Discussion Paper No. 709R

CONVERGENCE OF GAMES WITH ASYMMETRIC INFORMATION*

by

Kevin D. Cotter
Department of Economics
Northwestern University

Revised April 1987

Running Head: Convergence of games

Mailing address: Department of Economics
Northwestern University
2003 Sheridan Road
Evanston, IL 60201 U.S.A.

*Support from NSF grant IST-8609208 is acknowledged. An earlier version of this paper, titled "Nash Equilibrium with General Asymmetric Information," appeared as Discussion Paper #709, Center for Mathematical Studies in Economics and Management Science, Northwestern University. I am indebted to Maxwell Stinchcombe and an anonymous referee for pointing out an error in versions of Theorems 2.4 and 3.1 of that paper. Any remaining errors are my own.

Abstract

A general model of normal form games with uncertainty about payoffs and abstract player information is proposed, with an appropriate notion of convergence of game characteristics. Conditions analogous to those used by Milgrom and Weber are shown to be sufficient for players' expected payoff functions to be jointly continuous with respect to strategies. Under these conditions, the equilibrium correspondence is nonempty valued and upperhemicontinuous with respect to players' characteristics and the Boylan metric, but not the pointwise convergence metric, of information. When a player's information is nonatomic, the set of pure strategies is dense in the strategy space. Conditions for which every Nash equilibrium has a purification when each player's action space is infinite are given.

1. Introduction

In this paper the behavioral similarity of players' information in normal form games is characterized, using recent work which models information as an element of a well behaved metric space. Given a probability space which determines all exogenous uncertainty, Boylan (1971) defined the space of all possible information fields about the state of nature with a complete metric on that space. The economic implications of the Boylan metric has been studied by Allen (1986), who proved that the demand of a competitive utility maximizing consumer is continuous with respect to information about utility. An alternative is the pointwise convergence metric proposed by Cotter (1986), which is the weakest topology satisfying the above continuity of demand. One way of characterizing the behavioral similarity of games with asymmetric information is to identify the weakest topology (or metric) with respect to which game behavior is continuous. This metric can be used to study some fundamental properties of Nash equilibria which involve relationships between players' information and the resulting game. For example, robustness of game-theoretic properties such as the prevelance of games with pure strategy Nash equilibria can be studied. The stability of equilibria with respect to perturbations of game characteristics can also be examined. Some progress on the latter has been made in an elementary model by Fundenberg, Kreps, and Levine (1986) as a means of refining the definition of Nash equilibrium as an alternative to using perturbations of strategies [e.g., Kohlberg and Mertens (1986); Kreps and Wilson (1982)]. In addition, there are cases where information is required to be part of a metric space. This arises when information is part of a player's strategy, such as in signalling and information transmission, and when other players' information is part of a player's belief hierarchy, such as the one constructed by Mertens and Zamir (1985).

To study the relationship between players' information and their behavior a theory of normal form games with abstract information fields is presented in Section 2 of this paper. The key result gives conditions, analogous to those used by Milgrom and Weber (1981, Assumptions R1, R2) which are sufficient for the expected payoff function of any player to be continuous with respect to the strategies of all players. The proof, however, is complicated by the definition of players' information. Existence of Nash equilibria follows easily. In Section 3, the conditional expected value of each player's payoff is shown to be jointly continuous with respect all players' payoff functions, strategies, and information, when the Boylan metric of information is used. Therefore the Nash equilibrium correspondence is upperhemicontinuous. Using the pointwise convergence metric, the conditional expected value of each player's payoff is jointly continuous with respect to his own strategy, payoff function, and information, but not jointly continuous with respect to other players' strategies and information.

In Section 4, some results about pure strategies are proven. The set of pure strategies is shown to be a dense subset of the strategy space whenever the player's information field is nonatomic. In addition, the set of nonatomic information fields is shown to be dense with respect to the pointwise convergence metric, but not the Boylan metric, whenever the underlying probability space is nonatomic. Finally, some sufficient conditions are given for which every Nash equilibrium in behavioral strategies has a corresponding pure strategy equilibrium. These conditions are similar to, but more general than, Theorem 3 of Radner and Rosenthal (1982).

2. The model

Consider a game with a fixed finite set of players $I = \{1,2,\ldots,I\}$. Each player $i \in I$ has an underlying set of possible actions A_i, a compact metric space. Let $A = \prod_{i \in I} A_i$ be the space of joint actions of all players. All exogenous uncertainty in the game is generated by a common probability space $(\Omega, \mathcal{F}, \mu)$, where Ω is a set of possible states of nature, $\mathcal{F}$ is a countably generated σ-field of measurable subsets (events) of Ω, and μ is a probability measure on $\mathcal{F}$. This probability space affects players only through their payoff functions $v_i : \Omega \times A \to \mathbb{R}$. Finally, each player has some private information about Ω. Following Boylan (1971), let the space of information $\mathcal{F}^*$ be the set of all possible sub-σ-fields of $\mathcal{F}$ (i.e., measurable partitions of Ω) modulo null sets. Each player's information is some information field $\mathcal{Y}_i \in \mathcal{F}^*$.

This definition of information generalizes others which have been used in game theory. For example, information is defined by Radner and Rosenthal (1982) to be a random variable which is correlated with the player's payoff function. This restriction to random variables does not involve any practical loss of generality. However, the dependence of the game on the underlying information structure cannot be studied using random variables since changes in observed signals do not correspond to changes in the information they convey [Cotter (1987)]. This model also generalizes Milgrom and Weber (1981), who assumed that each player has a privately observable type space and that the information structure for all players is a joint probability distribution over all player type spaces. In the present model, each player's type is the observed expected payoff function conditional on the player's own information. As with the Radner-Rosenthal model, changes in information cannot be easily examined in the Milgrom-Weber model since varying the probability distribution over types changes not only players' information but the underlying structure of uncertainty as well. In particular, there is no meaning to better or finer information in their model. Defining a tighter probability distribution over types places increased weight on a particular type, so to define that as better information presumes that a particular type is the "correct" one.

To model a game with abstract private information, an appropriate definition of player strategies is needed. Two possibilities are suggested by the literature. Let $\mathcal{A}_i$ be the σ-field of Borel sets of A_i. Milgrom and Weber (1981) defined a distributional strategy to be a joint probability distribution over types and actions whose marginal distribution on types is the one given by the information structure. In this model, a distributional strategy is a joint probability distribution on $(A_i,\mathcal{A}_i) \times (\Omega,\mathcal{J}_i)$ whose marginal distribution on $(\Omega,\mathcal{J}_i)$ is μ restricted to $\mathcal{J}_i$. Such a definition requires that Ω be a metric space and $\mathcal{J}_i$ be contained in the Borel sets of Ω. In practice, this restriction does not entail much loss of generality since Ω can be taken to be the metric space of possible payoff functions, with its probability measure given by the map from states of nature to payoffs. However, distributional strategies are not suitable for the present model. An alternative, used by Radner and Rosenthal (1982) and this paper, is to define a behavioral strategy to be a function mapping the state space (or players' type spaces) into the set of probability distributions on $(A_i,\mathcal{A}_i)$ which is consistent with the player's information. The interpretation is that a player makes an observation based on the state, then chooses a mixed strategy over actions. In this model, a behavioral strategy is a function $s_i: \Omega \times \mathcal{A}_i \to \mathbb{R}$ such that for each $B \in \mathcal{A}_i$, $s_i(\cdot,B)$ is $\mathcal{J}_i$-measurable, and for a.e. ω, $s_i(\omega,\cdot)$ is a probability measure on A_i.

In this paper some basic properties of Nash equilibria of normal form games in behavioral strategies will be studied. Let S_i be the set of behavioral strategies for player i. To establish the existence of a Nash equilibrium in behavioral strategies using the well-known method of Glicksberg

(1952), a topology on S_i must be defined for which the expected payoff of each players is continuous with respect to all players' strategies. At the same time, S_i must be a compact, convex subset of a locally convex topological vector space. Radner and Rosenthal (1982) solved this problem by giving S_i the weak* topology. In this model, a sequence of strategies $\{s_i^n\}$ converges in the weak* topology if for every measurable function $g\colon \Omega \times A_i \to \mathbb{R}$ with $g(\omega,\cdot)$ continuous a.e. and $\int_\Omega \sup_{a_i \in A_i} |g(\omega,a)|\mu(d\omega)$ finite, $\int_\Omega [\int_{A_i} g(\omega,a)s_i^n(\omega,da)]\mu(d\omega)$ converges to $\int_\Omega [\int_{A_i} g(\omega,a)s_i(\omega,da)]\mu(d\omega)$. The weak* topology in the Radner-Rosenthal model is more easily constructed by their assumption that each player's action space is finite. When the action space is infinite, a more delicate treatment is required. Let $C(A_i)$ be the space of real continuous functions on A_i with the norm topology of uniform convergence. Then $C(A_i)$ is a separable Banach space with dual $M(A_i)$, the space of finite signed Borel measures on A_i with the duality $\langle h,\upsilon\rangle = \int_{A_i} h(a)\upsilon(da)$. Then $M(A_i)$ is also a separable Banach space with the variation norm $\|\upsilon\| = \sup_{B \subset A_i} |\upsilon(B)|$. In addition, $L^1(C(A_i)) = \{F\colon \Omega \to C(A_i) | F$ is Borel-measurable and $\|F\|_1 = \int_\Omega \|f(\omega)\|\mu(d\omega)$ is finite$\}$ is a separable Banach space [Neveu (1975, Proposition V-2-5)]. Then by Diestel and Uhl (1977, Theorem 1, p. 79; Theorem 1, p. 98), the dual of $L^1(C(A_i))$ is $L^\infty(M(A_i)) = \{G\colon \Omega \to M(A_i) | G$ is Borel-measurable and $\|G\|_\infty = \operatorname{ess\,sup}_{\omega \in \Omega} \|G(\omega)\|$ is finite$\}$ with the duality $\langle\langle F,G\rangle\rangle = \int_\Omega \langle F(\omega),G(\omega)\rangle\mu(d\omega)$. Give $L^\infty(M(A_i))$ the weak* topology, and let $S_i = \{s_i \in L^\infty(M(A_i)) | s_i(\omega)$ is a probability measure on A_i for a.e. $\omega\}$. S_i is closed and convex. By the Banach-Alaoglu theorem [Rudin (1973, p. 66)], S_i is compact, and also metrizable since $L^1(C(A_i))$ is separable [Rudin (1973, p. 68)].

Note that S_i is defined without reference to $\mathcal{Y}_i$-measurability. Strategies that are not $\mathcal{Y}_i$-measurable have no behavioral meaning, but such a formal restriction is inconvenient when studying the relationship between players' information and the game. Theorem 2.1 shows that when a player faces the expected value of his payoff function conditional on his information and other players' strategies, any strategy is payoff equivalent to the projection of that strategy on his information. Therefore the player may be assumed to choose a $\mathcal{G}_i$-measurable strategy without making any such formal restriction.

<u>Theorem 2.1</u>: Let $u \in L^1(C(A_i))$, $s \in L^\infty(M(A_i))$, and $\mathcal{G}_i \in \mathcal{F}^*$. Then the conditional expectations $E[u|\mathcal{G}_i] \in L^1(C(A_i))$ and $E[s|\mathcal{G}_i] \in L^\infty(M(A_i))$ exist and uniquely satisfy $\int_G E[s|\mathcal{G}_i]d\mu = \int_G s d\mu$ and $\int_G E[u|\mathcal{G}_i]d\mu = \int_G u d\mu$ for all $G \in \mathcal{G}_i$, where all integrals are Bochner integrals [Diestel and Uhl (1977, pp. 44-45)]. In addition, $\langle\langle E[u|\mathcal{G}_i],s\rangle\rangle = \langle\langle u,E[s|\mathcal{G}_i]\rangle\rangle$.

<u>Proof:</u> The existence of vector-valued conditional expectation follows from Theorem 4 of Diestel and Uhl (1977, p. 123), which also implies that $E[u|\mathcal{G}_i] \in L^1(C(A_i))$. Since [using Proposition V-2-5 of Neveu (1975)]
$\sup_{B\subset A_i} |E[s|\mathcal{G}_i](\omega)(B)| \leq E[\sup_{B\subset A}|s(\cdot,B)|\,|\mathcal{G}_i](\omega) \leq E[\|s\|_\infty|\mathcal{G}_i](\omega) = \|s\|_\infty$ a.e.,
it follows that $\|E[s|\mathcal{G}_i]\|_\infty \leq \|s\|_\infty$, so the former lies in $L^\infty(M(A_i))$.

To show $\langle\langle E[u|\mathcal{G}_i],s\rangle\rangle = \langle\langle u,E[s|\mathcal{G}_i]\rangle\rangle$, let $\{u^k\}$ be a sequence of simple functions increasing to $E[u|\mathcal{G}_i]$, with $u^k = \sum_{\ell=1}^{L_k} f_\ell^k \cdot I_{G_\ell^k}$, $f_\ell^k \in C(A_i)$ for each k and ℓ, and $G_\ell^k \in \mathcal{G}_i$, where I_F is the indicator function of F. Then using Proposition V-2-5 of Neveu (1975) again, $\langle u^k,E[s|\mathcal{G}_i]\rangle = \sum_{\ell=1}^{L_k} I_{G_\ell^k}\langle f_\ell^k,E[s|\mathcal{G}_i]\rangle$ $= \sum_{\ell=1}^{L_k} I_{G_\ell^k} E[\langle f_\ell^k,s\rangle|\mathcal{G}_i] = E[\langle u^k,s\rangle|\mathcal{G}_i]$ since $G_\ell^k \in \mathcal{G}_i$. This implies $\langle\langle u^k,E[s|\mathcal{G}_i]\rangle\rangle = \langle\langle u^k,s\rangle\rangle$. Then $|\langle\langle E[u|\mathcal{G}_i],E[s|\mathcal{G}_i]\rangle\rangle - \langle\langle E[u|\mathcal{G}_i],s\rangle\rangle|$

$\leq |\langle\langle E[u|\mathcal{G}_i]-u^k,E[s|\mathcal{G}_i]\rangle\rangle| + |\langle\langle E[u|\mathcal{G}_i]-u^k,s\rangle\rangle|$ which goes to 0 by dominated convergence. An identical argument shows that $\langle\!\langle E[u|\mathcal{G}_i], E[s|\mathcal{G}_i]\rangle\!\rangle = \langle\!\langle u, E[s|\mathcal{G}_i]\rangle\!\rangle$, completing the proof. Q.E.D.

Corollary 2.2: For every $s \in S_i$, $E[s|\mathcal{G}_i] \in S_i$. Furthermore, if $\{s^k\} \subset S_i$ is a sequence converging weak* to s, then $\{E[s^k|\mathcal{G}_i]\}$ converges weak* to $E[s|\mathcal{G}_i]$ for every $\mathcal{G}_i \in \mathcal{F}^*$.

Proof: The first part follows from the definition of conditional expectation. Let $\{s^k\}$ be as in the statement of the result. Then by Theorem 2.1, for every $u \in L^1(C(A_i))$, $\langle\!\langle u, E[s^k|\mathcal{G}_i]\rangle\!\rangle = \langle\!\langle E[u|\mathcal{G}_i], s^k\rangle\!\rangle$, which then converges to $\langle\!\langle E[u|\mathcal{G}_i], s\rangle\!\rangle = \langle\!\langle u, E[s|\mathcal{G}_i]\rangle\!\rangle$, completing the proof. Q.E.D.

Let $S = \prod_{i\in I} S_i$. The payoff function of player i is $v_i \in L^1(C(A))$, or alternatively, a function $v_i\colon \Omega \times A \to \mathbb{R}$ which is Borel-measurable such that $v_i(\omega,\cdot)$ is continuous for a.e. ω and $\int_\Omega \sup_{a\in A}|v_i(\omega,a)|\mu(d\omega)$ is finite. The payoff function can be defined in terms of strategies. Let $S_{-i} = \prod_{j\neq i} S_j$ and $A_{-i} = \prod_{j\neq i} A_j$ with generic elements s_{-i} and a_{-i} respectively. Define $\pi_i\colon S \to \mathbb{R}$ to be the induced payoff function

$$\pi_i(s) = \int_\Omega \int_{A_i} E[\int_{A_{-i}} v_i(\cdot,a_i,a_{-i})s_{-i}(\cdot,da_{-i})|\mathcal{G}_i](\omega)\cdot s_i(\omega,da_i)\mu(d\omega)$$

$$= \int_\Omega \int_{A_i} \{\int_{A_{-i}} v_i(\omega,a_i,a_{-i})s_{-i}(\omega,da_{-i})\}E[s_i|\mathcal{G}_i](\omega,da_i)\mu(d\omega).$$

It is easy to show that π_i is separately continuous in each player's strategy s_j. However, π_i need not be jointly continuous in all players' strategies. Joint discontinuity can occur even when each player's action space is finite and payoff functions are independent of the state of nature, as demonstrated by Example 2 of Milgrom and Weber (1981). Another example is given here for later reference.

Example 2.3: Suppose there are two players, with $A_1 = A_2 = \{1,2\}$. Consider a pure coordination game in which each player's payoff is one if $a_1 = a_2$ and zero otherwise. Note that payoffs do not depend on the state of nature. Let Ω be the unit square with its Borel sets and generic element (ω_1,ω_2). Define μ_1 to be Lebesgue measure on Ω, μ_2 Lebesgue measure on the diagonal of Ω, and let the probability measure on Ω be $\mu = (\mu_1 + \mu_2)/2$. Let consumer 1 have the information field generated by the first coordinate of Ω, and consumer 2 the information field generated by the second coordinate of Ω. Let $s_1^n \in S_1$ be defined by $s_1^n(\omega_1)$ the point mass on 1 for $\omega_1 \cdot 2^n$ odd, and the point mass on 0 otherwise. Let $s_2^n \in S_2$ be defined by $s_2^n(\omega_2)$ the point mass on 0 for $\omega_2 \cdot 2^n$ odd and the point mass on 1 otherwise. Then coordination is perfectly achieved on the diagonal, and randomly otherwise, so $\pi_i(s_1^n,s_2^n) = 3/4$ for each n and i. Then s_1^n and s_2^n both converge to the strategy s which assigns probability 1/2 to both actions, regardless of the state of nature. Since $\pi_i(s,s) = 1/2$, π_i is not continuous.

Note that the above example amounts to a discontinuity of the degree of coordination of strategies. A sequence of pairs of strategies can be perfectly coordinated, while the joint limit is completely uncoordinated. To assure continuity, the ability of players to perfectly coordinate strategies must be limited. The assumptions used by Radner and Rosenthal (1982) and Milgrom and Weber (1981) both exclude possibilities for perfect coordination. Milgrom and Weber (1981, Assumption R2) required that the joint probability distribution over the product space of all player types be absolutely continuous with respect to the product of the marginal distributions. Radner and Rosenthal (1982, Theorems 2 and 3) imposed more severe assumptions. They required that the set of players' information fields $\{\mathcal{G}_j | j=1,\ldots,I\}$ be independent, and that each player's information about other players' payoffs is a finite partition.

The next theorem extends the results of Radner and Rosenthal with conditions similar to those used by Milgrom and Weber. By defining information separately from the payoffs of players, these results make clear the extent to which players' information, as opposed to their payoff functions, must be restricted. The requirement that $v_i \in L^1(C(A))$ for each i is identical to the assumption of absolutely continuous payoffs used by Milgrom and Weber (1981, Assumption R1). Condition (AC) below is analogous to their assumption of absolutely continuous information (R2). For each player i, define the probability measure μ_i to be μ restricted to $\mathcal{G}_i \vee \sigma(v_i)$, and for every other player j, define μ_j to be μ restricted to $\mathcal{G}_j$. Then, loosely speaking, μ is absolutely continuous with respect to the "product measure" $\mu_1 \times \mu_2 \times \ldots \times \mu_I$. This statement is not quite correct since there is no product space of player types. The absence of independently defined type spaces requires a more complicated proof than the one used by Milgrom and Weber.

<u>Theorem 2.4</u>: π_i is continuous on S whenever $v_i \in L^1(C(A))$ and for every $\varepsilon > 0$ there exists $\delta > 0$ such that

$$\text{(AC)} \quad \text{For } F = \bigcup_{k=1}^{K} \{[\bigcap_{j \neq i} G_j^k] \cap (G_i^k \cap V_i^k)\} \text{ with } G_j^k \in \mathcal{G}_j \text{ for each}$$

$j \in I$ and $V_i^k \in \sigma(v_i)$, $\sum_{k=1}^{K} [\prod_{j \neq i} \mu(G_j^k) \cdot \mu(G_i^k \cap V_i^k)] < \delta$ implies $\mu(F) < \varepsilon$.

<u>Proof</u>: The first step is to expand the state space and the information

fields on a set of measure zero such that events from different information fields are never disjoint. Let $\Omega_0 = \prod_{j\in I} \Omega$ with the product σ-field, and $\mu(\Omega_0) = 0$. Define the state space $\Omega' = \Omega \cup \Omega_0$ and define μ on it accordingly. For each $\mathcal{G}_j$, $j \neq i$, define a sub-σ-field $\mathcal{H}_j$ on Ω' as follows. For each $G_j \in \mathcal{G}_j$ let $H_j = G_j \cup (\Omega\times\ldots\times\Omega\times G_j\times\Omega\times\ldots\times\Omega)$ with the latter G_j in the j^{th} place and $(\Omega\times\ldots\Omega\times G_j\times\Omega\times\ldots\times\Omega) \subset \Omega_0$. Then define $\mathcal{H}_j$ to be the sub-σ-field generated by all events of that form. Define $\mathcal{H}_i$ similarly using $\mathcal{G}_i \vee \sigma(v_i)$. All changes from the original state space are on sets with μ-probability zero, so the expected payoff functions are unchanged on Ω'.

The second step is to define a new probability measure ν on Ω' with respect to which the information fields $\{\mathcal{H}_j | j \in I\}$ are independent. For each j let $\{\mathcal{H}_j^n\}_n$ be a sequence of measurable finite partitions on Ω' which increase to $\mathcal{H}_j$, with the restriction that for each $H_j^n \in \mathcal{H}_j^n$, $j \in I$, it is the case that $\bigcap_{j\in I} H_j^n \neq \emptyset$. Define $\mathcal{H}^n = \bigvee_{j\in I} \mathcal{H}_j^n$. On $\mathcal{H}^n$ define the measure ν by $\nu(\bigcap_{j\in I} H_j^n) = \prod_{j\in I} \mu(H_j^n)$. To show ν is consistently defined for all n, let $H_j^n \in \mathcal{H}_j^n$ for each j, and write $H_j^n = \bigcup_{k_j} H_{j,k_j}^{n+1}$ where $H_{j,k_j}^{n+1} \in \mathcal{H}_j^{n+1}$. Then using ν as defined on $\mathcal{H}^{n+1}$, $\nu(\bigcap_{j\in I} H_j^n) = \sum_{k_1,\ldots,k_I} \nu(\bigcap_{j\in I} H_{j,k_j}^{n+1}) = \sum_{k_1,\ldots,k_I} \prod_{j\in I} \mu(H_{j,k_j}^{n+1})$ $= \prod_{j\in I} \mu(H_j^n)$, so the definition of ν is consistent for each n. Then ν is uniquely defined, countably additive on the field $\bigcup_n \mathcal{H}^n$, and for each $H_j \in \bigcup_n \mathcal{H}_j^n$, $\nu(H_j) = \mu(H_j)$. By Cotter (1984, Lemma 24, p. 33), $\mathcal{H}^n$ is increasing to $\mathcal{H} = \bigvee_{j\in I} \mathcal{H}_j$. Therefore by Diestel and Uhl (1977, Theorem III.5.8), ν has a unique extension to $\mathcal{H}$. Let $H_j \in \mathcal{H}_j$ for $j \in I$. By Chung (1974, Theorem 8.1.1), given ε there exists n and $H_j^n \in \mathcal{H}_j^n$ such that $\nu(H_j^n \,\Delta\, H_j) < \varepsilon$ for each j. Since $\nu(\bigcap_{j\in I} H_j^n) = \prod_{j\in I} \nu(H_j^n)$ for each n, it follows by continuity of the intersection operation [Dunford and Schwartz (1957, Theorem A, p. 168)] that $\nu(\bigcap_{j\in I} H_j) = \prod_{j\in I} \nu(H_j)$, so $\{\mathcal{H}_j | j \in I\}$ are independent with respect to ν. For the remainder of this proof, take Ω' to be the state space with the σ-field $\mathcal{H}$, restricting both μ and ν to $\mathcal{H}$.

The third step is to show using (AC) that μ is absolutely continuous with respect to ν. Choose $\varepsilon > 0$, and let $\delta > 0$ be as in (AC). Let $F \in \mathcal{H}$ with $\nu(F) < \delta$. By Chung (1974, Theorem 8.1.1) there exists n and $F^n \in \mathcal{H}^n$ such that $\nu(F^n) < \delta$ and $\mu(F \Delta F^n) < \varepsilon$. Then $\mu(F^n) < \varepsilon$, so $\mu(F) < 2\varepsilon$. Therefore μ is absolutely continuous with respect to ν. Note that some of the "dummy states" added to the original state space may have positive ν-probability.

The fourth step is to show that with respect to ν, $\{\mathcal{H}_j | j \neq i\}$ and $\sigma(v_i)$ are independent relative to $\mathcal{G}_i$. Denote the conditional expected utility operators with respect to μ and ν by $E_\mu[\cdot|\mathcal{G}_i]$ and $E_\nu[\cdot|\mathcal{G}_i]$ respectively, and similarly for conditional probabilities. Let $G_j \in \mathcal{H}_j$ for each j, and $V_i \in \sigma(v_i)$. Then for each $G_i \in \mathcal{G}_i$,

$$\int_{G_i} I_{[\bigcap_{j\neq i} G_j]\cap V_i}\, d\nu = \nu[\{\bigcap_{j\neq i} G_j\} \cap (G_i \cap V_i)] = \prod_{j\neq i} \nu(G_j)\nu(G_i \cap V_i)$$

$$= \int_{G_i} P_\nu[V_i|\mathcal{G}_i]\cdot\{\prod_{j\neq i} \nu(G_j)\}\, d\nu = \int_{G_i} P_\nu[V_i|\mathcal{G}_i]\cdot\{\prod_{j\neq i} P_\nu[G_j|\mathcal{G}_i]\}\, d\nu$$

so $P_\nu[F_i|\mathcal{G}_i]\cdot\{\prod_{j\neq i} P_\nu[G_j|\mathcal{G}_i]\} = P_\nu[\{\bigcap_{j\neq i} G_j\} \cap V_i|\mathcal{G}_i]$. Therefore $\{\mathcal{H}_j\}_{j\neq i}$ and $\sigma(v_i)$ are independent relative to $\mathcal{G}_i$, with respect to ν.

The fifth step is to rewrite player i's expected payoff function in terms of conditional expectations on ν. Let f be the Radon-Nikodym derivative $d\mu/d\nu$. Choose $\varepsilon > 0$. There exists $W \in \mathcal{H}$ such that for some M, $f(\omega) < M$ for all $\omega \in W$, $\{v_i(\omega,\cdot)|\omega \in W\}$ is uniformly equicontinuous with $\sup_{\substack{a\in A\\ \omega\in W}} |v_i(\omega,a)| < M$ and $\int_{W^c} \sup_{a\in A}\{v(\omega,a)\}\mu(d\omega) = \int_{W^c} f(\omega)\sup_{a\in A}\{v(\omega,a)\}\nu(d\omega) < \varepsilon$.

Then on f restricted to W, there exists some function $\hat{f} = \sum_{k=1}^{K} b_k I_{W_k}$ such that for each k, $W_k = \{\bigcap_{j\in I} G_j^k\} \cap V_i^k$ with $G_j^k \in \mathcal{H}_j$ for each j and $V_i^k \in \sigma(v_i)$ with $|f(\omega) - \hat{f}(\omega)| < \varepsilon/M$ for each $\omega \in W$, reducing W on a set of measure zero if necessary. Let s_i, $s_i' \in S_i$, s_{-i}, $s_{-i}' \in S_{-i}$, and define $w \in L^1(\Omega', \mathcal{H}, \mu; \mathbb{R})$ by

$$w(\omega) = \int_{A_i}\int_{A_{-i}} v_i(\omega,a_i,a_{-i})\{s_{-i}(\omega,da_{-i}) - s_{-i}'(\omega,da_{-i})\}E_\mu[s_i - s_i' | \quad _i](\omega,da_i)$$

Then $|\pi_i(s) - \pi_i(s')| = |\int_{\Omega'} f(\omega)w(\omega)\nu(d\omega)| \leqslant 2\varepsilon + \sum_{k=1}^{K} b_k |\int_{\Omega'} I_{W_k} w(\omega)\nu(d\omega)|$.

Consider the k^{th} term in the above sum. By Theorem 2.1, and the fact that μ and ν coincide on $\mathcal{H}_i$, it equals

$$\int_{G_i^k}\int_{A_i} E_\nu[\int_{A_{-i}} I_{W_k} v_i(\bullet,a_i,a_{-i})\{s_{-i}-s_{-i}'\}(\bullet,da_{-i})|\mathcal{G}_i]\bullet\{s_i-s_i'\}(\omega,da_i)\mu(d\omega) \tag{1}$$

$$= \int_{G_i^k}\int_{A_i}\int_{A_{-i}} E_\nu[I_{V_i^k} v_i | \quad _i](\omega,a)\{\prod_{j\neq i} E_\nu[I_{G_j^k}(s_j-s_j')|\mathcal{G}_i](\omega,da_j)\}\{s_i-s_i'\}(\omega,da_i)d\mu \tag{1'}$$

by the argument of step 4. Since, with respect to ν, $\mathcal{H}_j$ is independent of $\mathcal{H}_i$ for each $j \neq i$ and s_j, s_j' are $\mathcal{H}_j$-measurable, it follows that $E_\nu[I_{G_j^k}(s_j-s_j')|\mathcal{G}_i] = E_\nu[I_{G_j^k}(s_j-s_j')] = \bar{s}_j^k - \bar{s}_j'^k$ (say), so (1) equals

$$\int_{G_k^i}\int_{A_i}\int_{A_{-i}} E_\nu[I_{V_i^k} v_i |\mathcal{G}_i](\omega,a)\bullet\{\prod_{j\neq i} [\bar{s}_j^k - \bar{s}_j'^k](da_j)\}\bullet[s_i - s_i'](\omega,da_i)\mu(d\omega) \tag{1''}$$

Then by Parthasarathy (1967, Theorem II.6.8) and Cotter (1986, Lemma 4.2), given $\varepsilon > 0$ there exists for each j a finite collection $\{\hat{s}_j^{\ell_j} | \ell_j = 1,2,\ldots,L\} \subset M(A_j)$ such that for every $\bar{s}_j \in M(A_j)$, there exists ℓ_j such that for all $a_{-j} \in A_{-j}$ and $\omega \in W$,

$$|\int_{A_j} E_\nu[I_{V_i^k} v_i|\mathcal{G}_i](\omega,a_j,a_{-j}) \{\prod_{j\neq i} [\bar{s}_j(da_j) - \hat{s}_j^{\ell_j}(da_j)]\}| < \varepsilon/(Kb_k). \tag{2}$$

For each $\bar{\ell} = (\ell_j)_{j \neq i}$ define $g^{\bar{\ell}} \in L^1\big(C(A_i)\big)$ by $g^{\bar{\ell}}(\omega, a_i) = \int_{A_{-i}} E_\nu[I_{V_i^k} v_i | \mathcal{G}_i](\omega, a_{-i}, a_i) \cdot \{\prod_{j \neq i} \hat{s}_j^{\ell_j}(da_j)\}$. By Corollary 2.2, s'_{-i} and s_{-i} can be made close enough for each j so that for some $\bar{\ell}$, and all a_i and $\omega \in W$,

$$\left| \int_{A_{-i}} E_\nu[I_{V_i^k} v_i | \mathcal{G}_i](\omega, a_i, a_{-i}) \cdot \{\prod_{j \neq i} [\hat{s}_j^{\ell_j}(da_{-j}) - \bar{s}_j(da_j)]\} \right| < \varepsilon/(Kb_k). \qquad (3)$$

In addition, s_i and s'_i can be chosen close enough so that for all k and $\bar{\ell}$,

$$\left| \int_{G_i^k} \int_{A_i} g^{\bar{\ell}}(\omega, a_i)\{s_i(\omega, da_i) - s'_i(\omega, da_i)\}\mu(d\omega) \right| < \varepsilon/(Kb_k). \qquad (4)$$

Then $|\pi_i(s) - \pi_i(s')| < 5\varepsilon$ by use of the triangle inequality and equations (1'', 3, 4), completing the proof. Q.E.D.

<u>Corollary 2.5</u>: A Nash equilibrium exists whenever the conditions of Theorem 2.4 hold.

<u>Proof</u>: Use the standard existence proof of Glicksberg (1952). Q.E.D.

Note that the above results hold even if the players have asymmetric beliefs about Ω, where $(\Omega, \mathcal{F})$ is common knowledge and each player i has a unique probability measure μ_i. In order for the definition of "almost everywhere" to be consistent across players, the events of probability zero must be common knowledge, so $\mu_1, \ldots, \mu_n$ must be mutually absolutely continuous.

The necessity of condition (AC) is an open question. If (AC) does not hold, then the first two steps in the above proof remain valid, but μ is no longer absolutely continuous with respect to ν. It seems likely that a construction based on Example 2.3 can be made to show discontinuity.

3. Convergence of player characteristics

One advantage to modelling information as an explicit variable of the game as in Section 2 is in studying properties of the relationship between players' information and their resulting behavior. A basic property of this relationship is its continuity with respect to some appropriate topology of information. A topology which satisfies this property provides a meaningful description of the game-theoretic similarity of game characteristics. In addition, topologizing the space of games would provide rigorous means for studying perturbations of game characteristics, and allow the study of stability of equilibria with respect to perturbations. Several topologies on the space of information $\mathcal{F}^*$ have been defined, including the Boylan metric [Boylan (1971)] and the pointwise convergence metric [Cotter (1986)]. A sequence of information fields $\{\mathcal{G}_n\}$ converges in the pointwise convergence metric if and only if for every $f \in L^1(\mathbb{R})$, $\lim_{n\to\infty} \| E[f|\mathcal{G}_n] - E[f|\mathcal{G}] \| = 0$, while it converges with respect to the Boylan metric if and only if the latter convergence is uniform over all f which are uniformly bounded a.e. Define $\mathcal{V} = \{v_i \in L^1(C(A)) \mid$ the set $\{v_i(\omega,\cdot)\}$ is uniformly equicontinuous and bounded over a.e. $\omega\}$. This condition was used in analyzing both metrics of information by Cotter (1986) and Allen (1983) respectively.

Give $\mathcal{F}^*$ the Boylan metric unless stated otherwise, and define $\pi_i : S \times (\mathcal{F}^*)^I \times \mathcal{V} \to \mathbb{R}$ as in Section 2.

Theorem 3.1: The payoff function π_i is continuous over all $(\{\mathcal{G}_j | j \in I\}, v_i)$ which satisfy (AC).

Proof: For each $j \in I$ let $\{\mathcal{G}_j^n\} \subset \mathcal{F}^*$ be a sequence converging to $\mathcal{G}_j$ and $\{v_i^n\} \subset \mathcal{V}$ be a sequence converging to $v_i \in \mathcal{V}$. It suffices to show that

$\{\pi_i(s,\{\mathcal{G}_j^n|j\in I\},v_i^n)\}$ converges to $\pi_i(s,\{\mathcal{G}_j|j\in I\},v_i)$ uniformly in $s\in S$. Choose $(s_i,s_{-i})\in S$, and define $g^n\in L^1(C(A_i))$ by $g_n(\omega) = \int_{A_{-i}} v_i(\omega,a_i,a_{-i})s_{-i}(\omega,da_{-i})$ and define $g\in L^1(C(A_i))$ similarly, so both depend on s_{-i}. Note that g_n converges to g. Let $\|\cdot\|$ be the norm of the space $L^1(C(A_i))$, so $|\pi_i(s_i,s_{-i},\mathcal{G}_i^n,v_i^n) - \pi_i(s_i,s_{-i},\mathcal{G}_i,v_i)|$ $\leq \|E[g^n|\mathcal{G}_i^n] - E[g|\mathcal{G}_i]\| \leq \|g_n - g\| + \|E[g|\mathcal{G}_i^n] - E[g|\mathcal{G}_i]\|$. Choose $\varepsilon > 0$, and let $\delta > 0$ be such that for $\rho(a_i,a_i') < \delta$, it follows that for a.e. ω and all a_{-i} that $|v_i(\omega,a_i,a_{-i}) - v_i(\omega,a_i',a_{-i})| < \varepsilon$. Then for all s_{-i}, $|g(\omega,a_i) - g(\omega,a_i')| < \varepsilon$. Cover A_i by balls of radius δ with centers $\{a_i^1,a_i^2,\ldots,a_i^K\}$. Then choosing a_i^k to be in the same ball as a_i,

$$\begin{aligned}\|E[g|\mathcal{G}_i^n] - E[g|\mathcal{G}_i]\| \leq &\int_\Omega \sup_{a_i\in A_i} |E[g(\cdot,a_i)|\mathcal{G}_i^n] - E[g(\cdot,a_i^k)|\mathcal{G}_i^n]|\,d\mu \\ &+ \int_\Omega \sup_k |E[g(\cdot,a_i^k)|\mathcal{G}_i^n] - E[g(\cdot,a_i^k)|\mathcal{G}_i]|\,d\mu \\ &+ \int_\Omega \sup_{a_i\in A_i} |E[g(\cdot,a_i^k)|\mathcal{G}] - E[g(\cdot,a_i)|\mathcal{G}]|\,d\mu\end{aligned}$$

The first and third terms are less than ε for all n, while the second term can be made less than ε for sufficiently large n for all s_{-i}. Q.E.D.

<u>Corollary 3.2</u>: The Nash equilibrium correspondence is upperhemicontinuous over all player characteristics satisfying (AC).

Based on the results in Cotter (1986), it is surpising that a similar result does not hold for the pointwise convergence metric.

Example 3.3: Consider the game with two players in Example 2.3. For each i and n let $\mathcal{Y}_i^n$ be the information field generated by the random variable which equals 0 if the integer part of $\omega_i \cdot 2^n$ is even, and 1 otherwise. By Cotter (1986, Example 3.4), $\{\mathcal{Y}_i^n\}$ converges pointwise to the trivial information field (note that it does not converge in the Boylan metric). Let s_i^n be the strategy defined in Example 2.3. Suppose player i has the information field $\mathcal{Y}_i^n$. Condition (AC) is satisfied for each n and for the limit. Then $\pi_i(s_1^n, s_2^n, \mathcal{Y}_i^n) = 3/4$ for each n, but $\pi_i(s_1, s_2, \mathcal{Y}_i) = 1/2$, so π_i is not continuous with respect to pointwise convergence of information.

This result is disappointing since some of the properties of the pointwise convergence metric would be very useful in studying games with asymmetric information. The most important ones are the separability of the space of information and the denseness of the set of finite partitions. Neither of these properties hold for the Boylan metric. In particular, it may be difficult to model players' beliefs about other players' information, since such beliefs are most easily modelled as probability distributions on a separable metric space. Nevertheless, these results identify the Boylan metric as the key for studying similarity of games. Boylan convergence of information is less cumbersome and more transparent than the convergence of game characteristics required by Milgrom and Weber for their model (1981, Theorem 2, conditions iii, iv). In particular, the latter convergence concept does not appear to be topological, and does not permit convergence of information and payoffs separately.

Note that for fixed characteristics of other players and fixed s_{-i}, $\pi_i(\cdot, s_{-i}, \cdot, \cdot)$ is continuous over all $s_i \in S_i$ and all v_i, $\{\mathcal{Y}_j | j \in I\}$ which satisfy (AC). This follows immediately from Cotter (1986, Theorem 4.3).

4. Approximation by pure strategies

The use of mixed strategies in game theory, including related concepts such as distributional and behavioral strategies, has been criticized as useless in practice because mixed strategies either are not observed or are behaviorally meaningless. In defense, Milgrom and Weber (1981) argued that the the set of pure strategies is large enough to include all observable behavior, but a larger set is needed to obtain the required compactness of the strategy space. To support that argument, their Theorem 4 states that if a player's information is nonatomic, then the set of pure strategies is a dense subset of all distributional strategies for that player. A similar result for behavioral strategies is proven below. Let $S_i^P = \{s_i^P \in S_i | s_i^P(\omega)$ is a point mass on A_i for a.e. $\omega\}$ be the set of pure strategies, in which every state of nature is mapped into a single action. Then if the player's information field is atomless, any strategy can be approximated to any degree by a pure strategy. The idea is that a player can recover "almost all" randomization opportunities by merely randomizing over observable states of nature.

<u>Theorem 4.1</u>: If $\mathcal{G}_i$ is nonatomic, then S_i^P (in fact, the set of all simple pure strategies) is dense in the set of $\mathcal{G}_i$-measurable strategies.

<u>Proof</u>: Let $u^1, u^2, \ldots, u^k \in L^1(C(A_i))$ be $\mathcal{G}_i$-measurable simple functions, and let $\{G_1, G_2, \ldots, G_L\}$ be a $\mathcal{G}_i$-measurable partition of Ω with respect to which each u^k is measurable, so $u^k = \sum_{\ell=1}^L f_\ell^k \cdot I_{G_\ell}$ with $f_\ell^k \in C(A_i)$ for each k and ℓ. Then for $s_i \in S_i$, $\langle\langle u^k, s_i \rangle\rangle$ equals

$$\sum_{\ell=1}^{L} \int_{G_\ell} \int_{A_i} f_\ell^k(a_i) s_i(\omega, da_i)\mu(d\omega) = \sum_{\ell=1}^{L} \int_{A_i} f_\ell^k(a_i)\nu_\ell(da_i)$$

where $\upsilon_\ell = \int_{G_\ell} s_i(\omega)\mu(d\omega)$ is a measure on A_i.

By Carathéodory's theorem [Royden (1967 p. 321)] there exists a measure algebra isomorphism Φ_ℓ of $(A_i, \mathcal{A}_i, \upsilon_\ell)$ into the unit interval with its Borel sets and Lebesgue measure. Let μ_i be the restriction of μ to $\mathcal{G}_i$ and $\mathcal{G}_{i\ell}$ be the restriction of $\mathcal{G}_i$ to G_ℓ. Since μ_i is nonatomic, there exists, with an abuse of notation, a measure algebra isomorphism Ψ_ℓ of $(G_\ell, \mathcal{G}_{i\ell}, \mu_i)$ onto the unit interval with Borel sets and Lebesgue measure. Then $\Psi^{-1} o \Phi_\ell$ is a measure algebra isomorphism of $(A_i, \mathcal{A}_i, \upsilon_\ell)$ into $(G_\ell, \mathcal{G}_{i\ell}, \mu_i)$. By Theorem 15.11 of Royden (1967), there exists $G'_\ell \subset G_\ell$ with $\mu(G'_\ell) = \mu(G_\ell)$ and a measurable function $\hat{a}_\ell : G'_\ell \to A_i$ such that for every $B \in \mathcal{A}_i$, $\mu o \hat{a}_\ell^{-1}(B) = \upsilon_\ell(B)$. Therefore, for every k, $\int_{A_i} f^k_\ell(a)\upsilon_\ell(da) = \int_{G_\ell} f^k_\ell(\hat{a}(\omega))\mu(d\omega)$. Define $s^p_i \in S^p_i$ where, for $\omega \in G_\ell$, $s^p_i(\omega)$ is the point mass at $\hat{a}_j(\omega)$. Then $\langle\langle u^k, s_i \rangle\rangle = \langle\langle u^k, s^p_i \rangle\rangle$ for every k. Since the set of simple functions is strongly dense in $L^1(C(A_i))$, the proof is complete. Q.E.D.

Theorem 4.1 is useful only to the extent that nonatomic information is typical. The next two results give a partial answer to that question.

<u>Theorem 4.2</u>: If $(\Omega, \mathcal{F}, \mu)$ is nonatomic, then the set of nonatomic information fields is dense in $\mathcal{F}^*$ with respect to the pointwise convergence metric.

<u>Proof</u>: Since the set of finite partitions of the state space is dense in $\mathcal{F}^*$ [Cotter (1986, Proposition 2.3)], it suffices to show that given $D_1, \ldots, D_K \in \mathcal{F}$ and any finite partition $\mathcal{G} \in \mathcal{F}^*$, there exists a nonatomic information field $\mathcal{G}' \in \mathcal{F}^*$ such that $P[D_k|\mathcal{G}] = P[D_k|\mathcal{G}']$ for each k, where $P[D|\mathcal{G}] = E[I_D|\mathcal{G}]$. Without loss of generality assume that $\{D_1, \ldots, D_K\}$ is a disjoint partition of Ω.

Define the increasing sequence $\{\mathcal{G}_n\}$ inductively as follows. Let $\mathcal{G}_1 = \mathcal{G}$ and write $\mathcal{G}_n = \sigma\{G_1,\ldots,G_J\}$, the latter forming a disjoint partition of Ω. For each i and j let $B_{kj} \subset D_k \cap G_j$ such that $\mu(B_{kj}) = \mu(D_k \cap G_j)/2$ [Chung (1974, Exercise 23, p. 31)]. Then let $G_{j,1} = \bigcup_{k=1}^{K} B_{kj}$ and $G_{j,2} = G_j \cap G_{j,1}^c$ so $\mu(G_{j,1} \cap D_k) = \mu(B_{kj})$. Define $\quad_{n+1} = \sigma\{G_{1,1},G_{1,2},\ldots,G_{J,1},G_{J,2}\}$, so $P[D_k|\mathcal{G}_{n+1}] = P[D_k|\mathcal{G}_n]$ for each k. Then $\{\mathcal{G}_n\}$ is an increasing sequence of finite partitions, and therefore converges pointwise by Proposition 2.2 of Cotter (1986) to $\mathcal{G}'$ (say). Then $P[D_k|\mathcal{G}] = P[D_k|\mathcal{G}']$ for each k, so it remains only to be shown that $\mathcal{G}'$ is nonatomic. Let $G \in \mathcal{G}'$. By Theorem 8.1.1 of Chung (1974), given $\varepsilon > 0$ there exists n and $G_\varepsilon \in \mathcal{G}_n$ such that $\mu(G \,\Delta\, G_\varepsilon) < \varepsilon$, where $G \,\Delta\, G_\varepsilon = (G \cap G_\varepsilon^c) \cup (G^c \cap G_\varepsilon)$. Let $\varepsilon < \mu(G)$ so that $\mu(G \cap G_\varepsilon) > 0$. Using the construction of the previous paragraph, let G' be either $G \cap G_\varepsilon \cap (\bigcup_{j=1}^{J} G_{j,1})$ or $G \cap G_\varepsilon \cap (\bigcup_{j=1}^{J} G_{j,2})$; one of these allows $\mu(G') > 0$. Since $\bigcup_{j=1}^{J} G_{j,1} \in \mathcal{G}_n$, it follows that $\mu(G') < \mu(G)$. Therefore $\mathcal{G}'$ is nonatomic, completing the proof. Q.E.D.

<u>Theorem 4.3</u>: If $\mathcal{F}^*$ is given the Boylan metric d, then the set of nonatomic information fields is never dense in $\mathcal{F}^*$.

<u>Proof</u>: By Fact 9.3 of Allen (1983), for all $\mathcal{G} \in \mathcal{F}^*$, $d(\mathcal{G},\{\Omega,\emptyset\}) \geqslant \sup\{\|\mu(F) - P[F|\mathcal{G}]\| \,|\, A \in \mathcal{F}\} \geqslant \sup\{\|\mu(G) - I_G\| \,|\, G \in \mathcal{G}\} = \sup\{\mu(G)[1 - \mu(G)] \,|\, G \in \mathcal{G}\}$. If $\mathcal{G}$ is nonatomic, the latter equals 1/4. Q.E.D.

Another argument mitigating the use of mixed strategies is that in many games, an exact Nash equilibrium in pure strategies exists. Milgrom and Weber (1981, Theorem 5) and Radner and Rosenthal (1982, Theorems 1 and 2) identify similar but restrictive conditions for which a Nash equilibrium in pure strategies exists. An analogous result for this model is proven below.

Following the terminology of Radner and Rosenthal, a purification of the Nash equilibrium vector $s \in S$ is $s^P \in S^P = \prod_{i \in I} S_i^P$ such that for each i,

$$\pi_i(s, v_i, \mathcal{J}_i) = \pi_i(s^P, v_i, \mathcal{J}_i) \tag{5}$$

$$E[s_i] = E[s_i^P] \tag{6}$$

$$s^P \text{ is a Nash equilibrium.} \tag{7}$$

Radner and Rosenthal (1982, Theorem 1) give conditions for which every Nash equilibrium has a purification. One requirement is that each A_i is finite. That proof applies directly to this model with the same conditions. If A_i is infinite, a different condition on v_i may be used instead, which in turn requires a separate proof. Note that the weaker condition (AC) of Theorem 2.7 is not sufficient. A counterexample is provided by Radner and Rosenthal (1982, Example 1). The necessity of condition (c), even when (a) and (b) hold, is demonstrated by their Example 3.

<u>Theorem 4.5</u>: Suppose, for each i,

(a) $\{\mathcal{J}_j | j \neq i\}$ are independent of $\mathcal{J}_i \vee \sigma(v_i)$

(b) $\mathcal{J}_i$ is nonatomic,

(c) either v_i is simple or A_i is finite.

Then every Nash equilibrium has a purification.

<u>Proof</u>: The case of A_i finite follows from the proof of Radner and Rosenthal (1982, Theorem 1) without modification. Suppose conditions (a) - (c) hold, and v_i is simple. Write $v_i = \sum_{\ell=1}^{L} I_{F_\ell} u_{i\ell}$ and $\bar{u}_{i\ell}(a_i) =$

$\int_{A_{-i}} u_{i\ell}(a_i, a_{-i}) \cdot \bar{s}_{-i}(da_{-i})$, so

$$\int_{A_{-i}} E[v_i|\mathcal{G}_i](\omega,a_i,a_{-i})\cdot\bar{s}_{-i}(da_{-i}) = \sum_{\ell=1}^{L} P[F_\ell|\mathcal{G}_i](\omega)\bar{u}_{i\ell}(a_i). \qquad (8)$$

Define, for each ℓ, $\nu_\ell = \int_\Omega P[F_\ell|\mathcal{G}_i](\omega)s_i(\omega)\mu(d\omega)$ to be a measure on A_i. Then (8) equals $\sum_{\ell=1}^{L} \int_{A_i} \bar{u}_{i\ell}(a_i)\nu_\ell(da_i)$. Construct $\hat{a}_\ell$ and s_i^p as in the proof of Theorem 4.1, and let $s^p = (s_1^p, s_2^p,\ldots,s_N^p)$. To verify (6), let $B \subset A_i$. Then $E[s_i^p](B) = = \sum_{\ell=1}^{L} \int_{F_\ell} s_i^p(\omega,B)\mu(d\omega) = \sum_{\ell=1}^{L} \int_{F_\ell} I_{\hat{a}^{-1}(B)}(\omega)\mu(d\omega) = \sum_{\ell=1}^{L} \mu(\hat{a}^{-1}(B))$
$= \sum_{\ell=1}^{L} \nu_\ell(B) = \sum_{\ell=1}^{L} \int_\Omega P[F_\ell|\mathcal{G}_i](\omega)\cdot s_i(\omega,B)\mu(d\omega) = \int_\Omega s_i(\omega,B)\mu(d\omega) = E[s_i](B)$, proving (6). To show (5), note that each player's payoff depends only on the expected value of all other strategies, so replacing them with their purifications does not affect the expected payoffs by condition (6), which also demonstrates that s^p is a Nash equilibrium, verifying (7). Q.E.D.

References

B. Allen, Neighboring information and distributions of agents' characteristics under uncertainty, J. Math. Econ. 12 (1983), 63-101.

E. Boylan, Equiconvergence of martingales, Ann. Math. Stat. 42 (1971), 552-559.

K. Cotter, Agents' information as a market variable, unpublished Ph.D. dissertation, University of Minnesota, 1984.

K. Cotter, Similarity of information and behavior with a pointwise convergence topology, J. Math. Econ. 15 (1986), 25-38.

K. Cotter, Convergence of information, random variables, and noise, J. Math. Econ., forthcoming.

K. Chung, "A Course in Probability Theory," 2nd. ed., Academic Press, New York, 1974.

J. Diestel and J. Uhl, "Vector Measures," American Mathematical Society, Princeton, N.J., 1977.

N. Dunford and J. Schwartz, "Linear Operators, Part I," Interscience, New York, 1957.

D. Fudenberg, D. Kreps, and D. Levine, On the robustness of equilibrium refinements, IMSSS Discussion Paper #489, Stanford University, 1986.

I. Glicksberg, A further generalization of the Kakutani fixed-point theorem with applications to Nash equilibrium points, Proc. Amer. Math. Soc. 3 (1952), 170-174.

E. Kohlberg and J. Mertens, On the strategic stability of equilibria, Econometrica 54 (1986), 1003-1038.

D. Kreps and R. Wilson, Sequential equilibria, Econometrica 50 (1982), 863-894.

F. Mertens and S. Zamir, Formulation of Bayesian analysis for games with incomplete information, Inter. J. of Game Theory 14 (1985), 1-29.

P. Milgrom and R. Weber, Distributional strategies for games with incomplete information, CMSEMS Discussion Paper #428, Northwestern University, 1981.

J. Neveu, "Discrete Parameter Martingales," (T. Speed, Tr.), American Elsevier, New York, 1975.

K. Parthasarathy, "Probability Measures on Metric Spaces," Academic Press, New York, 1967.

R. Radner and R. Rosenthal, Private information and pure strategy equilibria, Math. Oper. Res. 7 (1982), 401-409.

H. Royden, "Real Analysis," 2nd. ed., Macmillan, New York, 1967.

W. Rudin, "Functional Analysis," McGraw-Hill, New York, 1973.