincides with the correlated equilibrium [Aumann (1987)]. In this specific setting we can conclude that Bayesian rational learning will lead to correlated equilibrium⁴.

3.5 Theorem

Theorem 3.5.1: Let $\{G^t: t \in T\}$ be a net of Bayesian games satisfying (a.3.1)–(a.3.4) and let $\{x^t: t \in T\}$ be a net in $\text{BNE}(G^t)$. Then we can extract a sequence $\{x^{n_t}: n_t = 1, 2, \dots\}$ from the net $\{x^t: t \in T\}$ such that x^{n_t} converges weakly to $x^* \in \text{BNE}(\bar{G})$.

3.6 Approximate Bayesian Nash equilibrium

Given an $\varepsilon > 0$, the strategy $x^* \in L_x$ is said to be an *approximate* or ε -BNE for the Bayesian game $G = \{(X_i, u_i, \mathcal{F}_i, \mu): i = 1, 2, \dots, n\}$ if for all i ,

$$v_i(x^*) \geq v_i(x_i, \tilde{x}_i^*) - \varepsilon \quad \text{for all } x_i \in L_{X_i}.$$

This concept of an approximate BNE has been widely discussed in the literature [e.g., Radner (1980)]. The justification of this notion is that it may be too costly to find the exact optimal response than a “nearly” optimal one [Radner (1980, p. 153)]. The latter may be viewed as a kind of bounded rationality. Denote by $\text{BNE}_\varepsilon(G)$ the set of all approximate or ε -BNE strategies for the game G . We now obtain the counterpart of Theorem 3.5.1 for the case of an approximate BNE.

Theorem 3.6.1: Let $\{G^t: t \in T\}$ be a net of Bayesian games satisfying (a.3.1)–(a.3.4) and let $\{x^t: t \in T\}$ be a net in $\text{BNE}_{\varepsilon_t}(G^t)$, where $\varepsilon_t \downarrow 0$. Then we can extract a sequence $\{x^{n_t}: n_t = 1, 2, \dots\}$ from the net $\{x^t: t \in T\}$ such that x^{n_t} converges weakly to $x^* \in \text{BNE}(\bar{G})$.

⁴ Nyarko (1992) has also shown that Bayesian learning leads to correlated equilibrium. His model, however, is different than ours. Moreover, Kalai-Lehrer (1993) have proved a non-myopic version of our Theorem 3.6.1 (see below). Their model is different than ours, specifically, they assume that the time horizon is discrete, each player's strategy set is finite and payoff functions are linear. Obviously, these assumptions are stronger than ours (compare with (a.3.1)–(a.3.4)).

Theorem 3.6.2: Let $\{G^t: t \in T\}$ be a net of Bayesian games satisfying (a.3.1)–(a.3.4) and let $x^* \in \text{BNE}(\bar{G})$. Then for each net $\{\varepsilon_t: t \in T\}$ bounded away from zero, there exists a net of strategies $\{x^t: t \in T\}$ in $\text{BNE}_{\varepsilon_t}(G^t)$ such that x^t converges (in the $L_1(\mu, Y)$ -norm) to x^* .

Theorem 3.6.1 indicates that even approximate BNE will converge to the (exact) full information BNE. This theorem gives as a corollary Theorem 3.5.1.⁵ It is important to note that Theorem 3.6.2 has an interesting interpretation. In particular, it shows that an (exact) full information BNE can be achieved by a path of plays by agents who have bounded rationality, i.e., this path of plays constitutes an approximate BNE for each period.

4. Pure strategy Bayesian Nash equilibrium convergence and approximation theorems

In this section will derive the counterparts of Theorem 3.5.1 and Theorems 3.6.1 and 3.6.2 for pure strategy Bayesian Nash equilibria for the game G . Previously the strategy set of each player i , i.e., X_i , was assumed to be a set-valued function from Ω to Y . We now set $Y = \mathbb{R}^m$. Now each \mathcal{F}_i will be a sub- σ -algebra of (Ω, \mathcal{F}) and the restriction of μ to \mathcal{F}_i will be still denoted by μ . We denote by $\text{ext } X_i$ the *extreme points* of X_i . (Recall that in this setting pure strategies are identified with extreme points.)

Formally, a *pure strategy Bayesian Nash equilibrium* for the game $G = \{X_i, u_i, \mathcal{F}_i, \mu\}: i = 1, 2, \dots, n\}$ is an $x^*: \Omega \rightarrow \prod_{i=1}^n Y_i$ such that each x_i^* is \mathcal{F}_i -measurable $x_i^*(\omega) \in \text{ext } X_i(\omega)$ μ -a.e. and for all i , ($i = 1, 2, \dots, n$),

$$\int u_i(\omega, x^*(\omega)) d\mu(\omega) \geq \int u_i(\omega, x_i(\omega), \tilde{x}_i^*(\omega)) d\mu(\omega)$$

for any \mathcal{F}_i -measurable function $x_i: \Omega \rightarrow Y_i$, $x_i(\omega) \in \text{ext } X_i(\omega)$ μ -a.e.

The following assumptions guarantee the existence of a pure strategy Bayesian Nash equilibrium for the game G [see Yannelis-Rustichini (1991)].

- (a.4.1) For each i , $(\Omega, \mathcal{F}_i, \mu)$ is a complete, atomless probability measure space,
- (a.4.2) For each i , $X_i: \Omega \rightarrow 2^{\mathbb{R}^m}$ is an integrably bounded, compact, convex, nonempty valued and \mathcal{F}_i -lower measurable correspondence,
- (a.4.3) For each i , and each $\omega \in \Omega$, $u_i(\omega, \cdot)$ is linear and continuous on $\prod_{j=1}^n Y_j = \mathbb{R}^{mn}$ and it is also integrably bounded.

As in Section 3.2 we can similarly recast the idea of learning. The only difference now is that the private information set of agent i at time t depends on the information he/she has acquired from past period Bayesian Nash equilibrium pure strategies still denoted by x^{t1}, x^{t2}, \dots . Everything else remains the same.

⁵ Despite the fact that Theorem 3.5.1 can be obtained as a corollary of Theorem 3.6.1 we have tried to separate the cases that Bayesian rational learning leads to the full information BNE (Theorem 3.5.1) and that ε -Bayesian rational (or bounded rational) learning leads to the full information BNE (Theorem 3.6.1), because we believe that both cases are of interest. Perhaps, the experimental work will indicate which kind of equilibrium is more appropriate. Some recent work in this direction is reported in the papers of El-Gamal-McKelvey-Palfrey (1992) and Rustichini-Villamil (1992 and 1992a).

Denote by $\text{PBNE}(G^t)$ the set of all *pure strategy Bayesian Nash equilibria* for G^t at time t , and by $\text{PBNE}(\bar{G})$ the set of all pure strategies Bayesian Nash equilibria for the full information Bayesian Nash equilibrium game \bar{G} .

Theorem 4.1: Let $\{G^t: t \in T\}$ be a net of Bayesian games satisfying (a.4.1)–(a.4.3) and let $\{x^t: t \in T\}$ be a net in $\text{PBNE}(G^t)$. Then we can extract a sequence $\{x^{n_i}: n_i = 1, 2, \dots\}$ from the net $\{x^t: t \in T\}$ such that x^{n_i} converges weakly to $x^* \in \text{PBNE}(\bar{G})$.

We conclude this section by mentioning that the counterparts of Theorems 3.6.1 and 3.6.2 for the case of pure strategies can be readily obtained. Since we have outlined the proofs of Theorems 3.6.1 and 3.6.2 in Section 6, we leave the proof for the case of pure strategies to the reader, in order to avoid repetition.

5. Proof of the theorems

We begin with a few observations. Note that since each u_i is integrably bounded and weakly continuous it follows from the Lebesgue dominated convergence theorem that each v_i is weakly continuous [Fact 4.2 in Yannelis-Rustichini (1991)]. Moreover, since each u_i is concave in the i^{th} coordinate so is v_i . For each i , ($i = 1, 2, \dots, n$) define $P_i: L_X \rightarrow 2^{L_{X_i}}$ by

$$P_i(x) = \{y_i \in L_{X_i}: v_i(y_i, \tilde{x}_i) > v_i(x)\}.$$

It follows from the weak continuity of v_i that P_i has a weakly open graph (i.e., the set $G_{P_i} = \{(x, y_i) \in L_X \times L_{X_i}: y_i \in P_i(x)\}$ is weakly open in $L_X \times L_{X_i}$). Also from the concavity of v_i in the i^{th} coordinate it follows that P_i is convex valued. Since each X_i is \mathcal{F}_i -lower measurable and compact valued, it has a measurable graph. By the Aumann measurable selection theorem there exists an \mathcal{F}_i -measurable function $f_i: \Omega \rightarrow Y$ such that $f_i(\omega) \in X_i(\omega)$ μ -a.e. Since X_i is integrably bounded $f_i \in L_1(\mu, Y)$ and therefore each set L_{X_i} is nonempty and so is $\prod_{i=1}^n L_{X_i} = L_X$. Clearly each L_{X_i}

is convex and therefore L_X is convex as well. By Diestel's theorem [see Theorem 3.1 in Yannelis (1991)] L_{X_i} is weakly compact subset of $L_1(\mu, Y)$. Hence the set L_X is weakly compact, convex and nonempty. Finally notice that for $t^1 \geq t^2$, (t^1, t^2 in T) we have that $L_{X_i^{t^2}} \subseteq L_{X_i^{t^1}}$, i.e., as information increases the strategy set of each player expands.

With all these preliminary observations out of the way we can begin the proof of Theorem 3.5.1.

5.1 Proof of Theorem 3.5.1

Let $\{x^t: t \in T\}$ be a net in $\text{BNE}(G^t)$, (i.e., $x^t \in L_{X_i^t}$ and $P_i(x^t) \cap L_{X_i^t} = \emptyset$ for all i , or equivalently $x^t \in L_{X_i^t}$ and for all i , $v_i(x^t) = \max_{y_i \in L_{X_i^t}} v_i(y_i, \tilde{x}_i^t)$). For simplicity let us

denote the net $\{x^t: t \in T\}$ by B . As we observed above the set $L_{X_i^t}$ is weakly compact and nonempty. Since for each $t \in T$, $x^t \in L_{X_i^t}$ and $L_{X_i^t}$ is weakly compact it follows that the weak closure of the set B denoted by $\omega\text{-cl } B$, is weakly compact. By the Eberlein-Smulian Theorem [Dunford-Schwartz (1958, p. 430)], $\omega\text{-cl } B$ is weakly

sequentially compact. Clearly the weak limit of x^t , denoted by x^* , belongs to ω -cl B . From Whitley's theorem [Aliprantis-Burkinshaw (1985, Lemma 10.12, p. 155)] we know that if $x^* \in \omega$ -cl B then there exists a sequence $\{x^{t_m} \in L_{X^{t_m}} : m = 1, 2, \dots\}$ such that x^{t_m} converges weakly to x^* . (For notational convenience we denote the above sequence simply by $\{x^m : m = 1, 2, \dots\}$.) Since $\mathcal{F}_i^t \subseteq \mathcal{F}_i^{t^1}$ for $t^1 \geq t$ we have that $L_{X^t} \subseteq L_{\bar{X}}$, and since each x^m is in $L_{X^m} \subseteq L_{\bar{X}}$, it follows that $x^* \in L_{\bar{X}}$. Hence, for each i , x_i^* is an \mathcal{F}_i -measurable selection from \bar{X}_i . To complete the proof we must show that

$$(5.1.1) \quad P_i(x^*) \cap L_{\bar{X}_i} = \phi \quad \text{for all } i.$$

Suppose otherwise, i.e., for some i , $P_i(x^*) \cap L_{\bar{X}_i} \neq \phi$. Choose $y_i \in P_i(x^*) \cap L_{\bar{X}_i}$, then $v_i(y_i, \tilde{x}_i^*) > v_i(x^*)$. Let

$$(5.1.2) \quad \varepsilon = v_i(y_i, \tilde{x}_i^*) - v_i(x^*) > 0.$$

For each m , ($m = 1, 2, \dots$) set $y_i^m = E[y_i | \mathcal{F}_i^m] \in L_{X_i^m}$. Note that

$$\begin{aligned} E[y_i | \mathcal{F}_i^m] &= E[E[y_i | \mathcal{F}_i^{m+1}] | \mathcal{F}_i^m] \\ &= E[y_i^{m+1} | \mathcal{F}_i^m]. \end{aligned}$$

Hence, $\{y_i^m, \mathcal{F}_i^m\}_{m=1}^\infty$ is a martingale in $L_{X_i^m} \subset L_1(\mu, Y)$ and by the martingale convergence theorem, y_i^m converges (in the $L_1(\mu, Y)$ norm) and thus weakly to y_i . It follows (recall that v_i is weakly continuous) that we can choose m_1 large enough so that for $m \geq m_1$ we have

$$|v_i(y_i, \tilde{x}_i^*) - v_i(y_i^m, \tilde{x}_i^m)| < \varepsilon/2 \quad \text{and} \quad |v_i(x^m) - v_i(x^*)| < \varepsilon/2.$$

Thus

$$\begin{aligned} |v_i(y_i, \tilde{x}_i^*) - v_i(y_i^m, \tilde{x}_i^m) + v_i(x^m) - v_i(x^*)| &\leq |v_i(y_i, \tilde{x}_i^*) - v_i(y_i^m, \tilde{x}_i^m)| \\ &\quad + |v_i(x^m) - v_i(x^*)| < \varepsilon/2 + \varepsilon/2. \end{aligned}$$

Then in view of (5.1.2) we have

$$v_i(y_i, \tilde{x}_i^*) - v_i(y_i^m, \tilde{x}_i^m) + v_i(x^m) - v_i(x^*) < v_i(y_i, \tilde{x}_i^*) - v_i(x_i^*)$$

and by rearranging we obtain

$$v_i(y_i^m, \tilde{x}_i^m) > v_i(x^m) \quad \text{for all } m \geq m_1,$$

a contradiction to the fact that x^m lies in $\text{BNE}(G^m)$. The above contradiction establishes the validity of (5.1.1), i.e., $P_i(x^*) \cap L_{\bar{X}_i} = \phi$ for all i , or equivalently $v_i(x^*) = \max_{y_i \in L_{X_i}} v_i(y_i, \tilde{x}_i^*)$ for all i . This completes the proof of Theorem 3.5.1.

5.2 Proof of Theorem 3.6.1

For each i , ($i = 1, 2, \dots, n$) and each $\varepsilon > 0$, define $P_i^\varepsilon: L_X \rightarrow 2^{L_{X^t}}$ by

$$P_i^\varepsilon(x) = \{y_i \in L_{X_i} : v_i(y_i, \tilde{x}_i) > v_i(x) + \varepsilon\}.$$

Let $\{x^t : t \in T\}$ be a net in $\text{BNE}_{e_t}(G^t)$, i.e., $x^t \in L_{X^t}$ and for all i , $P_i^{e_t}(x^t) \cap L_{X_i^t} = \phi$.

Adopting the argument in the proof of Theorem 3.5.1 we can extract a sequence $\{x^{n_t}: n_t = 1, 2, \dots\}$ from the net $\{x^t: t \in T\}$ such that x^{n_t} converges weakly to $x^* \in L_{\bar{X}}$. To complete the proof we must show that $x^* \in \text{BNE}(\bar{G})$, i.e.,

$$P_i(x^*) \cap L_{\bar{X}_i} = \phi \quad \text{for all } i.$$

Suppose otherwise, i.e., for at least one i , $P_i(x^*) \cap L_{\bar{X}_i} \neq \phi$. Let $y_i \in P_i(x^*) \cap L_{\bar{X}_i}$, then $v_i(y_i, \tilde{x}_i^*) - v_i(x^*) > 0$. Set

$$(5.2.1) \quad v_i(y_i, \tilde{x}_i^*) - v_i(x^*) = \delta.$$

For each m , ($m = 1, 2, \dots$) set $y_i^m = E[y_i | \mathcal{F}_i^m] = E[E[y_i | \mathcal{F}_i^{m+1}] | \mathcal{F}_i^m] = E[y_i^{m+1} | \mathcal{F}_i^m]$. Thus, $\{y_i^m, \mathcal{F}_i^m\}_{m=1}^\infty$ is a martingale in $L_{X_i^m} \subset L_1(\mu, Y)$ and by the martingale convergence theorem y_i^m converges (in the $L_1(\mu, Y)$ -norm and therefore weakly) to y_i . By the weak continuity of v_i we can choose m_1 large enough so that for $m \geq m_1$ we have that

$$|v_i(y_i, \tilde{x}_i^*) - v_i(y_i^m, \tilde{x}_i^m)| < \frac{\delta - \varepsilon_m}{2}$$

and

$$|v_i(x^m) - v_i(x^*)| < \frac{\delta - \varepsilon_m}{2}.$$

Recall that $\varepsilon_t \downarrow 0$ and so does the sequence ε_m . Thus,

$$\begin{aligned} |v_i(y_i, \tilde{x}_i^*) - v_i(y_i^m, \tilde{x}_i^m) + v_i(x^m) - v_i(x^*)| &\leq |v_i(y_i, \tilde{x}_i^*) - v_i(y_i^m, \tilde{x}_i^m)| \\ &\quad + |v_i(x^m) - v_i(x^*)| < \delta - \varepsilon_m. \end{aligned}$$

In view of (5.2.1) we have

$$v_i(y_i, \tilde{x}_i^*) - v_i(y_i^m, \tilde{x}_i^m) + v_i(x^m) - v_i(x^*) < v_i(y_i, \tilde{x}_i^*) - v_i(x^*) - \varepsilon_m$$

and by rearranging we obtain that

$$v_i(y_i^m, \tilde{x}_i^m) > v_i(x^m) + \varepsilon_m \quad \text{for all } m \geq m_1,$$

a contradiction to the fact that $x^m \in \text{BNE}_{\varepsilon_m}(G^m)$. Hence, we can conclude that $P_i(x^*) \cap L_{\bar{X}_i} = \phi$ for all i , i.e., $x^* \in \text{BNE}(\bar{G})$. This completes the proof of Theorem 3.6.1.

5.3 Proof of Theorem 3.6.2

Let x be in $\text{BNE}(\bar{G})$, i.e., $x \in L_{\bar{X}}$ and $P_i(x) \cap L_{\bar{X}_i} = \phi$ for all i . We will construct a net $\{x^t: t \in T\}$ in $\text{BNE}_{\varepsilon_t}(G^t)$ such that x^t converges (in the $L_1(\mu, Y)$ -norm) to x . For each i , ($i = 1, 2, \dots, n$) and each t in T set $x_i^t = E[x_i | \mathcal{F}_i^t]$. Note that

$$\begin{aligned} E[x_i | \mathcal{F}_i^t] &= E[E[x_i | \mathcal{F}_i^{t_1}] | \mathcal{F}_i^t] \quad \text{for } t_1 \geq t \\ &= E[x_i^{t_1} | \mathcal{F}_i^t] \quad \text{for } t_1 \geq t. \end{aligned}$$

Hence $\{x_i^t, \mathcal{F}_i^t\}_{t \in T}$ is a martingale in $L_{X_i} \subset L_1(\mu, Y)$ and by the martingale convergence theorem x_i^t converges in the $L_1(\mu, Y)$ -norm and hence weakly to x_i . To complete the proof we must show that $x^t = (x_1^t, \dots, x_n^t)$ lies in $\text{BNE}_{\varepsilon_t}(G^t)$ for any

net $\{\varepsilon_t: t \in T\}$ bounded away from zero such that $\varepsilon_t \downarrow \varepsilon$, ($\varepsilon > 0$). For each i , define the set $K_i = \{t \in T: P_i^{\varepsilon_t}(x^t) \cap L_{x^t} \neq \emptyset\}$. Notice that for $t \in T \setminus \bigcup_{i=1}^n K_i$ we have that $x \in \text{BNE}_{\varepsilon_t}(G^t)$. If for each i , K_i is a finite or empty set there is nothing to prove. Hence to complete the proof we need to show that for each i , K_i cannot be infinite. To this end suppose that for some i , K_i is infinite. Then the net $\{y_i^t: t \in K_i\}$ has the property that,

$$(5.3.1) \quad v_i(y_i^t, \tilde{x}_i^t) > v_i(x^t) + \varepsilon_t.$$

Note that $y_i^t \in L_{x^t} \subset L_{\tilde{x}_i}$ and the set $L_{\tilde{x}_i}$ is weakly compact. Hence, we can find a subnet still denoted by $\{y_i^t: t \in K_i\}$ such that y_i^t converges weakly to $y_i \in L_{\tilde{x}_i}$. By the weak continuity of v_i taking weak limits in the inequality (5.4.1) we obtain that $v_i(y_i, \tilde{x}_i) \geq v_i(x) + \varepsilon$. Thus, we can conclude that $y_i \in P_i(x) \cap L_{\tilde{x}_i}$ for some player i , a contradiction to the supposition that x lies in $\text{BNE}(\bar{G})$.

5.4 Proof of Theorem 4.1

We begin with some preparatory observations and facts. Denote by $L_{\text{ext } X_i}$ the set $\{x_i \in L_1(\mu, \mathbb{R}^m): x_i: \Omega \rightarrow \mathbb{R}^m \text{ is } \mathcal{F}_i\text{-measurable and } x_i(\omega) \in \text{ext } X_i(\omega) \mu\text{-a.e.}\}$. Define the mapping $\psi: L_1(\mu, \mathbb{R}^m) \rightarrow \mathbb{R}^m$ by $\psi(z) = \int z(\omega) d\mu(\omega)$. Denote the integral of the set-valued function $\text{ext } X_i: \Omega \rightarrow 2^{\mathbb{R}^m}$ by $\int \text{ext } X_i$ which in turn is equal to $\psi(L_{\text{ext } X_i}) = \{\psi(z): z \in L_{\text{ext } X_i}\}$. Set $\int \text{ext } X = \prod_{i=1}^n \int \text{ext } X_i$. Let $\int \text{ext } \tilde{X}_i = \prod_{i \neq i} \int \text{ext } X_j$ and denote the points of $\int \text{ext } \tilde{X}_i$ by \tilde{x}_i . Since by (a.4.3) for each $\omega \in \Omega$, $u_i(\omega, \cdot)$ is linear on $\prod_{j=1}^n Y_j = \mathbb{R}^{mn}$ the domain of the expected utility of each agent i , i.e., $g_i(x) = \int u_i(\omega, x(\omega)) d\mu(\omega)$ is now $\int X = \prod_{i=1}^n \int X_i$. However, we will show that the set $\int X$ is equal to $\int \text{ext } X$. To this end first note that since each X_i is compact and convex valued, by the standard Krein-Milman-Minkowski theorem we have that

$$(5.4.1) \quad \overline{\text{con}}(\text{ext } X_i(\omega)) = X_i(\omega) \mu\text{-a.e.}$$

By Theorem 5.3 in Himmelberg (1975) $\text{ext } X_i(\cdot)$ is lower measurable and so is $\overline{\text{con}} \text{ext } X_i(\cdot)$. Integrating (5.4.1) we obtain:

$$(5.4.2) \quad \int \overline{\text{con}}(\text{ext } X_i) = \int X_i.$$

Since by assumption the measure space $(\Omega, \mathcal{F}_i, \mu)$ is atomless by Theorem 3 of Aumann (1965) we have that:

$$(5.4.3) \quad \int \overline{\text{con}}(\text{ext } X_i) = \int \text{ext } X_i.$$

Combining (5.4.2) and (5.4.3) we have that for each i , $\int \text{ext } X_i = \int X_i$ and we conclude that $\int X = \int \text{ext } X$.

As in the proof of Theorem 3.5.1 for each i , define $P_i: \int \text{ext } X \rightarrow 2^{\int \text{ext } X_i}$ by

$$(5.4.4) \quad P_i(x) = \{y_i \in \int \text{ext } X_i: g_i(y_i, \tilde{x}_i) > g_i(x)\}.$$

It follows from the Lebesgue dominated convergence theorem that g_i is (norm)

continuous on $\int \text{ext } X$ (recall that by (a.4.3) for each $\omega \in \Omega$ $u_i(\omega, \cdot)$ is continuous and integrably bounded). Moreover, since by assumption for each $\omega \in \Omega$, $u_i(\omega, \cdot)$ is linear so is g_i . But since g_i is norm continuous and linear it is also weakly continuous [Dunford-Schwartz (Theorem 15, p. 422)]. Hence, P_i has a weakly open graph in $\int \text{ext } X_i \times \int \text{ext } X$. Finally note that by claim 5.1 in Yannelis-Rustichini (1991), $\int \text{ext } X$ is weakly compact and nonempty.

With all these preliminaries out of the way we can now complete the proof of Theorem 4.1. As a matter of fact the argument from now on is identical to the one adopted for the proof of Theorem 3.5.1 and we will only outline it. Let $\{x^t: t \in T\}$ be a net in $\text{PBNE}(G^t)$, i.e., $x^t \in \int \text{ext } X^t$ and for all i , $g_i(x^t) = \max_{y_i \in \int \text{ext } X_i^t} g_i(y_i, \hat{x}_i^t)$. As noted above, $\int \text{ext } X^t$ is weakly compact. Hence, by adopting the argument of Theorem 3.5.1 we can extract a sequence $\{z^m: m = 1, 2, \dots\}$ from the net $\{x^t: t \in T\}$ such that z^m converges weakly to $x^* \in \int \text{ext } X^t \subset \int \text{ext } \bar{X}$, i.e., for each i , x_i^* is \mathcal{F}_i -measurable selection from $\text{ext } \bar{X}_i$. To complete the proof one must show that $P_i(x^*) \cap \int \text{ext } \bar{X}_i = \emptyset$ for all i . An identical argument with that used in Theorem 3.5.1 can be now adopted to complete the proof of Theorem 4.1.

References

- Aliprantis, C. D., Burkinshaw, O.: Positive operators. New York: Academic Press 1985
- Aumann, R. J.: Integrals of set-valued functions. *J. Math. Anal. Appl.* **12**, 1–12 (1965)
- Aumann, R. J.: Correlated equilibria as an expression of Bayesian rationality. *Econometrica* **55**, 1–18 (1987)
- Blume, L., Easley, D.: Rational expectations and rational learning. mimeo 1992
- Debreu, G.: Integration of correspondences. *Proc. Fifth Berkeley Symp. Stat. Prob.* **II**, 351–372 (1967)
- Diestel, J., Uhl, J.: Vector measures. *Math. Survey* **15**. Providence, R.I.: A.M.S. 1977
- Dunford, N., Schwartz, J. T.: Positive operators. Part I. New York: Interscience 1958
- El-Gamal, M., McKelvey, R. D., Palfrey, T. R.: Learning in experimental games. *Econ. Theory* **4**, 901–922 (1994)
- Feldman, M.: An example of convergence to rational expectations with heterogeneous beliefs. *Int. Econ. Rev.* **28**, 635–650 (1987)
- Himmelberg, C. J.: Measurable relations. *Fundam. Math.* **LXXXVII**, 53–72 (1975)
- Jordan, J. S.: Bayesian learning in normal form games. mimeo 1991
- Kalai, E., Lehrer, E.: Rational learning leads to Nash equilibrium. *Econometrica* 1993 (to appear)
- Koutsougeras, L., Yannelis, N. C.: Learning in differential information economies with cooperative solution concepts: cores and values. mimeo 1992
- McKelvey, R. D., Page, T.: Common knowledge consensus, and aggregate information. *Econometrica* **54**, 109–127 (1986)
- Nyarko, Y.: Bayesian learning leads to correlated equilibria in normal form games. *Econ. Theory* **4**, 821–841 (1994)
- Radner, R.: Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *J. Econ. Theory* **22**, 136–154 (1980)
- Rustichini, A., Villamil, A. P.: Intertemporal pricing in retail markets: equilibrium price stickiness. mimeo 1992
- Rustichini, A., Villamil, A. P.: Learning stochastic demand in experimental markets. mimeo 1992a
- Yannelis, N. C.: Integration of Banach-valued correspondences. In: M. A. Khan and N. C. Yannelis (eds.) *Equilibrium theory in infinite dimensional spaces*. Berlin, Heidelberg, New York: Springer 1991
- Yannelis, N. C., Rustichini, A.: On the existence of correlated equilibria. In: M. A. Khan and N. C. Yannelis (eds.) *Equilibrium theory in infinite dimensional spaces*, Berlin, Heidelberg, New York: Springer 1991