SCAN TO PDF.

CARESS Working Paper #94-17
Payoff Continuity in Incomplete Information Games and Almost Uniform Convergence of Beliefs

by

Atsushi Kajii and Stephen Morris

UNIVERSITY of PENNSYLVANIA

Center for Analytic Research in Economics and the Social Sciences

McNEIL BUILDING, 3718 LOCUST WALK

PHILADELPHIA, PA 19104-6297

CARESS Working Paper #94-17

Payoff Continuity in Incomplete Information Games and Almost Uniform Convergence of Beliefs

Atsushi Kajii and Stephen Morris
Department of Economics
University of Pennsylvania

September 1994

Abstract

In single person decision problems, pointwise convergence of prior probabilities is sufficient to ensure continuity of equilibrium payoffs, since the set of states where conditional probabilities are badly behaved has small probability in the limit. But in many person decision problems - i.e. games - this argument breaks down: because of strategic interaction, arbitrarily small probability events matter. We show that "almost uniform convergence" of conditional probabilities is necessary and sufficient for continuity of payoffs.

1. Introduction

The paper studies the lower-hemicontinuity of equilibria of incomplete information games with respect to the information structure. To describe the closeness of information structures, we fix a countable state space, a game on that state space and a group of individuals' information partitions. Consider a sequence of probability distributions P^k on the state space which converges pointwise to some limit distribution P^∞. Each P^k completes the description of the information structure. So the continuity question becomes the following: in what additional sense must P^k converge to P^∞

to guarantee that for any equilibrium of the game under P^∞, there is a nearby equilibrium of the game under P^k.

As a benchmark, consider first the special case of a single person game, i.e. a decision problem under uncertainty. If the individual's conditional probabilities under P^k are close to those under P^∞, then actions which are optimal given P^∞ will be almost optimal given P^k. But as $k \to \infty$, the set of states where the individual's conditional probabilities vary significantly from those under P^∞ becomes arbitrarily small. Thus there is continuity of payoffs in the following sense: for sufficiently large k and for any decision rule optimal under P^∞, there is another decision rule which is almost optimal under P^k which gives almost the same (ex ante) expected utility. So pointwise convergence is already strong enough to guarantee payoff continuity.

It might appear that a similar logic would work in general - i.e. many person - games. Because of strategic interaction, however, arbitrarily small events may matter in games. Suppose player 1's conditional probabilities are badly behaved on a small event E_1. If player 2 has a good reason to believe that player 1 will behave very differently on a small event E_1, then player 2 may also want to behave differently on another event E_2, and then player 3 may want to behave differently on E_3, and so on. Although E_1 may be a small event, the union of this chain of events may be a large event. So not only must the probability of the set of states where conditional probabilities are badly behaved go to zero as $k \to \infty$, but also the whole union of the chain of events must become a small event.

Roughly speaking, the latter condition is equivalent to what we term "almost uniform convergence of conditional probabilities". The main result of this paper shows that almost uniform convergence of conditional probabilities is necessary and sufficient for payoff continuity.

We will discuss the significance of the "almost" qualifier and give an outline of the paper below. But it will be useful to first put this continuity result in the context of a recent literature on "almost common knowledge". Rubinstein's (1989) electronic mail game showed that small probability events could have large consequences because of strategic interaction. Morris, Rob and Shin (1993) characterized the properties of an information system which allowed small probability events to have large consequences. Carlsson and van Damme (1993) show that arbitrarily small amounts of noise lead to a discontinuity of the type described above. These results are important because they suggest that in strategic environments, economic models are not robust to their common knowledge assumptions. We would like to characterize what the relevant notion of closeness of information systems is for strategic environments.

Monderer and Samet [MS] (1990) and Stinchcombe (1988) have carried out an analogous exercise where they fix the prior probability measure and vary the information partitions. In some framework, the exercise we carry out must be a corollary of their exercises (and/or vica-versa). Indeed both the general structure of our argument and a number of steps in our proofs are directly inspired by MS. However, we believe our approach is intuitive and of independent interest. In the concluding discussion, we speculate on the exact relation.

The paper is organized as follows. In section 2, we introduce notation for the information system and incomplete information games. We also give two key lemmas confirming that closeness of conditional probabilities is the necessary and sufficient condition for strategies which are optimal under the limit distribution to be almost optimal under any distribution close enough to the limit.

In section 3, we consider a strong notion of continuity of actions. We say that there is continuity of actions if every equilibrium under the limit distribution is an ε-equilibrium close to the limit. We show that uniform convergence of conditional probabilities is necessary and sufficient for such continuity. This result is of interest because it demonstrates why uniform convergence of conditional probabilities is important; and it is independent of the number of players, i.e. the uniformity of convergence is necessary even when there is only one player and thus no strategic concerns.

In section 4, we consider a weaker notion of payoff continuity. We say there is continuity of payoffs if for every equilibrium under the limit distribution, there is an ε-equilibrium close to the limit which gives approximately the same ex ante payoffs. We show that if there is only one player, pointwise convergence is sufficient. But an example demonstrates that pointwise convergence is not sufficient if there is more than one player. On the other hand, another example demonstrates that uniform convergence of conditional probabilities is not necessary. What is required is that the set of states where conditional probabilities are well-behaved not only has high probability in the limit, but also is common p-belief, for some p close to 1, with high probability: that is, with high probability everyone believes with probability at least p that all conditional probabilities are well-behaved, everyone believes with probability at least p that everyone believes it, and so on *ad* *infinitum.* Thus we review results of Monderer and Samet (1989) characterizing common p-belief, give a formal definition of almost uniform convergence and show that it is necessary and sufficient for continuity of payoffs.

In the concluding section 5, we discuss some properties of almost uniform convergence, the significance of various assumptions, the possibility of uncountable state space extensions and the relation to the Monderer and Samet (1990) topology.

2. Preliminaries

Fix a countable state space Ω. There is a finite set of individuals $\mathcal{I} = \{1, .., I\}$. Individual $i \in \mathcal{I}$ has information partition $\mathcal{Q}_i$ of Ω. Write $Q_i(\omega)$ for the unique element of $\mathcal{Q}_i$ containing state ω. We will be concerned with a sequence of countably additive probability distributions, P^k. We write P^∞ for a limit distribution. We consider games where P^k represents players' common prior on the state space to study continuity properties of equilibria as $k \to \infty$.

We assume $P^k[Q_i(\omega)] \neq 0$, for all $\omega \in \Omega$, $i \in \mathcal{I}$, $k = 1, 2..., \infty$. Under this restriction conditional probabilities are well defined with

$$P^k[E|Q_i(\omega)] = \frac{P^k[E \cap Q_i(\omega)]}{P^k[Q_i(\omega)]} \tag{2.1}$$

This restriction is not innocuous, but it makes our analysis very clear and we believe that the benefits exceed the loss of generality. Denote by $\mathcal{P}$ the set of all countably additive probability measures that satisfy this restriction.

2.1. Convergence

A minimal requirement for any general statements about the continuity of the equilibria of incomplete information games will be that $P^k[\omega]$ converges to $P^\infty[\omega]$ at every $\omega \in \Omega$; that is, $P^k \to P^\infty$ *pointwise.* Throughout the paper we restrict attention to sequences which satisfy this property. Let us summarize some implications of pointwise convergence of probability measures in $\mathcal{P}$.

Fact 1. *If* $P^k \to P^\infty$ pointwise, *then* *(i) for any* $\delta > 0$, *there is a finite set* E^* *such that* $P^\infty[E^*] > 1 - \delta$, $P^k[E^*] > 1 - \delta$ *and* $|P^k[E^*] - P^\infty[E^*]| < \delta$ *for all large enough* k.

(ii) $P^k \to P^\infty$ *weakly.*

(iii) if E^k *is an increasing sequence of events with* $\bigcup_{k=1}^{\infty} E^k = \Omega$, *then* $P^k\left[E^k\right] \to 1$.

(iv) $P^k[E|Q_i(\omega)] \to P^\infty[E|Q_i(\omega)]$ *for every* E *and* ω.

We shall omit the straightforward proof. We just note that the countable additivity is crucial in (i) and (iii) and so is $P^\infty[Q_i(\omega)] > 0$ in (iv).

The limit behavior of *conditional* probabilities will play an important role. Fact 1(iv) showed that pointwise convergence of P^k ensures pointwise convergence of conditional probability $P^k[\omega|Q_i(\omega)]$. It is important that this is *not* true for uniform convergence, i.e. the condition that $P^k[\omega] \to P^\infty[\omega]$ uniformly in ω does not necessary imply $P^k[\omega|Q_i(\omega)] \to P^\infty[\omega|Q_i(\omega)]$ uniformly in ω (as examples 1 and 2 of section 4 will show). This is because although $P^\infty[Q_i(\omega)] > 0$ holds for all ω, it may be the case that $P^\infty[Q_i(\omega)]$ is not bounded away from zero i.e., there is a sequence ω^k such that $P^\infty\left[Q_i(\omega^k)\right] \to 0$. So the following concept is not implied by pointwise convergence, even in $\mathcal{P}$.

Definition 1. *(Uniform convergence of conditional probabilities [UCCP]) Conditional probabilities converge uniformly if for every* $\delta > 0$, *there exists* K *such that*

$$\left|P^k[E|Q_i(\omega)] - P^\infty[E|Q_i(\omega)]\right| \leq \delta, \text{ for all } E \subset \Omega,\ \omega \in \Omega,\ i \in \mathcal{I}, \text{ and } k \geq K. \tag{2.2}$$

When (2.2) holds, we shall write $P^k \to P^\infty$ *UCCP.*

2.2. Incomplete Information Games

Definition 2. *An* I *player incomplete information game* $[A, u]$ *fixes, for each player* i, *a finite set of actions,* A_i, *and a payoff function,* $u_i : A \times \Omega \to \Re$, *where* $A = A_1 \times \cdots \times A_I$ *and* $u = (u_1, ..., u_I)$.

An element a_i of A_i is called an action. A (mixed) strategy for player i is a Q_i-measurable function $\sigma_i : \Omega \to \Delta(A_i)$, where $\Delta(A_i)$ is the set of all probability measures on A_i. We denote by $\sigma_i(a_i|\omega)$ the probability that action a_i is chosen given ω under σ_i. A strategy profile is a function

$\sigma = (\sigma_i)_{i=1}^I$ where σ_i is a strategy for player i. We denote by $\sigma(a|\omega)$ the probability that action profile a is chosen given ω under σ. Write σ_{-i} for $(\sigma_j)_{j\neq i}$; write Σ_i for player i's mixed strategies and $\Sigma = \Sigma_1 \times .. \times \Sigma_I$. The payoff of strategy profile σ for player i under P^k is given by the expected utility; that is, $\sum_{\omega\in\Omega}\sum_{a\in A_i} u_i(a,\omega)\,\sigma(a|\omega)\,P^k[\omega]$. Abusing notation, we shall write $\sum_{\omega\in\Omega} u_i(\sigma,\omega)\,P^k[\omega]$ when no confusion should arise. With the natural linear structure and the pointwise convergence, the set Σ is convex and compact.

Definition 3. *A strategy σ_i is $\varepsilon-$optimal at ω under P^k given σ if for all $a_i \in A_i$,*

$$\sum_{\omega'\in Q_i(\omega)} \left(u_i\left(\sigma(\omega'),\omega'\right) - u_i\left(\left\{a_i,\sigma_{-i}(\omega')\right\},\omega'\right)\right) P^k\left[\omega'\left|Q_i(\omega)\right.\right] \geq -\varepsilon \tag{2.3}$$

Thus an action which gives a payoff significantly below the best response may be played with positive probability in an ε-optimal mixed strategy. But there it must be played with low probability. The following fact (used by Monderer and Samet (1990)) is a straightforward consequence of the definitions.

Fact 2. *If a pure strategy that assigns action a_i to ω is not ε_1-optimal at ω under P^k given σ, then action a_i is played with conditional probability no more than $\frac{\varepsilon_2}{\varepsilon_1}$ in any ε_2-optimal mixed strategy σ_i at ω under P^k given σ; that is, $\sigma_i(a_i|\omega) \leq \frac{\varepsilon_2}{\varepsilon_1}$*

Definition 4. *A strategy σ_i is $\varepsilon-$optimal if it is $\varepsilon-$optimal at every ω. Strategy profile σ is an ε-equilibrium for $[A,u]$ under P^k if each σ_i is $\varepsilon-$optimal given σ.*

This definition requires *interim* ε-optimality of *mixed* strategies. In section 5.2, we discuss how the results of the paper change completely with an *ex ante* notion, but would remain unchanged if *pure* strategy ε-optimality was required.

By convention, "optimal" means 0-optimal and "equilibrium" means 0-equilibrium. The following is a consequence of Debreu (1952):

Fact 3. *There exists an equilibrium in any game where the set of strategies is non-empty, compact and convex and each payoff function is continuous, and quasi-concave in σ_i.*

2.3. ε-Optimality and Conditional Probabilities

If σ_i is optimal under P^∞ given σ at ω, then intuitively σ_i will be ε-optimal under P^k given σ at ω if conditional probabilities are "close" at ω . To make this idea formal, it is useful first to identify those states ω where conditional probabilities under P^k are similar to those under P^∞. Formally:

$$\mathbf{A}_i^k(\delta) \equiv \left\{\omega : \left|P^k[E|Q_i(\omega)] - P^\infty[E|Q_i(\omega)]\right| \leq \delta \text{ , for all events } E \subset \Omega\right\}$$

$$\text{Also write } \mathbf{A}_*^k(\delta) \equiv \bigcap_{i \in \mathcal{I}} \mathbf{A}_i^k(\delta)$$

This δ-*conditional agreement set* does not depend on the choice of game; that is, it is a property of the "information system", rather than of any particular game. Note also that $\mathbf{A}_i^k(\delta)$ is a Q_i measurable set by construction. The convergence notions we discussed in the previous section are related to the limit behavior of the δ-conditional agreement set as follows.

Lemma 1. *If $P^k \to P^\infty$ pointwise, then for any $\delta > 0$, as $k \to \infty$, $P^k[\mathbf{A}_i^k(\delta)] \to 1$ for all $i \in \mathcal{I}$ and thus $P^k[\mathbf{A}_*^k(\delta)] \to 1$.*

Lemma 2. *$P^k \to P^\infty$ UCCP if and only if for each $\delta > 0$, there exists K such that for all $k \geq K$, $\mathbf{A}_*^k(\delta) = \Omega$.*

Proof. Lemma 2 follows from the definitions. For Lemma 1, observe that by facts 1(i) and 1(iv), for each i and ω, there exists a finite event $E_i^*(\omega)$ and integer $K_i(\omega)$ such that for all $k \geq K_i(\omega)$,

$$P^k\left[E_i^*(\omega)\,|Q_i(\omega)\right] > 1 - \tfrac{\delta}{3}$$

$$P^\infty\left[E_i^*(\omega)\,|Q_i(\omega)\right] > 1 - \tfrac{\delta}{3}$$

$$\left|P^k\left[E_i^*(\omega) \cap E\,|Q_i(\omega)\right] - P^\infty\left[E_i^*(\omega) \cap E\,|Q_i(\omega)\right]\right| < \tfrac{\delta}{3} \text{ for all } E \subset \Omega$$

Thus, for all $k \geq K_i(\omega)$ and $E \subset \Omega$, we have

$$\begin{aligned}\left|P^k\left[E\,|Q_i(\omega)\right] - P^\infty\left[E\,|Q_i(\omega)\right]\right| &\leq \left\{\begin{array}{c}\left|P^k\left[E_i^*(\omega) \cap E\,|Q_i(\omega)\right] - P^\infty\left[E_i^*(\omega) \cap E\,|Q_i(\omega)\right]\right| \\ \left|P^k\left[\Omega \setminus E_i^*(\omega)\,|Q_i(\omega)\right]\right| + \left|P^\infty\left[\Omega \setminus E_i^*(\omega)\,|Q_i(\omega)\right]\right|\end{array}\right\} \\ &< \delta\end{aligned}$$

Thus $\omega \in \mathbf{A}_i^k(\delta)$ for all $k \geq K_i(\omega)$ and $P^k[\mathbf{A}_i^k(\delta)] \geq P^k[\{\omega : K_i(\omega) \leq k\}] \to 1$ as $k \to \infty$, by Fact 1(iii). □

Intuitively, if σ_i is optimal against σ_{-i} under P^∞, then if $\omega \in \mathbf{A}_i^k(\delta)$ occurs, σ_i will not do so badly against σ_{-i} under P^k, since by the definition of $\mathbf{A}_i^k(\delta)$, player i's beliefs are close to those under P^∞. But notice that there is an inessential source of discontinuity due to the fact that any affine transformation of the utility function does not change preferences over strategies but does change utility levels and hence the bound ε which measures differences of utility level. We avoid this by requiring a bound on the payoffs. We say that a game $[A, u]$ is *bounded by* M if

$$\left|u_i(a, \omega) - u_i(a', \omega)\right| \leq M, \text{ for all } a, a' \in A,\ i \in \mathcal{I}, \text{ and } \omega \in \Omega.$$

Lemma 3. *If game $[A, u]$ is bounded by M and σ_i is optimal given σ under P^∞, then σ_i is $2\delta M$-optimal given σ under P^k at all states in $\mathbf{A}_i^k(\delta)$.*

Proof. Suppose σ_i is optimal given σ under P^∞; that is, for all $a_i \in A_i$, $\omega \in \Omega$;

$$\sum_{\omega' \in Q_i(\omega)} \left(u_i\left(\sigma(\omega'), \omega'\right) - u_i\left(\left\{a_i, \sigma_{-i}(\omega')\right\}, \omega'\right)\right) P^\infty\left[\omega' \left| Q_i(\omega)\right.\right] \geq 0 \tag{2.4}$$

Fix arbitrary a_i, and $\omega \in \mathbf{A}_i^k(\delta)$. Let $\mathbf{A}_1 = \left\{\omega' \in \mathbf{A}_i^k(\delta) : P^k\left[\omega' \left| Q_i(\omega)\right.\right] - P^\infty\left[\omega' \left| Q_i(\omega)\right.\right] \geq 0\right\}$ and $\mathbf{A}_2 = \mathbf{A}_i^k(\delta) \setminus \mathbf{A}_1$. Then:

$$\begin{aligned}
& \sum_{\omega' \in Q_i(\omega)} \left(u_i\left(\sigma(\omega'), \omega'\right) - u_i\left(\left\{a_i, \sigma_{-i}(\omega')\right\}, \omega'\right)\right) P^k\left[\omega' \left| Q_i(\omega)\right.\right] \\
= & \sum_{\omega' \in Q_i(\omega)} \left(u_i\left(\sigma(\omega'), \omega'\right) - u_i\left(\left\{a_i, \sigma_{-i}(\omega')\right\}, \omega'\right)\right) P^\infty\left[\omega' \left| Q_i(\omega)\right.\right] \\
& + \sum_{\omega' \in Q_i(\omega)} \left(u_i\left(\sigma(\omega'), \omega'\right) - u_i\left(\left\{a_i, \sigma_{-i}(\omega')\right\}, \omega'\right)\right) \left\{P^k\left[\omega' | Q_i(\omega)\right] - P^\infty\left[\omega' \left| Q_i(\omega)\right.\right]\right\}
\end{aligned}$$

Now the first term in the expression above is non-negative by (2.4). Now for the second term, note that $\omega \in \mathbf{A}_i^k(\delta)$ implies $Q_i(\omega) \subset \mathbf{A}_i^k(\delta)$, so:

$$\begin{aligned}
& \sum_{\omega' \in Q_i(\omega)} \left(u_i\left(\sigma(\omega'), \omega'\right) - u_i\left(\left\{a_i, \sigma_{-i}(\omega')\right\}, \omega'\right)\right) \left\{P^k\left[\omega' \left| Q_i(\omega)\right.\right] - P^\infty\left[\omega' \left| Q_i(\omega)\right.\right]\right\} \\
\geq & - \sum_{\omega' \in Q_i(\omega) \cap \mathbf{A}_1} M \left\{P^k\left[\omega' | Q_i(\omega)\right] - P^\infty\left[\omega' | Q_i(\omega)\right]\right\} \\
& - \sum_{\omega' \in Q_i(\omega) \cap \mathbf{A}_2} M \left\{P^\infty\left[\omega' | Q_i(\omega)\right] - P^k\left[\omega' | Q_i(\omega)\right]\right\} \\
= & -M\left\{P^k\left[\mathbf{A}_1 | Q_i(\omega)\right] - P^\infty\left[\mathbf{A}_1 | Q_i(\omega)\right]\right\} - M\left\{P^\infty\left[\mathbf{A}_2 | Q_i(\omega)\right] - P^k\left[\mathbf{A}_2 | Q_i(\omega)\right]\right\} \\
\geq & -2\delta M
\end{aligned}$$

where the last inequality follows by the definition of $\mathbf{A}_i^k(\delta)$, and this proves the lemma. □

Lemma 3 is tight in the following sense.

Lemma 4. *Suppose $\omega \notin \mathbf{A}_i^k(\delta)$. Then there exists a game $[A,u]$ bounded by 1 and an equilibrium σ of $[A,u]$ under P^∞, such that $\sigma_i(\omega)$ is not ε-optimal at ω given σ under P^k for any $\varepsilon \leq \delta$.*

Proof. $\omega \notin \mathbf{A}_i^k(\delta)$ implies that there exists an event E with

$$P^k[E|Q_i(\omega)] - P^\infty[E|Q_i(\omega)] > \delta \tag{2.5}$$

Fix any $\varepsilon \leq \delta$. We will construct a simple decision problem, where i's optimal action at ω under P^∞ is not ε-optimal under P^k. A decision problem is a special case of a game where players' payoffs do not depend on other players' actions (thus we may omit mention of a_{-i} in i's decision). Consider a game $[A,u]$ with $A_i = \{\alpha, \beta\}$ and $u_i(\alpha, \omega') = 0$ for all $\omega' \in \Omega$, and

$$u_i(\beta,\omega') = \begin{cases} 1 - P^\infty[E|Q_i(\omega)], & \text{if } \omega' \in E \cap Q_i(\omega) \\ -P^\infty[E|Q_i(\omega')], & \text{if } \omega' \in (\Omega \backslash E) \cap Q_i(\omega) \\ 0, & \text{if } \omega' \notin Q_i(\omega) \end{cases}$$

For $j \neq i$, set $u_j \equiv 0$. The decision problem is bounded by 1. An optimal strategy for i under P^∞ is $\sigma_i(\omega') = \alpha$, for all $\omega' \in \Omega$, since both α and β give expected payoff 0 at every state. But at state ω under P^k, this is not ε optimal, since

$$\begin{aligned} \sum_{\omega' \in Q_i(\omega)} u_i(\beta,\omega')P^k[\omega'|Q_i(\omega)] &= \left\{ \begin{array}{l} P^k[E|Q_i(\omega)]\{1 - P^\infty[E|Q_i(\omega)]\} \\ -\left\{1 - P^k[E|Q_i(\omega)]\right\} P^\infty[E|Q_i(\omega)] \end{array} \right\} \\ &= P^k[E|Q_i(\omega)] - P^\infty[E|Q_i(\omega)] \\ &> \delta, \text{by (2.5)} \\ &\geq \varepsilon \end{aligned}$$

□

3. Continuity of Actions

It is useful to consider a rather strong notion of continuity of actions. This will demonstrate why uniform convergence matters. The action continuity theorem will also be independent of the number of players. It is therefore a useful benchmark for understanding what goes wrong with payoff continuity with more than one player.

Definition 5. *Sequence P^k satisfies Approximate Lower hemicontinuity of Actions [ALA] if for every $\varepsilon > 0$ and $M > 0$, there exists K such that the following property holds for all $k \geq K$:*

For any game $[A, u]$ bounded by M, if σ is an equilibrium of $[A, u]$ under P^∞ then σ is an ε-equilibrium of $[A, u]$ under P^k.

In particular, σ_i is required to be ε-optimal not merely with high probability, but with probability one. Bound M can be set to be, say, 1, without loss of generality. But notice the bound M is chosen independent of the choice of game, so that ALP is a concept of uniform convergence with respect to games. This is appropriate for us since we want to characterize the convergence properties of information systems independently of the choice of game.

Theorem 1. *Suppose $P^k \rightarrow P^\infty$ pointwise. Then P^k satisfies ALA if and only if $P^k \rightarrow P^\infty$ UCCP, i.e., there is uniform convergence of conditional probabilities.*

Proof. For any $\varepsilon > 0$ and $M > 0$, set $\delta = \varepsilon/2M$ and let σ be any equilibrium of the P^∞ game. If UCCP holds, then (by lemma 2) $\mathbf{A}_i^k(\delta) = \Omega$ for all i and all sufficiently large k. But now lemma 3 ensures that σ_i is ε-optimal for all i at all states.

To prove the converse, we must show that whenever UCCP fails, we can find some $\varepsilon > 0$ and $M > 0$, such that there are an infinite number of k for which there exists a game $[A^k, u^k]$ bounded by M, an equilibrium σ^k of $[A^k, u^k]$ under P^∞, such that $\sigma_i^k(\omega)$ is not ε-optimal at some state ω^k under P^k. But note that if uniform convergence fails, then (by lemma 2) there exists $\varepsilon > 0$ such that for infinitely many k, $\mathbf{A}_*^k(\varepsilon) \neq \Omega$. Setting $M = 1$ and choosing any ω^k and i^k such that $\omega^k \in \Omega \backslash A_{i^k}^k(\varepsilon)$, lemma 4 proves the existence of the required game $[A^k, u^k]$ for each k. □

4. Continuity of Payoffs

The following is the key continuity property we shall study.

Definition 6. *Sequence P^k satisfies Approximate Lower hemicontinuity of Payoffs [ALP] if for every $\varepsilon > 0$, $\eta > 0$, and $M > 0$, there exists K such that the following property holds for all $k \geq K$: For any game $[A, u]$, bounded by M, if σ is an equilibrium of $[A, u]$ under P^∞, then there is a strategy profile σ' which is an ε-equilibrium of $[A, u]$ under P^k, such that*

$$\underset{i \in \mathcal{I}}{Max} \left| \sum_{\omega \in \Omega} \left(u_i \left(\sigma'(\omega), \omega \right) - u_i \left(\sigma(\omega), \omega \right) \right) P^k [\omega] \right| \leq \eta \tag{4.1}$$

In words, ALP holds if for any equilibrium σ under P^∞, there exists an ε-equilibrium σ' under P^k such that *ex ante* difference in utility level between σ and σ' *under* P^k is close for all individuals.

Remarks.

1. The closeness is *uniform* in games; like ALA, ALP describes uniform convergence with respect to games.

2. Since the difference is measured with respect to ex ante utility level, we can replace P^k with P^∞ as long as we are concerned with a pointwise convergent sequence.

3. Since the Nash equilibrium correspondence is not lower hemicontinuous with respect to payoffs, we must restrict ourselves to an approximate notion of continuity (i.e. allowing for ε-equilibria). Consider the following example. Suppose with probability $1 - \varepsilon$, payoffs are given by:

	L	R
U	1, 1	0, 0
D	0, 0	0, 0

With probability ε, payoffs are given by:

	L	R
U	1, 1	1, 0
D	0, 1	0, 0

Now consider the incomplete information game where neither player has any information. When $\varepsilon = 0$, there is an equilibrium where players choose (D, R). But for any $\varepsilon > 0$, the unique equilibrium has both players choosing (U, L). Thus there is no hope of getting an exact lower hemicontinuity result. Presumably we could replace our main result with an exact lower hemicontinuity result for generic payoffs.

4.1. Single Person Decision Problems

To understand when ALP obtains, let us start with single player games. Single player games are of course simple decision problems. In single player games, we naturally omit reference to σ_{-i}.

Theorem 2. *Suppose $P^k \to P^\infty$ pointwise and $I = 1$. Then ALP holds.*

Proof. Fix any $\varepsilon > 0$, $\eta > 0$, and $M > 0$. By lemma 3, for any game bounded by M, and strategy σ_1 optimal under P^∞, σ_1 is ε-optimal under P^k at all states in $\mathbf{A}_1^k(\frac{\varepsilon}{2M})$. By lemma 1, there exists K such that $P^k\left[\mathbf{A}_1^k\left(\frac{\varepsilon}{2M}\right)\right] > 1 - \frac{\eta}{M}$ and thus $P^k\left[\Omega \setminus \mathbf{A}_1^k\left(\frac{\varepsilon}{2M}\right)\right] < \frac{\eta}{M}$, for all $k \geq K$. Fix any such k and let $\sigma_1'(\omega) = \sigma_1(\omega)$ for all $\omega \in \mathbf{A}_1^k(\frac{\varepsilon}{2M})$, and any optimal strategy otherwise. Now σ' is ε-optimal under P^k and

$$\left|\sum_{\omega\in\Omega}\left(u_1\left(\sigma_1'(\omega),\omega\right) - u_1\left(\sigma_1(\omega),\omega\right)\right)P^k[\omega]\right| \leq P^k\left[\Omega \setminus \mathbf{A}_1^k\left(\frac{\varepsilon}{2M}\right)\right]\cdot M < \eta$$

□

This result is meant to be just an illustration of the potential problem. In fact, the result follows immediately from the theorem of maximum where σ is the choice variable and P is exogenous parameter, since the set of strategies is compact (with respect to pointwise convergence) and the utility function is continuous in σ and P with pointwise convergence. In particular, it also holds when $\varepsilon = 0$. If $I > 1$, this analogy fails since σ_{-1} is *endogenously* determined, which indicates that the many person case is more delicate.

4.2. Many Person Counterexample

We shall start with an example which shows that pointwise convergence is not sufficient for ALP in many player games. Suppose there are two players, Row (R) and Column (C). Each observes a non-negative integer, with probability under P^k given by the following table (where R observes the integer on the row, and C observes the column) where α^k is a constant, $\alpha^k \in (0,1)$:-

	0	1	2	3	..	$n-1$	n	$\cdots$
0	$\frac{1}{2}$	$\frac{1}{4}\alpha^k$	$\frac{1}{8}\alpha^k$	$\frac{1}{16}\alpha^k$		$\left(\frac{1}{2}\right)^n \alpha^k$	$\left(\frac{1}{2}\right)^{n+1} \alpha^k$	
1	0	$\frac{1}{4}(1-\alpha^k)$	$\frac{1}{8}(1-\alpha^k)$	0		0	0	
2	0	0	$\frac{1}{16}(1-\alpha^k)$	$\frac{1}{32}(1-\alpha^k)$		0	0	
3	0	0	0	$\frac{1}{64}(1-\alpha^k)$		0	0	
.								
$n-1$	0	0	0	0		$\left(\frac{1}{2}\right)^{2n-2}(1-\alpha^k)$	$\left(\frac{1}{2}\right)^{2n-1}(1-\alpha^k)$	
n	0	0	0	0		0	$\left(\frac{1}{2}\right)^{2n}(1-\alpha^k)$	
$\vdots$								

As $k \to \infty$, $\alpha^k \to 0$, so that we have a limit distribution P^∞:

	0	1	2	3	..	$n-1$	n	$\cdots$
0	$\frac{1}{2}$	0	0	0		0	0	
1	0	$\frac{1}{4}$	$\frac{1}{8}$	0		0	0	
2	0	0	$\frac{1}{16}$	$\frac{1}{32}$		0	0	
3	0	0	0	$\frac{1}{64}$		0	0	
$\vdots$								
$n-1$	0	0	0	0		$\left(\frac{1}{2}\right)^{2n-2}$	$\left(\frac{1}{2}\right)^{2n-1}$	
n	0	0	0	0		0	$\left(\frac{1}{2}\right)^{2n}$	
$\vdots$								

The example is an elaboration of the coordinated attack problem in the computer science literature (see Halpern (1986)) and the electronic mail game of Rubinstein (1989) [although we do not have an elegant interpretation of the information structure in terms of messages lost!]. This information system can be formally represented as follows:

Example 1. $\mathcal{I} = \{R, C\}$; $\Omega = \aleph_+^2$; $Q_i((n_R, n_C)) = \{(n'_R, n'_C) : n'_i = n_i\}$; $P^k[(0,0)] = \frac{1}{2}$; $P^k[(0,n)] = \left(\frac{1}{2}\right)^{n+1}\alpha^k$, $P^k[(n,n)] = \left(\frac{1}{2}\right)^{2n}(1-\alpha^k)$, $P^k[(n,n+1)] = \left(\frac{1}{2}\right)^{2n+1}(1-\alpha^k)$, *for all*

$n \geq 1$; $P^k[\omega] = 0$, *otherwise; where* $\alpha^k \to 0$; $P^\infty[(0,0)] = \frac{1}{2}$; $P^\infty[(n,n)] = \left(\frac{1}{2}\right)^{2n}$, $P^\infty[(n,n+1)] = \left(\frac{1}{2}\right)^{2n+1}$; $P^\infty[\omega] = 0$, *otherwise.*

Note that $P^k[\omega] \to P^\infty[\omega]$ uniformly with respect to Ω and $P^k[\omega|Q_i(\omega)] \to P^\infty[\omega|Q_i(\omega)]$ pointwise. But uniform convergence of conditional probabilities fails. For example, writing $h(n,k) \equiv P^k[(n,n)|Q_C((n,n))]$, we have (for $n \geq 2$)

$$h(n,k) = \frac{\left(\frac{1}{2}\right)^{2n}(1-\alpha^k)}{\left(\frac{1}{2}\right)^{2n}(1-\alpha^k) + \left(\frac{1}{2}\right)^{2n-1}(1-\alpha^k) + \left(\frac{1}{2}\right)^{n+1}\alpha^k}$$

Now as $k \to \infty$, $\alpha^k \to 0$ and so $h(n,k) \to \frac{1}{3}$ (for any fixed $n \geq 2$). But as $n \to \infty$, $h(n,k) \to 0$ (for any fixed k).

The ALP property fails in this example. Consider the following game. Each player has two possible actions $A_i = \{D, S\}$ (D for dangerous, S for safe). When R observes a message other than 0, payoffs are given by

	D	S
D	$2,2$	$-10,1$
S	$1,-10$	$1,1$

(4.2)

On the other hand, when R observes 0, payoffs are given by

	D	S
D	$-10,-10$	$-10,1$
S	$1,-10$	$1,1$

(4.3)

There is an equilibrium under P^∞ with $\sigma_i((n_R,n_C)) = S$, if $n_i = 0$; $\sigma_i((n_R,n_C)) = D$, if $n_i > 0$ (for $i = R$ and C), which yields ex ante utility $\frac{1}{2}\left(1-\alpha^k\right)(2) + \frac{1}{2}\left(1+\alpha^k\right)(1) = \frac{3}{2} - \frac{\alpha^k}{2}$ for player R under P^k, and which tends to $\frac{3}{2}$ as $k \to \infty$. But as we argue below, if ε is less than $\frac{5}{9}$, any ε-equilibrium has the property that D is chosen with less than probability $\frac{1}{3}$ at every state by both players. So any ε-equilibrium yields at most $\frac{4}{3}$ for each player (if payoffs are given by (4.2), then each player gets at most $\frac{1}{3}(2) + \frac{2}{3}(1) = \frac{4}{3}$; if payoffs are given by (4.2), then each player gets at most 1). Since $\frac{3}{2} > \frac{4}{3}$, ALP fails.

Let $\varepsilon < \frac{5}{9}$ and σ be any ε-equilibrium strategy for P^k.

1. First we claim that at any state $\omega = (n_R, n_C)$ for each player i, if the conditional probability of the opponent choosing D is less than $\frac{7}{9}$, then player i's expected payoff to choosing D is at most $-\frac{2}{3}$; that is, if $\sum_{\omega \in Q_i((n_R,n_C))} \sigma_j(D|\omega) P^k[\omega|Q_i((n_R,n_C))] < \frac{7}{9}$, then $\sum_{\omega \in Q_i((n_R,n_C))} u_i(D, \sigma_j(\omega)) P^k[\omega|Q_i((n_R,n_C))] \leq -\frac{2}{3}$. This is obviously true if payoffs are given by (4.3); if payoffs are given by (4.2), then $\frac{7}{9}(2)+\frac{2}{9}(-10) = -\frac{2}{3}$ is the most that can be hoped for. On the other hand, at any state (n_R, n_C) for each player i, choosing S yields 1 for sure. So if $\sum_{\omega \in Q_i((n_R,n_C))} \sigma_j(D|\omega) P^k[\omega|Q_i((n_R,n_C))] < \frac{7}{9}$, then the pure strategy D is not $\frac{5}{3}$−optimal against σ_j at (n_R, n_C), hence by fact 2, $\sigma_i(D|(n_R,n_C)) < \frac{\varepsilon}{5/3} < \frac{5/9}{5/3} = \frac{1}{3}$.

2. Now suppose that player i assigns probability at least $\frac{1}{3}$ to his opponent player j choosing action D with probability less than $\frac{1}{3}$; that is, $P^k\left[\left\{\omega : \sigma_j(D|\omega) < \frac{1}{3}\right\}\middle| Q_i(\omega)\right] \geq \frac{1}{3}$. Then he assigns probability at most $\frac{7}{9}$ $(= \frac{2}{3} + \frac{1}{3}(\frac{1}{3}))$ to his opponent choosing D. Thus (by step 1) he chooses D with probability at most $\frac{1}{3}$.

3. If C observes 0, she knows that R has observed 0 and thus that payoffs are given by (4.3). Thus for either player, the pure strategy D is not 11−optimal and so $\sigma_i(D|(0,n_j)) \leq \frac{\varepsilon}{11} < \frac{1}{3}$ by fact 2, for each $j \neq i$.

4. Notice that $P^k[(0,n_C)|Q_C((n_R,n_C))] \uparrow 1$ for all n_R, as $n_C \to \infty$. Fix N such that $P^k[(0,n_C)|Q_C((n_R,n_C))] > \frac{1}{3}$ for all $n_C > N$ and all n_R.

5. By steps 2,3, and 4, $\sigma_C(D|(n_R,n_C)) < \frac{1}{3}$ for all $n_C \geq N$, and for all n_R.

6. Note that when R observes $n_R \neq 0$, he assigns probability $\frac{1}{3}$ to the event of C observing n_C greater than n_R. Thus by step 5, if $n_R \geq N-1$, $P^k\left[\left\{\sigma_C(D|\omega) \leq \frac{1}{3}\right\}\middle| Q_R(\omega)\right] \geq \frac{1}{3}$. So, again by steps 2 and 4, we have $\sigma_R(D|(n_R,n_C)) < \frac{1}{3}$ for all $n_R \geq N-1$, and for all n_C.

7. Note that when C observes $n_C \neq 0$, he assigns probability at least $\frac{1}{3}$ to R observing either $n_R = 0$ or $n_R = n_C$. Thus by steps 3 and 6, if $n_C \geq N-1$, $P^k\left[\left\{\sigma_R(D|\omega) \leq \frac{1}{3}\right\}\middle| Q_C(\omega)\right] \geq \frac{1}{3}$. So, again by steps 2 and 4, we have $\sigma_C(D|(n_R,n_C)) < \frac{1}{3}$ for all $n_C \geq N-1$, and for all n_R.

8. The argument of steps 6 and 7 can be iterated to ensure that $\sigma_i(D|(n_R,n_C)) < \frac{1}{3}$ for every (n_R, n_C) for both $i = C, R$.

Morris, Rob and Shin (1993) examine more general forms of this type of "infection" argument. However, it is not the case that any failure of uniform convergence of conditional probabilities implies a failure of ALP. We know from theorem 2 that if there is one individual and pointwise convergence, then ALP must hold even when UCCP fails. Consider the following example. Suppose now there is only one individual, Column, who again observes a non-negative integer. But she does not observe another, binary, signal which is either *Up* or *Down*. Suppose P^k is given by

	0	1	2	.. ..	n
U	$\frac{1}{2}\alpha^k$	$\frac{1}{2}(1-\alpha^k)(1-r)$	$\frac{1}{2}(1-\alpha^k)(1-r)r$		$\frac{1}{2}(1-\alpha^k)(1-r)r^{n-1}$
D	$\frac{1}{2}(1-\alpha^k)$	$\frac{1}{2}\alpha^k(1-q)$	$\frac{1}{2}\alpha^k(1-q)q$		$\frac{1}{2}\alpha^k(1-q)q^{n-1}$

where $1 > q > r > 0$ and $0 < \alpha < 1$. P^k converges pointwise to a limit distribution P^∞:

	0	1	2	.. ..	n
U	0	$\frac{1}{2}(1-r)$	$\frac{1}{2}(1-r)r$		$\frac{1}{2}(1-r)r^{n-1}$
D	$\frac{1}{2}$	0	0		0

This can be represented formally as follows:

Example 2. *$\mathcal{I} = \{C\}$. $\Omega = \{U,D\} \times \aleph_+$. $Q_C((S,n)) = \{(U,n),(D,n)\}$ for all $(S,n) \in \Omega$. Let $P^k[(U,0)] = \frac{1}{2}\alpha^k$ and $P^k[(D,0)] = \frac{1}{2}(1-\alpha^k)$; let $P^k[(U,n)] = \frac{1}{2}(1-\alpha^k)(1-r)r^{n-1}$ and $P^k[(D,n)] = \frac{1}{2}\alpha^k(1-q)q^{n-1}$ for all $n \geq 1$, for some $1 > q > r > 0$; where $\alpha^k \to 0$ as $k \to \infty$; let $P^\infty[(U,n)] = \frac{1}{2}(1-r)r^{n-1}$, for all $n \geq 1$, $P^\infty[(D,0)] = \frac{1}{2}$, and $P^\infty[\omega] = 0$ otherwise.*

Again $P^k[\omega] \to P^\infty[\omega]$ uniformly with respect to Ω and $P^k[\omega|Q_i(\omega)] \to P^\infty[\omega|Q_i(\omega)]$ pointwise, but not uniformly: $P^k[(U,n)|Q_C((U,n))] = \frac{\frac{1}{2}(1-\alpha^k)(1-r)r^{n-1}}{\frac{1}{2}(1-\alpha^k)(1-r)r^{n-1}+\frac{1}{2}\alpha^k(1-q)q^{n-1}} \to 0$ (as $n \to \infty$), even though $P^\infty[(U,n)|Q_C((U,n))] = 1$, for all $n \geq 1$.

The sequence P^k of example 2 fails uniform convergence of conditional probabilities but satisfies ALP by theorem 2. So the ALP property is weaker than uniform convergence of conditional probability, but stronger than pointwise convergence. We shall characterize ALP in the next section, which is the main result of the paper, but it is useful to ask why the argument of Theorem 2 does not work with many players.

Say $I = 2$. By lemma 3, for any equilibrium σ under P^∞, σ_i is $\varepsilon \equiv 2\delta M$ optimal given σ under P^k at $\omega \in \mathbf{A}_i^k(\delta)$ for both i. We would like to find σ_1' that coincides with σ_1 on $\mathbf{A}_1^k(\delta)$, and is optimal against σ_2 outside $\mathbf{A}_1^k(\delta)$. Intuitively, if the probability $P^k\left[\mathbf{A}_*^k(\delta)\right]$ were sufficiently high, this construction would yield a ε-equilibrium. But since $\mathbf{A}_1^k(\delta) \neq \mathbf{A}_2^k(\delta)$ in general, σ_2 need not be $\varepsilon-$optimal against σ_1'. But if the conditional probability of the event $\mathbf{A}_1^k(\delta)$ is close enough to one at ω, then σ_2 will be $2\varepsilon-$optimal against σ_1' since $\sigma_1 = \sigma_1'$ on $\mathbf{A}_1^k(\delta)$. A symmetric argument works for player 2. Thus we would require not only that the prior probability of $\mathbf{A}_*^k(\delta)$ be high, but also that the conditional probability each individual attaches to the event $\mathbf{A}_*^k(\delta)$ is high at all $\omega \in \mathbf{A}_i^k(\delta)$. Recall that UCCP implies that the conditional probability of $\mathbf{A}_*^k(\delta)$ is 1 for large enough k (lemma 2), but clearly this is not necessary.

An alternative construction would be the following. Suppose there existed a high probability event $E \subset \mathbf{A}_*^k(\varepsilon)$ such that at all states in E, 1 assigned high probability to both E and $\mathbf{A}_*^k(\varepsilon)$. Then setting $\sigma_i' = \sigma_i$ on event E and anything optimal outside E would generate a $2\varepsilon-$equilibrium by the same argument as above. But for readers familiar with Monderer and Samet's (1989) notion of common p-belief, the existence of such an event E is equivalent to the requirement that, for some p close to one, the event $\mathbf{A}_*^k(\delta)$ is common $p-$belief with high probability. Thus we shall first give a brief review the idea of common $p-$belief.

4.3. Belief Operators

Define belief operators on events as follows. Let

$$B_i^{p,k}(E) = \{\omega : P^k[E|Q_i(\omega)] \geq p\}$$

$B_i^{p,k}(E)$ is the set of states where, given prior probability P^k, individual i believes event E with probability at least p. Note that by construction, whenever E is Q_i measurable, we have:

$$B_i^{p,k}(E \cap F) = E \cap B_i^{p,k}(F) \tag{4.4}$$

4.3.1. Common p-Belief

Define an "everyone believes" operator by

$$B_*^{p,k}(E) = \bigcap_{i \in \mathcal{I}} B_i^{p,k}(E)$$

From (4.4) we have:

$$B_*^{p,k}(E) \subset E \tag{4.5}$$

if $E = \bigcap_{i \in \mathcal{I}} E_i$, where each E_i is Q_i measurable.

Also define a "common p-belief" operator by

$$C^{p,k}(E) = \bigcap_{n \geq 1} \left[B_*^{p,k}\right]^n (E)$$

Definition 7. *An event E is common p-belief under P^k at state ω if $\omega \in C^{p,k}(E)$.*

From (4.5) and by the definition of $C^{p,k}$, it follows that:

$$C^{p,k}(E) = B_*^{p,k}\left(C^{p,k}(E)\right) \equiv \bigcap_{i \in \mathcal{I}} B_i^{p,k}(C^{p,k}(E)) \tag{4.6}$$

Monderer and Samet (1989) used (4.6) to show a relation between the iterative notion of common p-belief and the following fixed point notion. It thus generalizes the well-known connection between iterative and fixed point notions of *knowledge* (Aumann (1976)).

Definition 8. *Event E is p-evident under P^k if $E \subset B_*^{p,k}(E)$.*

Note that in particular, $C^{p,k}(E)$ is p-evident.

Theorem 3 (Monderer and Samet (1989)). *Event E is common p-belief under P^k at state ω if and only if there is an event F, p-evident under P^k, with $\omega \in F \subset B_*^{p,k}(E)$.*

4.4. Almost Uniform Convergence

The discussion at the end of the section 4.2 suggests the following weakening of uniform convergence:

Definition 9. *(Almost Uniform Convergence of Conditional Probabilities [AUCCP]) Conditional probabilities converge almost uniformly, if for each $\delta > 0$ and $p \in [0,1)$, there exists K such that for all $k \geq K$,*

$$P^k\left[C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right] \geq p \tag{4.7}$$

We write $P^k \to P^\infty$ AUCCP.

Example 1 failed this property. For any fixed $\delta > 0$, for any $p < 1$, and k sufficiently large, $\mathbf{A}_R^k(\delta) = \Omega$ and $\mathbf{A}_C^k(\delta)$ is of the form $\{(n_R, n_C) : n_C \leq N\}$; now for any p close to 1, $B_*^{p,k}\left(\mathbf{A}_*^k(\delta)\right) = \{(n_R, n_C) : Max\{n_R, n_C\} \leq N-1\}$; $\left[B_*^{p,k}\right]^2\left(\mathbf{A}_*^k(\delta)\right) = \{(n_R, n_C) : Max\{n_R, n_C\} \leq N-2\}$; $\left[B_*^{p,k}\right]^n\left(\mathbf{A}_*^k(\delta)\right) = \{(n_R, n_C) : Max\{n_R, n_C\} \leq N-n\}$; and so $C^{p,k}\left(\mathbf{A}_*^k(\delta)\right) = \emptyset$.

But example 2 satisfied this property since $C^{p,k}[\mathbf{A}_*^k(\delta)] = B_C^{p,k}\left(\mathbf{A}_C^k(\delta)\right) = \mathbf{A}_C^k(\delta) = \left\{(S,n) : n \geq N^k(\delta)\right\}$, where $N^k(\delta) \to \infty$ as $k \to \infty$; thus $\lim_{k\to\infty} P^k\left[C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right] = \lim_{k\to\infty} P^k\left[\mathbf{A}_C^k(\delta)\right] = 1$.

4.5. Main Result

Theorem 4. *Suppose $P^k \to P^\infty$ pointwise. Then P^k satisfies ALP if and only if $P^k \to P^\infty$ AUCCP.*

Proof. Suppose AUCCP holds. Fix $\delta > 0$ and $p \in [0,1)$ arbitrarily. Choose k such that $P^k\left[C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right] \geq p$. Write $F = C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)$ and $F_i = B_i^{p,k}(F)$. By (4.6), $F = \bigcap_{i\in\mathcal{I}} F_i$ and so $F \subset F_i$; by (4.5) $F \subset \mathbf{A}_*^k(\delta)$. So the following hold:

$$P^k\left[F\,|Q_i(\omega)\right] \geq p \text{ if } \omega \in F_i \tag{4.8}$$

$$P^k\left[(F_i \setminus F)\,|Q_i(\omega)\right] \leq 1-p \text{ if } \omega \in F_i \tag{4.9}$$

(4.8) is true by definition of $B_i^{p,k}(F)$; (4.9) is immediate from (4.4), (4.8) and $F = \bigcap_{i\in\mathcal{I}} F_i$.

Fix any game bounded by M, and let σ be any equilibrium of the game under P^∞. Now consider any strategy profile σ' with $\sigma_i'(\omega) = \sigma_i(\omega)$ for all i and $\omega \in F_i$. Pick any $\omega \in F$. Since $F \subset \mathbf{A}_*^k(\delta)$, by Lemma 3, σ_i' is $2\delta M$ optimal against σ at any $\omega \in F$; that is,

$$\sum_{\omega'\in Q_i(\omega)}\left(u_i\left(\left\{\sigma_i'(\omega), \sigma_{-i}(\omega')\right\}, \omega\right) - u_i\left(\left\{a_i, \sigma_{-i}(\omega')\right\}, \omega\right)\right) P^k\left[\omega'|Q_i(\omega)\right] \geq -2\delta M \text{ for any } a_i \in A_i \tag{4.10}$$

Given $\omega \in F$, $\sigma_{-i}' \neq \sigma_{-i}$ occurs with probability less than $1-p$ by (4.9). Since the game is bounded by M, we have

$$\sum_{\omega'\in Q_i(\omega)}\left|u_i\left(\left\{\sigma_i'(\omega), \sigma_{-i}(\omega')\right\}, \omega\right) - u_i\left(\left\{\sigma_i'(\omega), \sigma_{-i}'(\omega')\right\}, \omega\right)\right| P^k\left[\omega'|Q_i(\omega)\right] \leq (1-p)\,M \tag{4.11}$$

(4.10) and (4.11) imply that σ_i' is $2\delta M + (1-p)M$ optimal against σ' at any $\omega \in F_i$.

Now consider the constrained game (under P^k) where each player i is exogenously required to choose $\sigma_i(\omega)$ on F_i, but is free to choose any action outside F_i. An equilibrium σ^* of the modified game exists by fact 3. By the argument above, σ^* is an $(2\delta M + (1-p)M)$-equilibrium of the actual game. Moreover, $\sigma^* = \sigma$ on F, which occurs with probability more than p under P^k by construction and (4.8), hence the ex ante utility difference is at most $(1-p)M$. So we have ALP by choosing δ sufficiently small and p sufficiently close to one.

To establish the converse, we must show that if AUCCP fails, there exist $\varepsilon > 0$, $\eta > 0$ and $M > 0$ such that for infinitely many k, it is possible to construct a game $[A^k, u^k]$ bounded by M and an equilibrium σ^k of $[A^k, u^k]$ under P^∞, such that for every ε-equilibrium, σ', of $[A^k, u^k]$ under P^k, $\underset{i\in\mathcal{I}}{Max}\left|\sum_{\omega\in\Omega}\left(u_i(\sigma'(\omega),\omega) - u_i\left(\sigma^k(\omega),\omega\right)\right)P^k[\omega]\right| > \eta$.

Suppose $P^k \to P^\infty$ AUCCP does not hold; that is, there exist $\delta > 0$, $p < 1$ such that for infinitely many k,

$$P^k\left[C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right] < p$$

Then, by (4.6), for infinitely many k, there exists an individual $f(k) \in \mathcal{I}$ such that

$$P^k\left[B_{f(k)}^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)\right] < 1 - \frac{1-p}{I}$$

On the other hand, by lemma 1, there exists K such that for every $k > K$,

$$P^k\left[\Omega\setminus\mathbf{A}_{f(k)}^k(\delta)\right] < \frac{1-p}{2I}$$

From these two expressions, we have that for infinite many k,

$$P^k\left[B_{f(k)}^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)\cup\left(\Omega\setminus\mathbf{A}_{f(k)}^k(\delta)\right)\right] < 1 - \frac{1-p}{I} + \frac{1-p}{2I} = 1 - \frac{1-p}{2I} \tag{4.12}$$

Thus

$$P^k\left[\left(\Omega\setminus B_{f(k)}^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)\right)\cap\mathbf{A}_{f(k)}^k(\delta)\right] > \frac{1-p}{2I} \tag{4.13}$$

Fix any such k. We shall construct a game with the required property for this k. By definition of $\mathbf{A}_i^k(\delta)$, we can construct a set valued function $H_i : \Omega\setminus\mathbf{A}_i^k(\delta) \to 2^\Omega$ such that for all ω, we have (i) $H_i(\omega') = H_i(\omega)$ for all $\omega' \in Q_i(\omega)$; (ii) $H_i(\omega) \subset Q_i(\omega)$ and (iii) $P^k[H_i(\omega)|Q_i(\omega)] - P^\infty[H_i(\omega)|Q_i(\omega)] > \delta$.

Now consider the following game $[A, u]$. For each player i, let $A_i = \{x_i, z_i\}$. Define payoffs as follows:

$$u_i(\{x_i, a_{-i}\}, \omega) = \begin{cases} 0, & \text{if } a_{-i} = x_{-i}, \text{ for all } \omega \\ -4, & \text{if } a_{-i} \neq x_{-i}, \text{ for all } \omega \end{cases}$$

$$u_i(\{z_i, a_{-i}\}, \omega) = \begin{cases} P^\infty[(\Omega \setminus H_i(\omega)) | Q_i(\omega)] & ; \text{if } \omega \in \Omega \setminus \mathbf{A}_i^k(\delta) \text{ and } \omega \in H_i(\omega), \text{ for all } a_{-i} \\ -P^\infty[H_i(\omega) | Q_i(\omega)] & ; \text{if } \omega \in \Omega \setminus \mathbf{A}_i^k(\delta) \text{ and } \omega \notin H_i(\omega), \text{ for all } a_{-i} \\ p - 1, & ; \text{otherwise.} \end{cases}$$

Payoffs are bounded by 5, and the payoff to playing x_i is -4 times the conditional probability that there is at least one player who does not play x_{-i}. Notice that the payoff to playing z_i does not depend on the others' strategies at all. If $\omega \in \Omega \setminus \mathbf{A}_i^k(\delta)$, the payoff to z_i in the P^k game is:

$$\begin{aligned} \sum_{\omega' \in Q_i(\omega)} u_i(\{z_i, \sigma_{-i}(\omega')\}, \omega') P^k[\omega' | Q_i(\omega)] &= \begin{cases} P^k[H_i(\omega) | Q_i(\omega)] P^\infty[(\Omega \setminus H_i(\omega)) | Q_i(\omega)] \\ -P^k[(\Omega \setminus H_i(\omega)) | Q_i(\omega)] P^\infty[H_i(\omega) | Q_i(\omega)] \end{cases} \\ &= P^k[H_i(\omega) | Q_i(\omega)] - P^\infty[H_i(\omega) | Q_i(\omega)] \\ &> \delta \end{aligned}$$

In the P^∞ game, the payoff to z_i is 0.

Claim 1. *Every i choosing x_i at every ω is an equilibrium of the P^∞ game.*

This is clear since given that everyone else plays x_i, payoffs under P^∞ are 0 for x_i, but z_i yields $(p-1) < 0$ on $\mathbf{A}_i^k(\delta)$ and 0 on $\Omega \setminus \mathbf{A}_i^k(\delta)$.

Claim 2. *Fix any $\varepsilon < Min[\frac{\delta}{2}, \frac{1-p}{2}]$, and let σ be any ε-equilibrium of $[A, u]$. Then at all states in $\Omega \setminus B_i^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)$, player i chooses x_i with probability no greater than $\frac{1}{2}$; that is, if $\omega \in \Omega \setminus B_i^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)$, then $\sigma_i(x_i|\omega) \leq \frac{1}{2}$.*

We show this in several steps. Let σ be any ε-equilibrium under P^k.

Step 1. Choose any $\omega \in \Omega \backslash \mathbf{A}_i^k(\delta)$. Then under σ, the pure strategy of choosing x_i is not δ-optimal at ω for i, since choosing x_i yields non-positive utility where as the payoff to choosing z_i is $P^k[H_i(\omega)|Q_i(\omega)] - P^\infty[H_i(\omega)|Q_i(\omega)] > \delta$. Thus by fact 2, $\sigma(x_i|\omega) \leq \frac{\varepsilon}{\delta} < \frac{1}{2}$. Thus if $\omega \in \Omega \backslash \mathbf{A}_*^k(\delta)$, $\omega \in \Omega \backslash \mathbf{A}_j^k(\delta)$ for some j and $\sigma_j(x_j|\omega) < \frac{1}{2}$ for that j.

Step 2. Choose any $\omega \in \Omega \backslash B_i^{p,k}(\mathbf{A}_*^k(\delta))$; by definition of $B_*^{p,k}$, i assigns conditional probability less than p to the set $\mathbf{A}_*^k(\delta)$ and thus more than $1-p$ to $\Omega \backslash \mathbf{A}_*^k(\delta)$. By step 1, for i, at ω, the conditional probability of the event where some j plays x_j under σ is less than $\frac{1}{2}$ if $\Omega \backslash \mathbf{A}_*^k(\delta)$ occurs, so at ω, i's conditional probability of the event where some j does not play x_j under σ is more than $\frac{1}{2}(1-p)$. So the pure strategy of choosing x_i yields at most $-4 \cdot \frac{1}{2}(1-p) = -2(1-p)$, whereas the payoff to choosing z_i yields at least $-(1-p)$. This shows that x_i is not $(1-p)$-optimal at ω, so by fact 2, $\sigma_i(x_i|\omega) \leq \frac{\varepsilon}{1-p} < \frac{1}{2}$. But now for any $\omega \in \Omega \backslash B_*^{p,k}(\mathbf{A}_*^k(\delta))$, $\omega \in \Omega \backslash B_*^{p,k}\left(\mathbf{A}_j^k(\delta)\right)$ for some j and $\sigma_j(x_j|\omega) < \frac{1}{2}$ for than j.

Step 3. The claim is proved by iterating the argument of step 2. Consider the set $\Omega \Big\backslash \left[B_*^{p,k}\right]^2\left(\mathbf{A}_*^k(\delta)\right) \equiv \Omega \Big\backslash B_*^{p,k}\left(B_*^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)$. Replacing $\mathbf{A}_*^k(\delta)$ with $B_*^{p,k}(\mathbf{A}_*^k(\delta))$ in Step 2, we find that for all $\omega \in \Omega \Big\backslash B_*^{p,k}\left(B_*^{p,k}(\mathbf{A}_*^k(\delta))\right)$, $\sigma_j(x_j|\omega) < \frac{1}{2}$ for some j. Iterating, we have for all n, for all $\omega \in \Omega \Big\backslash \left[B_*^{p,k}\right]^n\left(\mathbf{A}_*^k(\delta)\right)$, $\sigma_j(x_j|\omega) < \frac{1}{2}$ for some j. Since $C^{p,k}\left(\mathbf{A}_*^k(\delta)\right) = \bigcap_{n\geq 1}\left[B_*^{p,k}\right]^n\left(\mathbf{A}_*^k(\delta)\right)$, we have that for all $\omega \in \Omega \Big\backslash C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)$, $\sigma_j(x_j|\omega) < \frac{1}{2}$ for some j. Finally, $\omega \in \Omega \Big\backslash B_i^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)$ implies $\sigma_i(x_i|\omega) < \frac{1}{2}$.

Let σ be any ε -equilibrium. If $\omega \in \left(\Omega \Big\backslash B_{f(k)}^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)\right) \cap \mathbf{A}_{f(k)}^k(\delta)$, by claim 2, $\sigma_{f(k)}\left(x_{f(k)}|\omega\right) < \frac{1}{2}$. So $\sum_{\omega' \in Q_{f(k)}(\omega)} u_{f(k)}(\sigma(\omega'),\omega') P^k\left[\omega'\left|Q_{f(k)}(\omega)\right.\right] \leq -\frac{1-p}{2}$. But now (writing x for the strategy profile where each player i always chooses x_i) we have

$$\begin{aligned} \underset{i\in\mathcal{I}}{Max}\left|\sum_{\omega\in\Omega}\left(u_i(\sigma(\omega),\omega) - u_i(x,\omega)\right)P^k[\omega]\right| &\geq \left|\sum_{\omega\in\Omega}\left(u_{f(k)}(\sigma(\omega'),\omega) - u_{f(k)}(x,\omega)\right)P^k[\omega]\right| \\ &> \frac{1-p}{2}P^k\left[\left(\Omega \Big\backslash B_{f(k)}^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)\right) \cap \mathbf{A}_{f(k)}^k(\delta)\right] \\ &\geq \frac{(1-p)^2}{4I} \text{ by (4.13)} \end{aligned}$$

So by setting $\varepsilon < Min[\frac{\delta}{2}, \frac{1-p}{2}]$, $M = 5$ and $\eta < \frac{(1-p)^2}{4I}$, we have shown that ALP fails. □

5. Discussion

5.1. Almost uniform convergence

How strong is almost uniform convergence? The strength comes from requiring both uniformity of convergence and requiring it not only for prior probabilities but also for conditional probabilities. A slight weakening comes from not requiring fully uniform convergence but only that the set of well-behaved states are common p-belief for some p close to 1, with high ex ante probability. The "almost common knowledge" literature, and example 1, suggest that common p-belief is sometimes a very stringent requirement.

There are circumstances where almost uniform convergence of conditional probabilities is equivalent to pointwise convergence of prior probabilities. Since pointwise convergence implies uniform convergence of conditional probabilities on a finite state space, we have:

Corollary 1. *If $P^k \to P^\infty$ pointwise, and Ω is finite, then UCCP holds and thus AUCCP, ALA and ALP hold.*

Fudenberg and Tirole (1992, page 567) show a closely related result: for a finite state space, if $P^k[E] \to 1$ as $k \to \infty$ and payoffs are always common knowledge in E, then a version of ALP holds. Our result makes clear why finiteness is key in such results: finiteness implies that all convergence is uniform with respect to states.

We know by theorems 2 and 4 that if there is only one individual, almost uniform convergence of conditional probabilities must be equivalent to pointwise convergence of prior probabilities. But let us briefly show this directly.

Lemma 5. *If $P^k \to P^\infty$ and $I = 1$, then AUCCP and thus ALP hold.*

Proof. Fix $\delta > 0$ and $p < 1$. If $I = 1$, then $C^{p,k}(E) = B_1^{p,k}(E)$ and so $P^k\left[C^{p,k}(\mathbf{A}_*^k(\delta))\right] = P^k\left[B_1^{p,k}(\mathbf{A}_1^k(\delta))\right]$. $P^k[\mathbf{A}_1^k(\delta) \backslash B_i^{p,k}(\mathbf{A}_1^k(\delta))] \leq \frac{p}{1-p} P^k[\sim \mathbf{A}_1^k(\delta)] = \frac{p(1-P^k[\mathbf{A}_1^k(\delta)])}{1-p}$; so
$P^k[B_i^{p,k}(\mathbf{A}_1^k(\delta))] \geq P^k[\mathbf{A}_1^k(\delta) \cap B_i^{p,k}(\mathbf{A}_1^k(\delta))] = P^k[\mathbf{A}_1^k(\delta)] - P^k[\mathbf{A}_1^k(\delta) \backslash B_i^{p,k}(\mathbf{A}_1^k(\delta))]$
$> P^k[\mathbf{A}_1^k(\delta)] - \frac{p(1-P^k[\mathbf{A}_1^k(\delta)])}{1-p} = \frac{P^k[\mathbf{A}_1^k(\delta)]-p}{1-p}$. By lemma 1, pointwise convergence implies $P^k[\mathbf{A}_1^k(\delta)] \to 1$ as $k \to \infty$, so that $P^k\left[C^{p,k}(\mathbf{A}_*^k(\delta))\right] > p$ for sufficiently large k. □

In fact, it is straightforward to give a generalization of both corollary 1 and lemma 5 which shows that as long as at most one individual has an infinite number of information sets in his partition, then pointwise convergence implies AUCCP and thus ALP.

5.2. Alternative Notions of ε–Optimality

An ex ante definition of (mixed strategy) ε-optimality would require that for all $i \in \mathcal{I}$ and $\sigma_i' \in \Sigma_i$,

$$\sum_{\omega \in \Omega} \left(u_i\left(\sigma(\omega), \omega\right) - u_i\left(\left\{\sigma_i'(\omega), \sigma_{-i}(\omega)\right\}, \omega\right)\right) P^k\left[\omega\right] \geq -\varepsilon$$

Pointwise convergence would be sufficient for ALP if we used ex ante ε-optimality in the ALP definition. But an interim notion is more natural in an incomplete information game. Notice that the results of this paper do not seem to parallel any general payoff continuity results for games because the interim ε-optimality notion is special to incomplete information games. A working paper version of Monderer and Samet (1989) emphasized the difference between ex ante and interim notions of ε-optimality.

A pure strategy definition of (interim) ε-optimality would say that a strategy σ_i is ε–optimal at ω under P^k given σ if for all a_i such that $\sigma_i(a_i) > 0$ and all $a_i' \in A_i$,

$$\sum_{\omega' \in Q_i(\omega)} \left(u_i\left(a_i, \omega'\right) - u_i\left(\left\{a_i', \sigma_{-i}(\omega')\right\}, \omega'\right)\right) P^k\left[\omega' \left| Q_i(\omega)\right.\right] \geq -\varepsilon \qquad (5.1)$$

Our results are true as stated under this alternative definition. Indeed, the proofs are considerably simpler.

5.3. Alternative Notions of Closeness of Payoffs

The main result of the paper is sensitive to the precise formulation of the notion of continuity. But the techniques would generalize quite easily to alternative formulations. For example, if the notion of closeness in the definition of ALP (i.e. equation (4.1)) was replaced with the following notion

$$\left| \sum_{\omega \in \Omega} \left(u_i\left(\sigma'(\omega), \omega\right) - u_i\left(\sigma(\omega), \omega\right)\right) P^k\left[\omega\right] \right| \leq \eta, \text{ for } \textit{some } i \in \mathcal{I}$$

then we would have to alter the key equation (4.7) in the definition of almost uniform convergence to require

$$P^k\left[B_i^{p,k}\left(C^{p,k}\left(\mathbf{A}_*^k(\delta)\right)\right)\right] \geq p \text{ for } some\ i \in \mathcal{I}$$

The main theorem would then hold for almost the same proof.

5.4. Uncountable case

The insights of this paper would appear to be applicable to problems with continuous signals and thus uncountable state spaces. Carlsson and van Damme (1993) study the equilibria of 2-player 2-action games where each player observes payoffs with a small amount of noise. With no noise, some of the complete information games would have two strict Nash equilibria. But with positive but arbitrarily small noise, the unique (Bayesian) equilibrium has each player only choosing the risk dominant of the two strict Nash equilibria. This result would appear to be analogous to our theorem 4: there is weak convergence of the prior probability measures (as the noise goes to zero); ALP fails; and, for the "canonical" choice of conditional probability, almost uniform convergence of conditional probabilities fails.

The technical problem with generalizing to uncountable state spaces was our requirement that $P^k[Q_i(\omega)]$ was always strictly positive. This allowed us to ignore problems of the indeterminate conditional probabilities; for instance, the δ-agreement set $\mathbf{A}_i^k(\delta)$ will depend on the choice of conditional probabilities. In Kajii and Morris (1994), we develop an uncountable state space analogue of Monderer and Samet's (1989) characterization of common p-belief, where belief operators act on classes of equivalent events. Roughly speaking, it is shown that there is no loss of generality in fixing any conditional probabilities, once it is understood that any statement on ω means that the statement is true with probability one. Thus the extension of the results of this paper would require only a careful checking of measurability conditions and regularity conditions ensuring existence of equilibrium.

5.5. Embedding: Changing Priors versus Changing Partitions

The standard way of comparing information structures is to fix the probability distribution on the state space and vary individuals' information partitions. Monderer and Samet [MS] (1990) and Stinchcombe (1988) have given continuity results for Bayesian equilibria and characterized topologies on information structures in that framework. The MS result also had to do with a high ex ante probability of the set of states where conditional probabilities are close. What is the contribution of our approach?

We believe that our changing priors approach is more intuitive because notions of closeness of priors are easier to understand that notions of closeness of information partitions. But more importantly, we believe that both exercises must be equivalent and that the versatile and analytically simple approach of this paper captures quite generally the issue of the closeness of information systems. There should be a general equivalence relation between a model written as a comparison of information partitions with fixed prior and our framework (and probably vica versa).

But to establish the formal, mathematical meaning of "equivalence" is beyond the scope of this paper. So we shall just give a general recipe to compare different partitions in our framework.

An information structure is a tuple $(S, \pi, (\mathcal{P}_i))$ where:

- S is a countable state space
- π is a countably additive probability measure on S.
- $\mathcal{P}_i$ is a partition of S, for each agent $i = 1, ..., I$.

An information structure $(\Omega, P, (\mathcal{F}_i))$ is said to be in *canonical form* if

1. $\Omega = S_0 \times S_1 \times \cdots \times S_I$ where each $S_i = \aleph$.
2. for each i, $\mathcal{F}_i = \{\{(s_0, ..., s_I) : s_{-i} \in \Omega_{-i}\} : s_i \in S_i\}$

In canonical form $(\Omega, P, (\mathcal{F}_i))$, a state $\omega \in \Omega$ is generated according to P, and agent i observes the ith coordinate of ω. Notice that in canonical form, there is no choice of partition, and the difference in information is captured completely by the prior P.

For a given information structure $(S, \pi, (\mathcal{P}_i))$, define an information structure $(\Omega, P, (\mathcal{F}_i))$ in canonical form as follows:

- $\Omega = S \times \underbrace{S \times \cdots \times S}_{I \text{ times}}$
- $P[\omega] \equiv P[(s_0, s_1, ..., s_I)] = \pi(s_0) \cdot \pi[s_1 | \mathcal{P}_1(s_0)] \cdot \cdots \cdot \pi[s_I | \mathcal{P}_I(s_0)]$

where for any $s \in S$, $\mathcal{P}_i(s)$ denotes the member of partition $\mathcal{P}_i$ that contains s, and

$$\pi[s_i | \mathcal{P}_i(s_0)] = \begin{cases} \frac{\pi(s_i)}{\pi(\mathcal{P}_i(s_0))} \text{ if } s_0 \in P(s_i) \\ 0 \text{ otherwise.} \end{cases}$$

So $\pi[s_0 | \mathcal{P}_i(s_i)]$ is the (natural) conditional probability of s_0 when agent i learns that a state in partition $\mathcal{P}(s_i)$ has been realized. In words, P is generated by a two stage process where nature first chooses a "true" state s_0 according to π, and then it chooses a "signal" s_i for each player i according to the conditional probability given s_0. Since $\sum_{s_i} \pi[s_i | s_0] = 1$, P is a countably additive probability measure on Ω, and by construction,

$$P[\bar{s}_0 | Q_i(\omega)] = \frac{\sum_{\{\omega : s_0 = \bar{s}_0\}} P[\omega]}{P[Q_i(\omega)]} = \frac{\pi[s_0]}{\pi[\mathcal{P}_i(s_0)]} = \pi[s_0 | P_i]$$

This construction is intended to demonstrate that our representation captures a canonical form. Actually translating our results - for example, into the framework of Monderer and Samet (1990) - would no doubt be rather complicated. In particular, when looking at results about games, we must be careful to keep track of payoff relevant states. In the definition of ALP, the continuity is required for all games defined on Ω. A more general construction might consider only games depending on some partition representing payoff-relevance.

References

[1] R. Aumann, Agreeing to disagree, *Annals of Statistics* 4 (1976), 1236-1239.

[2] H. Carlsson and E. van Damme, Global games and equilibrium selection, *Econometrica* 61:5 (1993), 989-1018.

[3] G. Debreu, A social equilibrium existence theorem, *Proceedings of the National Academy of Sciences*, 38:886-893.

[4] D. Fudenberg and J. Tirole, Game Theory, Cambridge: M.I.T. Press 1991.

[5] J. Halpern, Reasoning about knowledge: An overview. *Theoretical Aspects of Reasoning about Knowledge: proceedings of the 1986 conference.* Los Altos: Morgan Kaufman.

[6] A. Kajii and S. Morris, Common p-Belief: the General Case, CARESS Working Paper #94-15.

[7] D. Monderer and D. Samet, Approximating common knowledge with common beliefs. *Games and Economic Behavior* 1 (1989), 170-190.

[8] D. Monderer and D. Samet, Proximity of Information in Games with Incomplete Information, 1990.

[9] S. Morris, R. Rob and H. Shin (1993), Risk dominance and stochastic potential. CARESS Working Paper #93-15, forthcoming in *Econometrica.*

[10] A. Rubinstein, The electronic mail game: Strategic behavior under almost common knowledge. *American Economic Review* 79 (1989), 385-391.

[11] M. Stinchcombe, Approximate Common Knowledge, University of California at San Diego (1988).

This is a list of recent CARESS Working Papers. A complete list (dating from inception of the series) can be obtained by writing to:

Mrs. Kelly A. Quinn
CARESS
3718 Locust Walk
McNeil Building
Philadelphia, PA 19104-6297

92-01R "Normal Form Structures in Extensive Form Games" by George J. Mailath, Larry Samuelson and Jeroen M. Swinkels
Revised and final version is forthcoming in the Journal of Economic Theory.
92-02 "R&D, Investment and Industry Dynamics" by Saul Lach and Rafael Rob
92-03 "Sales, Uncertainty and the Determinants of Investments" by Rafael Rob
92-04 "Market Games with Asymmetric Information: The Private Information Core" by Beth Allen
92-05 "Incentives in Market Games with Asymmetric Information: Approximate (NTU) Cores in Large Economies" by Beth Allen
92-06R "Evolution of Equilibria in the Long Run: A General Theory and Applications" by Michihiro Kandori and Rafael Rob
92-07 "Committing to a Mechanism: Resilient Allocation Rules for Bilateral Trade" by Roger Lagunoff
*92-08 "Jet Symposium on Evolutionary Game Theory: Introduction" by George J. Mailath
Revised and final version appears in Journal of Economic Theory, 57, (August 1992), pp. 259-277.
92-09 "When does Information Lead to Trade? Efficiency and Trade with Heterogeneous Prior Beliefs and Asymmetric Information" by Stephen Morris
Revised and final version is forthcoming in Econometrica.
92-10 "Towards a Theory of Incentive Compatibility for Economies with Incomplete Market Structures" by Yves Younès
92-11 "Incentives in Market Games with Asymmetric Information: The Value" by Beth Allen
92-12 "On the Social Stability of Coalitional Property Rights Regimes" by Gerhard Glomm and Roger Lagunoff
*92-13 "Finite Bubbles with Short Sale Constraints and Asymmetric Information" by Franklin Allen, Stephen Morris and Andrew Postlewaite
Revised and final version appears in Journal of Economic Theory
*92-14 "Barter and Monetary Exchange Under Private Information" by Seve Williamson and Randall Wright
Revised and final version forthcoming in American Economic Review

92-15 "Market Games with Asymmetric Information: The Core with Finitely many States of the World" by Beth Allen
92-16 "Induction and Bounded Rationality in Repeated Games" by In-Koo Cho and Akihiko Matsui
92-17 "Dynamic Consistency and the Value of Information" by Stephen Morris
92-18 "Learning and the Ramsey Policy" by In-Koo Cho and Akihiko Matsui
92-19 "Coevolution to Mixed Strategies" by Beth Allen
92-20 "Fiat Money, Barter, and Optimal Monetary Policy with Capital" by Fumio Hayashi and Akihiko Matsui
92-21 "On Modeling Cheap Talk in Bayesian Games" by Steven A. Matthews and Andrew Postlewaite
92-22 "Acceptance Games and Protocols" by Stephen Morris
92-23 "The Re-sale Premium for Assets in General Equilibrium" by Stephen Morris
92-24 "Revising Beliefs and Knowledge: A Generalization of Blackwell's Theorem" by Stephen Morris
92-25 "Sufficiently Specialized Economies have Nonempty Cores" by Roger D. Lagunoff
92-26 "Inequality and the Social Stability of Economies with Collective Property Rights" by Gerhard Glomm and Roger D. Lagunoff
92-27 "Revising Knowledge: A Decision Theoretic Approach" by Stephen Morris
92-28 "Adverse Selection, Reputation, and Firms in Professional Service Markets" by Jaewoo Ryoo
92-29 "Constrained Suboptimality with Many Agents" by Atsushi Kajii
*92-30 "Search, Matching, and Unions" by Kenneth Burdett and Randall Wright
Revised and final version forthcoming in *Proceedings of the Third Symposium on Panel Data and Labor Market Dynamics*, eds. by H. Bunzel, P. Jensen and N. Westergard-Nielson, Amsterdam, North-Holland.
92-31 "The Burden of Proof in a Game of Persuasion" by Hyun Song Shin
92-32 "How much Common Belief is Necessary for a Convention?" by Hyun Song Shin and Timothy Williamson
92-33 "Multiple Tasks in the Principal-Agent Model" by Dirk Bergemann
92-34 "Duality Principles on Intertemporal Preference and Capital Accumulation" by H. Uzawa
93-01 "Perpetual Randomness in Evolutionary Economics" by George J. Mailath
93-02 "Noisy Bayes Updating and the Value of Information" by Stephen Morris and Hyun Song Shin
93-03 "Buyers and Sellers" by Kenneth Burdett, Melvyn Coles, Nobuhiro Kiyotaki and Randall Wright
93-04 "AUSI Expected Utility: An Anticipated Utility Theory of Relative Disappointment Aversion" by Simon Grant and Atsushi Kajii

93-05 "Bandwagon Effects and Long Run Technology Choice" by Michihiro Kandori and Rafael Rob

93-06 "A Note on the Incentives to Agglomerate Under Demand Uncertainty" by Rafael Rob

93-07 "A Primer on Functional Differential Equations with Applications to Economics" by Paul J. Zack

93-08 "An Endogenous Rate of Time Preference, the Penrose Effect, and Dynamic Optimality of Environmental Quality" by Hirofumi Uzawa

93-09 "Bid Ask Spreads with Two Sided Private Information" by Stephen Morris

93-10 "Risk, Uncertainty and Hidden Information" by Stephen Morris, Rafael Rob and Hyun Song Shin

93-11 "Risk Dominance and Stochastic Potential " by Stephen Morris, Rafael Rob and Hyun Song Shin
Revised and final version is forthcoming in Econometrica

93-12 "Structural Indifference in Normal Form Games" by George J. Mailath, Larry Samuelson and Jeroen M. Swinkels

93-13 "Search, Bargaining, Money and Prices" by Alberto Trejos and Randall Wright

93-14 "Agitators and Free Riders on the Path to Cooperation: The Evolution of Cooperation in Mechanisms for Public Projects" by Roger D. Lagunoff and Akihiko Matsui

93-15 "Depth of Knowledge and the Effect of Higher Order Uncertainty" by Stephen Morris, Andrew Postlewaite and Hyun Song Shin
Revised and final version is forthcoming in Economic Theory.

93-16 "Monetary Exchange Under Private Information: A Note" by Alberto Trejos

93-17 "Relative Income Concerns and the Rise in Married Women's Employment" by David Neumark and Andrew Postlewaite

93-18R "A Cardinal Characterization of the Rubinstein-Safra-Thomson Axiomatic Bargaining Theory" by Simon Grant and Atsushi Kajii

93-19 "Evolution and Rationalizability" by Akihiko Matsui

93-20 "Strictness, Evolutionary Stability and Repeated Games with Common Interest" by Dieter Balkenborg

93-21 "On Generic Pareto Improvement in Competitive Economies with Incomplete Asset Structures" by Alessandro Citanna and Antonio Villanacci

*93-22 "Search, Evolution and Money" Randall Wright
Revised and final version forthcoming in the Journal of Economic Dynamics and Control

*93-23 "A Note on Sunspot Equilibria in Search Models of Fiat Money" by Randall Wright
Revised and final version forthcoming in the Journal of Economic Theory

93-24 "Demand and Supply in New Markets: Diffusion with Bilateral Learning" by Nikolaos Vettas

93-25 "Short-Run Independence of Monetary Policy under Pegged Exchange Rates and Effects of Money on Exchange Rates and Interest Rates" by Lee E. Ohanian and Alan C. Stockman

93-26 "Majority Rule and the Public Provision of Health Care" by Miguel Gouveia

93-27 "Pareto Improving Financial Innovation in Incomplete Markets" by David Cass and Alessandro Citanna

93-28 "On the Dynamic Selection of Mechanisms as Clubs for the Provision of a Public Project" by Roger D. Lagunoff

94-01 "Expected Utility and Case-Based Reasoning" by Akihiko Matsui

94-02 "Sequential Stratified Sampling" by Edward J. Green and Ruilin Zhou

94-03 "Bargaining, Boldness and Nash Outcomes" by Simon Grant and Atsushi Kajii

94-04 "Learning and Strategic Pricing" by Dirk Bergemann and Juuso Valimaki

94-05 "Evolution in Mechanisms for Public Projects" by Roger D. Lagunoff and Akihiko Matsui (previous version 93-14)

94-06 "Constrained Suboptimality in Incomplete Markets: A General Approach and Two Applications" by Alessandro Citanna, Atsushi Kajii and Antonio Villanacci

94-07 "Pareto Improving Financial Innovation in Incomplete Markets" by David Cass and Alex Citanna (pervious version 93-27)

94-08 "Commodity Money Under Private Information" by Yiting Li

94-09 "Generic Local Uniqueness in the Walrasian Model: A Pedagogical Note" by Marcos de Barros Lisboa

94-10 "Bargaining-Induced Transaction Demand for Fiat Money" by Merwan Engineer and Shouyong Shi

94-11 "Politico-Economic Equilibrium and Economic Growth" by Per Krusell, Vincenzo Quadrini and José- Víctor Ríos-Rull

94-12 "On the Evolution of Pareto Optimal Behavior in Repeated Coordinates Problems" by Roger D. Lagunoff

94-13 "Evolution and Endogenous Interactions" by George J. Mailath, Larry Samuelson and Avner Shaked

94-14 "How Proper is Sequential Equilibrium?" by George J. Mailath, Larry Samuelson and Jeroen M. Swinkels

94-15 "Common p-Belief: The General Case" by Atsushi Kajii and Stephen Morris

94-16 "Impact of Public Announcements on Trade in Financial Markets" by Stephen Morris and Hyun Song Shin

94-17 "Payoff Continuity in Incomplete Information Games and Almost Uniform Convergence of Beliefs" by Atsushi Kajii and Stephen Morris