JOURNAL OF ECONOMIC THEORY **82**, 267–276 (1998)
ARTICLE NO. ET982418

Payoff Continuity in Incomplete Information Games

Atsushi Kajii*

Institute of Policy and Planning Sciences, University of Tsukuba, 1-1-1 Tennodai, Tsukubashi Ibaraki-Ken 305, Japan
akajii@shako.sk.tsukuba.ac.jp.

and

Stephen Morris†

Departement of Economics, University of Pennssylvania, Philadelphia, Pennsylvania 19104
smorris@econ.sas.upenn.edu.

Received May 11,1995; revised March 4, 1998

An incomplete information game is defined by a probability distribution μ over a type space and payoff functions u. Probability distribution μ' is *strategically close* to μ if, for any bounded payoff functions u and any equilibrium of the game (μ, u), there exists an approximate equilibrium of the game (μ', u) under which all players get approximately the same payoffs. This note shows that two probability distributions are strategically close if and only if (1) they assign similar ex ante probability to all events; and (2) with high ex ante probability, it is approximate common knowledge that they assign similar *conditional* probabilities to all events. *Journal of Economic Literature* Classification Numbers: C72, D82. © 1998 Academic Press

1. INTRODUCTION

An incomplete information game is described by a probability distribution over a type space and payoff functions specifying each player's payoffs. How does the set of equilibrium payoffs change as the probability distribution over types changes? In a single person game (i.e., a decision problem), it is straightforward to show that weak convergence of a sequence of

* Financial support from the Nomura Foundation for Social Sciences is gratefully acknowledged.

† Financial support from NSF grant 9709601 and the Alfred P. Sloan Foundation is gratefully acknowledged.

0022-0531/98 $25.00

probability distributions is sufficient to ensure continuity of ex ante equilibrium (i.e., optimal) payoffs. Thus small probability events have only a small impact on ex ante payoffs. But in many player games, small probability events may have a large impact on ex ante equilibrium payoffs (even when players are allowed to choose ε-best responses given their types). Formally, there is a failure of lower hemicontinuity of the (interim) ε-equilibrium correspondence. The problem is that an ex ante low probability event may have a high ex ante probability impact via players' *conditional* beliefs: with high probability, some player may think it likely that some other player thinks it likely that ... the small probability event is true.

In this note, we describe the weakest topology on probability distributions that is sufficient to restore lower hemicontinuity. Two probability distributions are close in this topology if they are close in the weak topology (i.e., they assign similar ex ante probabilities to all events); and, in addition, with high ex ante probability it is approximate common knowledge that they assign similar *conditional* probabilities to all events.

Our result provides a notion of the proximity of information, where information is represented by probabilities on a fixed type space. The strategic closeness question was first studied by Monderer and Samet [5], but with information represented by partitions and fixed probabilities over states. Our result follows Monderer and Samet in showing the importance of approximate common knowledge for proximity of information; we structure the argument to emphasize the parallels with Monderer and Samet's analysis, but we do not know the precise connection. We believe that the changing probabilities approach of this note allows the basic ideas of this literature to be presented in a clear and tractable way.

Our purpose is to present our results in the simplest possible framework. In the discussion following the results, we describe how our results would be changed with alternative modelling choices, as well as the relation to other work.

2. PRELIMINARIES

Fix a finite set of players, $\mathscr{I} = \{1, ..., I\}$, with $I \geqslant 2$; players' finite action sets, $\{A_i\}_{i \in \mathscr{I}}$, with $\# A_i \geqslant 2$ for all $i \in \mathscr{I}$; a countably infinite set of types for each player, $\{T_i\}_{i \in \mathscr{I}}$, and a countably infinite set S of other uncertainty. Uncertainty is described by the state space $\Omega = T_1 \times \cdots \times T_I \times S$. The action space is $A = A_1 \times \cdots \times A_I$. We shall use notation such as t_i, t_{-i}, a_i, a_{-i} with standard meanings. If $\mu \in \Delta(\Omega)$ is a probability distribution over the state space, write $\mu(t_i)$ for the marginal probability of μ on T_i; if $\mu(t_i) > 0$, write $\mu(t_{-i}, s \mid t_i) \equiv \mu(t, s)/\mu(t_i)$ and $\mu(E \mid t_i) \equiv \sum_{((t_i, t_{-i}), s) \in E} \mu(t_{-i}, s \mid t_i)$, for any event $E \subseteq \Omega$.

2.1. *Incomplete Information Games*

An incomplete information game (μ, u) now consists of a probability distribution $\mu \in \Delta(\Omega)$, and players' utility functions, $u \equiv \{u_i\}_{i \in \mathscr{I}}$, where each $u_i: A \times \Omega \to \mathfrak{R}$. The game is *bounded by* M if $|u_i(a, \omega) - u_i(a', \omega)| \leqslant M$ for all $a, a' \in A$, $\omega \in \Omega$ and $i \in \mathscr{I}$.

A (mixed) strategy for player i is a function $\sigma_i: T_i \to \Delta(A_i)$. We write $\sigma(t) \equiv \{\sigma_i(t_i)\}_{i \in \mathscr{I}}$ and extend utility functions to mixed strategies in the usual way. Write $v_i[a_i, \sigma_{-i}; \mu, u; t_i]$ for type t_i's interim expected payoff if he chooses action a_i and other players follow strategy profile σ_{-i} in game (μ, u), i.e.,

$$v_i[a_i, \sigma_{-i}; \mu, u; t_i] = \sum_{(t_{-i}, s) \in T_{-i} \times S} \mu(t_{-i}, s \mid t_i)\, u_i((a_i, \sigma_{-i}(t_{-i})), (t, s)),$$

and write $V_i[\sigma; \mu, u]$ for player i's ex ante expected payoff; i.e.,

$$\begin{aligned} V_i[\sigma; \mu, u] &= \sum_{t_i \in T_i} \mu(t_i) \sum_{a_i \in A_i} (\sigma_i(t_i)(a_i))\, v_i[a_i, \sigma_{-i}; \mu, u; t_i] \\ &= \sum_{(t, s) \in \Omega} \mu(t, s)\, u_i(\sigma(t), (t, s)). \end{aligned}$$

Definition 1. Strategy profile $\sigma = \{\sigma_i\}_{i \in \mathscr{I}}$ is an (interim) ε-equilibrium of incomplete information game (μ, u) if for all $i \in \mathscr{I}$, all $t_i \in T_i$ with $\mu(t_i) > 0$, all a_i with $\sigma_i(t_i)(a_i) > 0$ and all $a_i' \in A_i$,

$$v_i[a_i, \sigma_{-i}; \mu, u; t_i] \geqslant v_i[a_i', \sigma_{-i}; \mu, u; t_i] - \varepsilon.$$

Write $\mathscr{E}^{\varepsilon}(\mu, u)$ for the set of ε-equilibria of (μ, u); a (Bayesian Nash) equilibrium is a 0-equilibrium; write $\mathscr{E}(\mu, u)$ for the set of equilibria of (μ, u).

2.2 *p-Belief*

For any event E, write $B_\mu^p(E)$ for the set of states where all players believe event E with probability at least p (under μ); for convenience, say that any event is believed for any p by a zero probability type. Thus

$$B_\mu^p(E) \equiv \{(t, s) \in \Omega : \forall i \in \mathscr{I}, \mu(t_i) > 0 \Rightarrow \mu(E \mid t_i) \geqslant p\}.$$

Iterating this operator gives

$$C_\mu^p(E) \equiv \bigcap_{n \geqslant 1} [B_\mu^p]^n(E).$$

Event E is *p-evident* if whenever it is true, each individual i believes it with probability at least p.

DEFINITION 2. E is p-evident (under μ) if $E \subseteq B_{\mu}^{p}(E)$.

Following Monderer and Samet [4], it can be shown:

LEMMA 3. *$C_{\mu}^{p}(E) = B_{\mu}^{p}(C_{\mu}^{p}(E))$ and thus $C_{\mu}^{p}(E)$ is p-evident under μ.*

2.3. *Topologies on Probability Distributions*

The weak topology captures one intuitive notion of the closeness of probability distributions. In the countable state space setting of this paper, weak convergence is equivalent to the convergence of probabilities uniformly over events. Thus define d_0 by the rule

$$d_0(\mu, \mu') = \sup_{E \subseteq \Omega} |\mu(E) - \mu'(E)|.$$

We will require extra conditions on conditional probabilities; write $\mathscr{A}_{\mu,\mu'}(\delta)$ for the set of states where every player has similar conditional probabilities about events under μ and μ'

$$\mathscr{A}_{\mu,\mu'}(\delta) = \left\{ (t, s) \in \Omega : \begin{array}{l} \text{for all } i \in \mathscr{I}, \mu(t_i) > 0, \mu'(t_i) > 0, \\ \text{and } |\mu(E \mid t_i) - \mu'(E \mid t_i)| \leqslant \delta \text{ for all } E \subseteq \Omega \end{array} \right\}.$$

We show that for players to behave similarly under μ and μ', not only $\mathscr{A}_{\mu,\mu'}(\delta)$ must be a high probability event (for some small δ), but also $C_{\mu'}^{1-\delta}(\mathscr{A}_{\mu,\mu'}(\delta))$ must be a high probability event. So let

$$d_1(\mu, \mu') = \inf\{\delta : \mu'(C_{\mu'}^{1-\delta}(\mathscr{A}_{\mu,\mu'}(\delta))) \geqslant 1 - \delta\},$$

and

$$d^*(\mu, \mu') = \max\{d_0(\mu, \mu'), d_1(\mu, \mu'), d_1(\mu', \mu)\}.$$

Clearly, d^* is non-negative and symmetric, and $d^*(\mu, \mu') = 0$ if and only if $\mu = \mu'$. Since d^* does not necessarily satisfy the triangle inequality, d^* is not a metric. However, d^* generates a topology in the following sense: a generalized sequence $\{\mu^k : k \in K\}$, where K is a directed index set with partial order $\succ$, converges to μ if and only if for any $\varepsilon > 0$, there is a $\bar{k} \in K$ such that $k \succ \bar{k}$ implies $d^*(\mu^k, \mu) < \varepsilon$.

3. RESULTS

3.1. *Sufficiency*

First we show that if $d^*(\mu, \mu')$ is small, μ and μ' are strategically close.

LEMMA 4. *Suppose that event $E \subseteq \mathscr{A}_{\mu,\mu'}(\varepsilon_1)$ and E is $(1-\varepsilon_2)$-evident under μ'. If σ is an equilibrium of (μ, u) and u is bounded by M, there exists a $(4\varepsilon_1 + 2\varepsilon_2)M$-equilibrium σ' of (μ', u) with $\sigma'(t) = \sigma(t)$ at all $(t, s) \in E$.*

Proof. We will construct such a σ'. Let $\hat{T}_i = \{t_i \in T_i : (t, s) \in E$ for some $(t_{-i}, s) \in T_{-i} \times S\}$. Consider the modified version of (μ, μ') where players are required to play according to σ at all states (t, s) in E (i.e., player i's strategy σ_i' must satisfy $\sigma'(t_i) = \sigma_i(t_i)$ if $t_i \in \hat{T}_i$). Consider any equilibrium σ' of the modified game (this exists by standard arguments). If $t_i \notin \hat{T}_i$, then $\sigma_i'(t_i)$ is a best response to σ'_{-i} for type t_i in (μ', u) (by construction). If $t_i \in \hat{T}_i$, we have, for any $a_i \in A_i$,

$$|v_i[a_i, \sigma_{-i}; \mu', u; t_i] - v_i[a_i, \sigma_{-i}; \mu, u; t_i]| \leqslant 2\varepsilon_1 M,$$

because $E \subseteq \mathscr{A}_{\mu,\mu'}(\varepsilon_1)$; and

$$|v_i[a_i, \sigma'_{-i}; \mu', u; t_i] - v_i[a_i, \sigma_{-i}; \mu', u; t_i]| \leqslant \varepsilon_2 M,$$

because E is $(1-\varepsilon_2)$-evident. Thus,

$$|v_i[a_i, \sigma'_{-i}; \mu', u; t_i] - v_i[a_i, \sigma_{-i}; \mu, u; t_i]| \leqslant (2\varepsilon_1 + \varepsilon_2) M. \tag{1}$$

Now suppose $t_i \in \hat{T}_i$ and $\sigma_i'(t_i)[\hat{a}_i] > 0$. By construction of σ_i' and for any $a_i' \in A_i$,

$$v_i[\hat{a}_i, \sigma'_{-i}; \mu, u; t_i] \geqslant v_i[a_i', \sigma_{-i}; \mu, u; t_i]. \tag{2}$$

So,

$$\begin{aligned} v_i[\hat{a}_i, \sigma'_{-i}; \mu', u: t_i] &\geqslant v_i[\hat{a}_i, \sigma_{-i}; \mu, u; t_i] - (2\varepsilon_1 + \varepsilon_2) M, && \text{by (1)}, \\ &\geqslant v_i[a_i', \sigma_{-i}; \mu, u; t_i] - (2\varepsilon_1 + \varepsilon_2) M, && \text{by (2)}, \\ &\geqslant v_i[a_i', \sigma'_{-i}; \mu', u; t_i] - (4\varepsilon_1 + 2\varepsilon_2) M, && \text{by (1)}. \quad \blacksquare \end{aligned}$$

PROPOSITION 5. *Suppose that $d^*(\mu, \mu') \leqslant \delta$. Then if σ is an equilibrium of (μ, u) and u is bounded by M, there exists a $6\,\delta M$-equilibrium σ' of (μ', u) with*

$$|V_i[\sigma; \mu, u] - V_i[\sigma'; \mu', u]| \leqslant 3\,\delta M \qquad \text{for all} \quad i \in \mathscr{I}.$$

Proof. If $d^*(\mu, \mu') \leqslant \delta$, then $d_0(\mu, \mu') \leqslant \delta$ and thus

$$|\mu(E) - \mu'(E)| \leqslant \delta \qquad \text{for all} \quad E \subseteq \Omega; \tag{3}$$

also $d_1(\mu, \mu') \leqslant \delta$ and thus

$$\mu'[C_{\mu'}^{1-\delta}(\mathscr{A}_{\mu,\mu'}(\delta))] \geqslant 1 - \delta. \tag{4}$$

Now let σ be any equilibrium of (μ, u). By Lemma 3, $C_{\mu'}^{1-\delta}(\mathscr{A}_{\mu,\mu'}(\delta)$ is $(1-\delta)$-evident (under μ'). So by Lemma 4 (with $E = C_{\mu'}^{1-\delta}(\mathscr{A}_{\mu,\mu'}(\delta))$ and $\varepsilon_1 = \varepsilon_2 = \delta$) there exists a $6\delta M$-equilibrium of (μ', u) with $\sigma'(t) = \sigma(t)$ at all $(t, s) \in C_{\mu'}^{1-\delta}(\mathscr{A}_{\mu,\mu'}(\delta)$. Now (4) implies

$$|V_i(\sigma'; \mu', u) - V_i(\sigma; \mu', u)| \leqslant \delta M$$

and (3) implies

$$|V_i(\sigma; \mu', u) - V_i(\sigma; \mu, u)| \leqslant 2\,\delta M. \quad \blacksquare$$

3.2. *Necessity*

Conversely, we shall show that $d^*(\mu, \mu')$ must be small if μ and μ' are strategically close. To do this, we construct u and an equilibrium σ of (μ, u) such that no approximate equilibrium of (μ', u) is close to σ when $d^*(\mu, \mu')$ is large. We do this first for the trivial case when $d_0(\mu, \mu')$ is large and then for the more interesting case when $d_0(\mu, \mu')$ is small but $d_1(\mu, \mu')$ is large.

PROPOSITION 6. *If $d_0(\mu, \mu') > \delta$, there exists u bounded by* 1 *and an equilibrium σ of (μ, u) such that every δ-equilibrium of (μ', u), σ', satisfies $V_i(\sigma; \mu, u) - V_i(\sigma'; \mu', u) > \delta$ for all i.*

Proof. If $d_0(\mu, \mu') > \delta$, there exists $E \subseteq \Omega$ such that $\mu(E) - \mu'(E) > \delta$. Consider the degenerate game where

$$u_i(a, \omega) = \begin{cases} 1, & \text{if} \quad \omega \in E \\ 0, & \text{otherwise} \end{cases}.$$

Clearly, $V_i(\sigma; \mu, u) - V_i(\sigma'; \mu, u) > \delta$ for *any* strategy profiles σ and σ'. $\blacksquare$

To deal with the case where $d_1(\mu, \mu')$ is large, we will first construct a class of games illustrating why ex ante equilibrium payoffs may be very different even when $d_0(\mu, \mu')$ is small. We will use this construction in proving the next Lemma.

The Infection Game. Label actions $A_i = \{x_i, y_i, ...\}$, and for each $i \in \mathscr{I}$, fix nonempty type subsets $\{\hat{T}_i\}_{i \in \mathscr{I}}$, each $\hat{T}_i \subseteq T_i$, and some non empty events F_{t_i}, each $F_{t_i} \subseteq \Omega$, for every $t_i \in \hat{T}_i$. Write $\hat{E} = \hat{T}_1 \times \cdots \times \hat{T}_I \times S$; let

$$u_i((x_i, a_{-i}), (t, s)) = 0;$$

$$u_i((y_i, a_{-i}), (t, s)) = \begin{cases} 1 - \mu(F_{t_i} \mid t_i), \text{ if } t_i \notin \hat{T}_i \text{ and } (t, s) \in F_{t_i} \\ -\mu(F_{t_i} \mid t_i), \text{ if } t_i \notin \hat{T}_i \text{ and } (t, s) \notin F_{t_i} \\ -\varepsilon, \text{ if } t_i \in \hat{T}_i \text{ and } a_{-i} = x_{-i} \\ 2, \text{ if } t_i \in \hat{T}_i \text{ and } a_{-i} \neq x_{-i} \end{cases};$$

where $\varepsilon > 0$, and, for all $z_i \notin \{x_i, y_i\}$, $u_i((z_i, a_{-i}), (t, s)) = -1$, for all $a_{-i} \in A_i$ and $(t, s) \in \Omega$.

The game has the following interpretation. All actions other than x and y are strictly dominated for all players, and are thus irrelevant. If a player is a "strategic type" (i.e., $t_i \in \hat{T}_i$), then he will choose x if other players choose x, and y if other players choose y. If a player is a "committed type" (i.e., $t_i \notin \hat{T}_i$), then his best response is independent of the actions of other players. Note that u is bounded by 3.

Game (μ, u) has an equilibrium σ under which all types of all players always choose action x_i: the expected payoff to each type of each player is 0; the expected payoff to deviating to y_i is 0 for committed type $t_i \notin \hat{T}_i$ or $-\varepsilon$ for standard type $t_i \in \hat{T}_i$.

But now consider probability distribution μ' with $\mu'(F_{t_i} \mid t_i) - \mu(F_{t_i} \mid t_i) > \varepsilon$ for all $t_i \notin \hat{T}_i$. If $t_i \notin \hat{T}_i$, then the payoff to y_i is:

$$\mu'[F_{t_i} \mid t_i](1 - \mu[F_{t_i} \mid t_i]) + (1 - \mu'[F_{t_i} \mid t_i])(-\mu[F_{t_i} \mid t_i])$$
$$= \mu'[F_{t_i} \mid t_i] - \mu[F_{t_i} \mid t_i] > \varepsilon.$$

Thus y_i is the unique ε-best response irrespective of others' actions; that is, at any ε-equilibrium of (μ', u), type $t_i \notin \hat{T}_i$ plays y_i. So at least one player chooses y at each state (t, s) in $\Omega \backslash \hat{E}$.

Now consider type t_i with $\mu'(\hat{E} \mid t_i) < 1 - \varepsilon$. With at least probability ε, at least one of the opponents is playing y. So payoff to y_i is at least $2\varepsilon - (1 - \varepsilon)\varepsilon > \varepsilon$. Thus y_i is unique ε-best response for t_i in any ε-equilibrium of (μ', u). By definition, at each state (t, s) in $\Omega \backslash B_{\mu'}^{1-\varepsilon}(\hat{E})$, at least one player i has $\mu'(\hat{E} \mid t_i) < 1 - \varepsilon$, hence we conclude that at least one player i chooses y_i at each state in $\Omega \backslash B_{\mu'}^{1-\varepsilon}(\hat{E})$, i.e., the strategic types in $\Omega \backslash B_{\mu'}^{1-\varepsilon}(\hat{E})$ are "infected" by the committed types. We can then look at type t_i with $\mu'(B_{\mu'}^{1-\varepsilon}(\hat{E}) \mid t_i) < 1 - \varepsilon$, and the argument iterates. So at every state in $\Omega \backslash C_{\mu'}^{1-\varepsilon}(\hat{E})$, at least one player chooses action y_i. Thus every type t_i with $\mu'(C_{\mu'}^{1-\varepsilon}(\hat{E}) \mid t_i) < 1 - \varepsilon$ has interim payoff more than ε in any ε-equilibrium.

Note that $d_0(\mu, \mu')$ is irrelevant for the above construction; in particular, it is possible to have $d_0(\mu, \mu')$ small (which implies that $\mu'(\hat{E})$ is close to 1), but still have $\mu'(C_{\mu'}^{1-\varepsilon}(\hat{E}))$ much smaller than 1. In this case, equilibrium payoffs of the game (μ, u) (which include all players having ex ante payoff 0 by choosing action x in equilibrium) may be very different from equilibrium payoffs of the game (μ', u) (where in every equilibrium, with high probability, each player gets payoff at least ε by choosing action y). In the following Proposition, we verify that $d_1(\mu, \mu')$ large is *exactly* the condition we need to construct $\hat{E}$ with these properties.

PROPOSITION 7. *If $d_1(\mu, \mu') > \delta$, there exists u bounded by 3 and an equilibrium σ of (μ, u) such that every δ-equilibrium of (μ', u), σ', satisfies*

$$V_i[\sigma'; \mu', u] - V_i[\sigma; \mu, u] > \frac{\delta^2}{I} \qquad \text{for some} \quad i \in \mathscr{I}.$$

Proof. By construction, $\mathscr{A}_{\mu,\mu'}(\delta)$ has the form $\hat{T}_1 \times \cdots \times \hat{T}_I \times S$, where $t_i \notin \hat{T}_i$ holds if and only if there is event $F_{t_i} \subseteq \Omega$ with $\mu'(F_{t_i} \mid t_i) - \mu(F_{t_i} \mid t_i) > \delta$. Set $\hat{E} = \mathscr{A}_{\mu,\mu'}(\delta)$, and consider the infection game constructed for these events (with $\varepsilon = \delta$). In the equilibrium σ of (μ, u) where x is always played, $V_i(\sigma; \mu, u) = 0$ for all $i \in \mathscr{I}$. Since $d_1(\mu, \mu') > \delta$, $\mu'[C_{\mu'}^{1-\delta}(\hat{E})] < 1 - \delta$. By Lemma 3, $C_{\mu'}^{1-\delta}(\hat{E}) = B_{\mu'}^{1-\delta}(C_{\mu'}^{1-\delta}(\hat{E})) = \bigcap_{j \in \mathscr{I}} \Omega \backslash Z_j$, where $Z_j = \{(t, s) : \mu'(C_{\mu'}^{1-\delta}(\hat{E}) \mid t_j) < 1 - \delta\}$. So there exists a player i with $\mu'(Z_i) > \delta/I$. In any equilibrium σ of (μ', u), player i's interim payoff is more than δ at all states in Z_i by construction. His interim payoff is at least 0 at all states in $\Omega \backslash Z_i$. Thus his ex ante payoff is more than $(\delta/I)\delta + (1 - \delta/I)\, 0 = \delta^2/I$. ∎

3.3. *Summary of Results*

To sum up, Proposition 5 showed that if $d^*(\mu, \mu')$ is small, then ex ante equilibrium payoffs are close in (μ, u) and (μ', u). If $d^*(\mu, \mu')$ is large, then *either* $d_0(\mu, \mu')$ is large, in which case Proposition 6 shows that ex ante payoffs may be very different; *or* $d_1(\mu, \mu')$ (or $d_1(\mu', \mu)$) is large, in which case Proposition 7 shows that ex ante payoffs may be very different.

Our results may be summarized more formally as follows. Write

$$\phi(\mu, \mu'; u, \varepsilon) = \sup_{\sigma \in \mathscr{E}(\mu, u)} \; \inf_{\sigma' \in \mathscr{E}^{\varepsilon}(\mu', u)} \; \max_{i \in \mathscr{I}} |V_i[\sigma'; \mu', u] - V_i[\sigma; \mu, u]|$$

and

$$\phi^*(\mu, \mu'; u, \varepsilon) = \max\{\phi(\mu, \mu'; u, \varepsilon), \phi(\mu', \mu; u, \varepsilon)\}.$$

So $\phi^*(\mu, \mu'; u, \varepsilon)$ measures the strategic closeness of μ and μ'.

PROPOSITION 8. *$d^*(\mu^k, \mu) \to 0 \Leftrightarrow \phi^*(\mu^k, \mu; u, \varepsilon) \to 0$ for all bounded u and $\varepsilon > 0$.*

Proof. [$\Rightarrow$] By Proposition 5, if u is bounded by M and $\varepsilon \geqslant 6Md^*(\mu^k, \mu)$, $\phi^*(\mu^k, \mu; u, \varepsilon) \leqslant 3Md^*(\mu^k, \mu))$. Thus $\sigma^*(\mu^k, \mu, u, \varepsilon) \to 0$, for all bounded u and $\varepsilon > 0$ if $d^*(\mu^k, \mu) \to 0$.

[$\Leftarrow$] If $d^*(\mu', \mu) > \delta$, then by Propositions 6 and 7 we can find u bounded by 3 such that either $\phi(\mu', \mu; u, \varepsilon) > \delta^2/I$ or $\phi(\mu, \mu'; u, \varepsilon) > \delta^2/I$, and thus $\phi^*(\mu', \mu; u, \varepsilon) > \delta^2/I$. ∎

Thus the topology generated by d^* is the weakest such that $\phi^*(\mu^k, \mu; u, \varepsilon) \to 0$, for all bounded u and $\varepsilon > 0$, for any convergent sequence.

4. DISCUSSION

The Relation to Monderer and Samet [4]. Our work follows Monderer and Samet [4] who considered an alternative system of perturbing an information system. They fixed a state space and probability distribution and considered variations in the countable partitions of that state space that players observe. Our characterization of the proximity of information has a similar flavor to Monderer and Samet's, but we have not been able to establish a direct comparison. By considering a fixed type space, we exogenously determine which types in the information systems correspond to each other. In the Monderer and Samet approach, it is necessary to work out how to identify types in the two information systems. Thus we conjecture that two information systems are close in Monderer and Samet's sense if and only if the types in their construction can be labelled in such a way that the information systems are close in our sense.

Upper Hemicontinuity. This note addresses a lower hemicontinuity question. In the discrete state space setting of this paper, the weak topology is sufficient for upper hemicontinuity. Milgrom and Weber [3] analyze upper hemicontinuity in general state spaces.

Alternative Notions of Approximate Equilibrium.

— It is crucial that the notion of ε-equilibrium is *interim*. We require that each player's action choice be within ε of a best response, *contingent on his realized type*. If we required only ex ante ε-equilibrium, we would be allowing players to make choices that are very far from being best responses (even if only with ex ante small probability). If we had used ex ante ε-equilibrium in our construction, the weak topology generated by d_0 would be sufficient to generate payoff continuity in the countable state space setting considered in this paper.[1]

— We require *pure strategy* interim ε-equilibria; that is, we require that each pure strategy played with strictly positive probability (no matter how small) be within ε of a best response. By contrast, both Monderer and Samet [5] and an earlier version of this work (Kajii and Morris [2]) used

[1] In countable state spaces, the weak topology is equivalent to the strong topology. Engl [1] shows that the strong topology is sufficient and (for uncountable state spaces) necessary for lower hemicontinuity of the ex ante ε-equilibrium correspondence.

the weaker notion of mixed strategy interim ε-equilibria: any type's mixed strategy must generate an interim expected payoff (contingent on his type) within ε of a best response but might involve playing (with small probability) actions that give payoffs a long way from a best response. The results are unaffected by this distinction, but the arguments are simpler under the approach described here.

Sensitivity to Small Probability Events. We have seen that weak convergence is not enough for strategic closeness in general. But at some μ, weak convergence *is* sufficient. Say that probability distribution μ is *insensitive* to small probability events if $d_0(\mu^k, \mu) \to 0$ implies $d^*(\mu^k, \mu) \to 0$. Using the fact that convergence of conditional probabilities is uniform on a finite event, it can be shown that a necessary and sufficient condition for μ to be insensitive is that it can be approximated on a finite subset of Ω; that is, for all $\varepsilon > 0$, there exists a *finite* $(1-\varepsilon)$-evident (under μ) event E such that $\mu(E) \geqslant 1-\varepsilon$. Then it is straightforward to identify a number of sufficient conditions for μ to be insensitive:

— μ has *finite support*: $\{t \colon \mu(t) > 0\}$ is finite.

— μ is *independent* over types: for all $t \in T$, $\sum_{s \in S} \mu(t, s) = \prod_{i \in \mathscr{I}} \mu(t_i)$.

— μ is *perfectly correlated* over types: for all $(t, s) \in \Omega$, $\mu(t, s) > 0 \Rightarrow t_i = \cdots = t_I$.

REFERENCES

1. G. Engl, Lower hemicontinuity of the Nash equilibrium correspondence, *Games Econom. Behav.* **9** (1995), 151–160.
2. A. Kajii and S. Morris, "Payoff Continuity in Incomplete Information Games," CARESS Working Paper Paper 94-17, University of Pennsylvania, Philadelphia, 1994.
3. P. Milgrom and R. Weber, Distributional strategies for games with incomplete information, *Mat. Oper. Res.* **10** (1985), 619–632.
4. D. Monderer and D. Samet, Approximating common knowledge with incomplete information, *Games and Econom. Behav.* **1** (1989), 170–190.
5. D. Monderer and D. Samet, Proximity of information in games with common beliefs, *Math. Oper. Res.* **21** (1996), 707–725.