Common Learning

Martin W. Cripps[*], Jeffrey C. Ely[†]
George J. Mailath[‡], Larry Samuelson [§]

March 1, 2006
PRELIMINARY

Abstract

Suppose that agents observing informative signals about an unknown parameter will eventually learn the true parameter. Must the parameter eventually be common knowledge? We provide an example showing that it need not, and establish sufficient conditions under which the parameter eventually will be common knowledge.

1 Introduction

Standard models of Bayesian learning provide conditions under which an agent who receives a stream of signals about an unknown parameter will form posterior beliefs that converge to the truth. When these conditions are met, an agent who must attain a certain confidence in the value of the parameter in order to profitably exploit an opportunity can be assured of eventually doing so.

The issue becomes more complicated when joint actions are required to exploit an opportunity. Suppose that two agents can each attain a payoff of 1 if they coordinate on an action whose identify depends on an unknown parameter, as shown in Figure 1. Miscoordination brings a payoff of $-c$, while inaction yields a payoff of zero. The players

[*]Olin School of Business, Washington University in St. Louis
[†]Department of Economics, Northwestern University
[‡]Department of Economics, University of Pennsylvania
[§]Department of Economics, University of Wisconsin

	A	B	W
A	$1, 1$	$-c, -c$	$-c, 0$
B	$-c, -c$	$-c, -c$	$-c, 0$
W	$0, -c$	$0, -c$	$0, 0$

Parameter θ

	A	B	W
A	$-c, -c$	$-c, -c$	$-c, 0$
B	$-c, -c$	$1, 1$	$-c, 0$
W	$0, -c$	$0, -c$	$0, 0$

Parameter θ'

Figure 1: Payoffs from a potential joint opportunity, with actions A, B, or wait (W) available to each agent in each period.

receive a stream of signals that allow them to draw inferences about the parameter. Their task in each period is to choose action A, action B, or wait (W) until the next period. Other things equal, they would prefer to exploit the opportunity sooner rather than later.

Under what circumstances will the agents be able to make good use of their opportunity? Choosing action A dominates inaction if an agent attaches probability at least $\frac{c}{c+1} \equiv C$ to the event that the parameter is θ and the other player chooses A. One possibility is then for the agents to each adopt the rule of choosing A as soon as the posterior probability they attach to parameter θ reaches C. Given such behavior, however, each player could increase their expected payoff by instead choosing A as soon as they attach probability at least C to the event that their opponent attaches probability at least C to parameter θ. Having adopted this latter rule, however, each agent is then better off playing A when they attach probability at least C to the event that their opponent attaches probability at least C to the event that their opponent attaches probability at least C to parameter θ. Continuing in this fashion, coordinating on A in state θ requires that arbitrarily long strings of the form "there is probability C that my opponent thinks there is probability C that my opponent thinks ... there is probability C that the parameter is θ. The parameter must then become common C-belief.

Now suppose that various forms of this opportunity arise, characterized by different values of the miscoordination penalty c. What does it take to ensure that all of these opportunities can be exploited? The information process must be such that the parameter eventually becomes common 1-belief.

We refer to a situation in which the parameter becomes common 1-belief as "common learning" and say that the agents commonly learn the parameter in this case. In some cases, of course, the agents' signals will be sufficiently uninformative that they will have no hope of learning the value of the parameter, much less commonly learning this value. Suppose instead that the signals are rich enough that the players eventually (except for exceptional cases of probability zero) each learn the parameter. Must the parameter then be commonly learned?[1] We show in this paper that common learning can fail, even though the information process is such that both agents learn the true state with probability one. Hence, higher order beliefs matter—learning need not imply common learning. We also provide sufficient conditions for common learning to occur.

The issue of common learning appears in a variety of dynamic relationships:

Strategic Experimentation. Bolton and Harris (1999) study firms involved in dynamic information-gathering, each drawing inferences from their own observations and from observations of their rival. If the firms are sufficiently patient, each will eventually learn the underlying state. Will this state be commonly learned? In Bolton and Harris, the information observed by the firms is public and hence their posterior beliefs are common knowledge, ensuring that the state eventually becomes commonly learned. Suppose, however, that the firms can observe the outcomes of their own research and development experiments but can only observe whether their rival is still pursuing its research (or not). Posterior beliefs are then no longer common. While we can be certain that sufficiently patient firms will learn the state, it remains an open question whether the state will be commonly learned.[2]

Callibration. Foster and Vohra (1999) and Rustichini (1998) examine models in which an agent privately observing a stream of data will eventually be able to make good predictions of how the stream will continue. If there are many such agents, will it eventually be commonly learned that the agents are predicting well? This remains an open question.

Reputation with Imperfect Monitoring. Consider a repeated game of imperfect monitoring, with uncertainty about the type of one player, referred to as player 1. As the game proceeds, the other players will continually update their posterior beliefs about player 1's type. In general, the imperfect monitoring will ensure that their posterior beliefs

[1]Common 1-belief is often (somewhat informally) referred to as common knowledge. In these terms our question is the following. If both agents eventually know the value of the parameter, must it eventually be common knowledge?

[2]The equilibria of this game are the subject of current research (e.g., Moscarini and Squintani (2004), Rosenberg, Solan and Vieille (2005)). Similar issues arise if firms are making private observations on a state-dependent demand function, as in Bergemann and Välimäki (1997).

are private. Conditions are known under which player 1's reputation will eventually vanish, in the sense that the other players learn 1's type. However, we do not know whether this type ever becomes commonly learned.[3]

Folk Theorems with an Unknown Payoff Matrix. Consider a repeated game in which payoffs are generated by a process conditioned on an unknown parameter, with the agents making noisy and private observations of the payoffs. Establishing that payoffs eventually become commonly learned may be an important first step in establishing a folk theorem.[4]

The following section presents a general model and the tools needed for subsequent arguments, most notably the notion of common q-belief. Section 3 derives our main result, establishing sufficient conditions for common learning. Section 4 presents a counter example to common learning.

2 A General Framework

2.1 A Model of Multi-Agent Learning

Time is discrete and periods are denoted $t = 0, 1, 2, ...$ Before date zero, nature selects a parameter θ from the finite set Θ according to the prior distribution $\{p^\theta\}_{\theta\in\Theta}$.

Conditional on θ, there is a stochastic process $x \equiv \{x_t\}_{t=0}^{\infty}$ that generates a signal profile $x_t \in X$ for every period. Our notation does not make explicit the dependence of x on θ. We identify a state as a parameter and a sequence of signal profiles, with the set of states given by $\Omega \equiv \Theta \times X^\infty$. We use P to denote the measure on Ω induced by nature's move and the processes $(x)_{\theta\in\Theta}$, and use $E[\cdot]$ to denote expectations with respect to this measure. P^θ denotes the measure conditional on a given parameter and $E^\theta[\cdot]$ denotes expectations with respect to this measure.

[3]Cripps, Mailath and Samuelson (2004) avoid this issue, while showing that the reputation must vanish, by using player 1's knowledge of player 2's actions, and the implications for 1's actions, to argue that otherwise a contradictory implication arises for beliefs about player 1 and player 1's actions. However, it should be possible to assess whether player 1 can maintain a reputation by examining only beliefs and without relying on 1's knowledge of 2's actions or 1's induced actions.

[4]Gossner and Vieille (2003) and Wiseman (2005) examine folk theorems for the case in which the parameter generating payoffs is commonly known.

There are 2 agents in the game, denoted $n = 1, 2$. The signal profile in period t is denoted by $x_t \equiv (x_{1t}, x_{2t}) \in X$, where $x_{nt} \in X_n$ is the component of the period t signal that is observed by agent n. A partial history for agent n is denoted by $h_{nt} \equiv (x_{n0}, x_{n1}, ..., x_{nt-1})$. On occasion it is convenient to refer to the stochastic process $x_n \equiv \{x_{nt}\}_{t=0}^{\infty}$ that generates agent n's signals.

We let $H_{nt} \equiv (X_n)^t$ denote the space of partial histories for agent n and let $\{\mathcal{H}_{nt}\}_{t=0}^{\infty}$ denote the filtration induced on Ω by agent n's partial histories. The $\mathcal{H}_{nt}$-measurable random variables $p_{nt}(\theta) \in [0, 1]$, for $\theta \in \Theta$, is defined by

$$p_{nt}(\theta) \equiv E[\mathbf{1}_{\theta} \mid \mathcal{H}_{nt}].$$

Thus, $p_{nt}(\theta)$ denotes agent n's posterior on the parameter at the start of period t. The random variables $\{p_{nt}(\theta)\}_{t=0}^{\infty}$ are a bounded martingale with respect to the measure P, and so the agents' priors converge almost surely (Billingsley (1986, Theorem 35.4)).

We will assume throughout that each agent can individually learn the state. At this point, we simply make this as an assumption. We replace this assumption below with sufficient conditions on the signal generating process.

Assumption 1 (Learnability) *For all n and all $\theta \in \Theta$, $p_{nt}(\theta) \to 1$ P^{θ}-almost surely.*

Assumption 1 ensures that the limiting random variable, to which $p_n t(\theta)$ converges, places unitary probability on the true state θ. Intuitively, Assumption 1 holds when each agent's signals, conditional on the state, do not become uninformative as t increases. Our aim in this paper is to discuss the additional conditions that must be imposed to ensure not just that each agent learns the state, but that the agents commonly learn the state.

2.2 Common q-Belief

For any event $F \subset \Omega$ the random variable $E[\mathbf{1}_F \mid \mathcal{H}_{nt}]$ is the probability agent n attaches to F given her information at time t. We define

$$B_{nt}^q(F) \equiv \{\omega \in \Omega \mid E[\mathbf{1}_F \mid \mathcal{H}_{nt}] > q\}.$$

Thus $B_{nt}^q(F)$ is the set of states of the world where at time t agent n attaches at least probability q to event F. Assumption 1 implies (but is not equivalent to) $P^{\theta}(B_{nt}^q(\{\theta\})) \to 1$ as $t \to \infty$ for all $q \in (0, 1)$.

The event that F is *q-believed* at time t, denoted by $B_t^q(F)$, occurs if every agent attaches at least probability q to F, that is,

$$B_t^q(F) \equiv \cap_n B_{nt}^q(F).$$

The event that F is *common q-belief* at date t is

$$C_t^q(F) \equiv \cap_{k\geq 1}[B_t^q]^k(F).$$

Hence, on $C_t^q(F)$, the event F is q-believed and this event is itself q believed and so on.

The parameter θ is common q-belief at time t on the event $C_t^q(\theta)$. We say that the agents commonly learn that the parameter θ if, for any probability q, there is a time such that, with high probability when the parameter is θ, it is common q-belief at all subsequent times that the parameter is θ:

Definition 1 (Common Learning) *The agents commonly learn parameter $\theta \in \Theta$ if for each q there exists a T such that for all $t > T$,*

$$P^\theta(C_t^q(\theta)) \geq q.$$

The agents commonly learn Θ if they commonly learn each $\theta \in \Theta$.

Because $C_t^q(\theta) \subset B_{nt}^q(\theta)$, common learning implies individual learning, i.e., the common learning of θ implies that $p_{nt}(\theta) \to 1$ with P^θ probability one.

An event F is *q-evident* at time t if it is q-believed when it is true, that is,

$$F \subset B_t^q(F).$$

Proposition 1 (Monderer and Samet (1989)) *F is common q-belief at $\omega \in \Omega$ at time t if and only if there exists an event $F' \subset \Omega$ such that F' is q-evident at time t and $\omega \in F' \subset B_t^q(F)$.*

Corollary 1 *The agents commonly learn Θ if, for any $\theta \in \Theta$ and $q \in (0,1)$, there exists a sequence of events $\tilde{F}_t$ and a period T such that for all $t > T$, $\tilde{F}_t$ is q-evident, θ is q-believed on $\tilde{F}_t$, and $P^\theta(\tilde{F}_t) > q$.*

2.3 Perfect Correlation and Independence

This section considers two special cases, involving signals that are perfectly correlated and signals that are completely independent. In either of these two opposite extremes, it is relatively easy to establish common learning.

Suppose first that the signal process is public, as is commonly assumed in the literature. This ensures that $p_{nt}(\theta) = p_{n't}(\theta)$ for all n, n', θ and t, and hence that beliefs are always common knowledge. Assumption 1 thus immediately implies common learning.

At the other extreme, we have the case of independent signals. Here, the fact that each agent learns the state while learning nothing about other agents' signals ensures common learning.

Proposition 2 *Let Assumption 1 hold and suppose that for each $\theta \in \Theta$, the stochastic processes $\{x_{nt}\}_{t=0}^{\infty}$ and $\{x_{n't}\}_{t=0}^{\infty}$ are independent. Then the agents commonly learn Θ.*

Proof. Our task is to show that under a given parameter θ and for any $q < 1$, the event that θ is common q-belief occurs with at least probability q for all sufficiently large t. From Corollary 1, it is sufficient to find a sequence of events $\tilde{F}_t$ and time T such that for all $t > T$, $\tilde{F}_t \subset B_t^q(\tilde{F}_t)$ ($\tilde{F}$ is q-evident), $\tilde{F}_t \subset B_t^q(\theta)$ (θ is q-believed on $\tilde{F}_t$), and $P^\theta(\tilde{F}_t) > q$.

Let $\tilde{F}_t \equiv \{\theta\} \cap B_t^{\sqrt{q}}(\theta)$. Because $\tilde{F}_t \subset B_t^{\sqrt{q}}(\theta) \subset B_t^q(\theta)$, state θ is q-believed on $\tilde{F}_t$.

We next argue that $\tilde{F}_t$ is q-evident. We must show that $\tilde{F}_t \subset B_{nt}^q(\tilde{F}_t)$ for all n and for any t sufficiently large. By construction, $\tilde{F}_t \subset B_{nt}^{\sqrt{q}}(\theta)$, and hence agent n trivially attaches probability at least q (indeed, probability 1) to the state being in $B_{nt}^{\sqrt{q}}(\theta)$. It then suffices to show that on the set $\tilde{F}_t(\theta)$, agent n attaches at least probability q to the event $B_{n't}^{\sqrt{q}}(\theta) \cap \{\theta\}$, $n' \neq n$, since we would then have $\tilde{F}_t \subset \left(B_n^1(B_{nt}^{\sqrt{\theta}}(\theta)) \cap B_n^q(B_{n't}^{\sqrt{\theta}}(\theta) \cap \{\theta\}) \right) \subset B_n^q(\tilde{F}_t)$. By Assumption 1, we can choose T sufficiently large that $P^\theta(B_{nt}^{\sqrt{q}}(\theta)) > \sqrt{q}$ for all n and all $t > T$. The conditional independence of agents' signals implies that, given θ, h_{nt} is uninformative about others' signals, and hence $P^\theta(B_{n't}^{\sqrt{q}}(\theta) \mid \mathcal{H}_{nt}) > \sqrt{q}$. But, at any state in $\tilde{F}_t$ it is the case that $P(\theta \mid \mathcal{H}_{nt}) > \sqrt{q}$. Multiplying the previous two inequalities gives the needed result that on the set $\tilde{F}_{nt}(\theta)$, agent n attaches at least probability q to the event $B_{n't}^{\sqrt{q}}(\theta) \cap \{\theta\}$, $n' \neq n$.

Finally, we need to show that $P^\theta(\tilde{F}_t) > q$. Independence implies $P^\theta(\tilde{F}_t) = \prod_{n'} P^\theta(B_{n't}^{\sqrt{q}}(\theta)) > q$, where the final inequality follows from our choice of T. This completes the proof. ∎

The role of independence in this argument is to ensure that agent n's signals provide n with no information about n''s signals. One would expect common learning to be more likely the more information n has about n', so that n has a good idea of n''s beliefs. However, when agent n receives no information about n', as in the case of independent signals, agent n eventually thinks it quite likely that n' has learned the parameter. In addition, we can place a lower bound on the rate of growth of n's beliefs about n''s beliefs. This suffices to establish common learning. When signals are correlated, n's signals will often provide useful information about n''s, accelerating the rate at which n learns about n' and reinforcing common learning. Sometimes, however, atypical signal realizations will cause n's beliefs about n' to go badly wrong. This poses an obstacle that may disrupt common learning.

This suggests that we may be able to ensure common learning, without independence, as long as the dependence is such that signals leading one agent too far astray, in terms of beliefs about the other agent, do not occur too often. The next section pursues this intuition to establish our main result.

3 Conditions for Common Learning

3.1 Assumptions

To make things notationally less demanding, we will focus in this subsection on the case where there are only two parameter values. Nothing of any import in the result below rests on this restriction, whose relaxation requires only the frequent addition of phrases of the form, "for any θ, θ'."

Assumption 2 *Agents 1 and 2 have finite signal sets, I and J respectively. Conditional on θ, the signal profile process, $\{x_t\}_{t=0}^{\infty}$ is independent and identically distributed across t.*

We let $\pi_\theta = (\pi_\theta^{ij})_{i=1,j=1}^{I\ \ J} \in \Delta(I \times J)$ denote the process generating the agents' signals conditional on θ for every period t. That is, π_θ^{ij} is the probability $(x_{1t}, x_{2t}) = (i, j)$ for parameter θ and every t. We use $\phi_\theta(i) \equiv \sum_j \pi_\theta^{ij}$, or $\phi_\theta \equiv (\phi_\theta(i))_{i=1}^I$, to denote the marginal probability of agent 1's signal i and use $\psi_\theta(j) = \sum_i \pi_\theta^{ij}$, or $\psi_\theta \equiv (\psi_\theta(j))_{j=1}^J$, to denote the marginal probability of agent 2's signal j.

Assumption 3 (Learning)

(3.1) For every pair θ and θ', there exist signals i and j such that $\phi_\theta^i/\phi_{\theta'}^i \neq 1$ and $\psi_\theta^j/\psi_{\theta'}^j \neq 1$.

(3.2) For all i, j, and θ, $\phi_\theta^i > 0$ and $\psi_\theta^j > 0$.

Note that the first part is equivalent to assumption 1. The second part is merely for convenience.

3.2 Preliminary Results: Frequencies are Enough

Let f_t^{ij} denote the number of instances in which agent 1 has received the signal i and agent 2 received the signal j before period t. Let $f_{2t}^j \equiv \sum_i f_t^{ij}$ and $f_{1t}^i \equiv \sum_j f_t^{ij}$. Let the empirical frequencies of the signals be denoted by vectors $\hat{\phi}_t \equiv (f_{1t}^i/t)_{i \in I}$ and $\hat{\psi}_t \equiv (f_{2t}^j/t)_{j \in J}$. Given Asssumption 3.2, there exists $b > 1$ such that for all i, j, θ and θ', $\phi_\theta^i/\phi_{\theta'}^i > \frac{1}{b}$ and $\psi_\theta^j/\psi_{\theta'}^j > \frac{1}{b}$.

For any $x \in \mathbf{R}^N$ we will use $\|x\|$ to denote the variation norm of x, that is, $\|x\| \equiv \frac{1}{2}\sum_{i \in I} |x_i|$.

We define a pair of matrices that will play a key role in the analysis. Let M_1 be a $I \times J$ matrix whose ijth element is $\frac{\pi_\theta^{ij}}{\phi_\theta(i)}$, i.e. the conditional probability in state θ of signal j given signal i. With this definition, at any date t, when agent 1 has observed signals with empirical frequencies $\hat{\phi}_t$, agent 1's expectation of the empirical frequencies observed by agent 2 is given by the matrix product $\hat{\phi}_t M_1$. Similarly, define M_2 to be the $J \times I$ matrix with ijth element $\frac{\pi_\theta^{ij}}{\psi_\theta(j)}$. Note that $\hat{\phi}_t \cdot M_1 M_2$ gives agent 1's expectation of agent 2's expecation of the empirical frequencies observed by agent 1.

The matrix obtained by the product $M_{12} := M_1 M_2$ is formally equivalent to a Markov transition on the set I of signals for agent 1. Much of our analysis will be based on this analogy.[5] To begin with, the transition matrix M_{12} partitions the set I into *recurrent classes*. Two signals i and i' belong to the same recurrent class iff the probability of transition from i to i' (in some finite number of steps) is positive. We let $(R_k)_{k=1}^K$ denote the collection of recurrent classes, and we implicitly re-order the elements of I so that the recurrent classes are grouped together and in the order of their indices.

Similarly, the matrix $M_{21} := M_2 M_1$ is a Markov transition on the set J. Observe that there is a one-to-one correspondence between the recurrent classes of M_{21} and M_{12}. Indeed the relation which associates the class containing i with the class containing j iff $\pi_\theta^{ij} > 0$ is a bijection. It is convenient therefore to group the elements of J by their recurrent classes in the same order as was done with I. Also we will use the notation R_k to refer to the kth recurrent class in either I or J when the context is clear. Viewing M_{12} as a Markov chain leads to the following observation which is central to our analysis.

Lemma 1 *There exists $r < 1$ a natural number n such that for all $k \in \{1, \ldots, K\}$ and for all $\mu, \mu' \in \Delta R_k$,*

$$\|\mu(M_{12})^n - \mu'(M_{12})^n\| < r\|\mu - \mu'\|$$

and similarly for $(M_{21})^n$.

[5] Samet did something like this.

Proof: First note that under M_{12}, there is a positive probability of transition from any signal i to itself, that is M_{12} is aperiodic. And, by definition, the restriction of M_{12} to any given recurrent class is irreducible and hence ergodic. Thus, because signals are grouped by their recurrent classes, there exists a natural number n such that $(M_{12})^n$ has the block-diagonal form. The blocks are the non-zero n-step transition probabilities between signals within a recurrent class. Note that $\mu \in \Delta R_k$ implies that the product of μ with $(M_{12})^n$ is just the product of μ with the kth block of $(M_{12})^n$. Because it has all non-zero entries, a standard result from the theory of Markov processes implies that the kth block is a contraction mapping. In particular, there exists an $r < 1$ such that the displayed inequality in the statement of the lemma holds for all $\mu, \mu' \in \Delta R_k$. ■

The preceding result suggests the outline of our proof of common learning. For example, suppose that $K = 1$ so that there is a unique recurrent class, and suppose that the value of n given by Lemma 1 is equal to 1. Then when agent 1's frequencies are close to their expected values in some state θ, i.e. $\|\hat{\phi}_t - \phi_\theta\| < \delta$, it will follow that agent one expects that agent two expects that agent one's frequencies are even closer, i.e. $\|\hat{\phi}_t(M_{12}) - \phi_\theta\| = \|\hat{\phi}_t(M_{12}) - \phi_\theta(M_{12})\| < r\delta$. We will show below (Lemma 3) that for large enough t whenever the empirical frequencies are close to their expected values in state θ, an agent assigns high probability to θ. Furthermore, for large enough t, one agent's expectation of another agents frequencies is approximately correct with high probability (Lemma 4). It will follow from these two facts that agent one assigns high probability to agent two assigning high probability to agent one believing in θ. This is not far from implying that θ is common p-belief for high p.[6]

Lemma 2 *For any $\varepsilon > 0$, $P^\theta(\|\hat{\phi}_t - \phi_\theta\| < \varepsilon) \to 1$*

Proof: Law of Large Numbers. ■

Learning the state

Lemma 3 *There exist $\delta > 0$ and $0 < \beta < 1$ and a function $k(t)$ such that $k(t) \to 1$ such that*

$$P_t(\theta | h_{1t}) \geq k(t)$$

for all h_{1t} such that for all k, $\|\hat{\phi}_t^k - \phi_\theta^k\| < \delta$ and $\beta < \frac{\hat{\phi}_t(R_k)}{\phi_\theta(R_k)} < \beta^{-1}$. Likewise for player 2.

[6]Common p-belief does not follow immediately from this statement because in general "iterated" p-belief (one believes that two believes ... and two believes that one believes ...) is strictly weaker than common p-belief (*both* believe that both believe...). See Morris. Therefore even in this special case of $K = n = 1$, the proof requires some additional steps.

Proof: Define the log-likelihood ratio as follows

$$\lambda_t = \frac{\mathrm{Prob}(\theta|h_t)}{1-\mathrm{Prob}(\theta|h_t)}.$$

We analyze the evolution of λ_t and show that β and δ can be chosen to ensure that $\lambda_t \to \infty$. Given the i.i.d. process, λ_t can be expressed as follows.

$$\lambda_t = \lambda_0 + \sum_{s=0}^{t-1} \log\left(\frac{\phi_\theta(i)}{\phi_{\theta'}(i)}\right)$$

where λ_0 is the likelihood ratio at time zero, i.e. the prior. We will approximate the last term. Let

$$H = \mathbf{E}_\theta\left(\log\frac{\phi_\theta}{\phi_{\theta'}}\right)$$

denote the relative entropy of ϕ_θ and $\phi_{\theta'}$.

$$\begin{aligned}
\left|\sum_{s=0}^{t-1}\log\left(\frac{\phi_\theta(i)}{\phi_{\theta'}(i)}\right) - tH\right| &= \left|\sum_i f_{1t}^i \log\left(\frac{\phi_\theta(i)}{\phi_{\theta'}(i)}\right) - t\sum_i \phi_\theta(i)\log\left(\frac{\phi_\theta(i)}{\phi_{\theta'}(i)}\right)\right| \\
&= t\left|\sum_i(\hat{\phi}_t(i)-\phi_\theta(i))\log\left(\frac{\phi_\theta(i)}{\phi_{\theta'}(i)}\right)\right| \\
&\le t\sum_i\left|(\hat{\phi}_t(i)-\phi_\theta(i))\log\left(\frac{\phi_\theta(i)}{\phi_{\theta'}(i)}\right)\right| \\
&\le t\log b\|\hat{\phi}_t-\phi_\theta\|
\end{aligned}$$

Thus,

$$\lambda_t \ge \lambda_0 + t\left(H - \log b\|\hat{\phi}_t - \phi_\theta\|\right)$$

and we can show that $\lambda_t \to \infty$ by showing that δ and β can be chosen to ensure $\log b\|\hat{\phi}_t - \phi_\theta\| < H$. For this, it is enough to observe that the mapping

$$\left(\left\{\hat{\phi}_t(R_k)\right\}_k, \left\{\hat{\phi}_t^k\right\}_k\right) \to \sum_k\sum_{i\in k}\left|\hat{\phi}_t(R_k)(\hat{\phi}_t^k)(i) - \phi_\theta(i)\right| = \|\hat{\phi}_t - \phi_\theta\|$$

is continuous and obtains the value zero when $\hat{\phi}_t(R_k) = \phi_\theta(R_k)$ and $\hat{\phi}_t^k = \phi_\theta^k$. ∎

Inferring opponent's history

Lemma 4 *For any $\varepsilon_1 > 0$, $\varepsilon_2 > 0$, there exists T such that for all $t > T$,*

$$P^\theta(\|\hat{\phi}_t M_1 - \hat{\psi}_t\| < \varepsilon_1 | h_{1t}) > 1 - \varepsilon_2 \tag{1}$$

$$P^\theta(\|\hat{\psi}_t M_2 - \hat{\phi}_t\| < \varepsilon_1 | h_{2t}) > 1 - \varepsilon_2 \tag{2}$$

for every h_t.

Proof: Conditional on θ and h_{1t}, agent 2's signals are independent across time but *not* identically distributed. In period s, given signal i_s, we have the conditional distribution $(\pi_\theta^{i_s j}/\phi_\theta^{i_s})_j$ over agent 2's signals. The average of the probability that agent 2 observes signal j over the t periods $\{0, 1, \dots, t-1\}$, conditional on h_{1t} is

$$\bar{\psi}_t^j := \frac{1}{t} \sum_{s=0}^{t-1} \frac{\pi_\theta^{i_s j}}{\phi_\theta^{i_s}} = \sum_i \hat{\phi}_t^i \frac{\pi_\theta^{ij}}{\phi_\theta^i}.$$

Thus, $\hat{\phi}_t M_1$ gives the vector of these average probabilities.

Consider the probability, conditional on θ, and h_{1t}, that agent one will believe with high probability that agent two's empirical frequencies, $\hat{\psi}_t$, are close to the average probabilities.

$$P^\theta(\|\hat{\phi}_t M_1 - \hat{\psi}_t\| \leq \varepsilon_1 | h_{1t})$$

We will construct a lower bound for this probability. Obviously,

$$P^\theta(\|\hat{\phi}_t M_1 - \hat{\psi}_t\| \leq \varepsilon_1 | h_{1t}) > 1 - \sum_{j=1}^J P^\theta \left(\left| \bar{\psi}_t^j - \hat{\psi}_t^j \right| > \frac{\varepsilon_1}{J} | h_{1t} \right). \tag{3}$$

Now we consider a hypothetical stochastic process, with independent and identically distributed random variables in each of periods $0, \dots, t-1$, where each such random variable produces the signal j in each period $\{0, \dots, t-1\}$ with probability $\bar{\psi}_t^j$. Let the empirical frequencies generated by this process be denoted by $\eta_t \in \Delta(J)$. The true process, generating frequencies $\hat{\psi}_t$, attaches the same average probability to each signal j over periods $0, \dots, t-1$ as does the fictitious process, but the true process does not have identical distributions.

We use this fictitious process to find an upper bound on the terms in the sum in (3). By Hoeffding (1956 Theorem 4 p. 718) , the true process is more concentrated about its mean than is the hypothetical process, that is[7]

$$P^\theta(|\bar{\psi}_t^j - \eta_t^j| > \varepsilon_1/N) > P^\theta(|\bar{\psi}_t^j - \hat{\psi}_t^j| > \varepsilon_1/N | h_{1t}), \qquad j = 1, 2, ..., J.$$

Applying this upper bound to (3), we have

$$P^\theta(\|\hat{\phi}_t M_1 - \hat{\psi}_t\| \leq \varepsilon_1 | h_{1t}) > 1 - \sum_{j=1}^J P^\theta \left(\left| \bar{\psi}_t^j - \eta_t^j \right| > \frac{\varepsilon_1}{J} | h_{1t} \right). \tag{4}$$

[7]The intuition for this result is that 100 flips of a $(p, 1-p)$ coin generates a more dispersed distribution than 100 flips of 100 biased coins with average probability of heads equal p.

The event $|\bar{\psi}_t^j - \eta_t^j| > \varepsilon_1/N$ is the probability that the empirical frequency of a Bernoulli process is far from its mean. By Cramér's Theorem (Shiryaev (1987, p.68)), therefore,

$$\Pr\left(|\bar{\psi}_t^j - \bar{\eta}_t^j| > \varepsilon_1/N\right) \leq 2e^{-2t\varepsilon^2/J^2}.$$

Using this bound in (4), we have

$$P^\theta(\|\hat{\phi}_t M_1 - \hat{\psi}_t\| \leq \varepsilon_1 | h_{1t}) > 1 - 2Je^{-2t\varepsilon^2/J^2}.$$

Becuase this formula holds for any history h_{1t}, we can choose t large enough so that the right-hand side is less than ε_2 and the statement of the Lemma follows. ■

3.3 Common Learning

Our main result establishes convergence to common q-belief.

Proposition 3 *Under Assumptions 2 and 3, the agents commonly learn* Θ*.*

Proof: We fix an arbitrary state θ and define a sequence of events F_t and show that for large enough t it has the three requisite properties. First, F_t will have high probability conditional on state θ. Second, each agent will assign high conditional probability to θ at any history consistent with F_t. Finally, F_t will be p-evident for high p.

Take $\delta > 0$ and $0 < \beta < 1$ as given by Lemma 3. Pick $\varepsilon > 0$ such that $r\delta < \delta - 2n\varepsilon$ where r and n are given by Lemma 1. For each date t, we define the event F_t as follows. First, for each $k \in 1, \ldots K$

$$F_{1t}^k(0) = \{\|\hat{\phi}_t^k - \phi_\theta^k\| < \delta\} \qquad F_{2t}^k(0) = \{\|\hat{\psi}_t^k - \psi_\theta^k\| < \delta\} \tag{5}$$

Next, for any $l \in \{1, \ldots, n\}$ and for each k,

$$F_{1t}^k(2l-1) = \{\|\hat{\phi}_t^k (M_{12})^{l-1} M_1 - \psi_\theta^k\| < \delta - (2l-1)\varepsilon\}$$
$$F_{1t}^k(2l) = \{\|\hat{\phi}_t^k (M_{12})^{l} - \phi_\theta^k\| < \delta - 2l\varepsilon\}$$

and likewise for agent two:

$$F_{2t}^k(2l-1) = \{\|\hat{\psi}_t^k (M_{21})^{l-1} M_2 - \phi_\theta^k\| < \delta - (2l-1)\varepsilon\}$$
$$F_{2t}^k(2l) = \{\|\hat{\psi}_t^k (M_{21})^{l} - \psi_\theta^k\| < \delta - 2l\varepsilon\}$$

and write

$$F_{1t}^k = \cap_{l=0}^{2n-1} F_{1t}^k(l) \qquad F_{2t}^k = \cap_{l=0}^{2n-1} F_{2t}^k(l)$$

Finally, the events

$$G_{1t} = \{\beta < \frac{\hat{\phi}_t(R_k)}{\phi_\theta(R_k)} < \beta^{-1}\} \qquad G_{2t} = \{\beta < \frac{\hat{\psi}_t(R_k)}{\psi_\theta(R_k)} < \beta^{-1}\}$$

and we define $F_{1t} = \cap_k F_{1t}^k \cap G_{1t}$ and $F_{2t} = \cap_k F_{2t}^k \cap G_{2t}$, and $F_t = F_{1t} \cap F_{2t}$.

In view of Lemma 3 it follows immediately from the definition that for any $q < 1$, we have $F_t \subset B_t^q(\theta)$ for all t sufficiently large. Next we show that $P^\theta(F_t) > q$ for large enough t. By Lemma 2 it is enough to show that there is an open set U of signal frequencies such that $U \subset F_t$ for every t.

First, note that for each k, the mappings

$$\hat{\phi}_t^k \to \|\hat{\phi}_t^k - \phi_\theta^k\|$$

and

$$\hat{\phi}_t(R_k) \to \frac{\hat{\phi}_t(R_k)}{\phi_\theta(R_k)}$$

are continuous in $\hat{\phi}_t$ and take values 0 and 1 respectively when $\hat{\phi}_t = \phi_\theta$. thus, there exists ε_0 such that

$$\{\|\hat{\phi}_t - \phi_\theta\| < \varepsilon_0\} \subset \left[F_{1t}^0 \cap G_{1t}\right]$$

and similarly for player two. Next, let[8]

$$z = \min\{\|M_1\|, \|M_2\|, \|M_{12}\|, \|M_{21}\| \ldots, \|(M_{12})^n\|, \|(M_{21})^n\|\}$$

Then, for example,

$$\{\|\hat{\phi}_t^k - \phi_\theta^k\| < \frac{\delta - (2l-1)\varepsilon}{z}\} \subset F_{1t}^k(2l-1)$$

Thus F_t is an intersection of sets of frequencies, each of which includes an open neighborhood of $(\phi_\theta, \psi_\theta)$. The intersection of these neighborhoods is included in F_t and by Lemma 2 has probability approaching one conditional on state θ.

Finally we show that for any q, F_t is q-evident when t is sufficiently large. The idea is the following. The event $\cap_k F_{2t}^k(1)$ implies that agent two expects that $\cap_k F_{1t}^k(0)$ holds. Moreover, given G_{2t} he also expects that G_{1t} holds because in state θ, $G_{1t} = G_{2t}$. This is

[8] Recall that for a linear operator L, the norm $\|L\|$ is defined by $\|L\| = \sup_x \frac{\|Lx\|}{\|x\|}$.

becuase, by our definition of recurrent classes, in state θ whenever two sees a signal j in R_k, it is certain that one has seen a signal i in R_k. We then see by induction that $\cap_k F^k_{1t}(l)$ implies that agent one expects that $\cap_k F^k_{2t}(l-1)$ holds.

However, F_t yields by direct argument only these statements of length $2n-1$ and less: the events $F^k_{1t}(2n)$ and $F^k_{2t}(2n)$ are not included in the definition of F_t. At this point, we employ Lemma 1 to show that these events are nevertheless *implied* by F_t, closing the cycle.

$$F_t \subset F^k_{1t}(2n) \cap F^k_{2t}(2n) \quad \text{for all } k \text{ and } t. \tag{6}$$

When we note that $\phi^k_\theta(M_{12})^n = \phi^k_\theta$ and $\psi^k_\theta(M_{21})^n = \psi^k_\theta$, the above follows immediately from Lemma 1 and our choice of ε and n.

Thus, F_{1t} implies that agent one "expects" that F_{2t} holds. Formally, when F_{1t} holds, and agent one's expectation of the frequencies observed by agent two is approximately correct, then also F_{2t} holds as well.

$$F_{1t} \cap \{\|\hat{\phi}_t M_1 - \hat{\psi}_t\| < \frac{\varepsilon \min_j \psi_\theta(j)}{\beta z}\} \subset \cap_k F^k_{2t} \tag{7}$$

We now set out to prove (7). Fix k. First, for each $l = 1, \ldots, n$

$$\begin{aligned}
F^k_{1t}(2l) \cap \{\|\hat{\phi}^k_t M_1 - \hat{\psi}^k_t\| < \frac{\varepsilon}{z}\} &\subset F^k_{1t}(2l) \cap \{\|\hat{\phi}^k_t (M_{12})^l - \hat{\psi}^k_t (M_{21})^{l-1} M_2\| < \varepsilon\} \\
&= \{\|\hat{\phi}^k_t (M_{12})^l - \phi^k_\theta\| < \delta - 2l\varepsilon\} \cap \{\|\hat{\phi}^k_t (M_{12})^l - \hat{\psi}^k_t (M_{21})^{l-1} M_2\| < \varepsilon\} \\
&\subset \{\|\hat{\psi}^k_t (M_{21})^{l-1} M_2 - \phi^k_\theta\| < \delta - (2l-1)\varepsilon\} \\
&= F^k_{2t}(2l-1)
\end{aligned}$$

where the last inclusion is a consequence of the triangle inequality. Similarly $F^k_{1t}(2l-1) \cap \{\|\hat{\phi}^k_t M_1 - \hat{\psi}^k_t\| < \frac{\varepsilon}{z}\} \subset F^k_{2t}(2(l-1))$. Thus, $F_{1t} \cap \bigcap_k \{\|\hat{\phi}^k_t M_1 - \hat{\psi}^k_t\| < \frac{\varepsilon}{z}\} \subset \cap_k F^k_{2t}$. We can now prove (7) by showing that for all k

$$G_{1t} \cap \{\|\hat{\phi}_t M_1 - \hat{\psi}_t\| < \frac{\varepsilon \min_j \psi_\theta(j)}{\beta z}\} \subset \{\|\hat{\phi}^k_t M_1 - \hat{\psi}^k_t\| < \frac{\varepsilon}{z}\} \tag{8}$$

To that end, let M_1(j) denote the jth column of M_1 and

$$\|\hat{\phi}^k_t M_1 - \hat{\psi}^k_t\| = \sum_{j \in R_k} \left| \frac{\hat{\phi}_t}{\hat{\phi}_t(R_k)} \cdot M_1(j) - \frac{\hat{\psi}_t(j)}{\hat{\psi}_t(R_k)} \right|$$

which, because $\hat{\phi}_t(R_k) = \hat{\psi}_t(R_k)$

$$= \frac{1}{\hat{\phi}_t(R_k)} \sum_{j \in R_k} \left| \hat{\phi}_t M_1(j) - \hat{\psi}_t(j) \right|$$
$$\leq \frac{1}{\hat{\phi}_t(R_k)} \|\hat{\phi}_t M_1 - \hat{\psi}_t\|$$

and when G_{1t} holds,

$$< \frac{\beta}{\min_j \psi_\theta(j)} \|\hat{\phi}_t M_1 - \hat{\psi}_t\|$$

which proves (8) and hence (7).

We can now conclude the proof of q-evidence. Pick $p < 1$ so that $p^2 + p - 1 > q$. We have already shown that for all sufficienty large t,

$$F_{1t} \subset B_{1t}^p(\theta). \tag{9}$$

Hence by Lemma 4,

$$F_{1t} \subset B_{1t}^{p^2}(\{\|\hat{\phi}_t M_1 - \hat{\psi}_t\| < \frac{\varepsilon}{\beta z}\})$$

and clearly

$$F_{1t} \subset B_{1t}^{p^2}(F_{1t} \cap \{\|\hat{\phi}_t M_1 - \hat{\psi}_t\| < \frac{\varepsilon}{\beta z}\})$$

which by (7) implies

$$F_{1t} \subset B_{1t}^{p^2}(F_{1t} \cap \bigcap_k F_{2t}^k) \tag{10}$$

In state θ, agent one observes a signal in R_k if and only if agent two does as well. It follows that $G_{1t} \cap \theta = G_{2t}$. Together with (9), this implies

$$F_{1t} \subset G_{1t} \subset B_{1t}^p(G_{2t}).$$

Putting together (10) and (3.3) we have

$$F_{1t} \subset B_{1t}^{p^2+p-1}(F_t)$$

for all sufficiently large t. A similar argument applies for agent two and thus $F_t \subset B_t^{p^2+p-1}(F_t) \subset B_t^q(F_t)$. ■

4 A Counter Example to Common Learning

In this counter example there will be two equally likely parameter values, $\theta \in \{0,1\}$, and two agents, with agent 1 perfectly informed about the parameter. Nevertheless, there will not be common learning.

The idea in the example is to ensure that the probability agent 2 attaches to the event $\theta = 1$ is contained in the set $p_t \in \{t^{-1}, 1 - t^{-1}\}$ at period t, but that agent 1 is very confused about which of these beliefs agent 2 has. This confusion comes from a signal process that mimics the characteristics of the faulty message exchange in Rubinstein's (1989) email game. In each period of our example, agent 1 receives a signal from the set $\{0, 1, \ldots\}$, while agent 2 either receives the same signal or a signal one higher. We can interpret these signals as the number of messages a agent has received, in a process in which agent 1 first sends a message to agent 2, each successful message prompts a confirmation, and each message is lost with probability β_t.

The remainder is routine. The signal structure ensures that regardless of the realized signal profile, one of the agents will assign probability close to $\frac{1}{2}$ to the event that the other agent assigns probability close to $\frac{1}{2}$ to the event that ... agent 2 assigns probability t^{-1} to state 1. This will be true even when both agents are almost certain the state is 1.

In constructing this example, we rely heavily on the fact that (in contrast to the sufficient conditions for common learning presented in the preceding section) the signals are generated by a Markov process with a countable number of states. In essence, the state space is $\cup_{t=1}^{\infty}\{t^{-1}, 1 - t^{-1}\}$).

One might conjecture that if the space of signals were finite, then the uncountable number of states in the generating Markov process would lose their force and common learning could be ensured (given individual learning). However, the example can be modified to have only two signals at each date, with the agents still not commonly learning. The idea is to consider longer and longer finite truncations of the Rubinstein process, and to have these processes unfold one message per period. With this modification, common learning fails, though in a less dramatic way. In most periods, a message will get through and beliefs will be almost common-knowledge. However, with probability 1 there will be a future date at which a message gets lost, in which case the agents will be nowhere close to common belief. Hence, even a limitation to a finite message set does not suffice for common learning if the Markov process generating the messages has infinite states.

We will proceed by inductively defining the signals the agents receive and the resulting priors. It simplifies the notation to assume that the parameter value is realized in period 0 and agent 1 is informed of the parameter value in period 1. Signals begin to arrive, and we begin to track of the agents' posteriors, in period 2.

Let $p(h_{2t})$ denote agent 2's posterior that the parameter is $\theta = 1$ given the history of signals h_{2t}. We will construct a signalling structure to ensure that in period $t > 1$ agent 2's posteriors are contained in the set $\{t^{-1}, 1 - t^{-1}\}$.

In period 2, $p(h_{22}) = 1/2 \in \{t^{-1}, 1 - t^{-1}\}$. Now suppose that in period $t > 2$ agent 2's priors are in $\{t^{-1}, 1 - t^{-1}\}$. We will define the signal structure so that in period $t + 1$ agent 2's priors are either $(t + 1)^{-1}$ or $1 - (t + 1)^{-1}$. In period t, agent 1 will observe a signal that is a non-negative integer $x_{1t} \in \{0, 1, ...\}$ and agent 2 will observe a correlated signal that is either the same as agent 1's or one greater, so that $x_{2t} \in \{x_{1t}, x_{1t} + 1\}$. The joint distribution of these signals depends both upon the state and upon agent 2's current posterior, p, and a time-dependent parameter $\beta_t = (t^3 - 1)^{-1}$, as follows:

$$P^\theta((x_{1t}, x_{2t}) \mid p(h_{2t}) = p) = \begin{cases} 1 - \gamma_\theta^t(p), & \text{if } (x_{1t}, x_{2t}) = (0, 0); \\ \gamma_\theta^t(p)\beta_t(1 - \beta_t)^{x_{2t}+x_{1t}-1}, & \text{if } x_{2t} \in \{x_{1t}, x_{1t} + 1\}; \\ 0, & \text{otherwise.} \end{cases}$$

We will choose the functions $\gamma_\theta^t(\cdot)$ for $\theta = 0, 1$ so that when agent 2 observes $x_{2t} = 0$ he will revise his beliefs to $p(h_{2t+1}) = (t + 1)^{-1}$. This is ensured by the condition:

$$\frac{p(1 - \gamma_1^t(p))}{p(1 - \gamma_1^t(p)) + (1 - p)(1 - \gamma_0^t(p))} = \frac{1}{t + 1}. \tag{11}$$

In addition, we choose the functions $\gamma_\theta^t(\cdot)$ to ensure that when agent 2 observes $x_{2t} > 0$ he will revise his beliefs to $p(h_{2t+1}) = 1 - (t + 1)^{-1}$. This is ensured by the condition:

$$\frac{p\gamma_1^t(p)}{p\gamma_1^t(p) + (1 - p)\gamma_0^t(p)} = 1 - \frac{1}{t + 1}. \tag{12}$$

Straightforward algebra shows that for $p \in \{t^{-1}, 1 - t^{-1}\}$ there exists solutions $\gamma_\theta^t(p) \in ((t + 1)^{-3}, 1 - (t + 1)^{-3})$, for $\theta = 0, 1$, to the equations (11), (12).

This describes what happens to agent 2's beliefs, but does not yet confirm that agent 2 actually learns the state given this structure. However, this follows immediately from the fact that his posteriors in period t have the support $\{t^{-1}, 1 - t^{-1}\}$. When $\theta = 1$, there cannot be positive probability attached to the limiting posterior being zero, and conversely in state zero.[9] This completes our description of agent 2's beliefs.

[9]Hence, when $\theta = 1$, the values γ_θ^t must converge to one, so that the signal x_{2t} and hence posterior t^{-1} becomes increasingly unlikely.

Now consider agent 1's beliefs. Her posteriors about the state are trivial, since she knows the value of θ. Let us focus on agent 1's information about agent 2's posterior p. Since agent 2's beliefs are concentrated on $\{t^{-1}, 1-t^{-1}\}$, it suffices to describe 1's beliefs about 2's beliefs to define $\phi \equiv P(p_t = 1-t^{-1}|\mathcal{H}_{1t})$. Conditional on ϕ and the parameter θ, let $\mu^t(\phi, \theta)$ denote agent 1's posterior probability that in date t agent 2 will see no signals, $x_{2t} = 0$, and (as a result) agent 2 has a posterior $p_{t+1} = (t+1)^{-1}$. We can calculate $\mu^t(\phi, \theta)$ as

$$\mu^t(\phi, \theta) = \phi\left(1 - \gamma_\theta^t\left(1 - \frac{1}{t}\right)\right) + (1-\phi)\left(1 - \gamma_\theta^t\left(\frac{1}{t}\right)\right).$$

The bounds $\gamma_\theta^t(p) \in ((t+1)^{-3}, 1-(t+1)^{-3})$ imply that $\mu^t(\phi, \theta) > (t+1)^{-3}$.

Suppose that agent 1 sees no signals himself, $x_{1t} = 0$, then the updated probability that agent sees $x_{2t} = 0$ is given by

$$\frac{\mu^t(\phi, \theta)}{\mu^t(\phi, \theta) + (1-\mu^t(\phi, \theta))\beta_t} \geq \frac{1}{1 + (t^3 - 1)\beta_t} = \frac{1}{2},$$

where the first inequality substitutes our bound on $\mu^t(\phi, \theta)$ and the equality substitutes for β_t. Thus, regardless of the previous history, conditional on $x_{1t} = 0$ agent 1 assigns probability at least $\frac{1}{2}$ to the event that agent 2 assigns probability $\frac{1}{2}$ to state 1.

Finally, to show that common learning fails, consider an arbitrary signal pair (x_{1t}, x_{2t}) and the higher order beliefs of agent $x_{1t} + x_{2t} + 1 \pmod 2$. For example, suppose this is agent 2. Whatever the current agent-2 posterior p_t and whatever the parameter θ, the conditional probability that agent 1 has seen $x_{1t} = x_{2t}$ is

$$\frac{\gamma_\theta^t(p_t)(1-\beta_t)^{2x_{2t}-1}\beta_t}{\gamma_\theta^t(p_t)(1-\beta_t)^{2x_{2t}-1}\beta_t + \gamma_\theta^t(p_t)(1-\beta_t)^{2x_{2t}}\beta_t} = \frac{1}{2-\beta_t} \geq \frac{1}{2}.$$

Therefore, regardless of any other beliefs, agent 2 assigns probability at least $\frac{1}{2}$ to $x_{1t} = x_{2t}$. Likewise, agent 1 would assign probability at least $\frac{1}{2}$ to agent 2 having observed $x_{2t} = x_{1t} - 1$. Now by the usual contagion argument, we conclude that with probability 1 there exists an integer k such that the k-iterated statement "agent i assigns probability at least $\frac{1}{2}$ to agent $3-i$ assigning probability at least $\frac{1}{2}$ to $\ldots$, agent 2 assigning probability $\frac{1}{t}$ to state zero" is true at date t.

References

Bergemann, D. and J. Välimäki (1997) "Market Diffusion with Two-Sided Learning", *RAND J. Econ* **28**, 773–795.

Billingsley, P. (1986) *Probability and Measure, Second Edition*, John Wiley.

Bolton, P. and C. Harris (1999) "Strategic Experimentation", *Econometrica* **67**, 349–374.

Brémaud, P. (1999) *Markov Chains*, Springer-Verlag, New York.

Cripps, M. W., G. M Mailath and L. Samuelson (2004) "Disappearing Private Reputations in Long Run Relationships", PIER Working Paper 04-031, Department of Economics, University of Pennsylvania.

Cover, T.M. and J. A. Thomas (1991) *Elements of Information Theory*, John Wiley.

den Hollander, F. (2000) *Large Deviations*, American Mathematical Society.

Fill J.A. (1991) "Eigenvalue Bounds on Convergence to Stationarity for Nonreversible Markov Chains with an Application to the Exclusion Process", *Annals of Applied Probability* **1**, 62–87.

Foster D. and R. Vohra (1999) "Regret in the On-Line Decision Problem", *Games and Economic Behavior* **29**, 7–35.

Gossner, O. and N. Vieille (2003) "Strategic Learning in Games with Symmetric Information", *Games and Economic Behavior* **42**, 25-47.

Hoeffding, W. (1956) "On the Distribution of the Number of Successes in Independent Trials", *Annals of Mathematical Statistics* **27**, 713–721.

Monderer, D. and D. Samet (1989) "Approximating Common Knowledge with Common Beliefs", *Games and Economic Behavior* **1**, 170–190.

Moscarini, G. and F. Squintani (2004) "Competitive Experimantation with Private Information", *Cowles Foundation* DP 1489.

Rosenberg, D., E. Solan and N. Vieille (2005) *"Timing Games with Informational Externalities"*, mimeo MEDS .

Rubinstein, A. (1989) "The Electronic Mail Game: Strategic Behavior Under Almost Common Knowledge", *American Economic Reveiw* **79**, 385–391.

Rustichini, A. (1999) "Minimizing Regret: The General Case", *Games and Economic Behavior* **29**, 224–243.

Shiryaev, A.N. (1987) *Probability*, Springer-Verlag, New York.

Wiseman, T. (2005) "A Folk Theorem for Games with Unknown Payoff Distributions", *Econometrica* forthcoming.