Common Learning*

Martin W. Cripps† Jeffrey C. Ely‡ George J. Mailath§
Larry Samuelson¶

August 22, 2006

Abstract

Consider two agents who learn the value of an unknown parameter by observing a sequence of *private* signals. The signals are independent and identically distributed across time but not necessarily agents. Does it follow that the agents will *commonly* learn its value, i.e., that the true value of the parameter will become (approximate) common-knowledge? We show that the answer is affirmative when each agent's signal space is finite and show by example that common learning can fail when observations come from a countably infinite signal space.

Keywords: Common learning, common belief, private signals, private beliefs.

JEL Classification Numbers: D82, D83.

*Mailath and Samuelson thank the National Science Foundation (grants SES-0350969 and SES-0549946, respectively) for financial support.

†University College London and Washington University in St. Louis, cripps@olin.wustl.edu
‡Northwestern University, jeffely@northwestern.edu
§Yale University and University of Pennsylvania, george.mailath@yale.edu
¶University of Wisconsin, LarrySam@ssc.wisc.edu

1 Introduction

Consider two agents who learn the value of an unknown parameter by observing a sequence of *private* signals. The signals are independent and identically distributed across time but not necessarily agents. Does it follow that the agents will *commonly* learn its value, i.e., that the true value of the parameter will become (approximate) common-knowledge? We show that the answer is affirmative when each agent's signal space is finite and show by example that common learning can fail when observations come from a countably infinite signal space.

This is an important question for a number of reasons. Common learning is precisely the condition that ensures efficient outcomes in dynamic coordination problems in which agents learn the appropriate course of action privately over time. For example, suppose the two agents have the possibility of profitably coordinating on an action, but that the action depends on an unknown parameter. In every period $t = 0, 1, \ldots,$ each agent receives a signal. The agent can then choose action A, action B, or to wait (W) until the next period. Simultaneous choices of A when the parameter is θ_A or B when it is θ_B bring payoffs of 1 each. Lone choices of A or B or joint choices that do not match the parameter bring a payoff of $-c < 0$ and cause the investment opportunity to disappear. Waiting is costless. Figure 1 summarizes these payoffs.

Under what circumstances do there exist nontrivial equilibria of this investment game, i.e., equilibria in which the agents do not always wait? Choosing action A is optimal for an agent in some period t only if the agent attaches probability at least $\frac{c}{c+1} \equiv q$ to the joint event that the parameter is θ_A and the other agent chooses A. Now consider the set of histories $\mathscr{A}$ at which both agents choose A. At any such history, each agent ℓ must assign probability at least q to $\mathscr{A}$, that is $\mathscr{A}$ must be *q-evident* (Monderer and Samet, 1989). Furthermore, at any history in $\mathscr{A}$, each agent ℓ must assign probability at least q to the parameter θ_A. But this pair of conditions is equivalent to the statement that θ_A is *common q-belief*—the existence of histories at which there is common q-belief in θ_A is a necessary condition for eventual coordination in this game. Conversely, the possibility of common

	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

	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_B

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

q-belief is sufficient for a nontrivial equilibrium, as it is an equilibrium for each agent ℓ to choose A on the q-evident event on which θ_A is common q-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? It suffices that the information process be such that the parameter eventually becomes arbitrarily close to common 1-belief.

Beyond coordination problems, common learning is a potentially important tool in the analysis of dynamic games with incomplete information. In the equilibria of these games, players typically learn over time about some unknown parameter. Examples include reputation models such as Cripps, Mailath, and Samuelson (forthcoming), where one player learns the "type" of the other, and experimentation models such as Wiseman (2005), where players are learning about their joint payoffs in an attempt to coordinate on some (enforceable) target outcome. Characterizing equilibrium in these games requires analyzing not only each player's beliefs about payoffs, but also her beliefs about the beliefs of others and how these higher-order beliefs evolve. Existing studies of these models have imposed strong assumptions on the information structure in order to keep the analysis tractable. We view our research as potentially leading to some general tools for studying common learning in dynamic games.

In general, the relationship between individual and common learning is subtle. However, there are two special cases in which individual learning immediately implies common learning. When the signals are public then beliefs are trivially common-knowledge. At the opposite extreme, common learning occurs when the agents' signal processes are stochastically independent and so (conditional on the parameter) each learns nothing about the other's beliefs (Proposition 2).

Apart from these extreme cases, when the signals are private and not independent, the following difficulty must be addressed. If the signals are correlated, and if the realized signal frequencies for agent 1 (say) are sufficiently close to the population frequencies under the parameter θ, then 1 will be confident that θ is the value of the parameter. Moreover, he will be reasonably confident that 2 will have observed a frequency that leads to a similar degree of confidence in θ. However, if 1's frequency is "just" close enough to lead to some fixed degree of confidence, then 1 may not be confident that 2's realized frequency leads to a similar degree of confidence: while 2's frequency may be close to 1's frequency, it may be on the "wrong side" of the boundary for the required degree of confidence.

If the set of signals is finite, the distribution of one agent's signals, conditional on the other agent's signal, has a Markov chain interpretation. This allows us to appeal to a contraction mapping principle in our proof of common learning, ensuring that if agent 1's signals are on the "right side" of a confidence boundary then so must be 1's beliefs about 2's signals. In contrast, with a countably infinite signal space, the corresponding Markov chain interpretation lacks the relevant contraction mapping structure and common learning may fail.

While we have described the model as one in which the agents begin with a common prior over the set of parameters, we explain in Remark 3 how our analysis sheds light on agents who initially disagree but *converge* on a common belief through a process of learning. Indeed, we can allow agents to begin the process with arbitrary higher-order beliefs over the parameter space. As long as each agent attaches some minimum probability to each parameter, and this is common knowledge, the agents will commonly learn the parameter and hence approach a common posterior over the distribution of signals.

2 A Model of Multi-Agent Learning

2.1 Individual Learning

Time is discrete and periods are denoted by $t = 0, 1, 2, \ldots$. Before period zero, nature selects a parameter θ from the finite set Θ according to the prior distribution p.

For notational simplicity, we restrict attention to 2 agents, denoted $\ell = 1$ (he) and 2 (she). Our positive results (Propositions 2 and 3) hold for arbitrary finite number of agents (see Remarks 2 and 4).

Conditional on θ, a stochastic process $\zeta^\theta \equiv \{\zeta_t^\theta\}_{t=0}^\infty$ generates a signal profile $z_t \equiv (z_{1t}, z_{2t}) \in Z_1 \times Z_2 \equiv Z$ for each period t, where Z_ℓ is the set of possible period-t signals for agent $\ell = 1, 2$. For each $\theta \in \Theta$, the signal process $\{\zeta_t^\theta\}_{t=0}^\infty$ is independent and identically distributed across t. We let $\zeta_\ell^\theta \equiv \{\zeta_{\ell t}\}_{t=0}^\infty$ denote the stochastic process generating agent ℓ's signals. When convenient, we let $\{\theta\}$ denote the event $\{\theta\} \times Z^\infty$ that the parameter value is θ, and we often write θ rather than $\{\theta\}$ when the latter appears as an argument of a function.

A state consists of a parameter and a sequence of signal profiles, with the set of states given by $\Omega \equiv \Theta \times Z^\infty$. We use P to denote the measure on Ω induced by the prior p and the signal processes $(\zeta^\theta)_{\theta \in \Theta}$, and use $E[\,\cdot\,]$ to denote expectations with respect to this measure. Let P^θ denote the measure conditional on a given parameter and $E^\theta[\,\cdot\,]$ expectations with respect to this measure.

A period-t history for agent ℓ is denoted by $h_{\ell t} \equiv (z_{\ell 0}, z_{\ell 1}, \ldots, z_{\ell t-1})$. We let $H_{\ell t} \equiv (Z_\ell)^t$ denote the space of period-t histories for agent ℓ and let $\{\mathscr{H}_{\ell t}\}_{t=0}^\infty$ denote the filtration induced on Ω by agent ℓ's histories. The random variables $\{P(\theta \mid \mathscr{H}_{\ell t})\}_{t=0}^\infty$, giving agent ℓ's beliefs about the parameter θ at the start of each period, are a bounded martingale with respect to the measure P, for each θ, and so the agents' beliefs converge almost surely (Billingsley, 1979, Theorem 35.4). For any state ω, $h_{\ell t}(\omega) \in \mathscr{H}_{\ell t}$ is the agent ℓ period-t history induced by ω. As usual, $P(\theta \mid \mathscr{H}_{\ell t})(\omega)$ is often written $P(\theta \mid h_{\ell t}(\omega))$ or $P(\theta \mid h_{\ell t})$ when ω is understood.

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

$$B^q_{\ell t}(F) \equiv \{\omega \in \Omega : E[\mathbf{1}_F \mid \mathscr{H}_{\ell t}](\omega) \geq q\}.$$

Thus, $B^q_{\ell t}(F)$ is the set of states where at time t agent ℓ attaches at least probability q to event F.

Definition 1 (Individual Learning) *Agent ℓ* learns *parameter θ if conditional on parameter θ, agent ℓ's posterior on θ converges in probability to* 1*, i.e., if for each $q \in (0,1)$ there is T such that for all $t > T$,*

$$P^\theta(B^q_{\ell t}(\theta)) > q. \tag{1}$$

Agent ℓ learns Θ *if ℓ learns each $\theta \in \Theta$.*

Individual learning is equivalent to

$$\lim_{t\to\infty} P^\theta(B^q_{\ell t}(\theta)) = 1, \qquad \forall q \in (0,1). \tag{2}$$

Remark 1 We have formulated individual learning using convergence in probability rather than almost sure convergence to facilitate the comparison with common learning. Convergence in probability is in general a weaker notion than almost sure convergence. However, since $P(\theta \mid \mathscr{H}_{\ell t})$ converges almost surely to some random variable, (2) is equivalent to $P(\theta \mid \mathscr{H}_{\ell t}) \to 1$ P^θ-a.s.

♦

We assume that each agent individually learns the parameter—there is no point considering common learning if individual learning fails. Our aim is to identify the additional conditions that must be imposed to ensure not just that each agent learns the parameter, but that the agents commonly learn the parameter.

2.2 Common Learning

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

$$B_t^q(F) \equiv B_{1t}^q(F) \cap B_{2t}^q(F).$$

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

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

Hence, on $C_t^q(F)$, the event F is q-believed and this event is itself q-believed and so on. We are interested in common belief as a measure of approximate common-knowledge because, as shown by Monderer and Samet (1989), it is common belief that ensures continuity of behavior in incomplete-information games.

A related but distinct notion is that of *iterated* q-belief. The event that F is iterated q-belief is defined to be

$$I_t^q(F) \equiv B_{1t}^q(F) \cap B_{2t}^q(F) \cap B_{1t}^q B_{2t}^q(F) \cap B_{2t}^q B_{1t}^q(F) \cap \dots$$

Morris (1999, Lemma 14) shows that iterated belief is (possibly strictly) weaker than common belief:

Lemma 1 (Morris) $C_t^q(F) \subset I_t^q(F)$.

See Morris (1999, p. 388) for an example showing the inclusion can be strict.

The parameter θ is common q-belief at time t on the event $C_t^q(\theta)$. We say that the agents commonly learn 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 2 (Common Learning) *The agents* commonly learn *parameter $\theta \in \Theta$ if for each $q \in (0,1)$ there exists a T such that for all $t > T$,*

$$P^{\theta}(C_t^q(\theta)) > q.$$

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

Common learning is equivalent to

$$\lim_{t\to\infty} P^{\theta}(C_t^q(\theta)) = 1, \qquad \forall q \in (0,1).$$

Because $C_t^q(\theta) \subset B_{\ell t}^q(\theta)$, common learning implies individual learning (recall (2)).

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

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

Our primary technical tool links common q-belief and q-evidence. Monderer and Samet (1989, Definition 1 and Proposition 3) show:

Proposition 1 (Monderer and Samet) *F' is common q-belief at $\omega \in \Omega$ and 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 and only if for all $\theta \in \Theta$ and $q \in (0,1)$, there exists a sequence of events F_t and a period T such that for all $t > T$,*

(i) θ *is q-believed on F_t at time t,*

(ii) $P^{\theta}(F_t) > q$*, and*

(iii) F_t *is q-evident at time t.*

2.3 Special Cases: Perfect Correlation and Independence

We are primarily interested in private signals that are independently and identically distributed over time, but not identically or independently across agents. We begin, however, with two special cases to introduce some basic ideas.

Suppose first the signals are public, as commonly assumed in the literature. Then agent ℓ knows everything there is to know about $\hat{\ell}$'s beliefs, and we have $P(\theta \mid \mathscr{H}_{1t}) = P(\theta \mid \mathscr{H}_{2t})$ for all θ and t—and hence beliefs are always common. Individual learning then immediately implies common learning.

At the other extreme, we have independent signals. Here, the fact that agent ℓ learns *nothing* about agent $\hat{\ell}$'s signals ensures common learning.

Proposition 2 *Suppose each agent learns* Θ *and that for each* $\theta \in \Theta$*, the stochastic processes* $\{\zeta_{1t}^{\theta}\}_{t=0}^{\infty}$ *and* $\{\zeta_{2t}^{\theta}\}_{t=0}^{\infty}$ *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. We let $F_t \equiv \{\theta\} \cap B_t^{\sqrt{q}}(\theta)$ and verify that F_t satisfies the sufficient conditions for common learning provided in Corollary 1

(i) Because $F_t \subset B_t^{\sqrt{q}}(\theta) \subset B_t^q(\theta)$, parameter θ is q-believed on F_t at time t.

(ii) To show $P^\theta(F_t) > q$, note that independence implies $P^\theta(F_t) = \prod_\ell P^\theta(B_{\ell t}^{\sqrt{q}}(\theta))$. By (1), we can choose T sufficiently large that $P^\theta(B_{\ell t}^{\sqrt{q}}(\theta)) > \sqrt{q}$ for all ℓ and all $t > T$ and hence $P^\theta(F_t) > q$.

(iii) To show that F_t is q-evident, we must show that $F_t \subset B_{\ell t}^q(F_t)$ for $\ell = 1,2$. By construction, $F_t \subset B_{\ell t}^{\sqrt{q}}(\theta)$. Since $B_{\ell t}^{\sqrt{q}}(\theta) \in \mathscr{H}_{\ell t}$, on F_t agent ℓ attaches probability 1 to the state being in $B_{\ell t}^{\sqrt{q}}(\theta)$ and we have

$$\begin{aligned}
B_{\ell t}^q(F_t) =& \{\omega : E[\mathbf{1}_{B_{\ell t}^{\sqrt{q}}(\theta)} \mathbf{1}_{B_{\hat{\ell} t}^{\sqrt{q}}(\theta)\cap\{\theta\}} \mid \mathscr{H}_{\ell t}] \geq q\} \\
=& \{\omega : \mathbf{1}_{B_{\ell t}^{\sqrt{q}}(\theta)} E[\mathbf{1}_{B_{\hat{\ell} t}^{\sqrt{q}}(\theta)\cap\{\theta\}} \mid \mathscr{H}_{\ell t}] \geq q\} \\
=& B_{\ell t}^{\sqrt{q}}(\theta) \cap B_{\ell t}^q(B_{\hat{\ell} t}^{\sqrt{q}}(\theta) \cap \{\theta\}).
\end{aligned}$$

Thus, it suffices to show that on the set F_t, agent ℓ attaches at least probability q to the event $B^{\sqrt{q}}_{\hat{\ell}t}(\theta) \cap \{\theta\}$, $\hat{\ell} \neq \ell$. As above, (1) allows us to choose T sufficiently large that $P^{\theta}(B^{\sqrt{q}}_{\ell t}(\theta)) > \sqrt{q}$ for all ℓ and all $t > T$. The conditional independence of agents' signals implies that, given θ, agent ℓ's history is uninformative about $\hat{\ell}$'s signals, and hence $P^{\theta}(B^{\sqrt{q}}_{\hat{\ell}t}(\theta) \mid \mathscr{H}_{\ell t}) > \sqrt{q}$.[1] But, on F_t, we have $P(\theta \mid \mathscr{H}_{\ell t}) > \sqrt{q}$. Consequently, again on F_t

$$P(B^{\sqrt{q}}_{\hat{\ell}t}(\theta) \cap \{\theta\} \mid \mathscr{H}_{\ell t}) = P^{\theta}(B^{\sqrt{q}}_{\hat{\ell}t}(\theta) \mid \mathscr{H}_{\ell t}) P(\theta \mid \mathscr{H}_{\ell t}) > q, \tag{3}$$

and we have the desired result. ■

Remark 2 (Arbitrary finite number of agents) The proof of Proposition 2 covers an arbitrary finite number of agents once we redefine F_t as $\{\theta\} \cap B^{\sqrt[n]{q}}_t(\theta)$, where n is the number of agents.

♦

The role of independence in this argument is to ensure that agent ℓ's signals provide ℓ with no information about $\hat{\ell}$'s signals. Agent ℓ thus not only learns the parameter, but eventually thinks it quite likely that $\hat{\ell}$ has also learned the (same) parameter (having no evidence to the contrary). In addition, we can place a lower bound, uniform across agent ℓ's histories, on how confident agent ℓ is that $\hat{\ell}$ shares ℓ's confidence in the parameter (see (3)). This suffices to establish common learning.

One would expect common learning to be more likely the more information ℓ has about $\hat{\ell}$, so that ℓ has a good idea of $\hat{\ell}$'s beliefs. When signals are correlated, ℓ's signals will indeed often provide useful information about $\hat{\ell}$'s, accelerating the rate at which ℓ learns about $\hat{\ell}$ and reinforcing common learning. Clearly this

[1] Since conditional probabilities are only unique for P-almost all states, the set F_t depends upon the choice of version of the relevant conditional probabilities. In the proof, we have selected the constant function $P^{\theta}(B^{\sqrt{q}}_{\hat{\ell}t}(\theta))$ as the version of $P^{\theta}(B^{\sqrt{q}}_{\hat{\ell}t}(\theta) \mid \mathscr{H}_{\ell t})$. For other versions of conditional probabilities, the definition of F_t must be adjusted to exclude appropriate zero probability subsets.

is the case for perfect correlation, but perhaps surprisingly, intermediate degrees of correlation can generate information that may disrupt common learning. The danger is that agent 1 may have observed signal frequencies "just" close enough to lead to some fixed degree of confidence in the value of the parameter, but in the process may have received evidence that 2's frequencies are on the "wrong side" of her corresponding boundary, even though quite close to it. We show this by example in Section 4.

3 Sufficient Conditions for Common Learning

3.1 Common Learning

For our positive result, we assume that the signal sets are finite.

Assumption 1 (Finite Signal Sets) *Agents* 1 *and* 2 *have finite signal sets,* I *and* J *respectively.*

We use I and J to also denote the cardinality of sets I and J, trusting the context will prevent confusion.

We denote the probability distribution of the agents' signals conditional on θ by $(\pi^\theta(ij))_{i\in I, j\in J} \in \Delta(I\times J)$. Hence, $\pi^\theta(ij)$ is the probability that $(z_{1t}, z_{2t}) = (i, j)$ for parameter θ and every t. For each $\theta \in \Theta$, let

$$\begin{aligned} & I^\theta \equiv \{i \in I : \textstyle\sum_j \pi^\theta(ij) > 0\} \\ \text{and} \qquad & J^\theta \equiv \{l \in J : \textstyle\sum_i \pi^\theta(ij) > 0\} \end{aligned}$$

be the sets of signals that appear with positive probability under parameter θ. Denote $\left(\pi^\theta(ij)\right)_{i\in I^\theta, j\in J^\theta}$ by Π^θ.

We define $\phi^\theta(i) \equiv \sum_j \pi^\theta(ij)$ to denote the marginal probability of agent 1's signal i and $\psi^\theta(j) = \sum_i \pi^\theta(ij)$ to denote the marginal probability of agent 2's signal j. We let $\phi^\theta = (\phi^\theta(i))_{i\in I^\theta}$ and $\psi^\theta = (\psi^\theta(j))_{j\in J^\theta}$ be the row vectors of expected frequencies of the agents' signals under parameter θ. Notice that we restrict attention to those signals that appear with positive probability under parameter θ in defining the vectors ϕ^θ and ψ^θ.

Given Assumption 1, the following is equivalent to (1).

Assumption 2 (Individual Learning) *For every pair θ and θ', the marginal distributions are distinct, i.e. $\phi^\theta \neq \phi^{\theta'}$ and $\psi^\theta \neq \psi^{\theta'}$.*

Our main result is:

Proposition 3 *Under Assumption 1 and Assumption 2, the agents commonly learn Θ.*

Remark 3 (The role of the common prior and agreement on π^θ) Though we have conserved on notation by presenting Proposition 3 in terms of a common prior, the analysis applies with little change to a setting where the two agents have different but commonly known priors. Indeed, the priors need not be commonly known—it is enough that there be a commonly known bound on the minimum probability any parameter receives in each agent's prior. We can modify Lemma 3 to still find a neighborhood of signals frequencies in which every "type" of agent *i* will assign high probability to the true parameter. The rest of the proof is unchanged.

Our model also captures settings in which the agents have different beliefs about the conditional signal-generating distributions $(\pi^\theta(ij))_{i\in I, j\in J}$. In particular, such differences of opinion can be represented as different beliefs about a parameter ϕ^θ that determines the signal-generating process given θ. The model can then be reformulated as one in which agents are uncertain about the joint parameter (θ, ϕ^θ) (but know the signal-generating process conditional on this parameter) and our analysis applied.

Our work is complementary to Acemoglu, Chernozhukov, and Yildiz (2006), who consider environments in which even arbitrarily large samples of common data may not reconcile disagreements in agents' beliefs. Acemoglu, Chernozhukov, and Yildiz (2006) stress the possibility that the agents in their model may not know the signal-generating process $(\pi^\theta(ij))_{i\in I, j\in J}$, but we have just argued that this is not an essential distinction in our context. The key difference is that the signal-generating processes considered by Acemoglu, Chernozhukov, and Yildiz (2006) need not suffice for individual learning. In our context, it is unsurprising that common learning need not hold when individual learning fails.

♦

3.2 Outline of the Proof

Let $f_t(ij)$ denote the number of periods in which agent 1 has received the signal i and agent 2 received the signal j before period t. Defining $f_{2t}(j) \equiv \sum_i f_t(ij)$ and $f_{1t}(i) \equiv \sum_j f_t(ij)$, the realized frequencies of the signals are given by the row vectors $\hat{\phi}_t \equiv (f_{1t}(i)/t)_{i \in I}$ and $\hat{\psi}_t \equiv (f_{2t}(j)/t)_{j \in J}$. Finally, let $\hat{\phi}_t^{\theta} = (f_{1t}(i)/t)_{i \in I^{\theta}}$ denote the realized frequencies of the signals that appear with positive probability under parameter θ, with a similar convention for $\hat{\psi}^{\theta}$.

The main idea of the proof is to classify histories in terms of the realized frequencies of signals observed and, for given $q \in (0,1)$, to identify events such as $B_{1t}^q(\theta)$ and $B_{1t}^q(B_{2t}^q(\theta))$ with events exhibiting the appropriate frequencies.

Section 3.4 develops the tools required for working with frequencies. The analysis begins with an open neighborhood of frequencies within which each agent will assign high probability to parameter θ. Indeed, Lemma 3 shows that there is a $\delta > 0$ so that whenever 1's observed frequency distribution $\hat{\phi}_t$ is within a distance δ of ϕ^{θ}, his marginal signal distribution under θ, the posterior probability he assigns to θ approaches one over time. Let $F_{1t}(0)$ denote this δ-neighborhood of ϕ^{θ},

$$F_{1t}(0) \equiv \left\{ \omega : \left\| \hat{\phi}_t^{\theta} - \phi^{\theta} \right\| < \delta \right\}.$$

By the weak law of large numbers, the probability under θ that the realized frequency falls in $F_{1t}(0)$ converges to one (Lemma 4).

Next, we consider the set of frequencies that characterize the event that 1 assigns high probability to θ *and* to 2 assigning high probability to θ. This involves three steps.

STEP 1: Since the event we are interested in implies that 1 assigns high probability to θ, we can approximate 1's beliefs about 2 by his beliefs conditional on θ being the true parameter.

STEP 2: We now introduce an object that plays a central role in the proof, the $I^\theta \times J^\theta$ matrix M_1^θ whose ijth element is $\frac{\pi^\theta(ij)}{\phi^\theta(i)}$, i.e. the conditional probability under parameter θ of signal j given signal i. At any date t, when agent 1 has realized frequency distribution $\hat{\phi}_t$, his estimate (expectation) of the frequencies observed by agent 2 *conditional on parameter* θ is given by the matrix product

$$\hat{\phi}_t^\theta M_1^\theta.$$

The corresponding matrix for agent two, denoted M_2^θ, is the $J^\theta \times I^\theta$ matrix with jith element $\frac{\pi^\theta(ij)}{\psi^\theta(j)}$.

We now make a key observation relating ϕ^θ, ψ^θ, M_1^θ, and M_2^θ. Let D_1^θ be the $I^\theta \times I^\theta$ diagonal matrix with ith diagonal element $(\phi^\theta(i))^{-1}$ and let e be a row vector of 1's. It is then immediate that

$$\phi^\theta M_1^\theta = \phi^\theta D_1^\theta \Pi^\theta = e\Pi^\theta = \psi^\theta. \tag{4}$$

A similar argument implies

$$\psi^\theta M_2^\theta = \phi^\theta. \tag{5}$$

Note that the product $\hat{\phi}_t^\theta M_1^\theta M_2^\theta$ gives agent 1's expectation of agent 2's expectation of the frequencies observed by agent 1 (conditional on θ). Moreover, $M_{12}^\theta \equiv M_1^\theta M_2^\theta$ is a Markov transition matrix on the set I^θ of signals for agent 1.[2] Section 3.3 collects some useful properties of this Markov process.

From (4), the continuity of the linear map M_1^θ implies that whenever 1's frequencies are in a neighborhood of ϕ^θ, we are assured that 1 expects that 2's frequencies are in the neighborhood of ψ^θ, and hence that 2 assigns high probability to θ. Of course, "expecting" that 2 assigns high probability to θ is not the same as

[2] This perspective is inspired by Samet (1998).

assigning high probability to it, and we must account for the error in 1's estimate of 2's frequencies, leading to the third step.

STEP 3: We need to bound the probability of any large error in this estimate. Lemma 5 shows that conditional on θ, there is a time T after which the probability that 2's realized frequencies are more than some given ε away from 1's estimate ($\hat{\phi}_t^\theta M_1^\theta$) is less than ε. A crucial detail here is that this bound applies uniformly across all histories for 1. There is thus a neighborhood of ϕ^θ such that if 1's frequency $\hat{\phi}_t^\theta$ falls in this neighborhood for sufficiently large t, then agent 1 assigns high probability to the event that 2 assigns high probability to θ. Let $F_{1t}(1)$ denote this neighborhood, which we can equivalently think of as a neighborhood of ψ^θ into which $\hat{\phi}_t^\theta M_1^\theta$ must fall, that is,

$$F_{1t}(1) \equiv \left\{\omega : \left\|\hat{\phi}_t^\theta M_1^\theta - \psi^\theta\right\| < \delta - \varepsilon\right\},$$

where ε is small and determined below.

For sufficiently large t, the intersection $F_{1t}(0) \cap F_{1t}(1) \equiv F_{1t}$ is contained in $B_{1t}^q(B_{1t}^q(\theta) \cap B_{2t}^q(\theta)) = B_{1t}^q(\theta) \cap B_{1t}^q(B_{2t}^q(\theta))$, providing the first steps toward common learning. However, in order to show q-common belief, we need to show that *all* orders of iterated (joint) q-belief can be obtained on neighborhoods of ϕ^θ and ψ^θ, and common learning requires in addition these neighborhoods have high probability. Rather than attempting a direct argument, we apply Corollary 1.

Suppose (for the sake of exposition) that every element of M_{12}^θ is strictly positive. In that case, M_{12}^θ is a contraction when viewed as a mapping on ΔI^θ, a property critical to our argument. Hence, for some $r \in (0,1)$, if 1's frequencies are within δ of ϕ^θ, then 1's prediction of 2's prediction of 1's frequencies are within $r\delta$ of ϕ^θ. Consequently, iterating B_{1t}^q and B_{2t}^q does not lead to "vanishing" events.

Fix θ and a period t large. A natural starting point would be to try $F_{1t} \cap F_{2t}$ (where F_{2t} is defined similarly for agent 2 to F_{1t}) as a candidate for F_t in Corollary 1. But since we also need F_t to be likely under θ, we intersect these sets with the event $\{\theta\}$ so that $F_t \equiv F_{1t} \cap F_{2t} \cap \{\theta\}$.

Observe that $\hat{\phi}_t^\theta \in F_{1t}(0)$ for all $\omega \in F_t$ by construction. It is also intuitive (and indeed true) that θ is q-believed on F_t at time t and that $P^\theta(F_t) > q$ for sufficiently large t. It remains to verify that the set F_t is q-evident at time t, that is, $F_t \subset B_t^q(F_t) = B_{1t}^q(F_t) \cap B_{2t}^q(F_t)$. It suffices to argue that

$$F_{1t} \cap \{\theta\} \subset B_{1t}^q(F_{1t} \cap F_{2t} \cap \{\theta\})$$

(the argument is symmetric for agent 2).

We first note that $F_{\ell t} \in \mathscr{H}_{\ell t}$ (i.e., agent ℓ knows the event $F_{\ell t}$ in period t). Next, a straightforward application of the triangle inequality yields

$$F_{1t}(1) \cap \hat{F}_{1t}(1) \subset F_{2t}(0),$$

where $\hat{F}_{1t}(1) \equiv \left\{\omega : \left\|\hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta\right\| < \varepsilon\right\}$ is the event that 2's realized frequencies are close to 1's estimate. Note that the event $\hat{F}_{1t}(1)$ may not be known by either agent (i.e., we may have $\hat{F}_{1t}(1) \notin \mathscr{H}_{\ell t}$ for $\ell = 1, 2$).

Since M_2^θ is a stochastic matrix, for all $\omega \in \hat{F}_{1t}(1)$, we have $\left\|\hat{\phi}_t^\theta M_1^\theta M_2^\theta - \hat{\psi}^\theta M_2^\theta\right\| < \varepsilon$. We now set ε small enough that $r\delta < \delta - 2\varepsilon$. Since M_{12}^θ is a contraction with fixed point ϕ^θ (see (4) and (5)), we have, again from the triangle inequality,

$$F_{1t}(0) \cap \hat{F}_{1t}(1) \subset F_{2t}(1).$$

Hence, $F_{1t} \cap \hat{F}_{1t}(1) \subset F_{2t}$, and so

$$F_{1t} \cap \hat{F}_{1t}(1) \cap \{\theta\} \subset F_{2t} \cap \{\theta\}.$$

But, from Lemma 5 (recall step 3) we know that $\{\theta\} \subset B_{1t}^q(\hat{F}_{1t}(1) \cap \{\theta\})$ for large t. Consequently,

$$F_{1t} \cap \{\theta\} \subset B_{1t}^q(F_{1t} \cap \hat{F}_{1t}(1) \cap \{\theta\}) \subset B_{1t}^q(F_{1t} \cap F_{2t} \cap \{\theta\}),$$

and we are done.

The proof of Proposition 3 must account for the possibility that some elements of M_{12}^{θ} may not be strictly positive. However, as we show in Lemma 2, since M_{12}^{θ} is irreducible when restricted to a recurrence class, some power of this restricted matrix is a contraction. The proof proceeds as outlined above, with the definition of $F_{\ell t}$ now taking into account the need to take powers of M_{12}^{θ}.

Remark 4 (Arbitrary finite number of agents) The restriction to two agents simplifies the notation, but the result holds for any finite number of agents. We illustrate the argument for three agents (and keep the notation as similar to the two agent case as possible). Denote agent 3's finite set of signals by K. The joint probability of the signal profile $ijk \in I \times J \times K$ under θ is $\pi^{\theta}(ijk)$. In addition to the marginal distributions ϕ^{θ} and ψ^{θ} for 1 and 2, the marginal distribution for 3 is φ^{θ}. As before, M_1^{θ} is the $I^{\theta} \times J^{\theta}$ matrix with ijth element $\sum_k \pi^{\theta}(ijk)/\phi^{\theta}(i)$ (and similarly for M_2). For the pair $1-3$, we denote by N_1^{θ} the $I^{\theta} \times K^{\theta}$ matrix with ikth element $\sum_j \pi^{\theta}(ijk)/\phi^{\theta}(i)$ (and similarly for N_3^{θ}). Finally, for the pair $2-3$, we have analogous definitions for the matrices Q_2^{θ} and Q_3^{θ}. As before, ϕ^{θ} is a stationary distribution of $M_1^{\theta}M_2^{\theta}$, but now also of $N_1^{\theta}N_3^{\theta}$; similar statements hold for ψ^{θ} and the transitions $M_2^{\theta}M_1^{\theta}$ and $Q_2^{\theta}Q_3^{\theta}$, as well as for φ^{θ} and the transitions $N_3^{\theta}N_1^{\theta}$ and $Q_3^{\theta}Q_2^{\theta}$.

Suppose (as in the outline and again for exposition only) that every element of the various Markov transition matrices is non-zero, and let $r < 1$ now be the upper bound on the modulus of contraction of the various contractions. The argument of the outline still applies, once we redefine $F_{1t}(1) \equiv \{\omega : \left\|\hat{\phi}_t^{\theta} M_1^{\theta} - \psi^{\theta}\right\| < \delta - \varepsilon\} \cap \{\omega : \left\|\hat{\phi}_t^{\theta} N_1^{\theta} - \varphi^{\theta}\right\| < \delta - \varepsilon\}$ and $\hat{F}_{1t}(1) \equiv \{\omega : \left\|\hat{\phi}_t^{\theta} M_1^{\theta} - \hat{\psi}^{\theta}\right\| < \varepsilon\} \cap \{\omega : \left\|\hat{\phi}_t^{\theta} N_1^{\theta} - \hat{\varphi}^{\theta}\right\| < \varepsilon\}$ (with similar definitions for the other two agents).

♦

3.3 Preliminary Results: Expectations about Expectations

We summarize here some important properties of the Markov chains induced by the transition matrices M_{12}^{θ} and M_{21}^{θ}.

Remark 5 (Markov Chains) From (4) and (5), the vector ϕ^θ is a stationary distribution for M^θ_{12} and ψ^θ is a stationary distribution for $M^\theta_{21} \equiv M^\theta_2 M^\theta_1$. Moreover, the matrix $M^\theta_{12} D^\theta_1 = D^\theta_1 \Pi^\theta D^\theta_2 [\Pi^\theta]^T D^\theta_1$ is obviously symmetric and has a nonzero diagonal (where D^θ_2 is the diagonal matrix whose jth diagonal element is $(\psi^\theta(j))^{-1}$ for $j \in J^\theta$). This first property implies that the Markov process M^θ_{12} with initial distribution ϕ^θ is reversible.[3] Consequently, the process has ϕ^θ as a stationary distribution when run backward as well as forward, and hence (since $\phi^\theta(i) > 0$ for all $i \in I^\theta$) has no transient states. The second property implies that M^θ_{12} has a nonzero diagonal and hence is aperiodic.

♦

Remark 6 (Recurrent Classes) Two signals i and i' belong to the same *recurrence class* under the transition matrix M^θ_{12} if and only if the probability of a transition from i to i' (in some finite number of steps) is positive.[4] We let $(R^\theta_1(k))^K_{k=1}$ denote the collection of recurrence classes, and we order the elements of I^θ so that the recurrence classes are grouped together and in the order of their indices. This is a partition of I^θ because (from Remark 5) there are no transient states. Similarly, the matrix $M^\theta_{21} \equiv M^\theta_2 M^\theta_1$ is a Markov transition on the set J^θ that we can partition into recurrence classes $(R^\theta_2(k))^K_{k=1}$.

Define a mapping ξ from $(R^\theta_1(k))^K_{k=1}$ to $(R^\theta_2(k))^K_{k=1}$ by letting $\xi(R^\theta_1(k)) = R^\theta_2(k')$ if there exist signals $i \in R^\theta_1(k)$ and $j \in R^\theta_2(k')$ with $\pi^\theta(ij) > 0$. Then ξ is a bijection (as already reflected in our notation). It is convenient therefore to group the elements of J^θ by their recurrence classes in the same order as was done with I^θ. We use the notation $R^\theta(k)$ to refer to the kth recurrence class in either I^θ or J^θ when the context is clear. This choice of notation also reflects the equalities of

---

[3] As $M^\theta_{12} D^\theta_1$ is symmetric, the detailed balance equations at ϕ^θ hold, i.e.,

$$\phi^\theta(i) M^\theta_{12}(ii') = \phi^\theta(i') M^\theta_{12}(i'i)$$

(Brémaud, 1999, page 81).

[4] Since the Markov process has no transient states, if the probability of a (finite-step) transition from i to i' is positive, then the probability of a (finite-step) transition from i' to i is also positive.

the probabilities of $R_1^\theta(k)$ and $R_2^\theta(k)$ under θ, that is

$$\phi^\theta(R_1^\theta(k)) \equiv \sum_{i\in R_1^\theta(k)} \phi^\theta(i) = \sum_{j\in R_2^\theta(k)} \psi^\theta(j) \equiv \psi^\theta(R_2^\theta(k)). \tag{6}$$

Since agent 1 observes a signal in $R_1^\theta(k)$ under parameter θ if and only if agent 2 observes a signal in $R_2^\theta(k)$, conditional on θ the realized frequencies of the recurrence classes also agree.

♦

Let $\gamma^{\theta k}$ denote a probability distribution over I^θ that takes positive values only on the kth recurrence class $R^\theta(k)$, and denote the set of such distributions by $\Delta R^\theta(k)$.

Lemma 2 *There exist $r<1$ and a natural number n such that for all $k\in\{1,\dots,K\}$ and for all $\gamma^{\theta k},\tilde{\gamma}^{\theta k}$ in $\Delta R^\theta(k)$*[5]

$$\left\|\gamma^{\theta k}(M_{12}^\theta)^n - \tilde{\gamma}^{\theta k}(M_{12}^\theta)^n\right\| \leq r\left\|\gamma^{\theta k} - \tilde{\gamma}^{\theta k}\right\| \tag{7}$$

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

Proof. We have noted that M_{12}^θ is aperiodic. By definition, the restriction of M_{12}^θ to any given recurrence class is irreducible and hence ergodic. Thus, because signals are grouped by their recurrence classes, there exists a natural number n such that $(M_{12}^\theta)^n$ has the block-diagonal form with each block containing only strictly positive entries. The blocks consist of the non-zero n-step transition probabilities between signals within a recurrence class. The product of $\gamma^{\theta k}$ with $(M_{12}^\theta)^n$ is just the product of $\gamma^{\theta k}$ restricted to $R^\theta(k)$ with the kth block of $(M_{12}^\theta)^n$. Because it has all non-zero entries, the kth block is a contraction mapping (Stokey and Lucas, 1989, Lemma 11.3). In particular, there exists an $r<1$ such that (7) holds. ■

[5]For any $x\in\mathbb{R}^N$, $\|x\|\equiv\sum_{k=1}^N |x_k|$ is the variation norm of x.

3.4 Preliminary Results: Frequencies are Enough

Let $\hat{\phi}^{\theta k}$ denote the distribution over I^{θ} obtained by conditioning $\hat{\phi}$ on the kth recurrence class $R^{\theta}(k)$ (for those cases in which $\hat{\phi}^{\theta}(R^{\theta}(k)) > 0$), and let $\phi^{\theta k}$, $\psi^{\theta k}$, and $\hat{\psi}_t^{\theta k}$ be analogous.

Our first result shows that if agent 1's signal frequencies are sufficiently close to those expected under θ, the posterior probability he attaches to parameter θ approaches one.

Lemma 3 *There exist $\delta \in (0,1)$, $\beta \in (0,1)$, and a sequence $\xi : \mathbb{N} \to [0,1]$ with $\xi(t) \to 1$ such that*

$$P(\theta \mid h_{1t}) \geq \xi(t)$$

for all $\theta \in \Theta$ and h_{1t} satisfying $P(\theta \mid h_{1t}) > 0$, $\left\|\hat{\phi}_t^{\theta k} - \phi^{\theta k}\right\| < \delta$ for all k, and $\beta < \frac{\hat{\phi}_t^{\theta}(R^{\theta}(k))}{\phi^{\theta}(R^{\theta}(k))} < \beta^{-1}$ for all k. An analogous result holds for agent 2.

Proof. Fix a parameter θ and $\tilde{\delta} < \min_{i,\theta}\{\phi^{\theta}(i) : \phi^{\theta}(i) > 0\}$. Then $\left\|\hat{\phi}_t^{\theta k} - \phi^{\theta k}\right\| < \tilde{\delta}$ for all k only if $\hat{\phi}_t$ puts strictly positive probability on every signal $i \in I^{\theta}$. For θ' and h_{1t} with $P(\theta' \mid h_{1t}) > 0$, define the ratio

$$\lambda_{1t}^{\theta\theta'} \equiv \log \frac{P(\theta \mid h_{1t})}{P(\theta' \mid h_{1t})} = \log \frac{\phi^{\theta}(i_{t-1}) P(\theta \mid h_{1t-1})}{\phi^{\theta'}(i_{t-1}) P(\theta' \mid h_{1t-1})}.$$

We now show that β and $\delta \leq \tilde{\delta}$ can be chosen so that there exists $\eta > 0$ with the property that

$$\lambda_{1t}^{\theta\theta'} \geq \lambda_{10}^{\theta\theta'} + t\eta \qquad \forall \theta' \neq \theta$$

for all $\theta' \in \Theta$ and histories h_{1t} for which $\left\|\hat{\phi}_t^{\theta k} - \phi^{\theta k}\right\| < \tilde{\delta}$ for all k and for which $\lambda_{1t}^{\theta\theta'}$ is defined. Notice that $\lambda_{10}^{\theta\theta'} = \frac{p(\theta)}{p(\theta')}$ is the log-likelihood ratio at time zero, that is, the ratio of prior probabilities.

Our choice of $\tilde{\delta}$, implying that every signal $i \in I^{\theta}$ has appeared in the history h_{1t}, ensures that $P(\theta' | h_{1t}) > 0$ (and hence $\lambda_{1t}^{\theta\theta'}$ is well defined) only if $I^{\theta} \subset I^{\theta'}$. This in turn ensures that the following expressions are well defined (in particular, having nonzero denominators). Because signals are distributed independently and identically across periods, $\lambda_{1t}^{\theta\theta'}$ can be written as

$$\lambda_{1t}^{\theta\theta'} = \lambda_{10}^{\theta\theta'} + \sum_{s=0}^{t-1} \log\left(\frac{\phi^{\theta}(i_s)}{\phi^{\theta'}(i_s)}\right).$$

We find a lower bound for the last term. Let

$$H^{\theta\theta'} \equiv E^{\theta}\left(\log\frac{\phi^{\theta}(i)}{\phi^{\theta'}(i)}\right) > 0$$

denote the relative entropy of ϕ^{θ} with respect to $\phi^{\theta'}$. Then,

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

for $b = \max_{i,\theta,\theta'\in\Theta}\left\{\frac{\phi^{\theta}(i)}{\phi^{\theta'}(i)} : \phi^{\theta}(i) > 0\right\}$. By Assumption 2, $b > 1$. Thus,

$$\lambda_{1t}^{\theta\theta'} \ge \lambda_{10}^{\theta\theta'} + t\left(H^{\theta\theta'} - \log b\left\|\hat{\phi}_t^{\theta} - \phi^{\theta}\right\|\right).$$

We now argue that $\delta \le \tilde{\delta}$ and β can be chosen to ensure $H^{\theta\theta'} - \log b\left\|\hat{\phi}_t^{\theta} - \phi^{\theta}\right\| > \eta$ for all θ, θ' and some $\eta > 0$. For this, it is enough to observe that the mapping

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

is continuous and equals zero if and only if $\hat{\phi}_t^\theta(R^\theta(k)) = \phi^\theta(R^\theta(k))$ and $\hat{\phi}_t^{\theta k} = \phi^{\theta k}$ for all k.

We thus have δ and β such that for θ and h_{1t} satisfying the hypotheses of the lemma and θ' with $P(\theta' \mid h_{1t}) > 0$, it must be that $\lambda_{1t}^{\theta\theta'} \geq \lambda_{10}^{\theta\theta'} + t\eta$ and hence

$$\frac{p(\theta')}{p(\theta)} \geq \frac{P(\theta'|h_{1t})}{P(\theta|h_{1t})} e^{t\eta}.$$

Noting that this inequality obviously holds for θ' with $P(\theta' \mid h_{1t}) = 0$, we can sum over $\theta' \neq \theta$ and rearrange to obtain

$$\frac{P(\theta \mid h_{1t})}{1 - P(\theta \mid h_{1t})} \geq \frac{p(\theta)}{1 - p(\theta)} e^{t\eta},$$

giving the required result. ■

We next note that with high probability, observed frequencies match their expected values. Together with Lemma 3, this implies that each agent learns Θ.

Lemma 4 *For all $\varepsilon > 0$ and θ, $P^\theta(\|\hat{\phi}_t^\theta - \phi^\theta\| < \varepsilon) \to 1$ and $P^\theta(\|\hat{\psi}_t^\theta - \psi^\theta\| < \varepsilon) \to 1$ as $t \to \infty$.*

Proof. This follows from the Weak Law of Large Numbers (Billingsley, 1979, p. 86). ■

We now show that each agent believes that, conditional on any parameter θ, his or her expectation of the frequencies of the signals observed by his or her opponent is likely to be nearly correct. Recall that $\hat{\phi}_t^\theta M_1^\theta$ is agent 1's expectation of 2's frequencies $\hat{\psi}_t^\theta$ and that $\hat{\psi}_t^\theta M_2^\theta$ is agent 2's expectation of 1's frequencies $\hat{\phi}_t^\theta$.

Lemma 5 *For any $\varepsilon_1 > 0$, $\varepsilon_2 > 0$, there exists T such that for all $t > T$ and for every h_t with $P^\theta(h_t) > 0$,*

$$P^\theta\left(\left\|\hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta\right\| < \varepsilon_1 \mid h_{1t}\right) > 1 - \varepsilon_2 \tag{8}$$

and

$$P^{\theta}\left(\left\|\hat{\psi}_t^{\theta} M_2^{\theta} - \hat{\phi}_t^{\theta}\right\| < \varepsilon_1 \mid h_{2t}\right) > 1 - \varepsilon_2. \tag{9}$$

Proof. We focus on (8); the argument for (9) is identical. Defining $\bar{\psi}_t^{\theta} \equiv \hat{\phi}_t^{\theta} M_1^{\theta}$, the left side of (8) is bounded below:

$$P^{\theta}\left(\left\|\hat{\phi}_t^{\theta} M_1^{\theta} - \hat{\psi}_t^{\theta}\right\| < \varepsilon_1 \,\middle|\, h_{1t}\right) \geq 1 - \sum_{j \in J^{\theta}} P^{\theta}\left(\left|\bar{\psi}_t^{\theta}(j) - \hat{\psi}_t^{\theta}(j)\right| \geq \frac{\varepsilon_1}{J^{\theta}} \,\middle|\, h_{1t}\right). \tag{10}$$

Conditional on θ and h_{1t}, agent 2's signals are independently, but *not* identically, distributed across time. In period s, given signal i_s, agent 2's signals are distributed according to the conditional distribution $(\pi^{\theta}(i_s j)/\phi^{\theta}(i_s))_j$. However, we can bound the expression on the right side of (10) using a related process obtained by averaging the conditional distributions. The average probability that agent 2 observes signal j over the t periods $\{0, 1, \ldots, t-1\}$, conditional on h_{1t} is

$$\frac{1}{t} \sum_{s=0}^{t-1} \frac{\pi^{\theta}(i_s j)}{\phi^{\theta}(i_s)} = \sum_i \hat{\phi}_t(i) \frac{\pi^{\theta}(ij)}{\phi^{\theta}(i)} = \bar{\psi}_t^{\theta}(j),$$

agent 1's expectation of the frequency that 2 observed j.

Consider now t independent and identically distributed draws of a random variable distributed on J^{θ} according to the "average" distribution $\bar{\psi}_t^{\theta} \in \Delta(J^{\theta})$; we refer to this process as the average process. Denote the frequencies of signals generated by the average process by $\eta_t \in \Delta(J^{\theta})$. The process generating the frequencies $\hat{\psi}_t$ attaches the same average probability to each signal j over periods $0, \ldots, t-1$ as does the average process, but does not have identical distributions (as we noted earlier).

We use the average process to bound the terms in the sum in (10). By Hoeffding (1956, Theorem 4, p. 718), the original process is more concentrated about its mean than is the average process, that is,[6]

$$\tilde{P}\left(\left|\bar{\psi}_t^\theta(j)-\eta_t(j)\right| \geq \frac{\varepsilon_1}{J^\theta}\right) \geq P^\theta\left(\left|\bar{\psi}_t^\theta(j)-\hat{\psi}_t^\theta(j)\right| \geq \frac{\varepsilon_1}{J^\theta} \,\middle|\, h_{1t}\right), \qquad j \in J^\theta,$$

where $\tilde{P}$ is the measure associated with the average process. Applying this upper bound to (10), we have

$$P^\theta\left(\left\|\hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta\right\| < \varepsilon_1 \,\middle|\, h_{1t}\right) \geq 1 - \sum_{j\in J^\theta} \tilde{P}\left(\left|\bar{\psi}_t^\theta(j)-\eta_t(j)\right| \geq \frac{\varepsilon_1}{J^\theta}\right). \tag{11}$$

The event $\{|\bar{\psi}_t^\theta(j)-\eta_t(j)| > \varepsilon_1/J^\theta\}$ is the event that the realized frequency of a Bernoulli process is far from its mean. By a large deviation inequality ((42) in Shiryaev (1996, p. 69)),

$$\tilde{P}\left(|\bar{\psi}_t^\theta(j)-\eta_t(j)| > \frac{\varepsilon_1}{J^\theta}\right) \leq 2e^{-2t\varepsilon_1^2/(J^\theta)^2}.$$

Using this bound in (11), we have

$$P^\theta\left(\left\|\hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta\right\| < \varepsilon_1 \,\middle|\, h_{1t}\right) \geq 1 - 2J^\theta e^{-2t\varepsilon_1^2/(J^\theta)^2}.$$

This inequality holds for any history h_{1t}. We can thus choose t large enough so that the right-hand side is less than ε_2 and the statement of the lemma follows. ■

3.5 Proof of Proposition 3

We fix an arbitrary parameter θ and define a sequence of events F_t (suppressing notation for the dependence of F_t on θ), and show that F_t has the three requisite properties from Corollary 1 for sufficiently large t.

[6]For example, 100 flips of a $(p, 1-p)$ coin generates a more dispersed distribution than $100p$ flips of a $(1,0)$ coin and $100(1-p)$ flips of a $(0,1)$ coin.

The event F_t. Let $\delta \in (0,1)$ and $\beta \in (0,1)$ be the constants identified in Lemma 3. Pick $\varepsilon > 0$ such that $r\delta < \delta - 2n\varepsilon$ where $r < 1$ and n are identified in Lemma 2. For each date t, we define the event F_t as follows.

First, we ask that agent 1's realized frequency of signals from I^θ and 2's from J^θ be close to the frequencies expected under θ. For each k, define the events

$$F_{1t}^k(0) \equiv \left\{\omega : \left\|\hat{\phi}_t^{\theta k} - \phi^{\theta k}\right\| < \delta\right\} \tag{12}$$

$$\text{and} \qquad F_{2t}^k(0) \equiv \left\{\omega : \left\|\hat{\psi}_t^{\theta k} - \psi^{\theta k}\right\| < \delta\right\}. \tag{13}$$

Lemma 2 ensures that $\|\hat{\phi}_t^{\theta k}(M_{12}^\theta)^n - \phi^{\theta k}(M_{12}^\theta)^n\|$ will then be smaller than δ on $F_{1t}^k(0)$. We define our event so that the same is true for all powers of M_{12}^θ between 0 and n. Hence, for any $l \in \{1,\ldots,n\}$ and for each k, let

$$F_{1t}^k(2l-1) \equiv \left\{\omega : \left\|\hat{\phi}_t^{\theta k}(M_{12}^\theta)^{l-1}M_1^\theta - \psi^{\theta k}\right\| < \delta - (2l-1)\varepsilon\right\} \tag{14}$$

$$\text{and} \qquad F_{1t}^k(2l) \equiv \left\{\omega : \left\|\hat{\phi}_t^{\theta k}(M_{12}^\theta)^{l} - \phi^{\theta k}\right\| < \delta - 2l\varepsilon\right\}. \tag{15}$$

Similarly, for agent 2,

$$F_{2t}^k(2l-1) \equiv \left\{\omega : \left\|\hat{\psi}_t^{\theta k}(M_{21}^\theta)^{l-1}M_2^\theta - \phi^{\theta k}\right\| < \delta - (2l-1)\varepsilon\right\} \tag{16}$$

$$\text{and} \qquad F_{2t}^k(2l) \equiv \left\{\omega : \left\|\hat{\psi}_t^{\theta k}(M_{21}^\theta)^{l} - \psi^{\theta k}\right\| < \delta - 2l\varepsilon\right\}. \tag{17}$$

Next, define the events

$$F_{1t} \equiv \bigcap_{k=1}^{K}\bigcap_{\kappa=0}^{2n-1} F_{1t}^k(\kappa) \equiv \bigcap_{k=1}^{K} F_{1t}^k \equiv \bigcap_{\kappa=0}^{2n-1} F_{1t}(\kappa),$$

$$F_{2t} \equiv \bigcap_{k=1}^{K}\bigcap_{\kappa=0}^{2n-1} F_{2t}^k(\kappa) \equiv \bigcap_{k=1}^{K} F_{2t}^k \equiv \bigcap_{\kappa=0}^{2n-1} F_{2t}(\kappa),$$

$$[\theta] \equiv \left\{\omega \in \{\theta\} \times (I \times J)^\infty : P^\theta(h_{\ell t}) > 0,\ \ell \in \{1,2\}, t = 0,1,\ldots\right\},$$

and

$$G_t^\theta \equiv [\theta] \cap \left\{ \beta < \frac{\hat{\phi}_t(R^\theta(k))}{\phi^\theta(R^\theta(k))} < \beta^{-1},\ \forall k \right\} \equiv [\theta] \cap G_{1t} \tag{18}$$

$$= [\theta] \cap \left\{ \beta < \frac{\hat{\psi}_t(R^\theta(k))}{\psi^\theta(R^\theta(k))} < \beta^{-1},\ \forall k \right\} \equiv [\theta] \cap G_{2t}. \tag{19}$$

The equality of the two descriptions of G_t^θ follows from Remark 6. Finally, we define the event F_t,

$$F_t \equiv F_{1t} \cap F_{2t} \cap G_t^\theta.$$

In the analysis that follows, we simplify notation by using $\{\|\cdot\| < \varepsilon\}$ to denote the event $\{\omega : \|\cdot\| < \varepsilon\}$.

θ is q-believed on F_t. By definition $F_t \subset F_{1t}(0) \cap F_{2t}(0) \cap G_t^\theta$. Lemma 3 then implies that for any $q < 1$, we have $F_t \subset B_t^q(\theta)$ for all t sufficiently large.

F_t is likely under θ. If $\hat{\phi}^t = \phi^\theta$ and $\hat{\psi}^t = \psi^\theta$, then the inequalities (12)–(19) appearing in the definitions of the sets F_{1t}, F_{2t}, and G_t^θ are strictly satisfied (because $\phi^{\theta k} M_1^\theta = \psi^{\theta k}$ and $\psi^{\theta k} M_2^\theta = \phi^{\theta k}$ for each k). The (finite collection of) inequalities (12)–(19) are continuous in $\hat{\phi}^t$ and $\hat{\psi}^t$ and independent of t. Hence, (12)–(19) are satisfied for any $\hat{\phi}^t$ and $\hat{\psi}^t$ sufficiently close to ϕ^θ and ϕ^θ. We can therefore choose $\varepsilon^\dagger > 0$ sufficiently small such that

$$\{\|\hat{\phi}^t - \phi^\theta\| < \varepsilon^\dagger,\ \|\hat{\psi}^t - \psi^\theta\| < \varepsilon^\dagger\} \cap [\theta] \ \subset\ F_t, \qquad \forall t.$$

By Lemma 4, the P^θ-probability of the set on the left side approaches one as t gets large, ensuring that for all $q \in (0,1)$, $P^\theta(F_t) > q$ for all large enough t.

F_t is q-evident. We show that for any q, F_t is q-evident when t is sufficiently large. Recalling that ε and β were fixed in defining F_t, choose $\varepsilon_1 \equiv \varepsilon\beta \min_{j \in J^\theta} \psi^\theta(j)$. Note that $\varepsilon_1 / \hat{\psi}^\theta(R^\theta(k)) < \varepsilon$ on the events $F_{1t} \cap G_t^\theta$ and $F_{2t} \cap G_t^\theta$.

[STEP 1] The first step is to show that if the realized frequencies of agent 1's signals are close to their population frequencies under θ and his expectations of agent 2's frequencies are not too far away from agent 2's realized frequencies, then (conditional on θ) the realized frequencies of agent 2's signals are also close to their population frequencies under θ. In particular, we show

$$F_{1t} \cap G_t^\theta \cap \left\{ \left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \right\} \subset F_{2t}. \tag{20}$$

First, fix k and note that for each $l = 1, \ldots, n$,

$$\begin{aligned}
&F_{1t}^k(2l) \cap \{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \} \\
&\quad \subset F_{1t}^k(2l) \cap \{ \left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^l - \hat{\psi}_t^{\theta k} (M_{21}^\theta)^{l-1} M_2^\theta \right\| < \varepsilon \} \\
&\quad = \{ \left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^l - \phi^{\theta k} \right\| < \delta - 2l\varepsilon \} \cap \{ \left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^l - \hat{\psi}_t^{\theta k} (M_{21}^\theta)^{l-1} M_2^\theta \right\| < \varepsilon \} \\
&\quad \subset \{ \left\| \hat{\psi}_t^{\theta k} (M_{21}^\theta)^{l-1} M_2^\theta - \phi^{\theta k} \right\| < \delta - (2l-1)\varepsilon \} \\
&\quad = F_{2t}^k(2l-1).
\end{aligned} \tag{21}$$

The first inclusion uses the fact that $(M_{21}^\theta)^{l-1} M_2^\theta$ is a stochastic matrix. The equalities use the definitions of $F_{1t}^k(2l)$ and $F_{2t}^k(2l-1)$. The last inclusion is a consequence of the triangle inequality. Similarly, for $l = 1, \ldots, n$, we have

$$F_{1t}^k(2l-1) \cap \{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \} \subset F_{2t}^k(2(l-1)).$$

This suffices to conclude that

$$\bigcap_{\kappa=1}^{2n-1} F_{1t}^k(\kappa) \cap \{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \} \subset \bigcap_{\kappa=0}^{2n-2} F_{2t}^k(\kappa). \tag{22}$$

We next note that

$$\begin{aligned} F_{1t}^{k}(0)\cap\{\left\|\hat{\phi}_t^{\theta k}M_1^{\theta}-\hat{\psi}_t^{\theta k}\right\|<\varepsilon\} &\subset F_{1t}^{k}(2n)\cap\{\left\|\hat{\phi}_t^{\theta k}M_1^{\theta}-\hat{\psi}_t^{\theta k}\right\|<\varepsilon\} \\ &\subset F_{2t}^{k}(2n-1), \end{aligned} \tag{23}$$

where $F_{1t}^{k}(0)\subset F_{1t}^{k}(2n)$ is an implication of $\phi^{\theta k}(M_{12}^{\theta})^{n}=\phi^{\theta k}$, Lemma 2, and our choice of ε and n; while the second inclusion follows from (21) (for $l=n$). Combining (22)–(23) for $k=1,\ldots,K$, we have

$$F_{1t}\cap\bigcap_{k}\left\{\left\|\hat{\phi}_t^{\theta k}M_1^{\theta}-\hat{\psi}_t^{\theta k}\right\|<\varepsilon\right\}\subset F_{2t}. \tag{24}$$

As the matrix M_1^{θ} maps recurrence classes to recurrence classes, on G_t^{θ} we have that

$$\begin{aligned} \|\hat{\phi}_t^{\theta}M_1^{\theta}-\hat{\psi}_t^{\theta}\| &= \sum_k \|\hat{\phi}_t^{\theta}(R^{\theta}(k))\hat{\phi}_t^{\theta k}M_1^{\theta}-\hat{\psi}_t^{\theta}(R^{\theta}(k))\hat{\psi}_t^{\theta k}\| \\ &> \|\hat{\phi}_t^{\theta}(R^{\theta}(k))\hat{\phi}_t^{\theta k}M_1^{\theta}-\hat{\psi}_t^{\theta}(R^{\theta}(k))\hat{\psi}_t^{\theta k}\| \\ &= \hat{\psi}_t^{\theta}(R^{\theta}(k))\|\hat{\phi}_t^{\theta k}M_1^{\theta}-\hat{\psi}_t^{\theta k}\|, \end{aligned}$$

since $\hat{\phi}_t^{\theta}(R^{\theta}(k))=\hat{\psi}_t^{\theta}(R^{\theta}(k))$ on $[\theta]$ (recall Remark 6). Our choice of ε_1 then yields that, on $F_{1t}\cap G_t^{\theta}$,

$$\left\|\hat{\phi}_t^{\theta}M_1^{\theta}-\hat{\psi}_t^{\theta}\right\|<\varepsilon_1 \quad\Rightarrow\quad \varepsilon>\frac{\varepsilon_1}{\hat{\psi}_t^{\theta}(R^{\theta}(k))}>\|\hat{\phi}_t^{\theta k}M_1^{\theta}-\hat{\psi}_t^{\theta k}\|, \qquad \forall k.$$

Therefore

$$F_{1t}\cap G_t^{\theta}\cap\left\{\left\|\hat{\phi}_t^{\theta}M_1^{\theta}-\hat{\psi}_t^{\theta}\right\|<\varepsilon_1\right\}\subset F_{1t}\cap\bigcap_k\left\{\left\|\hat{\phi}_t^{\theta k}M_1^{\theta}-\hat{\psi}_t^{\theta k}\right\|<\varepsilon\right\},$$

and by (24) we have proved (20).

[STEP 2] We now conclude the proof of q-evidence. Pick $p\in\left(\sqrt{q},1\right)$ and set $\varepsilon_2=1-p$ in Lemma 5.

Consider the event $F_{1t} \cap G_t^{\theta}$. For t sufficiently large, given any history consistent with a state in $F_{1t} \cap G_t^{\theta}$, agent 1 attaches at least probability p to θ ($F_{1t} \cap G_t^{\theta} \subset B_{1t}^{p}(\theta)$) (Lemma 3). Conditional on θ we have, by Lemma 5, that for large t, agent 1 attaches probability at least p to $\left\|\hat{\phi}_t^{\theta} M_1^{\theta} - \hat{\psi}_t^{\theta}\right\| < \varepsilon_1$. Hence

$$F_{1t} \cap G_t^{\theta} \ \subset \ B_{1t}^{p^2}\left(\left\{\left\|\hat{\phi}_t^{\theta} M_1^{\theta} - \hat{\psi}_t^{\theta}\right\| < \varepsilon_1\right\} \cap [\theta]\right).$$

Since $F_{1t} \cap G_{1t}$ is measurable with respect to $\mathscr{H}_{1t}$ and $G_t^{\theta} = [\theta] \cap G_{1t}$, we have

$$F_{1t} \cap G_t^{\theta} \ \subset \ B_{1t}^{p^2}\left(F_{1t} \cap G_t^{\theta} \cap \left\{\left\|\hat{\phi}_t^{\theta} M_1^{\theta} - \hat{\psi}_t^{\theta}\right\| < \varepsilon_1\right\}\right),$$

and hence, from (20),

$$F_{1t} \cap G_t^{\theta} \subset B_{1t}^{p^2}\left(F_{1t} \cap F_{2t} \cap G_t^{\theta}\right) \ = \ B_{1t}^{p^2}(F_t). \tag{25}$$

A similar argument for agent 2 gives $F_{2t} \cap G_t^{\theta} \subset B_{2t}^{p^2}(F_t)$ and thus $F_t \subset B_t^{p^2}(F_t) \subset B_t^{q}(F_t)$ for sufficiently large t.

4 A Counterexample to Common Learning

This section presents an example in which Assumption 1 fails and common learning does not occur, although the agents do privately learn. There are two values of the parameter, θ' and θ'', satisfying $0 < \theta' < \theta'' < 1$. Signals are nonnegative integers. The distribution of signals is displayed in Figure 2.[7] If we set $\theta' = 0$ and $\theta'' = 1$, then we can view one period of this process as an instance of the signals in Rubinstein's (1989) electronic mail game, where the signal corresponds to the number of "messages" received.[8] It is immediate that the agents faced with

[7]It would cost only additional notation to replace the single value ε in Figure 2 with heterogeneous values, as long as the resulting analogue of (26) is a collection whose values are bounded away form 0 and 1.

[8]Rubinstein (1989) is concerned with whether a single signal drawn from this distribution allows agents to condition their action on the state, while we are concerned with whether an arbi-

Probability	Player-1 signal	Player-2 signal
θ	0	0
$\varepsilon(1-\theta)$	1	0
$(1-\varepsilon)\varepsilon(1-\theta)$	1	1
$(1-\varepsilon)^2\varepsilon(1-\theta)$	2	1
$(1-\varepsilon)^3\varepsilon(1-\theta)$	2	2
$(1-\varepsilon)^4\varepsilon(1-\theta)$	3	2
$(1-\varepsilon)^5\varepsilon(1-\theta)$	3	3
$\vdots$	$\vdots$	$\vdots$

Figure 2: The distribution of signals for the counterexample given parameter $\theta \in \{\theta', \theta''\}$, where $\varepsilon \in (0,1)$.

a sequence of independent draws from this distribution learn Θ. We now show that common learning does not occur.

What goes wrong when trying to establish common learning in this context, and how does this depend upon the infinite set of signals? Establishing common q-belief in parameter θ requires showing that if agent 1 has observed signals just on the boundary of inducing probability q that the parameter is θ, then agent 1 nonetheless believes 2 has seen signals inducing a similar belief (and believes that 2 believes 1 has seen such signals, and so on). In the case of finite signals, a key step in this argument is the demonstration that (an appropriate power of) the Markov transition matrix M_{12}^{θ} is a contraction. In the current case, the corresponding Markov process is not a contraction (though the marginal distribution is still stationary). As a result, agent ℓ can observe signals on the boundary of inducing probability q of state θ while believing that agent $\hat{\ell}$ has observed signals on the "wrong side" of this boundary.

The first step in our argument is to show that, regardless of what agents have

trarily large number of signals suffices to commonly learn the parameter. Interestingly, repeated observation of the original Rubinstein process (i.e., $\theta' = 0$ and $\theta'' = 1$) leads to common learning. In particular, consider the event F_t at date t that the state is θ' and no messages have ever been received. This event is $q(t)$-evident where $q(t)$ approaches 1 as t approaches infinity, since 1 assigns probability 1 and 2 assigns a probability approaching 1 to F_t whenever it is true.

observed, nth-order beliefs attach positive probability to agent 2 having observed larger and larger (and rarer and rarer) signals, as n gets larger (cf. (27) and (29) below). We then argue that agents attaching strictly positive nth-order belief to agent 2 having observed such extraordinarily rare signals will also attach strictly positive nth order-belief to another rare event—that agent 2 has *never* seen a zero signal (cf. (31)). Since zero signals are more likely under parameter θ'', this ensures a positive nth-order belief in agent 2's being being confident the parameter is θ', even when it is not, precluding common learning.

Let

$$q \equiv \min\left\{ \frac{\varepsilon(1-\theta'')}{\theta''+\varepsilon(1-\theta'')}, \frac{(1-\varepsilon)}{(2-\varepsilon)} \right\}. \tag{26}$$

Note that regardless of the signal observed by agent 1, he always believes with probability at least q that 2 has seen the same signal, and regardless of the signal observed by 2, she always believes with probability at least q that 1 has seen a higher signal.

We show that for all t sufficiently large there is (independently of the observed history) a finite iterated q-belief that θ' is the true parameter. This implies that θ'' can never be iterated p-believed for any $p > 1-q$, with Lemma 1 then implying that θ'' can never be common p-belief. That is, we will show that for t large enough, $B^q_{2t}(\theta') = \Omega$ and so $B^p_{2t}(\theta'') = \varnothing$ for all $p > 1-q$.

Define for each k, the event that agent ℓ observes a signal of at least k before time t:

$$D_{\ell t}(k) \equiv \{\omega : z_{\ell s} \geq k \text{ for some } s \leq t\}.$$

Note that $D_{\ell t}(0)$ is equal to Ω (the event that any t-length history occurs). For every $k \geq 0$ the definition of q implies:

$$D_{1t}(k) \subset B^q_{1t}(D_{2t}(k)),$$

and

$$D_{2t}(k-1) \subset B^q_{2t}(D_{1t}(k)),$$

which together imply

$$D_{2t}(k-1) \subset B^q_{2t} B^q_{1t}(D_{2t}(k)).$$

By induction, for all $0 \le m \le k$,

$$D_{2t}(m) \subset \left(B^q_{2t} B^q_{1t}\right)^{k-m}(D_{2t}(k)). \tag{27}$$

Now, for any K and any list $\left(k^1, k^2, \ldots, k^K\right)$, where $k^s \ge k^{s-1}$, define the event that agent ℓ observes distinct signals of at least k^s before time t,

$$D_{\ell t}\left(k^1, k^2, \ldots, k^K\right) \equiv \{\omega : \exists \text{ distinct } \tau_s \le t,\ \ s = 1, \ldots, K,\ \text{ s.t. } z_{\ell \tau_s} \ge k^s\}.$$

Note that for $K \le t$, $D_{\ell t}(0, k^2, \ldots, k^K) = D_{\ell t}(k^2, \ldots, k^K)$. Whenever agent 1 observes a signal k he knows that agent 2 has seen a signal at least $k-1$. Hence,

$$D_{1t}\left(k^1, k^2, \ldots, k^K\right) \subset B^q_{1t}(D_{2t}(k^1, k^2-1, k^3-1, \ldots, k^K-1))$$

and by similar reasoning

$$D_{2t}\left(k^1, k^2, \ldots, k^K\right) \subset B^q_{2t}(D_{1t}(k^1+1, k^2, k^3, \ldots, k^K)),$$

so that for all n, if $0 \le k^1 \le k^2 - 2n$, then

$$D_{2t}\left(k^1, k^2, \ldots, k^K\right) \subset \left(B^q_{2t} B^q_{1t}\right)^n D_{2t}(k^1+n, k^2-n, k^3-n, \ldots, k^K-n). \tag{28}$$

From (27),

$$\Omega = D_{2t}(0) \subset \left(B^q_{2t} B^q_{1t}\right)^{2^{t-1}} D_{2t}(2^{t-1}) \tag{29}$$

and, for $t \ge 2$, from (28),

$$D_{2t}(2^{t-1}) = D_{2t}(0, 2^{t-1}) \subset \left(B^q_{2t} B^q_{1t}\right)^{2^{t-2}} D_{2t}(2^{t-2}, 2^{t-2}). \tag{30}$$

Inserting (30) in (29) gives $\Omega \subset \left(B^q_{2t} B^q_{1t}\right)^{2^{t-1}+2^{t-2}} D_{2t}(2^{t-2}, 2^{t-2})$. Continuing in this fashion and noting that $2^{t-1}+2^{t-2}+\ldots+2^{t-t}=2^t-1$, we obtain

$$\Omega \subset \left(B^q_{2t} B^q_{1t}\right)^{2^t-1} D_{2t}(\underbrace{2^{t-t}, 2^{t-t}, \ldots, 2^{t-t}}_{t \text{ times}}) = \left(B^q_{2t} B^q_{1t}\right)^{2^t-1} D_{2t}(\underbrace{1, 1, \ldots, 1}_{t \text{ times}}). \quad (31)$$

Now choose t large enough so that after a t-length history in which signal 0 was never observed, agent 2 assigns probability at least q to θ', i.e.,[9]

$$D_{2t}(\underbrace{1, 1, \ldots, 1}_{t \text{ times}}) \subset B^q_{2t}(\theta').$$

Using (31), we then have $\Omega \subset \left(B^q_{2t} B^q_{1t}\right)^{2^t-1} B^q_{2t}(\theta')$ and hence have shown that for t large enough, regardless of the history, there cannot be iterated p-belief in θ'' for any $p > 1-q$, i.e. $I^p(\theta'') = \varnothing$. Now by Lemma 1, $C^p(\theta'') = \varnothing$.

References

ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2006): “Learning and Disagreement in an Uncertain World,” MIT, unpublished.

BILLINGSLEY, P. (1979): *Probability and Measure*. John Wiley and Sons, New York, 1st edn.

BRÉMAUD, P. (1999): *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer-Verlag, New York.

CRIPPS, M. W., G. J. MAILATH, AND L. SAMUELSON (forthcoming): “Disappearing Private Reputations in Long-Run Relationships,” *Journal of Economic Theory*.

[9] This is possible because after such a history $\frac{P(\theta'|h_{2t})}{1-P(\theta'|h_{2t})} = \frac{p(\theta')}{p(\theta'')}\left(\frac{1-\theta'}{1-\theta''}\right)^t$.

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.

MORRIS, S. (1999): “Approximate Common Knowledge Revisited,” *International Journal of Game Theory*, 28(3), 385–408.

RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge,” *American Economic Review*, 79(3), 385–391.

SAMET, D. (1998): “Iterated Expectations and Common Priors,” *Games and Economic Behavior*, 24(1/2), 131–141.

SHIRYAEV, A. N. (1996): *Probability*. Springer-Verlag, New York, second edn.

STOKEY, N., AND R. E. LUCAS, JR. (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA.

WISEMAN, T. (2005): “A Partial Folk Theorem for Games with Unknown Payoff Distributions,” *Econometrica*, 73(2), 629–645.