Macroeconomic Dynamics, **1**, 1997, 568–587. Printed in the United States of America.

LEARNING IN BAYESIAN GAMES BY BOUNDED RATIONAL PLAYERS I.

TAESUNG KIM
Seoul National University

NICHOLAS C. YANNELIS
University of Illinois at Urbana–Champaign

We study learning in Bayesian games (or games with differential information) with an arbitrary number of bounded rational players, i.e., players who choose approximate best response strategies [approximate Bayesian Nash Equilibrium (BNE) strategies] and who also are allowed to be completely irrational in some states of the world. We show that bounded rational players by repetition can reach a limit full information BNE outcome. We also prove the converse, i.e., given a limit full information BNE outcome, we can construct a sequence of bounded rational plays that converges to the limit full information BNE outcome.

Keywords: Bayesian Game, Bayesian Nash Equilibrium, Approximate BNE, Bounded Rational Play, Continuum of Players

1. INTRODUCTION

We study learning in Bayesian games (or games with differential information) by an arbitrary number of bounded rational players. The description of the Bayesian game is as follows. Let $(\Omega, \mathcal{F}, \mu)$ be a probability measure space interpreted as the set of states of the world. A Bayesian game is $G = \{(X_\alpha, u_\alpha, \mathcal{F}_\alpha, q_\alpha) : \alpha \in A\}$, where

1. A denotes a set of *players*,
2. $X_\alpha(\omega)$ is the set of *actions* available to agent α when the state is ω,
3. $u_\alpha(\omega, \cdot) : \prod_{a \in A} X_a(\omega) \to R$ is the state-dependent *utility function* of agent α,
4. $\mathcal{F}_\alpha$ denotes the *private information* of agent α,
5. q_α is the *prior* of agent α [q_α is a density function or Radon Nikodym derivative, i.e., $\int_{\omega \in \Omega} q_\alpha(\omega)\, d\mu(\omega) = 1$].

The strategy $\tilde{x}_\alpha$ of player α is an $\mathcal{F}_\alpha$-measurable function, where $\tilde{x}_\alpha(\omega) \in X_\alpha(\omega)$ μ-a.e. Denote by $E_\alpha(\omega)$ the element in $\mathcal{F}_\alpha$ containing ω. Given $E_\alpha(\omega)$,

With great sorrow, I announce the tragic death of my friend, Taesung Kim. He died March 13, 1997, from a heart attack at the age of 38. He is survived by his wife, Young-Eun, and his two daughters, Na-Yeon and Na-Gyung. Taesung was a very capable, sincere, and polite person. A gentleman like him will be definitely missed. Address correspondence to: Nicholas C. Yannelis, Department of Economics, University of Illinois at Urbana–Champaign, Champaign, IL 61820, USA; e-mail: nyanneli@uiuc.edu.

 1365-1005/97 $9.00 + .10

define player α's conditional expected utility $v_\alpha(\omega, \tilde{x})$ for the strategy profile $\tilde{x} = (\tilde{x}_\alpha)_{\alpha \in A}$ as

$$v_\alpha(\omega, \tilde{x}) = \int_{\omega' \in E_\alpha(\omega)} u_\alpha(\omega, \tilde{x}(\omega')) q_\alpha(\omega' \mid E_\alpha(\omega))\, d\mu(\omega'),$$

where $q_\alpha(\omega' \mid E_\alpha(\omega))$ denotes the conditional probability of ω', given $E_\alpha(\omega)$.

An *ε-Bayesian Nash equilibrium* (BNE) for G is a strategy profile $\tilde{x} = (\tilde{x}_\alpha)_{\alpha \in A}$ such that for all $\alpha \in A$, for all $\omega \in D$ with $\mu(D) \geq 1 - \varepsilon$,

$$v_\alpha(\omega, \tilde{x}) \geq v_\alpha(\omega, \tilde{x}_{-\alpha}, \tilde{y}_\alpha) - \varepsilon,$$

for all strategy y_α. The ε-Bayesian Nash equilibrium captures the idea of bounded rationality of players in the following sense: First, in come states whose measure is less than ε, players can be completely irrational. Second, even in states where agents are supposed to be rational, each player's strategy is required to be ε-best response against others' strategies. In that sense, we can call the *ε-BNE* outcome as *bounded rational play* and the player who chooses *ε-BNE* strategies *bounded rational*. We denote by $BNE_\varepsilon(G)$ the set of all *ε-Bayesian Nash equilibrium* for G.

Consider now the above game in a dynamic learning context. For each $t = 1, 2, \ldots$, let $G^t = \{(X_\alpha, u_\alpha, \mathcal{F}_\alpha^t, q_\alpha) \colon \alpha \in A\}$ denote the Bayesian game in period t, where

$$\mathcal{F}_\alpha^t = \mathcal{F}_\alpha^{t-1} \bigvee \sigma(\tilde{x}^{t-1}), \tag{1}$$

and $\sigma(\tilde{x}^{t-1})$ denotes the σ-algebra that the ε-Bayesian Nash equilibrium strategy $\tilde{x}^{t-1}$ of the previous period generated. $\mathcal{F}_\alpha^{t-1} \bigvee \sigma(\tilde{x}^{t-1})$ denotes the join, i.e., the smallest σ-algebra containing $\mathcal{F}_\alpha^{t-1}$ and $\sigma(\tilde{x}^{t-1})$. Consequently, for each player α and period t,

$$\mathcal{F}_\alpha^t \subseteq \mathcal{F}_\alpha^{t+1} \subseteq \mathcal{F}_\alpha^{t+2} \subseteq \cdots.$$

The above expression represents a learning process for player α. This learning process generates a sequence of Bayesian games $\{G^t : t \in T\}$ where $G^t = \{(X_\alpha, u_\alpha, \mathcal{F}_\alpha^t, q_\alpha) : \alpha \in A\}$.

Let us now define a *limit full information Bayesian game* as

$$\bar{G} = \{(X_\alpha, u_\alpha, \bar{\mathcal{F}}_\alpha, q_\alpha) \colon \alpha \in A\},$$

where

$$\bar{\mathcal{F}}_\alpha = \bigvee_{t \in T} \mathcal{F}_\alpha^t.$$

Our main objectives in this paper are to address the following questions:

(i) Can bounded rational players by repetition reach a limit full information BNE outcome? That is, if the sequence $\tilde{x}^t \in BNE_{\varepsilon t}(G^t)$ with $\varepsilon_t \to 0$, can we extract a subsequence whose limit $\tilde{x}^* \in BNE(\bar{G})$?

(ii) If $\tilde{x}^* \in BNE(\bar{G})$ (i.e., $\tilde{x}^*$ is a limit full information BNE strategy for $\bar{G}$), can we construct a sequence of bounded rational plays $\tilde{x}^t$ [i.e., $\tilde{x}^t \in BNE_\varepsilon(G^t)$] such that $\tilde{x}^t$ converges to $\tilde{x}^*$?

Our main results show that the answers to both questions are affirmative whether the number of players is finite or continuum.

We now compare our results to the ones in the Bayesian learning literature [e.g., Feldman (1987), Jordan (1991), Kalai and Lehrer (1993), Koutsougeras and Yannelis (1994), and Nyarko (1996)]. Jordan (1991) studied the iterative Bayesian learning process for the finite-player, finite-strategy normal form games. He proved that if players with common priors about others' payoffs use Bayesian learning in each period, myopic Bayesian Nash equilibrium of each period converges to the complete information Nash equilibrium of the normal form game. Kalai and Lehrer (1993) also showed that even nonmyopic Bayesian Nash equilibrium of the repeated game converges to the ε-Nash equilibrium of the complete information normal form game if players use Bayesian updating in each period.

The framework we use in this paper is more general than those of the above authors in the following sense: Players do not necessarily have common priors and they are bounded rational during the learning process. Because our framework is quite general, it may be the case that, in the limit, incomplete information still prevails. In other words, it could be the case that

$$\bar{\mathcal{F}}_\alpha(\omega) \subset \left(\bigvee_{\alpha \in A} \mathcal{F}_\alpha \right)(\omega),$$

where $\bar{\mathcal{F}}_\alpha(\omega)$ and $(\bigvee_{\alpha \in A} \mathcal{F}_\alpha)(\omega)$ are the smallest elements containing ω in $\bar{\mathcal{F}}_\alpha$ and $\bigvee_{\alpha \in A} \mathcal{F}_\alpha$, respectively. Therefore, our equilibrium in the limit is not necessarily the Nash equilibrium of the complete information game. It is a Bayesian Nash equilibrium of the limit full information game. However, if the learning through Bayesian Nash equilibrium strategy of each period reaches the complete information in the limit, i.e.,

$$\bar{\mathcal{F}}_\alpha(\omega) \supset \left(\bigvee_{\alpha \in A} \mathcal{F}_\alpha \right)(\omega),$$

then one of our convergence results (Theorem 1) is comparable to those of Jordan (1991) and Kalai and Lehrer (1993) because a Bayesian Nash equilibrium of the complete information game is just a Nash equilibrium.

Also note that we extend and generalize the previous results of Koutsougeras and Yannelis (1994) in two different directions. First, we have interim, not ex ante, expected utility function; and second, we allow for a continuum of players. Both improvements necessitate new arguments.

Finally, the results of Feldman (1987), Jordan (1991), Kalai and Lehrer (1993), and Nyarko (1996) are in the spirit of convergence to a limit full information *BNE* [question (i), above]. However, we also obtain the converse result, i.e., given a *BNE* strategy $\tilde{x}^*$ of a limit full information game $\bar{G}$, we can construct a sequence of ε-*BNE* strategy $\tilde{x}^t$ for a game G^t that converges to $\tilde{x}^*$. In other words, we can always approximate the limit full information *BNE* strategy by the repetition of bounded rational play. This may be interpreted as a kind of stability of the *BNE* outcome.

The rest of the paper proceeds as follows: Section 2 contains notation and definitions. Section 3 describes the Bayesian game with a finite number of players, and the main notions used in the paper are rigorously defined. In Section 4, we prove that bounded rational players by repetition will reach the limit full information *BNE* outcome and, conversely, given a limit full information *BNE* outcome, we can construct a sequence of bounded rational play that converges to the limit full information *BNE* outcome. Section 5 addresses the same questions as those in Section 4 but in a Bayesian game with a continuum of players. Finally, we have collected the technical results in the Appendix.

2. NOTATION AND DEFINITIONS

We begin with some notation and definitions.

Let 2^X denote the set of all non-empty subsets of the set X. If X and Y are sets, the *graph* of the set-valued function (or correspondence) $\varphi : X \to 2^Y$ is $G_\varphi = \{(x, y) \in X \times Y : y \in \varphi(x)\}$.

Let $(\Omega, \mathcal{F}, \mu)$ be a complete, finite measure space and Y be a separable Banach space. The correspondence $\varphi : \Omega \to 2^Y$ is said to have a *measurable graph* if

$$G_\varphi \in \mathcal{F} \bigotimes \mathcal{B}(Y),$$

where $\mathcal{B}(Y)$ denotes the Borel σ-algebra on Y and $\bigotimes$ denotes the product σ-algebra. The measurable function $f : \Omega \to Y$ is called a *measurable selection* of $\varphi : \Omega \to 2^Y$ if

$$f(\omega) \in \varphi(\omega) \quad \text{for} \quad \mu\text{-a.e. } \omega.$$

Let $(\Omega, \mathcal{F}, \mu)$ be a finite measure space and Y be a Banach space. Following Diestel and Uhl (1977), the function $f : \Omega \to Y$ is called *simple* if there exist $y_1, y_2, \ldots, y_n$ in Y and $E_1, E_2, \ldots, E_n$ in $\mathcal{F}$ such that

$$f = \sum_{i=1}^{n} y_i \chi E_i,$$

where $\chi_{E_i}(\omega) = 1$ if $\omega \in E_i$ and $\chi_{E_i}(\omega) = 0$ if $\omega \notin E_i$. A function $f : \Omega \to Y$ is called $\mathcal{F}$-*measurable* if there exists a sequence of simple functions $f_n : \Omega \to Y$ such that

$$\lim_{n\to\infty} \| f_n(\omega) - f(\omega)\| = 0 \quad \text{for} \quad \mu\text{-a.e. } \omega.$$

An $\mathcal{F}$-measurable function $f : \Omega \to Y$ is said to be *Bochner integrable* if there exists a sequence of simple functions $\{f_n : n = 1, 2, \ldots\}$ such that

$$\lim_{n\to\infty} \int_{\omega\in\Omega} \| f_n(\omega) - f(\omega)\| \, d\mu(\omega) = 0.$$

In this case, for each $E \in \mathcal{F}$, we define the *integral* of f, denoted by $\int_E f(\omega)\, d\mu(\omega)$, as

$$\lim_{n\to\infty} \int_E f_n(\omega)\, d\mu(\omega).$$

It can be shown [see Diestel and Uhl (1977, Theorem 2, p. 45)] that if $f : \Omega \to Y$ is an $\mathcal{F}$-measurable function, then f is Bochner integrable if and only if

$$\int_\Omega \| f(\omega)\|\, d\mu(\omega) < \infty.$$

It turns out to be important in our paper that the *Dominated Convergence Theorem* holds for Bochner integrable functions. In particular, if $f_n : \Omega \to Y (n = 1, 2, \ldots)$ is a sequence of Bochner integrable functions such that for μ-a.e. ω,

$$\lim_{n\to\infty} f_n(\omega) = f(\omega), \quad \text{and} \quad \| f_n(\omega)\| \leq g(\omega),$$

where $g : \Omega \to R$ is an integrable function, then f is Bochner integrable and

$$\lim_{n\to\infty} \int_{\omega\in\Omega} \| f_n(\omega) - f(\omega)\|\, d\mu(\omega) = 0.$$

The space of equivalence classes of Y-valued Bochner integrable functions $y : \Omega \to Y$, normed by

$$\|y\| = \int_\Omega \|y(\omega)\|\, d\mu(\omega),$$

is denoted by $L_1(\mu, Y)$. It is a standard result that, normed by the functional $\|\cdot\|$ above, $L_1(\mu, Y)$ becomes a Banach space [see Diestel and Uhl (1977, p. 50)].

A Banach space Y has the *Radon-Nikodym property* (RNP) with respect to the measure space $(T, \mathcal{T}, \nu)$ if for each ν-continuous vector measure $G: \mathcal{T} \to Y$ of bounded variation, there exists some $g \in L_1(\nu, Y)$ such that for all $E \in \mathcal{T}$,

$$G(E) = \int_E g(t)\, d\nu(t).$$

It is a standard result [Diestel and Uhl (1977)] that if Y^* (the norm dual of Y) has the RNP with respect to $(T, \mathcal{T}, \nu)$, then

$$(L_1(\nu, Y))^* = L_\infty(\nu, Y^*).$$

A correspondence $\varphi : \Omega \to 2^Y$ is said to be *integrably bounded* if there exists a function $h \in L_1(\mu, R)$ such that

$$\sup\{\|y\| : y \in \varphi(\omega)\} \leq h(\omega) \quad \text{for} \quad \mu\text{-a.e.}\, \omega.$$

If $\{A_n : n = 1, 2, \ldots\}$ is a sequence of non-empty subsets of a Banach space Y, we denote by $Ls A_n$ the set of its *limit superior* points, i.e.,

$$Ls A_n = \left\{ x \in Y : x = \lim_{k\to\infty} x_{n_k}, x_{n_k} \in A_{n_k} \quad \text{for} \quad k = 1, 2, \ldots \right\}.$$

3. THE GAME WITH DIFFERENTIAL INFORMATION

3.1. Bayesian Nash Equilibrium

Let $(\Omega, \mathcal{F}, \mu)$ be a complete, finite, separable measure space, where Ω denotes the set of states of the world and σ-algebra $\mathcal{F}$, the set of events. Let Y be a separable Banach space and A be a set of agents (which is any finite or infinite set).

A *Bayesian game* (or a *game with differential information*) is $G = \{(X_\alpha, u_\alpha, \mathcal{F}_\alpha, q_\alpha) : \alpha \in A\}$, where

1. $X_\alpha : \Omega \to 2^Y$ is the *action set-valued function* of agent t, where $X_\alpha(\omega)$ is the set of actions available to agent α when the state is ω;
2. for each $\omega \in \Omega$, $u_\alpha(\omega, \cdot) : \prod_{a\in A} X_a(\omega) \to R$ is the *utility function* of agent α, which can depend on the states;
3. $\mathcal{F}_\alpha$ is a sub σ-algebra of $\mathcal{F}$ which denotes the *private information* of agent α;
4. $q_\alpha : \Omega \to R_{++}$ is the *prior* of agent α [q_α is a density function or Radon Nikodym derivative, i.e., $\int_{\omega\in\Omega} q_\alpha(\omega)\, d\mu(\omega) = 1$].

Let L_{X_α} denote the set of all Bochner integrable and $\mathcal{F}_\alpha$-measurable selections from the action set-valued function $X_\alpha : \Omega \to 2^Y$ of agent α, i.e.,

$$L_{X_\alpha} = \{\tilde{x}_\alpha \in L_1(\mu, Y) : \tilde{x}_\alpha \quad \text{is} \quad \mathcal{F}_\alpha\text{-measurable and} \quad \tilde{x}_\alpha(\omega) \in X_\alpha(\omega)\ \mu\text{-a.e.}\ \omega\}.$$

The typical element of L_{X_α} is denoted as $\tilde{x}_\alpha$, and that of $X_\alpha(\omega)$ as $x_\alpha(\omega)$ (or x_α). Let $L_X = \prod_{a\in A} L_{X_a}$ and $L_{X_{-\alpha}} = \prod_{a\neq\alpha} L_{X_a}$. Given a Bayesian game G, a strategy for agent α is an element $\tilde{x}_\alpha$ in L_{X_α}.

Throughout the paper, we assume that for each $\alpha \in A$, there exists a finite or countable partition $\prod_\alpha$ of Ω. Moreover, the σ-algebra $\mathcal{F}_\alpha$ is generated by $\prod_\alpha$. For each $\omega \in \Omega$, let $E_\alpha(\omega)(\in\prod_\alpha)$ denote the smallest set in $\mathcal{F}_\alpha$ containing ω and we assume that for all α,

$$\int_{\omega'\in E_\alpha(\omega)} q_\alpha(\omega')\, d\mu(\omega') > 0.$$

For each $\omega \in \Omega$, the *conditional (interim) expected utility function* of agent α, $v_\alpha(\omega, \cdot, \cdot) : L_{X_{-\alpha}} \times X_\alpha(\omega) \to R$ is defined as

$$v_\alpha(\omega, \tilde{x}_{-\alpha}, x_\alpha) = \int_{\omega'\in E_\alpha(\omega)} u_\alpha(\omega, \tilde{x}_{-\alpha}(\omega'), x_\alpha) q_\alpha(\omega' \mid E_\alpha(\omega))\, d\mu(\omega'),$$

where

$$q_\alpha(\omega' \mid E_\alpha(\omega)) = \begin{cases} 0 & \text{if} \quad \omega' \notin E_\alpha(\omega) \\ \dfrac{q_\alpha(\omega')}{\int_{\tilde{\omega}\in E_\alpha(\omega)} q_\alpha(\tilde{\omega})\, d\mu(\tilde{\omega})} & \text{if} \quad \omega' \in E_\alpha(\omega). \end{cases}$$

The function $v_\alpha(\omega, \tilde{x}_{-\alpha}, x_\alpha)$ is interpreted as the conditional expected utility of agent α using the action x_α when the realized state is ω and the other agents employ the strategy profile $\tilde{x}_{-\alpha}$, where $\tilde{x}_{-\alpha}$ is an element of $L_{X_{-\alpha}}$.

A *Bayesian Nash equilibrium* for G [denoted by *BNE(G)*] is a strategy profile $\tilde{x}^* \in L_X$ such that for all $\alpha \in A$,

$$v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, \tilde{x}^*_\alpha(\omega)\big) = \max_{y_\alpha \in X_\alpha(\omega)} v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, y_\alpha\big)\ \mu\text{-a.e.}\ \omega.$$

Given an $\varepsilon > 0$, the strategy profile $\tilde{x}^*$ is said to be an *ε-Bayesian Nash equilibrium* for G [denoted by $BNE_\varepsilon(G)$] if there exists $E(\subset \Omega)$ with $\mu(E) < \varepsilon$ such that for all $\alpha \in A$, for all $\omega \in \Omega/E$,

$$v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, \tilde{x}^*_\alpha(\omega)\big) \geq v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, y_\alpha\big) - \varepsilon,$$

for all $y_\alpha \in X_\alpha(\omega)$.

3.2. Learning

Let $T = \{1, 2, \ldots\}$ denote the time horizon. For each $t \in T$, let G^t denote the Bayesian game at period t and $\sigma(\tilde{x}^t)$ denote the σ-algebra that the Bayesian Nash equilibrium $\tilde{x}^t$ generates. At each period, player α will have as private information the information that the past period ε-Bayesian Nash equilibrium strategies have generated. Hence, the information of player α at period $t+1$, $\mathcal{F}^{t+1}_\alpha$, is

$$\mathcal{F}^{t+1}_\alpha = \mathcal{F}^t_\alpha \bigvee \sigma(\tilde{x}^t), \tag{2}$$

where $\tilde{x}^t \in BNE(G^t)$ and $\mathcal{F}^t_\alpha \bigvee \sigma(\tilde{x}^t)$ denotes the join, i.e., the smallest σ-algebra containing $\mathcal{F}^t_\alpha$ and $\sigma(\tilde{x}^t)$. Consequently, for each player α and period t,

$$\mathcal{F}^t_\alpha \subseteq \mathcal{F}^{t+1}_\alpha \subseteq \mathcal{F}^{t+2}_\alpha \subseteq \cdots.$$

This represents a learning process for player α.

The learning process generates a sequence of Bayesian games $\{G^t : t \in T\}$, where $G^t = \{(X_\alpha, u_\alpha, \mathcal{F}^t_\alpha, q_\alpha) : \alpha \in A\}$. All terms in G^t are defined as before except that $\mathcal{F}^t_\alpha$ is the σ-algebra of player α at time t given by (2). Let

$$L_{X^t_\alpha} = \big\{\tilde{x}_\alpha \in L_1(\mu, Y) : \tilde{x}_\alpha \quad \text{is} \quad \mathcal{F}^t_\alpha\text{-measurable and} \quad \tilde{x}_\alpha(\omega) \in X_\alpha(\omega)\ \mu\text{-a.e.}\ \omega\big\}.$$

and $L_{X^t} = \prod_{\alpha \in A} L_{X^t_\alpha}$.

Let

$$\bar{\mathcal{F}}_\alpha = \bigvee_{t \in T} \mathcal{F}^t_\alpha,$$

where $\bar{\mathcal{F}}_\alpha$ is interpreted as the pooled information of player α over the entire time horizon T. The Bayesian game

$$\bar{G} = \{(X_\alpha, u_\alpha, \bar{\mathcal{F}}_\alpha, q_\alpha) : \alpha \in A\}$$

is called the *limit information Bayesian game*. $L_{\bar{X}}$ and $BNE(\bar{G})$ are defined for the Bayesian game $\bar{G}$ in an analogous way with those in the game G^t.

Note that $\bar{\mathcal{F}}_\alpha$ is not necessarily the same as the full information, $\bigvee_{\alpha \in A} \mathcal{F}_\alpha$, which is the pooled information over all players.

3.3. Examples

Next, we present two examples that elaborate the Bayesian learning results we derive in this paper.

Example 1

There are two states of the world ω_1 and ω_2, and two players:

$$\mathcal{F}_1 = \{\{\omega_1\}, \{\omega_2\}\},$$

$$\mathcal{F}_2 = \{\{\omega_1, \omega_2\}\}.$$

Both players have the common prior μ on the states

$$\mu(\omega_1) = 9/10, \qquad \mu(\omega_2) = 1/10.$$

Player 1, the row player, chooses T or B and player 2, the column player, chooses L or R. If state ω_1 is realized, then the game to be played is

	L	R
T	1, 1	1, 0
B	0, 0	0, 2

and if state ω_2 is realized, then the game is

	L	R
T	0, 1	0, 0
B	1, 0	1, 2

In period t, the Bayesian Nash equilibrium is that player 1 chooses T in state ω_1, and B in state ω_2, and player 2 chooses L regardless of the state.

Suppose that the state is ω_1. Then (T, L) will be actually played in period t. After the play in period t, player 2 knows that the state is ω_1 by observing the choice of T in period t by player 1. Then from period $t+1$ on, they will play (T, L), which is a Bayesian Nash equilibrium given the new information. This is in turn the Nash equilibrium for state ω_1.

Suppose that the realized state is ω_2. Then the actual play in period t will be (B, L) from Bayesian Nash equilibrium. From period $t+1$ on, player 2 knows that the state is ω_2 by observing the choice of B in period t by player 1. Therefore, from period $t+1$ on, they will play (B, R), which is a Bayesian Nash equilibrium, given the information. This is the Nash equilibrium for state ω_2.

This example is the one in which both players learn the true state, and the play path of the Bayesian Nash equilibrium converges to the complete information Nash equilibrium by Bayesian learning. However, the next example shows that in some cases, the true state is not revealed until the end, and the play path only converges to the Bayesian Nash equilibrium.

Example 2

The structure of the game is the same as the above example except the following: If state ω_1 is realized, then the game to be played is

	L	R
T	1, 1	1, 0
B	0, 0	0, 0

and if state ω_2 is realized, then the game is

	L	R
T	1, 1	1, 2
B	0, 0	0, 0

In period t, the Bayesian Nash equilibrium is that player 1 chooses T and player 2 chooses L regardless of the state.

Suppose that state ω_2 is realized. In this case, player 2 cannot tell the true state, even after observing the action in period t by player 1. So, the play path will be (T, L) forever, which is just a realization of a Bayesian Nash equilibrium. Notice that (T, R) is the complete information Nash equilibrium for state ω_2.

4. LEARNING IN FINITE BAYESIAN GAMES

We can now state the assumptions needed for our main theorems.

Assumption 1. $X_\alpha : \Omega \to 2^Y$ is a non-empty, convex, weakly compact valued and integrably bounded correspondence having an $\mathcal{F}_\alpha$-measurable graph, i.e., $G_{X_\alpha} \in \mathcal{F}_\alpha \bigotimes \mathcal{B}(Y)$.

Assumption 2.

(i) For each $\omega \in \Omega$, $u_\alpha(\omega, \cdot) : \prod_{a\in A} X_a(\omega) \to R$ is continuous where $X_a(\omega)$ is endowed with the weak topology.
(ii) For each $x \in \prod_{a\in A} Y_a$ with $Y_a = Y$, $u_\alpha(\cdot, x) : \Omega \to R$ is $\mathcal{F}$-measurable.
(iii) u_α is integrably bounded.

Remark. Under Assumptions 1 and 2, and provided that $u_\alpha(\omega, \cdot)$ is concave, Kim and Yannelis (1996) have shown that $BNE(G) \neq \emptyset$, and therefore $BNE_\varepsilon(G) \neq \emptyset$ as well.

We now state our first result that bounded rational play converges to a limit full information BNE.

THEOREM 1. *Let* $\{G^t : t \in T\}$ *be a sequence of Bayesian games satisfying Assumptions* 1 *and* 2 *and let* $\tilde{x}^t \in BNE_{\varepsilon t}(G^t)$, *where* $\varepsilon_t \to 0$. *Then, there exists a subsequence* $\{\tilde{x}^{t_n} : n = 1, 2, \ldots\}$ *of* $\{\tilde{x}^t : t \in T\}$ *such that* $\tilde{x}^{t_n}$ *converges weakly to* $\tilde{x}^* \in BNE(\bar{G})$.

Proof. First notice that, by Lemma A.3 in the Appendix, each $L_{X^t_\alpha}$ is weakly compact and so is $L_{X^t} = \prod_{a \in A} L_{X^t_\alpha}$. By Lemma A.3, $L_{\bar{X}}$ is weakly compact. Let $B = \{\tilde{x}^t : t \in T\}$. Because $\tilde{x}^t \in L_{X^t} \subset L_{\bar{X}}$ and $L_{\bar{X}}$ is weakly compact, it follows that the weak closure of B, denoted by $w - cl B$, is weakly compact. By the Eberlein–Smulian theorem [Dunford and Schwartz (1958, p. 430)], $w - cl B$ is weakly sequentially compact. Clearly, the weak limit of $\tilde{x}^t$, denoted by $\tilde{x}^*$, belongs to $w - cl B$. From Whitley's theorem [Aliprantis and Burkinshaw (1985, Lemma 10.12, p. 155)], we know that, if $\tilde{x}^* \in w - cl B$, then there exists a sequence $\{\tilde{x}^{t_n} : n = 1, 2, \ldots\}$ such that $\tilde{x}^{t_n}$ converges weakly to $\tilde{x}^*$. Because $\tilde{x}^{t_n} \in L_{X_{t_n}} \subset L_{\bar{X}}$, it follows that $\tilde{x}^* \in L_{\bar{X}}$ and we can conclude that $\tilde{x}^{t_n}_\alpha$ is $\bar{\mathcal{F}}_\alpha$-measurable.

Now fix $\alpha \in A$. Let $y_\alpha \in X_\alpha(\omega)$. We need to show that a.e. $\omega \in \Omega$,

$$v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, \tilde{x}^*_\alpha(\omega)\big) \geq v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, y_\alpha\big),$$

where

$$v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, y_\alpha\big) = \int_{\omega' \in E_\alpha(\omega)} u_\alpha(\omega, \tilde{x}_{-\alpha}(\omega'), y_\alpha) q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))\, d\mu(\omega').$$

Because $\tilde{x}^t \in BNE_{\varepsilon t}(G^t)$, there exists D_t ($\subset \Omega$) with $\mu(D_t) \geq 1 - \varepsilon_t$ such that for all $\alpha \in A$, for all $\omega \in D_t$,

$$v^t_\alpha\big(\omega, \tilde{x}^t_{-\alpha}, \tilde{x}^t_\alpha(\omega)\big) \geq v^t_\alpha\big(\omega, \tilde{x}^t_{-\alpha}, y_\alpha\big) - \varepsilon_t,$$

where

$$v^t_\alpha\big(\omega, \tilde{x}^t_{-\alpha}, y_\alpha\big) = \int_{\omega' \in E^t_\alpha(\omega)} u_\alpha\big(\omega, \tilde{x}^t_{-\alpha}(\omega'), y_\alpha\big) q_\alpha\big(\omega' \mid E^t_\alpha(\omega)\big)\, d\mu(\omega').$$

Let $D = Ls D_t$. Then, $\mu(D) = 1$ because $\varepsilon_t \to 0$. By taking a subsequence if necessary, we have that for all t, for all $\omega \in D$,

$$v^t_\alpha\big(\omega, \tilde{x}^t_{-\alpha}, \tilde{x}^t_\alpha(\omega)\big) \geq v^t_\alpha\big(\omega, \tilde{x}^t_{-\alpha}, y_\alpha\big) - \varepsilon_t. \tag{3}$$

Because $\tilde{x}^t$ converges weakly to $\tilde{x}^*$, it follows from Lemma A.1 in the Appendix that $v^t_\alpha(\omega, \tilde{x}^t_{-\alpha}, \tilde{x}^t_\alpha(\omega))$ converges to $v_\alpha(\omega, \tilde{x}^*_{-\alpha}, \tilde{x}^*_\alpha(\omega))$ and $v^t_\alpha(\omega, \tilde{x}^t_{-\alpha}, y_\alpha)$ converges to $v_\alpha(\omega, \tilde{x}^*_{-\alpha}, y_\alpha)$ as $t \to \infty$. Thus, it follows from (3) that a.e. $\omega \in \Omega$,

$$v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, \tilde{x}^*_\alpha(\omega)\big) \geq v_\alpha\big(\omega, \tilde{x}^*_{-\alpha}, y_\alpha\big),$$

and this completes the proof of the theorem. ∎

COROLLARY 1. *Let* $\{G^t : t \in T\}$ *be a sequence of Bayesian games satisfying Assumptions* 1 *and* 2 *and let* $\tilde{x}^t \in BNE(G^t)$. *Then, there exists a subsequence* $\{\tilde{x}^{t_n} : n = 1, 2, \ldots\}$ *of* $\{\tilde{x}^t : t \in T\}$ *such that* $\tilde{x}^{t_n}$ *converges weakly to* $\tilde{x}^* \in BNE(\overline{G})$.

Proof. The conclusion follows if we let $\varepsilon_t = 0$ for all t in Theorem 1. ■

Remark. Theorem 1 and Corollary 1 still hold if A is any infinite set. The proof remains unchanged. However, the proof of Theorem 2 fails if A is infinite.

THEOREM 2. *Let* $\{G^t : t \in T\}$ *be a sequence of Bayesian games satisfying Assumptions* 1 *and* 2 *and let* $\tilde{x}^* \in BNE(\overline{G})$. *Then, for any* $\varepsilon > 0$, *there exists* $\{\tilde{x}^t : t \in T\}$, *where* $\tilde{x}^t \in BNE_\varepsilon(G^t)$, *such that* $\tilde{x}^t$ *converges in* $L_1(\mu, Y)$-*norm to* $\tilde{x}^*$.

Proof. Let $\tilde{x}^t_\alpha = E[\tilde{x}^*_\alpha \mid \mathcal{F}^t_\alpha]$. Note that

$$\begin{aligned} E\left[\tilde{x}^*_\alpha \,\middle|\, \mathcal{F}^t_\alpha\right] &= E\left[E\left[\tilde{x}^*_\alpha \,\middle|\, \mathcal{F}^{t+1}_\alpha\right] \,\middle|\, \mathcal{F}^t_\alpha\right] \\ &= E\left[\tilde{x}^{t+1}_\alpha \middle| \mathcal{F}^t_\alpha\right]. \end{aligned}$$

Hence, $\{\tilde{x}^t_\alpha, \mathcal{F}^t_\alpha\}_{t=1}^\infty$ is a martingale in $L_{X^t_\alpha} \subset L_1(\mu, Y)$ and by the martingale convergence theorem, $\tilde{x}^t_\alpha$ converges in the $L_1(\mu, Y)$-norm (and hence weakly to $\tilde{x}^*_\alpha$). To complete the proof, it is enough to show that $\tilde{x}^t$ lies in $BNE_\varepsilon(G^t)$ for t big enough. Suppose, by way of contradiction, that for infinitely many t, there exists D_t with $\mu(D_t) \geq \varepsilon$ and $y^t_\alpha \in X^t_\alpha$ such that

$$v^t_\alpha\left(\omega, \tilde{x}^t_{-\alpha}, y^t_\alpha\right) \geq v^t_\alpha\left(\omega, \tilde{x}^t_{-\alpha}, \tilde{x}^t_\alpha(\omega)\right) + \varepsilon,$$

for all $\omega \in D_t$. Let $D = Ls\, D_t$. Then $\mu(D) \geq \varepsilon$. By taking a subsequence if necessary, we have that for all t there exists $y^t_\alpha \in X_\alpha(\omega)$ such that

$$v^t_\alpha\left(\omega, \tilde{x}^t_{-\alpha}, y^t_\alpha\right) \geq v^t_\alpha\left(\omega, \tilde{x}^t_{-\alpha}, \tilde{x}^t_\alpha(\omega)\right) + \varepsilon,$$

for all $\omega \in D$. Because $X_\alpha(\omega)$ is weakly compact, we can assume that y^t_α converges weakly to some $y^*_\alpha \in X_\alpha(\omega)$ by taking a subsequence if necessary. Then it follows from Lemma A.1 in the Appendix that for all $\omega \in D$,

$$v_\alpha\left(\omega, \tilde{x}^*_{-\alpha}, y^*_\alpha\right) \geq v_\alpha\left(\omega, \tilde{x}^*_{-\alpha}, \tilde{x}^*_\alpha(\omega)\right) + \varepsilon > v_\alpha\left(\omega, \tilde{x}^*_{-\alpha}, \tilde{x}^*_\alpha(\omega)\right),$$

which contradicts the fact that $\tilde{x}^* \in BNE(\overline{G})$. The above contradiction establishes the validity of our theorem. ■

5. LEARNING IN CONTINUUM BAYESIAN GAMES

In this section, we study the Bayesian game G with a measure space of agents. A Bayesian game with a measure space of agents $(A, \mathcal{A}, \nu)$ is $G = \{(X, u, \mathcal{F}_\alpha, q_\alpha) : \alpha \in A\}$, where

1. $X : A \times \Omega \to 2^Y$ is the *action set-valued function*, where $X(\alpha, \omega)$ is interpreted as the set of actions available to agent α when the state is ω;

2. for each $(\alpha, \omega) \in A \times \Omega$, $u(\alpha, \omega, \cdot, \cdot) : L_1(\nu, Y) \times X(\alpha, \omega) \to R$ is the *utility function*, where $u(\alpha, \omega, x, x_\alpha)$ is interpreted as the utility of agent α using action x_α when the state is ω and other players use the joint action x;
3. $\mathcal{F}_\alpha$ is a sub σ-algebra of $\mathcal{F}$ that denotes the *private information* of agent α;
4. $q_\alpha : \Omega \to R_{++}$ is the *prior* of agent α.

As before, let L_{X_α} denote the set of all Bochner integrable, $\mathcal{F}_\alpha$-measurable selections from the action set-valued function $X(t)$ of agent α, i.e.,

$$L_{X_\alpha} = \{\tilde{x}(\alpha) \in L_1(\mu, Y) : \tilde{x}(\alpha, \cdot) : \Omega \to Y \quad \text{is} \quad \mathcal{F}_\alpha\text{-measurable and}$$
$$\tilde{x}(\alpha, \omega) \in X(\alpha, \omega)\ \mu\text{-a.e.}\ \omega\}.$$

Let

$$L_X = \{\tilde{x} \in L_1(\nu, L_1(\mu, Y)) : \tilde{x}(\alpha) \in L_{X_\alpha} \quad \text{for} \quad \nu\text{-a.e.}\ \alpha\}.$$

In a Bayesian game with a measure space of agents, a *strategy* for agent α is an element in L_{X_α} and a *joint strategy profile* is an element in L_X. For each $(\alpha, \omega) \in A \times \Omega$, the conditional expected utility function of agent α, $v(\alpha, \omega, \cdot, \cdot) : L_X \times X(\alpha, \omega) \to R$ is defined as

$$v(\alpha, \omega, \tilde{x}, x_\alpha) = \int_{\omega' \in E_\alpha(\omega)} u(\alpha, \omega, \tilde{x}(\omega'), x_\alpha) q_\alpha(\omega' \mid E_\alpha(\omega))\, d\mu(\omega'),$$

where

$$q_\alpha(\omega' \mid E_\alpha(\omega)) = \begin{cases} 0 & \text{if} \quad \omega' \notin E_\alpha(\omega) \\ \dfrac{q_\alpha(\omega')}{\int_{\tilde{\omega} \in E_\alpha(\omega)} q_\alpha(\tilde{\omega})\, d\mu(\tilde{\omega})} & \text{if} \quad \omega' \in E_\alpha(\omega). \end{cases}$$

A *Bayesian Nash equilibrium* for G is a strategy profile $\tilde{x}^* \in L_X$ such that for ν-a.e. α, for μ-a.e. ω,

$$v(\alpha, \omega, \tilde{x}^*, \tilde{x}^*(\alpha, \omega)) = \max_{y \in X(\alpha, \omega)} v(\alpha, \omega, \tilde{x}^*, y).$$

We can now state the assumptions needed for the proof of the next theorem.

Assumption 3.

(i) $X : A \times \Omega \to 2^Y$ is a non-empty, convex, weakly compact valued and integrably bounded correspondence having an $\mathcal{A} \bigotimes \mathcal{F}$-measurable graph, i.e., $G_X \in \mathcal{A} \bigotimes \mathcal{F} \bigotimes \mathcal{B}(Y)$.
(ii) For each $\alpha \in A$, $X(\alpha, \cdot) : \Omega \to 2^Y$ has an $\mathcal{F}_\alpha$-measurable graph, i.e., $G_{X_{(\alpha)}} \in \mathcal{F}_\alpha \bigotimes \mathcal{B}(Y)$.

Assumption 4.

(i) For each $(\alpha, \omega) \in A \times \Omega$, $u(\alpha, \omega, \cdot, \cdot) : L_1(\nu, Y) \times X(\alpha, \omega) \to R$ is continuous where $L_1(\nu, Y)$ and $X(\alpha, \omega)$ are endowed with the weak topologies.

(ii) For each $(x, y) \in L_1(\nu, Y) \times Y$, $u(\cdot, \cdot, x, y) : A \times \Omega \to R$ is $\mathcal{A} \bigotimes \mathcal{F}$-measurable.
(iii) For each $\alpha \in A$, $u(\alpha, \cdot, \cdot, \cdot)$ is integrably bounded.

Assumption 5.

(i) Ω is a countable set.
(ii) The dual Y^* of Y has the RNP (see Section 2 for definition) with respect to $(A, \mathcal{A}, \nu)$.

Remark. Note that Assumptions 3 and 4 are the same as Assumptions 1 and 2. The only new assumption here is Assumption 5, which we need to prove the weak continuity of the expected utility function in Lemma A.2 in the Appendix. If Ω is uncountable and each agent's information partition is uncountable, then to prove the weak continuity of expected utility we need the following assumption:

Assumption 5′. For each $(\alpha, \omega, x_\alpha) \in A \times \Omega \times Y$, $u(\alpha, \omega, \cdot, x_\alpha) : L_1(\nu, Y) \to R$ is linear.

Assumption 5′ is rather a strong assumption but it is necessary to prove the weak continuity if Ω is uncountable [see, for example, Balder and Yannelis (1993)]. Instead, we use Assumption 5 for the result below.

Given an $\varepsilon > 0$, the strategy profile $\tilde{x}^*$ is said to be an *ε-BNE* for G if there exist $B(\subset A)$ and $E(\subset \Omega)$ with $\nu(B) < \varepsilon$, $\mu(E) < \varepsilon$ such that for all $\alpha \in A/B$, for all $\omega \in \Omega/E$,

$$v(\alpha, \omega, \tilde{x}^*, \tilde{x}^*(\alpha, \omega)) \geq v(\alpha, \omega, \tilde{x}^*, y) - \varepsilon,$$

for all $y \in X(\alpha, \omega)$.

THEOREM 3. *Let $\{G^t : t \in T\}$ be a sequence of Bayesian games satisfying Assumptions* 3–5 *and let $\tilde{x}^t \in BNE_{\varepsilon_t}(G^t)$, where $\varepsilon_t \to 0$. Then, there exists a subsequence $\{\tilde{x}^{t_n} : n = 1, 2, \ldots\}$ of $\{\tilde{x}^t : t \in T\}$ such that $\tilde{x}^{t_n}$ converges weakly to $\tilde{x}^* \in BNE(\bar{G})$.*

Proof. Let $B = \{\tilde{x}^t : t \in T\}$. Because $\tilde{x}^t \in L_{X^t} \subset L_{\bar{X}}$ and $L_{\bar{X}}$ is weakly compact, it follows that the weak closure of B, denoted by $w - clB$, is weakly compact. As in the proof of Theorem 1 we can extract a sequence $\{\tilde{x}^{t_n} : n = 1, 2, \ldots\}$ from B such that $\tilde{x}^{t_n}$ converges weakly to $\tilde{x}^*$. Because $\tilde{x}^{t_n} \in L_{X_{t_n}} \subset L_{\bar{X}}$, it follows that $\tilde{x}^* \in L_{\bar{X}}$, and therefore $\tilde{x}^{t_n}_\alpha$ is $\bar{\mathcal{F}}_\alpha$-measurable.

We need to show that for ν-a.e. α, for μ-a.e. ω,

$$v(\alpha, \omega, \tilde{x}^*, \tilde{x}^*(\alpha, \omega)) \geq v(\alpha, \omega, \tilde{x}^*, y),$$

for all $y \in X(\alpha, \omega)$, where

$$v(\alpha, \omega, \tilde{x}, y) = \int_{\omega' \in E_\alpha(\omega)} u(\alpha, \omega, \tilde{x}(\omega'), y) q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))\, d\mu(\omega').$$

Because $\tilde{x}^t \in BNE_{\varepsilon t}(G^t)$, there exist $B_t(\subset A)$ and $D_t(\subset \Omega)$ with $\nu(B_t) \geq 1 - \varepsilon_t$ and $\mu(D_t) \geq 1 - \varepsilon_t$ such that for all $\alpha \in B_t$, for all $\omega \in D_t$,

$$v^t(\alpha, \omega, \tilde{x}^t, \tilde{x}^t(\alpha, \omega)) \geq v^t(\alpha, \omega, \tilde{x}^t, y) - \varepsilon_t,$$

for all $y \in X(\alpha, \omega)$, where

$$v^t(\alpha, \omega, \tilde{x}^t, y) = \int_{\omega' \in E_\alpha^t(\omega)} u(\alpha, \omega, \tilde{x}^t(\omega'), y) q_\alpha\left(\omega' \mid E_\alpha^t(\omega)\right) d\mu(\omega').$$

Let $B = Ls B_t$ and $D = Ls D_t$. Then, $\nu(B) = 1$ and $\mu(D) = 1$ because $\varepsilon_t \to 0$. By taking a subsequence if necessary, we have that, for all t, for all $\alpha \in B$, for all $\omega \in D$,

$$v^t(\alpha, \omega, \tilde{x}^t, \tilde{x}^t(\alpha, \omega)) \geq v^t(\alpha, \omega, \tilde{x}^t, y) - \varepsilon_t, \tag{4}$$

for all $y \in X(\alpha, \omega)$. Because $\tilde{x}^t$ converges weakly to $\tilde{x}^*$, it follows from Lemma A.2 in the Appendix that $v^t(\alpha, \omega, \tilde{x}^t, \tilde{x}^t(\alpha, \omega)) \to v(\alpha, \omega, \tilde{x}^*, \tilde{x}^*(\alpha, \omega))$ and $v^t(\alpha, \omega, \tilde{x}^t, y) \to v(\alpha, \omega, \tilde{x}^*, y)$ as $t \to \infty$. So, it follows from equation (4) that for ν-a.e. α, μ-a.e. ω,

$$v(\alpha, \omega, \tilde{x}^*, \tilde{x}^*(\alpha, \omega)) \geq v(\alpha, \omega, \tilde{x}^*, y),$$

for all $y \in X(\alpha, \omega)$. ■

COROLLARY 2. *Let $\{G^t : t \in T\}$ be a sequence of Bayesian games satisfying Assumptions* 3–5 *and let $\tilde{x}^t \in BNE(G^t)$. Then, there exists a subsequence $\{\tilde{x}^{t_n} : n = 1, 2, \ldots\}$ of $\{\tilde{x}^t : t \in T\}$ such that $\tilde{x}^{t_n}$ converges weakly to $\tilde{x}^* \in BNE(\underline{G})$.*

Proof. The conclusion follows if we let $\varepsilon_t = 0$ for all t in Theorem 3. ■

THEOREM 4. *Let $\{G^t : t \in T\}$ be a sequence of Bayesian games satisfying Assumptions* 3–5 *and let $\tilde{x}^* \in BNE(\underline{G})$. Then, for any $\varepsilon > 0$, there exists $\{\tilde{x}^t : t \in T\}$, where $\tilde{x}^t \in BNE_\varepsilon(G^t)$, such that $\tilde{x}^t$ converges to $\tilde{x}^*$ in $L_1(\mu, Y)$-norm.*

Proof. Let $\tilde{x}^t(\alpha) = E[\tilde{x}^*(\alpha) | \mathcal{F}_\alpha^t]$. Note that

$$\begin{aligned} E\left[\tilde{x}^*(\alpha) \mid \mathcal{F}_\alpha^t\right] &= E\left[E\left[\tilde{x}^*(\alpha) \mid \mathcal{F}_\alpha^{t+1}\right] \mid \mathcal{F}_\alpha^t\right] \\ &= E\left[\tilde{x}^{t+1}(\alpha) \mid \mathcal{F}_\alpha^t\right]. \end{aligned}$$

Hence, $\{\tilde{x}^t(\alpha), \mathcal{F}_\alpha^t\}_{t=1}^\infty$ is a martingale in $L_{X^t(\alpha)} \subset L_1(\mu, Y)$ and by the martingale convergence theorem, $\tilde{x}^t(\alpha)$ converges in the $L_1(\mu, Y)$-norm and hence weakly to $\tilde{x}^*(\alpha)$. To complete the proof, it is enough to show that $\tilde{x}^t$ lies in $BNE_\varepsilon(G^t)$ for t big enough. Suppose, by way of contradiction, that for infinitely many t, there exist $B^t(\subset A)$ and $D^t(\subset \Omega)$ and $y^t(\alpha, \omega) \in X(\alpha, \omega)$ with $\nu(B_t) \geq \varepsilon$, $\mu(D^t) \geq \varepsilon$ such that

$$v^t(\alpha, \omega, \tilde{x}^t, y^t(\alpha, \omega)) \geq v^t(\alpha, \omega, \tilde{x}^t, \tilde{x}^t(\alpha, \omega)) + \varepsilon,$$

for all $\alpha \in B^t$, for all $\omega \in E^t$. Let $B = Ls B^t$ and $D = Ls D_t$. Then $\nu(B) \geq \varepsilon$ and $\mu(D) \geq \varepsilon$. By taking a subsequence if necessary, we have that, for all, t there exists $y^t(\alpha, \omega) \in X(\alpha, \omega)$ such that

$$v^t(\alpha, \omega, \tilde{x}^t, y^t(\alpha, \omega)) \geq v^t(\alpha, \omega, \tilde{x}^t, \tilde{x}^t(\alpha, \omega)) + \varepsilon,$$

for all $\alpha \in B$, for all $\omega \in D$. Because $X(\alpha, \omega)$ is weakly compact, we can assume that $y^t(\alpha, \omega)$ converges weakly to some $y^*(\alpha, \omega) \in X(\alpha, \omega)$ by taking a subsequence if necessary. Then it follows from Lemma A.2 in the Appendix that for all $\alpha \in B$, for all $\omega \in E$,

$$v(\alpha, \omega, \tilde{x}^*, y(\alpha, \omega)) \geq v(\alpha, \omega, \tilde{x}^*, \tilde{x}^*(\alpha, \omega)) + \varepsilon > v(\alpha, \omega, \tilde{x}^*, \tilde{x}^*(\alpha, \omega)),$$

which contradicts that $\tilde{x}^* \in BNE(\bar{G})$. ■

6. CONCLUSIONS

We study learning in Bayesian games with an arbitrary number of players. In particular, we show that in a very general setting (i.e., under quite weak assumptions) repeated bounded rational play will converge to a limit full information Bayesian Nash equilibrium. This result holds for not only a finite or countable number of players (Theorem 1), but even for a continuum (Theorem 3). Moreover, we show the converse, i.e., given a limit full information Bayesian Nash equilibrium strategy we can construct a sequence of bounded rational plays that converges to a limit full information Bayesian Nash equilibrium outcome. This result also holds for finitely many players (Theorem 2) and for a continuum of players (Theorem 4), as well.

REFERENCES

Aliprantis, C.D. & O. Burkinshaw (1985) *Positive Operators*. New York: Academic Press.

Balder, E.J. & N. C. Yannelis (1993) On the continuity of expected utility. *Economic Theory* 3, 625–643.

Diestel, J. & J. Uhl (1977) *Vector Measures*, Mathematical Surveys 15. Providence, RI: American Mathematical Society.

Dunford, N. & J.T. Schwartz (1958) *Linear Operators*, Pt. I. New York: Interscience.

Feldman, M. (1987) An example of convergence to rational expectations with heterogeneous beliefs. *International Economic Review* 28, 635–650.

Hiai, F. & H. Umegaki (1977) Integrals, conditional expectations and martingales of multivalued function. *Journal of Multivariate Analysis* 7, 149–182.

James, R. (1964) Weakly compact sets. *Transactions of the American Mathematical Society* 113, 129–140.

Jordan, J.S. (1991) Bayesian learning in normal form games. *Games and Economic Behavior* 3, 60–81.

Kalai, E. & E. Lehrer (1993) Rational learning leads to Nash equilibrium. *Econometrica* 61, 1019–1046.

Kim, T. & N.C. Yannelis (1996) Existence of equilibrium in Bayesian games with infinitely many agents. *Journal of Economic Theory* (forthcoming).

Koutsougeras, L. & N.C. Yannelis (1994) Convergence and approximation results for noncooperative Bayesian games: Learning theorems. *Economic Theory* 4, 843–858.

Nyarko, Y. (1996) Bayesian learning and convergence to Nash equilibria without common priors. *Economic Theory* (forthcoming).

Tulcea, A.I. & C.I. Tulcea (1969) *Topics in the Theory of Lifting*. New York: Springer–Verlag.

Yannelis, N.C. (1991) Integration of Banach-valued correspondences. In M.A. Khan & N.C. Yannelis (eds.), *Equilibrium Theory with Infinitely Many Commodities*. Heidelberg: Springer–Verlag.

APPENDIX

We begin the Appendix by proving the weak continuity of the conditional expected utility function of each agent. For each $\alpha \in A$, for each $\omega \in \Omega$, define $v^t_\alpha(\omega, \cdot, \cdot) : L_{X^t_{-\alpha}} \times X_\alpha \to R$ and $v_\alpha(\omega, \cdot, \cdot) : L_{\bar{X}_{-\alpha}} \times X_\alpha \to R$ as

$$v^t_\alpha\left(\omega, \tilde{x}^t_{-\alpha}, y^t_\alpha\right) = \int_{\omega' \in \Omega} u_\alpha\left(\omega, \tilde{x}^t_{-\alpha}(\omega'), y^t_\alpha\right) q_\alpha\left(\omega' \,\middle|\, E^t_\alpha(\omega)\right) d\mu(\omega'),$$

and

$$v_\alpha(\omega, \tilde{x}_{-\alpha}, y_\alpha) = \int_{\omega' \in \Omega} u_\alpha\left(\omega, \tilde{x}^t_{-\alpha}(\omega'), y_\alpha\right) q_\alpha\left(\omega' \,\middle|\, \bar{E}^t_\alpha(\omega)\right) d\mu(\omega').$$

LEMMA A.1. *Assume that for each $\omega\Omega$, $u_\alpha(\omega, \cdot) : \prod_{a \in A} X_a(\omega) \to R$ is weakly continuous. If $\tilde{x}^t$ converges weakly to $\tilde{x}^*$ and y^t_α converges weakly to y^*_α, then for each ω, $v^t_\alpha(\omega, \tilde{x}^t_{-\alpha}, y^t_\alpha)$ converges to $v_\alpha(\omega, \tilde{x}_{-\alpha}, y_\alpha)$.*

Proof. We prove this via two steps. First, we show:

Claim A.1. For each $a \in A$, for each $\omega \in \Omega$, the sequence $\{\tilde{x}^t_a(\omega)\}$ in X_a converges weakly to $\tilde{x}_a(\omega)$.

Proof of Claim A.1. Fix $\omega \in \Omega$. To prove the claim, we need to show that for all $y^* \in Y^*$, $y^*(\tilde{x}^t_a(\omega))$ converges to $y^*(\tilde{x}_a(\omega))$. Because $\prod_a = \{E^1_a, E^2_a, \ldots\}$ is a countable partition of Ω of agent s, $\tilde{x}^t_a$ and $\tilde{x}_a$ can be written as

$$\tilde{x}^t_a = \sum_{k=1}^{\infty} x^{n,k}_a \chi_{E^k_\alpha} \quad \text{and} \quad \tilde{x}_a = \sum_{k=1}^{\infty} x^k_a \chi_{E^k_a},$$

where $x^{n,k}_a, x^k_a \in X_a$. Note that for each $s \in T$, there exists a unique $E^{k(\omega)}_a \in \prod_a$ with $\omega \in E^{k(\omega)}_a$. Then,

$$\begin{aligned} y^*\left(\tilde{x}^t_\alpha(\omega)\right) &= \int_{\omega' \in \Omega} \tilde{x}^t_a(\omega) \frac{y^*}{\mu\left(E^{k(\omega)}_a\right)} \chi_{E^{k(\omega)}_a}(\omega')\, d\mu(\omega') \\ &= \int_{\omega' \in \Omega} \tilde{x}^t_a(\omega') \frac{y^*}{\mu\left(E^{k(\omega)}_a\right)} \chi_{E^{k(\omega)}_a}(\omega')\, d\mu(\omega') \qquad \textbf{(A.1)} \end{aligned}$$

because $\tilde{x}^t_a(\omega') = \tilde{x}^t_a(\omega)$ if $\omega' \in E^{k(\omega)}_a$. Note that

$$\frac{y^*}{\mu\left(E^{k(\omega)}_a\right)} \in L_\infty(\mu, Y^*) \quad \text{and} \quad \tilde{x}^t_a \in L_1(\mu, Y).$$

Because $\tilde{x}^t_a$ converges to $\tilde{x}_a$ weakly in $L_1(\mu, Y)$, equation (A.1) converges to

$$\int_{\omega' \in \Omega} \tilde{x}_a(\omega') \frac{y^*}{\mu\left(E^{k(\omega)}_a\right)} \chi_{E^{k(\omega)}_a}(\omega')\, d\mu(\omega') = y^*(\tilde{x}_a(\omega)).$$

Because the choice of $y^* \in Y^*$ is arbitrary, $\tilde{x}^t_a(\omega)$ converges weakly to $\tilde{x}_a(\omega)$. This proves Claim A.1.

Claim A.2. For each $\omega \in \Omega$,

$$\int_{\omega'} u_\alpha\left(\omega, \tilde{x}^t_{-\alpha}(\omega'), y^t_\alpha\right) q_\alpha\left(\omega' \,\middle|\, E^t_\alpha(\omega)\right) d\mu(\omega')$$

$$\to \int_{\omega'} u_\alpha(\omega, \tilde{x}_{-\alpha}(\omega'), y_\alpha) q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))\, d\mu(\omega').$$

Proof of Claim A.2. By Claim A.1, for each $a \in A$, for each $\omega \in \Omega$, $\tilde{x}^t_a(\omega)$ converges weakly to $\tilde{x}_a(\omega)$. By the continuity of $u_\alpha(\omega, \cdot, \cdot)$ with the given topologies, for each $\omega \in \Omega$, $u_\alpha(\omega, \tilde{x}^t_{-\alpha}(\omega), y^t_\alpha)$ converges to $u_\alpha(\omega, \tilde{x}_{-\alpha}(\omega), y_\alpha)$. Because $E^t_\alpha(\omega) \supset E^{t+1}_\alpha(\omega) \supset \bar{E}_\alpha(\omega)$ for all t, $\mu(E^t_\alpha(\omega)) \to \mu(\bar{E}_\alpha(\omega))$. So, by the definition of q_α, $q_\alpha(\omega' \mid E^t_\alpha(\omega)) \to q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))$ a.e. ω. Therefore, by the Lebesgue dominated convergence theorem,

$$\int_{\omega'} u_\alpha\left(\omega, \tilde{x}^t_{-\alpha}(\omega'), y^t_\alpha\right) q_\alpha\left(\omega' \mid E^t_\alpha(\omega)\right) d\mu(\omega')$$

$$\to \int_{\omega'} u_\alpha(\omega, \tilde{x}_{-\alpha}(\omega'), y_\alpha) q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))\, d\mu(\omega'). \qquad \blacksquare$$

For each $\alpha \in A$, for each $\omega \in \Omega$, define $v^t(\alpha, \omega, \cdot, \cdot) : L_1(\nu, L_{X^t}) \times X_\alpha(\omega) \to R$ and $v(\alpha, \omega, \cdot, \cdot) : L_1(\nu, L_{\bar{X}}) \times X_\alpha(\omega) \to R$ as

$$v^t\left(\alpha, \omega, \tilde{x}^t, y^t_\alpha\right) = \int_{\omega' \in \Omega} u\left(\alpha, \omega, \tilde{x}^t(\omega'), y^t_\alpha\right) q_\alpha\left(\omega' \,\middle|\, E^t_\alpha(\omega)\right) d\mu(\omega'),$$

and

$$v(\alpha, \omega, \tilde{x}, y_\alpha) = \int_{\omega' \in \Omega} u(\alpha, \omega, \tilde{x}(\omega'), y_\alpha) q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))\, d\mu(\omega').$$

LEMMA A.2. *Let $(A, \mathcal{A}, \nu)$ and $(\Omega, \mathcal{F}, \mu)$ be finite measure spaces, where Ω is a countable set. Let X be a weakly compact subset of the separable Banach space Y whose dual Y^* has the RNP (Radon–Nikodym property) with respect to $(T, \mathcal{T}, \nu)$. Assume that for each $\alpha \in A$, for each $\omega \in \Omega$, $u(\alpha, \omega, \cdot, \cdot) : L_1(\nu, X) \times X \to R$ is weakly continuous. If $\tilde{x}^t$ converges weakly to $\tilde{x}^*$ and y^t_α converges weakly to y^*_α, then for each ω, $v^t(\alpha, \omega, \tilde{x}^t_{-\alpha}, y^t_\alpha)$ converges to $v(\alpha, \omega, \tilde{x}_{-\alpha}, y_\alpha)$.*

Proof. We prove this via two steps. First, we show

Claim A.3. For each $\omega \in \Omega$, the sequence $\{\tilde{x}^t(\omega)\}$ in $L_1(\nu, X)$ converges weakly to $\tilde{x}(\omega)$.

Proof of Claim A.3. Fix $\omega \in \Omega$. To prove the claim, we need to show that for all $y^* \in [L_1(\nu, Y)]^* = L_\infty(\nu, Y^*)$ [by the RNP of Y^* with respect to $(A, \mathcal{A}, \nu)$],

$$\int_{\infty \in A} \tilde{x}^t_\alpha(\omega) y^*(\alpha)\, d\nu(\alpha) \quad \text{converges to} \quad \int_{\alpha \in A} \tilde{x}_\alpha(\omega) y^*(t)\, d\nu(t).$$

Because $\prod_\alpha = \{E^1_\alpha, E^2_\alpha, \ldots\}$ is a countable partition of Ω of agent t, $\tilde{x}^t_\alpha$ and $\tilde{x}_\alpha$ can be written as

$$\tilde{x}^t_\alpha = \sum_{k=1}^{\infty} x^{n,k}_\alpha \chi_{E^k_\alpha} \quad \text{and} \quad \tilde{x}_\alpha = \sum_{k=1}^{\infty} x^k_\alpha \chi_{E^k_\alpha},$$

where $x_\alpha^{n,k}$, $x_\alpha^k \in X$. Note that for each $t \in T$, there exists a unique $E_\alpha^{k(\omega)} \in \prod_\alpha$ with $\omega \in E_\alpha^{k(\omega)}$. Moreover, for each $t \in T$,

$$\mu\left(E_\alpha^{k(\omega)}\right) > \mu(\{\omega\}) > 0.$$

First, choose $y^* \in L_\infty(\nu, Y^*)$ such that

$$y^* = a^* \chi_{T_0}, \quad \text{where} \quad a^* \in Y^* \quad \text{and} \quad T_0 \in \mathcal{T}.$$

Then,

$$\begin{aligned}\int_{t\in T} \tilde{x}_\alpha^t(\omega) y^*(t)\, d\nu(t) &= \int_{t\in T_0} \tilde{x}_\alpha^t(\omega) a^*\, d\nu(t)\\ &= \int_{t\in T_0}\left[\int_{\omega'\in\Omega} \tilde{x}_\alpha^t(\omega)\frac{a^*}{\mu\left(E_\alpha^{k(\omega)}\right)}\chi_{E_\alpha^{k(\omega)}}(\omega')\, d\mu(\omega')\right] d\nu(t)\\ &= \int_{t\in T_0}\left[\int_{\omega'\in\Omega} \tilde{x}_\alpha^t(\omega')\frac{a^*}{\mu\left(E_\alpha^{k(\omega)}\right)}\chi_{E_\alpha^{k(\omega)}}(\omega')\, d\mu(\omega')\right] d\nu(t) \quad \textbf{(A.2)}\end{aligned}$$

because $\tilde{x}_\alpha^t(\omega') = \tilde{x}_\alpha^t(\omega)$ if $\omega' \in E_\alpha^{k(\omega)}$. Note that, for each $t \in T$,

$$\frac{a^*}{\mu\left(E_\alpha^{k(\omega)}\right)} \in L_\infty(\mu, Y^*) \quad \text{and} \quad \tilde{x}_\alpha^t \in L_1(\mu, Y).$$

Because $\mu(E_\alpha^{k(\omega)})$ is uniformly bounded from below by $\mu(\{\omega\})$, the mapping $t \mapsto [a^*/\mu(E_\alpha^{k(\omega)})]\chi_{E_\alpha^{k(\omega)}}$ is in $L_\infty(\nu, L_\infty(\mu, Y^*))$. Because $\tilde{x}^t$ converges weakly to $\tilde{x}$ in $L_1(\nu, L_1(\mu, Y))$ equation (A.2) converges to

$$\begin{aligned}\int_{t\in T_0}\left[\int_{\omega'\in\Omega} \tilde{x}_\alpha(\omega')\frac{a^*}{\mu\left(E_\alpha^{k(\omega)}\right)}\chi_{E_\alpha^{k(\omega)}}(\omega')\, d\mu(\omega')\right] d\nu(t) &= \int_{t\in T_0} \tilde{x}_\alpha(\omega) a^*\, d\nu(t)\\ &= \int_{t\in T} \tilde{x}_\alpha(\omega) y^*(t)\, d\nu(t),\end{aligned}$$

where the first equality holds because $\tilde{x}_\alpha^t(\omega') = \tilde{x}_\alpha^t(\omega)$ if $\omega' \in E_\alpha^{k(\omega)}$. So, for any simple function $y^* \in L_\infty(\nu, Y^*)$,

$$\int_{t\in T} \tilde{x}_\alpha^t(\omega) y^*(t)\, d\nu(t) \quad \text{converges to} \quad \int_{t\in T} \tilde{x}_\alpha(\omega) y^*(t)\, d\nu(t).$$

Next, let $y^* \in L_\infty(\nu, Y^*)$. Because $(T, \mathcal{F}, \nu)$ is a finite measure space, there exists a sequence of simple functions converging to y^* uniformly (recall the Egoroff theorem). Let $\varepsilon > 0$ be given and let $h \in L_1(\nu, Y)$ be a simple function such that

$$\|y^* - h\| < \frac{\varepsilon}{p} \quad \text{where} \quad p > \sup\left\{ \int_{t\in T} \|\tilde{x}^t_\alpha(\omega)\| \, d\nu(t), \int_{t\in T} \|\, \tilde{x}_\alpha(\omega)\, \| d\nu(t) : n = 1, 2, \ldots \right\}.$$

Then,

$$\left| \int_{t\in T} \tilde{x}^t_\alpha(\omega) y^*(t)\, d\nu(t) - \int_{t\in T} \tilde{x}_\alpha(\omega) y^*(t)\, d\nu(t) \right|$$

$$\leq \left| \int_{t\in T} \tilde{x}^t_\alpha(\omega)(y^*(t) - h(t))\, d\nu(t) \right| + \left| \int_{t\in T} \left(\tilde{x}^t_\alpha(\omega) - \tilde{x}_\alpha(\omega)\right) h(t)\, d\nu(t) \right|$$

$$+ \left| \int_{t\in T} \tilde{x}_\alpha(\omega)(h(t) - y^*(t))\, d\nu(t) \right|$$

$$\leq 2\varepsilon + \left| \int_{t\in T} \left(\tilde{x}^t_\alpha(\omega) - \tilde{x}_\alpha(\omega)\right) h(t)\, d\nu(t) \right|.$$

Because h is simple, we obtain

$$\lim_{n\to\infty} \left| \int_{t\in T} \left(\tilde{x}^t_\alpha(\omega) - \tilde{x}_\alpha(\omega)\right) h(t)\, d\nu(t) \right| = 0.$$

Thus, the above estimates imply that

$$\int_{t\in T} \tilde{x}^t_\alpha(\omega) y^*(t)\, d\nu(t) \quad \text{converges to} \quad \int_{t\in T} \tilde{x}_\alpha(\omega) y^*(t)\, d\nu(t),$$

for all $y^* \in L_\infty(\nu, Y^*)$. This proves Claim A.3.

Claim A.4. For each $\alpha \in A$, for each $\omega \in \Omega$,

$$\int_{\omega'} u\left(\alpha, \omega, \tilde{x}^t(\omega'), y^t_\alpha\right) q_\alpha\left(\omega' \;\middle|\; E^t_\alpha(\omega)\right) d\mu(\omega')$$

$$\to \int_{\omega'} u(\alpha, \omega, \tilde{x}(\omega'), y_\alpha) q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))\, d\mu(\omega').$$

Proof of Claim A.4. By Claim A.3, for each $\omega \in \Omega$, $\tilde{x}^t(\omega)$ converges weakly to $\tilde{x}(\omega)$. By the weak continuity of $u(\alpha, \omega, \cdot, \cdot)$, for each $\alpha \in A$, for each $\omega \in \Omega$, $u(\alpha, \omega, \tilde{x}^t(\omega), y^t_\alpha)$ converges to $u(\alpha, \omega, \tilde{x}(\omega), y_\alpha)$. Because $E^t_\alpha(\omega) \supset E^{t+1}_\alpha(\omega) \supset \bar{E}_\alpha(\omega)$ for all t, $\mu(E^t_\alpha(\omega)) \to \mu(\bar{E}_\alpha(\omega))$. So, by the definition of q_α, $q_\alpha(\omega' \mid E^t_\alpha(\omega)) \to q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))$ a.e. ω. Therefore, by the Lebesgue dominated convergence theorem,

$$\int_{\omega'} u\left(\alpha, \omega, \tilde{x}^t(\omega'), y^t_\alpha\right) q_\alpha\left(\omega' \middle| E^t_\alpha(\omega)\right) d\mu(\omega')$$

$$\to \int_{\omega'} u(\alpha, \omega, \tilde{x}(\omega'), y_\alpha) q_\alpha(\omega' \mid \bar{E}_\alpha(\omega))\, d\mu(\omega'). \qquad \blacksquare$$

The Lemma below is known as Diestel's theorem and several alternative proofs can be found in the literature. For completeness, we provide a proof [see also Yannelis (1991, p. 7) and the reference therein].

LEMMA A.3. *Let Y be a separable Banach space and $X : \Omega \to 2^Y$ be an integrably bounded, weakly compact, convex valued correspondence. Then*

$$L_X = \{\tilde{x} \in L_1(\mu, Y) : \tilde{x} \quad \textit{is} \quad \mathcal{F}\textit{-measurable and} \quad \tilde{x}(\omega) \in X(\omega)\ \mu\text{-a.e. } \omega\}$$

is weakly compact in $L_1(\mu, Y)$.

Proof. The proof is based on the celebrated theorem of James (1964). Note that the dual of $L_1(\mu, Y)$ is $L_\infty(\mu, Y^*_{w^*})$ where w^* denotes the w^*-topology, i.e., $L_1(\mu, Y)^* = L_\infty(\mu, Y^*_{w^*})$ [see, for instance, Tulcea and Tulcea (1969)]. Let x be an arbitrary element of $L_\infty(\mu, Y^*_{w^*})$. If we show that x attains its supremum on L_X, the result will follow from James' theorem [James (1964)]. Note that

$$\sup_{\psi \in L_X} \psi \cdot x = \sup_{\psi \in L_X} \int_{\omega \in \Omega} (\psi(\omega) x(\omega))\, d\mu(\omega)$$

$$= \int_{\omega \in \Omega} \sup_{\varphi \in X(\omega)} (\varphi \cdot x(\omega))\, d\mu(\omega),$$

where the second equality follows from Theorem 2.2 of Hiai and Umegaki (1977). Define $g : \Omega \to 2^Y$ as

$$g(\omega) = \{y \in X(\omega) : y \cdot x(\omega) = \sup_{\varphi \in X(\omega)} \varphi \cdot x(\omega)\}.$$

It follows from the weak compactness of $X(\omega)$ that for all $\omega \in \Omega$, $g(\omega)$ is non-empty. Define $f : \Omega \times Y \to R$ by

$$f(\omega, y) = \sup_{\varphi \in X(\omega)} \varphi \cdot x(\omega) - y \cdot x(\omega).$$

It is easy to see that for each ω, $f(\omega, \cdot)$ is continuous and for each y, $f(\cdot, y)$ is $\mathcal{F}$-measurable and hence $f(\cdot, \cdot)$ is jointly measurable. Then observe that $G_g = f^{-1}(0) \cap G_X$ and that because $f^{-1}(0)$ and G_X belong to $\mathcal{F} \bigotimes \mathcal{B}(Y)$, so does G_g. It follows from the Aumann measurable selection theorem that there exists an $\mathcal{F}$-measurable function $z : \Omega \to Y$ such that $z(\omega) \in g(\omega)$ μ-a.e. ω. Thus, $z \in L_X$ and we have

$$\sup_{\varphi \in L_X} \varphi \cdot x = \int_{\omega \in \Omega} (z(\omega) x(\omega))\, d\mu(\omega) = z \cdot x.$$

Because $x \in L_\infty(\mu, Y^*_{w^*})$ was chosen arbitrarily, we conclude that every element of $(L_1(\mu, Y))^*$ attains its supremum on L_X and this completes the proof of the fact that L_X is weakly compact. $\blacksquare$