Macroeconomic Dynamics, **2**, 1998, 141–155. Printed in the United States of America.

ARTICLES

LEARNING IN BAYESIAN GAMES BY BOUNDED RATIONAL PLAYERS II: NONMYOPIA

KONSTANTINOS SERFES AND NICHOLAS C. YANNELIS
University of Illinois at Urbana–Champaign

We generalize results of earlier work on learning in Bayesian games by allowing players to make decisions in a nonmyopic fashion. In particular, we address the issue of *nonmyopic* Bayesian learning with an arbitrary number of bounded rational players, i.e., players who choose approximate best-response strategies for the entire horizon (rather than the current period). We show that, by repetition, nonmyopic bounded rational players can reach a limit full-information nonmyopic Bayesian Nash equilibrium (NBNE) strategy. The converse is also proved: Given a limit full-information NBNE strategy, one can find a sequence of nonmyopic bounded rational plays that converges to that strategy.

Keywords: Bayesian Game, Nash Equilibrium, Nonmyopia, Bayesian Learning, Bounded Rationality

1. INTRODUCTION

The issue of *myopic* Bayesian learning by a finite number of bounded rational players has been addressed by Koutsougeras and Yannelis (1994). Recently, Kim and Yannelis (1997b) extended that work by allowing the number of bounded rational players to be arbitrary, i.e., any finite or infinite set or a continuum. Here, we drop the myopia assumption and allow the players to be *nonmyopic* i.e., to make decisions by taking into account the future.

In particular, the description of the model is as follows: Let (Ω, F, μ) be a probability measure space interpreted as the set of states of the world. Let T denote the *time horizon* and A the *set of players*. A *repeated Bayesian game* (or a repeated game with differential information) is a sequence of games $\{G^t : t \in T\}$ such that for each t, $G^t = \{(F^t_\alpha, X^t_\alpha, u_\alpha, q_\alpha) : \alpha \in A\}$, where

1. F^t_α denotes the *private information* of agent α in period t,
2. $X^t_\alpha(\omega)$ is the *set of actions* available to agent α in period t when the state is ω,
3. $u_\alpha(\omega, \cdot) : \Pi_{\alpha\in A} X^t_\alpha(\omega) \to R$ is the *utility function* of agent α,
4. q_α is the *prior* of agent α [q_α is a density function, or Radon-Nikodym derivative, such that, $\int_{\omega\in\Omega} q_\alpha(\omega)\, d\mu(\omega) = 1$].

Address correspondence to: Nicholas C. Yannelis, Department of Economics, University of Illinois at Urbana–Champaign, Champaign IL 61820, USA; e-mail: nyanneli@uiuc.edu.

© 1998 Cambridge University Press 1365-1005/98 $9.50

The strategy $x_\alpha = \{x_\alpha^t : t \in T\}$ of player α is a sequence $\{x_\alpha^t : t \in T\}$ where each component is F_α^t measurable and $x_\alpha^t(\omega) \in X_\alpha^t(\omega)$, μ-a.e. and for all $t \in T$. Given $E_\alpha^t(\omega)$, for each player α and for each strategy profile $x^t = \{x_\alpha^t : \alpha \in A\}$ in period t, define the *conditional expected utility* $v_\alpha(\omega, x^t)$ of player α as

$$v_\alpha(\omega, x^t) = \int_{\omega' \in E_\alpha^t(\omega)} u_\alpha(\omega, x^t(\omega')) q_\alpha\left(\omega' \,\middle|\, E_\alpha^t(\omega)\right) d\mu(\omega'),$$

where $q_\alpha(\omega' \mid E_\alpha^t(\omega))$ denotes the conditional probability of ω', given $E_\alpha^t(\omega)$.

We define player α's *total discounted expected utility* $U_\alpha(\omega, x)$ for the strategy profile $x = \{x^t : t \in T\}$ as

$$U_\alpha(\omega, x) = \sum_{t \in T} \delta^t v_\alpha(\omega, x^t),$$

where $\delta \in [0, 1)$ is the *discount factor.*

An ε*-nonmyopic Bayesian Nash equilibrium* [$\text{NBNE}_\varepsilon(G^*)$] is a strategy profile $x = \{x^t : t \in T\}$ such that, for all $\alpha \in A$ and for μ-a.e.,

$$U_\alpha(\omega, x) \geq U_\alpha(\omega, x_{-\alpha}, y_\alpha) - \varepsilon$$

for all strategies y_α.

The NBNE_ε captures the idea of a bounded rational player in the sense that each player chooses approximate or ε-best-response strategies by taking into account the future decisions. We call players who choose NBNE_ε equilibrium strategies bounded rational, or we say that the play is bounded rational.

Learning in this model takes place as follows: The private information of player α in period $t+1$, denoted by F_α^{t+1}, is given by

$$F_\alpha^{t+1} = F_\alpha^t \vee \sigma(x^t),$$

where x^t is the projection of a $\text{NBNE}(G^*)$ on the tth coordinate and $F_\alpha^t \vee \sigma(x^t)$ denotes the join, i.e., the smallest σ-algebra containing F_α^t and $\sigma(x^t)$. Consequently, for each player α and period t, we have

$$F_\alpha^t \subseteq F_\alpha^{t+1} \subseteq F_\alpha^{t+2} \subseteq \cdots. \tag{1}$$

Expression (1) represents a learning process for player α.

Let

$$\bar{F}_\alpha = \vee_{t \in T} F_\alpha^t, \tag{2}$$

where $\bar{F}_\alpha$ is the pooled information of player α over the entire horizon T. A Bayesian game

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

where $X_\alpha, u_\alpha, q_\alpha$ are defined as above and $\bar{F}_\alpha$ is given by (2), is called a *limit full-information Bayesian game*.

Notice that the above setting is more general than the one of Kim and Yannelis (1997b). In particular, by letting the discount be equal to zero, we are reduced to the Kim-Yannelis framework. The questions that we address (and for which we provide positive answers) are the following:

1. Can nonmyopic bounded rational players by repetition reach a limit full-information NBNE outcome?
2. Conversely, pick a NBNE strategy for a limit full-information game. Can we construct a sequence of bounded rational plays that converge to that strategy?

In a different setting and for a less general Bayesian game framework than ours, question 1 has been addressed by Kalai and Lehrer (1993) and Nyarko (1996). Question 2 is addressed for the first time in a nonmyopic setting.

The rest of the paper is organized as follows: Section 2 contains notation and definitions. Section 3 describes the Bayesian game with differential information with finitely many players. Section 4 proves the existence of a NBNE(G^*). In Section 5, we describe how learning takes place. In Section 6, we prove that nonmyopic bounded rational players will reach a limit full-information NBNE outcome and, conversely, given a limit full-information NBNE outcome, we can construct a sequence of bounded rational nonmyopic play that converges to the limit NBNE outcome. Section 7 addresses the same questions as those in Section 5 but in a Bayesian game with a continuum of nonmyopic players.

2. NOTATION AND DEFINITIONS

If X and Y are sets, the *graph* of the set-valued function (or correspondence), $\phi : X \to 2^Y$, is denoted by

$$G_\phi = \{(x, y) \in X \times Y : y \in \phi(x)\}.$$

Let (Ω, F, μ) be a complete, finite measure space, and let X be a separable Banach space. The set-valued function $\phi : \Omega \to 2^X$ is said to have a *measurable graph* if $G_\phi \otimes \beta(X)$, where $\beta(X)$ denotes the Borel σ-algebra on X and $\otimes$ denotes the product σ-algebra. The set-valued function $\phi : \Omega \to 2^X$ is said to be *lower measurable* or just *measurable* if for every open subset V of X, the set

$$\{\omega \in \Omega : \phi(\omega) \cap V \neq \emptyset\}$$

is an element of F.

Let (Ω, F, μ) be a finite measure space and let X be a Banach space. Following Diestel and Uhl (1977), the function $f : \Omega \to X$ is called *simple* if there exist $x_1, x_2, \ldots, x_n$ in X and $\alpha_1, \alpha_2, \ldots, \alpha_n$ in F such that $\sum_{i=1}^n x_i \chi_{\alpha_i}$ where $\chi_{\alpha_i}(\omega) = 1$ if $\omega \in \alpha_i$ and $\chi_{\alpha_i}(\omega) = 0$ if $\omega \notin \alpha_i$. A function $f : \Omega \to X$ is said to be

μ*-measurable* if there exists a sequence of simple functions $f_n : \Omega \to X$ such that $\lim_{n\to\infty} \| f_n(\omega) - f(\omega)\| = 0$ for almost all $\omega \in \Omega$. A μ-measurable function $f : \Omega \to X$ is said to be *Bochner integrable* if there exists a sequence of simple functions $\{ f_n : n = 1, 2, \ldots\}$ such that

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

In this case, for each $E \in F$, we define the integral to be

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

It can be shown [see Diestel and Uhl (1977), Theorem 2, p. 45] that, if $f : \Omega \to X$ is a μ-measurable function, then f is Bochner integrable if and only if $\int_\Omega \| f(\omega)\| \, d\mu(\omega) < \infty$.

For $1 \le p < \infty$, we denote by $L_p(\mu, X)$ the space of equivalence classes of X-valued Bochner integrable functions $x : \Omega \to X$ normed by

$$\|x\|_p = \left(\int_\Omega \|x(\omega)\|^p \, d\mu(\omega)\right)^{\frac{1}{p}}.$$

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

Let $X : \Omega \to 2^Y$ be a correspondence, where Y is a Banach space. Also, let $u : \Omega \times Y \to R$ be a real-valued function. Ω can be decomposed into an atomless part Ω_1 and a countable union of atoms Ω_2. The following result from Balder and Yannelis (1993, Theorem 2.8) states: Suppose that a.e. in Ω_1, $X(\omega)$ is convex and closed, $u(\omega, \cdot)$ is concave and upper semicontinuous on $X(\omega)$ and $u(\omega, \cdot)$ is integrably bounded. Suppose further that for all $\omega \in \Omega_2$, $X(\omega)$ is weakly closed and $u(\omega, \cdot)$ is weakly upper semicontinuous on $X(\omega)$. Then,

$$U(x) = \int_\Omega u(\omega, x(\omega))\, d\mu(\omega)$$

is weakly upper semicontinuous on the weakly closed set $L_X = \{y \in L_1(\mu, Y) : y(\omega) \in X(\omega)$ and y is F-measurable$\}$.

A corollary of the above result says that if Ω is countable and $X(\omega)$ is weakly closed and $u(\omega, \cdot)$ is weakly continuous, then U is weakly continuous as well.

A Banach space 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 [see Diestel and Uhl (1977)] that if Y^* (the norm dual of Y) has the RNP with respect to $(T, \mathcal{T}, \nu)$, then

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

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 $\{F_i : i \in I\}$ be a monotone increasing net of sub-σ-fields of F (i.e., $F_{i_1} \subseteq F_{i_2}$ for $i_1 \le i_2, i_1, i_2$ in I). A net $\{x_i : i \in I\}$ in $L_1(\mu, X)$ is *a martingale* if

$$E\left(x_i \mid F_{i_1}\right) = x_{i_1}, \qquad \forall i \ge i_1.$$

We denote the above martingale by $\{x_i, F_i\}_{i \in I}$. The proof of the following *martingale convergence theorem* can be found in Diestel and Uhl (1977, p. 126). A martingale $\{x_i, F_i\}_{i\in I}$ 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 $E(x \mid F_i) = x_i$ for all $i \in I$. Finally, recall [see, e.g., Diestel and Uhl (1977, p. 129)] that if the martingale $\{x_i, F_i\}_{i\in I}$ converges in the $L_1(\mu, X)$-norm to $x \in L_1(\mu, X)$, it also converges almost everywhere, i.e., $\lim_{i\to\infty} x_i = x$ almost everywhere.

3. THE GAME WITH DIFFERENTIAL INFORMATION

Let T be a countable set that denotes the *time horizon.* An element of T is denoted by t. Let (Ω, F, μ) be a complete, finite, separable measure space, where Ω denotes the set of states of the world and the σ-algebra 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 *repeated Bayesian game* (or a repeated game with differential information) is a sequence of games $\{G^t : t \in T\}$ such that for each t, $G^t = \{(F^t_\alpha, X^t_\alpha, u_\alpha, q_\alpha) : \alpha \in A\}$, where

1. F^t_α, is a sub-σ-algebra of F which denotes the *private information* of agent α in period t.
2. $X^t_\alpha : \Omega \to 2^Y$ is the *action set-valued function* of agent α, where $X^t_\alpha(\omega)$ is the set of actions of agent α, in period t, when the state is ω, which is F^t_α-measurable.[1]
3. For each $\omega \in \Omega$, $u_\alpha(\omega, \cdot) : \Pi_{\alpha\in A} X^t_\alpha(\omega) \to R$ is the *utility function* of agent α, which depends on the states.
4. $q_\alpha : \Omega \to R_{++}$ is the *prior* of agent α, [q_α is a density function or Radon-Nikodym derivative, such that $\int_{\omega\in\Omega} q_\alpha(\omega)\, d\mu(\omega) = 1$].

Let $L_{X^t_\alpha}$ denote the set of all Bochner-integrable and F^t_α-measurable selections from the action set-valued function $X^t_\alpha : \Omega \to 2^Y$ of agent α in period t, i.e.,

$$L_{X^t_\alpha} = \left\{x^t_\alpha \in L_1(\mu, Y) : x^t_\alpha \text{ is } F^t_\alpha\text{-measurable and } x^t_\alpha(\omega) \in X^t_\alpha(\omega),\ \mu\text{-a.e.}\right\}.$$

Let L_{X_α} be the product of $L_{X^t_\alpha}$ over all $t \in T$, i.e., $L_{X_\alpha} = \Pi_{t\in T} L_{X^t_\alpha}$. A typical element of L_{X_α} is denoted by $x_\alpha = \{x^t_\alpha : t \in T\}$ and is a sequence of strategies for player α over the entire horizon, where each element of the sequence belongs to $L_{X_\alpha^t}$, i.e., $x_\alpha^t \in L_{X_\alpha^t}$. A typical element of $X_\alpha^t(\omega)$ is denoted by $x_\alpha^t(\omega)$ and a typical element of $\Pi_{t\in T} X_\alpha^t(\omega)$ is denoted by $x_\alpha(\omega)$.

Let $L_X = \Pi_{\alpha\in A} L_{X_\alpha}$ and $L_{X_{-\alpha}} = \Pi_{a\neq\alpha} L_{X_a}$. A typical element of L_X is denoted by x and of $L_{X_{-\alpha}}$ by $x_{-\alpha}$. We endow all product spaces with the product topology.

Throughout the paper, we assume that for each $\alpha \in A$ and each $t \in T$, there exists a finite or countable partition P_α^t of Ω. Moreover, the σ-algebra F_α^t is generated by P_α^t. For each $\omega \in \Omega$ and $t \in T$, let $E_\alpha^t(\omega) \in P_\alpha^t$ denote the smallest set in F_α^t containing ω and assume that, for each α and for each t,

$$\int_{\omega'\in E_\alpha^t(\omega)} q_\alpha(\omega')\, d\mu(\omega') > 0.$$

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

$$v_\alpha\big(\omega, x_{-\alpha}^t, x_\alpha^t(\omega)\big) = \int_{\omega'\in E_\alpha^t(\omega)} u_\alpha\big(\omega, x_{-\alpha}^t(\omega'), x_\alpha^t(\omega')\big) q_\alpha\big(\omega' \mid E_\alpha^t(\omega)\big)\, d\mu(\omega'),$$

where

$$q_\alpha\big(\omega' \mid E_\alpha^t(\omega)\big) = \begin{cases} 0 & \text{if} \quad \omega' \notin E_\alpha^t(\omega) \\ \dfrac{q_\alpha(\omega')}{\int_{\tilde\omega\in E_\alpha^t(\omega)} q_\alpha(\tilde\omega)\, d\mu(\tilde\omega)} & \text{if} \quad \omega' \in E_\alpha^t(\omega). \end{cases}$$

The function $v_\alpha(\omega, x_{-\alpha}^t, x_\alpha^t(\omega))$ is interpreted as the conditional expected utility of agent α in period t, when he/she is using the action $x_\alpha^t(\omega)$, the realized state is ω, and the other agents employ the strategy profile $x_{-\alpha}^t$, where $x_{-\alpha}^t \in L_{X_{-\alpha}^t}$.

For each $\omega \in \Omega$, the *total discounted conditional (interim) expected utility* of agent α,

$$U_\alpha(\omega, \cdot, \cdot) : L_{X_{-\alpha}} \times \prod_{t\in T} X_\alpha^t(\omega) \to R$$

is defined as

$$U_\alpha(\omega, x_{-\alpha}, x_\alpha(\omega)) = \sum_{t\in T} \delta^t v_\alpha\big(\omega, x_{-\alpha}^t, x_\alpha^t(\omega)\big),$$

where $\delta \in [0, 1)$, is the *discount factor.*

The function $U_\alpha(\omega, x_{-\alpha}, x_\alpha(\omega))$ is interpreted as the total discounted expected utility of agent α, when he/she is using the sequence of strategies $x_\alpha(\omega)$, the realized state is ω, and the other agents employ the sequence-of-strategies profile $x_{-\alpha}$.

A *nonmyopic Bayesian Nash equilibrium* for $G^* = \{G^t : t \in T\}$ [denoted by NBNE(G^*)], is a strategy profile $x^* \in L_X$ such that, for all $\alpha \in A$,

$$U_\alpha\big(\omega, x_{-\alpha}^*, x_\alpha^*(\omega)\big) = \max_{y_\alpha\in\Pi_{t\in T} X_\alpha^t(\omega)} U_\alpha\big(\omega, x_{-\alpha}^*, y_\alpha(\omega)\big),\ \mu\text{-a.e.}$$

Given an $\varepsilon > 0$, the strategy profile $x^* \in L_X$ is said to be an ε-NBNE$_\varepsilon(G^*)$ if, for each $\alpha \in A$ and μ-a.e.,

$$U_\alpha\left(\omega, x^*_{-\alpha}, x^*_\alpha(\omega)\right) \geq U_\alpha\left(\omega, x^*_{-\alpha}, y_\alpha(\omega)\right) - \varepsilon$$

for all $y_\alpha(\omega) \in \Pi_{t\in T} X^t_\alpha(\omega)$.

4. EXISTENCE OF A NBNE(G^*)

We can now state the assumptions needed for the existence of an NBNE(G^*). First, we will establish the weak continuity of the total discounted expected utility. Once this is done, the existence of a NBNE(G^*) follows from [Kim and Yannelis (1997a) or Yannelis (1997)]. We need the following assumptions:

Assumption 1. $X^t_\alpha : \Omega \to 2^Y$ is a non-empty, convex, weakly compact-valued and integrably bounded correspondence, having an F^t_α-measurable graph.

Assumption 2.

(i) For each $\omega \in \Omega$ and for each $t \in T$, $u_\alpha(\omega, \cdot, \cdot) : \Pi_\alpha X^t_\alpha(\omega) \to R$ is continuous, where each $X^t_\alpha(\omega)$ is endowed with the weak topology.

(ii) For each $x \in \Pi_{\alpha\in A} Y_\alpha$, with $Y_\alpha = Y$, $u_\alpha(\cdot, x) : \Omega \to R$ is F-measurable.

(iii) For each $\omega \in \Omega$ and $x_{-\alpha} \in \Pi_{a\neq\alpha} X_a(\omega)$, $u_\alpha(\omega, x_{-\alpha}, \cdot) : X_\alpha(\omega) \to R$ is concave.

(iv) u_α is integrably bounded.

THEOREM 1. *Let $G^* = \{G^t : t \in T\}$ be a Bayesian game satisfying Assumptions* 1 *and* 2. *Then, there exists a NBNE for G^*.*

Proof. It follows from Kim and Yannelis (1997b, Lemma A.1) that the conditional expected utility $v_\alpha(\omega, x^t_{-\alpha}, x^t_\alpha(\omega))$ is weakly continuous for each $t \in T$. We need to show now that the total discounted expected utility,

$$U_\alpha(\omega, x_{-\alpha}, x_\alpha(\omega)) = \sum_{t\in T} \delta^t v_\alpha\left(\omega, x^t_{-\alpha}, x^t_\alpha(\omega)\right),$$

is weakly continuous as well.

Because the set T is countable, the desired result follows from Balder and Yannelis (1993, Corollary 2.9).[2]

Since U_α is weakly continuous, concave and F-measurable and the sets X^t_α satisfy Assumption 1, all conditions of the Kim and Yannelis (1997a) or Yannelis (1997) equilibrium existence theorem are satisfied and therefore we can conclude that a NBNE for G^* exists. ■

Remark. Because NBNE(G^*) $\subset$ NBNE$_\varepsilon(G^*)$, it also follows that NBNE$_\varepsilon$ $(G^*) \neq \emptyset$.

5. NONMYOPIC LEARNING

As we mentioned in the preceding sections, T denotes the time horizon. Agents enter the game having private information about the states of nature and they choose a sequence of actions to maximize the approximate (or ε) total discounted expected utility, given that the other players have chosen an ε-optimal strategy. At the end of each period, each player observes the equilibrium strategies of all the players. This observation generates new information, and agents refine their partitions. More formally, let $\sigma(x^t)$ denote the σ-algebra that the NBNE(G^*) generates in period t. Then, the information of player α in period $t+1$, denoted by F_α^{t+1}, is given by

$$F_\alpha^{t+1} = F_\alpha^t \vee \sigma(x^t),$$

where x^t is the projection of a NBNE(G^*) on the tth coordinate and $F_\alpha^t \vee \sigma(x^t)$ denotes the join, i.e., the smallest σ-algebra containing F_α^t and $\sigma(x^t)$. Consequently, for each player α,

$$F_\alpha^t \subseteq F_\alpha^{t+1} \subseteq F_\alpha^{t+2} \subseteq \cdots. \tag{3}$$

This represents a learning process for player α.

Now, let

$$\bar{F}_\alpha = \vee_{t\in T} F_\alpha^t, \tag{4}$$

where $\bar{F}_\alpha$ is interpreted as the pooled information of player α over the entire horizon T. A one-shot Bayesian game,

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

where $X_\alpha, u_\alpha, q_\alpha$ are defined as before and $\bar{F}_\alpha$ is given by (4), is called a *limit full-information Bayesian game*. $\bar{L}_X$ and NBNE($\bar{G}$) are defined for the Bayesian game $\bar{G}$ in a way analogous to that with L_X and NBNE(G^*) in the game G^*.

Note that $\bar{F}_\alpha$ may or may not be the same as the full-information $\vee_{\alpha\in A} F_\alpha$, which is the pooled information over all players.

6. NONMYOPIC LEARNING IN FINITE BAYESIAN GAMES

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

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

Proof. Let $\{x^t : t \in T\}$ be an element of NBNE$_\varepsilon(G^*)$. First, notice that, from Diestel's Theorem [see, e.g., Kim and Yannelis (1997b, Lemma A.3)], each $L_{X_\alpha^t}$ is weakly compact. From the measurability constraints, it follows that for each $t \in T$, $L_{X^t} \subset \bar{L}_X$. Because $x^t \in L_{X^t}$ for each t and $\bar{L}_X$ is weakly compact, it follows from the *Eberlein-Smulian Theorem* [Dunford and Schwartz (1958, p. 430)], that there exists a subsequence $\{x^{t_n} : n = 1, 2, \ldots\}$, such that x^{t_n} converges weakly to $x^* \in \bar{L}_X$. Hence, for each α, x_α^* is $\bar{F}_\alpha$-measurable.

Fix $\alpha \in A$ and $\omega \in \Omega$. Let $y_\alpha(\omega) \in X_\alpha(\omega)$ be a strategy in the limit full-information Bayesian game $\overline{G}$. We need to show that x^* is a NBNE for $\overline{G}$, i.e., that μ-a.e.,

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

Suppose by way of contradiction that for some player $\alpha \in A$ and for $D \subset \Omega$ with $\mu(D) > 0$, there exists $y_\alpha(\omega)$ such that, for all $\omega \in D$,

$$U_\alpha\big(\omega, x^*_{-\alpha}, y_\alpha(\omega)\big) > U_\alpha\big(\omega, x^*_{-\alpha}, x^*_\alpha(\omega)\big).$$

Let

$$\big[U_\alpha\big(\omega, x^*_{-\alpha}, y_\alpha(\omega)\big) - U_\alpha\big(\omega, x^*_{-\alpha}, x^*_\alpha(\omega)\big)\big] = \rho.$$

For each m, $(m = 1, 2, \ldots)$ and $\omega \in D$, set $y^m_\alpha = E[y_\alpha \mid F^m_\alpha]$. Note that

$$E\big[y_\alpha \,\big|\, F^m_\alpha\big] = E\big[E\big[y_\alpha \,\big|\, F^{m+1}_\alpha\big] \,\big|\, F^m_\alpha\big] = E\big[y^{m+1}_\alpha \,\big|\, F^m_\alpha\big].$$

Hence, $\{y^m_\alpha, F^m_\alpha\}^\infty_{m=1}$ is a martingale in $L_{X^t_\alpha} \subset L_1(\mu, Y)$ and by the martingale convergence theorem [see Diestel and Uhl (1977, Corollary 2, p. 126)] y^m_α converges [in the $L_1(\mu, Y)$-norm] and thus weakly to y_α.[3] It follows that (recall that U_α is weakly continuous) we can choose m_1 large enough so that, for $m \geq m_1$, we have[4]

$$\big|U_\alpha\big(\omega, x^*_{-\alpha}, y_\alpha(\omega)\big) - U_\alpha\big(\omega, x^m_{-\alpha}, y^m_\alpha(\omega)\big)\big| < \frac{\rho - \varepsilon}{2}$$

and

$$\big|U_\alpha\big(\omega, x^m_{-\alpha}, x^m_\alpha(\omega)\big) - U_\alpha\big(\omega, x^*_{-\alpha}, x^*_\alpha(\omega)\big)\big| < \frac{\rho - \varepsilon}{2},$$

where $(\rho - \varepsilon)/2 > 0$.

Thus,

$$\begin{aligned}&\big|U_\alpha\big(\omega, x^*_{-\alpha}, y_\alpha(\omega)\big) - U_\alpha\big(\omega, x^m_{-\alpha}, y^m_\alpha(\omega)\big) + U_\alpha\big(\omega, x^m_{-\alpha}, x^m_\alpha(\omega)\big)\\ &\quad - U_\alpha\big(\omega, x^*_{-\alpha}, x^*_\alpha(\omega)\big)\big| < \big|U_\alpha\big(\omega, x^*_{-\alpha}, y_\alpha(\omega)\big) - U_\alpha\big(\omega, x^m_{-\alpha}, y^m_\alpha(\omega)\big)\big|\\ &\quad + \big|U_\alpha\big(\omega, x^m_{-\alpha}, x^m_\alpha(\omega)\big) - U_\alpha\big(\omega, x^*_{-\alpha}, x^*_\alpha(\omega)\big)\big| < \frac{\rho - \varepsilon}{2} + \frac{\rho - \varepsilon}{2} = \rho - \varepsilon.\end{aligned}$$

Then, we have

$$\begin{aligned}&U_\alpha\big(\omega, x^*_{-\alpha}, y_\alpha(\omega)\big) - U_\alpha\big(\omega, x^m_{-\alpha}, y^m_\alpha(\omega)\big) + U_\alpha\big(\omega, x^m_{-\alpha}, x^m_\alpha(\omega)\big)\\ &\quad - U_\alpha\big(\omega, x^*_{-\alpha}, x^*_\alpha(\omega)\big) < \rho - \varepsilon,\end{aligned}$$

and by rearranging, we obtain

$$U_\alpha\big(\omega, x^m_{-\alpha}, y^m_\alpha\big) > U_\alpha\big(\omega, x^m_{-\alpha}, x^m_\alpha\big) + \varepsilon$$

for all $\omega \in D$ and for all $m \geq m_1$. Hence, we found a strategy $\{y^m_\alpha : m \geq m_1\}$ that is an approximate Bayesian Nash equilibrium for player α from period m_1

onward, while the other players have kept their strategies fixed. This contradicts the fact that $\{x^t : t \in T\} \in \mathrm{NBNE}_\varepsilon(G^*)$. ■

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

Proof. It follows if we set $\varepsilon = 0$ in the above proof. ■

THEOREM 3. *Let $G^* = \{G^t : t \in T\}$ be a Bayesian game satisfying Assumptions* 1 *and* 2 *and let x^* be an element of NBNE$(\bar{G})$. Then, for any $\varepsilon > 0$ and t' large enough, there exists $\{x^t : t \geq t'\}$ in NBNE$_\varepsilon(G^*)$, such that x^t converges in the $L_1(\mu, Y)$-norm to x^*.*

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

$$E\left[x^*_\alpha \mid F^t_\alpha\right] = E\left[E\left[x^*_\alpha \mid F^{t+1}_\alpha\right] \mid F^t_\alpha\right] = E\left[x^{t+1}_\alpha \mid F^t_\alpha\right].$$

Hence, $\{x^t_\alpha, F^t_\alpha\}_{t=1}^\infty$ is a martingale in $L_{X^t_\alpha} \subset L_1(\mu, Y)$ and, by the martingale convergence theorem, x^t_α converges in the $L_1(\mu, Y)$-norm to x^*. To complete the proof, it is enough to show that, for t large enough, $x^t \in \mathrm{NBNE}_\varepsilon(G^*)$. Suppose by way of contradiction that, for infinitely many t's, there exists D, with $\mu(D) > 0$ and $y^t_\alpha(\omega) \in \Pi_{t \geq t'} X^t_\alpha(\omega)$ such that

$$U_\alpha\left(\omega, x^t_{-\alpha}, y^t_\alpha(\omega)\right) > U_\alpha\left(\omega, x^t_{-\alpha}, x^t_\alpha(\omega)\right) + \varepsilon$$

for all $\omega \in D$. Because $X^t_\alpha(\omega) \subset X_\alpha(\omega)$ and the latter set is weakly compact, we can assume that $y^t_\alpha(\omega)$ converges weakly to some $y^*_\alpha(\omega)$ (by passing to a subsequence if necessary).[5] It follows from the weak continuity of U_α that, for all $\omega \in D$,

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

which contradicts the fact that $x^* \in \mathrm{NBNE}(\bar{G})$. ■

7. NONMYOPIC LEARNING IN BAYESIAN GAMES WITH A CONTINUUM OF PLAYERS

In this section, we study the Bayesian game $G^* = \{G^t : t \in T\}$, where the set of players is a measure space. A Bayesian game with a measure space of agents $(A, \mathcal{A}, \nu)$ is a sequence of games $\{G^t : t \in T\}$ such that, for each t, $G^t = \{(F^t_\alpha, X^t_\alpha, u_\alpha, q_\alpha) : \alpha \in A\}$, where

1. F^t_α, is a sub-σ-algebra of F which denotes the *private information* of agent α in period t.
2. $X^t : A \times \Omega \to 2^Y$, is the *action set-valued function*, where $X^t(\alpha, \omega)$ is the set of actions available to agent α in period t when the state is ω, which is F^t_α-measurable.

3. For each $(\alpha, \omega) \in A \times \Omega$, $u(\alpha, \omega, \cdot, \cdot) : L_1(\nu, Y) \times X^t(\alpha, \omega) \to R$ is the *utility function* of agent α, using action $x^t_\alpha(\omega)$, when the state is ω and the other players use the joint action x^t.
4. $q_\alpha : \Omega \to R_{++}$, is the *prior* of agent α [q_α is a density function, or Radon-Nikodym derivative, such that $\int_{\omega \in \Omega} q_\alpha(\omega)\, d\mu(\omega) = 1$].

As before, let $L_{X^t_\alpha}$ denote the set of all Bochner-integrable and F^t_α-measurable selections from the action set-valued function $X^t(\alpha, \omega)$ of agent α in period t, i.e.,

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

Let

$$L_{X^t} = \big\{x^t \in L_1(\nu, L_1(\mu, Y)) : x^t(\alpha) \in L_{X^t_\alpha},\ \nu\text{-a.e.}\big\}.$$

A typical element of L_{X^t} is denoted by x^t.

Let L_X be the product of L_{X^t} over all $t \in T$, i.e., $L_X = \Pi_{t \in T} L_{X^t}$. A typical element of L_X is denoted by $x = \{x^t : t \in T\}$ and is a sequence of strategy profiles over the entire horizon, where each element of the sequence belongs to L_{X^t}, i.e., $x^t \in L_{X^t}$. A strategy of agent α is an element of $L_{X_\alpha} = \Pi_{t \in T} L_{X^t_\alpha}$ denoted by $\{x^t(\alpha) : t \in T\}$. For each $(\alpha, \omega) \in A \times \Omega$, the conditional expected utility function of agent α, $v(\alpha, \omega, \cdot, \cdot) : L_{X^t} \times X^t(\alpha, \omega) \to R$ is defined as

$$v(\alpha, \omega, x^t, x^t(\alpha, \omega)) = \int_{\omega' \in E^t_\alpha(\omega)} u(\alpha, \omega, x^t(\omega'), x^t(\alpha, \omega)) q_\alpha\big(\omega' \,\big|\, E^t_\alpha(\omega)\big)\, d\mu(\omega'),$$

where

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

For each $\omega \in \Omega$, the *total discounted conditional (interim) expected utility* of player α, $U(\alpha, \omega, \cdot, \cdot) : L_X \times \Pi_{t \in T} X^t(\alpha, \omega) \to R$ is given by

$$U(\alpha, \omega, x, x(\alpha, \omega)) = \sum_{t \in T} \delta^t v(\alpha, \omega, x^t, x^t(\alpha, \omega)).$$

A *nonmyopic Bayesian Nash equilibrium* for G^* is a strategy profile $x^* \in L_X$ such that, for ν-a.e. and μ-a.e.,

$$U(\alpha, \omega, x^*, x^*(\alpha, \omega)) = \max_{y \in \Pi_{t \in T} X^t(\alpha, \omega)} U(\alpha, \omega, x^*, y(\alpha, \omega)).$$

We can now state the assumptions needed for the proof of the next theorem.[6]

Assumption 3.

(i) $X^t : A \times \Omega \to 2^Y$ is a non-empty, convex, weakly compact valued and integrably bounded correspondence having $A \otimes F$-measurable graph, i.e., $G_{X^t} \in \mathcal{A} \otimes F \otimes \mathcal{B}(Y)$.

(ii) For each $\alpha \in A$ and $t \in T$, $X^t(\alpha, \cdot) : \Omega \to 2^Y$ has an F^t_α-measurable graph, i.e., $G_{X(\alpha)} \in F^t_\alpha \otimes \mathcal{B}(Y)$.

Assumption 4.

(i) For each $(\alpha, \omega) \in A \times \Omega$ and $t \in T$, $u(\alpha, \omega, \cdot, \cdot) : L_1(\nu, Y) \times X^t(\alpha, \omega) \to R$ is continuous where $L_1(\nu, Y)$ and $X^t(\alpha, \omega)$ 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 $\mathcal{A} \otimes F$-measurable.

(iii) For each $\alpha \in A$, $u(\alpha, \cdot, \cdot, \cdot)$ is integrably bounded.

Assumption 5.

(i) The dual Y^* of Y has the RNP with respect to $(A, \mathcal{A}, \nu)$.

Given an $\varepsilon > 0$, the strategy profile $x^* \in L_X$ is said to be a NBNE$_\varepsilon(G^*)$ if, for μ-a.e. and ν-a.e.,

$$U(\alpha, \omega, x^*, x^*(\alpha, \omega) \geq U(\alpha, \omega, x^*, y(\alpha, \omega)) - \varepsilon$$

for all $y(\alpha, \omega) \in \Pi_{t \in T} X^t(\alpha, \omega)$.

THEOREM 4. *Let $G^* = \{G^t : t \in T\}$ be a Bayesian game satisfying Assumptions* 3–5 *and let $\{x^t : t \in T\}$ be a sequence in NBNE$_\varepsilon(G^*)$. Then, there exists a subsequence $\{x^{t_n} : n = 1, 2, \ldots\}$ of $\{x^t : t \in T\}$ such that $\{x^{t_n} : n = 1, 2, \ldots\}$ converges weakly to $x^* \in$ NBNE$(\bar{G})$.*

Proof. Let $\{x^t : t \in T\}$ be an element of NBNE$_\varepsilon(G^*)$. Recall that each $L_{X^t_\alpha}$ is weakly compact. From the measurability constraints, it follows that, for each $t \in T$, $L_{X^t} \subset \bar{L}_X$. Because $x^t \in L_{X^t}$ for each t and $\bar{L}_X$ is weakly compact, it follows from the *Eberlein-Smulian Theorem* [Dunford and Schwartz (1958, p. 430)] that there exists a subsequence $\{x^{t_n} : n = 1, 2, \ldots\}$ such that x^{t_n} converges weakly to $x^* \in \bar{L}_X$. Hence, for each α, x^*_α is $\bar{F}_\alpha$-measurable.

Fix $\alpha \in A$ and $\omega \in \Omega$. Let $y(\alpha, \omega) \in \bar{X}(\alpha, \omega)$ be a strategy in the limit full-information game. We need to show that x^* is a NBNE for $\bar{G}$, i.e., that μ-a.e. and ν-a.e.,

$$U(\alpha, \omega, x^*, x^*(\alpha, \omega)) \geq U(\alpha, \omega, x^*, y(\alpha, \omega))$$

for all $y(\alpha, \omega)$ in the limit full-information game.

Suppose by way of contradiction that, for some $M \subset A$ and $D \subset \Omega$ with $\nu(M) > 0$ and $\mu(D) > 0$, there exists $y(\alpha, \omega)$ in the limit full-information game such that, for all $\alpha \in M$ and $\omega \in D$,

$$U(\alpha, \omega, x^*, y(\alpha, \omega)) > U(\alpha, \omega, x^*, x^*(\alpha, \omega)).$$

Let

$$[U(\alpha, \omega, x^*, y(\alpha, \omega)) - U(\alpha, \omega, x^*, x^*(\alpha, \omega))] = \rho > 0.$$

For each m ($m = 1, 2, \ldots$) and $\omega \in D$, set $y^m_\alpha = E[y_\alpha | F^m_\alpha]$. Note that

$$E\left[y_\alpha \mid F^m_\alpha\right] = E\left[E\left[y_\alpha \mid F^{m+1}_\alpha\right] \mid F^m_\alpha\right] = E\left[y^{m+1}_\alpha \mid F^m_\alpha\right].$$

Hence, $\{y_\alpha^m, F_\alpha^m\}_{m=1}^\infty$ is a martingale in $L_{X_\alpha^t} \subset L_1(\mu, Y)$ and by the martingale convergence theorem y_α^m converges [in the $L_1(\mu, Y)$-norm] and thus weakly to y_α. It follows that (recall that U_α is weakly continuous) we can choose m_1 large enough so that, for $m \geq m_1$, we have[7]

$$|U(\alpha, \omega, x^*, y(\alpha, \omega)) - U(\alpha, \omega, x^m, y^m(\alpha, \omega))| < \frac{\rho_- \varepsilon}{2}$$

and

$$|U(\alpha, \omega, x^m, x^m(\alpha, \omega)) - U(\alpha, \omega, x^*, x^*(\alpha, \omega))| < \frac{\rho - \varepsilon}{2}.$$

Thus,

$$\begin{aligned}
&|U(\alpha, \omega, x^*, y(\alpha, \omega)) - U(\alpha, \omega, x^m, y^m(\alpha, \omega)) + U(\alpha, \omega, x^m, x^m(\alpha, \omega)) \\
&\quad - U(\alpha, \omega, x^*, x^*(\alpha, \omega))| < |U(\alpha, \omega, x^*, y(\alpha, \omega)) \\
&\quad - U(\alpha, \omega, x^m, y^m(\alpha, \omega))| + |U(\alpha, \omega, x^m, x^m(\alpha, \omega)) \\
&\quad - U(\alpha, \omega, x^*, x^*(\alpha, \omega))| < \frac{\rho - \varepsilon}{2} + \frac{\rho - \varepsilon}{2} = \rho - \varepsilon.
\end{aligned}$$

Then, we have

$$\begin{aligned}
&U(\alpha, \omega, x^*, y(\alpha, \omega)) - U(\alpha, \omega, x^m, y^m(\alpha, \omega)) + U(\alpha, \omega, x^m, x^m(\alpha, \omega)) \\
&\quad - U(\alpha, \omega, x^*, x^*(\alpha, \omega)) < \rho - \varepsilon,
\end{aligned}$$

and by rearranging, we obtain

$$U(\alpha, \omega, x^m, y^m(\alpha, \omega)) > U(\alpha, \omega, x^m, x^m(\alpha, \omega)) + \varepsilon$$

for all $\alpha \in M$ and $\omega \in D$ and for all $m \geq m_1$, a contradiction to the fact that $\{x^t\}_{t\in T} \in \text{NBNE}_\varepsilon(G^*)$. ■

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

Proof. It follows if we set $\varepsilon = 0$ in the above proof. ■

THEOREM 5. *Let $G^* = \{\underline{G}^t : t \in T\}$ be a Bayesian game satisfying Assumptions 3–5 and let x^* in NBNE$(\underline{G})$. Then, for any $\varepsilon > 0$ and t' large enough, there exists $\{x^t : t \geq t'\}$ in NBNE$_\varepsilon(G^*)$ such that x^t converges in the $L_1(\mu, Y)$-norm to x^*.*

Proof. Let $x_\alpha^t = E[x_\alpha^* \mid F_\alpha^t]$. Note that

$$E\left[x_\alpha^* \mid F_\alpha^t\right] = E\left[E\left[x_\alpha^* \mid F_\alpha^{t+1}\right] \mid F_\alpha^t\right] = E\left[x_\alpha^{t+1} \mid F_\alpha^t\right].$$

Hence, $\{x_\alpha^t, F_\alpha^t\}_{t=1}^\infty$ is a martingale in $L_{X_\alpha^t} \subset L_1(\mu, Y)$ and, by the martingale convergence theorem, x_α^t converges in the $L_1(\mu, Y)$-norm to x^*. To complete the proof, it is enough to show that, for t large enough, $x^t \in \text{NBNE}_\varepsilon(G^*)$. Suppose by way of contradiction that, for infinitely many t', there exists $M \subset A$ and $D \subset \Omega$, with $\nu(M) > 0$, $\mu(D) > 0$, and $y^t(\alpha, \omega) \in \Pi_{t \geq t'} X^t(\alpha, \omega)$ such that

$$U(\alpha, \omega, x^t, y^t(\alpha, \omega)) > U(\alpha, \omega, x^t, x^t(\alpha, \omega)) + \varepsilon$$

for all $\alpha \in M$ and $\omega \in D$.[8] Because $X^t(\alpha, \omega) \subset X(\alpha, \omega)$ and the latter set is weakly compact, we can assume that $y^t(\alpha, \omega)$ converges weakly to some $y^*(\alpha, \omega)$, by passing to a subsequence if necessary. Then, it follows from the weak continuity of $U(\alpha)$ that, for all $\alpha \in M$ and $\omega \in D$,

$$U(\alpha, \omega, x^*, y^*(\alpha, \omega)) \geq U(\alpha, \omega, x^*, x^*(\alpha, \omega)) + \varepsilon > U(\alpha, \omega, x^*, x^*(\alpha, \omega)),$$

which contradicts the fact that $x^* \in \text{NBNE}(\bar{G})$. ■

8. CONCLUSIONS

We showed that the assumption of myopia can be disregarded from Bayesian learning at no real cost. In particular, players can choose strategies by taking into account future actions and still by repetition the nonmyopic bounded rational players can reach a limit full-information Bayesian Nash equilibrium outcome. The converse is also true, i.e., given a limit full-information Bayesian Nash equilibrium outcome, we can construct a sequence of nonmyopic bounded rational plays converging to the limit full-information Bayesian Nash outcome.

NOTES

1. X^t_α depends on t only through the measurability constraint.
2. See also Section 2.
3. y^m_α converges to $y^\infty_\alpha = E[y_\alpha \mid \bar{F}_\alpha]$, which is equal to y_α because, by construction, y_α is in the limit full-information game and hence is $\bar{F}_\alpha$ -measurable.
4. By $y^m_\alpha(\omega)$, we mean $\{y^m_\alpha(\omega) : m \geq m_1\}$ and similarly for $x^m_{-\alpha}$.
5. We say that there exists a time period t' such that the strategy of a player is a NBNE from that period onward. Thus, when we write y^t_α, we really mean that $\{y^t_\alpha : t \geq t'\}$. The same applies to $x^t_{-\alpha}$.
6. Notice that the assumptions below also guarantee the existence of a NBNE for G^* by appealing to Kim and Yannelis (1997a) and by recalling that, for each $(\alpha, \omega) \in A \times \Omega$, $U(\alpha, \omega, \cdot, \cdot)$ is weakly continuous. The latter follows directly from Kim and Yannelis (1997b, Lemma A.2) and Balder and Yannelis (1993, Corollary 2.9).
7. By $y^m_\alpha(\omega)$, we mean $\{y^m_\alpha(\omega) : m \geq m_1\}$ and similarly for x^m.
8. We say that there exists a time period t' such that the strategy of a player is a NBNE from that period onward. Thus, when we write y^t_α, what we really mean is $\{y^t_\alpha\}_{t \geq t'}$. The same applies to $x^t_{-\alpha}$.

REFERENCES

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*, vol. 15, Mathematical Surveys. Providence, RI: American Mathematical Society.
Dunford, N. & J.T. Schwartz (1958) *Linear Operators*, pt. 1. New York: Interscience.

Kalai, E. & E. Lehrer (1993) Rational learning leads to Nash equilibrium. *Econometrica* 61, 1019–1045.

Koutsougeras, L. & N.C. Yannelis (1994) Convergence and approximation results for noncooperative Bayesian games: Learning theorems. *Economic Theory* 4, 843–857.

Kim, T. & N.C. Yannelis (1997a) (in press) Existence of equilibrium in Bayesian games with infinitely many agents. *Journal of Economic Theory* 77, 330–353.

Kim, T. & N.C. Yannelis (1997b) Learning in Bayesian games by bounded rational players I. *Macroeconomic Dynamics* 1, 568–587.

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

Yannelis, N.C. (1997) (in press) On the existence of a Bayesian Nash equilibrium. In Y. Abramovich, E. Avgerinos & N.C. Yannelis (eds.), *Functional Analysis Economic Theory.* Heidelberg: Springer.