Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 11-012

"Common Learning with Intertemporal Dependence"

by

Martin W. Cripps, Jeffrey C. Ely,
George J. Mailath, and Larry Samuelson

Common Learning with Intertemporal Dependence*

Martin W. Cripps
Department of Economics
University College London
London WC1E 6BT UK
m.cripps@ucl.ac.uk

Jeffrey C. Ely
Department of Economics
Northwestern University
Evanston, IL 60208 USA
jeffely@northwestern.edu

George J. Mailath
Department of Economics
University of Pennsylvania
Philadelphia, PA 19105 USA
gmailath@econ.upenn.edu

Larry Samuelson
Department of Economics
Yale University
New Haven, CT 06520 USA
Larry.Samuelson@yale.edu

May 19, 2011

Abstract Consider two agents who learn the value of an unknown parameter by observing a sequence of private signals. Will the agents *commonly* learn the value of the parameter, i.e., will the true value of the parameter become approximate common-knowledge? If the signals are independent and identically distributed across time (but not necessarily across agents), the answer is yes (Cripps, Ely, Mailath, and Samuelson, 2008). This paper explores the implications of allowing the signals to be dependent over time. We present a counterexample showing that even extremely simple time dependence can preclude common learning, and present sufficient conditions for common learning.

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

JEL Classification Numbers: D82, D83.

*We thank Stephen Morris for helpful comments. Cripps thanks the Cowles Foundation and the Economic and Social Research Council (UK) via ELSE and Mailath and Samuelson thank the National Science Foundation (grants SES-0648780 and SES-0850263, respectively) for financial support.

Contents

Common Learning with Intertemporal Dependence

1 Introduction

Coordinating behavior requires people to have beliefs that are not too different. Differences in beliefs may arise when agents must learn about their environment in order to identify the appropriate action on which to coordinate. Suppose two agents would like to jointly choose an action that depends on the value of an unknown underlying parameter, and that each agent observes a sequence of private signals sufficiently informative as to ensure she will (almost surely) learn the parameter value. Successful coordination requires that the agents (at least approximately) *commonly* learn the parameter value—agent 1 must attach sufficiently high probability not only to a particular value of the parameter, but also to the event that agent 2 attaches high probability to this value, and to the event that agent 2 attaches high probability to the event that agent 1 attaches high probability to this value, and so on.

Cripps, Ely, Mailath, and Samuelson (2008) show that common learning obtains when the private signals are drawn from finite sets and the signal distributions, conditional on the parameter, are independent over time. A counter-example, based on Rubinstein's (1989) email game, shows that common learning can fail when signal sets are infinite.

We consider common learning when the signal distributions are *not* independent over time. We are ultimately interested in applications in which agents' signals are affected by their actions, as in repeated games. We then cannot expect signals to be independent over time. This paper takes an intermediate step toward this goal, studying a tractable class of time-dependent signal processes. The signals are generated by an exogenously-specified, parameter-dependent hidden Markov process. In a hidden Markov process, there is a finite Markov chain determining the current state, that in turn determines the distribution of the agents' private signals. The signals may not be perfectly informative about the current state and so the state is hidden from the agents.

We begin in Section 3 with an example showing that even a seemingly tiny touch of intertemporal dependence can preclude common learning. Section 4 shows that if the hidden state becomes public infinitely often with probability one, then again we have common learning. For example, there may be a public signal that is uniquely associated with a single recurrent hidden state.

Sections 5.2 and 5.3 then present two sets of separation conditions on the hidden Markov process that suffice for common learning when there is no public information. When signals are generated by a hidden Markov process, learning in general calls for the agent to use the frequencies of the signals she observes and the intertemporal structure of these observations to draw inferences about the realized hidden states and so the parameter. However, drawing inferences about the likely realizations of hidden states is a notoriously difficult problem. Section 5.2 offers a "relative separation" condition, that the signal distributions generated by the different parameter values are not "too close" to one another, that is expressed solely in terms of expected signal frequencies, saying nothing about intertemporal patterns. We then establish common learning via an argument independent of agent's inferences about the history of realized states in the hidden Markov process.

The condition offered in Section 5.2 is quite strong. Section 5.3 offers a weaker separation condition. However, we must then supplement this condition with an additional assumption, intuitively requiring that unusual realizations of the states in the hidden Markov process cannot be too likely as explanations of observed signal frequencies. The two sets of conditions are thus not nested. We view the sufficient condition of Section 5.2 as being more demanding, though it has the advantage of being more concise and more intuitive, as well as more straightforward to verify.

Cripps, Ely, Mailath, and Samuelson (2008) is the obvious point of departure for this work. Like Cripps, Ely, Mailath, and Samuelson (2008), the primarily technical tool is the characterization of common p-belief offered by Monderer and Samet (1989).[1]

2 Common Learning

2.1 The Model

Nature first selects a parameter θ from the set $\Theta = \{\theta', \theta''\}$ according to the prior distribution p. There are two agents, denoted by $\ell = 1, 2$, who observe signals in the periods $t = 0, 1, 2, \ldots$.

Conditional on θ, the agents observe signals generated by a hidden Markov process (Ephraim and Merhav, 2002). We let X denote the finite set of states for this Markov process. The state in period zero, x_0, is generated by a parameter-dependent measure $\iota^\theta \in \Delta(X)$. The subsequent parameter-dependent transition probabilities are denoted by $\pi^\theta = \{\pi^\theta_{xx'}\}$, where $\pi^\theta_{xx'}$ is the probability that the Markov process is in state x' in some period t, given state x in period $t-1$ and parameter θ.

The agents do not observe the state of the Markov process. However, each agent ℓ observes in each period t a private signal $z_{\ell t} \in Z_\ell$. We assume each Z_ℓ is finite and let $z_t = (z_{1t}, z_{2t}) \in Z_1 \times Z_2 = Z$.

The signal profile z_t is independent across periods conditional on the parameter and the hidden state. The joint distribution of z conditional on x and θ is denoted by $\phi^{x\theta}$, so that the probability that $z_t = z$ is $\phi^{x\theta}_z$. We similarly denote the probability that $z_{\ell t} = z_\ell$ by $\phi^{x\theta}_{z_\ell}$ and denote the corresponding marginal distribution by $\phi^{x\theta}_\ell$. We denote the marginal distribution on agent ℓ's signals induced by the distribution $\phi^{x\theta}_\ell$ by ϕ^θ_ℓ and the ergodic distribution (when defined) over states by ξ^θ.

[1] Steiner and Stewart (2011) examine a setting in which (in its simplest form) one agent is informed of the parameter and the other receives signals whose distribution shares the essential features of our example in Section 3. Steiner and Stewart (2011) establish a common learning result for this setting, and then focus on the conditions under which adding communication to the model either preserves or disrupts the common learning.

A state of the world $\omega \in \Omega$ consists of a parameter and a sequence of hidden states and signal profiles, and hence $\Omega = \Theta \times X^\infty \times Z^\infty$. We use P (respectively P^θ) to denote the measure on Ω induced by the prior p (resp. parameter θ), the state process $(\iota^\theta, \pi^\theta)$ and the signal process ϕ^θ. We let $E[\,\cdot\,]$ and $E^\theta[\,\cdot\,]$ denote the expectations with respect to these measures. We abuse notation by often writing θ or $\{\theta\}$ for the event $\{\theta\} \times X^\infty \times Z^\infty$, so that θ and $\{\theta\}$ denote both a value of the parameter and an event in Ω.

A period-t history for agent ℓ is denoted by $h_{\ell t} = (z_{\ell 0}, z_{\ell 1}, \ldots, z_{\ell t-1}) \in H_{\ell t} = (Z_\ell)^t$; $\{\mathcal{H}_{\ell t}\}_{t=0}^\infty$ denotes the filtration induced on Ω by agent ℓ's histories. The random variables $P(\theta \mid \mathcal{H}_{\ell t})$, giving agent ℓ's posteriors on 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).

2.2 Common Learning

For any event $F \subset \Omega$, the $\mathcal{H}_{\ell t}$-measurable random variable $P(F \mid \mathcal{H}_{\ell t})$ is the probability agent ℓ attaches to F given her information at time t.

We define $B^q_{\ell t}(F)$ to be the set of states for which at time t agent ℓ attaches at least probability q to the event F:

$$B^q_{\ell t}(F) := \{\omega \in \Omega : P(F \mid \mathcal{H}_{\ell t})(\omega) \geq q\}.$$

Recalling that a state ω is an element of $\Theta \times X^\infty \times Z^\infty$, the set $B^q_{\ell t}(F)$ can be thought of as the set of t-length private histories $h_{\ell t}$ at which agent ℓ attaches at least probability q to the event F (since agent ℓ knows whether $B^q_{\ell t}(F)$ has occurred (i.e., $B^q_{\ell t}(F) \in \mathcal{H}_{\ell t}$)).

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

$$B^q_t(F) := B^q_{1t}(F) \cap B^q_{2t}(F).$$

Note that while agent 1 knows whether $B^q_{1t}(F)$ has occurred, he need *not* know whether $B^q_{2t}(F)$, and so $B^q_t(F)$, has occurred. The event that F is *common q-belief* at date t is

$$C^q_t(F) := B^q_t(F) \cap B^q_t(B^q_t(F)) \cap \cdots = \bigcap_{n \geq 1} [B^q_t]^n (F).$$

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

We say 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 1 (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 $t > T$,*

$$P^\theta(C^q_t(\theta)) > q.$$

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

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

2.3 Sufficient Conditions for Common Learning

The countable collection of events $\left\{[B_t^q]^n(\theta)\right\}_{n\geq 1}$ can be cumbersome to work with, and it is often easier to approach common learning with the help of a characterization in terms of q-evident events. An event F is *q-evident* at time t if it is q-believed when it is true, that is,

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

From Monderer and Samet (1989, Definition 1 and Proposition 3), we have:

Proposition 1 *The event 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 $\omega \in F \subset B_t^q(F')$.*

We use the following immediate implication:

Corollary 1 *The agents commonly learn θ if and only if for all $q \in (0,1)$, there exists a sequence of events $\{F_t\}_t$ and a period T such that for all $t > T$,*

(i) $P^\theta(F_t) > q$,

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

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

It is straightforward to establish common learning when the signals are independent across players. More precisely, suppose that for each t, the private signal histories h_{1t} and h_{2t} are (conditionally on θ) independent.[2] Applying Corollary 1 to the events $F_t = B_t^{\sqrt{q}}(\theta)$ then shows that common learning holds when agents individually learn (Cripps, Ely, Mailath, and Samuelson, 2008, Proposition 2). This simple argument does not rely on finite signal and parameter spaces, being valid for arbitrary signal and parameter spaces.

The relationship between individual and common learning is more subtle when the signal histories are not conditionally independent across agents. Cripps, Ely, Mailath, and Samuelson (2008) study the case where the signals are conditionally (on θ) independent over time, rather than being determined by a hidden Markov process, but with arbitrary correlation between different agent's signals within a period. Individual learning is then equivalent to the marginal distributions of the private signals being distinct. Cripps, Ely, Mailath, and Samuelson's (2008) main result is:

Proposition 2 *Suppose that signals are conditionally (on θ) independently distributed across time (so that $\pi^\theta_{xx'} = \pi^\theta_{x^\dagger x'}$ for all x, $x^\dagger$, x', and θ), and that the agents individually learn (so that the marginal distributions are distinct, i.e., $\phi_\ell^{\theta'} \neq \phi_\ell^{\theta''}$ for all ℓ). Then the agents commonly learn Θ.*

[2] In this case, the intertemporal dependence can be quite general, and need not be described by a finite state hidden Markov process.

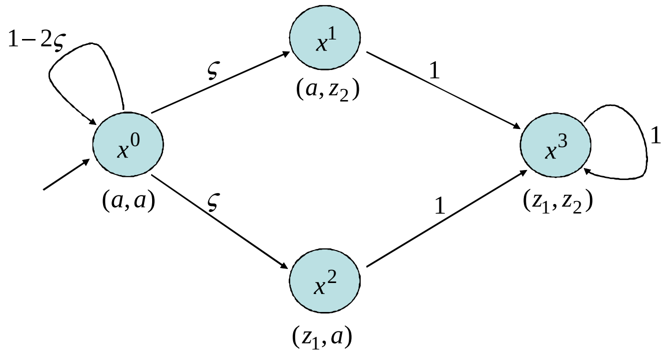

Figure 1: The hidden Markov process for our example, where $0 < \zeta < \frac{1}{2}$. The probabilities on the state transitions are indicated above the transitions, and the signal realizations possible in each state are indicated below the state, with $z_1, z_2 \in \{b, c\}$. The process begins in state x^0, and in each period stays in x^0 with probability $1 - 2\zeta$, and transits with equal probability to either x^1 or x^2.

This result requires the agents' signal spaces to be finite. Cripps, Ely, Mailath, and Samuelson (2008, Section 4) provide an example showing that common learning can fail when signals are conditionally independent across time (but not agents), but drawn from infinite sets.

3 An Example with No Common Learning

We present here an example in which intertemporal dependence in the signal distributions prevents common learning. There are two values of the parameter θ, given by θ' and θ'' with $0 \leq \theta' < \theta'' \leq 1$ and a hidden Markov process with four states, x^k, $k = 0, 1, 2$, and 3. There are three signals, denoted by a, b, and c, i.e., $Z_\ell = \{a, b, c\}$. State x^0 is the initial state, and invariably generates the signal pair (a, a). The signal distributions in the other three states are independent across agents, conditional on the state and parameter, and are given by (for $\ell = 1, 2$ and $j = 1, 2, 3$)

$$\phi_\ell^{x^j \theta} = \begin{cases} (1, 0, 0) & \text{if } j = \ell, \\ (0, \theta, 1 - \theta) & \text{otherwise} \end{cases}. \tag{1}$$

Figure 1 illustrates the Markov process and specifies the transition probabilities. In the special case where $\theta' = 0$ and $\theta'' = 1$, we have essentially the "clock" scenario of Halpern and Moses (1990, p. 568). We refer to the case presented here as the "noisy clock" example.

The process begins in state x^0, generating an uninformative a signal for each agent, and generating a string of such signals as long as it remains in state x^0. However, with probability 1 (under both θ' and θ''), the Markov process eventually makes a transition to either state x^1 or x^2 (each transition being equally likely), with state x^ℓ generating signal a for agent ℓ and either b or c for the other agent. The Markov process then necessarily moves to state x^3, at which point no further a signals are observed. Here, each player independently draws signal b with probability θ and signal c with probability $1-\theta$, so that the subsequent frequencies of signals b and c reveal the parameter. We thus have individual learning.

The agents do not commonly learn the parameter. Instead, in reasoning reminiscent of Rubinstein's (1989) email game, an agent who has seen a string of a signals (before switching to either signal b or signal c, and never subsequently observing another a) knows that the other agent has observed either one more or one less a signal. This sets off an infection argument, with the agents forming iterated beliefs that attach significant probability to ever-longer strings of a signals, culminating in a belief that one agent has seen nothing but a's. But then that agent has learned nothing, precluding common learning. More formally:

Lemma 1 *In the noisy clock example, at any history $h_{\ell T}$, each agent has iterated q-belief that the other agent has observed at least $T-1$ periods of a, where $q = (1-2\zeta)/(2-2\zeta)$.*

Proof. Fix T, and let $A_{\ell t}$ be the event that agent ℓ has observed precisely $t > 1$ signal a's in the history $h_{\ell T}$, for $t < T$. Then given agent ℓ has observed t signal a's, he knows that agent $\tilde{\ell}$ has seen either one more a or one less. We have

$$\begin{aligned} P^\theta(A_{\ell t} \cap A_{\tilde{\ell} t-1}) &= (1-2\zeta)^{t-2}\zeta \\ \text{and} \qquad P^\theta(A_{\ell t} \cap A_{\tilde{\ell} t+1}) &= (1-2\zeta)^{t-1}\zeta, \end{aligned}$$

and so

$$P^\theta(A_{\tilde{\ell} t+1} \mid A_{\ell t}) = \frac{1-2\zeta}{2-2\zeta} = q.$$

Thus, conditional on observing $A_{\ell t}$, agent ℓ attaches at least probability q to the event $A_{\tilde{\ell} t+1}$. Or, in the language of belief operators,

$$A_{\ell t} \subset B^q_\ell(A_{\tilde{\ell} t+1}), \qquad A_{\tilde{\ell} t+1} \subset B^q_{\tilde{\ell}}(A_{\ell t+2}), \qquad \cdots .$$

Iterating, we get

$$\begin{aligned} A_{\ell T} &\subset [B^q_\ell B^q_{\tilde{\ell}}]^{\frac{T-t-1}{2}} B^q_\ell A_{\tilde{\ell} T}, && \text{for } T-t \text{ odd,} \\ \text{and} \qquad A_{\ell T} &\subset [B^q_\ell B^q_{\tilde{\ell}}]^{\frac{T-t}{2}} A_{\ell T}, && \text{for } T-t \text{ even.} \end{aligned}$$

■

Lemma 1 implies that each agent has iterated q-belief that that the other player's posterior on θ is close to his prior, and so common learning fails (Morris, 1999, Lemma 14).

This example generalizes to one in which the signal distributions have full support in each state. Suppose that, in each state, with probability $1-9\varepsilon$, the signals are distributed as in (1) and Figure 1, and with probability 9ε, there is a uniform draw from the set of joint signals $\{aa, ab, ac, ba, bb, bc, ca, cb, cc\}$. We again have a failure of common learning. Let $\tilde{\tau}$ be the first date at which the process is not in state x^0. There exists $\eta > 0$ such that at any time t and conditional on $\tilde{\tau} > \tau$ for any $\tau < t$, there is probability at least η that agent 2 observes a history h_{2t} such that $\Pr^\theta(\tilde{\tau} > \tau + 1 \mid h_{2t}) > \eta$ (Appendix A.1 contains a proof). The same statement holds reversing the roles of agents 1 and 2, and so there is iterated η-belief in $\tilde{\tau} = t$. Since the signals are uninformative about the parameter in state x^0, there is then iterated η-belief that the agents do not learn the parameter.

4 Resets: A Block-Based Condition for Common Learning

Our first positive result requires that there is a public signal "0" that reveals some recurrent state $\bar{x}$ of the hidden Markov process. Either both agents observe the signal 0 or neither do, signal 0 is observed with unitary probability in state $\bar{x}$, and signal 0 is never observed in another state. As a result, the hidden state becomes public infinitely often with probability one. We refer to this as a "reset," since observing signal 0 allows an agent to begin a new process of forming expectations about the other agent's signals.

The periodic *public* identification of the state breaks an agent's private history into "blocks" of consecutive signals, with a new block beginning each time the signal 0 is observed. The string of signals within each block can be viewed as a single signal, drawn from a countably infinite (since block lengths are unbounded) set of signals. By the Markov property (and the common knowledge nature of the signal 0), the strings of signals observed within a block are independent across blocks. We have thus transformed a model of time-dependent signals selected from a finite set to a model where, by time t, the agents will have observed a random number of time-independent signals selected from a countable set of block signals.

Moving to time-independent signals is useful because it allows us to apply a result from Cripps, Ely, Mailath, and Samuelson (2008). At the same time, we have lost the finite-signal-set assumption used in our earlier positive result. Nevertheless, the length of each block is common knowledge, precluding the infections in beliefs that can disrupt common learning. However, the unbounded block lengths give rise to a second difficulty, namely unbounded likelihood ratios—arbitrarily long blocks of private signals can be arbitrarily informative.

Applying the arguments from Cripps, Ely, Mailath, and Samuelson (2008) to histories where *all* block lengths are less than some constant $\mathfrak{c}$ does yield a sequence of self-evident events. The events in the sequence restrict, for each block length, the frequency of different blocks to be in an appropriate neighborhood of a distribution over blocks of signals. However, since arbitrarily long blocks arise eventually with probability one, the probability of the events in the sequence converges to zero asymptotically. To obtain common learning on a sequence of events requires that the sequence accommodate arbitrarily long blocks of signals. The sequence of events we use allows for arbitrarily long blocks, but restricts, for each block length less than $\mathfrak{c}$, the frequency of different blocks to be in the appropriate neighborhood. A key idea is that we do not restrict the frequency of different blocks greater than $\mathfrak{c}$, but we only consider histories on which the average length of all blocks observed (including those longer than $\mathfrak{c}$) is close to its expected value. This ensures that atypical long blocks have only a small effect on beliefs, and so do not upset the individual learning and self-evidence implied by the restriction on block lengths less than $\mathfrak{c}$.

4.1 Assumptions and Common Learning

We make four assumptions on the signal processes that determine the measure P^{θ}. Our first assumption is that the process on the hidden states is ergodic:

Assumption 1 (Ergodicity) *For all θ, the hidden Markov process π^{θ} is aperiodic and irreducible. The implied stationary distribution on X is denoted by ξ^{θ}.*

The second assumption is technical, while the last two are substantive.

We work with signal-generating processes with the property that no signal reveals the true parameter with probability one. This simplifies the analysis by eliminating nuisance cases, such as cases in which likelihood ratios are undefined.[3] A full-support assumption would ensure this, but we cannot literally invoke a full-support assumption in the presence of resets, since the definition of a reset ensures that if agent ℓ observes signal 0, agent $\tilde{\ell}$ cannot observe another signal. The following assumption accommodates resets while still conveniently excluding cases in which the posterior jumps to one. Under part 1 of this assumption, all of an agent's signals occur with positive probability under both parameters, that is, $\sum_x \xi_x^{\theta}\phi_{z_\ell}^{x\theta} > 0$ for all θ, ℓ, z_ℓ. This avoids the possibility that beliefs might jump to unity because agent ℓ's signals have different supports for different parameter values. Similarly if $P^{\theta'}(z_{\ell t} = z_\ell | z_{\ell,t-1} = z'_\ell) = 0$ but $P^{\theta''}(z_{\ell t} = z_\ell | z_{\ell,t-1} = z'_\ell) > 0$, then observing such a transition will cause agent ℓ's posterior on θ'' to jump to unity, and the second part of the assumption precludes this possibility.

Assumption 2 (Common Support)

1. *$\sum_x \xi_x^{\theta}\phi_{z_\ell}^{x\theta} > 0$ for all ℓ, θ and all $z_\ell \in Z_\ell$.*
2. *$\pi_{xx'}^{\theta'} = 0$ if and only if $\pi_{xx'}^{\theta''} = 0$.*

The next two assumptions provide the key ingredients for our result. The first of these assumptions is necessary and sufficient to ensure that each agent can learn

[3] For example, we would have to augment the definition of relative entropy in (3) by specifying a value for those cases in which the probability in the denominator is zero.

the parameter, hence providing the individual learning condition of our desired "individual learning implies common learning" result.[4]

We need two pieces of notation. First, when signals are time-dependent, correlations between current and future signals convey information about the value of the underlying parameter. The probability that an arbitrary ordered pair of signals $z_{\ell t} z_{\ell,t+1}$ is realized under P^θ is

$$P^\theta(z_{\ell t} z_{\ell,t+1}) = \sum_x \xi_x^\theta \phi_{z_{\ell t}}^{x\theta} \sum_{x'} \pi_{xx'}^\theta \phi_{z_{\ell,t+1}}^{x'\theta}. \tag{2}$$

Second, given two distributions, p and q defined on a common outcome space $\mathcal{E}$, their *relative entropy*, or *Kullback-Leibler distance*, is given by

$$H(p\|q) = \sum_{e\in\mathcal{E}} p(e) \log \frac{p(e)}{q(e)}. \tag{3}$$

The Kullback-Leibler distance is always nonnegative, and equals zero only when $p = q$ (Cover and Thomas, 1991, Theorem 2.6.3). However, it is not a metric, since it is not symmetric and does not satisfy the triangle inequality.

Assumption 3 (Identification) *For $\ell \in \{1, 2\}$ and $\theta \neq \tilde{\theta}$, there exists $\beta_\ell > 0$ such that*

$$H\left(P^\theta(z_{\ell t} z_{\ell,t+1}) \middle\| P^{\tilde{\theta}}(z_{\ell t} z_{\ell,t+1})\right) = E^\theta\left(\log \frac{P^\theta(z_{\ell t} z_{\ell,t+1})}{P^{\tilde{\theta}}(z_{\ell t} z_{\ell,t+1})}\right) = \beta_\ell.$$

The final assumption is the reset condition: there exists a public signal identifying a state in the hidden Markov process.

Assumption 4 (Resets) *There exists a state $\bar{x} \in X$ and a signal $0 \in Z_1 \cap Z_2$ such that*

$$\phi_{z_\ell}^{x\theta} = \begin{cases} 1 & \textit{if } x = \bar{x},\ z_\ell = 0 \\ 0 & \textit{if } x \neq \bar{x},\ z_\ell = 0. \end{cases}$$

The signal 0 is a public signal that reveals the hidden state $\bar{x}$: either both agents observe it or neither do, and it is never observed in a state other than $\bar{x}$. It is without loss of generality to assume that the signal 0 appears with probability 1 in state $\bar{x}$, since otherwise we could split x into two states, one featuring signal 0 with probability 1 and one featuring signal 0 with probability 0. The pair of zero signals is also denoted 0.

These assumptions suffice for common learning:

Proposition 3 *If the signal process satisfies Assumptions 2–4, then the agents commonly learn the parameter θ.*

[4] Assumption 3 is necessary and sufficient for identifying the parameter when the hidden Markov process is irreducible (as assumed here). See Ephraim and Merhav (2002, p. 1439) and the references therein for details and conditions identifying the parameter in more general models.

4.2 Proof of Proposition 3: Preliminary Considerations

4.2.1 Blocks of Data

For a given history (h_{1t}, h_{2t}), define $\tau_1, \tau_2, ..., \tau_{N+1}$ to be the times (in order) that the agents observed the public signal 0. We use N to denote the (random) number of completed blocks observed before time t and use $n = 1, ..., N$ to count these block-signals, suppressing the dependence of N on t in the notation. Let $\bar{z}_\ell^o$ denote the block of signals observed up to and including the first zero signal. We define $\bar{z}_{\ell n}$ to be the block of signals observed between the n^{th} and $n+1^{\text{st}}$ zero signal (if there are no such signals $\bar{z}_{\ell n} = \varnothing$). Finally, we define $\bar{z}_\ell^e$ to be the block of signals after the last public signal (and the empty set again if they do not exist). That is,

$$\begin{aligned}
&\bar{z}_\ell^o = (z_{\ell 0}, z_{\ell 1}, \ldots, z_{\ell,\tau_1}) = (z_{\ell 0}, z_{\ell 1}, \ldots, z_{\ell \tau_1 - 1}, 0), \\
&\bar{z}_{\ell n} = (z_{\ell,\tau_n+1}, z_{\ell,\tau_n+2}, \ldots, z_{\ell,\tau_{n+1}-1}), \\
\text{and} \quad &\bar{z}_\ell^e = (z_{\ell,\tau_{N+1}+1}, z_{\ell,\tau_{N+1}+2}, \ldots, z_{\ell t}).
\end{aligned} \tag{4}$$

We use $b_s = (z_1, z_2, \ldots, z_s) \in B_s$ to denote a generic block of non-zero signal profiles of length s, where $B_s = B_{1s} \times B_{2s}$ and $B_{\ell s} = (Z_\ell \setminus \{0\})^s$.[5] The countable set of all possible agent-ℓ signal blocks is $B_\ell = \bigcup_{s=0}^\infty B_{\ell s}$ (where $B_{\ell 0} = \varnothing$). We define $\zeta^\theta(b_s)$ to be the θ-probability that a given block of data b_s occurs between two zero signals, that is,

$$\zeta^\theta(b_s) = P^\theta(z_1, z_2, \ldots, z_s, z_{s+1} = 0 \mid z_0 = 0), \\
\forall b_s = (z_1, z_2, \ldots, z_s) \in B := \cup_{s=0}^\infty B_s. \tag{5}$$

The probability that a zero signal immediately follows another zero signal is $\zeta^\theta(\varnothing)$. The measure ζ^θ is uniquely defined by the transition and signal probabilities. The Markov process is stationary and the zero signal is realized infinitely often with probability one. Therefore,

$$\sum_s \sum_{b_s \in B_s} \zeta^\theta(b_s) = 1.$$

Order the set of possible blocks of signals each player can receive by length, beginning with the shortest block (the empty set or 0-block), then the possible 1-block signals ordered arbitrarily, and so on. We refer to a given signal block for agent 1 as $b_{si} \in B_{1s}$, so that b_{si} denotes the i-th element of the set of s-blocks. We perform a similar operation on agent 2's blocks writing them as $b_{sj} \in B_{2s}$, where j ranges over all s-blocks for agent 2. This notation implies that any $b \in B$ can be referred to as a triple sij where s is the (public) length of the block and b_{si} (b_{sj}) is the data agent 1 (agent 2) observed. The marginals are

$$\zeta_{si}^\theta = \sum_j \zeta_{sij}^\theta \text{ and } \zeta_{sj}^\theta = \sum_i \zeta_{sij}^\theta, \quad \text{where } \zeta_{sij}^\theta = \zeta^\theta(b_{sij}).$$

[5] It is possible that not all such blocks occur with positive probability when preceded and followed by zero signals (that is blocks of length $s+2$ of the form $(0, \bar{z}_s, 0)$).

We summarize an agent's history, $h_{\ell t}$, by an initial block, $\bar{z}_\ell^o \in B_\ell$, a terminal block, $\bar{z}_\ell^e \in B_\ell$, and a potentially large but random number N of full blocks in B_ℓ. The data collected by the agent 1 is summarized by a triple $(\bar{z}_1^o, \bar{z}_1^e, f_{si}^t)$ where $f_{si}^t \in \mathbb{N}$ records the number of observations of the block b_{si} by agent 1 before time t. Similarly, agent 2's data is summarized by $(\bar{z}_2^o, \bar{z}_2^e, f_{sj}^t)$.

The process generating signals is ergodic, so there is an exponential upper bound on the arrival times of the zero signal. Denote by σ the time of first observation of the zero signal, that is, $\sigma = \min\{t \geq 0 : z_t = 0\}$. The ergodicity of the hidden Markov process and Assumption 4 imply that for any state $x \neq \bar{x}$ and any θ there is a strictly positive probability of moving within $|X|-1$ steps from state x to state $\bar{x}$, at which point signal 0 necessarily appears. Let $\rho > 0$ be the minimum such probability, where we minimize over the $|X|$ possible initial states,

$$\rho = \min_x P^\theta\{\sigma < \infty \mid x_0 = x\} = \min_x P^\theta\{\sigma \leq |X| - 1 \mid x_0 = x\}. \tag{6}$$

Starting from anywhere, therefore, the probability that state $\bar{x}$ is visited in the next $|X|-1$ periods is at least ρ. Hence $(1-\rho)P^\theta(\sigma \geq t \mid x_0) \geq P^\theta(\sigma \geq t + |X| \mid x_0)$. A simple calculation then gives

$$P^\theta(\sigma \geq t \mid x_0) \leq \frac{\lambda^t}{(1-\rho)}, \quad \text{where } \lambda = (1-\rho)^{1/|X|} < 1. \tag{7}$$

We note the following for future reference.

Lemma 2 *The expected time to the first realization of the zero signal is finite, as is the expected length of full blocks. The variance of the length of full blocks is also finite.*

Proof. The expected time till the first realization of the zero signal satisfies

$$\begin{aligned}
E^\theta(\sigma \mid x_0) &= \sum_{s=1}^{\infty} s\left[P^\theta(\sigma \geq s \mid x_0) - P^\theta(\sigma \geq s+1 \mid x_0)\right] \\
&= \sum_{s=1}^{\infty} sP^\theta(\sigma \geq s \mid x_0) - \sum_{s=2}^{\infty}(s-1)P^\theta(\sigma \geq s \mid x_0) \\
&= \sum_{s=1}^{\infty} P^\theta(\sigma \geq s \mid x_0) \leq \sum_{s=1}^{\infty} \lambda^s/(1-\rho) = \frac{\lambda}{(1-\lambda)(1-\rho)} < \infty.
\end{aligned}$$

Since the minimum in (6) is taken over all x, including $\bar{x}$, this calculation also bounds this expected length of a full block (take $x_0 = \bar{x}$). A similar argument shows that the variance is finite. ■

4.2.2 Posterior Beliefs

We now show that the agents' posterior beliefs can be written as a function of the frequencies of agents' blocks of data.

Agent ℓ's posterior at time t, $p^{\theta}_{\ell t}$, is the $h_{\ell t}$-measurable random variable describing the probability that agent ℓ attaches to the parameter θ at time t given the observed data. From Bayes' rule, we have

$$\mathcal{L}_{\ell t} := \log \frac{P^{\theta'}(h_{\ell t})}{P^{\theta''}(h_{\ell t})} = \log \frac{p^{\theta'}_{\ell t}}{1 - p^{\theta'}_{\ell t}} - \log \frac{p^{\theta'}_{0}}{1 - p^{\theta'}_{0}}. \tag{8}$$

Repeatedly conditioning on the arrival times τ_m of the zero signal for the first equality, then applying the Markov assumption and the fact that 0 signals are public for the second equality, a substitution from (5) for the third, and finally, defining n^t_{si} for the number of observations of block b_{si} in t periods, gives

$$\begin{aligned}
\mathcal{L}_{1t} &= \log \frac{P^{\theta}(h_{1\tau_1})P^{\theta}(h_{1\tau_2} \mid h_{1\tau_1}) \cdots P^{\theta}(h_{1\tau_{N+1}} \mid h_{1\tau_N})P^{\theta}(h_{1t} \mid h_{1\tau_{N+1}})}{P^{\tilde\theta}(h_{1\tau_1})P^{\tilde\theta}(h_{1\tau_2} \mid h_{1\tau_1}) \cdots P^{\tilde\theta}(h_{1\tau_{N+1}} \mid h_{1\tau_N})P^{\tilde\theta}(h_{1t} \mid h_{1\tau_{N+1}})} \\
&= \log \frac{P^{\theta}(h_{1\tau_1})P^{\theta}(h_{1\tau_2} \mid z_{\tau_1} = 0) \cdots P^{\theta}(h_{1\tau_{N+1}} \mid z_{\tau_N} = 0)P^{\theta}(h_{1t} \mid z_{\tau_{N+1}} = 0)}{P^{\tilde\theta}(h_{1\tau_1})P^{\tilde\theta}(h_{1\tau_2} \mid z_{\tau_1} = 0) \cdots P^{\tilde\theta}(h_{1\tau_{N+1}} \mid z_{\tau_N} = 0)P^{\tilde\theta}(h_{1t} \mid z_{\tau_{N+1}} = 0)} \\
&= \log \frac{P^{\theta}(\bar z^o_1)\zeta^{\theta}(\bar z_{11}) \cdots \zeta^{\theta}(\bar z_{1N})P^{\theta}(\bar z^e_1 \mid z_{\tau_{N+1}} = 0)}{P^{\tilde\theta}(\bar z^o_1)\zeta^{\tilde\theta}(\bar z_{11}) \cdots \zeta^{\tilde\theta}(\bar z_{1N})P^{\tilde\theta}(\bar z^e_1 \mid z_{\tau_{N+1}} = 0)} \\
&= \log \frac{P^{\theta}(\bar z^o_1)}{P^{\tilde\theta}(\bar z^o_1)} + \log \frac{P^{\theta}(\bar z^e_1)}{P^{\tilde\theta}(\bar z^e_1)} + \sum_{si} n^t_{si} \log \frac{\zeta^{\theta}_{si}}{\zeta^{\tilde\theta}_{si}}.
\end{aligned}$$

We exclude from the summation in the last line any signal profiles b_{si} that occur with zero probability (under all θ).

Recall that N is the number of completed blocks observed by time t. We can write agent 1's beliefs as a sum of independent random variables: the log-likelihood of the data before the first zero, the log-likelihood of the data after the last zero and the empirical measure

$$\hat\zeta^t_{si} = n^t_{si}/N$$

of the block data. That is,

$$\log \frac{p^{\theta}_{1t}}{1 - p^{\theta}_{1t}} = \log \frac{p^{\theta}_{0}}{1 - p^{\theta}_{0}} + \log \frac{P^{\theta}(\bar z^o_1)}{P^{\tilde\theta}(\bar z^o_1)} + \log \frac{P^{\theta}(\bar z^e_1)}{P^{\tilde\theta}(\bar z^e_1)} + N \sum_{si} \hat\zeta^t_{si} \log \frac{\zeta^{\theta}_{si}}{\zeta^{\tilde\theta}_{si}}. \tag{9}$$

This equation expresses the posterior, p^{θ}_{1t}, in terms of the data $(\bar z^o_1, \bar z^e_1, f^t_{si})$ (an equivalent argument holds for agent 2).

4.2.3 The Sequence of Events

We now describe the class of events we use to establish common learning. The events we consider depend on a mixture of private and public information. The public information is the lengths of blocks. We require the initial and terminal blocks to be not too long. We also require that the average length of the completed blocks be close to its expected length. This allows the agent to observe some long blocks, but prevents rare events having particularly perverse effects on the agents' private beliefs. The private event is that the agents' observations of block signals of length less than some number $\mathfrak{c}$ are close to their expected frequencies.

The first event $I_t(\mathfrak{b})$ is the public event that the initial and terminal blocks are not long. The parameter $\mathfrak{b} \in \mathbb{N}$ bounds the length of the initial and terminal blocks:

$$I_t(\mathfrak{b}) = \{(h_{1t}, h_{2t}) : \max\{\tau_1, t - \tau_N\} \leq \mathfrak{b} \ \}.$$

The second event $M_t(\alpha, \theta)$ is the public event that the average length of the blocks that are complete is close to the expected length of blocks under the relevant parameter. The average length of the completed blocks is $\sum_{si} s\hat{\zeta}^t_{si} = \sum_{sj} s\hat{\zeta}^t_{sj}$. The expected length is $\sum_{si} s\zeta^\theta_{si} = \sum_{sj} s\zeta^\theta_{sj}$. The parameter α determines how close the mean block length is to its expected value:

$$M_t(\alpha, \theta) = \left\{ (h_{1t}, h_{2t}) : \left| \sum_{sk} s(\hat{\zeta}^t_{sk} - \zeta^\theta_{sk}) \right| < \alpha, \ k = i, j \right\}.$$

To define the private event that agents' observed signal frequencies are close to their expected values, we first consider the following $\mathfrak{c}$-truncation of the model. In every period $N = 0, 1, 2, \ldots$, each agent observes one of a finite number of block signals, drawn from $\{b_{si}\}_{s \leq \mathfrak{c}} \cup \{b^*\}$, for agent 1 and $\{b_{sj}\}_{s \leq \mathfrak{c}} \cup \{b^*\}$ for agent 2. The pair (b_{si}, b_{sj}) is selected in each period from the ζ-distribution of block signals generated by θ, with the signal b^* replacing any block longer than $\mathfrak{c}$. The joint distribution of the agents' signals in the $\mathfrak{c}$-truncation is

$$\begin{aligned} P^\theta(b_{sij}) = \zeta^\theta_{sij} &=: \varphi^\theta_{sij}, \qquad \forall s \leq \mathfrak{c}; \text{ and} \\ 1 - \sum_{s > \mathfrak{c}, i, j} \zeta^\theta_{sij} &=: \varphi^\theta_{b^*}, \qquad \text{otherwise.} \end{aligned}$$

We use $\varphi^\theta_{si} := \sum_j \varphi^\theta_{sij}$ and $\varphi^\theta_{sj} := \sum_i \varphi^\theta_{sij}$ to denote the agents' marginals for signals and $(\hat{\varphi}^N_{si}, \hat{\varphi}^N_{sj}, \hat{\varphi}^N_{b^*})$ to denote the agents' empirical measure at time N.

Cripps, Ely, Mailath, and Samuelson (2008, Proposition 3) covers the $\mathfrak{c}$-truncation model, and so there is common learning in that model. From the proof of that proposition, we have the following for the $\mathfrak{c}$-truncation model:[6]

For all $\varepsilon > 0$, there exists $\delta \in (0, \varepsilon)$ and a sequence of events $\{F^\theta_N(\varepsilon)\}^\infty_{N=0}$ given by $F^\theta_N(\varepsilon) = \{\theta\} \cap G^\theta_{1N}(\varepsilon) \cap G^\theta_{2N}(\varepsilon)$, where $G^\theta_{\ell N}(\varepsilon)$ is an event on ℓ's private signal profiles, such that, for all $N \in \mathbb{N}$,

$$((\hat{\varphi}^N_{si})_{s \leq \mathfrak{c}}, \hat{\varphi}^N_{b^*}) \in G^\theta_{1N}(\varepsilon) \Longrightarrow \sum_{s \leq \mathfrak{c}, i} |\hat{\varphi}^N_{si} - \varphi^\theta_{si}| < \varepsilon,\text{[7]} \tag{P1}$$

[6]Condition (P1) follows from Cripps, Ely, Mailath, and Samuelson (2008, (13)–(14)) (modulo different notation) while (P2) is established in Cripps, Ely, Mailath, and Samuelson (2008, page 926).

[7]Since frequencies sum to one, the consequent of (P1) implies $\sum_{s \leq \mathfrak{c}, i} |\hat{\varphi}^N_{si} - \varphi^\theta_{si}| + |\hat{\varphi}^N_{b^*} - \varphi^\theta_{b^*}| < 2\varepsilon$; and a similar comment applies to the antecedent of (P2).

and

$$\sum_{s\leq\mathfrak{c},i} |\hat{\varphi}^N_{si} - \varphi^\theta_{si}| < \delta \Longrightarrow ((\hat{\varphi}^N_{si})_{s\leq\mathfrak{c}}, \hat{\varphi}^N_{b^*}) \in G^\theta_{1N}(\varepsilon), \tag{P2}$$

with a similar property holding for agent 2. Moreover, for all $q \in (0,1)$, there is a $T_\varepsilon \in \mathbb{N}$ such that, for all $N > T_\varepsilon$,

$$F^\theta_N(\varepsilon) \text{ is } q\text{-evident.} \tag{P3}$$

Returning to the untruncated model, we apply Corollary 1 to the intersection of $I_t(\mathfrak{b})$ (initial and terminal blocks are not too long), $M_t(\alpha,\theta)$ (the average block length is close to its expectation), $N > T_\varepsilon$ (there are more than T_ε completed blocks), and the event that agents' frequencies for blocks are in the sets $G^\theta_{\ell N}(\varepsilon)$ defined above. To make this precise, recall that $\hat{\zeta}^t_{si}$ denotes the frequency of the signal blocks b_{si} received by agent 1 over the N completed blocks observed in the first t periods, and denote the frequency of the N blocks observed longer than $\mathfrak{c}$ in the first t periods by $\hat{\zeta}^t_{b^*}$. Then

$$\widetilde{G}^\theta_{1t}(\varepsilon) := \{(h^t_1, h^t_2) : ((\hat{\zeta}^t_{si})_{s\leq\mathfrak{c}}, \hat{\zeta}^t_{b^*}) \in G^\theta_{1N}(\varepsilon)\},$$

and *mutatis mutandis* $\widetilde{G}^\theta_{2t}(\varepsilon)$. Define

$$\widetilde{F}^\theta_t := \{\theta\} \cap I_t(\mathfrak{b}) \cap M_t(\alpha,\theta) \cap \{N \geq T_\varepsilon\} \cap \widetilde{G}^\theta_{1t}(\varepsilon) \cap \widetilde{G}^\theta_{2t}(\varepsilon). \tag{10}$$

To complete the specification of the sequence $\{\widetilde{F}^\theta_t\}_t$, we must specify ε, $\mathfrak{b}$, α, and $\mathfrak{c}$. We take $\mathfrak{b} = \log t$, and the values of ε, α and $\mathfrak{c}$ are determined in Lemma 5 below.

4.3 Proof of Proposition 3: Common Learning

We define some events that are helpful in constructing bounds on $\widetilde{F}^\theta_t$. Define

$$B_{\ell t}(\mathfrak{b},\alpha,\mathfrak{c},\varepsilon',\theta) = I_t(\mathfrak{b}) \cap M_t(\alpha,\theta) \cap S_{\ell t}(\mathfrak{c},\varepsilon',\theta), \tag{11}$$

where

$$S_{\ell t}(\mathfrak{c},\varepsilon',\theta) = \left\{ h_{\ell t} : \sum_{s\leq\mathfrak{c},i} s|\hat{\zeta}^t_{si} - \zeta^\theta_{si}| < \mathfrak{c}\varepsilon' \right\}.$$

The event $S_{\ell t}(\mathfrak{c},\varepsilon',\theta)$ is the private event that all the signal blocks shorter than $\mathfrak{c}$ for agent ℓ occur with close to the expected frequency under the relevant parameter. Note that we do not require that the signal blocks be shorter than $\mathfrak{c}$.

We bound $\widetilde{G}^\theta_{\ell t}(\varepsilon)$ by events $S_{\ell t}(\mathfrak{c},\varepsilon',\theta)$ for different values of ε'. In particular, taking $\varepsilon' = \varepsilon$, for all $((\hat{\phi}^t_{si})_{s\leq\mathfrak{c},i}, \hat{\phi}^t_{b^*}) \in \widetilde{G}^\theta_{1t}(\varepsilon)$ (with a similar comment for player 2), we have, from (P1),

$$\sum_{s\leq\mathfrak{c},i} |\hat{\zeta}^t_{si} - \zeta^\theta_{si}| < \varepsilon,$$

and hence

$$\mathfrak{c}\varepsilon > \sum_{s\leq \mathfrak{c},i} \mathfrak{c}|\hat{\zeta}^t_{si} - \zeta^\theta_{si}| > \sum_{s\leq \mathfrak{c},i} s|\hat{\zeta}^t_{si} - \zeta^\theta_{si}|,$$

and so $\widetilde{G}^\theta_{1t}(\varepsilon) \subset S_{1t}(\mathfrak{c}, \varepsilon, \theta)$ and hence $\widetilde{F}^\theta_t \subset B_{\ell t}(b, \alpha, \mathfrak{c}, \varepsilon, \theta)$.

Similarly, taking $\varepsilon' = \delta$, we have $\sum_{s\leq \mathfrak{c},i} s|\hat{\zeta}^t_{si} - \zeta^\theta_{si}| < \delta$ implies $\sum_{s\leq \mathfrak{c},i} |\hat{\zeta}^t_{si} - \zeta^\theta_{si}| < \delta$, and so (by (P2)) $S_{1t}(\mathfrak{c}, \delta, \theta) \subset \widetilde{G}^\theta_{1t}(\varepsilon)$. Hence, we can bound $\tilde{F}^\theta_t$ by

$$\begin{aligned} \{\theta\} \cap \{N \geq T_\varepsilon\} \cap B_{1t}(\mathfrak{b}, \alpha, \mathfrak{c}, \delta, \theta) \cap B_{2t}(\mathfrak{b}, \alpha, \mathfrak{c}, \delta, \theta) \\ \subset \widetilde{F}^\theta_t \subset B_{1t}(\mathfrak{b}, \alpha, \mathfrak{c}, \varepsilon, \theta) \cap B_{2t}(\mathfrak{b}, \alpha, \mathfrak{c}, \varepsilon, \theta). \end{aligned} \tag{12}$$

We use the first bound in (12) to show that $\widetilde{F}^\theta_t$ is likely and the second to show that the parameter is learned on $\widetilde{F}^\theta_t$.

4.3.1 The Event is Likely

To show that the events $\{\widetilde{F}^\theta_t\}_t$ are likely, we begin by showing that the events $B_{\ell t}(\log t, \alpha, \mathfrak{c}, \varepsilon', \theta)$ occur with high probability under the parameter θ for arbitrary values of the parameters ε', α, and $\mathfrak{c}$.

Lemma 3 *Given $\alpha > 0$, $\varepsilon' \in (0, 1)$, $\mathfrak{c} \in \mathbb{N}$ and Assumptions 2–4,*

$$P^\theta(B_{1t}(\log t, \alpha, \mathfrak{c}, \varepsilon', \theta)) \to 1, \qquad \text{as } t \to \infty.$$

Proof. We need to verify that with P^θ-probability one as t increases:

1. $\log t \geq \max\{\tau_1, t - \tau_N\}$,
2. $\left|\sum_{si} s(\hat{\zeta}^t_{si} - \zeta^\theta_{si})\right| < \alpha$, and
3. $\sum_{s\leq \mathfrak{c},i} s|\hat{\zeta}^t_{si} - \zeta^\theta_{si}| < \varepsilon'$.

Verification of 1: The probability that it takes more than $\mathfrak{b}$ periods for the first zero signal to arrive is at most $\lambda^{\mathfrak{b}}/(1-\rho)$ (by (7)). Thus the probability that the first condition holds is at least $1 - 2\lambda^{\log t}/(1-\rho)$. This tends to one as t becomes arbitrarily large.

Verification of 2: First we show that $N \to \infty$ (the number of the complete blocks tends to infinity) with probability one as $t \to \infty$. The probability of no zero signals in $\sqrt{t}$ periods is bounded above by $\lambda^{\sqrt{t}}/(1-\rho)$ (from (7)). The probability that over t periods divided into blocks of $\sqrt{t}$ periods there is at least one zero in each block is at least $1 - \sqrt{t}\lambda^{\sqrt{t}}/(1-\rho)$. This tends to one as $t \to \infty$. Thus the number of blocks tends to infinity as $t \to \infty$ with P^θ-probability one.

The length of each block is independent and identically distributed under θ (by the strong Markov property). We have shown (in Lemma 2) that its distribution has a finite mean and variance. By the Weak Law of Large Numbers, therefore, the probability that the average block length $\sum_{si} s\hat{\zeta}^t_{si}$ is more than α away from the expected block length $\sum_{si} s\zeta_{si}^{\theta}$ tends to zero as the number of blocks increases to infinity (a P^{θ}-probability one event).

Verification of 3: There are finitely many block signals i for each block length s. Let n_s denote the number of such signals. We have restricted attention to $s \leq \mathfrak{c}$, so it suffices to prove

$$\left|\hat{\zeta}_{si}^{t} - \zeta_{si}^{\theta}\right| < \frac{\varepsilon'}{\mathfrak{c} s n_s}, \qquad \forall i, s \leq \mathfrak{c}.$$

The Weak Law of Large Numbers applies to the random variable that indicates whether the block b_{si} occurred. Thus for any given si, the probability that $\left|\hat{\zeta}_{si}^{t} - \zeta_{si}^{\theta}\right| < \varepsilon'/\mathfrak{c} s n_s$ tends to one. This then applies to all si. ∎

We now argue that $\widetilde{F}_t^{\theta}$ is likely under θ, for sufficiently large t. Note that we can choose t sufficiently large for the event $\{N \geq T_{\varepsilon}\}$ to have probability arbitrarily close to one (see the proof of Lemma 3). Hence, from Lemma 3, we have,

$$\lim_{t\to\infty} P^{\theta}\left(B_{1t}(\log t, \alpha, \mathfrak{c}, \underline{\delta}, \theta) \cap B_{2t}(\log t, \alpha, \mathfrak{c}, \underline{\delta}, \theta) \cap \{N \geq T_{\varepsilon}\}\right) = 1. \tag{13}$$

Combining with (10)–(12) and (13), we have

$$\lim_{t\to\infty} P^{\theta}(\widetilde{F}_t^{\theta}) = 1.$$

4.3.2 The Parameter is Learned

Our next task is to show that θ is q-believed on $\widetilde{F}_t^{\theta}$. From (12), it suffices to show that the agents learn the parameter θ on the event $B_{\ell t}(\mathfrak{b}, \alpha, \mathfrak{c}, \varepsilon, \theta)$. Assumption 3 is sufficient to ensure that observed block frequencies identify the parameter. We let $[\zeta_{si}^{\theta}]$ denote player 1's distribution of block-signals, with $[\zeta_{sj}^{\theta}]$ denoting player 2's distribution.

Lemma 4 *Given Assumptions 1 and 3, there exists $\beta > 0$ such that, for $\theta \neq \tilde{\theta} \in \{\theta', \theta''\}$,*

$$\beta < H\left([\zeta_{si}^{\tilde{\theta}}] \middle\| [\zeta_{si}^{\theta}]\right), H\left([\zeta_{sj}^{\tilde{\theta}}] \middle\| [\zeta_{sj}^{\theta}]\right).$$

Proof. Since $\sum_i \zeta_{ti}^{\theta}$ is the probability that a completed block has length t, the distribution of arrival times of the zero signal is determined by ζ^{θ}. Similarly, $P^{\theta}(z_{\ell t} z_{\ell,t+1} = z_{\ell} z'_{\ell} \mid \bar{x})$, the probability that the pair of signals $z_{\ell} z'_{\ell}$ is observed in period t, conditional on the hidden Markov process starting in period -1 in state $\overline{x}$,[8] is determined by ζ^{θ}. From Assumption 1 and (2), we have that $\lim_{t\to\infty} P^{\theta}(z_{\ell t} z_{\ell,t+1} = z_{\ell} z'_{\ell} \mid \bar{x}) = P^{\theta}(z_{\ell} z'_{\ell})$.

Suppose the statement of the lemma is false, and $H([\zeta_{si}^{\tilde{\theta}}] \| [\zeta_{si}^{\theta}]) = 0$ (an identical argument applies if it fails for agent 2). Then, $[\zeta_{si}^{\theta'}] = [\zeta_{si}^{\theta''}]$ for all si, and so

[8]Recall that the first signal of a completed block of signals is the realization following a period in which after the 0 signal, and so state $\overline{x}$, is observed.

$P^{\theta'}(z_{\ell t} z_{\ell,t+1} = z_\ell z'_\ell \mid \bar{x}) = P^{\theta''}(z_{\ell t} z_{\ell,t+1} = z_\ell z'_\ell \mid \bar{x})$, for all t, and all pairs $z_\ell z'_\ell$. But Assumption 3 implies that there is at least one pair of signals for which

$$P^{\theta'}(z_\ell z'_\ell) \neq P^{\theta''}(z_\ell z'_\ell),$$

a contradiction. ■

In the next lemma we show that learning occurs on $B_{\ell t}(\log t, \alpha, \mathfrak{c}, \varepsilon, \theta)$. While learning on the intersection of $S_{\ell t}(\mathfrak{c}, \varepsilon, \theta)$ and the event that all blocks are of length $\mathfrak{c}$ or less is straightforward, arbitrarily long blocks may preclude learning. However, $B_{\ell t}$ also requires that the average block length be approximately correct, and so for $\mathfrak{c}$ sufficiently large, the arbitrarily long blocks are sufficiently infrequent that the learning cannot be overturned. In making this argument another feature of the block structure is important: a block's effect on learning is proportional to its length. Thus we can bound the informativeness of the long blocks by controlling their average length.

Lemma 5 *If Assumptions 2–4 hold, then there exists α, $\varepsilon \in (0,1)$, $\mathfrak{c} \in \mathbb{N}$ and a sequence $\xi : \mathbb{N} \to [0,1]$ with $\lim_{t\to\infty} \xi(t) = 1$ such that*

$$p^\theta_{\ell t} = P(\theta \mid h_{\ell t}) \geq \xi(t),$$

for all θ and all $h^t_\ell \in B_{\ell t}(\log t, \alpha, \mathfrak{c}, \varepsilon, \theta)$.

Proof. We prove for $\ell = 1$. Choose ε and α sufficiently small and $\mathfrak{c}$ sufficiently large so that

$$-\frac{\beta}{2} < \left(\alpha + 2\mathfrak{c}\varepsilon + \frac{2\lambda^{\mathfrak{c}+1}}{(1-\rho)}\left(\frac{1+\mathfrak{c}(1-\lambda)}{(1-\lambda)^2}\right)\right)\log\nu, \tag{14}$$

where β is given by Lemma 4, ρ is defined in (6), λ in (7), and $\nu > 0$ is a lower bound on all positive "observable" transition probabilities, that is,

$$\nu = \min_{\theta, x, x', z_1} \left\{\pi^\theta_{xx'}\phi^{x'\theta}_{z_1} : \pi^\theta_{xx'}\phi^{x'\theta}_{z_1} > 0\right\}.$$

We now bound the probability of the initial and terminal blocks $\bar{z}^o_1$ and $\bar{z}^e_1$. On histories in $B_{1t}(\mathfrak{b}, \alpha, \mathfrak{c}, \varepsilon, \theta)$, the initial and terminal block last less than $\mathfrak{b}$ periods. By Assumption 2, the stochastic processes under θ' and θ'' have common support and this support is finite when restricted to the first $\mathfrak{b}$ periods. Hence,

$$\begin{aligned}
P^\theta(\bar{z}^o_1) &= \sum_{(x_0, x_1, \ldots, x_{\tau_1})} \prod_{m=1}^{\tau_1} \phi^{x_m\theta}_{z_{1m}} P^\theta(x_m \mid x_{m-1}) P^\theta(x_0)\phi^{x_0\theta}_{z_{10}} \\
&\geq \min_{(x_0, x_1, \ldots, x_{\tau_1})} \prod_{m=1}^{\tau_1} \phi^{x_m\theta}_{z_{1m}} \pi^\theta_{x_{m-1}x_m} P^\theta(x_0)\phi^{x_0\theta}_{z_{10}} \\
&\geq \phi^{x_0\theta}_{z_{10}} P^\theta(x_0) \left(\min_{\theta, xx', z_1} \phi^{x'\theta}_{z_1}\pi^\theta_{xx'}\right)^{\mathfrak{b}}
\end{aligned}$$

$$= \phi_{z_{10}}^{x_0\theta} P^\theta(x_0)\nu^{\mathfrak{b}}.$$

Hence, there is a lower bound on the probabilities of all positive probability outcomes in the first $\mathfrak{b} = \log t$ periods. Letting $\phi_{z_{10}}^{x_0\theta} P^\theta(x_0) = K$, we have $P^{\theta'}(\bar{z}_1^o) > K\nu^{\log t}$. We similarly have that $P^{\theta'}(\bar{z}_1^e) > K'\nu^{\log t}$, for a different constant K'. A substitution into (9) then gives

$$\log \frac{p_{1t}^\theta}{1 - p_{1t}^\theta} \geq \log \frac{p_0^\theta}{1 - p_0^\theta} + 2\log t \log \nu + \log KK' + N \sum\nolimits_{si} \hat{\zeta}_{si}^t \log \frac{\zeta_{si}^\theta}{\zeta_{si}^{\tilde{\theta}}}. \tag{15}$$

We now argue that we can approximate the summation on the right side by $\sum_{si} \zeta_{si}^\theta \log(\zeta_{si}^\theta / \zeta_{si}^{\tilde{\theta}}) > \beta$, and hence show the log likelihood grows linearly in N. A similar calculation to the one bounding $P^\theta(\bar{z}_1^o)$ above implies that $\zeta_{si}^\theta / \zeta_{si}^{\tilde{\theta}} \leq \nu^{-s}$. Hence,

$$\begin{aligned}\left| \sum_{si} \hat{\zeta}_{si}^t \log \frac{\zeta_{si}^\theta}{\zeta_{si}^{\tilde{\theta}}} - \sum_{si} \zeta_{si}^\theta \log \frac{\zeta_{si}^\theta}{\zeta_{si}^{\tilde{\theta}}} \right| &\leq \sum_{si} \left| \hat{\zeta}_{si}^t - \zeta_{si}^\theta \right| \left| \log \frac{\zeta_{si}^\theta}{\zeta_{si}^{\tilde{\theta}}} \right| \\ &\leq \log(\nu^{-1}) \sum\nolimits_{si} s \left| \hat{\zeta}_{si}^t - \zeta_{si}^\theta \right|. \end{aligned} \tag{16}$$

On the set B_{1t}, the sum of these differences for $s \leq \mathfrak{c}$ of is bounded. So, on B_{1t},

$$\begin{aligned}\sum\nolimits_{si} s \left| \hat{\zeta}_{si}^t - \zeta_{si}^\theta \right| &\leq \mathfrak{c}\varepsilon + \sum\nolimits_{s>\mathfrak{c},i} s \left| \hat{\zeta}_{si}^t - \zeta_{si}^\theta \right| \\ &\leq \mathfrak{c}\varepsilon + \sum\nolimits_{s>\mathfrak{c},i} s\hat{\zeta}_{si}^t + \sum\nolimits_{s>\mathfrak{c},i} s\zeta_{si}^\theta. \end{aligned} \tag{17}$$

We now construct an upper bound for the right side of (17) that holds on the event B_{1t}. The public event that the mean lengths are close can be re-written as

$$\left| \sum\nolimits_{s\leq\mathfrak{c},i} s(\hat{\zeta}_{si}^t - \zeta_{si}^\theta) + \sum\nolimits_{s>\mathfrak{c},i} s(\hat{\zeta}_{si}^t - \zeta_{si}^\theta) \right| < \alpha.$$

On S_{1t}, the private event for agent 1, we have $\left| \sum_{s\leq\mathfrak{c},i} s(\hat{\zeta}_{si}^t - \zeta_{si}^\theta) \right| < \mathfrak{c}\varepsilon$. Combining these two inequalities,

$$\left| \sum\nolimits_{s>\mathfrak{c},i} s(\hat{\zeta}_{si}^t - \zeta_{si}^\theta) \right| < \alpha + \mathfrak{c}\varepsilon.$$

Hence $\sum_{s>\mathfrak{c},i} s\hat{\zeta}_{si}^t \leq \sum_{s>\mathfrak{c},i} s\zeta_{si}^\theta + \alpha + \mathfrak{c}\varepsilon$. Substituting into (17), we get

$$\sum\nolimits_{si} s \left| \hat{\zeta}_{si}^t - \zeta_{si}^\theta \right| \leq \alpha + 2\mathfrak{c}\varepsilon + 2\sum\nolimits_{s>\mathfrak{c},i} s\zeta_{si}^\theta.$$

Using (7) for the third inequality, we have

$$\sum_{s>\mathfrak{c},i} s\zeta_{si}^\theta \;\leq\; \sum_{s>\mathfrak{c}} sP^\theta(\sigma = s) \;\leq\; \sum_{s>\mathfrak{c}} sP^\theta(\sigma \geq s)$$

$$\le (1-\rho)^{-1} \sum_{s>\mathfrak{c}} s\lambda^s,$$

$$= \frac{\lambda^{\mathfrak{c}+1}}{(1-\rho)} \left(\frac{1+\mathfrak{c}(1-\lambda)}{(1-\lambda)^2} \right).$$

This allows us to rewrite the bound in (17) on S_{1t} as

$$\sum_{si} s \left| \hat{\zeta}^t_{si} - \zeta^\theta_{si} \right| \le \alpha + 2\mathfrak{c}\varepsilon + \frac{2\lambda^{\mathfrak{c}+1}}{(1-\rho)} \left(\frac{1+\mathfrak{c}(1-\lambda)}{(1-\lambda)^2} \right).$$

From (16), Lemma 4, and (14), we then have

$$\sum_{si} \hat{\zeta}^t_{si} \log \frac{\zeta^\theta_{si}}{\zeta^{\tilde{\theta}}_{si}} \ge \sum_{si} \zeta^\theta_{si} \log \frac{\zeta^\theta_{si}}{\zeta^{\tilde{\theta}}_{si}} + \left(\alpha + 2\varepsilon + \frac{2\lambda^{\mathfrak{c}+1}}{(1-\rho)} \left(\frac{1+\mathfrak{c}(1-\lambda)}{(1-\lambda)^2} \right) \right) \log \nu$$

$$\ge \beta + \left(\alpha + 2\varepsilon + \frac{2\lambda^{\mathfrak{c}+1}}{(1-\rho)} \left(\frac{1+\mathfrak{c}(1-\lambda)}{(1-\lambda)^2} \right) \right) \log \nu \ge \beta/2.$$

A final substitution into (15) then gives

$$\log \frac{p^\theta_{1t}}{1-p^\theta_{1t}} \ge \log \frac{p^\theta_0}{1-p^\theta_0} + 2 \log t \log \nu + \log KK' + N\beta/2.$$

It only remains to show that N (the number of completed blocks) increases linearly in t on B_{1t}. (This swamps any effect of the $\log t$ in the other term.) But (on B_{1t}) the total length of the completed blocks is at least $t - 2\mathfrak{b}$ and the average block length is $\sum_{si} s\hat{\zeta}^t_{si} \le \sum_{si} s\zeta^{\theta'}_{si} + \alpha$. Hence, on B_{1t},

$$N \ge \frac{t - 2\mathfrak{b}}{\sum_{si} s\hat{\zeta}^t_{si}} \ge \frac{t - 2\log t}{\sum_{si} s\zeta^\theta_{si} + \alpha},$$

completing the proof of the lemma. ∎

4.3.3 The Event is q-Evident

To show that $\widetilde{F}^\theta_t$ is q-evident, it is sufficient to show that

$$P(\{\theta\} \cap \widetilde{G}^\theta_{2t}(\varepsilon) \mid h_{1t}) > q,$$

for all h_{1t} in the event $I_t(\log t) \cap M_t(\alpha, \theta) \cap \{N \ge T_\varepsilon\} \cap \widetilde{G}^\theta_{1t}(\varepsilon)$. But when $N > T_\varepsilon$, this is ensured by (P3), since the only aspects of blocks longer than $\mathfrak{c}$ relevant to $\widetilde{F}^\theta_t$ are their publicly known length and number.

5 Separation: Frequency-Based Conditions for Common Learning

5.1 Learning from Intertemporal Patterns

Agents draw inferences about the value of the parameter from the frequencies of the signals they observe, and from the intertemporal pattern of signals across periods (see footnote 4). As part of their inference procedure, agents will make inferences about the history of hidden states that has generated their history of signals. However, the problem of calculating the posterior probabilities of hidden-state histories is notoriously difficult (Ephraim and Merhav, 2002, p.1573). Moreover, we are interested in common learning, which also requires each agent to infer the signal history of the other agent. Even in the simplest setting of temporally independent signals, there are signal histories on which an agent learns and yet does not believe that the other agent learns. Hence, common learning occurs on a subset of the histories on which individual learning occurs. The trick is to identify the "right" tractable subset. Tractability forces us to focus on events defined by signal frequencies alone. On such events, for common learning, agents need only infer frequencies of the other agent's signals, *not* their temporal structure.

Suppose the hidden Markov process is ergodic, and denote by $\psi_\ell^\theta := \sum_x \xi_x^\theta \phi_\ell^{x\theta}$ the ergodic distribution over agent ℓ's signals. Following our analysis of resets and Cripps, Ely, Mailath, and Samuelson (2008), the natural events to use to prove common learning are neighborhoods of ψ_ℓ^θ. The difficulty with using such events is that individual learning effectively requires an agent to make inferences about the evolution of the hidden states, and the relationship between these states and the signals.

Accordingly, we first investigate the possibility of common learning on large events. Recall that, conditional on the hidden state x and parameter θ, $\phi_\ell^{x\theta}$ is the distribution over agent ℓ's signals. The convex hull of these distributions under each parameter are denoted by

$$\Phi_\ell^{\theta'} = \mathrm{co}\{\phi_\ell^{x\theta'} : x \in X\} \quad \text{and} \quad \Phi_\ell^{\theta''} = \mathrm{co}\{\phi_\ell^{x\theta''} : x \in X\}. \tag{18}$$

The events we study are neighborhoods of $\Phi_\ell^{\theta'}$ and $\Phi_\ell^{\theta''}$. Individual learning on these events is not guaranteed. For example, in Figure 2, if the data $\hat{\phi}_\ell^t$ were observed, the agent would not infer θ', since $\hat{\phi}_\ell^t$ is more likely to be have been generated by a hidden state history close to the ergodic distribution from θ''. In order to obtain individual learning on such crude events, the two events must be significantly separated (Assumption 7). When the sets are significantly separated, the parameter is identified on the relevant convex hull, and common learning is almost immediate. As a byproduct, we obtain learning even when the hidden Markov process is not irreducible, since the agents are able to learn without needing to make inferences about the hidden states.

We then turn to learning on the neighborhoods of the ergodic distribution of signals. This makes it easier to verify that agents learn (and Assumption 7 implies Assumption 8), but complicates the verification that the events are q-evident.

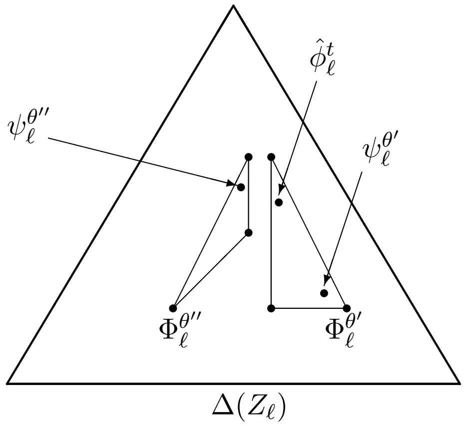

Figure 2: The convex hulls of the signal distributions generated by the various states of the hidden Markov process under parameters θ and θ', $\Phi^{\theta'}_\ell$ and $\Phi^{\theta''}_\ell$, are disjoint, but individual learning is not ensured. In particular, individual learning of θ' does not occur at the empirical frequency $\hat{\phi}^t_\ell$, since it is more likely to have occurred under θ'' and a hidden state distribution close to the ergodic distribution $\xi^{\theta''}$ than under θ' and a hidden state distribution far from the ergodic distribution $\xi^{\theta'}$.

5.2 Convex Hulls

This section establishes that if the distributions of signals for different states are sufficiently close together for a given parameter, and these sets of distributions are sufficiently far apart for different parameters, then there is common learning. We refer to this as a "relative separation" condition. This condition is intuitive and relatively straightforward to check, but it is demanding (since it must deal with the issue illustrated in Figure 2). Section 5.3 presents sufficient conditions for common learning that are less demanding but more cumbersome.

5.2.1 Common Learning on Convex Hulls

Our first assumption is that the signals agents observe under $P^{\theta'}$ and $P^{\theta''}$ have full support.

Assumption 5 (Full Support Signals) $\phi^{x\theta}_{z_\ell} > 0$*, for all* ℓ*,* θ*,* z_ℓ *and* x*.*

We also assume that, conditional on each hidden state, the agents' private signals are independent.

Assumption 6 (Conditional Independence) $\phi^{x\theta}_{z_1 z_2} = \phi^{x\theta}_{z_1}\phi^{x\theta}_{z_2}$*, for all* (z_1, z_2)*,* x*, and* θ*.*

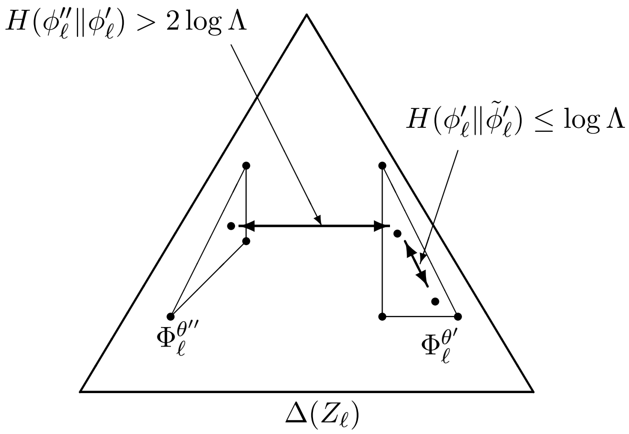

Figure 3: Illustration of Assumption 7. The distance between any two points within $\Phi_\ell^{\theta'}$ (and similarly for $\Phi_\ell^{\theta''}$) is less than Λ. Relative separation requires the distance between any point in $\Phi_\ell^{\theta'}$ and any point in $\Phi_\ell^{\theta''}$ to be greater than 2Λ, using relative entropy to measure distance.

Define a parameter Λ that bounds the diversity of the probabilities $(\phi_\ell^{x\theta})_{x\in X} \subset \Delta(Z_\ell)$:

$$\Lambda := \inf\{\Lambda' \geq 1 : \Lambda' \phi_\ell^{x\theta} \geq \phi_\ell^{x'\theta}, \quad \forall x, x', \theta, \ell\}. \tag{19}$$

That is, the factor Λ increases the probabilities $\phi_\ell^{x\theta}$ enough to make them greater than $\phi_\ell^{x'\theta}$ for any other hidden state x'. The factor Λ is well defined, since the supports of the signal distributions are the same (Assumption 5).

Assumption 7 (Relative Separation) *For Λ given by (19), and for $\ell = 1, 2$,*

$$\min\left\{ \max_{\phi''\in\Phi_\ell^{\theta''}} \min_{\phi'\in\Phi_\ell^{\theta'}} H(\phi'\|\phi''), \quad \max_{\phi'\in\Phi_\ell^{\theta'}} \min_{\phi''\in\Phi_\ell^{\theta''}} H(\phi''\|\phi') \right\} > 2\log\Lambda. \tag{20}$$

Though relative entropy is not a metric, the term on the left can be interpreted as a Hausdorff-like measure of the distance between the two sets $\Phi_\ell^{\theta'}, \Phi_\ell^{\theta''} \subset \Delta(Z_\ell)$. The construction of Λ ensures $H(\phi\|\tilde{\phi}) \leq \log\Lambda$ for all $\phi, \tilde{\phi} \in \Phi_\ell^{\theta'}$ and all $\phi, \tilde{\phi} \in \Phi_\ell^{\theta''}$. So (20) can be interpreted as requiring the distance between $\Phi_\ell^{\theta'}$ and $\Phi_\ell^{\theta''}$ be more than twice the distance across them (see Figure 3).[9]

Relative separation, with full support and conditional independence are sufficient for common learning.

Proposition 4 *If Assumptions 5–7 hold, then θ is commonly learned.*

[9] If the sets $\Phi_\ell^{\theta''}, \Phi_\ell^{\theta'}$ had a non-empty intersection the minimizer in (20) will be a point in the intersection, and so the maximum is no larger than $\log\Lambda$.

5.2.2 Proof of Proposition 4: Common Learning

Let $n^t_{xz_\ell}$ denote the number of periods $s < t$ in which $(x_s, z_{\ell s})$ occurs. We also let $n^t_x = \sum_{z_\ell} n^t_{xz_\ell}$ and $n^t_{z_\ell} = \sum_x n^t_{xz_\ell}$ denote the marginal frequencies of states and signals respectively. The time-t empirical measures of the agent's signals and the hidden states are then

$$\begin{aligned} \hat{\phi}^t_\ell &\in \Delta(Z_\ell), \qquad \hat{\phi}^t_{z_\ell} := n^t_{z_\ell}/t; \\ \text{and } \hat{\xi}^t &\in \Delta(X), \qquad \hat{\xi}^t_x := n^t_x/t. \end{aligned}$$

We are interested in the case that the empirical measure of the private signals observed by each agent are close to the convex hulls $\Phi^{\theta'}_\ell$ and $\Phi^{\theta''}_\ell$ in (18). The event we show is common q-belief in state θ for small $\varepsilon > 0$ is

$$F^\theta_t(\varepsilon) = \left\{ \omega \in \Omega \ : \hat{\phi}^t_1 \in \Phi^\theta_1(\varepsilon),\ \hat{\phi}^t_2 \in \Phi^\theta_2(\varepsilon) \ \right\}, \tag{21}$$

where $\Phi^\theta_\ell(\varepsilon) = \{\hat{\phi}_\ell \in \Delta(Z_\ell) : \left|\hat{\phi}_\ell - \phi_\ell\right| < \varepsilon \text{ for some } \phi_\ell \in \Phi^\theta_\ell\}$.

We again follow the agenda set out in Corollary 1 The first result is that the event $F^\theta_t(\varepsilon)$ has asymptotically P^θ-probability 1.

Lemma 6 *For all $\varepsilon > 0$, $P^\theta(F^\theta_t(\varepsilon)) \to 1$ as $t \to \infty$.*

Proof. For fixed θ and an arbitrary sequence of hidden states $x_0, x_1, x_3, \ldots$, $P^\theta(\hat{\phi}^t_\ell \in \Phi^\theta_\ell(\varepsilon) \mid x_0, x_1, \ldots) \to 1$ as $t \to \infty$. Taking expectations over the sequences of hidden states then yields the result. ∎

The next lemma (proved in Appendix A.2) verifies that the parameter θ is learned on $F^\theta_t(\varepsilon)$. The proof exploits Assumption 7 to show that signal frequencies close to Φ^θ_ℓ are much more likely to have arisen under parameter θ than under $\tilde{\theta}$ irrespective of the sequence of hidden states, and hence the posterior attached to θ by agent ℓ converges to one, i.e., agent ℓ learns parameter θ.

Lemma 7 *If Assumptions 5 and 7 hold, then there exists $\varepsilon' > 0$ such that for $\varepsilon \in (0, \varepsilon')$, for all $\eta > 0$ there exists T such that for all $t > T$ and all $h_{\ell t} \in F^\theta_t(\varepsilon)$,*

$$P(\theta \mid h_{\ell t}) > 1 - \eta.$$

Finally, the q-evidence of $F^\theta_t(\varepsilon)$ will follow almost immediately from individual learning (Lemma 7) and Lemma 6, since inferences about the hidden states play no role in determining whether the histories are in $F^\theta_t(\varepsilon)$.

Lemma 8 *If Assumptions 5–7 hold, then for any $q < 1$ there exists $\varepsilon' > 0$ such that for $\varepsilon \in (0, \varepsilon')$, there exists T such that for all $t > T$, the event $F^\theta_t(\varepsilon)$ is q-evident.*

Proof. For any h_{1t},

$$\begin{aligned} P(\hat{\phi}_2^t \in \Phi_2^\theta(\varepsilon) \mid h_{1t}) &\geq P^\theta(\hat{\phi}_2^t \in \Phi_2^\theta(\varepsilon) \mid h_{1t}) P(\theta \mid h_{1t}) \\ &= \sum_{x^t \in X^t} P^\theta(\hat{\phi}_2^t \in \Phi_2^\theta(\varepsilon) \mid x^t, h_{1t}) P^\theta(x^t \mid h_{1t}) P(\theta \mid h_{1t}) \\ &= \sum_{x^t \in X^t} P^\theta(\hat{\phi}_2^t \in \Phi_2^\theta(\varepsilon) \mid x^t) P^\theta(x^t \mid h_{1t}) P(\theta \mid h_{1t}), \end{aligned}$$

where the last equality follows from Assumption 6.

Fix $q \in (0,1)$. There exists a time T' such that for all $t \geq T'$, and all $x^t \in X^t$, $P^\theta(\hat{\phi}_2^t \in \Phi_2^\theta(\varepsilon) \mid x^t) > \sqrt{q}$.[10] From Lemma 7, there exists T'' such that for all $t \geq T''$, $P(\theta \mid h_{1t}) > \sqrt{q}$ for $h_{1t} \in F_t^\theta(\varepsilon)$. Combining these two inequalities, we conclude that for all $t \geq \max\{T', T''\}$, $P(\hat{\phi}_2^t \in \Phi_2^\theta(\varepsilon) \mid h_{1t}) > q$ for all $h_{1t} \in F_t^\theta(\varepsilon)$, and so (since the same argument holds for agent 2), $F_t^\theta(\varepsilon)$ is q-evident. ■

5.3 Average Distributions

The relative separation condition of Assumption 7 is strong, requiring that every possible signal distribution under θ' (i.e., generated by an *arbitrary* distribution over hidden states) is far away from every possible signal distribution under θ''. This section begins with a weaker separation condition (Assumption 8) that places restrictions on neighborhoods of the average signal distributions (rather than their convex hulls) under the two parameters, requiring them to differ by an amount that is related to the differing distributions of the hidden Markov process induced by the two parameters. This weakened separation condition comes at a cost, in that we must then introduce an addition assumption (Assumption 9). This assumption requires that, for a given value of the parameter, likely explanations of anomalous signal realizations must not rely too heavily on unlikely realizations of the underlying hidden states.

We denote the ε-neighborhood of the stationary signal distribution by $\breve{\Phi}_\ell^\theta(\varepsilon) := \{\phi_\ell \in \Delta(Z_\ell) : \|\phi_\ell - \psi_\ell^\theta\| < \varepsilon\}$.[11] An advantage of working with an assumption on average signal frequencies rather than the convex hulls of signal distributions, is that we can use a sequence of smaller events (defined in terms of neighborhoods of average frequencies $\breve{\Phi}_\ell^\theta(\varepsilon)$, rather than neighborhoods of convex hulls, $\Phi_\ell^\theta(\varepsilon)$) when applying Corollary 1 to establish common learning. In particular, this makes it easier to verify that the agents learn. However, it is now harder to show that an agent

[10] Suppose not. Then, there is some q such that for all T', there is a $t \geq T'$ and x^t such that $P^\theta(\hat{\phi}_\ell^t \notin \Phi_\ell^\theta(\varepsilon) \mid x^t) > 1 - \sqrt{q}$. But, since X is finite, there is then a single state $x \in X$ such that the event that the frequency of signals in the periods in which x is realized is more than ε distant from $\phi_2^{x\theta}$ has P^θ-probability at least $1 - \sqrt{q}$, for arbitrarily large T'. But this is ruled out by the Weak Law of Large Numbers.

[11] In this section the norm used will always be the variation norm, defined by $\|\zeta\| := (1/2)\sum_w |\zeta_w|$.

expects the opponent's observations to lie in the target set, and so the argument for q-evidence is more involved. This latter argument relies on constructing bounds on probabilities using large deviations arguments.

We maintain the assumptions that the signal distributions have full support under each parameter value (Assumption 5) and that the distributions are conditionally (on the hidden state and parameter) independent (Assumption 6). We complement this with two assumptions. These assumptions are designed to ensure that if agent ℓ's signals are in $\breve{\Phi}^{\theta}_{\ell}(\varepsilon)$ for some small ε, then (1) the agent's posteriors p^{θ} to tend to one and (2) agent ℓ believes that ℓ''s signals are in $\breve{\Phi}^{\theta}_{\ell'}(\varepsilon)$. The first of these implications will ensure learning, and the second q-evidence, leading to common learning.

We begin by introducing a function

$$A^{\theta}(\tilde{\xi}) := \sup_{\substack{v\in\Delta(X)\\ v\gg 0}} \sum_{x'} \tilde{\xi}_{x'} \log \frac{v_{x'}}{\sum_{\tilde{x}} v_{\tilde{x}}\pi^{\theta}_{\tilde{x}x'}}. \tag{22}$$

The function A^{θ} is nonnegative, strictly convex, and $A^{\theta}(\xi^{\theta}) = 0$ uniquely.[12] Also, A^{θ} increases as $\tilde{\xi}$ moves away from ξ^{θ}. We can interpret A^{θ} as a measure of how far the distribution $\tilde{\xi}$ over hidden states is from the ergodic distribution ξ^{θ}. In particular, fixing $\tilde{\xi}$, we can ask how much less likely v is under distribution ξ than is $v^T\pi^{\theta}$, and $A^{\theta}(\xi)$ takes the maximal difference as the measure of the distance of $\tilde{\xi}$ from ξ. It is this function that will capture the role of hidden states in our conditions.

Next, given a distribution ξ on the hidden states and a signal distribution $\hat{\phi}^t_{\ell}$ from a history $h_{\ell t}$, we allocate the signals to states of the hidden Markov process, under the assumption that the hidden states themselves appear in proportions ξ. Any such allocation can be interpreted as an explanation of the observed signal frequency (data) in terms of the underlying hidden state realizations. An allocation determines a collection of conditional distributions $(\hat{\phi}^x_{\ell})_{x\in X}$, where $\hat{\phi}^x_{\ell} := (n^t_{xz_{\ell}}/n^t_x)_{z_{\ell}} \in \Delta(Z_{\ell})$, $n^t_{xz_{\ell}}$ is the number of observations of the xz_{ℓ} pair in the allocation, and $n^t_x = \sum_{z_{\ell}} n^t_{xz_{\ell}}$. The set of all possible such allocations, a convex linear polytope in the space $\Delta(Z)^{|X|}$, is the set of possible explanations of the data. For arbitrary $\phi_{\ell} \in \Delta(Z_{\ell})$ and $\xi \in \Delta(X)$, the set of possible explanations is

$$J_{\ell}(\phi_{\ell},\xi) := \left\{(\phi^x_{\ell})_{x\in X} \in \Delta(Z)^{|X|} : \phi_{\ell} = \textstyle\sum_x \xi_x \phi^x_{\ell}\right\}. \tag{23}$$

Our first assumption is designed to ensure individual learning on a neighborhood $\breve{\Phi}^{\theta}_{\ell}(\varepsilon)$ of the stationary distribution ψ^{θ}_{ℓ}.

[12]To get some intuition, observe that $A^{\theta}(\xi) = \sup_v H(\xi\|v^T\pi^{\theta}) - H(\xi\|v)$. Choosing $v = \xi$ ensures the second term is zero so $A^{\theta} \geq 0$. If ξ is a stationary measure for π^{θ}, then $v^T\pi^{\theta}$ is closer to ξ than v is, so A^{θ} cannot be strictly positive. For more on this function, see den Hollander (2000, Theorem IV.7, p.45).

Assumption 8 *For $\theta \neq \tilde{\theta}$, for $\ell = 1, 2$, there exists $\bar{\varepsilon} > 0$ such that for all $\phi_\ell \in \check{\Phi}^{\theta}_{\ell}(\bar{\varepsilon})$,*

$$-\min_{\substack{(\hat{\phi}^x_\ell)_x \in \\ J_\ell(\phi_\ell, \xi^\theta)}} \sum_{x, z_\ell} \xi^\theta_x \hat{\phi}^x_{z_\ell} \log \phi^{x\theta}_{z_\ell} < \min_{\xi \in \Delta(X)} \left\{ A^{\tilde{\theta}}(\xi) - \max_{\substack{(\hat{\phi}^x_\ell)_x \in \\ J_\ell(\phi_\ell, \xi)}} \sum_{x, z_\ell} \xi_x \hat{\phi}^x_{z_\ell} \log \phi^{x\tilde{\theta}}_{z_\ell} \right\}. \quad (24)$$

This condition is implied by the Assumption 7: In particular, since $A^\theta \geq 0$, a sufficient condition for (24) is that, for all $\xi \in \Delta(X)$,

$$\min_{(\hat{\phi}^x_\ell)_x \in J_\ell(\psi^\theta_\ell, \xi^\theta)} \sum_{x, z_\ell} \xi^\theta_x \hat{\phi}^x_{z_\ell} \log \phi^{x\theta}_{z_\ell} > \max_{(\hat{\phi}^x_\ell)_x \in J_\ell(\psi^\theta_\ell, \xi)} \sum_{x, z_\ell} \xi_x \hat{\phi}^x_{z_\ell} \log \phi^{x\tilde{\theta}}_{z_\ell}.$$

It can only make the minimum smaller if each signal observed is matched to the hidden state for which it is least likely. Similarly, it can only make the maximum larger if each signal observed is matched to the hidden state for which it is most likely. Hence a sufficient condition for the above is $\sum_{z_\ell} \psi^\theta_{z_\ell} \log \underline{\phi}^\theta_{z_\ell} > \sum_{z_\ell} \psi^\theta_{z_\ell} \log \bar{\phi}^{\tilde{\theta}}_{z_\ell}$, where this notation is defined just before (A.4). However, this last inequality is implied by (A.4), which is ensured by Assumption 7.

Suppose we are given a parameter θ' and a collection of signals whose frequencies ϕ_ℓ match (up to ε) the expected signal distribution under θ' (that is, $\phi_\ell \in \check{\Phi}^{\theta'}_\ell(\varepsilon)$). In order for agent ℓ to learn θ', observing ϕ_ℓ under θ' should be much more likely than under θ'', a property implied by Assumption 8. The likelihood of ϕ_ℓ depends on how the signals are allocated to hidden states.

The expression on the left of (24) bounds (from below) the probability of observing ϕ_ℓ, under θ' and the most likely distribution of hidden states $\xi^{\theta'}$, where we construct the bound by asking: what is the least likely way of allocating the signals to the hidden states consistent with ϕ_ℓ?

The expression on the right of the inequality bounds (from above) the probability of ϕ_ℓ, under θ'', where we construct the bound by asking: what is the most likely way of allocating the signals to the hidden states consistent with ϕ_ℓ? Importantly, as illustrated by our discussion of Figure 2, for ϕ_ℓ far from $\psi^{\theta''}_\ell$, this allocation requires trading off the probability "costs" of

1. likely realizations of hidden states and unlikely realizations of the signals

against

2. unlikely realizations of hidden states and likely realizations of the signals.

Recall that the $A^{\theta''}$ function captures the cost of specifying the distribution of hidden states that is different from the stationary distribution (since the expression on the left of the inequality is calculated at the stationary distribution $\xi^{\theta'}$, the analogous term does not appear).

Our second condition ensures q-evidence of the event we study below.

Assumption 9 *For $\ell = 1, 2$ and all θ and some $\varepsilon^\dagger \in (0, \bar{\varepsilon})$ and all ϕ_ℓ such that $\|\phi_\ell - \psi_\ell^\theta\| < \varepsilon^\dagger$,*

$$-\min_{J_\ell(\phi_\ell,\xi^\theta)} \sum_{x,z_\ell} \xi_x^\theta \hat{\phi}_{z_\ell}^x \log \phi_{z_\ell}^{x\theta} < \min_{\{\|\xi-\xi^\theta\|\geq f(\Lambda)\varepsilon^\dagger\}} \left\{ A^\theta(\xi) - \max_{J_\ell(\phi_\ell,\xi)} \sum_{x,z_\ell} \xi_x \hat{\phi}_{z_\ell}^x \log \phi_{z_\ell}^{x\theta} \right\}. \tag{25}$$

where $f(\Lambda) := (2/\log \Lambda)^{1/2}$ and Λ is defined in (19).

In order to demonstrate q-evidence of an event of the form $\breve{\Phi}_1^{\theta'}(\varepsilon^\dagger) \cap \breve{\Phi}_2^{\theta'}(\varepsilon^\dagger)$, we need to show that agent 1 is confident that 2's private signal frequencies are in $\breve{\Phi}_2^{\theta'}(\varepsilon^\dagger)$ when $\phi_1 \in \breve{\Phi}_1^{\theta'}(\varepsilon^\dagger)$. By Assumption 6, 1's inferences about 2's private signals are determined by 1's inferences about the hidden states. This explains why Assumption 9 can imply q-evidence of the relevant event even though it only involves characteristics of agent 1. In particular, since 1 learns θ' on $\breve{\Phi}_1^{\theta'}(\varepsilon^\dagger)$, if 1 is sufficiently confident that the hidden state distribution is close to its stationary distribution under θ', 1 will be confident that 2's private signal frequencies are in $\breve{\Phi}_2^{\theta'}(\varepsilon^\dagger)$.

Assumption 9 essentially requires, given θ', the private signal frequency ϕ_ℓ is more likely to have been realized from the stationary distribution of states $\xi^{\theta'}$ (the left side of (25)) than from some state distribution *not* in some neighborhood of $\xi^{\theta'}$ (the right side of (25)). Since the probability trade-offs faced by agent 1 mimic those described above, the form of (25) is very close to that of (24). In particular, deviations from the ergodic distribution are penalized at rate A^θ (the right side). Notice that Assumption 8 compares various explanations of the data under different values of the parameter, while Assumption 9 is comparing explanations based on the same value of the parameter. The parameter $\Lambda \geq 1$ measures the dissimilarity of the signal distributions for different states under the parameter θ (with $\Lambda = 1$ if the signal distributions are identical under the different states). The factor $f(\Lambda) > 0$ is a decreasing function of Λ with $\lim_{\Lambda\to 1} f(\Lambda) = \infty$. As one would expect, this constraint becomes weaker as the signal distributions in each state become more similar.

Proposition 5 *Common learning holds under assumptions 5, 6, 8, and 9.*

The proof again takes us through the agenda set out in Corollary 1. The event we will show to be common p-belief is the event that the empirical measure of the private signals observed by each agent are close to their expected values under the parameter.

Define

$$\breve{F}_{\varepsilon t}^\theta := \left\{ \omega \in \Omega \ : \hat{\phi}_1^t \in \breve{\Phi}_1^\theta(\varepsilon),\ \hat{\phi}_2^t \in \breve{\Phi}_2^\theta(\varepsilon) \right\}. \tag{26}$$

The event that we show is common p-belief for parameter θ is $\breve{F}_{\varepsilon^\dagger t}^\theta$ where $\varepsilon^\dagger > 0$ is from Assumption 9.

We first show that the event occurs with sufficiently high probability.

Lemma 9 *For all $\varepsilon > 0$, $P^\theta(\breve{F}_{\varepsilon t}^\theta) \to 1$ as $t \to \infty$.*

The (omitted) proof is a straightforward application of the ergodic theorem (Brémaud, 1999, p. 111).

The next step is that the parameter is individually learned on the event $\breve{F}^{\theta}_{\varepsilon^{\dagger}t}$.

Lemma 10 *Suppose Assumptions 5, 6, and 8 hold. For all $q \in (0,1)$, $\varepsilon < \bar{\varepsilon}$, θ, and ℓ, there exists T such that for all $t \geq T$, $P(\theta \mid h_{\ell t}) > q$ for all $h_{\ell t} \in \breve{F}^{\theta}_{\varepsilon t}$.*

Finally, we show that if agent 1's signals are in $\breve{F}^{\theta}_{\varepsilon^{\dagger}t}$, then she attaches arbitrarily high probability to agent 2's signals being in $\breve{F}^{\theta}_{\varepsilon^{\dagger}t}$ — that is, q-evidence. This proof proceeds in two steps. First we show that if agent 1 believes agent 2's signals are not in $\breve{F}^{\theta}_{\varepsilon^{\dagger}t}$ then she must also believe that the hidden distribution $\hat{\xi}$ is a long way from its stationary distribution under θ, because if it were close to ξ^{θ}, the independence of signals alone would ensure 2's signals were in $\breve{F}^{\theta}_{\varepsilon^{\dagger}t}$. The second step is to use our earlier bounds to characterize the probability agent 1 believes $\hat{\xi}$ is far from ξ^{θ} when she has seen a history h_{1t} consistent with the event $\breve{F}^{\theta}_{\varepsilon^{\dagger}t}$.

Lemma 11 *Suppose Assumptions 5, 6, 8, and 9 hold. For all $q \in (0,1)$ and θ, the set $\breve{F}^{\theta}_{\varepsilon^{\dagger}t}$ is q-evident under the parameter θ for t sufficiently large.*

A Appendix: Proofs

A.1 A Full Support Example with No Common Learning

Suppose the hidden Markov process π is described by the state transitions in Figure 1. The private signal distribution in state $x \in \{x^0, x^1, x^2, x^3\}$ is given by (1) and Figure 1 with probability $1-9\varepsilon$, and by a uniform draw from $\{aa, ab, ac, ba, bb, bc, ca, cb, cc\}$ with probability 9ε.

Let $\tilde{\tau}$ be the first date at which the process is not in state x^0. The following lemma implies that there exists $\eta > 0$ such that at any time t and conditional on $\tilde{\tau} > \tau$ for any $\tau < t$, there is probability at least η that agent ℓ observes a history $h_{\ell t}$ such that $P^{\theta}(\tilde{\tau} > \tau + 1 \mid h_{2t}) > \eta$ (take $\eta = \min\{\eta_1, \eta_2\}$). We can iterate this argument to obtain iterated η-belief at time t that the process is still in state x^0, precluding common learning.

Lemma A.1 *For $\varepsilon > 0$ sufficiently small, there exists $\eta_1, \eta_2 > 0$ such that for all times τ and $t > \tau$, and all ℓ,*

$$P^{\theta}\left(\left\{h_{\ell t} : P^{\theta}(\tilde{\tau} > \tau + 1 \mid h_{\ell t}) > \eta_1\right\} \middle| \, \tilde{\tau} > \tau\right) > \eta_2.$$

.

Proof. Fix τ. For any $t > \tau$, define the event

$$E_t := \left\{h_{\ell t} : \forall \tau' \leq \tau + 1, \; \#\{s : z_{\ell s} = a, \; \tau + 1 - \tau' < s \leq \tau + 1\} \geq \tfrac{2}{3}\tau'\right\}.$$

We first argue that for all $h_{\ell t} \in E_t$, $P^{\theta}(\tilde{\tau} > \tau + 1 \mid h_{\ell t})$ is bounded away from 0 independently of t and the particular history $h_{\ell t}$, giving η_1. Then we argue that, conditional on the hidden Markov process still being in the state x^0 at time τ (i.e., $\tilde{\tau} > \tau$), E_t has probability bounded away from 0 independently of t, giving η_2 and completing the proof.

Observe that for all $h_{\ell t} \in E_t$, $P^\theta(\tilde{\tau} > \tau + 1 \mid h_{\ell t})$ is bounded away from 0 independently of t and $h_{\ell t}$ if and only if there exists an upper bound independent of t and $h_{\ell t}$ for

$$\frac{1 - P^\theta(\tilde{\tau} > \tau + 1 \mid h_{\ell t})}{P^\theta(\tilde{\tau} > \tau + 1 \mid h_{\ell t})} = \frac{\sum_{s=1}^{\tau+1} P^\theta(\tilde{\tau} = s \mid h_{\ell t})}{P^\theta(\tilde{\tau} > \tau + 1 \mid h_{\ell t})}. \tag{A.1}$$

Fix t and a history $h_{\ell t} \in E_t$. For fixed s, we have

$$\frac{P^\theta(\tilde{\tau} = s \mid h_{\ell t})}{P^\theta(\tilde{\tau} > \tau + 1 \mid h_{\ell t})} = \frac{P^\theta(h_{\ell,(\tau+2,t)} \mid h_{\ell,\tau+2}, \tilde{\tau} = s) P^\theta(h_{\ell,\tau+2} \mid \tilde{\tau} = s) P^\theta(\tilde{\tau} = s)}{P^\theta(h_{\ell,(\tau+2,t)} \mid h_{\ell,\tau+2}, \tilde{\tau} > \tau + 1) P^\theta(h_{\ell,\tau+2} \mid \tilde{\tau} > \tau + 1) P^\theta(\tilde{\tau} > \tau + 1)}, \tag{A.2}$$

where $h_{\ell,(\tau+2,t)}$ is the history of signals observed by agent ℓ in periods $\{\tau+2, \ldots, t-1\}$.

Let n_a and n_z be the number of a and $z \in \{b, c\}$ signals observed in periods $\{s+1, \ldots, \tau+1\}$ of $h_{\ell t}$, respectively. Since $h_{\ell t} \in E_t$, we have $n_a \geq 2(\tau - s + 1)/3$, so that $n_a - n_z \geq (\tau - s + 1)/3$.

In periods before s, the hidden Markov process is in state x^0, and so the probabilities of the signals in those periods are identical on the events $\{\tilde{\tau} > \tau\}$ and $\{\tilde{\tau} = s\}$, allowing us to cancel the common probabilities in the first s periods. In period s, the hidden Markov process is either in state x^1 or in state x^2, and we bound the probability of the signal in the numerator in that period by 1, and use 3ε as the lower bound in the denominator. In periods after s and before $\tau + 2$, signal b in state x^3 has probability $\theta(1 - 9\varepsilon) + 3\varepsilon$, while signal c has probability $(1 - \theta)(1 - 9\varepsilon) + 3\varepsilon$. These two probabilities are bounded above by $1 - 6\varepsilon$, the probability of a in state x^0. Thus,

$$\begin{aligned}
\frac{P^\theta(h_{\ell,\tau+2} \mid \tilde{\tau} = s) P^\theta(\tilde{\tau} = s)}{P^\theta(h_{\ell,\tau+2} \mid \tilde{\tau} > \tau + 1) P^\theta(\tilde{\tau} > \tau + 1)} &< \frac{(3\varepsilon)^{n_a}(1 - 6\varepsilon)^{n_z}}{3\varepsilon(1 - 6\varepsilon)^{n_a}(3\varepsilon)^{n_z}} \frac{P^\theta(\tilde{\tau} = s)}{P^\theta(\tilde{\tau} > \tau + 1)} \\
&= \frac{(3\varepsilon)^{n_a}(1 - 6\varepsilon)^{n_z}}{3\varepsilon(1 - 6\varepsilon)^{n_a}(3\varepsilon)^{n_z}} \frac{(1 - 2\zeta)^{s-1} 2\zeta}{(1 - 2\zeta)^{\tau+1}} \\
&= \frac{2\zeta}{3\varepsilon(1 - 2\zeta)} \frac{(3\varepsilon)^{n_a - n_z}}{(1 - 6\varepsilon)^{n_a - n_z}} \frac{1}{(1 - 2\zeta)^{\tau - s + 1}}.
\end{aligned} \tag{A.3}$$

For $\varepsilon > 0$ sufficiently small,

$$\kappa := \frac{(3\varepsilon)^{\frac{1}{3}}}{(1 - 6\varepsilon)^{\frac{1}{3}}(1 - 2\zeta)} < 1.$$

and so the left side of (A.3) is bounded above by

$$\frac{2\zeta}{3\varepsilon(1 - 2\zeta)^2} \kappa^{\tau - s + 1}.$$

We then note that, for $s \leq \tau + 1$,

$$\frac{P^\theta(h_{\ell,(\tau+2,t)} \mid h_{\ell,\tau+2}, \tilde{\tau} = s)}{P^\theta(h_{\ell,(\tau+2,t)} \mid h_{\ell,\tau+2}, \tilde{\tau} > \tau + 1)} \leq \max_{t', h_{\ell,(\tau+2,t')}, x'} \frac{P^\theta(h_{\ell,(\tau+2,t')} \mid x_{\tau+2} = x^3)}{P^\theta(h_{\ell,(\tau+2,t')} \mid x_{\tau+2} = x')}$$

is bounded. Hence, the left sides of (A.2) and therefore (A.1)are bounded above by a geometric series, and so have an upper bound independent of t and $h_{\ell t}$.

Now we show that the probability of the event E_t, conditional on the hidden state being x^0 at time τ, is bounded away from zero. Given that we are conditioning on the state being x_0 at time τ, it is convenient to show that the probability of the event

$$\tilde{E}_t := \left\{h_{\ell t} : \forall \tau' \leq \tau,\ \#\{s : z_{\ell s} = a,\ \tau + 1 - \tau' < s \leq \tau\} \geq \tfrac{2}{3}\tau'\right\},$$

is bounded away from zero, and then to extend the result to E_t by noting that probability of an a signal in period $\tau+1$, conditional on being in state x^0 in period τ, is at least $(1-2\zeta)(1-9\varepsilon)$.

Conditional on being in state x^0 at time τ, the distribution of agent ℓ's signals is identical and independently distributed through time, and so $\tilde{E}_t$ has the same probability as the event

$$\widehat{E}_t := \left\{h_{\ell t} : \forall \tau' \leq \tau,\ \#\{s : z_{\ell s} = a,\ 0 \leq s < \tau'\} \geq \tfrac{2}{3}\tau'\right\}.$$

Moreover, $\widehat{E} \subset \widehat{E}_t$, where

$$\widehat{E} := \left\{\{z_{\ell,s}\}_{s=0}^\infty \in Z_\ell^\infty :\ \#\{s : z_{\ell s} = a,\ 0 \leq s < t\} \geq \tfrac{2}{3}t\right\}$$

is the collection of outcome paths of agent ℓ signals for which every history $h_{\ell t}$ has at least a fraction two thirds of a's. The proof is complete once we show that $\widehat{E}$ has strictly positive probability, conditional on $x_t = x^0$ for all t.

Let X_k be a random walk on the integers described by

$$X_{k+1} = \begin{cases} X_k + 1, & \text{with probability } p_1 = (1-6\varepsilon)^3, \\ X_k, & \text{with probability } p_2 = 3(1-6\varepsilon)^2 6\varepsilon, \\ X_k - 1, & \text{with probability } p_3 = 3(1-6\varepsilon)(6\varepsilon)^2, \text{ and} \\ X_k - 2, & \text{with probability } p_4 = (6\varepsilon)^3, \end{cases}$$

with initial condition $X_0 = 1$. The process $\{X_k\}$ tracks the fraction of a signals over successive triples of signal realizations at periods $3k$ as follows:

1. if the triple aaa is realized, $X_{k+1} = X_k + 1$,
2. if a single non-a is realized in the triple, $X_{k+1} = X_k$,
3. if two non-a's are realized in the triple, $X_{k+1} = X_k - 1$, and
4. if only non-a's are realized in the triple, $X_{k+1} = X_k - 2$.

An outcome that begins with the triple aaa and for which X_k is always a strictly positive integer is in $\widehat{E}$. Hence, it is enough to argue that the probability that $\{X_k\}$ is always strictly positive is strictly positive, when ε is small. This is most easily seen by considering the simpler random walk $\{Y_k\}$ given by

$$Y_{k+1} = \begin{cases} Y_k + 1, & \text{with probability } p_1, \\ Y_k - 2, & \text{with probability } 1 - p_1, \end{cases}$$

with initial condition $Y_0 = 1$. Clearly, $\Pr(X_k \geq 1, \forall k \mid X_0 = 1) \geq \Pr(Y_k \geq 1, \forall k \mid Y_0 = 1)$. Moreover, for $p_1 \neq \frac{2}{3}$, every integer is a transient state for $\{Y_k\}$. Finally, if $p_1 > \frac{2}{3}$ (which is guaranteed by ε small), $\Pr(Y_k \geq 1, \forall k \mid Y_0 = 1) > 0$. ■

A.2 Common Learning on Convex Hulls: Proof of Lemma 7

We need to show that the posteriors, $P(\theta \mid h_{\ell t})$, converge to one on $F_t^{\theta}(\varepsilon)$ (since posteriors are a martingale, almost sure convergence is immediate). It is sufficient to show that

$$\frac{P(\theta' \mid h_{\ell t})}{P(\theta'' \mid h_{\ell t})} \frac{p(\theta'')}{p(\theta')} = \frac{P^{\theta'}(h_{\ell t})}{P^{\theta''}(h_{\ell t})} \to \infty,$$

for all $h_{\ell t} \in F_t^{\theta'}(\varepsilon)$ as $t \to \infty$. Denoting the hidden state history by $x^t = (x_0, x_1, \ldots, x_{t-1}) \in X^t$, we have

$$\begin{aligned} \frac{P^{\theta'}(h_{\ell t})}{P^{\theta''}(h_{\ell t})} &= \frac{\sum_{x^t \in X^t} P^{\theta'}(h_{\ell t} \mid x^t) P^{\theta'}(x^t)}{\sum_{x^t \in X^t} P^{\theta''}(h_{\ell t} \mid x^t) P^{\theta''}(x^t)} \\ &\geq \frac{\min_{x^t \in X^t} P^{\theta'}(h_{\ell t} \mid x^t)}{\max_{x^t \in X^t} P^{\theta''}(h_{\ell t} \mid x^t)} \\ &= \frac{\min_{x^t \in X^t} \prod_{s=0}^{t} \phi_{z_{\ell s}}^{x_s \theta'}}{\max_{x^t \in X^t} \prod_{s=0}^{t} \phi_{z_{\ell s}}^{x_s \theta''}}. \end{aligned}$$

The last line calculates $P^{\theta}(h_{\ell t} \mid x^t)$: conditional on a state history x^t, with state x_s at time s, the probability of the signal z_ℓ is $\phi_{z_\ell}^{x_s \theta}$. Define the maximum and minimum probability of the signal z_ℓ under the parameter θ:

$$\bar{\phi}_{z_\ell}^{\theta} = \max_{x \in X} \phi_{z_\ell}^{x\theta} \quad \text{and} \quad \underline{\phi}_{z_\ell}^{\theta} = \min_{x \in X} \phi_{z_\ell}^{x\theta}.$$

As we can do the maximization and minimization above term by term, and taking logs allows us to write the product as a summation, we have

$$\log \frac{P^{\theta'}(h_{\ell t})}{P^{\theta''}(h_{\ell t})} \quad \geq \quad \sum_{s=0}^{t-1} \log \frac{\underline{\phi}_{z_{\ell s}}^{\theta'}}{\bar{\phi}_{z_{\ell s}}^{\theta''}} = \sum_{z_\ell} n_{z_\ell}^t \log \frac{\underline{\phi}_{z_\ell}^{\theta'}}{\bar{\phi}_{z_\ell}^{\theta''}} = t \sum_{z_\ell} \hat{\phi}_{z_\ell}^{\ell t} \log \frac{\underline{\phi}_{z_\ell}^{\theta'}}{\bar{\phi}_{z_\ell}^{\theta''}}.$$

Since $h_{\ell t} \in F_t^{\theta'}(\varepsilon)$, to establish the Lemma, it is sufficient to show that for ε sufficiently small,

$$0 < \min_{\phi \in \Phi_\ell^{\theta'}(\varepsilon)} \sum_{z_\ell} \phi_{z_\ell} \log \frac{\underline{\phi}_{z_\ell}^{\theta'}}{\bar{\phi}_{z_\ell}^{\theta''}}. \tag{A.4}$$

Fix $\bar{\Lambda} > \Lambda$ such that (20) continues to hold as a strict inequality with $\bar{\Lambda}$ replacing Λ. Choose $\varepsilon' > 0$ sufficiently small that for all $\varepsilon \in (0, \varepsilon')$,

$$\bar{\Lambda}\underline{\phi}_{z_\ell}^{\theta} \geq \bar{\phi}_{z_\ell}^{\theta} + (1+\bar{\Lambda})\varepsilon, \qquad \forall x, x', \theta, z_\ell, \ell, \tag{A.5}$$

and

$$\max_{\phi'' \in \Phi_\ell^{\theta''}(\varepsilon)} \min_{\phi' \in \Phi_\ell^{\theta'}(\varepsilon)} H(\phi' \| \phi'') > 2 \log \bar{\Lambda}. \tag{A.6}$$

From (A.5), $\phi'_{z_\ell}\bar{\Lambda}^{-1} \leq \underline{\phi}_{z_\ell}^{\theta'}$ for all $\phi' \in \Phi_\ell^{\theta'}(\varepsilon)$ and $\phi''_{z_\ell}\bar{\Lambda} \geq \bar{\phi}_{z_\ell}^{\theta''}$ for all $\phi'' \in \Phi_\ell^{\theta''}(\varepsilon)$. Thus

$$\min_{\phi' \in \Phi_\ell^{\theta'}(\varepsilon)} \sum_{z_\ell} \phi'_{z_\ell} \log \frac{\underline{\phi}_{z_\ell}^{\theta'}}{\bar{\phi}_{z_\ell}^{\theta''}} \geq \min_{\phi' \in \Phi_\ell^{\theta'}(\varepsilon)} \sum_{z_\ell} \phi'_{z_\ell} \log \frac{\phi'_{z_\ell}}{\bar{\Lambda}^2 \phi''_{z_\ell}}, \qquad \forall \phi'' \in \Phi_\ell^{\theta''}(\varepsilon).$$

Maximizing the right side over $\phi'' \in \Phi_\ell^{\theta''}(\varepsilon)$ we get

$$\min_{\phi \in \Phi_\ell^{\theta'}(\varepsilon)} \sum_{z_\ell} \phi_{z_\ell} \log \frac{\underline{\phi}_{z_\ell}^{\theta'}}{\bar{\phi}_{z_\ell}^{\theta''}} \geq \max_{\phi'' \in \Phi''_\ell(\varepsilon)} \min_{\phi' \in \Phi_\ell^{\theta'}(\varepsilon)} H(\phi' \| \phi'') - 2 \log \bar{\Lambda}.$$

The right side is positive by (A.6), and so (A.4) holds.

A.3 Common Learning from Average Distributions

A.3.1 Preliminaries and a Key Bound

The frequencies of pairs $x_s x_{s+1}$ of successive hidden states determines the probabilities we are interested in. We first derive an expression for $P^\theta(h_{\ell t} \cap x^t)$ in terms of these hidden pairs. Let $u_{\ell s} := (x_s, z_{\ell s}) \in X \times Z_\ell =: U$ be a complete description of agent ℓ's data generating process at time s. Denote by $n_{u_\ell u'_\ell}^t$ the number of occurrences of the ordered pair $u_\ell u'_\ell$, under the convention of periodic boundary conditions ($u_{\ell t} = u_{\ell 0}$).[13] We write $\hat{P}^t(u_\ell u'_\ell)$ for the empirical pair probability measure $(t^{-1} n_{u_\ell u'_\ell}^t)$. Since the process generating $\{u_{\ell t}\}$ is Markov, we can explicitly calculate the probability of $h_{\ell t} \cap x^t$ as

$$P^\theta(h_{\ell t} \cap x^t) = P^\theta(u_\ell^t) = \frac{P^\theta(u_{\ell 0})}{P^\theta(u_{\ell 0} \mid u_{\ell, t-1})} \prod_{u_\ell u'_\ell} P^\theta(u'_\ell \mid u_\ell)^{n_{u_\ell u'_\ell}^t}$$

[13] This guarantees that the marginal distributions of u_ℓ and u'_ℓ agree. See den Hollander (2000, §2.2) for more on the empirical pair-measure.

(where the denominator $P^\theta(u_{\ell 0} \mid u_{\ell,t-1})$ only appears if it is nonzero, in which case its presence is implied by the periodic boundary condition), and so

$$\begin{aligned} P^\theta(h_{\ell t} \cap x^t) &= O(1) \exp\left\{ \sum_{u_\ell u'_\ell} n^t_{u_\ell u'_\ell} \log P^\theta(u'_\ell \mid u_\ell) \right\} \\ &= O(1) \exp\left\{ t \sum_{u_\ell u'_\ell} \hat{P}^t(u_\ell u'_\ell) \log P^\theta(u'_\ell \mid u_\ell) \right\}. \end{aligned} \tag{A.7}$$

Thus, the frequencies of successive pairs of states and signals determines the likelihood of $P^\theta(h_{\ell t} \cap x^t)$. If one wanted to infer the hidden state x^t the conditional distribution of x^t given $h_{\ell t}$ is determined by the frequencies of the pairs $(x_s, z_{\ell s})$ and $(x_{s+1}, z_{\ell,s+1})$.

The subset of possible $\hat{P}^t$, empirical pair measures at time t, consistent with the observed history $h_{\ell t}$ of private signals is

$$L_t(h_{\ell t}) := \left\{ \hat{P}^t \in \Delta(U^2) : \exists x^t \text{ s.t. } (n^t_{u_\ell u'_\ell}) = t\hat{P}^t \text{ under } (x^t, h^t_\ell) \right\}.$$

We are now in a position to state and prove a key bound for both Lemmas 10 and 11, where A^θ is the function defined in (22).

Lemma A.2 *Suppose that agent ℓ observed a history, $h_{\ell t}$, of private signals at time t. For all $\mathcal{X}^* \subset \Delta(X)$,*

$$\begin{aligned} t^{-1} \log P^\theta(\{\hat{\xi}^t \in \mathcal{X}^*\} \cap h_{\ell t}) - t^{-1} O(\log t) \\ \le - \inf_{\hat{\xi}^t \in \mathcal{X}^*} \left\{ A^\theta(\hat{\xi}^t) - \max_{J_\ell(\hat{\phi}^t_\ell, \hat{\xi}^t)} \sum_{x, z_\ell} \hat{\xi}^t_x \hat{\phi}^x_{z_\ell} \log \phi^{x\theta}_{z_\ell} \right\}. \end{aligned} \tag{A.8}$$

Proof. We consider the probability that the signals $h_{\ell t}$ occurred for each history of the hidden state,

$$P^\theta(\{\hat{\xi}^t \in \mathcal{X}^*\} \cap h_{\ell t}) = \sum_{\{x^t \,:\, \hat{\xi}^t \in \mathcal{X}^*\}} P^\theta(h_{\ell t} \cap x^t). \tag{A.9}$$

We do the summation (A.9) in two stages: first summing over sets (or classes) of x^t's and then summing over the sets. We bound above the probability of these sets and then use the fact that the number of sets grows polynomially in t to bound this sum.

The set of state histories x^t that (when combined with the signal history $h_{\ell t}$) generate any particular empirical pair-measure $\hat{P}^t \in L(h_{\ell t})$ is

$$R_t(\hat{P}^t, h_{\ell t}) := \left\{ x^t : (n^t_{u_\ell u'_\ell}) = t\hat{P}^t \text{ under the history } (x^t, h_{\ell t}) \right\}. \tag{A.10}$$

Partitioning X^t using the sets $R_t(\hat{P}^t, h_{\ell t})$, we rewrite the sum in (A.9) as

$$P^\theta(\{\hat{\xi}^t \in \mathcal{X}^*\} \cap h_{\ell t}) = \sum_{\{x^t \,:\, \hat{\xi}^t \in \mathcal{X}^*\}} P^\theta(h_{\ell t} \cap x^t) = \sum_{\substack{\hat{P}^t \in L_t(h_{\ell t}),\\ \hat{\xi}^t \in \mathcal{X}^*}} \sum_{x^t \in R_t(\hat{P}^t, h_{\ell t})} P^\theta(h_{\ell t} \cap x^t).$$

On $R_t(\hat{P}^t, h_{\ell t})$, the value of $P^\theta(h_{\ell t} \cap x^t)$ is constant (by (A.7)) (up to $O(1)$ effects). Hence a substitution from (A.7) gives

$$\begin{aligned} & P^\theta(\{\hat{\xi}^t \in \mathcal{X}^*\} \cap h_{\ell t}) \\ = & \sum_{\substack{\hat{P}^t \in L_t(h_{\ell t}),\\ \hat{\xi}^t \in \mathcal{X}^*}} \left| R_t(\hat{P}, h_{\ell t}) \right| O(1) \exp\left\{ (t-1) \sum_{u_\ell u'_\ell} \hat{P}^t(u_\ell u'_\ell) \log P^\theta(u'_\ell \mid u_\ell) \right\}. \end{aligned} \tag{A.11}$$

(We use $|\cdot|$ to denote the number of elements in a set.)

It remains to estimate the number of histories x^t with the property that, when combined with the signal history $h_{\ell t}$, they generate a fixed frequency of the pairs $(u_{\ell s} u_{\ell, s+1})$. That is, we need to estimate the cardinality of the set $R_t(\hat{P}^t, h_{\ell t})$ for different values of $\hat{P}^t \in L_t(h_{\ell t})$.

Generate sequences x^t by taking the current state u_s and choosing a successor state u_{s+1} (which determines x_{s+1}) consistent with next period's signal $z_{\ell, s+1}$. There are $n^t_{u_s z_{\ell, s+1}}$ such transitions from $u = u_s$ to $u' = u_{s+1}$, and so there are $\prod_{u_\ell z'_\ell}(n^t_{u_\ell z'_\ell}!)$ choices. This double counts some histories (permuting the $n^t_{uu'}$ transitions from u to u' does not change the history), so we divide by the factor $\prod_{u_\ell u'_\ell}(n^t_{u_\ell u'_\ell}!)$. Hence, we have the upper bound

$$\left| R_t(\hat{P}^t, h^t_\ell) \right| \le \frac{\prod_{u_\ell z'_\ell}(n^t_{u_\ell z'_\ell}!)}{\prod_{u_\ell u'_\ell}(n^t_{u_\ell u'_\ell}!)}.$$

This upper bound is not tight, since it also includes impossible histories (there is no guarantee that it is possible to move to z_{s+2} from the successor pair (z_{s+1}, x)).

Applying Stirling's formula,

$$\begin{aligned} \left| R_t(\hat{P}^t, h_{\ell t}) \right| & \le \exp\left\{ \sum_{u_\ell z'_\ell} \log(n^t_{u_\ell z'_\ell}!) - \sum_{u_\ell u'_\ell} \log(n^t_{u_\ell u'_\ell}!) \right\} \\ & \sim O(1) \exp\left\{ \tfrac{1}{2} \sum_{u_\ell z'_\ell} \log n^t_{u_\ell z'_\ell} - \tfrac{1}{2} \sum_{u_\ell u'_\ell} \log n^t_{u_\ell u'_\ell} \right. \\ & \qquad \left. + \sum_{u_\ell z'_\ell} n^t_{u_\ell z'_\ell} \log(n^t_{u_\ell z'_\ell}) - \sum_{u_\ell u'_\ell} n^t_{u_\ell u'_\ell} \log(n^t_{u_\ell u'_\ell}) \right\} \\ & = O(t) \exp \sum_{u_\ell u'_\ell} n^t_{u_\ell u'_\ell} \log \frac{n^t_{u_\ell z'_\ell}}{n^t_{u_\ell u'_\ell}} \end{aligned}$$

$$= \quad -O(t)\exp\left\{t\sum_{u_\ell u'_\ell}\hat{P}^t(u_\ell u'_\ell)\log\hat{P}^t(x' \mid u_\ell, z'_\ell)\right\}.$$

Combining this with (A.11), we obtain

$$P^\theta(\{\hat{\xi}^t \in \mathcal{X}^*\} \cap h_{\ell t}) \le O(t) \sum_{\substack{\hat{P}^t \in L_t(h_{\ell t}),\\ \hat{\xi}^t \in \mathcal{X}^*}} \exp\left\{t\sum_{u_\ell u'_\ell}\hat{P}^t(u_\ell u'_\ell)\log\frac{P^\theta(u'_\ell \mid u_\ell)}{\hat{P}^t(x' \mid u_\ell, z'_\ell)}\right\}.$$

The number of terms in $L_t(h_{\ell t})$ is bounded above by $(t+1)^{|U|^2}$ and so only grows polynomially in t.[14] Applying this upper bound to the number of terms in the summation and multiplying it by the largest term (and rewriting the argument of the log using Bayes' rule) yields

$$\begin{aligned}
&P^\theta(\{\hat{\xi}^t \in \mathcal{X}^*\} \cap h_{\ell t})\\
&\quad \le O(t)(t+1)^{|U|^2} \sup_{\substack{\hat{P}^t \in_t L(h_{\ell t}),\\ \hat{\xi}^t \in \mathcal{X}^*}} \exp\left\{t\sum_{u_\ell u'_\ell}\hat{P}^t(u_\ell u'_\ell)\log\frac{P^\theta(u'_\ell \mid u_\ell)}{\hat{P}^t(u'_\ell \mid u_\ell)}\frac{\hat{P}^t(u_\ell, z'_\ell)}{\hat{P}^t(u_\ell)}\right\}.
\end{aligned}$$

Taking logarithms and dividing by t, we get

$$t^{-1}\log P^\theta(\{\hat{\xi}^t \in \mathcal{X}^*\} \cap h^t_\ell) - t^{-1}O(\log t) \tag{A.12}$$

$$\begin{aligned}
&\le \sup_{\substack{\hat{P}^t \in L_t(h_{\ell t})\\ \hat{\xi}^t \in \mathcal{X}^*}} \sum_{u_\ell u'_\ell}\hat{P}^t(u_\ell u'_\ell)\log\frac{P^\theta(u'_\ell \mid u_\ell)}{\hat{P}^t(u'_\ell \mid u_\ell)} + \sum_{u_\ell z'_\ell}\hat{P}(u_\ell, z'_\ell)\log\frac{\hat{P}^t(u_\ell, z'_\ell)}{\hat{P}^t(u_\ell)}\\
&= \sup_{\substack{\hat{P}^t \in L_t(h_{\ell t})\\ \hat{\xi}^t \in \mathcal{X}^*}} E^{\hat{P}^t}\left(\log\frac{\pi^\theta_{xx'}\phi^{x'\theta}_{z'_\ell}}{\hat{P}^t(u'_\ell \mid u_\ell)} + \log\hat{P}^t(z'_\ell \mid u_\ell)\right) \qquad \text{(A.13)}\\
&= \sup_{\substack{\hat{P}^t \in L_t(h_{\ell t})\\ \hat{\xi}^t \in \mathcal{X}^*}} E^{\hat{P}^t}\left(\log\frac{\pi^\theta_{xx'}}{\hat{P}^t(x' \mid u_\ell)} + \log\frac{\phi^{x'\theta}_{z'_\ell}}{\hat{P}^t(z'_\ell \mid x', u_\ell)} + \log\hat{P}^t(z'_\ell \mid u_\ell)\right)\\
&= \sup_{\substack{\hat{P}^t \in L_t(h_{\ell t})\\ \hat{\xi}^t \in \mathcal{X}^*}} E^{\hat{P}^t}\left(\log\frac{\pi^\theta_{xx'}}{\hat{P}^t(x' \mid u_\ell)} + \log\frac{\hat{P}^t(z'_\ell \mid u_\ell)}{\hat{P}^t(z'_\ell \mid x', u_\ell)} + \log\phi^{x'\theta}_{z'_\ell}\right).
\end{aligned}$$

Above we decompose the first term of (A.13) into two parts (using Bayes' Rule). Now we write out the expectations in full, which allows us to write the first two terms in the above as relative entropies of conditional distributions:

$$\sum_{u_\ell u'_\ell}\hat{P}^t(u_\ell u'_\ell)\left(-\log\frac{\hat{P}^t(x' \mid u_\ell)}{\pi^\theta_{xx'}} - \log\frac{\hat{P}^t(z'_\ell \mid x', u_\ell)}{\hat{P}^t(z'_\ell \mid u_\ell)} + \log\phi^{x'\theta}_{z'_\ell}\right) \tag{A.14}$$

[14]There are at most $(t+1)^{|U|^2}$ elements in this summation: the number in the (u, u')th entry can take at most $t+1$ values. See Cover and Thomas (1991, Theorem 12.1.1, p.280).

$$= -\sum_{u_\ell} \hat{P}^t(u_\ell) H\left([\hat{P}^t(x' \mid u_\ell)]_{x'} \,\Big\|\, [\pi^\theta_{xx'}]_{x'} \right)$$

$$- \sum_{u_\ell, x'} \hat{P}^t(u_\ell, x') H\left([\hat{P}^t(z'_\ell \mid x' u_\ell)]_{z'_\ell} \,\Big\|\, [\hat{P}^t(z'_\ell | u_\ell)]_{z'_\ell} \right) + \sum_{x, z_\ell} \hat{P}^t(x, z_\ell) \log \phi^{x\theta}_{z_\ell}.$$

Since

$$\sum_{z_\ell} \hat{P}^t(z_\ell \mid x) \hat{P}^t(x' \mid u_\ell) = \sum_{z_\ell} \hat{P}^t(x' \mid x, z_\ell) \hat{P}^t(z_\ell \mid x) = \hat{P}^t(x' \mid x) := \hat{\pi}^t_{xx'},$$

from the convexity of relative entropy, the first term in (A.14) is less than

$$-\sum_x \hat{P}^t(x) H\left([\hat{\pi}^t_{xx'}]_{x'} \,\|\, [\pi^\theta_{xx'}]_{x'} \right).$$

Writing H^* for the middle (relative entropy) of (A.14) we now have the following upper bound for (A.12):

$$-\inf_{\substack{\hat{P}^t \in L_t(h_{\ell t}) \\ \hat{\xi}^t \in \mathcal{X}^*}} \sum_{x, z_\ell} \hat{P}^t(x, z_\ell) \left\{ H\left([\hat{\pi}^t_{xx'}]_{x'} \,\|\, [\pi^\theta_{xx'}]_{x'} \right) + H^* - \log \phi^{x\theta}_{z_\ell} \right\}.$$

As H^* is a relative entropy it is non-negative so excluding it only weakens the bound.

The infimum over $\hat{P}^t \in L_t(h_{\ell t})$ can be taken by first minimizing over $\hat{\pi}_{xx'}$ subject to the requirement that $\hat{\xi}^t$ is the marginal distribution of x' given the marginal distribution $\hat{\xi}^t$ of x (so that $\hat{\xi}^t$ is the stationary distribution of the Markov chain with transition probabilities $\hat{\pi}$). By a version of Sanov's Theorem for empirical pair-measures (den Hollander, 2000, Theorem IV.7, p.45),

$$\inf_{\hat{\pi}} \sum_x \hat{\xi}_x H\left([\hat{\pi}_{xx'}]_{x'} \,\Big\|\, [\pi^{\theta''}_{xx'}]_{x'} \right) = \sup_{v \in \mathbb{R}^{|X|}_{++}} \sum_{x'} \hat{\xi}_{x'} \log \frac{v_{x'}}{\sum_{\tilde{x}} v_{\tilde{x}} \pi^{\theta''}_{\tilde{x}x'}} = A^{\theta''}(\hat{\xi}).$$

We thus have the following upper bound for (A.12):

$$-\inf_{\hat{\xi}^t \in \mathcal{X}^*} \left\{ A^\theta(\hat{\xi}^t) - \max_{J_\ell(\hat{\phi}^t_\ell, \hat{\xi}^t)} \sum_{x, z_\ell} \hat{\xi}^t_x \hat{\phi}^x_{z_\ell} \log \phi^{x\theta}_{z_\ell} \right\}.$$

■

A.3.2 Proof of Lemma 10

We prove that θ' is individually learned (θ'' is identical). It is sufficient to show that

$$\frac{P^{\theta'}(h_{\ell t})}{P^{\theta''}(h_{\ell t})} \to \infty, \tag{A.15}$$

on the event $\breve{F}^{\theta'}_{\varepsilon t}$ as $t \to \infty$, and that the divergence is uniform in the histories on $\breve{F}^{\theta'}_{\varepsilon t}$. We prove this by constructing an upper bound for the numerator and a lower bound for the denominate of (A.15) that only depend on the event $\breve{F}^{\theta'}_{\varepsilon t}$.

We define $X^{\theta'}_{\nu}$ to be the set of hidden state frequencies ξ within ν of their stationary distribution under θ':

$$X^{\theta'}_{\nu} := \{\xi \in \Delta(X) : \|\xi - \xi^{\theta'}\| < \nu\}.$$

For any signal distribution ϕ_ℓ, $K^{\theta'}_{\nu\ell}(\phi_\ell)$ is the set of pairs of state frequencies $\xi \in X^{\theta'}_{\nu}$ and conditional signal distributions $(\phi^x_\ell)_x$ with the property that the conditional signal distributions $(\phi^x_\ell)_x$ are consistent with ϕ_ℓ and ξ:

$$K^{\theta'}_{\nu\ell}(\phi_\ell) := \left\{(\xi, (\phi^x_\ell)_{x\in X}) : \xi \in X^{\theta'}_{\nu},\ (\phi^x_\ell)_x \in J_\ell(\phi_\ell, \xi)\right\}.$$

Given the distribution, $\hat{\phi}^t_\ell$, of signals from $h_{\ell t}$, the event $K^{\theta'}_{\nu\ell t}(\hat{\phi}^t_\ell)$ is the event that the realized frequencies $\hat{\xi}^t$ of hidden states are both close to their stationary distribution and the associated realized conditional frequencies $(\hat{\phi}^x_\ell)_x$ are consistent with the observed history of private signals, i.e., $(\hat{\xi}^t, (\hat{\phi}^x_\ell)_x) \in K^{\theta'}_{\nu\ell}(\hat{\phi}^t_\ell)$.

We begin by providing a lower bound for the numerator in (A.15), where we write $x^t \in X^{\theta'}_{\nu}$ if the implied frequency over hidden states from the hidden state history x^t is in $X^{\theta'}_{\nu}$:

$$\begin{aligned} P^{\theta'}(h_{\ell t}) &\geq \sum_{x^t \in X^{\theta'}_{\nu}} P^{\theta'}(h_{\ell t} \mid x^t) P^{\theta'}(x^t) \\ &\geq \min_{\tilde{x} \in X^{\theta'}_{\nu}} P^{\theta'}(h_{\ell t} \mid \tilde{x}^t) P^{\theta'}(x^t \in X^{\theta'}_{\nu}). \end{aligned}$$

Here we take the minimum of one of the terms in the summation and take it outside the sum. This inequality can be rewritten as

$$t^{-1} \log P^{\theta'}(h_{\ell t}) \geq t^{-1} \log P^{\theta'}(x^t \in X^{\theta'}_{\nu}) + \min_{(\hat{\xi}, (\hat{\phi}^x_\ell)) \in K^{\theta'}_{\nu\ell t}(\hat{\phi}^t_\ell)} \sum_x \hat{\xi}_x \sum_{z_\ell} \hat{\phi}^x_{z_\ell} \log \phi^{x\theta'}_{z_\ell},$$

using $t^{-1} \log P^{\theta'}(h_{\ell t} \mid x^t) = \sum_x \hat{\xi}_x \sum_{z_\ell} \hat{\phi}^x_{z_\ell} \log \phi^{x\theta'}_{z_\ell}$, where $(\hat{\xi}, (\hat{\phi}^x_\ell))$ are the relevant frequencies in $(x^t, h_{\ell t})$. As $P^{\theta'}(x^t \in X^{\theta'}_{\nu}) \to 1$, we can simplify this to

$$t^{-1} \log P^{\theta'}(h_{\ell t}) \geq O(t^{-1}) + \min_{K^{\theta'}_{\nu\ell t}(\hat{\phi}^t_\ell)} \sum_{x, z_\ell} \hat{\xi}_x \hat{\phi}^x_{z_\ell} \log \phi^{x\theta'}_{z_\ell}. \tag{A.16}$$

Now we combine this with the bound from Lemma A.2. In particular, using the bound (A.8) with $\mathcal{X}^* = \Delta(X)$ on the denominator and (A.16) on the numerator, we obtain a bound on the ratio given by

$$t^{-1} \log \frac{P^{\theta'}(h_{\ell t})}{P^{\theta''}(h_{\ell t})} \geq O(t^{-1}) + \min_{K^{\theta'}_{\nu\ell t}(\hat{\phi}^\ell)} \sum_{x, z_\ell} \hat{\xi}_x \hat{\phi}^x_{z_\ell} \log \phi^{x\theta'}_{z_\ell} - t^{-1} O(\log t)$$

$$+ \inf_{\hat{\xi}^t \in \Delta(X)} \left\{ A^{\theta''}(\hat{\xi}^t) - \max_{J_\ell(\hat{\phi}_\ell^t, \hat{\xi}^t)} \sum_{x, z_\ell} \hat{\xi}_x^t \hat{\phi}_{z_\ell}^x \log \phi_{z_\ell}^{x\theta''} \right\}.$$

A sufficient condition for (A.15) is therefore that there exist $\varrho > 0$ such that for t sufficiently large,

$$\min_{(\hat{\xi}, \hat{\phi}_\ell^x) \in K_{\nu\ell t}^{\theta'}(\hat{\phi}_\ell^t)} \sum_{x,z} \hat{\xi}_x \hat{\phi}_{z_\ell}^x \log \phi_{z_\ell}^{x\theta'}$$
$$+ \inf_{\hat{\xi}^t \in \Delta(X)} \left\{ A^{\theta''}(\hat{\xi}^t) - \max_{J_\ell(\hat{\phi}_\ell^t, \hat{\xi}^t)} \sum_{x, z_\ell} \hat{\xi}_x^t \hat{\phi}_{z_\ell}^x \log \phi_{z_\ell}^{x\theta''} \right\} > \varrho. \quad \text{(A.17)}$$

For ε small, any $h_{\ell t} \in \breve{\Phi}_\ell^{\theta'}(\varepsilon)$ implies a signal distribution close to $\psi_\ell^{\theta'}$, and for ν small, every state distribution in $X_\nu^{\theta'}$ is close to $\xi^{\theta'}$. Hence, for ν and ε sufficiently small, $K_{\nu\ell t}^{\theta'}(\hat{\phi}_\ell^t)$ is close to $J_\ell(\psi_\ell^{\theta'}, \xi^{\theta'})$ and (A.17) is implied by

$$\min_{J_\ell(\psi_\ell^{\theta'}, \xi^{\theta'})} \sum_{x, z_\ell} \xi_x^{\theta'} \hat{\phi}_{z_\ell}^x \log \phi_{z_\ell}^{x\theta'} + \min_{\tilde{\xi} \in \Delta(X)} \left[A^{\theta''}(\tilde{\xi}) - \max_{J_\ell(\psi_\ell^{\theta'}, \tilde{\xi})} \sum_{x, z_\ell} \tilde{\xi}_x \hat{\phi}_{z_\ell}^x \log \phi_{z_\ell}^{x\theta''} \right] > 0.$$

This is clearly ensured by Assumption 8.

A.3.3 Proof of Lemma 11

We prove that the set $\breve{F}_{\varepsilon^\dagger t}^{\theta'}$ is q-evident under θ' for t sufficiently large (the argument for the other parameter is identical). To establish this, it is sufficient to show that for all $q \in (0,1)$, there exists a T such that for all $t \geq T$, if $\hat{\phi}_1^t \in \breve{\Phi}_1^{\theta'}(\varepsilon^\dagger)$, then agent 1 attaches probability at least q to θ' (proved in Lemma 10) *and* probability at least q to $\hat{\phi}_2^t \in \breve{\Phi}_2^{\theta'}(\varepsilon^\dagger)$. We, therefore, consider agent 1's beliefs about agent 2's signals.

The first step is to show that to characterize agent 1's beliefs about agent 2's signals it is sufficient to characterize her beliefs about the hidden states. Agent 2's signal in period s is sampled from $\phi_2^{x_s\theta}$. Conditional on x^t, therefore, agent 2's signals are independently (but not identically) distributed across time and we can apply Cripps, Ely, Mailath, and Samuelson (2008, Lemma 3) to deduce that there exists a $\kappa > 0$ such that for all $\gamma > 0$,

$$P^{\theta'}\Big(\|\hat{\phi}_2^t - \textstyle\sum_x \hat{\xi}_x^t \phi_2^{x\theta}\| > \gamma \Big| x^t \Big) < \kappa e^{-t\gamma^2}, \qquad \forall x^t.$$

Hence, conditional on x^t, agent 1 makes a very small error in determining $\hat{\phi}_2^t$. This inequality holds for all x^t and also holds conditioning on the full history (x^t, h_{1t}), because (conditional on x^t) agent 2's signals are independent of h_{1t}. If we define $G^t := \{\omega : \|\hat{\phi}_2^t - \sum_x \hat{\xi}_x^t \phi_2^{x\theta}\| > \gamma\}$, then for all h_{1t}

$$P^{\theta'}\left(\|\hat{\phi}_2^t - \textstyle\sum_x \hat{\xi}_x^t \phi_2^{x\theta}\| > \gamma \Big| h_{1t} \right) = \sum_{x^t} P^{\theta'}(x^t \mid h_{1t}) P^{\theta'}(G^t \mid h_{1t}, x^t)$$

$$
\begin{aligned}
&= \sum_{x^t} P^{\theta'}(x^t|h_{1t})P^{\theta'}(G^t \mid x^t) \\
&< \kappa e^{-t\gamma^2},
\end{aligned}
$$

where the last line substitutes the previous inequality. The triangle inequality can be used to bound the gap between $\hat{\phi}_2^t$ and its unconditional expected value, $\psi_2^{\theta'}$, by two terms. One measures the gap between $\hat{\phi}_2^t$ and its expected value conditional on x^t. The other measures the gap between its unconditional expected value and its expectation conditional on x^t:

$$
\|\hat{\phi}_2^t - \psi_2^{\theta'}\| \leq \|\hat{\phi}_2^t - \textstyle\sum_x \hat{\xi}_x^t \phi_2^{x\theta'}\| + \|\textstyle\sum_x \hat{\xi}_x^t \phi_2^{x\theta'} - \psi_2^{\theta'}\|.
$$

Conditional on h_{1t}, we have an upper bound on the probability that the first of the terms on the right side is bigger than γ for all h_{1t}. This probability, therefore, is also the probability that the the second term on the right side is close to a bound on the left side, and so assuming $\gamma < \varepsilon^\dagger$:

$$
P^{\theta'}\Big(\|\hat{\phi}_2^t - \psi_2^{\theta'}\| > \varepsilon^\dagger \Big| h_{1t}\Big) < \kappa e^{-t\gamma^2} + P^{\theta'}\Big(\Big\|\textstyle\sum_x \hat{\xi}_x^t \phi_2^{x\theta'} - \psi_2^{\theta'}\Big\| > \varepsilon^\dagger - \gamma \Big| h_{1t}\Big). \tag{A.18}
$$

The next step in the argument is to describe how close the hidden state distributions, $\hat{\xi}^t$, need to be to their expected values for agent 1 to believe 2's signals are in $\breve{\Phi}_2^{\theta'}(\varepsilon^\dagger)$. The summation in the right side of (A.18) can be written as $M\hat{\xi}^t$ and the term $\psi_2^{\theta'}$ can be written as $M\xi^{\theta'}$, where M is the $|Z| \times |X|$ matrix with columns $\phi_2^{x\theta'}$. In the variation norm, Dobrushin's inequality (Brémaud, 1999, p.236) implies

$$
\Big\|\textstyle\sum_x \hat{\xi}_x^t \phi_2^{x\theta'} - \psi_2^{\theta'}\Big\| = \|M\hat{\xi}^t - M\xi^{\theta'}\| \leq \|\hat{\xi}^t - \xi^{\theta'}\| \max_{\tilde{x}\bar{x}} \Big\|\phi_2^{\tilde{x}\theta'} - \phi_2^{\bar{x}\theta'}\Big\|. \tag{A.19}
$$

However, Pinsker's inequality (Cesa-Bianchi and Lugosi, 2006, p. 371), i.e., $\|a - b\| \leq \sqrt{H(a\|b)/2}$, implies

$$
\max_{\tilde{x}\bar{x}} \|\phi_2^{\tilde{x}\theta'} - \phi_2^{\bar{x}\theta'}\| \leq \max_{\tilde{x}\bar{x}} \left(\frac{1}{2}\sum_{z_2} \phi_{z_2}^{\tilde{x}\theta'} \log \frac{\phi_{z_2}^{\tilde{x}\theta'}}{\phi_{z_2}^{\bar{x}\theta'}}\right)^{1/2} \leq \left(\frac{1}{2}\log\Lambda\right)^{1/2}, \tag{A.20}
$$

where $\Lambda > 1$ was defined in (19). Define$f(\Lambda) := (2/\log\Lambda)^{1/2}$. Applying (A.19) in (A.18) gives

$$
P^{\theta'}\Big(\|\hat{\phi}_2^t - \psi_2^{\theta'}\| > \varepsilon^\dagger \Big| h_{1t}\Big) < \kappa e^{-t\gamma^2} + P^{\theta'}\Big(\|\hat{\xi}^t - \xi^{\theta'}\| > (\varepsilon^\dagger - \gamma)f(\delta)\Big| h_{1t}\Big)
$$

for all h_{1t}. (Notice that as the signal distributions become closer and so $f(\Lambda) \to \infty$, the last term approaches zero, and so it is easy for 1 to infer 2's signals, because as the conditional distributions $\phi_2^{x\theta'}$ become more similar, it less important to infer the hidden states accurately.) Re-writing this in terms of our earlier definitions,

$$
P^{\theta'}\Big(\hat{\phi}_2^t \notin \breve{\Phi}_2^{\theta'}(\varepsilon^\dagger)\Big| h_{1t}\Big) < \kappa e^{-t\gamma^2} + P^{\theta'}\Big(\hat{\xi}^t \notin X_{(\varepsilon^\dagger-\gamma)f(\Lambda)}^{\theta'}\Big| h_{1t}\Big). \tag{A.21}
$$

We now use our previous bounds to estimate the probability on the right side of (A.21) for some $\hat{\phi}_1 \in \breve{\Phi}_1^{\theta'}(\varepsilon^\dagger)$. First, from Bayes' rule we have

$$\log P^{\theta'}\Big(\hat{\xi}^t \notin X^{\theta'}_{(\varepsilon^\dagger-\gamma)f(\Lambda)}\Big|\, h_{1t}\Big) = \log \frac{P^{\theta'}\Big(\Omega \setminus X^{\theta'}_{(\varepsilon^\dagger-\gamma)f(\Lambda)}, h_{1t}\Big)}{P^{\theta'}(h_{1t})}. \tag{A.22}$$

From Lemma A.2, we use (A.8) to bound the numerator,

$$\begin{aligned} t^{-1}\log P^{\theta'}\Big(\Omega \setminus X^{\theta'}_{(\varepsilon^\dagger-\gamma)f(\Lambda)}, h_{1t}\Big) - t^{-1}O(\log t) \\ \leq - \inf_{\hat{\xi}^t \in \Omega\setminus X^{\theta'}_{(\varepsilon^\dagger-\gamma)f(\Lambda)}} \left\{ A^{\theta'}(\hat{\xi}^t) - \max_{J_\ell(\hat{\phi}^t_\ell,\hat{\xi}^t)} \sum_{x,z_\ell} \hat{\xi}^t_x \hat{\phi}^x_{z_\ell} \log \phi^{x\theta'}_{z_\ell} \right\}. \end{aligned}$$

This infimum is finite by Assumption 5: If the signals did not have full support it might be impossible to generate the history $h_{1t} \in \breve{F}^{\theta'}_{\varepsilon^\dagger t}$ from the hidden histories in $\Omega \setminus X^{\theta'}_{(\varepsilon^\dagger-\gamma)f(\Lambda)}$.

We use (A.16) to bound the denominator of (A.22) in the same way as it was used to derive (A.17). Substituting the bounds on the fraction (A.22) into (A.21), therefore, provides an upper bound on the probability that agent 1 believes that agent 2's signals are not in the set $\breve{F}^{\theta'}_{\varepsilon^\dagger t}$. That is,

$$P^{\theta'}\left(h_{2t} \notin \breve{F}^{\theta'}_{\varepsilon^\dagger t} \mid h_{1t} \in \breve{F}^{\theta'}_{\varepsilon^\dagger t}\right) = P^{\theta'}\left(\hat{\phi}^t_2 \notin \breve{\Phi}^{\theta'}_2(\varepsilon^\dagger) \mid h_{1t}\right) \leq \kappa e^{-t\gamma^2} + \kappa' e^{-t\mathfrak{H}},$$

where $\kappa' > 0$ is polynomial in t and

$$\mathfrak{H} := \min_{K^{\theta'}_{\nu 1t}(\hat{\phi}^t_1)} \sum_{x,z_1} \hat{\xi}_x \hat{\phi}^x_{z_1} \log \phi^{x\theta'}_{z_1} + \inf_{\Omega\setminus X^{\theta'}_{(\varepsilon^\dagger-\gamma)f(\Lambda)}} \left\{ A^{\theta'}(\hat{\xi}^t) - \max_{J_\ell(\hat{\phi}^t_\ell,\hat{\xi}^t)} \sum_{x,z_\ell} \hat{\xi}^t_x \hat{\phi}^x_{z_\ell} \log \phi^{x\theta'}_{z_\ell} \right\}.$$

If we can show $\mathfrak{H} > 0$ for all $\hat{\phi}_1 \in \breve{\Phi}^{\theta'}_1(\varepsilon^\dagger)$, then we have proved the lemma. By choosing ν small the terms $\hat{\xi}^t_x$ in the first sums can be made arbitrarily close to $\xi^{\theta'}_x$. We can also choose $\gamma = ct^{-1/3} \to 0$ as $t \to \infty$. So a sufficient condition for the above is

$$\min_{J_1(\hat{\phi}^t_1,\xi^{\theta'})} \sum_{x,z_1} \xi^{\theta'}_x \hat{\phi}^x_{z_1} \log \phi^{x\theta'}_{z_1} + \inf_{\|\xi-\xi'\|>\varepsilon^\dagger f(\Lambda)} \left\{ A^{\theta'}(\xi) - \max_{J_1(\hat{\phi}^t_1,\xi)} \sum_{x,z_1} \xi_x \hat{\phi}^x_{z_1} \log \phi^{x\theta'}_{z_1} \right\} > 0$$

for all $\hat{\phi}^t_1 \in \breve{\Phi}^{\theta'}_1(\varepsilon^\dagger)$. Assumption 9 thus implies that $\mathfrak{H} > 0$. The proof is completed by observing that the bound $\mathfrak{H}$ is independent of the details of the history h_{1t}. In particular, the order of the polynomial terms in T is determined by the number of state-signal pairs, and not the specific history.

References

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.

Cesa-Bianchi, N., and G. Lugosi (2006): *Prediction, Learning, and Games.* Cambridge University Press.

Cover, T. M., and J. A. Thomas (1991): *Elements of Information Theory.* John Wiley & Sons, Inc., New York.

Cripps, M. W., J. C. Ely, G. J. Mailath, and L. Samuelson (2008): "Common Learning," *Econometrica*, 76(4), 909–933.

den Hollander, F. (2000): *Large Deviations.* American Mathematical Society, Providence, Rhode Island.

Ephraim, Y., and N. Merhav (2002): "Hidden Markov Processes," *IEEE Transactions on Information Theory*, 48(6), 1518–1569.

Halpern, J. Y., and Y. Moses (1990): "Knowledge and Common Knowledge in a Distributed Environment," *Journal of the ACM*, 37, 549–587.

Monderer, D., and D. Samet (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1(2), 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.

Steiner, J., and C. Stewart (2011): "Communication, Timing, and Common Learning," *Journal of Economic Theory*, 146(1), 230–247.