No. 255 (85-4)

Information Patterns and Nash Equilibria in Extensive Games II

by

Pradeep Dubey
and
Mamoru Kaneko

March 1985

Information Patterns and Nash Equilibria

in Extensive Games II*

by

Pradeep Dubey

Department of Economics, Indian Statistical Institute

New Delhi, India

Department of Economics, University of Illinois at Urbana-Champaign

Champaign, IL 61820, USA

and

Mamoru Kaneko

Institute of Socio-Economic Planning, University of Tsukuba

Ibaraki-ken 305, Japan

ABSTRACT: In Part I of this paper we introduced extensive games with a non-atomic continuum of players. It was shown that the Nash plays (outcomes) are invariant of the information patterns on the game, provided that no player's unilateral change in moves can be observed by others. This led to an enormous reduction in the Nash plays of the these games, as exemplified in the anti-folk theorem. Our concern in this sequel is to develop a finite version of these results.

Key words: ε-Nash equilibrium; replication of an extensive game; primitive Nash play.

---

*This work was supported by O.N.R. Grant No. N00014-77-C-0518 issued under Contract Authority NR047-006.

1. INTRODUCTION[1]

In Part I of this paper we introduced extensive games with a non-atomic continuum of players. It was found (Propositions 4.1 and 4.2 of Part I) that the Nash plays are invariant of the information patterns on the game, provided that no player's unilateral changes in moves can be observed by others. This led to an enormous reduction in the Nash plays of these games, as exemplified in the anti-folk theorem (Section 6 of Part I). Our concern in this sequel is to develop a finite version of these results. The problem is that information is discontinuous with moves. Indeed the example in Section 2 of this paper openly flouts the non-atomic result. It consists of two sequences of games $n\Gamma \underset{n}{\to} \Gamma$, $n\Gamma_c \underset{n}{\to} \Gamma_c$. Here $n\Gamma_c$, Γ_c are obtained from $n\Gamma$, Γ by coarsening information sets; and the finite-player games $n\Gamma$, $n\Gamma_c$ "converge" with n to the non-atomic games Γ , Γ_c . But when the Nash plays of $n\Gamma$, $n\Gamma_c$ are computed they are found to diverge. Thus the non-atomic results of Part I which says in particular that the Nash plays of Γ and Γ_c coincide, is called into question.

The conundrum becomes clear if we notice that, in point of fact, the games $n\Gamma$ do <u>not</u> converge to Γ . This is because of the information conditions. In Γ no single player can affect the integral (average) of bids by the very assumption that he is a point in a non-atomic continuum. In sharp contrast, in every n --no matter how large n is-- any unilateral deviation of bids by a player does change the average and can be precisely observed by others. The large n case therefore does <u>not</u> reflect the non-atomic assumption as far as information is concerned. Intuitively one would think that that assumption should translate into a finite game to mean:

---

[1]The notation and terminology of Part I is carried over into this paper.

very small changes in the average cannot be observed by anyone. This, in turn, may be imagined to stem from an intrinsic grid on the scales of measurement; or it may simply be thought of as a behavioral postulate of inertia in players' reactions (there is a positive lower bound on the change that must occur for anyone to react). Once we make such a postulate, everything falls into place. But, in doing so, we are forced to break away from the standard notion of a Nash Equilibrium (N.E.) and to alter it in order to take this lower bound into account. So we introduce "ε-N.E.'s" in Section 3 (where $\varepsilon \equiv$ the lower bound). Then a finite version of the non-atomic result of Part I becomes available (Proposition 1). It can be transformed into a formal convergence theorem (Propositon 2). The problem here is to construct a natural model of a sequence of extensive games that is "convergent". One such is suggested is Section 4, and it may be of some interest beyond the use to which it is put in this paper.

With the introduction of this bound, however, a cat is now let out of the bag. Consider the standard "convex case", i.e., one in which the set of moves is convex at any position and each player's payoff is concave in his own moves. It includes a large class of "normal form" games, and a fortiori the games obtained by repeating them, provided suitable payoffs are chosen (e.g. discounted sums, lim inf of the average). Then a curious result (Proposition 3) occurs: the set of plays achieved at ε-N.E.'s is independent of the information conditions and of $\varepsilon > 0$. Moreover, this is true regardless of the number of players. The upshot (Section 5.2) is an "anti-folk theorem" for finite-player games: for any $\varepsilon > 0$, the ε-N.E. plays of the repeated game Γ^{∞} coincide with the N.E. plays of the minimal information variant of Γ^{∞}. Thus a certain delicateness in the folk theorem is brought to light. It dramatically breaks down with the slightest coarsening of informaton, i.e., making $\varepsilon > 0$, no matter how small.

The bound on a player's capacity of observation so far pertained only to others' moves. It is impelling to extend it to all observation. In Section 6 we impose bounds also on what a player can see of his own strategies and payoffs. Propositions 1 and 2 can essentially be retrieved with the obvious modifications. This time the anti-folk theorem breaks down. But it breaks in a manner which is continuous with these two additional bounds (Proposition 4). If their magnitude is small, we still get a diluted version of it. And, in any case, the folk theorem is far from getting reinstated.

2. AN EXAMPLE

To highlight the problem of the discontinuity of information it might be best to examine an example in detail. There are two types of traders who exchange two commodities through a trading-post. The initial endowments are (1,0) and (0,1) for the two types, and all of them have the same utility function $u(x_1, x_2) = \sqrt{x_1 x_2}$. A move of a trader is to bid a quantity of his commodity for sale in the trading-post. Those of the first type move simultaneously at the start of the game. The second type of traders can find out the average quantity bid of commodity 1 before they make their moves, also simultaneously, i.e., without knowledge of the move of anyone of their own type. The total amount received of commodity 1(2) is then disbursed to traders of type 2(1) in proportion to their bids. If any one type bids a total of zero then no trade takes place and the bids are returned. We examine a replication sequence $\{n\Gamma\}_{n=1}^{\infty}$ where the game $n\Gamma$ has n traders of each type. The limit of this sequence appears to be the game Γ in which there is a continuum of each type. But this is not quite true because information is discontinuous in going from $n\Gamma$ to Γ, with the result that a direct asymptotic version of Propositions 5 and 5* of Part I cannot be obtained.

Formally the player-set in $n\Gamma$ is $nN = nL \cup nM$, where $nL = \{1, \ldots, n\}$ and $nM = \{n+1, \ldots, 2n\}$. The set of positions is

$$X = \{x_0\} \cup [0,1]^{nL} \cup ([0,1]^{nL} \times [0,1]^{nM}) ;$$

and

$$\pi(x_0) = nL , \quad \pi(x) = nM \quad \text{for all} \quad x \in [0,1]^{nL} ;$$

$$\pi(x) = \phi \quad \text{for} \quad x \in [0,1]^{nL} \times [0,1]^{nM} ;$$

$$S_i^{x_0} = [0,1] \quad \text{for} \quad i \in nL ;$$

$$S_i^{x} = [0,1] \quad \text{for} \quad i \in nM \quad \text{and} \quad x \in [0,1]^{nL} ;$$

$$I_i = \{\{x_0\}\} \quad \text{for} \quad i \in nL ;$$

$$I_i = \{\{s^{x_0} \in [0,1]^{nL} : \frac{1}{n}\sum_{j \in nL} s_j^{x_0} = \alpha\} : \alpha \in [0,1]\} \quad \text{for} \quad i \in nM .$$

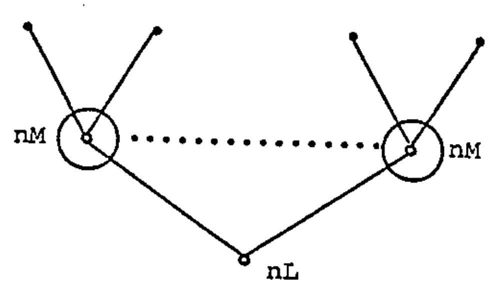

Figure 1. The game $n\Gamma$

For any $\sigma = (s^{x_0}, \{s^x\}_{x\in[0,1]} nL)$ in $\Sigma(n\Gamma)$, the final holding ψ_i of player i is determined by $\xi(\sigma)$ ($= (s^{x_0}, s^{x_1})$ with $x_1 = (s^{x_0})$) by the rule:

$$\psi_i(\sigma) = \begin{cases} \left(1 - s_i^{x_0}, \dfrac{s_i^{x_0}}{\sum_{\ell\in nL} s_\ell^{x_0}} \cdot \sum_{\ell\in nM} s_\ell^{x_1}\right) & \text{if } \sum_{\ell\in nL} s_\ell^{x_0} > 0 \text{ and } \sum_{\ell\in nM} s_\ell^{x_1} > 0 \\ (1,0) & \text{otherwise} \end{cases}$$

if $i \in nL$; and

$$\psi_i(\sigma) = \begin{cases} \left(\dfrac{s_i^{x_1}}{\sum_{\ell\in nM} s_\ell^{x_1}} \cdot \sum_{\ell\in nL} s_\ell^{x_0}, 1 - s_i^{x_1}\right) & \text{if } \sum_{\ell\in nL} s_\ell^{x_0} > 0 \text{ and } \sum_{\ell\in nM} s_\ell^{x_1} > 0 : \\ (0,1) & \text{otherwise} \end{cases}$$

if $i \in nM$. The payoffs to the players are of course the utilities of their final holdings.

Each $n\Gamma$ has some trivial <u>inactive</u> N.E.'s at which any one type bids nothing and the other type bids arbitrarily. All these lead to the same final allocation (1,0), (0,1) as the initial one.

<u>All</u> other Nash plays of $n\Gamma$ are given by:

s^{x_0} is arbitrary;

$$s_i^x = \begin{cases} \frac{n-1}{2n-1} & \text{if} \quad x \in I_i(x_1) \ , \quad x_1 = (s^{x_0}) \\ 0 & \text{otherwise} \end{cases}$$

if $i \in nM$. Thus the set of Nash allocations (i.e. those produced at some N.E.) of $n\Gamma$ is:

$$\left\{ (x_1, \ldots, x_n, x_{n+1}, \ldots, x_{2n}) : x_i = \left(1 - \alpha_i, \frac{\alpha_i}{\sum_{\ell \in nL} \alpha_\ell} \cdot \frac{n(n-1)}{2n-1}\right) \right.$$

$$\text{for } i \in nL \ , \quad x_i = \left(\frac{1}{n} \sum_{i \in nL} \alpha_i, \ 1 - \frac{n-1}{2n-1}\right) \text{ for } i \in nM \text{ , where}$$

$$\left. \text{the } \alpha_i \in [0,1] \text{ are arbitrary and } \sum_{i \in nL} \alpha_i > 0 \right\} \cup \{\text{initial endowments}\} .$$

Now consider the game $n\Gamma_c$ obtained by assuming that players of type 2 observe nothing, i.e., $I_i = [0,1]^{nL}$ for $i \in nM$.

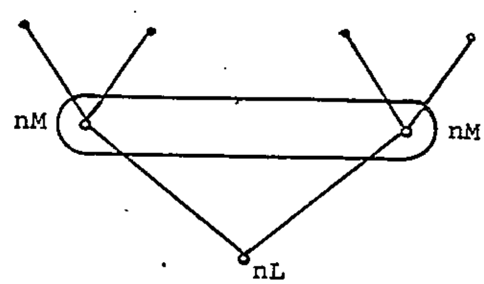

Figure 2. The game $n\Gamma_c$

Besides the inactive N.E.'s, $n\Gamma_c$ has a unique active N.E.:

$$
\begin{aligned}
s_i^{x_0} &= \frac{n-1}{2n-1} \quad \text{if} \quad i \in nL ; \\
s_i^{x} &= \frac{n-1}{2n-1} \quad \text{for all} \quad x \in I_i \ , \quad \text{if} \quad i \in nM .
\end{aligned}
\tag{1}
$$

Thus the Nash allocations of $n\Gamma_c$ are either given by the initial allocation or by:

$$
x_i = \begin{cases} \left(1 - \frac{n-1}{2n-1}, \frac{n-1}{2n-1}\right) & \text{if} \quad i \in nL \\ \left(\frac{n-1}{2n-1}, 1 - \frac{n-1}{2n-1}\right) & \text{if} \quad i \in nM \end{cases}
$$

Thus it is clear that the N.E.'s of $n\Gamma$ and $n\Gamma_c$ do not converge as $n \to \infty$.

On the other hand, if we look directly at the limit game[1] Γ (in which there is a continuum of each type) then, by Proposition 5 (or 5*) of Part I, Γ and Γ_c have the same N.E.'s (in terms of the plays, i.e., allocations, produced). Indeed one can easily compute that these consist of the inactive N.E.'s and the unique active N.E. in which all traders bid 1/2 to wind up with the final holding (1/2, 1/2) (which is also the "competitive allocation" of the underlying economy.

[1]which can be formalized in the obvious way.

3. BOUNDED CAPACITY OF OBSERVATION

(with respect to others' moves)

The example of Section 2 shows us that we need to do some "doctoring" to the Nash plays of $n\Gamma$ in order to develop an asymptotic version of the non-atomic result. In Γ an arbitrary unilateral change in strategy by a player does not affect the integral of moves anywhere and therefore it cannot be observed by others. We translate this condition into a finite setting quite directly to mean: very small changes made by a player (say of "size" less than ε for some positive ε), cannot be seen by others. Then we are lead to the notion of an "ε-N.E.", which is simply an N.E. modified by this constraint on observation.

To formulate this precisely let Γ stand for an arbitrary extensive game as in Section 2 of Part I, except that now (for simplicity) there are no chance moves. Recall that Γ is given by:

$$\Gamma = (N,\ X,\ \pi,\ \{S^x\}_{x\in X},\ \Phi,\ \{h_i\}_{i\in N},\ \{I_i\}_{i\in N}) .$$

Throughout we assume that N is finite. There is a metric space D with a distance denoted d , and each S_i^x ($x \in X$ and $i \in \pi(x)$) is a nonempty subset of D . Without loss of generality, we may assume that there are no ending positions in Γ . This is so because we can attach a fictitious path of infinite length to each of them. Furthermore, for the purposes of this paper, it will be convenient to take these fictitious paths to be identical, with the same move at each step, whose distance from any "real" move in Γ is ∞ .

Define the metric δ on X by the rule:

if $x = (s^{x_0}, s^{x_1}, \ldots, s^{x_m})$ and $y = (r^{y_0}, r^{y_1}, \ldots, r^{y_k})$ (with $x_0 \equiv y_0$), then:

$$\delta(x,y) = \begin{cases} \max\limits_{0\le\ell\le m} \sum\limits_{i\in\pi(x_\ell)} d(s_i^{x_\ell}, r_i^{y_\ell}) & \text{if } m = k \text{ and } \pi(x_\ell) = \pi(y_\ell) \\ & \text{for } \ell = 0, 1, \ldots, m\,; \\ \infty & \text{otherwise.} \end{cases}$$

In other words, if the paths leading to x and y (from the start x_0 of the game) involve identical player-sets at each stage, then the distance between x and y is the maximum over all relevant pairs of moves. For all other cases it is infinite. Let $p = (s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots)$ and $\tilde{p} = (r^{y_0}, r^{y_1}, \ldots, r^{y_m}, \ldots)$ be two plays. Then

$$\delta^*(p,\tilde{p}) = \sup_{\ell=0,1\ldots} \delta(x_\ell, y_\ell)$$

constitutes a metric on $\Lambda(\Gamma\_)$ (since we assume that there is no chance move, all outcome are, in fact, plays). That is, the distance between two plays is simply the supremum of the distances between corresponding positions that lie on the plays. In particular, note that if two plays originally had different lengths then, even though our convention of attaching fictitious paths makes them both identically long, the distance between them is nevertheless infinite.

Let $\sigma \in \Sigma(\Gamma)$ and let $\xi(\sigma) = (s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots)$. For $\varepsilon > 0$ and $\tau_i \in \Sigma_i(\Gamma)$, $\xi^\varepsilon(\sigma|\tau_i) = (r^{y_0}, r^{y_1}, \ldots, r^{y_m}, \ldots)$ is, intuitively speaking, the play that results from player i's deviation to τ_i, <u>after</u> accounting for the fact that the others cannot observe any change (in positions) of distance less than ε . Precisely, $(r^{y_0}, \ldots, r^{y_m}, \ldots)$ is defined as follows (with $x_0 \equiv y_0$):

if $i \in \pi(y_\ell)$ then $r_i^{y_\ell} = \tau_i(y_\ell)$, for all $\ell = 0, 1, \ldots$

for all $j \neq i$, $j \in \pi(y_k)$ and any $k = 0, 1, \ldots$

$$r_j^{y_k} = \begin{cases} s_j^{x_k} & \text{if (a) } \delta(x^k, y^k) < \varepsilon \text{ and} \\ & \text{(b) for } 0 \le \ell \le k \text{ , } s_j^{x_\ell} \in S_j^{y_\ell} \text{ whenever } j \in \pi(x_\ell) \\ \sigma_j(y_k) & \text{otherwise.} \end{cases}$$

Consider any $\varepsilon > 0$. A $\sigma \in \Sigma(\Gamma)$ will be called an <u>ε-modified Nash Equilibrium</u> of Γ , and denoted <u>ε-N.E.</u>, if

$$h_i(\xi^{\varepsilon}(\sigma|\tau_i)) \le h_i(\xi(\sigma))$$

for all $\tau_i \in \Sigma_i(\Gamma)$ and all $i \in N$. This is exactly like an ordinary N.E. except now each player assumes that if he makes very small unilateral deviations (within ε), then the others will not be able to observe it. Therefore they will continue to make the same moves as before, <u>even though</u> they had planned to do otherwise in their strategies. Once the deviation exceeds ε then, of course, others moves are according to their strategies.

The nothing of an ε-N.E. enables us to smooth out the discontinuity of information in going from finite-player games to games with a continuum. The key idea is that if there is a large number of small players who observe some kind of a "macro" aggregate of everyone's moves, then a unilateral deviation will affect the aggregate only slightly. In a continuum this effect is nil. The nothing of an ε-N.E. therefore becomes relevant to the continuum phenomenon. (See Section 4 where we build an asymptotic version of the continuum result of Part I of this paper using ε-N.E.)

For the first proposition we need to define the notion of an <u>ε-inner play</u> of Γ for $\varepsilon > 0$. This is a play $p = (s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots)$ in $\Lambda(\Gamma_-)$ with the property that for all $\ell = 0, 1, \ldots$

$$\left.\begin{Bmatrix} \delta(x_\ell, y) < \varepsilon \\ \\ i \in \pi(x_\ell) \end{Bmatrix}\right. \Rightarrow s_i^{x_\ell} \in S_i^y \quad .$$

In other words, the move employed by any player in the play p remains feasible for small changes in positions around p. If σ is an ε-N.E., and if $\xi(\sigma)$ is an ε-inner play, then we say that σ is a <u>(ε)-N.E.</u> If a play p is produced by a (ε)-N.E. σ, i.e., $p = \xi(\sigma)$, then p is called a <u>(ε)-N.E. play</u>.

For any game Γ let $\Delta(\Gamma)$ be the class of all games obtained from Γ by varying only the information patterns of Γ (in accordance with the conditions in Section 2 of Part I).

<u>Proposition 1.</u> Assume (a) σ is an (ε)-N.E. of Γ; and (b) $\delta^*(\xi^\varepsilon(\sigma|\tau_i), \xi(\sigma)) < \varepsilon$ for all $\tau_i \in \Sigma_i(\Gamma)$ and $i \in N$. Then $\xi(\sigma)$ is a (ε)-N.E. play of $\tilde{\Gamma}$ for every $\tilde{\Gamma} \in \Delta(\Gamma)$.

In the example of Section 2, take $\Gamma = n\Gamma$ and $\tilde{\Gamma} = n\Gamma_c$. If $n > 1/\varepsilon$, Proposition 1 applies, and we see that any (ε)-N.E. play of $n\Gamma$ is also an (ε)-N.E. play of $n\Gamma_c$. But $n\Gamma_c$ has a unique (ε)-N.E. play, and this is simply the N.E. play of $n\Gamma_c$ given in (1). Thus the (ε)-N.E. play of $n\Gamma$ is also unique and coincides with the N.E. play of $n\Gamma_c$. From the proposition it is needed that $n > 1/\varepsilon$; but, in fact, (ε)-N.E. plays of $n\Gamma$ are invariant of ε in this example, on account of its convex structure, as will be clarified in Section 5.

Proof of Proposition 1. Since N is finite, there exists a $\tilde{\sigma} \in \Sigma(\tilde{\Gamma})$ such that $\xi(\tilde{\sigma}) = \xi(\sigma)$. It suffices to show that $\tilde{\sigma}$ is an (ε)-N.E. of $\tilde{\Gamma}$. If this were not so, then for some $\tilde{\tau}_i \in \Sigma_i(\tilde{\Gamma})$ we would have $h_i(\xi^{\varepsilon}(\tilde{\sigma}|\tilde{\tau}_i)) > h_i(\xi(\tilde{\sigma}))$. Let $\{I_i^t\}_{t=1}^K$, where K could be ∞ , be the information sets of i through which the play $\xi^{\varepsilon}(\tilde{\sigma}|\tilde{\tau}_i)$ passes. Define $\tau_i \in \Sigma_i(\Gamma)$ by

$$\tau_i(x) = \begin{cases} \tilde{\tau}_i(x) & \text{if } x \in I_i^t \text{ for some } t \\ \text{arbitrary otherwise.} \end{cases}$$

We will prove that $\xi^{\varepsilon}(\tilde{\sigma}|\tilde{\tau}_i) = \xi^{\varepsilon}(\sigma|\tau_i)$. Put $\xi^{\varepsilon}(\tilde{\sigma}|\tilde{\tau}_i) = (s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots)$, $\xi^{\varepsilon}(\sigma|\tau_i) = (r^{y_0}, r^{y_1}, \ldots, r^{y_m}, \ldots)$ and $\xi(\sigma) = (w^{z_0}, w^{z_1}, \ldots, w^{z_m}, \ldots)$ where $x_0 \equiv y_0 \equiv z_0$. Make the inductive assumption that $x_\ell = y_\ell$ for $\ell \le k$. (For $k = 1$ this is obvious.) By (b) $\delta(z_k, y_k) < \varepsilon$, hence also $\delta(z_k, x_k) < \varepsilon$. Also by (a) and (b), $w_j^{z_\ell} \in S_j^{y_\ell}$ $(= S_j^{x_\ell})$ for all $\ell \le k$ and $j \in \pi(z_\ell)$ $(= \pi(y_\ell) = \pi(x_\ell))$. By the definition of $\xi^{\varepsilon}(\ldots)$ we have $s_j^{x_k} = w_j^{z_k}$ and $r_j^{y_k} = w_j^{z_k}$, therefore $s_j^{x_k} = r_j^{y_k}$ for $j \in \pi(z_k)\setminus\{i\}$. If $i \in \pi(z_k) = \pi(y_k) = \pi(x_k)$ then $\tilde{\tau}_i(x_k) = \tau_i(x_k) = \tau_i(y_k)$. Thus $s^{x_k} = r^{y_k}$ which proves that $x_{k+1} = y_{k+1}$.

But then we have $h_i(\xi^{\varepsilon}(\sigma|\tau_i)) = h_i(\xi^{\varepsilon}(\tilde{\sigma}|\tilde{\tau}_i)) > h_i(\xi(\tilde{\sigma})) = h_i(\xi(\sigma))$ contradicting that σ is a (ε)-N.E. of Γ . □

It is often the case that an ε-N.E. is also automatically a (ε)-N.E. In particular this is so for layered games $\Gamma = (N, X, \ldots)$, i.e. those in which, for any fixed ℓ, $\pi(x)$ is constant and so is S_i^x $(i \in \pi(x))$ for $x \in X_\ell$ = the ℓ^{th}-layer of $\Gamma = \{x \in X : x$ is at a distance[1] ℓ from the start x_0 of $\Gamma\}$. Clearly in a layered game every play is ε-inner for any ε, thus its ε-N.E.'s and (ε)-N.E.'s coincide.

Denote the (ε)-N.E. plays of Γ by $\eta_\varepsilon(\Gamma)$. Then as in Proposition 2.2 in Part I, we see that

$$\text{if } \Gamma \prec \tilde{\Gamma} \text{, then } \eta_\varepsilon(\Gamma) \subset \eta_\varepsilon(\tilde{\Gamma}) . \tag{4}$$

Thus, under refinement of information, the (ε)-N.E. plays grow in a nested manner. Proposition 1 gives conditions under which the growth is stopped.

4. CONVERGENCE

Proposition 1 already has within it the germ of a convergence result. We already saw how it worked on the example of Section 2. The main problem is to precisely formulate the notion of a "convergent sequence" of extensive games. For simplicity, we will focus on the case of replication, i.e., when the r^{th} game has r clones of each of the players of the first, original game. Also players will be assumed to observe the average of others' moves. (See, however, end of Section.)

Fix $\hat{\Gamma} = (N, \hat{X}, \pi, \{\hat{S}^x\}_{x \in X}, \Phi, \{h_i\}_{i \in N}, \{I_i\}_{i \in N})$. A sequence[2] $p_i = (r_i^{x_0}, r_i^{x_1}, \ldots, r_i^{x_m}, \ldots)$, where $i \in N$, will be called compatible with a play $p = (s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots) \in \Lambda(\hat{\Gamma}_-)$ if

---

[1] If $y = (s^{x_0}, \ldots s^{x_\ell})$ it is said to be at a distance $\ell+1$ from x_0.

[2] Let us adopt the convention: if $i \notin \pi(x)$ then (a) $S_i^x = \{\phi\}$ and, accordingly, (b) $s_i^x = \phi$ for $s^x \in S^x$.

$$\text{for all } \ell \geq 0 \text{ , } r_i^{x_\ell} \in \hat{S}_i^{x_\ell} \text{ .} \tag{5}$$

I.e., p_i is any list of possible moves of player i along the positions in p. Define $P_i = \{(p, p_i) : p \in \Lambda(\hat{\Gamma}\_)$ and p_i is compatible with $p\}$. In order to define replication we need to assume that there are functions $h_i : P_i \to R$ which satisfy

$$\overline{h}_i((s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots), (s_i^{x_0}, s_i^{x_1}, \ldots, s_i^{x_m}, \ldots)) \tag{6}$$
$$= h_i((s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots)).$$

We further assume

each $\hat{S}_i^x$ $(i \in \pi(x))$ is a convex set of some fixed normed vector space V (over the reals). (7)

In the sequel, $\hat{\Gamma}$ will serve as the ambient game within which replication is defined

First suppose

a subset $S_i^x \subset \hat{S}_i^x$ is exogenously specified for each $x \in \hat{X}$ and $i \in \pi(x)$. (S_i^x could well be finite.) (8)

The game Γ (to be replicated) is <u>embedded</u> in $\hat{\Gamma}$, i.e., it is obtained from $\hat{\Gamma}$ by reducing the moves to S_i^x :

$$\Gamma = (N, X, \pi, \{S^x\}_{x \in X}, \Phi, \{h_i\}_{i \in N}, \{I_i\}_{i \in N}) \text{ .}$$

Here $X \subset \hat{X}$; and π , Φ , h_i and I_i are to be viewed as defined by the appropriate restriction.

The r^{th} replica game $r\Gamma$ (with $1\Gamma \equiv \Gamma$) will be denoted:

$$r\Gamma = (rN,\ rX,\ r\pi,\ \{rS^x\}_{x\epsilon rX},\ r\Phi,\ \{rh_i\}_{i\epsilon rN},\{rI_i\}_{i\epsilon rN}.$$

There is a map $r\beta$ that relates the r^{th} replica game $r\Gamma$ to the ambient game $\hat{\Gamma}$. At each position x in $r\Gamma$ there are r clones of the player-set at the position $r\beta(x)$ in $\hat{\Gamma}$. Averaging moves made by r clones of player i in $r\Gamma$ is, by our convexity assumption (7), a move by i in $\hat{\Gamma}$. Thus we can work with $\hat{\Gamma}$ throughout the process of replication. The map $r\beta$ from rX to $\hat{X}$ satisfies (for $x, y \in rX$):

$$rN = \bigcup_{i\epsilon N}[i]_r \quad \text{where} \quad [i]_r = \{(i,1),\ \ldots,\ (i,r)\}\ ; \tag{9}$$

$$r\pi(x) = \bigcup_{i\epsilon\pi\cdot r\beta(x)} [i]_r\ ; \tag{10}$$

$$rS^x_{(i,t)} = S_i^{r\beta(x)} \quad \text{for } i \in \pi\cdot r\beta(x) \quad \text{and} \quad 1 \le t \le r\ ; \tag{11}$$

$$x \succ_r y \longleftrightarrow r\beta(x) \succ r\beta(y)\ , \quad \text{and} \quad r\beta(rx_0) = x_0\ ; \tag{12}$$

$$\text{if} \quad s^x \in S^x \quad \text{then} \quad r\beta(x,\ s^x) = (\hat{x},\ w^{\hat{x}})\ , \quad \text{where} \quad r\beta(x) = \hat{x}$$

$$\text{and} \quad w^{\hat{x}} = \left(\frac{1}{r}\sum_{t=1}^{r} s^x_{(i,t)}\right)_{i\epsilon\pi\cdot r\beta(x)} . \tag{13}$$

Here $\succ$, $\succ_r$ are the partial orders (see Section 2 of Part I) induced by Φ , $r\Phi$ on $\hat{X}$, rX . Also, by (7),

$$\left(\frac{1}{r}\sum_{t=1}^{r} s^{x}_{(i,t)}\right)_{i\in\pi\cdot r\beta(x)} \in \prod_{i\in\pi\cdot r\beta(x)} S^{r\beta(x)} .$$

This defines the extensive form $r\Gamma_-$, i.e., all of $r\Gamma$ minus payoffs and information sets.

That $r\Gamma_-$ exists and is unique can be verified by building up $r\Gamma_-$ layer by layer, starting from the 0^{th} layer, and using (10)-(13). At the start rx_0 of $r\Gamma$ the player-set is $r\pi(rx_0) = \bigcup_{i\in\pi(x_0)} \{(i,1),\ldots,(i,r)\}$ by (12) and (10). Each $rS^{rx_0}_{(i,t)} = S^{x_0}_i$ by (11). For any choice of moves by the player in $r\pi(rx_0)$, the average $\frac{1}{r}(\ldots)$ yields a point in $\prod_{i\in\pi(x_0)} S^{x_0}_i$ (since each $S^{x_0}_i$ is convex), i.e., it yields a position $\hat{x}$ in $\hat{X}_1$. Now $\pi(\hat{x})$ cloned r times furnishes (by (10) and (13)) the set of players who move at $(s^{x_0}_{(i,t)})_{(i,t)\in r\pi(rx_0)} \in r\beta^{-1}(\hat{x}) \in rX_1$ in $r\Gamma_-$, etc.

It remains to define the payoffs and information sets of $r\Gamma$. Note that $r\beta$ yields a map $r\beta^* : \Lambda(r\Gamma_-) \to \Lambda(\hat{\Gamma}_-)$ by the rule:

$$r\beta^*(r^{y_0}, r^{y_1}, \ldots, r^{y_m}, \ldots) = (s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots)$$

$$\text{if } \ r\beta(r^{y_0}, \ldots, r^{y_\ell}) = (s^{x_0}, \ldots, s^{x_\ell}) \ \text{ for } \ \ell = 0, 1, \ldots.$$

(Here, of course, $y_0 \equiv rx_0$ by assumption.)

Then if $p = (s^{rx_0}, s^{x_1}, \ldots, s^{x_m}, \ldots) \in \Lambda(r\Gamma_-)$, the payoff function $rh_{(i,t)}$ is defined by

$$rh_{(i,t)}(p) = \bar{h}_i(r\beta^*(p), (s^{x_0}_{(i,t)}, s^{x_1}_{(i,t)}, \ldots, s^{x_m}_{(i,t)}, \ldots)). \tag{14}$$

I.e., a player's payoff depends only on his personal moves and the average of others' moves. Finally, define player (i,t)'s $((i,t) \in [i]_r))$ information pattern by

$$rI_{(i,t)} = \{\{x \in rX : r\beta(x) \in I_i(r\beta(y))\} : y \in rX\} . \qquad (15)$$

This says that, as we replicate, a player's information pattern is defined via Γ on the averages of moves in $r\Gamma$. This completes the definition of $r\Gamma$.

Difine δ on $\hat{X}$ as in Section 2, using the norm $\| \ \|$ on V. This in turn gives us metrics $r\delta$ on rX $(r = 1, 2, \ldots)$ by

$$r\delta(x,y) = \delta(r\beta(x), r\beta(y)) .$$

For the convergence result we will require uniform boundedness of moves

$$\sup\{\|\alpha\| : \alpha \in \hat{S}_i^x \text{ for some } x \in \hat{X} \text{ and } i \in \pi(x)\} \leq B < \infty . \qquad (16)$$

Proposition 2. Suppose $\hat{\Gamma}$ satisfies (16) and Γ is embedded in $\hat{\Gamma}$. Then if $r > B/\varepsilon$ the plays produced at (ε)-N.E.'s of $\tilde{\Gamma}$ are invariant of $\tilde{\Gamma} \in \Delta(r\Gamma)$.

Proof. Straightforward, using Proposition 1.

Neither the finite-type assumption on players, nor the fact that they can observe only averages, seem crucial to a convergence result. In general one would start with an ambient non-atomic game in which every finite game of the sequence is embedded. The observation could be on the distribution of strategies. Proposition 1 lends itself to this general set up also. We have not, however, worked out a detailed picture.

5. THE CONVEX CASE

5.1. Primitive Nash Plays

When the game has a convex structure (see (21) and (22) below), there is an enormous simplification in the set of (ε)-N.E. They become independent of and also of the information patterns. Thus we can compute (ε)-N.E. of any game by going to one of its maximally coarse information-variants. The upshot of this is an anti-folk theorem (Section 5.2) for finite, convex games.

Although our setting will remain general, the definitions and results of this section are perhaps best illustrated for layered games, to which we will refer throughout.

Put $\mathcal{C}_i(\Gamma) = \{q : q$ is compatible with some $\tilde{p} \in \Lambda(\Gamma_-)\}$. Consider $q = (r_i^{y_0}, r_i^{y_1}, \ldots, r_i^{y_m}, \ldots)$ in $\mathcal{C}_i(\Gamma)$ and $p = (s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots)$ in $\Lambda(\Gamma_-)$. We will say that the pair (q,p) is play-producing if (intuitively) the moves of i given by q, along with those of others as given by p, together suffice to produce a play in $\Lambda(\Gamma_-)$. Precisely, we require:

There exist a play $(w^{z_0}, w^{z_1}, \ldots, w^{z_m}. \ldots)$ in $\mathcal{P}(\Gamma)$ such that (for $\ell = 0, 1, \ldots$)

$$\text{(a)} \quad j \in \pi(z_\ell) \setminus \{i\} \rightarrow w_j^{z_\ell} = s_j^{x_\ell} \tag{17}$$

$$\text{(b)} \quad i \in \pi(z_\ell) \rightarrow w_i^{z_\ell} = r_i^{y_\ell} .$$

Such a play, if it exists, is necessarily unique and will be denoted $p \oplus q$. Define, for $p = (s^{x_0}, s^{x_1}, \ldots, s^{x_\ell}, \ldots) \in \Lambda(\Gamma_-)$

$$\mathcal{C}_i(p) = \{q \in \mathcal{C}_i(\Gamma) : (q,p) \text{ is play-producing}\} . \tag{18}$$

Note

$\mathcal{C}_i(p)$ is not empty because it always contains

$$(s_i^{x_0}, s_i^{x_1}, \ldots, s_i^{x_m}, \ldots) ; \tag{19}$$

if Γ is a layered game, then for any $\tilde{p} \in \Lambda(\Gamma_-)$

$$\mathcal{C}_i(\tilde{p}) = \mathcal{C}_i(\Gamma) = \mathop{\times}_{\ell=0}^{\infty} S_i^{x_\ell} . \tag{20}$$

Recall here our convention: if $i \notin \pi(x_\ell)$ then $S_i^{x_\ell} = \{\phi\}$; and $s_i^{x_\ell} = \phi$ for $s^{x_\ell} \in S^{x_\ell}$. Identify ϕ with the origin of some vector space $W \neq V$. Then each $S_i^{x_\ell} \subset V$ or W , hence $\mathcal{C}_i(p)$ may be viewed as a subset of the vector space $\mathop{\times}_{\ell=0}^{\infty} U^\ell$, where $U^\ell = V$ if $i \in \pi(x_\ell)$ and $U^\ell = W$ if $i \notin \pi(x_\ell)$. We assume, for $p \in \Lambda(\Gamma_-)$ and $i \in N$:

$$\mathcal{C}_i(p) \text{ is convex} \tag{21}$$

$$\hat{h}_i : \mathcal{C}_i(p) \to R \text{ , given by } \hat{h}_i(q) = h_i(p \oplus q) \text{ , is concave.} \tag{22}$$

Note that (21) holds if Γ is a layered game and each S_i^x (for $x \in X$, $i \in \pi(x)$) is a convex subset of V .

A play $p \in P(\Gamma)$ will be called a <u>primitive Nash play</u> if, for any $i \in N$,

$$h_i(p) \geq h_i(p \oplus q) \quad \text{for all} \quad q \in C_i(p) \ . \tag{23}$$

This is a Nash-like condition, except that player i is restricted to only those moves that are play-producing with respect to others' fixed moves on p. For instance the N.E. plays of $n\Gamma_c$ in the example of Section 2 are clearly primitive Nash plays.

Note that primitive Nash plays are, by definition, invariant of the information condition. Thus Proposition 3 below shows that, if $\varepsilon > 0$, ε-N.E. plays of convex games do not proliferate with refinement of information <u>even in the finite-player case</u>.

<u>Proposition 3</u>. Assume that (16), (21) and (22) hold for Γ. For $\varepsilon > 0$ let p be an ε-N.E. play for Γ i.e., a play produced by an ε-N.E. Then p is a primitive Nash play.

<u>Proof</u>. If p is not a primitive Nash play, then

$$h_i(p \oplus q) > h_i(p)$$

for some $q \in C_i(p)$. By (21), (22) and (19)[1],

$$h_i(\lambda(p \oplus q) + (1-\lambda)p) \geq \lambda h_i(p \oplus q) + (1-\lambda)h_i(p) > h_i(p) \ , \tag{24}$$

for any $0 < \lambda < 1$. Choose $\tilde{\lambda}$ sufficiently close to 0 to ensure that:

$$\delta^*(\tilde{\lambda}(p \oplus q) + (1-\tilde{\lambda})p, \ p) < \varepsilon \ . \tag{25}$$

---

[1] Condition (19) implies $p = p \oplus p_i$ where $p_i = (u_i^{x_0}, u_i^{x_1}, \ldots, u_i^{x_m}, \ldots)$ for $p = (u^{x_0}, u^{x_1}, \ldots, u^{x_m}, \ldots)$.

This can be done because of (16). Suppose $q = (r_i^{y_0}, \ldots, r_i^{y_m}, \ldots)$ and $p = (s^{x_0}, \ldots, s^{x_m}, \ldots)$. Then, by (21), $\tilde{q} = (\tilde{\lambda} r_i^{y_0} + (1-\tilde{\lambda}) s_i^{x_0}, \ldots, \tilde{\lambda} r_i^{y_m} + (1-\tilde{\lambda}) s_i^{x_m}, \ldots)$ is in $C_i(p)$; and clearly $(p \oplus \tilde{q}) = \tilde{\lambda}(p \oplus q) + (1-\tilde{\lambda})p$. Let $(p \oplus \tilde{q}) = (w^{z_0}, \ldots, w^{z_m}, \ldots)$. Construct a strategy τ_i of player i as follows:

$$\tau_i(x) = \begin{cases} \tilde{\lambda} r_i^{y_\ell} + (1-\tilde{\lambda}) s_i^{x_\ell} & \text{if } x = z_\ell \\ \text{arbitrary otherwise.} & \end{cases}$$

Denoting by $\sigma \in \Sigma(\Gamma)$ a strategy selection that produced $p = \xi(\sigma)$, we see that, by (25),

$$\xi^{\varepsilon}(\sigma | \tau_i) = \tilde{\lambda}(p \oplus q) + (1-\tilde{\lambda})p .$$

Then, by (24),

$$h_i(\xi^{\varepsilon}(\sigma | \tau_i)) > h_i(p)$$

contradicting that σ is an ε-N.E. of Γ.

Q.E.D.

<u>Corollary</u>. Let Γ be a layered game satisfying (16), (21) and (22). Denote the set of its primitive Nash plays of Γ by $N(\Gamma)$. Then the set of ε-N.E. plays of $\tilde{\Gamma}$ is invariant of $\tilde{\Gamma} \in \Delta(\Gamma)$ and of $\varepsilon > 0$, and coincides with $N(\Gamma)$, i.e.,

$$\eta_\varepsilon(\tilde{\Gamma}) = N(\Gamma)$$

for all $\tilde{\Gamma} \in \Delta(\Gamma)$, all $\varepsilon > 0$.

Proof. First observe that $N(\Gamma_1) = N(\Gamma_2)$ for $\Gamma_1, \Gamma_2 \in \Delta(\Gamma)$ because the definition of primitive Nash plays does not depend on the information pattern. By Proposition 3, $\eta^\varepsilon(\tilde{\Gamma}) \subset N(\tilde{\Gamma})$, thus

$$\eta_\varepsilon(\tilde{\Gamma}) \subset N(\Gamma) \quad \text{for} \quad \tilde{\Gamma} \in \Delta(\Gamma) .$$

Now take any $p = (s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots) \in N(\Gamma)$. Define strategies σ in $\tilde{\Gamma}$ by $\sigma_i(y) = s_i^{x_m}$ if y and x_m are in the same layer. It is easy to check that σ is an ε-N.E. of $\tilde{\Gamma}$ and yields the play p . □

This can be applied to the example of Section 2. We get $\eta_\varepsilon(n\Gamma) = N(\Gamma) = \eta_\varepsilon(n\Gamma_c)$ for all n and all $\varepsilon > 0$. Thus replication was not really necessary for an asymptotic version of the non-atomic result!

Since Γ is layered, there is a coarse game $\Gamma_{\bar{c}}$ in which the information sets of player i in $\Gamma_{\bar{c}}$ are the sets $\{x \in X^k : i \in \pi(x)\}$, $k = 1, \ldots,$ where X^k is the sets of positions whose distance from the start x_0 are k . We easily see :

$$\eta_\varepsilon(\Gamma_{\bar{c}}) = \eta_0(\Gamma_{\bar{c}}) \quad \text{for all} \quad \varepsilon \geq 0 .$$

Hence if Γ , besides being layered, satisfies (16), (21) and (22), then we get

$$\eta_\varepsilon(\Gamma) = \eta_0(\Gamma_{\bar{c}}) \quad \text{for all} \quad \varepsilon > 0 . \tag{26}$$

Moreover, by Proposition 2.2 of Part I of this paper, $\eta_0(\Gamma_{\bar{c}}) \subset \eta_0(\Gamma)$, so ε - N.E. plays of Γ are N.E. plays of Γ . More generally, when Γ is not layered, we must replace ε-N.E. by (ε)-N.E. for this to be true.

5.2. The Anti-Folk Theorem for Finite Games

Let Γ be a normal form game, i.e., $\pi(x_0) = N$ and $\pi(s^{x_0}) = \phi$ for all $s^{x_0} \in S^{x_0}$. Assume, for any $i \in N$ and any $s^{x_0} \in S^{x_0}$:

(a) $S_i^{x_0}$ is convex

(b) The map $f : S_i^{x_0} \to \mathbb{R}$ given by $f(t_i) = h_i(s^{x_0}|t_i)$ is concave.

Denote by Γ^∞ the infinite repetition of Γ . We assume that each player can observe at least the entire past history of his own moves and payoffs. The payoffs in Γ^∞ are either discounted sums or lim inf of the average.

Thus Γ^∞ satisfies all the hypotheses of the corollary in Section 5.1. Hence, by (26),

$$\eta_\varepsilon(\Gamma^\infty) = \eta_0(\Gamma_{\bar{c}}^\infty) = N(\Gamma^\infty) \ ,$$

which is the anti-folk theorem. In the case of discounted sums as payoffs, $\eta_0(\Gamma_{\bar{c}}^\infty)$ can be sharply characterized by:

$$(s^{x_0}, s^{x_1}, \ldots, s^{x_m}, \ldots) \in \eta_0(\Gamma_{\bar{c}}^\infty) \leftrightarrow \text{each } s^{x_\ell} \text{ is an N.E. of } \Gamma, \ \ell = 0, 1, \ldots \ .$$

6. BOUDS ON ALL OBSERVATIONS

It is impelling to consider the case when the bound on a player's capacity of observation pertains not only to (*) others' moves, but also to (**) his own payoffs and (***) his own strategies.

To take (***) into account we could make the implicit assumption that $d(\alpha,\beta) > \varepsilon$ for distinct α, β in S_i^x.

In the presence of (**), and without the constraint of (*) ((***) may, but need not, be in the background), we have: $\sigma \in \Sigma(\Gamma)$ is a <u>$[\tilde{\varepsilon}]$-N.E.</u> of Γ if

$$h_i(\xi(\sigma|\tau_i)) \geq h_i(\xi(\sigma)) + \tilde{\varepsilon} \tag{27}$$

for all $\tau_i \in \Sigma_i(\Gamma)$ and $i \in N$. If the ε-constraint of (*) (as in Section 3) is added, then we are left with <u>$[\tilde{\varepsilon}]$-(ε)-N.E.</u>'s which is defined exactly as in (27) but with $\xi(\sigma|\tau_i)$ replaced by $\xi^{\varepsilon}(\sigma|\tau_i)$ and the added requirement that $\xi(\sigma)$ be ε-inner. The interpretation of these modified N.E.'s is clear.

Denote the set of plays produced at (ε_1)-N.E.'s, $[\varepsilon_2]$-N.E.'s, $[\varepsilon_2]$-(ε_1)-N.E.'s of Γ by $\eta_{\varepsilon_1}(\Gamma)$, $\eta^{\varepsilon_2}(\Gamma)$, $\eta_{\varepsilon_1}^{\varepsilon_2}(\Gamma)$. ($\eta_{\varepsilon_1}(\Gamma)$ was introduced earlier.) Then

$$\eta_{\varepsilon_1}(\Gamma) \subset \eta_{\varepsilon_1}^{\varepsilon_2}(\Gamma) \tag{28}$$

If $\Gamma_1 \prec \Gamma_2$ then $\eta_{\varepsilon}^{\varepsilon}(\Gamma_1) \subset \eta_{\varepsilon_1}^{\varepsilon_2}(\Gamma_2)$ (29)

Under the hypotheses of Proposition 1, with ε_1 in place of ε in (b), we have: if $p(z)$ is a $[\varepsilon_2]$-(ε_1)-N.E. play of $\tilde{\Gamma}$, then it is a $[\varepsilon_2]$-(ε_1)-N.E. play of $\tilde{\Gamma}$ for every $\tilde{\Gamma} \in \Delta(\Gamma)$. (30)

Under the hypotheses of Proposition 2, if $r > B/\varepsilon_1$, then (31)

$$\eta_{\varepsilon_1}^{\varepsilon_2}(r\Gamma) = \eta_{\varepsilon}^{\varepsilon_2}(\tilde{\Gamma}) \quad \text{for any} \quad \tilde{\Gamma} \in \Delta(r\Gamma)$$

If Γ is layered, and satisfies (16), (21),(22): (32)

$$\eta_{\varepsilon_1}^{\varepsilon_2}(\Gamma) = \eta^{\varepsilon_2}(\Gamma_{c*}) = \eta^{\varepsilon_2}(\Gamma_{\bar{c}})$$

(28) follows from the definition. (29), (30) and (31) can be established in the same manner as (4), Proposition 1, Proposition 2. (32) is straightforward to check.

In the presence of (**) or (***) (or both), Proposition 3, and thus the anti-folk theorem of Section 5.2, breaks down. We will now show (Proposition 4 below) that it breaks in a continuous way as a function of the two additional bounds ε_2, ε_3 introduced here. Thus if their magnitude is small, an approximate version of the anti-folk theorem still obtains.

We start with an underlying game $\hat{\Gamma} = (N, \hat{X}, \pi, \{\hat{S}^x\}_{x \in \hat{X}}, \ldots)$. A game Γ will be called <u>ε-embedded</u> in $\hat{\Gamma}$ if it is embedded in $\hat{\Gamma}$ (see Section 4), and satisfies (denoting $\Gamma = (N, X, \pi, \{S^x\}_{x \in X}, \ldots)$):

for any $\alpha \in \hat{S}_i^x$ there is a $\beta \in S_i^x$ with $\|\alpha-\beta\| \le \varepsilon$. (33)

Put $G_\varepsilon = \{\Gamma : \Gamma$ is ε-embedded in $\hat{\Gamma}\}$, $N(\varepsilon_1, \varepsilon_2, \varepsilon_3) = \bigcup_{\Gamma \in G_{\varepsilon_3}} \eta_{\varepsilon_1}^{\varepsilon_2}(\Gamma)$.

Since $\hat{\Gamma}$ obviously belongs to G_ε, G_ε is not empty.

We shall examine the continuity properties of the set $N(\varepsilon_1, \varepsilon_2. \varepsilon_3)$ in the variables ε_2, ε_3. Note that this set covers $\eta_{\varepsilon_1}^{\varepsilon_2}(\Gamma)$ for any

$\Gamma \in G_{\varepsilon_3}$, and that $N(\varepsilon_1, 0, 0) = \eta_{\varepsilon_1}(\hat{\Gamma})$. The assumption on payoffs:

$$h_i : P(\hat{\Gamma}) \to \mathbb{R} \text{ is continuous for each } i \in N \tag{34}$$

will be needed.

Proposition 4.

(i) If $0 \leq \varepsilon_1$, $0 \leq \varepsilon_2^* \leq \varepsilon_2$, $0 \leq \varepsilon_3^* \leq \varepsilon_3$, then;

$$N(\varepsilon_1, \varepsilon_2^*, \varepsilon_3^*) \subset N(\varepsilon_1, \varepsilon_2, \varepsilon_3) .$$

(ii) Suppose that $\hat{\Gamma}$ is a layered game satisfying (16),(21),(22), (34). Then, for fixed $\varepsilon_1 > 0$,

(a) $N(\varepsilon_1, 0, 0) = \bigcap_{\substack{\varepsilon_2 > 0 \\ \varepsilon_3 > 0}} N(\varepsilon_1, \varepsilon_2, \varepsilon_3)$;

(b) if, furthermore, each $\hat{S}_i^x$ is closed then $N(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ is upper and lower semi-continuous in ε_2 , ε_3 at $(\varepsilon_1, 0, 0)$.

Proof. (i) Consider a play p in $N(\varepsilon_1, \varepsilon_2^*, \varepsilon_3^*)$. Then there is a game Γ in $G_{\varepsilon_3^*}$ such that for some $[\varepsilon_2^*]$-(ε_1)-N.E. σ of Γ , $\xi(\sigma) = p$. Since $\varepsilon_3 \geqq \varepsilon_3^*$, $G_{\varepsilon_3} \subset G_{\varepsilon_3^*}$, which implies $\Gamma \in G_{\varepsilon_3}$. Furtheremore, since $\varepsilon_2 \geqq \varepsilon_2^*$,

$$h_i(\xi(\sigma)) + \varepsilon_2 \geqq h_i(\xi(\sigma)) + \varepsilon_2^* \geqq h_i(\xi^{\varepsilon_1}(\sigma|\tau_i))$$

for all $\tau_i \in \Sigma_i(\Gamma)$ and $i \in N$.

(ii)(a) It will suffice to show that evey p in $\bigcap_{\substack{\varepsilon_2>0\\ \varepsilon_3>0}} N(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ is a primitive Nash play of $\hat{\Gamma}$, for then $N(\varepsilon_1, 0, 0) = \eta_{\varepsilon_1}(\hat{\Gamma})$ $= N(\hat{\Gamma}) = \bigcap_{\substack{\varepsilon_2>0\\ \varepsilon_3>0}} N(\varepsilon_1, \varepsilon_2, \varepsilon_3)$, using (i) and the corollary of Section 5.

We will show this by contradiction. Suppose some p in $\bigcap_{\substack{\varepsilon_2>0\\ \varepsilon_3>0}} N(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ is not a primitive Nash play. Then there is a $q \in \mathcal{C}_i(p)$ such that $h_i(p \oplus q) > h_i(p)$. By (21), (22),

$$h_i(\lambda(p \oplus q)) + (1-\lambda)\ p \geq \lambda h_i(p \oplus q) + (1-\lambda)h_i)(p) > h_i(p)$$
$$\text{for any } 0 < \lambda < 1.$$

Choose $\tilde{\lambda}$ sufficiently close to 0 to ensure that

$$\delta^*(\tilde{\lambda}(p \oplus q) + (1-\tilde{\lambda})p,\ p) < \varepsilon_1/2\ . \tag{35}$$

This can be done because of (16). Put $(p \oplus \tilde{q}) = \lambda(p \oplus q) + (1-\lambda)p$. By (33) and (34), there exists an $\varepsilon_3 > 0$ such that for any Γ in G_{ε_3}, we can find a $\tilde{\tilde{q}} \in \mathcal{C}_i(p)$ satisfying

$$\delta^*(p \oplus \tilde{\tilde{q}},\ p \oplus \tilde{q}) < \varepsilon_1/2 \text{ and } h_i(p \oplus \tilde{\tilde{q}}) > h_i(p)\ . \tag{36}$$

It follows from (35) and (36) that

$$\delta^*(p \oplus \tilde{\tilde{q}},\ p) \leq \delta^*(p \oplus \tilde{\tilde{q}},\ p \oplus \tilde{q}) + \delta^*(p \oplus \tilde{q},\ p) < \varepsilon_1\ . \tag{37}$$

Again choose ε_2 sufficiently close to 0 to ensure (using (34)) that

$$h_i(p \oplus \tilde{\tilde{q}}) > h_i(p) + \varepsilon_2 . \tag{38}$$

Take any game $\Gamma \in G_{\varepsilon_3}$ such that $p \in \Lambda(\Gamma\_)$. Let $(p \oplus \tilde{\tilde{q}}) = (s^{x_0}, s^{x_1}, \ldots)$. Then choose $\tau_i \in \Sigma_i(\Gamma) \in S^i$ in such a way as to satisfy:

$$\sigma_i(x) = \begin{cases} s^x{}_i & \text{if } i \in \pi(x) \text{ and } x \in \{x_0, x_1, \ldots\} \\ \text{arbitrary otherwise.} \end{cases} \tag{39}$$

Clearly this can be donw. Consider any σ in $\Sigma(\Gamma)$ for which $\xi(\sigma) = p$. Then we have by (37) and (39), $\xi(\sigma|\tau_i) = (p \oplus \tilde{\tilde{q}})$, i.e.,

$$\delta^*(\xi(\sigma|\tau_i), p) < \varepsilon_1 .$$

By (38)

$$h_i(\xi(\sigma|\tau_i) = h_i(p \oplus \tilde{\tilde{q}}) > h_i(p) + \varepsilon_2 .$$

Therefore σ is not a $[\varepsilon_2]$-(ε_1)-N.E. of Γ, i.e., $p \notin N(\varepsilon_1, \varepsilon_2, \varepsilon_3)$. This is a contradiction.

(ii)(b) Lower semi-continuity is straightforward. Upper semi-continuity can be proved by the method used in (ii)(a).

Q.E.D.

REFERENCE

P.Dubey and M.Kaneko, Information Patterns and Nash Equilibria in Extensive Games I, Mathematical Social Sciences 8 (1984).