

Bayesian Learning Leads to Correlated Equilibria in Normal Form Games Author(s): Yaw Nyarko Source: *Economic Theory*, Vol. 4, No. 6, Symposium: Bounded Rationality and Learning (Oct., 1994), pp. 821-841 Published by: Springer Stable URL: <u>http://www.jstor.org/stable/25054810</u> Accessed: 09/05/2010 19:21

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/page/info/about/policies/terms.jsp. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/action/showPublisher?publisherCode=springer.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Springer is collaborating with JSTOR to digitize, preserve and extend access to Economic Theory.



Bayesian learning leads to correlated equilibria in normal form games*

Yaw Nyarko

Department of Economics, New York University, New York, NY10003, USA

Received: June 8, 1992; revised version February 12, 1993

Summary. Consider an infinitely repeated normal form game where each player is characterized by a "type" which may be unknown to the other players of the game. Impose only two conditions on the behavior of the players. First, impose the Savage (1954) axioms; i.e., each player has some beliefs about the evolution of the game and maximizes its expected payoffs at each date given those beliefs. Second, suppose that any event which has probability zero under one player's beliefs also has probability zero under the other player's beliefs. We show that under these two conditions limit points of beliefs and of the empirical distributions (i.e., sample path averages or histograms) are correlated equilibria of the "true" game (i.e., the game characterized by the true vector of types).

1. Introduction

Consider an infinitely repeated normal form game where each player is characterized by a "type" which may be unknown to the other players of the game. Impose only two conditions on the behavior of the players. First, impose the Savage (1954) axioms; i.e., suppose that each player has some beliefs about the evolution of the game and maximizes its expected payoffs in each period given those beliefs. Second, suppose that any event which has probability zero under one player's beliefs also has probability zero under the other player's beliefs. We show that under these two conditions there is "convergence" to a correlated equilibrium of the "true" game (i.e., the game characterized by the true vector of types): Limit points of beliefs about play AND the sample path averages (i.e., the empirical distribution) of play are equilibria of the true game.

The concepts of correlated and especially Nash equilibrium are very important in economics. It is therefore necessary to determine when the optimizing behavior

^{*} I gratefully thank Professor Jim Jordan for very many conversations. I thank Professors J-P Benoit, Lawrence Blume, David Easley and Ehud Kalai for very helpful comments. I am also very grateful to both the C. V. Starr Center and the Challenge Fund at New York University for their generosity.

of players subject to imperfect information leads to equilibrium behavior over time. This paper therefore provides answers to the question of how robust the Nash and correlated equilibrium concepts are, and whether, when initially out of "equilibrium", players can "learn" their way to equilibrium. This research is therefore a continuation of that began by Blume and Easley (1984). In answering these questions of course, to obtain a "good" result one should impose conditions which are as weak as possible. The players are facing very complicated decision problems. The less is the amount of prior knowledge and coordination we model the players with, the better (more "realistic") is the model. Hence we have attempted to do away with, as much as is possible, all common prior, independence of types and finiteness assumptions.

The only consistency condition we impose on players is that they agree on probability zero events. This condition is of course much weaker than the usual Harsanyi (1968) common prior assumption. We refer to our much weaker consistency condition as the Generalized Harsanyi condition (GH). (For an example of what could go wrong when condition (GH) is violated see Nyarko, 1991a.) The conditions we use in our analysis are also much weaker than those of Kalai and Lehrer (1990), who also obtain results on the convergence to Nash equilibrium. Within our framework their assumptions essentially require a finite or countable infinity of types, and is violated in applications where the type space is say an interval of the real line. Indeed, our main example in Section 2 satisfies all of our assumptions but violates those of Kalai and Lehrer (1990).

This paper is a generalization of earlier papers: Jordan (1991a and b) first studied the model considered here. However, the Jordan papers assumed that types are "independent." Nyarko (1992) also studied that model under a type-independence assumption, but relaxed the common prior assumption used in the Jordan papers. Jordan (1991a and b) and Nyarko (1992) all concluded that when the types are independent, convergence is to the Nash equilibrium of the true game. The concept of a correlated equilibrium was introduced by Aumann (1974) and is the natural extension of the concept of a Nash equilibrium, allowing players' actions to be correlated by some external signals, say. This paper allows for types which are not "independent" and hence allows for correlated equilibrium behavior even in the limit over time. When we assume that types are independent (or, stronger yet, when they are common knowledge) then the results of this paper imply convergence to Nash equilibrium. Since our method of proof is different from Jordan's we obtain an alternate route for obtaining those results. The Jordan papers provided results on convergence of beliefs. One may ask: What do the beliefs have to do with the actual play of the game? We provide here results on both beliefs AND the empirical distributions, and show that beliefs and actual sample path empirical distributions merge. Hence the latter is over time an equilibrium distribution.

Following Harsanyi (1968) and Mertens and Zamir (1985) a "type" represents a player's utility function **AND** that player's beliefs about others in the game. Hence, almost by definition, a player's beliefs should be allowed to depend upon its type. Hence we believe that the relaxation of the independence of types assumption of Jordan (1991) and Nyarko (1992) is important. However, with this relaxation our limit points are in general correlated, as opposed to Nash, equilibria.

To perhaps illustrate the strength of our results consider the following. Take any finite action normal form game where no player has a strictly dominated action. Impose our condition that players agree on probability zero events and are optimizing. Impose no other condition. Fix any **finite** time horizon. Then within that finite time horizon **any** play is possible, given some beliefs of players. Agreement on probability zero events does not preclude any behavior in **finite** time. However, after a "long" time there will be "a lot of" agreement as to the future play of the game. (See Theorem 8.2 for details of this.) Hence in the limit there will be total agreement as to the future play, and hence that play must be an equilibrium (either Nash or, when types are not independent, correlated).

To further illustrate our results consider a game with a unique mixed strategy Nash equilibrium. Suppose there is no imperfect information on types. Then our results show that beliefs and the sample path empirical distributions converge to the **mixed** strategy Nash equilibrium of the game. This is true despite the fact that each player may be choosing a **pure** strategy at each date. This therefore provides a rationale for the use of mixed strategies in terms solely of play under Bayesian Rationality. This of course is related to the much earlier arguments of Harsanyi (1973).

As an example where our Generalized Harsanyi condition (GH) (i.e., agreement on probability zero events) is violated, consider the standard versions of fictitious play. In fictitious play, each player believes that others are choosing strategies according to a **fixed** but unknown distribution, independently across time. Each player optimizes given these beliefs. Hence, each player is actually choosing a strategy which is highly time-dependent. However, each player believes the others are choosing time-independent actions. Such models violate our condition (GH). Such models are therefore able to generate non-convergence to Nash (or correlated) equilibrium, which is not possible in under condition (GH).

In Section 2 below we provide examples which illustrate ALL of the major results of this paper. The rest of the paper is devoted to formally stating the insights of the examples. Concluding remarks are provided in Section 11. All proofs are relegated to the appendix.

The results of this paper are for a model where players at each date maximize their within-period expected utility (i.e., the discount factor is zero). The results and techniques of this paper carry over immediately to the positive discount factor case. However in that model there are some subtle and important details on the definition of a correlated equilibrium. We therefore leave the full treatment of that case to a forthcoming paper.

This paper is related to the "Bayesian" models where uncertainty may be "reduced" imperfect information over an unknown parameter. This includes Blume and Easley (1984), Easley and Kiefer (1982), Feldman (1987) and Nyarko (1991b). However, such a "reduction" is not possible in the model of this paper, and the techniques used here are very different. Convergence to a correlated equilibrium in a very different model has also been obtained by Koutsougeras and Yannelis (1992).

2. Examples

2.1. Consider the following 3-person game payoff matrix (used by Aumann (1974) in a slightly different context):

0, 1, 3	0, 0, 0	2, 2, 2	0, 0, 0	0, 1, 0	0, 0, 0
1, 1, 1	1,0,0	2, 2, 0	2, 2, 2	1, 1, 1	1, 0, 3

Player A chooses the row (Top or Bottom), B chooses the column (Left or Right) and C chooses the matrix (First, Middle or Third).

Using dominance arguments it is easy to see that in any Nash equilibrium to the above normal form game, Player A chooses Bottom, Player B chooses Left column and player C randomizes (with any probabilities) between the First and Third matrices.

Let ω be a realization from infinitely many independent and identical cointossing experiments where an outcome from {Heads, Tails} is chosen with equal probability. Hence ω is an element of {Head, Tails}^{∞}. Players A and B are told of the realization of ω . We may consider player A's "type" to be $\tau_A = \omega$, Player B's "type" to be $\tau_B = \omega$, so that $\tau_A = \tau_B$. Player C is not told of ω but knows how it is chosen (i.e., C knows the distribution of ω). Player C has a trivial type space (consisting of a singleton element, say, representing "no information").

Let ω_n denote the *n*-th coordinate of ω . In particular $\omega_n \in \{\text{Heads, Tails}\}$. Consider the following strategies for the players: At each date *n*, Player A of type $\tau_A = \omega$ chooses Top at date *n* if $\omega_n =$ Heads and Bottom if $\omega_n =$ Tails. Player B of type $\tau_B = \omega$ chooses the Left Column if $\omega_n =$ Heads and Right if $\omega_n =$ Tails. Player C chooses the Middle matrix all the time. It should be easy to see that if any player believes the others are choosing actions in the manner just described then it is optimal for that player to choose actions in the manner described above for that player. In particular, each player is choosing a best response at each date to its beliefs about the other players.

Under this behavior, observe that Middle matrix is chosen at each date. This is not a Nash equilibrium action for the true game. Hence we have,

Observation 1: The actions or play of the players need NOT converge to a NASH equilibrium of the true game.

Consider the beliefs of players about the date n play, conditional on the history of the game from date 1 through date n-1. These beliefs will assign probability one to Middle matrix being played. Middle is not part of any Nash equilibrium. In particular we have,

Observation 2: The **beliefs** of players about the future play of the game conditional on the past (either conditional or not conditional on players own realized types) need **NOT** converge to a **Nash** equilibrium of the true game.

This example appears to contradict the conclusions of Jordan (1991a and b) and Nyarko (1992), where convergence of beliefs to a **Nash** equilibrium was proved. However, notice that the Players' beliefs about the types of others are not independent of their own type; indeed, we have extreme dependence with $\tau_A = \tau_B$. Hence the independence assumptions used in the just-mentioned papers are violated. This example, and in particular observation 2 above, shows that when that independence assumption is violated the conclusions of those papers fail. However, notice that the behavior we have described in this example is actually a correlated equilibrium. Players A and B use the outcomes of the coin-tosses to coordinate their actions across (Top, Left) and (Bottom, Right). The outcomes (Top, Left, Middle) and (Bottom, Right, Middle) with probability 1/2 each constitutes a correlated equilibrium distribution. In particular, suppose a "Principal" (or correlating device) "suggests" that Player A should play action Top. Then A knows, in the correlated equilibrium, that B and C will play Left and Middle respectively, so it is optimal for A to follow the suggestion of the Principal. Similarly for B. In the correlated equilibrium the only action that will be suggested to player C is the action Middle. C assigns equal probability to A and B choosing action pairs (Top, Left) and (Bottom, Right). Hence following the suggestion of the Principal is optimal. This verifies that (Top, Left, Middle) and (Bottom, Right, Middle) with probability 1/2 each is indeed a correlated equilibrium distribution.

Note further that this distribution is also the beliefs of each player about the next period play of the game *not conditioning on own types*. In particular, if we asked each player to predict the outcome of the future of the game conditional on only the history of the game but *not* conditional on their own realized type then each player would predict the play to be (Top, Left, Middle) and (Bottom, Right, Middle) with probability one half each. This illustrates one of the main results of this paper (Theorem 9.1 below), namely,

Result 1: Beliefs of players about the future of the game conditional upon the past of the game but **not conditional upon own types** converge to a **correlated** equilibrium distribution for the true game.

Beliefs about the future of the game conditional upon the past and conditional upon own types do not converge to the set of correlated equilibria. Indeed, A's belief about the date N play of the game conditional upon A's realized type and the history preceding date N is either that (Top, Left, Middle) will occur with probability one or that (Bottom, Right, Middle) will occur with probability one. (Which will occur is of course determined completely by the date N coordinate, ω_N , of A's realized type ω .) However, neither of the outcomes (Top, Left, Middle) with probability one or (Bottom, Right, Middle) with probability one is a correlated equilibrium since in either case Player C will be choosing a sub-optimal action. In particular we have,

Observation 3: Beliefs of players about the future of the game conditional upon the past of the game **and conditional upon own types** need **not** converge to a **correlated** equilibrium distribution for the true game.

Let us now look at actual play again. We may invoke the strong law of large numbers to conclude that for almost every sample path, in each sufficiently long history or play of the game the outcome (Top, Left, Middle) will occur for approximately as many periods as the outcome (Bottom, Right, Middle). In particular the average number of times each outcome will occur will in the limit be equal to 1/2 for almost every sample path. This outcome is the same as the limit point of beliefs of players (see Result 1 above), and in particular is a correlated equilibrium for the true game. This illustrates the second main result of this paper (Theorem 9.6): We define the empirical distribution of play to be the distribution (or histogram) obtained by taking the average number of occurrences of each action in the past history.

Result 2: The empirical distribution of play converges to a correlated equilibrium distribution for the true game. Further, the beliefs of players **not** conditioning on types and the empirical distribution of play converge to the same limit point over time.

2.2. Example: The need for sub-sequential limit points. The example of Section 2.1 may have given the impression that we will be proving the **convergence** of beliefs and empirical distributions. In the example of that section such convergence did indeed take place. However, for a general game the beliefs about the future of the game need not even converge. Consider the following coordination game with two players, Player A (the row player) and Player B (the column player).

1, 1	0,0
0,0	1, 1

Suppose that at every even period the players choose actions (Top, Left) while at every odd period they choose the actions (Bottom, Right). Each player will then be best responding given their beliefs about the other. However each player's beliefs about the future of the game conditional on the past do not converge. Instead the beliefs have two limit points: (Top, Left) along the sub-sequence of even dates and (Bottom, Right) along the sub-sequence of odd dates. Hence result 1 of example 2.1, will actually be stated in terms of sub-sequential limits: Any limit point of **Beliefs** of players about the future of the game conditional upon the past of the game but **not conditional upon own types** is a **correlated** equilibrium distribution for the true game.

We now illustrate that for some games the empirical distribution need not converge either. Indeed, fix any sequence of numbers $\{x_n\}_{n=1}^{\infty}$ taking values of either 0 or 1 such that the averages $\sum_{n=1}^{N} x_n/N$ do not converge but oscillate between being arbitrarily close to 0 and arbitrarily close to 1 infinitely often. (This can of course be done by choosing x_n to be equal to 0 for a long time, then equal to 1 for an even longer time then equal to 0 for a yet longer time, etc.) Suppose now that players choose actions (Top, Left) at each date *n* where $x_n = 1$ and choose actions (Bottom, Right) otherwise. Then it should be clear that the empirical distribution of play does not converge, but instead oscillates between arbitrarily high average for (Top, Left) to an arbitrarily low average.

However, consider now a sub-sequence of dates where beliefs about next period play conditional on the past converges. In this simple example this will be a sub-sequence of dates where either (Top, Left) is played at each date or where (Bottom, Right) is played at each date. Along such a sub-sequence of dates the empirical distribution will converge. The limit point of the empirical distribution along any such sub-sequence is of course is either (Top, Left) with probability one (i.e., with limiting average equal to one) or (Bottom, Right) with probability one. Either of these will constitute a correlated (actually Nash) equilibrium.

Hence result 2 of example 2.1, will actually be stated in terms of sub-sequential limits. Our result will say, loosely speaking (and see Section 9 for details) the following: Fix a sample path. Suppose that along a sub-sequence of dates **beliefs** of players about the future of the game conditional upon the past of the game but **not** conditional upon own types converges to some limit point, v say. We know from Result 1 that v must be a correlated equilibrium for the true game. Our result states that the empirical distribution of play, constructed along that sub-sequence (i.e., using only observations on that sub-sequence) also converges to the limit point v.

2.3. What if types are common knowledge? Suppose that all utility parameters are common knowledge and that there is no uncertainty about "types". Without imperfect information on types, beliefs conditional on types and beliefs not conditional on types are the same. Also there are no types to allow correlation in actions of players. Hence (sub-sequential) limits of beliefs conditional on types must be Nash as (opposed to correlated equilibria). The same of course is true of the sample path averages (empirical distributions) along the convergent sub-sequences of beliefs. More is true: Since player *i* knows its own actions, *i*'s beliefs conditional on *i*'s type equals *i*'s actual play. Our previous assertion therefore implies that limit points of actual play and not merely beliefs about play are Nash equilibria. (See Section 10 for the details.)

3. Some terminology

I is the **finite** set of players. Given any collection of sets $\{S_i\}_{i \in I}$, we define $S \equiv \prod_{i \in I} S_i$ and $S_{-i} \equiv \prod_{j \neq i} S_j$. Given any collection of functions $f_i: S_i \to Y_i$ for $i \in I$, $f_{-i}: S_{-i} \to Y_{-i}$ is defined by $f_{-i}(s_{-i}) \equiv \prod_{j \neq i} f_j(s_j)$. The Cartesian product of metric spaces will always be endowed with the product topology. Let S be any metric space. $\mathcal{P}(S)$ is the set of probability measures on (Borel) subsets of S. Unless otherwise stated (and we will!) the set $\mathcal{P}(S)$ will be endowed with the weak topology. If S is a cartesian product of the form $\prod_{i \in I} S_i$, then $\mathcal{P}_0(S)$ will denote the set of all product measures on $\prod_{i \in I} S_i$. Given any $v \in \mathcal{P}(S)$ we let v(ds) denote integration: $\int h(s)v(ds)$ is the integral of the real-valued function h on S with respect to v. If S is a Cartesian product $S = Y \times Z$ we let v(dy) denote integration over Y with respect to the marginal of v on Y. The latter will often be denoted by Marg_V v. \Re denotes the real line.

4. The basic structure

4.1. Following Jordan (1991a) we have the following basic structure of the game. *I* is the **finite** set of players. S_i represents the **finite** set of actions available to player *i* at each date $n = 1, 2, ...; S \equiv \prod_{i \in I} S_i$. Even though the action space S_i is independent of the date we shall sometimes write S_i as S_{in} when we seek to emphasize the set of action choices at **date n**. $S^N \equiv \prod_{r=1}^N S$ and $S^\infty \equiv \prod_{r=1}^\infty S$ are the set of date N and infinite histories, respectively. S^N and S^∞ are endowed with their respective product

topologies. s^N and s^∞ will denote generic elements of S^N and S^∞ respectively. s^0 will denote the null history, (at date 0, when there is no history)! Perfect recall is assumed; in particular, at date *n* when choosing the date *n* action s_{in} the player *i* will have information on $s^{n-1} = \{s_1, \ldots, s_{n-1}\}$.

4.2. Payoff functions. Each player *i* has an attribute vector which is some element θ_i of the set θ_i . The attribute vector will represent the parameter of its utility function unknown to other players in the game. $u_i: \theta_i \times S \to \Re$ is player *i*'s (within period or instantaneous) utility function which depends upon its attribute vector, θ_i , as well as the vector of actions, $s \in S$, chosen by all the players. We assume that u_i is continuous and uniformly bounded on it's domain. We shall suppose that θ_i is a compact subset of finite dimensional Euclidean space. This is without loss of generality since the set of joint actions, S, is assumed finite. We suppose that the players know the functional forms of each player's utility function, $\{u_i\}_{i \in I}$. Each player *i* knows its own attribute vector θ_i but does not necessarily know those of other players, θ_{-i} .

5. Equilibria for the complete information problem

5.1. Nash equilibria. Recall that $\mathscr{P}_0(S)$ is the set of all product measures over $S = \prod_{i \in I} S_i$. Define for each $i \in I$ and $\{\theta\}_{i \in I} \in \theta$,

$$ND_{i}(\theta) \equiv \{ v \in \mathscr{P}_{0}(S) : s_{i} \in \operatorname{Argmax} \int u_{i}(\theta_{i}, .., s_{-i}) v(ds_{-i}) \text{ for } v \text{-a.e. } s_{i} \in S_{i} \} \text{ and } ND \equiv \bigcap_{i \in I} ND_{i}(\theta_{i})$$

 $ND(\theta)$ is the set of Nash equilibrium distributions, where all players are best responding to the others.

5.2. Correlated equilibria. Correlated equilibria are typically defined using "correlating devices" which involve expanding the space of uncertainty and endowing players with information partitions representing the correlated signals received by those players. (See Aumann (1974).) We shall use the following definition which involves conditioning on actions. This is equivalent to the definitions that involve information partitions.

Fix any $i \in I$, $\theta = \{\theta_i\}_{i \in I} \in \theta$ and $Q \in \mathscr{P}(S)$. Define $Y_i(\theta_i, Q)$ to be the set of all actions, $s_i \in S_i$, which are optimal for player *i* when the actions of other players are governed by $Q(\cdot|s_i)$. Also define $\Psi_i(\theta_i)$ to be the set of all distributions of play, Q, for which the above statement is true for all date one actions excluding possibly a set with Q-probability zero. I.e.,

$$Y_i(\theta_i, Q) \equiv \{s_i^* \in S_i : s_i^* \in \operatorname{Argmax} \int u_i(\theta_i, \dots, s_{-i}) Q(ds_{-i} | s_i^*)$$
(5.3)

and

$$C_i(\theta_i) \equiv \{ Q \in \mathscr{P}(S) : Q(Y_i(\theta_i, Q)) = 1 \}.$$
(5.4)

Finally, define

$$C(\theta) = \bigcap_{i \in I} C_i(\theta_i) \quad \text{where } \theta = \{\theta_i\}_{i \in I}.$$
(5.5)

 $C(\theta)$ is the set of all correlated equilibrium distributions of play. Any Q in $C(\theta)$ is such

that for all $i \in I$ and for Q-almost almost all $s_i \in S_i$ the following is true: if the "Principal suggests" the action s_i to player *i* and player *i* believes that the other players are choosing actions according to $Q(.|s_i)$ then s_i is an optimal action for that player. It is these "suggestions" which allow for correlation of the players' actions.

It should be clear that any correlated equilibrium distribution which satisfies some kind of independence across players should be a Nash equilibrium. Indeed, fix any $\theta \in \Theta$ and $Q \in C(\theta)$. Suppose that under Q each player's action is independent. I.e., suppose that Q is a product measure so that for all $s = \{s_i\}_{i \in I} \in S$, $Q(\{s\}) = \prod_{i \in I} Q(\{s_i\})$. Then it should be clear that Q is a Nash equilibrium **distribution**; i.e., $Q \in ND(\theta)$.

6. The imperfect information problem

6.1. The type space. Player $i \in I$ may be any "type" in a type space T_i , assumed to be a complete and separable metric space. We set $T \equiv \prod_{i \in I} T_i$. Player *i*'s type, τ_i , specifies, among other things, that player's attribute vector, θ_i . We let $\theta_i(\tau_i)$ denote the attribute vector of player *i*, type τ_i ; $\theta_i: T_i \rightarrow \theta_i$ is therefore the projection mapping from the type space T_i representing the *i*-th player's attribute vector into Θ_i .

6.2. Bayesian strategy processes without common priors. We fix a collection of probability distributions $\{\mu_i\}_{i \in I}$ over $T \times S^{\infty}$. μ_i represents the ex ante beliefs player *i* has about the evolution of the game **before** *i* has realized its own type. The ex post beliefs are therefore $\mu_i(\cdot|\tau_i)$. The collection of measures, $\{\mu_i\}_{i \in I} \in \prod_{i \in I} \mathscr{P}(T \times S^{\infty})$, is a Bayesian Strategy Process (BSP) for the game with (not necessarily common) priors if for each $i \in I$,

$$\mu_{i}(\{(\tau, s^{\infty}) \in T \times S^{\infty} | \text{ for all } N, s_{iN+1} \text{ maximizes} \\ \int u_{i}(\theta_{i}(\tau_{i}), ., s_{-iN+1}) d\mu_{i}(ds_{-iN+1} | s^{N}, s_{iN+1}, \tau_{i})\}) = 1.$$
(6.3)

Condition (6.3) requires that at each date N given player *i*'s beliefs about the evolution of the game, $\mu_i(.|s^N, s_{iN+1}, \tau_i)$, player *i* maximizes its expected utility. (6.3) by itself does not imply that under *i*'s beliefs about the game other players $j \neq i$ are maximizing their expected utility. (However, this latter assertion will hold under condition (GH) which will be introduced in Section 7.)

Note that we have **not** ruled out correlation in players' choice of actions. In particular, we do **not** impose the following assumption:

$$\mu_i(ds_{N+1}|s^N,\tau) = \prod_{j \in I} \mu_i(ds_{jN+1}|s^N,\tau_i).$$
(6.4)

Condition (6.4) is indeed a natural assumption to impose, and it holds if players choose actions at each date simultaneously. Condition (6.4) is used by Jordan (1991a and b) and Nyarko (1992) to obtain a result on the convergence to Nash equilibrium behavior. Since our main result will be convergence to correlated equilibrium, we have no need for (6.4). We will impose (6.4) in Section 10 when we indicate how the earlier results on convergence to Nash equilibria are special cases of our results on the convergence to correlated equilibria. We separate all the independence assumptions so that it becomes clear what assumptions are used for the various conclusions.

General conditions for the existence of BSP's where correlated actions are allowed is provided in Cotter (1991) and Yannelis and Rustichini (1991).

7. The generalized Harsanyi consistency condition

We will impose the following condition on the beliefs of players, $\{\mu_i\}_{i \in I}$, which requires that, ex ante, the players agree on probability zero events. The Harsanyi (1968) common prior assumption requires $\mu_i = \mu_j$ for all *i* and *j*. Our condition (GH) below is therefore a generalization of the Harsanyi assumption. The common prior assumption is used by Jordan (1991a and b).

Given any two probability measures μ' and μ'' on some (measure) space Ω , we say that μ' is absolutely continuous with respect to μ'' if for all (measurable) subsets D of Ω , $\mu'(D) > 0$ implies that $\mu''(D) > 0$. We then write $\mu' \ll \mu''$. We say that μ' and μ'' are mutually absolutely continuous with respect to each other if $\mu' \ll \mu''$ and $\mu'' \ll \mu'$.

7.1. Condition (GH): The measures $\{\mu_i\}_{i \in I}$ in $\mathscr{P}(T \times S^{\infty})$ are mutually absolutely continuous with respect to each other.

Condition (GH) does not require the ex post probabilities, $\mu_i(\cdot|\tau_i)$ and $\mu_j(\cdot|\tau_j)$, to be mutually absolutely continuous. We shall use the following much weaker version of condition (GH):

7.2. Condition (GGH). There exists a measure μ^* over $T \times S^{\infty}$ such that for all $i \in I$, μ^* is absolutely continuous with respect to μ_i .

One may wish to interpret μ^* as the "true" distribution of the types and play while μ_i is player i's beliefs. Any event which has positive probability under μ^* in condition (GGH) will have strictly positive probability under μ_i for each $i \in I$. The converse however need not be true under the weaker condition (GGH). When condition (GGH) holds we shall state our results in terms of the measure μ^* ; that condition should therefore be thought of as providing such a measure. If condition (GH) holds, to obtain condition (GGH) we may take the measure μ^* to be equal to any of the μ_i 's or indeed any measure over $T \times S^{\infty}$ which is mutually absolutely continuous with respect to any (and therefore all) of the μ_i 's; e.g., μ^* may be taken to be the average measure $\sum_{i \in I} \mu_i / (\#I)$.

7.3. Remark. As will soon become apparent, the principal use of condition (GH) or (GGH) is to ensure agreement in the limit about **play** of the game; (in particular its main use will be to prove Theorem 8.2 below). Hence, for all the main results of this paper, we may replace conditions (GH) and (GGH) above with assumptions which require only absolute continuity of the marginals of μ_i on S^{∞} and not necessarily over all of $T \times S^{\infty}$. In particular, beliefs of players about types are by themselves unimportant. We use assumptions (GH) and (GGH) as stated above because it is expositionally more convenient.

7.4. Remark. Let μ^* denote the true distribution of play for the various playertypes. The absolute continuity assumption of Kalai and Lehrer (1990) requires that for each *i* in *I* and for each $\tau = \{\tau_i\}_{i \in I} \in T$, $\mu^*(.|\tau) \ll \mu_i(\cdot|\tau_i)$. In particular, their assumption requires ex post absolute continuity. Our condition (GGH) is weaker and requires only ex ante absolute continuity. In particular example 2.1 obeys condition (GH) but violates the Kalai and Lehrer assumptions. For their assumption to hold the set of ("behavior equivalent classes") of types must be finite or countably infinite. (See Nyarko 1991b for details.)

8. (GH) Implies that beliefs about the future "merge"

The following Lemma follows immediately from Blackwell and Dubins (1963) theorem on "Merging of Opinions": Let $\mu_i(ds_{n+1}|s^n)$ denote the probability distribution over the date n + 1 actions conditional on the "past," s^n , with respect to the measure μ_i . The norm $\|\cdot\|$ denotes the total variation norm on S; i.e., given $p, q \in \mathscr{P}(S)$,

$$||p|| \equiv \operatorname{Sup}_{E}|p(E) - q(E)|$$
 (8.1)

where the supremum is over (measurable) subsets E of S.

The theorem below implies that for each *i* and *j* in *I*, the beliefs of the players about the future of the game conditional on the past, $\{\mu_i(ds_{n+1}|s^n)\}_{n=1}^{\infty}$ and $\{\mu_j(ds_{n+1}|s^n)\}_{n=1}^{\infty}$ have the same limiting behavior and share the same limit points along any subsequence of dates in each sample path. Observe that these conditional probabilities are **not** conditioned on players' types.

8.2. Theorem (Blackwell and Dubins). Suppose that the measures $\{\mu_i\}_{i \in I}$ on $T \times S^{\infty}$ obey condition (GGH) and let μ^* be as in that condition. Define

$$W = \{(\tau, s^{\infty}) \in T \times S^{\infty} : \lim_{n \to \infty} \|\mu_i(ds_{n+1}|s^n) - \mu_j(ds_{n+1}|s^n)\| = 0 \,\forall i, j \in I\}.$$
 (8.3)

Then

$$\mu^*(W) = 1. \tag{8.4}$$

9. Convergence to correlated equilibria

We now show that the probability distribution of the future of the game conditional on the past, $\mu_i(ds_{N+1}|s^N)$, converges to the set $C(\theta)$ of correlated equilibrium distributions for the true attribute vector θ . Recall that $\mathscr{P}(S)$ is endowed with the topology of weak convergence. Given any subset C of $\mathscr{P}(S)$ and any sequence of measures $\{q_n\}_{n=1}^{\infty}$ in $\mathscr{P}(S)$ we write $q_n \twoheadrightarrow C$ if any limit point of the sequence of measures lies in the set C.

9.1. Theorem (Beliefs converge to correlated equilibria). Let $\{\mu_i\}_{i \in I}$ be a BSP. Suppose condition (GGH) holds and let μ^* be as in that condition. Define

$$G \equiv \{(\tau, s^{\infty}): \mu_i(ds_{N+1} | s^N) \twoheadrightarrow C(\theta(\tau)) \text{ for all } i \in I\}$$
(9.2)

Then $\mu^*({G \cap W}) = 1$ where W is as in (8.3). (A sketch of the main idea of the proof and the intuition behind it is given in the appendix, as is the proof itself.)

9.3. The convergence of the empirical distributions (i.e., sample path averages). In Theorem 9.6 below we relate the limit point of beliefs (as discussed in the previous sub-section) to sample path empirical distributions (i.e., the distributions obtained by taking a histogram or sample path average of occurrences of the different vectors of joint actions of players). Related results on the convergence of empirical distributions are also reported in Jordan (1992). Our results are as follows:

Fix a sample path and suppose that along that sample path the beliefs about the future conditional on the past, but NOT conditional on own types, $\mu_i(ds_{n+1}|s^n)$, **converges** as $n \to \infty$. Let us denote that limit point by $v \in \mathscr{P}(S)$. From Theorem 8.2

v is independent of *i*. Part (a) of our result (Theorem 9.6 below) concludes that the empirical distribution also converges to v. In particular, fix any $s^* \in S$, and define $1_n(s^*)$ to equal one if $s_n = s^*$ and $1_n(s^*) = 0$ otherwise. Then $\sum_{n=1}^N 1_n(s^*)/N$ converges to $v(\{s^*\})$ as $N \to \infty$. From the results of the previous sub-section we know that any limit point of beliefs is a correlated equilibrium. Hence, we may conclude that on any sample path where these beliefs converge, the empirical distribution is in the limit a correlated equilibrium distribution for the true game.

We know from the example of Section 2.2 that beliefs about the future given the past do not necessarily converge. Part (b) of our result (Theorem 9.6) handles this case. It states that beliefs and empirical distributions "merge" along convergent subsequences. The result states, loosely speaking, the following: Fix any sample path. Suppose that along a sub-sequence of dates beliefs about the future given the past, $\mu_i(ds_{n+1}|s^n)$, converges to some limit point, $v \in \mathscr{P}(S)$. Then the empirical distribution constructed along that sub-sequence also converges to v. Since the limit points of beliefs are correlated equilibria we conclude that those limit points of the empirical distributions are also correlated equilibria.

Notice that in the above we used the language "loosely speaking." Even though the statement above is intuitively correct there is a technicality that must be taken care of. In particular, when forming the empirical distributions along a sub-sequence of dates, we need to take a "rich" enough set of dates in that sub-sequence. The formal statement of part (b) of our result is as follows. Fix a sample path and suppose there is a sub-sequence of dates such that along that sub-sequence the beliefs, $\mu_i(ds_{n+1}|s^n)$, converge to some $v \in \mathscr{P}(S)$. Fix any arbitrarily small neighborhood of v, and refer to this neighborhood as Λ . (For measurability reasons we restrict attention to neighborhoods which are intervals (or rectangles) with rational end-points or (vertices).) Construct the empirical distribution from only the periods where the beliefs, $\mu_i(ds_{n+1}|s^n)$, lie in the neighborhood Λ . Then our result states that the empirical distribution thus constructed will lie in the neighborhood Λ . The example below shows what could go wrong if we take "too thin" a sub-sequence of dates.

9.4. Example. There are two players, A and B. Player A chooses the row: Top, Middle or Bottom. Player B chooses the column: Left, Center or Right. The payoff matrix is as below:

-1,1	1, -1	0,0
1, -1	-1,1	0,0
0,0	0,0	1,1

Notice that if we exclude the actions Bottom and Right, the payoff matrix is that of a "matching pennies" game.

Suppose the play of the game is as follows: On each even date the players play the unique mixed strategy Nash equilibrium of the matching pennies part of the game: I.e., Player A uses the mixed strategy which randomizes with equal probability over the actions Top and Middle and B uses the mixed strategy which randomizes with equal probability over the actions Left and Center. On each odd date Player A chooses action Bottom while B chooses action Right. It is easy to see that under this behavior each player is best-responding to the other.

There are two limit points of beliefs of players about the future of the game: One limit point assigns probability of 1/4 to each of the four vectors of actions in the matching pennies part of the game. Another limit point of beliefs assigns probability one to the vector (Bottom, Right). In obtaining convergence of the empirical distribution along sub-sequences it should be clear what sub-sequences to choose from. Either we take the sub-sequence of even dates or the sub-sequence of odd dates.

If, however, we take too "thin" a sub-sequence we may not obtain convergence of the empirical distribution to the limit point of beliefs. In particular, fix a sample path and consider the sub-sequence of dates where the action vector (Top, Left) occurs. On almost every sample path there will indeed be a sub-sequence of such occurrences. Of course along that sub-sequence the empirical distribution will always assign probability one to the action vector (Top, Left). The empirical distribution along that sub-sequence therefore does not converge to the limit point of beliefs. Since (Top, Left) is not a Nash or correlated equilibrium, the empirical distribution also does not converge to the set of correlated equilibrium distributions. Our chosen sub-sequence is "too thin."

The use of neighborhoods Λ as described in Section 9.2 will enable us to pick out the even or odd sub-sequences of dates. Suppose we take Λ to be a small neighborhood around the probability which assigns mass of 1/4 to each of the vectors on the matching pennies part of the game. Players' beliefs will lie in this neighborhood at precisely the sub-sequence of even dates. If Λ is a small neighborhood of the probability which assigns point mass to {Bottom, Right} then we pick out the sub-sequence of odd dates. The use of these neighborhoods provides us with a "rich" enough sub-sequence of dates.

9.5. The statement and proof of Theorem 9.6 below is really a probability-theoretic result and does not use anywhere the fact that players are optimizing. The result states that beliefs (i.e., conditional probabilities of the future given the past) and sample path empirical distributions "merge" over time (i.e., have the same subsequential limits). The theorem is stated for any general $\mu \in \mathscr{P}(T \times S^{\infty})$. It is easy to see how this theorem and condition (GH) or (GGH) prove all the assertions made above. We now state and prove this theorem.

We use the following notation in the result below. Given any event E let 1_E represent the indicator function on E (equal to one if E holds and equal to zero otherwise).

9.6. Theorem. Let μ be any measure on $T \times S^{\infty}$. Then the following is true for each sample path $(\tau, s^{\infty}) \in T \times S^{\infty}$ (outside of a set with μ -probability zero):

- (a) Suppose that for some $v \in \mathscr{P}(S)$, along the given sample path $\lim_{n \to \infty} \mu(ds_{n+1}|s^n) = v$. Fix any $s^* \in S$. Then, along the sample path the average number of times s^* occurs is in the limit equal to $v(\{s^*\})$. I.e., $\lim_{N \to \infty} \sum_{n=1}^{N} 1_{\{s_n = s^*\}}/N = v(\{s^*\})$.
- (b) Suppose that for some $v \in \mathscr{P}(S)$, along the given sample path there is a sub-sequence of dates $\{n(k)\}_{k=1}^{\infty}$ such that along the sub-sequence $\lim_{k \to \infty} \mu(ds_{n(k)+1}|s^{n(k)}) = v$. Fix any $s^* \in S$. Let Λ be any arbitrarily small non-degenerate

closed interval with rational endpoints which contains $v({s^*})$. Then

$$\lim_{N \to \infty} \sum_{n=1}^{N} \mathbf{1}_{\{s_n = s^*\}} z_{n-1} / \sum_{n=1}^{N} z_{n-1} \in \Lambda.$$

where

$$z_n \equiv 1_{\{\mu(\{s_{n+1} = s^*\} | s^n) \in \Lambda\}}.$$
(9.7)

10. Independence of types implies convergence to Nash equilibria

In obtaining the convergence to Nash equilibria result of Jordan (1991a and b) and Nyarko (1992) the assumption below is used (in addition to (6.4) above). This assumption requires that each player i's beliefs about the types of others is independent of player i's type:

10.1. A type-independence assumption. Define $\pi_i \equiv \text{Marg}_T \mu_i$. Then π_i is a product measure on the type space T; i.e., $\pi_i = \prod_{i \in I} [\text{Marg}_{T_i} \pi_i]$

The example provided in Section 2.1 obeys (6.4) but violates (10.1). Hence without (10.1) the result on the convergence to **Nash** (as opposed to correlated) equilibria obtained in the Jordan (1991) and Nyarko (1992) does not hold. An immediate corollary of our Theorem 9.1 is that if players choose actions independently and the priors over types are independent (i.e., if (6.4) and (10.1) hold) then beliefs of players converge to a Nash equilibrium distribution. It is easy to see why this is so. Under (10.1) and (6.4), each player *i*'s beliefs about date N + 1 play conditional on date N history will be a product measure; i.e., $\mu_i(ds_{N+1}|s_N) = \prod_{j \in I} \mu_i(ds_{jN+1}|s_N)$. Any limit point of beliefs will also have this property. We know from the previous section that any limit point of beliefs is a correlated equilibrium. As explained at the end of Section 5, any correlated equilibrium distribution with the above mentioned product measure property is necessarily a Nash equilibrium distribution. Hence limit points of beliefs are Nash equilibria. We therefore obtain the conclusions of Jordan (1991a and b) and Nyarko (1992) as special cases of Theorem 9.1:

10.2. Corollary. Let $\{\mu_i\}_{i \in I}$ be a BSP and suppose condition (GGH) holds. Also suppose that (6.4) and (10.1) hold. Recall that the set W is defined in (8.3). Define

$$\overline{G} \equiv \{(\tau, s^{\infty}): \mu_i(ds_{N+1} | s^N) \twoheadrightarrow \mathsf{ND}(\theta(\tau)) \text{ for all } i \in I\}$$

Then $\mu^*(\overline{G} \cap W) = 1$.

10.3. Remark. Of course, from Theorem 9.6 we may conclude that under the conditions of Corollary 10.2 the empirical distribution (constructed along the appropriate sub-sequences) are also Nash (as opposed to correlated) equilibria.

10.4. Model with types common knowledge. Suppose now that there is only one vector of possible types (or, alternatively, that the vector of types is common knowledge). Then, trivially, condition (10.1) holds. If players choose actions simultaneously (so that (6.4) holds) then Corollary 10.2 implies that any limit point of beliefs is a Nash equilibrium. Since there is only one vector of types, each player's beliefs about it's own play not conditioning own types is equal to beliefs

conditioning on types which in turn is equal to actual play of that player. Hence **actual play**, and not merely beliefs about play, converges to a Nash equilibrium. I.e.,

10.5. Corollary. Let $\{\mu_i\}_{i \in I}$ be a BSP and suppose condition (GGH) holds. Also suppose that (6.4) holds and the type space T is a singleton (or equivalently the true vector of types is common knowledge.) Define μ^{**} to be the **true** distribution of play, i.e., that induced by $\{\mu_i\}_{i \in I}$. Define

$$G^{**} \equiv \{(\tau, s^{\infty}): \mu^{**}(ds_{N+1} | s^N) \twoheadrightarrow \mathrm{ND}(\theta(\tau)) \text{ for all } i \in I\}$$

Then $\mu^*(G^{**} \cap W) = 1$ where W is as in (8.3).

11. Concluding Remark

We quote from Aumann (1987):

"The equilibrium concept of Nash... is without doubt the single game-theoretic tool that is most applied in economics. Yet...a little reflection leads to some puzzlement as to why and under what conditions players...might be expected to play such an equilibrium."

The axioms of Savage (1954) imply Bayesian Rationality; i.e., players maximize subjective expected utility. Aumann (1987) argues that Bayesian Rationality (without common priors) is equivalent to a **subjective** correlated equilibrium. We show that the repeated play of a game where Bayesian Rationality is assumed and there is agreement on probability zero events (and nothing else!) leads to an **objective** correlated equilibrium. We therefore have a justification of such equilibrium in terms of the behavior of players over time under Bayesian rationality.

12. Appendix: The proofs

Proof of Theorem 9.1.: The proof will proceed in three steps: In Step I we use condition (GGH) to relate convergence for any player *i* on sample paths with μ_i -probability one to convergence for **all** players, on sets with μ^* -probability one. In particular, we first prove (in lemma 9.1.1 below) that³ it suffices to show that for each *i*, player i's beliefs about the next period play of the game **Not** conditioning on *i*'s type τ_i , i.e., $\mu_i(ds_{n+1}|s^n)$, converge to the set of correlated distributions of one-period play where *i* is best responding, $C_i(\theta_i)$. This will be the only part of the proof using condition (GGH).

We will define G_i to be the set where the above mentioned convergence occurs. In Step II of our proof we use the separating hyperplane theorem to show that on any sample path **outside** of G_i there will exist two disjoint sets (which we will denote by B and Φ), which can be separated by a vector x, such that for infinitely many dates n, beliefs conditional only on history (and **not** on own types) lie in B while beliefs conditional on history **and** own types belongs to Φ (i.e., $\mu_i(ds_{-in+1}|s^n, s_{in+1}) \in B$ and $\mu_i(ds_{-in+1}|s^n, s_{in+1}, \tau_i) \in \phi$).

In Step III of our proof we show the following: The set of sample paths where infinitely often beliefs conditional on type lie one set, B, and beliefs not conditional on types lie in another set, ϕ , is a set of sample paths with zero probability if the

sets B and ϕ are disjoint and can be separated. This part of the proof is stated with some generality in the hope that it will be of independent interest.

Step II and Step III imply that $\mu_i(G_i) = 1$. Step I implies that this suffices to prove the theorem. Now the details!

Step I: Recall that the set G is defined in (9.2) and W in (8.3). We seek to show that

$$\mu^{*}(\{(\tau, s^{\infty}) \in G \cap W\}) = 1.$$
(12.1)

Define

$$G_i \equiv \{(\tau, s^{\infty}): \mu_i(ds_{n+1}|s^n) \twoheadrightarrow C_i(\theta_i(\tau_i))\}$$
(12.2)

where, recall, the convergence \rightarrow is as defined in the beginning of Section 9. We have the following:

Lemma 9.1.1. To prove (12.1) it suffices to prove that for each $i \in I$, $\mu_i(G_i) = 1$ where G_i is as in (12.2).

Proof. On G_i any convergent sub-sequence of $\{\mu_i(ds_{n+1}|s^n)\}_{n=1}^{\infty}$ converges to the set $C_i(\theta_i(\tau_i))$. From the definition of the set W in (8.3), on W the measures $\mu_i(ds_{n+1}|s^n)$ and $\mu_j(ds_{n+1}|s^n)$ become closer and closer to each other as $n \to \infty$ for each i and j in I. Hence it is easy to see that on $G_i \cap W$,

$$\mu_j(ds_{n+1}|s^n) \twoheadrightarrow C_i(\theta_i) \quad \text{for all } j \in I.$$
(12.3)

Hence on $\bigcap_{i \in I} G_i \cap W$, (12.3) holds for each *i* and *j* in *I*. Since $C(\theta) = \bigcap_{i \in I} C_i(\theta_i)$, we conclude that on $\bigcap_{i \in I} G_i \cap W$, $\mu_j(ds_{n+1}|s^n) \twoheadrightarrow C(\theta)$ for all $j \in I$. Hence $\bigcap_{i \in I} G_i \cap W$ is a subset of *G*. For (12.1) it therefore suffices to show that $\mu^*(\bigcap_{i \in I} G_i \cap W) = 1$. From Theorem 8.2, $\mu^*(W) = 1$. Hence it remains only to show that $\mu^*(G_i) = 1$ for each *i* in *I*. For this, from condition (GGH) it is easy to see that it suffices to show that $\mu_i(G_i) = 1$ for each $i \in I$. #

Step II: The separation argument. From now on fix an $i \in I$. Fix any x and y in $\Re^{S_{-i}}$. Let $x \cdot y$ denote their inner product. The vector x shall be called a rational vector if each of its coordinates is a rational number. Any probability over S_{-i} is an element of $\Re^{S_{-i}}$; it will be called a rational probability if it is a rational vector. Let $Z \equiv \{1, 2, ...\}$, the set of all positive integers. Define for $x \in \Re^{S_{-i}}$ and $k \in Z$,

 $B(x, 1/k) \equiv$ the open ball with center x and radius 1/k; (12.4)

$$\mathfrak{B} \equiv \{B(x, 1/k) | x \in \mathfrak{R}^{S-i} \text{ and } x \text{ is rational and } k \in \mathbb{Z}\};$$
(12.5)

$$\phi_{\mathbf{x}} \equiv \{ q \in \mathfrak{R}^{\mathbf{S}_{-i}} | q \cdot \mathbf{x} \ge 0 \}; \tag{12.6}$$

$$\boldsymbol{\Phi} \equiv \{ \boldsymbol{\phi}_x | x \in \Re^{S_{-i}} \text{ and } x \text{ is rational} \};$$
(12.7)

and define the cartesian product,

 $J \equiv \{(x, B, \phi, k) : x \in \Re^{S_{-i}} \text{ and } x \text{ is rational, } B \in \mathfrak{B}, \phi \in \Phi \text{ and } k \in \mathbb{Z}\}$ (12.8)

Note that the "index" set J is a countable set.

Define $P(\tau_i, s_i^*)$ to be the set of all probability distributions of play of the other players against which the action $s_i^* \in S_i$ is a best response for the player *i* of type τ_i ; i.e.,

$$P(\tau_i, s_i^*) \equiv \{q_{-i} \in \mathscr{P}(S_{-i}) | s_i^* \in \operatorname{Argmax} \int u_i(\theta_i(\tau_i), \dots, s_{-i}) q_{-i}(ds_{-i})\}$$
(12.9)

Bayesian learning

Recall that the set of probability measures is endowed with the weak topology. (Indeed, since here we assume that S is finite, this reduces to the usual topology on the unit simplex.) Further, expected utility is (weakly) continuous in probabilities. It is therefore easy to check that $P(\tau_i, s_i^*)$ is compact and convex. We therefore conclude from the separating hyperplane theorem that if $v_{-i} \notin P(\tau_i, s_i^*)$ then v_{-i} and $P(\tau_i, s_i^*)$ can be "separated" by some vector x. More precisely, the following can easily be shown:

Lemma 9.1.2. Fix any $\tau_i \in T_i$, $s_i^* \in S_i$ and $v_{-i} \in \mathscr{P}(S_{-i})$. Suppose that $v_{-i} \notin P(\tau_i, s_i^*)$. Then there exists a $(x, B, \phi, k) \in J$ such that $v_{-i} \in B$, $P(\tau_i, s_i^*) \subseteq \phi$ and "x separates the sets B and ϕ by at least 1/k"; i.e.,

$$b.x \le -1/k < 0 \le p \cdot x$$
 for all $b \in B$ and $p \in \phi$. (12.10)

(It should be easy to see that for lemma 9.1.2 we do not have to worry about the possibility that $P(\tau_i, s_i^*)$ is empty; (12.10) still holds in that case.)

Recall that the set G_i is defined in (12.2). In the following lemma we show that if $(\tau, s^{\infty}) \notin G_i$, then there are two sets, B and ϕ , which can be separated in the sense of (12.10) such that following is true infinitely often: Beliefs not conditional upon types, $\mu_i(ds_{-in+1}|s^n, s_{in+1})$, lie in the set B, and beliefs conditional on types, $\mu_i(ds_{-in+1}|s^n, s_{in+1}, \tau_i)$, lie in ϕ .

Lemma 9.1.3. For μ_i -a.e. $(\tau, s^{\infty}) \notin G_i$ the following is true: There exists a $(x, B, \phi, k) \in J$ such that (12.10) holds and an $s_i^* \in S_i$ such that if we define

$$D \equiv \{\tau_i' \in T_i | P_i(\tau_i', s_i^*) \subseteq \phi\}$$

and

$$H_n \equiv \{ (\tau', s'^{\infty}) \in T \times S^{\infty} | s'_{in+1} = s_i^* \text{ and } \mu_i (ds_{-in+1} | s'^n, s'_{in+1}) \in B \}$$

then

i. for infinitely many dates n, (τ, s^{∞}) lies in the set H_n ; and ii. $\tau_i \in D$.

Remark 9.1.3': Let the sets D and H_n be as defined in Lemma 9.1.3 above. If $(\tau', s'^{\infty}) \in H_n$ then player *i* type τ'_i chooses the action $s'_{in+1} = s^*_i$ at date n + 1. Since μ_i is part of a BSP, player *i*'s action is optimal (μ_i -a.e.). Hence $\mu_i(ds_{-in+1}|s'^n, s'_{in+1}, \tau'_i) \in P_i(\tau'_i, s^*_i)$. If $\tau'_i \in D$ then $P_i(\tau'_i, s^*_i) \subseteq \phi$. Hence μ_i -a.e.,

$$\mu_i(ds_{-in+1}|s^{\prime n},s_{in+1}^{\prime},\tau_i^{\prime})\in\phi\quad\text{on }H_n\cap\{\tau_i\in D\}.$$

Lemma 9.1.3 therefore implies that for infinitely many dates beliefs conditional on types lie in the set ϕ .

Proof of Lemma 9.1.3. We begin with an immediate implication of the conditional Borel-Cantelli lemma. Fix any $B \in \mathfrak{B}$ and $s_i^* \in S_i$. For ease of exposition we define

$$\mu_{in}(\cdot|s_i^*) \equiv \mu_i(ds_{-in+1}|s^n, s_{in+1} = s_i^*).$$
(12.11)

Define the events $H'_n \equiv \{s_{in+1} = s_i^*\}$ and $H''_n \equiv \{\mu_{in}(\cdot|s_i^*) \in B\}$. Then $H_n = H'_n \cap H''_n$. Let $1_{\{\cdot\}}$ denote the indicator function on an event $\{\cdot\}$ (equal to one on the event and zero otherwise). The event H''_n is measurable with respect to the σ -algebra generated by $\{s^n, s_{in+1}\}$. Using this fact for the third inequality below we obtain,

$$\sum_{n=1}^{\infty} \mu_i(H_n|s^n) = \sum_{n=1}^{\infty} \mu_i(H'_n \cap H''_n|s^n) = \sum_{n=1}^{\infty} \mu_i(H''_n|H'_n, s^n) \mu_i(H'_n|s^n)$$
$$= \sum_{n=1}^{\infty} 1_{(H''_n)} \mu_i(H'_n|s^n) = \sum_{n=1}^{\infty} 1_{(\mu_i \cap (\cdot|s^*_i) \in B)} \mu_i(\{s_{in+1} = s^*_i\}|s^n). \quad (12.12)$$

Since H_n is measurable with respect to the σ -algebra generated by s^{n+1} , the (conditional) Borel-Cantelli lemma (see, e.g., Chow et al., 1971, p. 26) implies that on a set of sample paths with μ_i -probability one

"whenever (12.12) is infinite the event H_n occurs for infinitely many n." The μ_i probability one set of sample paths where the statement in quotes above holds may depend on B and s_i^* . However since \mathfrak{B} and S_i are both countable there exists a set of sample paths with μ_i -probability one where the statement in quotes holds for all $B \in \mathfrak{B}$ and $s_i^* \in S_i$ simultaneously.

Fix any $(\tau, s^{\infty}) \notin G_i$. Then there exists a subsequence of dates, $\{n(t)\}_{i=1}^{\infty}$ and a measure $v \in \mathscr{P}(S)$ such that $\mu_i(ds_{n(t)+1}|s^{n(t)}) \to v$ as $t \to \infty$ and $v \notin C_i(\theta_i(\tau_i))$. $v \notin C_i(\theta_i(\tau_i))$ implies that there exists an $s_i^* \in S_i$ such that $v(s_i^*) > 0$ and such that if we define $v_{-i} \equiv v_{-i}(ds_{-i}|s_i = s_i^*)$ then s_i^* is not optimal for the player-type τ_i against v_{-i} . In particular, $v_{-i} \notin P_i(\tau_i, s_i^*)$. Given this s_i^* and v_{-i} , we may use lemma 9.1.2 to obtain a $(x, B, \phi, k) \in J$ such that (12.10) holds, $v_{-i} \in B$ and $\tau_i \in D$ (where the set D is as defined in the statement of this lemma). In particular part (ii) of this lemma holds.

Now, $\mu_i(ds_{n(t)+1}|s^{n(t)}) \rightarrow v$ and $v(s_i^*) > 0$ imply that for all t sufficiently large,

$$\mu_i(\{s_{in(t)+1} = s_i^*\} | s^{n(t)}) > \nu(s_i^*)/2 > 0$$
(12.13)

and also that

$$\mu_{in(t)}(\cdot | s_i^*) \to v(\cdot | s_i = s_i^*) \equiv v_{-i}$$
(12.14)

Since v_{-i} lies in the open ball *B* we conclude from (12.14) that for all *t* sufficiently large, $\mu_{in(t)}(\cdot|s_i^*) \in B$. This fact and (12.13) imply that for all dates n(t) for *t* sufficiently large the expression in the summation in the right-hand-side of (12.12) will exceed $v(s_i^*)/2$. Hence the summation in (12.12) is infinite. The discussion above on conditional Borel-Cantelli lemma then implies that part (i) of this lemma holds for μ_i -a.e. $(\tau, s^{\infty}) \notin G_i$.

Step III: In our next result we conclude that we may assign probability zero to the event that there is a type such that infinitely often beliefs conditional on that type lie in one set while beliefs not conditional on type lie in another set if those two sets can be separated in the sense of (12.10). We state this result in a general framework in the hope that it may be of independent interest.

Lemma 9.1.4. Let (Ω, \Im, μ) be a probability space, with Ω a complete and separable metric space, \Im its Borel subsets and μ a probability measure over (Ω, \Im) . Let $\{\Im_n\}_{n=1}^{\infty}$ be sequence of non-decreasing sub- σ -fields of \Im with $\Im_{\infty} = V_{n=1}^{\infty} \Im_n$. Let $\{x_n\}_{n=1}^{\infty}$ be a uniformly bounded sequence of real-valued random variables on (Ω, \Im) with x_n measurable with respect to \Im_n for all *n*. Let ξ be a random variable on (Ω, \Im) taking values in a complete and separable metric space \widehat{T} . Let *D* be a (Borel measurable) subset of \widehat{T} . Let $\{H_n\}_{n=1}^{\infty}$ be a sequence of events such that for each *n*, H_n is measurable with respect to \Im_n . Let k > 0 be a constant. Suppose that μ -a.e.

Bayesian learning

the following is true:

- (a) $E[x_{n+1}|\Im_n,\xi] \ge 0$ on the event $H_n \cap \xi^{-1}(D)$; and
- (b) $E[x_{n+1}|\mathfrak{I}_n] \leq -1/k$ on the event H_n .

Then $\mu(\{\xi \in D \text{ and } H_n \text{ occurs infinitely often}\}) = 0.$

Proof. Let $1_D \equiv 1_{\{\xi \in D\}}$ denote the indicator function on D (equal to one if $\xi \in D$ and zero otherwise). We first show that μ -a.e.,

 $\lim_{n\to\infty} \{ E[\mathbf{1}_D E[x_{n+1}|\mathfrak{I}_n,\xi]|\mathfrak{I}_n] - E[x_{n+1}|\mathfrak{I}_n] E[\mathbf{1}_D|\mathfrak{I}_n] \} = 0 \ \mu\text{-a.e.} \quad (12.15)$

Iterated use of the elementary properties of conditional expectations results in the following:

$$E[1_{D}E[x_{n+1}|\mathfrak{I}_{n},\xi]|\mathfrak{I}_{n}] = E[E[1_{D}x_{n+1}|\mathfrak{I}_{n},\xi]|\mathfrak{I}_{n}] = E[1_{D}x_{n+1}|\mathfrak{I}_{n}]$$

= $E[E[1_{D}x_{n+1}|\mathfrak{I}_{n+1}]|\mathfrak{I}_{n}] = E[x_{n+1}E[1_{D}|\mathfrak{I}_{n+1}]|\mathfrak{I}_{n}];$ (12.16)

and

$$E[x_{n+1}|\mathfrak{I}_n]E[\mathfrak{I}_D|\mathfrak{I}_n] = E[x_{n+1}E[\mathfrak{I}_D|\mathfrak{I}_n]|\mathfrak{I}_n]$$
(12.17)

Define

$$\Delta_{n} \equiv x_{n+1} (E[1_{D} | \mathfrak{I}_{n+1}] - E[1_{D} | \mathfrak{I}_{n}])$$
(12.18)

From the martingale convergence theorem one may show that $E[1_D|\mathfrak{I}_n]$ converges μ -a.e. as $n \to \infty$ (to $E[1_D|\mathfrak{I}_\infty]$). Hence $\Delta_n \to 0$ as $n \to \infty$ μ -a.e. This is turn implies (via another application of the martingale convergence theorem – see, e.g., Chung, 1974, Theorem 9.4.8) that

$$\lim_{n \to \infty} E[\Delta_n | \Im_n] = 0 \quad \mu\text{-a.e.}$$
(12.19)

It is easy to see that $E[\Delta_n | \Im_n]$ is equal to the right hand side of (12.16) minus the right hand side of (12.17). Since the expression in brackets in (12.15) is the left hand side of (12.16) minus the left hand side of (12.17), (12.19) implies (12.15).

From conditions (a) and (b) of the hypothesis of this lemma, on H_n the first term in (12.15) is non-negative and the first part of the second term, $E[x_{n+1}|\mathfrak{I}_n]$, is no greater than -1/k. From the martingale convergence theorem the second part of the second term, $E[\mathfrak{I}_D|\mathfrak{I}_n]$, converges to $E[\mathfrak{I}_D|\mathfrak{I}_\infty]$. (All of the above statements are of course true μ -a.e.) Hence (12.15) implies that on the event where H_n occurs infinitely often,

$$0 = \operatorname{limsup}_{n \to \infty} \left\{ E[\mathbf{1}_{D} E[\mathbf{x}_{n+1} | \mathfrak{I}_{n}, \xi] | \mathfrak{I}_{n}] - E[\mathbf{x}_{n+1} | \mathfrak{I}_{n}] E[\mathbf{1}_{D} | \mathfrak{I}_{n}] \right\}$$

$$\geq (1/k) E[\mathbf{1}_{D} | \mathfrak{I}_{\infty}] \geq 0.$$
(12.20)

Hence $E[1_D|\mathfrak{I}_{\infty}] = 0$ on the event where H_n infinitely often. This then implies the conclusion of the lemma by integrating over the set of sample paths where H_n occurs infinitely often. //

Proof of Theorem 9.1. (cont'd). From Lemma 9.1.1 it suffices to show that for all $i \in I$.

$$\mu_i(\{(\tau, s^{\infty}) \notin G_i\}) = 0.$$
 (12.21)

Fix an $i \in I$. From lemma 9.1.3 for μ_i -a.e. $(\tau, s^{\infty}) \notin G_i$, there exists an element of $J \times S_i$

satisfying the conclusions (i)–(ii) of lemma 9.1.3. Hence to show (12.21) it suffices to show that the event where these conclusions hold has μ_i -probability zero. For this, since $J \times S_i$ is countable, it suffices to show that for each *fixed* element of $J \times S_i$ the event where these conclusions hold has μ_i -probability zero.

So fix an element $(x, B, \phi, k, s_i^*) \in J \times S_i$ and suppose that (12.10) holds. Let the sets D and H_n be as in Lemma 9.1.3. We seek to show that

$$\mu_i(\{\tau_i \in D \text{ and } H_n \text{ occurs infinitely often}\}) = 0.$$
(12.22)

Define $x_{n+1} = x(s_{-in+1})$ and $\mathfrak{I}_n = \sigma(\{s^n, s_{in+1}\})$ (i.e., the σ -algebra generated by the indicated variables). Then x_n is \mathfrak{I}_n -measurable. From remark 9.1.3' and (12.10), on the event $H_n \cap \{\tau_i \in D\}$,

$$E[x_{n+1}|\mathfrak{I}_n,\tau_i] \equiv \mu_i(ds_{-in+1}|s^n,s_{in+1},\tau_i)\cdot x \ge 0 \quad \text{for all } \tau_i \in D. \quad (12.23a)$$

Further on H_n , $\mu_i(ds_{-in+1}|s'^n, s'_{in+1}) \in B$ so (12.10) implies that

$$E[x_{n+1}|\mathfrak{I}_n] \equiv \mu_i(ds_{-in+1}|s^n, s_{in+1}) \cdot x \le -1/k.$$
 (12.23b)

Lemma 9.1.4 then implies (12.22). //

Proof of Theorem 9.6. Fix any probability μ over $T \times S^{\infty}$. Let \mathscr{L} denote the class of closed intervals in [0,1] with rational end-points. Note that the class \mathscr{L} is countable. Then we have the following claim:

Claim: For each (τ, s^{∞}) , excluding possibly a set with zero μ -probability, the following is true: For all $A \in \mathscr{L}$ and $s^* \in S$, if we define $z_{n-1} \equiv 1_{\{\mu(\{s_n = s^*\} \mid s^{n-1}) \in A\}}$, then on the event where $\sum_{n=1}^{\infty} z_{n-1} = \infty$,

$$\lim_{N \to \infty} \sum_{n=1}^{N} \mathbb{1}_{\{s_n = s^*\}} z_{n-1} / \sum_{n=1}^{N} z_{n-1} \in \Lambda.$$
 (12.24)

Proof of Claim: Fix any $\mu \in \mathscr{P}(T \times S^{\infty})$. Since \mathscr{L} and S are both countable it suffices to prove that (12.24) holds μ -a.e. for a **fixed** $\Lambda \in \mathscr{L}$ and $s^* \in S$; so fix any such Λ and s^* . Let z_{n-1} be as in the claim. Define $y_n \equiv 1_{\{s_n = s^*\}}$. Whenever $z_{n-1} = 1$, $E(y_n | s^{n-1}) \equiv \mu(\{s_n = s^*\} | s^{n-1}) \in \Lambda$ so $E(y_n | s^{n-1}) \cdot z_{n-1} \in \Lambda$. Therefore, on the event where $\sum_{n=1}^{\infty} z_{n-1} > 0$,

$$\operatorname{Lim}_{N \to \infty} \sum_{n=1}^{N} E[y_n | s^{n-1}] \cdot z_{n-1} / \sum_{n=1}^{N} z_{n-1} \in \Lambda.$$
(12.25)

Define $\xi_n \equiv y_n - E(y_n|s^{n-1})$ where the expectation is taken with respect to the conditional probability $\mu(\cdot|s^{n-1})$. Then $E[\xi_n|s^{n-1}] = 0$ and $|\xi_n| \le 1$. Using the Martingale Convergence Theorem and Kronecker's lemma it is easy to show the following (see Taylor, 1974): On the event where $\sum_{n=1}^{\infty} z_{n-1} = \infty$,

$$\operatorname{Lim}_{N \to \infty} \sum_{n=1}^{N} \xi_{n} \cdot z_{n-1} / \sum_{n=1}^{N} z_{n-1} = 0$$
 (12.26)

(The Taylor (1974) argument uses independence of ξ_n but this is not required. Indeed, define $M_N \equiv \sum_{n=1}^N [z_{n-1}\xi_n/\sum_{t=1}^n z_{t-1}]$ whenever the denominator is non-zero (and set $M_N = 0$ otherwise). Since $E[\xi_n|s^{n-1}] = 0$, $\{M_N\}_{N=1}^{\infty}$ is a Martingale sequence with respect to information generated by the partial history of actions. Since $|\xi_n| \leq 1$, $E[M_N^2] \leq E \sum_{n=1}^N z_{n-1}/(\sum_{t=1}^n z_{n-1})^2$. One may then proceed just as in Taylor (1974): Lemma 1 of Taylor (1974) implies that $EM_N^2 \leq 2$; the Martingale convergence theorem implies M_N converges; Kronecker's lemma then implies (12.26).)

From (12.25) and (12.26) we may conclude that (12.24) holds. //

Proof of Theorem 9.6 (Cont'd): Part (b) of the Theorem is the same as the claim above. To prove part (a) of the Theorem we proceed as follows. Fix any sample path and let v as in part (a) of the Theorem. Fix any s^* in the support of v. Let $\{\Lambda^k\}_{k=1}^{\infty}$ be a sequence of intervals in \mathscr{L} monotonically decreasing to the point $v(\{s^*\})$. Fix any integer k. If $\mu(\{s_{n+1} = s^*\} | s^n)$ converges to $v(\{s^*\})$ then $\mu(\{s_{n+1} = s^*\} | s^n) \in \Lambda^k$ for all n sufficiently large. In particular, if z_n is defined as in the claim above but with $\Lambda = \Lambda_k$, then $z_n = 1$ for all n sufficiently large. Hence, the claim implies that for each k, $\lim_{N\to\infty} \sum_{n=1}^{N} 1_{\{s_n = s^*\}} / N \in \Lambda_k$. Taking limits as $k \to \infty$ then proves part (a) of the Theorem.

References

- Aumann, R.: Subjectivity and correlation in randomized strategies. J. Math. Econ. 1, 67-96 (1974)
- Aumann, R.: Correlated equilibrium as an expression of Bayesian rationality. Econometrica 55, 1–18 (1987)
- Blackwell, D., Dubins, L.: Merging of opinions with increasing information. Ann. Math. Stat. **38**, 882–886 (1963)
- Blume, L., Easley, D.: Rational expectations equilibrium: An alternative approach. J. Econ. Theory 34, 116-129 (1984)
- Chow, Y. S., Robbins, H., Siegmund, D.: Great expectations: The theory of optimal stopping. Boston: Houghton Mifflin Co. 1971
- Cotter, K.: Correlated equilibrium in games with type-dependent strategies. J. Econ. Theory 54, 48–68 (1991)
- Chung, K. L.: A course in probability theory. New York: Academic Press 1974
- Easley, D., Kiefer, N.: Controlling a stochastic process with unknown parameters. Econometrica, 56, 1045-1064 (1988)
- Feldman, M.: An example of convergence to rational expectations with heterogeneous beliefs. Int. Econ. Rev. 28, 635–650 (1988)
- Harsanyi, J. C.: Games with incomplete information played by Bayesian players, Parts I, II, III. Management Sci. 14, 159-182, 320-334, 486-502 (1967/1968)
- Harsanyi, J. C.: Games with randomly distributed payoffs: A new rationale for mixed strategy equilibrium points. Int. J. Game Theory 2, 1–23 (1973)
- Jordan, J. S.: Bayesian learning in normal form games. Games Econ. Behav. 3, 60-81 (1991a)
- Jordan, J. S.: Bayesian learning in repeated games. Manuscript, University of Minnesota 1991b
- Jordan, J. S.: Bayesian learning in games: a non-Bayesian perspective. Manuscript, University of Minnesota 1992
- Kalai, E., Lehrer, E.: Bayesian learning and Nash equilibrium. Manuscript, Northwestern University 1990
- Koutsougeras, L., Yannelis N.: Convergence and approximation results for non-cooperative Bayesian games: learning theorems. Econ. Theory 4, 843–858 (1994)
- Nyarko, Y.: Learning in mis-specified models and the possibility of cycles. J. Econ. Theory 55, 416–427 (1991a)
- Nyarko, N.: The convergence of Bayesian belief hierarchies. C. V. Starr Center Working Paper No. 91-50, New York University 1991b
- Nyarko, Y.: Bayesian learning without common priors and convergence to Nash equilibria. Manuscript, New York University 1992
- Savage, L.: The foundations of statistics. New York: Wiley 1954
- Yannelis, N., Rustichini, A.: On the existence of correlated equilibria. In: Khan, M., Yannelis, N. (eds.) Equilibrium theory with infinitely many commodities, pp. 268–280. Berlin, Heidelberg, New York: Springer-Verlag 1991