Economic Theory 11, 643–655 (1998)

Economic Theory

Exposita Notes

Bayesian learning and convergence to Nash equilibria without common priors★

Yaw Nyarko
Department of Economics, New York University, 269 Mercer Street, New York, NY10003, USA

Received: March 3, 1995; revised version: February 17, 1997

Summary. Consider an infinitely repeated game where each player is characterized by a "type" which may be unknown to the other players in the game. Suppose further that each player's belief about others is independent of that player's type. Impose an absolute continuity condition on the *ex ante* beliefs of players (weaker than mutual absolute continuity). Then any limit point of beliefs of players about the future of the game conditional on the past lies in the set of Nash or Subjective equilibria.

Our assumption does not require common priors so is weaker than Jordan (1991); however our conclusion is weaker, we obtain convergence to subjective and not necessarily Nash equilibria. Our model is a generalization of the Kalai and Lehrer (1993) model. Our assumption is weaker than theirs. However, our conclusion is also weaker, and shows that limit points of beliefs, and not actual play, are subjective equilibria.

JEL Classification Numbers: C70, C73, D81, D82, D83, D84.

1 Introduction

1.1 Discussion of main ideas

Consider a finite collection of players in an infinitely repeated game. Suppose that each player is characterized by a "type" which is not necessarily known to the other players of the game. Impose only two conditions on players. First, suppose each player obeys the Savage (1954) axioms: In particular, each player has a prior probability belief over the set of types of all the

★ I am deeply indebted to Professor Jim Jordan who first interested me in the problem studied in this paper. I thank Professors Beth Allen, J-P Benoit, L. Blume, David Easley and Ehud Kalai for very helpful comments. I am also very grateful to the C.V. Starr Center, the Challenge Fund and the Presidential Fellowship at New York University for their generosity. I am also grateful to two anonymous referees for their comments and conjectures.

players as well as the actions over time each player-type will choose. Each player then maximizes her expected utility for the infinite horizon game given that belief. Second, suppose that the beliefs of players are such that if one player assigns probability zero to an event, then all other players assign probability zero to that event. We show that under conditions slightly weaker than these two, the beliefs of players converge to the set of Nash or subjective equilibria.

The motivation of this paper is the same as that of Blume and Easley (1984) much earlier: We seek to determine conditions under which players initially not in equilibrium can "learn" their way to an equilibrium. The results reported in this paper are in one sense a generalization of the results of Jordan (1991 and 1995). Jordan studied the same model as presented here, but assumed that players' priors are the same. We generalize the results of Jordan by showing that the convergence still occurs when we weaken the common prior assumption by requiring merely that players' "uncommon" priors are mutually absolutely continuous with respect to each other; i.e., they assign probability zero to the same events. (The condition we use, which we refer to as condition (GGH), is weaker than this.) Without common priors, however, when the discount factor is positive we obtain convergence to subjective as opposed to Nash equilibria, so our conclusion, i.e., the convergence result, is weaker than that of Jordan. Our condition (GGH) is one of the weakest in the literature under which a convergence result is possible. Without it we do not believe much can be said about the players' limiting behavior. For an example of what could go wrong without (GGH) see Nyarko (1991). Like Jordan (1991 and 1995) our results shall state that *beliefs* of players converge to the set of Nash or subjective equilibria. The *actual* play over time does not necessarily converge to the set of such equilibria. In Section 1.2 we provide an example to illustrate all this.

Kalai and Lehrer (1993a) (henceforth KL93) obtain results on the convergence of *actual* play to the set of subjective equilibria under an absolute continuity condition. In Section 4, we provide an analogue to the KL93 assumption, in the context of the model with many types, requiring their assumption to hold "for almost every" vector of types. We refer to this assumption as (KL). We show that condition (KL) is *stronger than* (GGH), the condition we use in this paper. The example in Section 1.2 will satisfy (GGH) but will fail (KL).

Like Jordan (1991 and 1995) we shall suppose that each player's belief about the other players is independent of that player's own type. In Nyarko (1994a) this assumption is relaxed and shows that without the type-independence assumption, the limit points of beliefs are not necessarily Nash equilibria; instead they are correlated equilibria. Nyarko (1994a) is one of the few papers which gives a convergence to correlated equilibrium result. However that paper applies only to the zero discount factor problems. It remains an open question whether the correlated equilibrium result extends to the positive discount factor problem. Indeed, it is not clear what the appropriate concept of correlated equilibrium should be used for that problem. The main result of this paper, although it allows for positive discount factors, is of little help in directly extending that paper[1].

The Nyarko (1994a) paper has a second result: the convergence of sample path averages of true play (i.e., the empirical distributions) have the same limit points as beliefs. When combined with the main result of this paper, we obtain a link between the convergence of beliefs result here and the convergence of empirical distributions. Jordan (1996) has also independently obtained related results on the convergence of sample path averages.

1.2 Example (*coin-tossing*)

To illustrate our result on the convergence of beliefs as opposed to actual play we present the following example. Consider a game with two players A and B. Player A (resp. B) has two actions to choose from at each date, TOP and BOTTOM (resp. LEFT and RIGHT). The stage-game payoffs are as below:

		Player B	
		LEFT	RIGHT
Player A	TOP	1,1	0,0
	BOTTOM	0,0	1,1

Players have a discount factor of 0 (i.e., seek to maximize stage game payoffs in each period).

Let τ^A be a realization from infinitely many independent and identical coin-tossing experiments where an outcome from {HEADS, TAILS} is chosen with equal probability. Hence τ^A is an element of {HEADS, *TAILS*}$^\infty$. Let τ^B be another realization from an i.i.d sequence of coin-tosses, {HEADS, TAILS}$^\infty$, which is independent of the sequence from which τ^A was obtained. At date 0 Player A is told of τ^A (and is not informed about τ^B) and player B is told of τ^B (and is not informed about τ^A). We may consider τ^A to be player A's "type" and τ^B to be player B's type. Suppose that each player knows how the types are drawn. Consider the following play of the game. At date n Player A looks at the n-th coordinate of his sequence of coin-tosses. If it is a HEADS he plays his first action, TOP; if it is a TAILS he plays his second action, BOTTOM. Similarly, if the n-th element of τ^B is HEADS player B plays the action LEFT at date n, otherwise he plays action

[1] It should also be remarked that the method of proof of the two convergence results are also radically different. Nyarko (1994a) uses a separating hyperplane argument. The result of this paper uses the original Jordan (1995) result and Blackwell and Dubins (1963) Theorem. It should be emphasized that the techniques used in the zero discount factor problem will not work in positive discount factor case. Further, with zero discount factor and independent types the convergence is to Nash equilibria, while here it is to subjective equilibria.

RIGHT. Suppose further that each player knows that the other is choosing actions via this rule. The beliefs of each player about the other assign equal probability to each action. Hence the vector of beliefs of players (not conditioning on types) is a Nash equilibrium. However, on each sample path (TOP, RIGHT) and (BOTTOM, LEFT) will be played infinitely often. Note that these pairs of actions are *not* Nash equilibria of the stage game. Hence limit points of beliefs, and not actual play, are Nash equilibria. □

2 The model

2.1 Some terminology

I is the *finite* set of players. Given any collection of sets[2] $\{X_i\}_{i\varepsilon I}$, we define $X \equiv \prod_{i\varepsilon I} X_i$ and $X_{-i} \equiv \prod_{j\neq i} X_j$. Given any collection of functions $f_i : X_i \to Y_i$ for $i\varepsilon I$, $f_{-i} : X_{-i} \to Y_{-i}$ is defined by $f_{-i}(x_{-i}) \equiv \prod_{j\neq i} f_j(x_j)$. The cartesian product of metric spaces will always be endowed with the product topology. Let X be any metric space. We let $\mathscr{P}(X)$ denote the set of probability measures defined over (Borel) subsets of X. Unless otherwise stated the set $\mathscr{P}(X)$ will be endowed with the weak topology. If X is a cartesian product $X = YZ$, we denote by $\text{Marg}_Y\, v$ the marginal of v on Y. Suppose for $k = 1, 2, \ldots, K$, v_k is a probability measure on a complete and separable metric space X_k. Then we let $\prod_{k=1}^K v_k$ denote the product measure of $\{v_k\}_{k=1}^K$ on the cartesian product $\prod_{k=1}^K X_k$.

2.2 Actions and strategies

I is the *finite* set of players. S_i represents the *finite* set of actions available to player i at each date $n = 1, 2, \ldots$; also $S \equiv \prod_{i\varepsilon I} S_i$. Even though the action space S_i is independent of the date we shall sometimes write S_i as S_{in} and S as S_n when we seek to emphasize the set of action choices at *date n*. We define $S^N \equiv \prod_{n=1}^N S_n$ and $S^\infty \equiv \prod_{n=1}^\infty S_n$, the set of date N and infinite histories, respectively. s^0 will denote the null history, (at date 0, when there is no history)! In summary, s or S with a "superscript" (e.g., s^N) denotes the history, while with a "subscript" (e.g., s_n) denotes the current period. *Perfect recall* is assumed: at date n when choosing the date n action s_{iN}, player i will have information on $s^{N-1} = \{s_1, \ldots, s_{N-1}\}$. We define $F_{iN} \equiv \{f_{iN} : S^{N-1} \to \mathscr{P}(S_i)\}$; $F_i \equiv \prod_{N=1}^\infty F_{iN}$ and $F \equiv \prod_{i\varepsilon I} F_i$. F_i is the set of all behavior

[2] In the interest of ease of exposition the following conventions will be adopted in the paper without explicitly mentioning them: All spaces in this paper will be metric spaces. For such a space, any generic subset of a space will be assumed to be a Borel subset, any generic function will be assumed to be Borel measurable and any generic probability will be assumed to be defined on Borel subsets. These qualifiers will usually not be repeated in the text. All statements in this paper involving conditional probabilities will usually be independent of the version chosen. Hence we shall make statements like "*the* conditional probability ..." when we should be saying "*any version of* the conditional probability." Such versions will also be assumed to be regular (i.e., they are Borel measurable in their conditioning arguments). $\mathbb{R}$ denotes the real line.

strategies for player i. F_{iN} is endowed with the topology of pointwise convergence; F_i and F are endowed with their respective product topologies. The mapping $m : F \to \mathscr{P}(S^\infty)$ defines the probability distribution $m(f)$ on S^∞ resulting from the behavior strategy profile f; i.e., induced by the following transition equation:

$$m(f)\left(D|s^N\right) \equiv f_{N+1}\left(s^N\right)(D) \qquad \forall D \subseteq S_{N+1} \ . \tag{2.1}$$

2.3 The type space and payoffs

Each player $i\varepsilon I$ has an *attribute vector* which is some element θ_i of the set Θ_i. $u_i : \Theta_i \times S \to \mathbb{R}$ is player i's within-period *continuous and bounded* utility function which depends upon her attribute vector, θ_i, as well as the vector of actions, $s\varepsilon S$, chosen by the players. Each player i knows her own attribute vector θ_i but does not necessarily know those of other players, θ_{-i}. We shall suppose that Θ_i is a *compact* subset of finite dimensional Euclidean space. Player i has a discount factor which is a (known) *continuous* function, $\delta_i : \Theta_i \to [0, 1)$, of the player i's attribute vector. We define $U_i : \Theta_i \times S^\infty \to \mathbb{R}$ by $U_i(\theta_i, s^\infty) = \sum_{n=1}^{\infty} [\delta_i(\theta_i)]^{n-1} u_i(\theta_i, s_n)$ where $s^\infty \equiv \{s_n\}_{n=1}^{\infty}$ and $V_i : \theta_i \times F \to \mathbb{R}$ by $V_i(\theta_i, f) = \int_{S^\infty} U_i(\theta_i, s^\infty) dm(f)$.

Each player i is characterized by a type, τ_i. The type τ_i will specify, among other things, that player's attribute vector, θ_i. We let T_i denote the set of possible types of player i, and set $T \equiv \prod_{i\varepsilon I} T_i$. We let $\theta_i(\tau_i)$ denote the attribute vector of player i of type τ_i. We assume that $\theta_i(\tau_i)$ is *continuous* in τ_i; (by appropriately defining the type space this can be shown to be without loss of generality). The type space will be assumed to be a *complete and separable metric space*. We define

$$\Omega \equiv T \times F \times S^\infty \ . \tag{2.2}$$

2.4 The prior beliefs of players

We shall say that a probability measure $\mu\,\varepsilon\,\mathscr{P}(T \times F \times S^\infty)$ *respects the probability* $m(f)$ if $m(f)$ is a version of the conditional probability, $\mu(.|\tau, f)$, of μ over S^∞ conditional on $(\tau, f)\varepsilon T \times F$. An *ex ante subjective belief* for a player i is any probability μ_i over $T \times F \times S^\infty$ which respects the probability $m(f)$. We will use the following assumptions on μ_i:

$$\mu_i\left(M_i \cup M_i^0\right) = 1 \tag{2.3}$$

where

$$M_i \equiv \left\{ (\tau, f, s^\infty)\varepsilon\Omega : \delta_i(\theta_i) > 0 \text{ and } f_i \text{ maximizes} \right.$$
$$\left. \int_{F_{-i}} V_i(\theta_i(\tau_i), ., f_{-i}) d\mu_{\rm i}(.|\tau_i) \right\}$$

and

$$M_i^0 \equiv \left\{ (\tau, f, s^\infty)\varepsilon\Omega : \delta_i(\theta_i) = 0 \text{ and } \forall n, s_{in+1} \text{ maximizes} \right.$$

$$\left. \int_{S_{-i}} u_i(\theta_i(\tau_i), ., s_{-in+1}) d\mu_i(.|s^n, \tau_i) \right\} .$$

$$\text{Marg}_F \mu_i(.|\tau) = \prod_{j\varepsilon I} \text{Marg}_{F_j} \mu_i(.|\tau_j) \text{ for } \mu_i \text{ almost every } \tau . \quad (2.4)$$

$\text{Marg}_T \mu_i$ is a product measure on the type space T;
i.e.,

$$\text{Marg}_T \mu_i = \prod_{j\varepsilon I} \left[\text{Marg}_{T_j} \ \mu_i \right] . \quad (2.5)$$

Assumption (2.3) requires that with μ_i probability one each player i is maximizing subjective expected discounted sum of utilities. Whenever the discount factor is equal to zero (i.e., on the set M_i^0 above) player i will be required to maximize expected utility at each date. (2.3) does not imply that under i's belief about the game other players $j \neq i$ are maximizing their expected utility. Assumption (2.4) says that other than through their types, players have no way of correlating the choice of their behavior strategies. Assumption (2.5) is a restriction on the beliefs of players over the type space and requires them to be product measures[3].

Given any two probability measures μ' and μ'' on some (measure) space Ω, μ' is absolutely continuous with respect to μ'' if for all (measurable) subsets D of $\Omega, \mu'(D) > 0$ implies that $\mu''(D) > 0$. We then write $\mu' \ll \mu''$. μ' and μ'' are mutually absolutely continuous if $\mu' \ll \mu''$ and $\mu'' \ll \mu'$. Fix a collection of *ex ante* subjective beliefs of players, $\{\mu_i\}_{i\varepsilon I}$. Fix any measure μ^* over Ω which respects the probability $m(f)$. Then μ^* and $\{\mu_i\}_{i\varepsilon I}$ obey the *generalized Harsanyi consistency condition*, (GGH), if

$$\text{(GGH)} \qquad \text{Marg}_{T_i \times S^\infty} \mu^* \ll \text{Marg}_{T_i \times S^\infty} \mu_i \qquad \text{for all } i\varepsilon I . \quad (2.6)$$

The usual Harsanyi (1967) consistency condition requires $\mu_i = \mu_j$ for all $i, j\varepsilon I$, in which case by setting $\mu^* = \mu_i$ we obtain (GGH). Elsewhere (Nyarko (1997a)) we have defined Condition (GH) to be where μ_i and μ_j are mutually absolutely continuous with respect to each other. (GGH) above generalizes this latter condition (hence the name "GGH") by first requiring merely absolute (and not mutual absolute) continuity, and this with respect to marginals. The measure μ^* in (GGH) is often interpreted as the "true distribution."

[3] Assumptions (2.4) and (2.5) imply that for each i, μ_i is a product measure over $\prod_{j\varepsilon I} T_j \times F_j$ (which is part of the definition of a Bayesian Strategy Process in Jordan (1995)). Nyarko (1994a) has shown that without (2.5) in general there is no longer convergence of beliefs to Nash equilibria; rather the convergence is to *correlated* equilibria.

2.5 *Kuhn strategic representation of* (*type unconditional*) *beliefs* (*KSRBs*)

For each $i \varepsilon I$, we define equivalence class relation, $\sim$, on F_i as follows: For each f_i and $f_i' \varepsilon F_i$, $f_i \sim f_i'$ if for all $f_{-i} \varepsilon F_{-i}$, $m(f_i, f_{-i}) = m(f_i', f_{-i})$. Let $F_i \sim$ denote the set of equivalence classes of $\sim$. From Kuhn's (1953) Theorem and (Aumann (1964) for the infinite horizon case) we may conclude that there is a function $\kappa_i : (\mathscr{P})F_i \rightarrow F_i \sim$ such that for any (mixed strategy) $\varphi_i \varepsilon \mathscr{P}(F_i)$ and any $f_i \varepsilon \kappa(\varphi_i)$ and any $f_{-i} \varepsilon F_{-i}$, the probability distribution on S^∞ induced by φ_i and f_{-i} is equal to $m(f_i, f_{-i})$. The behavior strategy $f_i \varepsilon \kappa(\varphi_i)$ is said to be *realization equivalent* to the mixed strategy φ_i.

Fix any collection of *ex ante* subjective beliefs of players $\{\mu_i\}_{i \varepsilon I}$. Fix any

$$f_j^{*i} \ \varepsilon \ \kappa(\text{Marg}_{F_j} \mu_i(.)) \qquad \text{and define} \qquad f^{*i} \equiv \{f_j^{*i}\}_{j \varepsilon I} \ . \tag{2.7}$$

From (2.4) and (2.5) μ_i is a product measure over $\prod_{j \varepsilon I} [T_j \times F_j]$. Consider an outside observer with belief $\mu_i \varepsilon \mathscr{P}(\Omega)$ who never observes the types of the players. Such an outside observer will have a belief $\prod_{j \varepsilon I} \text{Marg}_{F_j} \mu_i$ over F. Denote this by $\mu_i(df)$ and observe that it is realization equivalent to f^{*i}. Hence we may refer to f^{*i} as the *Kuhn strategic representation of beliefs* (*or KSRB*) of $\mu_i(df)$. We stress here that this representation does *not* condition on types!

The *shift operator* $\sigma_{iN} : S^N \times F_i \rightarrow F_i$ is defined by setting for each $s^N \varepsilon S^N$ and $f_i \varepsilon F_i$, $\sigma_{iN}(s^N, f_i) \equiv f_i'$ where the date n coordinate of f_i' is defined by $f_{in}'(.) \equiv f_{iN+n}(s^N, .)$. If f^{*i} is the KSRB of the outside observer with belief μ_i who does not observe players' types, then $f_N^{*i} \equiv \sigma_{iN}(s^N, f^{*i})$ is the KSRB of that outside observer's belief about the "future" following the date N history s^N.

2.6 *Nash and subjective equilibria*

Fix any attribute vector $\theta \equiv \{\theta_i\}_{i \varepsilon I} \varepsilon \Theta$. Define $\forall i \varepsilon I$

$$N_i(\theta_i) \equiv \left\{ f = \prod_{j \varepsilon I} f_j \varepsilon F : f_i \varepsilon \ \text{argmax} \ V_i(\theta_i, f_{-i}, \cdot) \right\} \ ;$$

$$N(\theta) \equiv \cap_{i \varepsilon I} N_i(\theta_i) \ ;$$

and

$$\text{ND}(\theta) \equiv \{ v \varepsilon \mathscr{P}(S^\infty) \text{: there exists an } f \varepsilon N(\theta) \text{ with } v = m(f) \} \ .$$

$N(\theta)$ is the set of Nash equilibrium behavior strategy profiles for the complete information game with *fixed attribute* vector θ. ND(θ) is the set of all distributions of play that can be generated by some Nash equilibrium behavior strategy profile. ND(θ) is also equal to the set of all distributions that are induced by some *subjective equilibrium* profile of strategies. (For definition of latter see Battigali et al. (1988) and (1992) or Kalai-Lehrer (1993b).)

3 Convergence to Nash equilibria

Given any $\mu \varepsilon \mathscr{P}(\Omega)$, and any date N, let $\mu(.|s^N)$ be (any fixed version of) the conditional probability of μ given the history s^N. Let $\mu_N(ds^{N++}|s^N)$ denote the probability distribution over the "future", $s^{N++} \equiv \{s_{N+1}, s_{N+2}, \ldots\} \varepsilon S^\infty$ conditional on the "past," s^N, induced by the measure $\mu(.|s^N)$. The norm $\|.\|$

denotes the total variation norm on $\mathscr{P}(S^\infty)$; i.e., given $p, q \varepsilon \mathscr{P}(S^\infty)$, $\|p - q\| \equiv \text{Sup}_E |p(E) - q(E)|$, where the supremum is over (Borel measurable) subsets E of S^∞. Let $\{x_n\}_{n=1}^\infty$ be a sequence in some metric space X. Let χ be any subset of X. We write $x_n \to^c \chi$ if every cluster point of $\{x_n\}_{n=1}^\infty$ lies in the set χ. Recall that $\mathscr{P}(S^\infty)$ is endowed with its weak topology.

We are now ready to state our main result in Theorem 3.1. The set W in Theorem 3.1 is the event where the players' beliefs about the future, $s^{N++} \equiv \{s_{N+1}, s_{N+2}, \ldots\}$, given the past, s^N, (and *not* conditioning on types) "merge" with μ^* (and hence with each other) as the date N tends to infinity. GD is the set where limit points of each player's beliefs about the future conditional on the past (not conditioning on own-types) play like some Nash equilibrium. G_i is the set where cluster points of the continuation strategies of player i's KSRB of beliefs not conditioning on own types, f^{*i}, lie in the set $N_i(\theta_i)$. In particular[4], if $g = (g_i, g_i) \varepsilon F$ is such a cluster point then g is a best-response to g_i for the player with attribute vector θ_i. Theorem 3.1 states that each of the sets W, GD and $\{G_i\}_{i\varepsilon I}$ have probability one.

Theorem 3.1. Let $\{\mu_i\}_{i\varepsilon I}$ be a collection of *ex ante* subjective beliefs of players and suppose that they satisfy (2.3)–(2.5). Let $\mu^* \varepsilon \mathscr{P}(\Omega)$ and suppose that μ^* and $\{\mu_i\}_{i\varepsilon I}$ obey (GGH). Define $\forall i \varepsilon I$,

$$G_i \equiv \left\{(\tau, f, s^\infty)\varepsilon\Omega : \sigma_N\left(s^N, f^{*i}\right) \to {}^c N_i(\theta_i)\right\} \quad \text{and} \quad G \equiv \cap_{i\varepsilon I} G_i \ ;$$

$$\text{GD}_i \equiv \left\{(\tau, f, s^\infty)\varepsilon\Omega : \mu_{iN}\left(ds^{N++}|S^N\right) \to {}^c \text{ND}(\theta(\tau))\right\} \text{ and } \text{GD} \equiv \cap_{i\varepsilon I} \text{GD}_i \ ;$$

and

$$W \equiv \left\{(\tau, f, s^\infty)\varepsilon\Omega : \lim_{N\to\infty} \left\|\mu_{iN}\left(ds^{N++} \mid s^N\right) - \mu_N^*\left(ds^{N++} \mid s^N\right)\right\| = 0 \ \forall i, j\varepsilon I\right\} \ .$$

Then

$$\mu^*(W \cap \text{GD} \cap G) = 1 \ . \tag{3.1}$$

4 Comparison with Kalai and Lehrer (1993a)

Consider again our model with I players each having beliefs represented by the ex ante subjective beliefs $\{\mu_i\}_{i\varepsilon I}$ defined in Section 3 and obeying best-response and independence conditions, (2.3)–(2.5). Let μ^* denote the induced *true* distribution over Ω, with some given distribution of initial types. This of course captures the KL93 *model* when extended to a situation with many types, as is indeed described in section 6 of their paper. Our formulation is slightly more general since we do not assume that $\{\mu_i\}_{i\varepsilon I}$ is a Bayesian-Nash equilibrium, and we have not placed any restrictions on the type space.

---

[4] Note that $G \neq \cap_{i\varepsilon I}\{(\tau, f, s^\infty)\varepsilon\Omega : \sigma_N(s^N, f^{*i}) \to {}^c N(\theta)\}$. Instead, the former set contains the latter, usually strictly. In particular G is *not* the set where continuation of KSRB of beliefs are Nash equilibria. The difference is the same as the difference between Nash and subjective Nash equilibria, and is because on G players are allowed to have different (limit) beliefs about play off the equilibrium path.

The natural extention of the KL93 *assumption* to the model with many types is:

$$\text{(KL) a. } \mu^*(T_{\text{KL}-T}) = 1 \text{ where } T_{\text{KL}-T} \equiv \cap_{i\varepsilon I}\{\tau = (\tau_i, \tau_{-i})\varepsilon T \mid \text{Marg}_{S^\infty}\mu^*(. \mid \tau) \ll \text{Marg}_{S^\infty}\mu_i(. \mid \tau_i)\} \ .$$

$$\text{b. } \text{Marg}_{T_i}\mu^* \ll \text{Marg}_{T_i}\mu_i \quad \forall i\varepsilon I \ .$$

Part (b) of (KL) is a technical condition which is required to rule out unimportant probability zero issues on the type space[5]. Part (a) of (KL) delivers, with part (b), the KL93 conclusions for μ^*-almost every τ.

Proposition 4.1: (KL) implies (GGH).

The examples in sections 1.2 and 5 show that (GGH) can be true while (KL) fails. Hence Proposition 4.1 above implies that *(KL) is strictly stronger than (GGH)*.

5 Some related recent literature

Koutsougeras and Yannelis (1994) have provided convergence results related to those presented here. For a survey of this and other related papers we refer the reader to Nyarko et al. (1994b). The model of this paper is immune to the criticisms of Nachbar (1996), which, I believe, are aimed more at the Kalai and Lehrer paper. Nachbar's main argument is that "prediction" of the future play of the game may be in opposition to "optimization", and that both prediction and optimization are often possible only if one assumes equilibrium from the beginning. Since, in this paper, we have a fairly weak absolute continuity condition (GGH), our formulation does not imply the stronger definition of prediction as used in Nachbar's paper. Indeed, by pointing out possible problems with the stronger notion of prediction (which are implicit in stronger notions of absolute continuity), the Nachbar paper actually puts our condition (GGH) in a much better light. It should also be remarked that the principal assumption I use in this paper is the absolute condition (GGH) which, loosely speaking, requires that set of sample paths which has positive probability under the true distribution, also has positive probability under each player's beliefs. Nowhere is there an assumption of equilibrium or coordination of beliefs, other than through (GGH)! Further discussion of the relation between this paper and those of Jordan (1991 and 1995), Nachbar (1996) and Kalai-Lehrer (1993) appear in Nyarko (1997).

---

[5] Failure of KL-Tb would make it hard to interpret μ^* as the "true distribution." Suppose (KL)-b. fails. Then μ^* will assign positive probability to a set of own types $\bar{T}_i \subseteq T_i$ which μ_i assigns zero probability to. What does player i do when he observes such an own type, $\tau_i\varepsilon\bar{T}_i$? On the set $\bar{T}_i$ player i could behave irrationally and still be consistent with (2.3). (One could then ask that (2.3) be strengthened to require, instead, the maximization to hold conditional on a given τ_i *for all* τ_i, rather than "*with ex ante probability one*" as stated. However this then introduces potential problems, including measurability issues.)

Lehrer and Smorodinsky (1997) also obtain a convergence to Nash result. They begin with the model of Jordan (1995) and *add* a "separation" condition on beliefs which is weaker than that of KL93. In our paper we begin with the Jordan model and *remove* conditions. In particular, the Jordan model is ours with the additional requirement of common priors, $\mu_i = \mu_j$ for all $i, j \varepsilon I$ (from which condition (GGH) follows trivially). Our condition (GGH) is therefore strictly weaker than the assumptions of Lehrer and Smorodinsky[6]. Sandroni (1995a,b) has also obtained related conditions for convergence to Nash equilibrium, none of which is weaker than our condition[7] (GGH).

6 Conclusion

We have provided a condition under which there is convergence to Nash or Subjective equilibria in infinitely repeated games. Our result is that beliefs, and not actual play, converge. We have argued that when the type space is big (uncountably so), then this is often the best conclusion that can be obtained. A connection between convergence of beliefs and convergence of actual empirical distributions of play (i.e., sample path averages) may be obtained by using the results of this paper in conjunction with the results of Nyarko (1994a) or Jordan (1996).

Appendix

Proof of Theorem 3.1. We shall show first that

$$W \cap G \subseteq \mathrm{GD} \ . \tag{8.1}$$

This part of the proof will follow directly from the definitions of the sets W, G and GD. We will next show that

$$\mu^*(W \cap G) = 1 \ . \tag{8.2}$$

This second part of the proof uses (GGH) and the original Jordan (1995) result with common priors. It should be obvious that (8.1) and (8.2) imply the conclusion of Theorem 3.1.

[6] I conjecture that the Lehrer and Smorodinsky result can be generalized to handle models without common priors. However, even in that case their assumption will then be at best merely incomparable with (and not weaker) than ours since even under common priors example 1.5 satisfies our condition (GGH) but fails their "separation" condition. To see this note that their separation condition implies convergence of actual play to Nash equilibrium, which we know is not the case in that example.

[7] Sandroni (1995a) introduces the notion of almost absolute continuity (which is a simple variation of KL93's absolute continuity) and shows that almost absolute continuity is necessary and sufficient for convergence to Nash equilibrium.

We now show (8.1). Fix any sample path $(\tau, f, s^{\infty})\varepsilon W \cap G$. Fix any $i'\varepsilon I$ and let $\{n(k)\}_{k=1}^{\infty}$ be a sub-sequence of dates such that $\lim_{k\to\infty} \mu_{i'N(k)}\left(ds^{N(k)++} \mid s^{N(k)}\right)$ exists and is equal to some $v_{\infty}\varepsilon\mathscr{P}(S^{\infty})$. Since $(\tau, f, s^{\infty})\varepsilon W, \lim_{k\to\infty} \mu_{jN(k)}\left(ds^{N(k)++} \mid s^{N(k)}\right) = v_{\infty}$ for all $j\varepsilon I$. For ease of exposition define along the given sample path for each $i\varepsilon I, f_N^{*i} \equiv \sigma_N(s^N, f^{*i})$. Since F is compact in its weak topology we can extract a further sub-sequence $\{N(k_t)\}_{t=1}^{\infty}$ of the subsequence $\{N(k)\}_{k=1}^{\infty}$ such that *for each* $i\varepsilon I$,

$$\lim_{t\to\infty} f_{N(k_t)}^{*i} \quad \text{exists and equals some } f_{\infty}^{*i}\varepsilon F_i \ . \tag{8.3}$$

Note that (i) from the definition of a KSRB we conclude that $m\left(f_{\infty}^{*i}\right) = v_{\infty}$ for all $i\varepsilon I$ and (ii) from the obvious continuity of the payoff function, $f_{\infty}^{*i}\varepsilon\, N_i(\theta_i)$. Consider the collection $\left\{f_{\infty}^{*i}\right\}_{i\varepsilon I}$, where for each $i, f_{\infty}^{*i} = \left\{f_{i,\infty}^{*i}, f_{-i,\infty}^{*i}\right\}$ and where $f_{i,\infty}^{*i}$ denotes i's strategy and $f_{-i,\infty}^{*i}$ is player i's belief. Properties (i) and (ii) above are actually the defining relations of what is referred to in the literature as a subjective equilibrium. Standard arguments then imply that there exists a Nash equilibrium which plays like it: i.e., $v_{\infty}\varepsilon\text{ND}(\theta)$. We have therefore shown the following: For any sample path $(\tau, f, s^{\infty})\varepsilon W \cap G$, if for some $i'\varepsilon I$ along a subsequence of dates $\mu_{i'N}(ds^{N++} \mid s^N)$ converges to some $v_{\infty}\varepsilon\mathscr{P}(S^{\infty})$ then $v_{\infty}\varepsilon\text{ND}(\theta(\tau))$. This proves (8.1).

We now prove (8.2). A careful reading of the proof of main result of Jordan (1995) indicates that what is indeed shown is that $\mu_i(G_i) = 1$ (and we spell out the details below). From its definition, one can easily see that the set G_i "depends only on (τ_i, s^{∞})" in the following sense: there exists a Borel-measureable subset $\mathbf{\Phi}_i$ of $T_i \times S^{\infty}$ such that $G_i = \{(\tau, f, s^{\infty})\varepsilon\Omega \mid (\tau_i, s^{\infty})\varepsilon\mathbf{\Phi}_i\}$. Hence the marginal of μ_i on $T_i \times S^{\infty}$ assigns probability one to $\mathbf{\Phi}_i$. From (GGH) the same is true of μ^*. Hence $\mu^*(G_i) = 1, \forall i\varepsilon I$. From (GGH) and the Blackwell-Dubins (1962) merging of opinions result, we conclude that $\mu^*(W) = 1$. Combining these equalities then proves (8.2). This completes the proof of Theorem 3.1. □

Theorem (Jordan, 1995). Fix any $i\varepsilon I$ and suppose that $\mu_i\varepsilon\, \mathscr{P}(\Omega)$ obeys conditions (2.3)–(2.5). Then $\mu_i(G_i) = 1$.

Proof of Propostion 4.1: For[8] ease of exposition we adopt the convention in footnote 8 below. Suppose that (KL) is true. Fix any $i\varepsilon I$ and any $D \subseteq T_i \times S^{\infty}$. Define for each $\tau_i\varepsilon T_i$, the set $D(\tau_i) \equiv \{s^{\infty}\varepsilon S^{\infty} \mid (\tau_i s^{\infty})\varepsilon D\}$ and $C \equiv \{\tau_i\varepsilon T_i \mid \mu_i(D(\tau_i) \mid \tau_i) = 0\}$. Suppose that $\mu_i(D) = 0$. Then $\mu_i(C) = 1$. (KLb) then implies that $\mu^*(C) = 1$. (KLa) in turn implies the event $\{\tau\varepsilon T \mid \mu^*(D(\tau_i) \mid \tau) = 0\}$ has μ^* probability one.

So,

[8] If μ is a probability over a product space $X_1 \times X_2$ and $M \subseteq X_1$ we let $\mu(M)$ denote $\mu(M \times X_2)$.

$$0 = \int_T \mu^*(D(\tau_i) \mid \tau) d\mu^* = \int_{T_i} \left[\int_{T_{-i}} \mu^*(D(\tau_i) \mid (\tau_i, \tau_{-i})) d\mu^* \right] d\mu^*$$
$$= \int_{T_i} \mu^*(D(\tau_i) \mid \tau_i) d\mu^* = \mu^*(D) \ .$$

We have therefore shown that $\mu_i(D) = 0$ implies that $\mu^*(D) = 0$. (GGH) therefore holds. □

References

1. Aumann, R.: Mixed and Behavior Strategies in Infinite Extensive Games. Advances in Game Theory. Annals of Mathematical Studies **5**, 627–650 (1964)
2. Battigalli, P., Guaitoli, D.: Conjectured Equilibrium and Rationalizability in a macroeconomic game with incomplete information. Bocconi University (1988)
3. Battigalli, P., Gilli, M., Molinari, C.: Learning and Convergence to Equilibrium in Repeated Strategic Interactions: An Introductory Survey. Ricerche Economiche (1992)
4. Blackwell, D., Dubins L.: Merging of Opinions with Increasing Information. Annals of Mathematical Statistics **38**, 882–886 (1963)
5. Blume, L., Easley, D.: Rational Expectations Equilibrium: An Alternative Approach. Journal of Economic Theory **34**, 116–129 (1984)
6. Harsanyi, J. C.: Games with Incomplete Information Played by Bayesian Players, Parts I, II, III, Management Science, vol. 14, 3, 5, 7., pp. 159–182, 320–334, 486–502 (1967, 1968)
7. Jordan, J. S.: Bayesian Learning in Normal Form Games, Games and Economic Behavior **3**, 60–81 (1991)
8. Jordan, J. S.: Bayesian Learning in Repeated Games. Games and Economic Behavior **9 (1)**, 8–20 (1995)
9. Jordan, J. S.: Bayesian Learning in Games: A NonBayesian Perspective In: The Dynamics of Norms, C. Bicchieri, R. Jeffrey, and B. Skyrms, (eds.). Cambridge: Cambridge University Press 1996
10. Kalai, E., Lehrer E.: Rational Learning Leads to Nash Equilibrium. Econometrica **61**, 1019–1046 (1993a)
11. Kalui, E., Lehrer, E.: Subjective Equilibrium in Repeated Games. Econometrica **61**, 1231–1240 (1993b)
12. Koutsougeras, L., Yannelis, N. C.: Convergence and Approximation Results for non-cooperative Bayesian Games: Learning Theorems. Economic Theory **4 (6)**, 843–858 (1994)
13. Kuhn, H.: Extensive Form Games and the Problem of Information. In: Contributions to the Theory of Games, Vol. II, Annals of Mathematics Studies, 28, H. W. Kuhn and A. W. Tucker (eds.). Princeton: Princeton University Press, pp. 193–216 (1953)
14. Lehrer, E., Smorodinsky, R.: Repeated Large Games with Incomplete Information. Games and Economic Behavior. **18** (1), 116–134 (1994)
15. Nachbar, J.: Prediction, Optimization, and Learning in Repeated Games. Forthcoming, Econometrica (1997)
16. Nyarko, Y.: Learning in Mis-Specified Models and the Possibility of Cycles. Journal of Economic Theory **55**, 416–427 (1991)
17. Nyarko, Y.: Bayesian Learning in Leads to Correlated Equilibria in Normal Form games. Economic Theory **4 (6)**, 821–841 (1994a)
18. Nyarko, Y.: Bounded Rationality and Learning, with N. Yannelis and M. Woodford. Economic Theory **4 (6)**, 811–820 (1994b)
19. Nyarko, Y.: Convergence in Economic Models with Bayesian Hierarchies of Beliefs. Journal of Economic Theory, forthcoming (1997a)
20. Nyarko, Y.: The Truth is in the Eyes of the Beholder: or Equilibrium in Beliefs and Rational Learning in Games. Manuscript, New York University (1997b)

21. Sandroni, A.: Does Rational Learning Lead to Nash Equilibrium in Finitely Repeated Games? Manuscript (1995a)
22. Sandroni, A.: Necessary and Sufficient Conditions for Convergence to Nash Equilibrium: The almost Absolute Continuity Hypothesis. Manuscript, University of Pennsylvannia (1995b)
23. Savage, L.: The Foundations of Statistics. New York: Wiley 1954