Chapter 6

Learning

We now extend the Bayesian framework described in 2.1 to accommodate learning from a sequence of signals. Section 6.3 asks whether an agent will eventually learn the state. Section 6.4 asks whether agents with different prior beliefs will eventually hold similar beliefs. Section 6.5 asks whether agents with different priors expect their disagreement to reduce given information (thus studying a second-order belief). Section 6.6 asks whether agents will commonly learn, i.e., whether agents will eventually believe that other agents believe that they ... have learned the state.

6.1 Preliminaries

Let (Θ, d_Θ) be a complete separable metric space endowed with its Borel σ-algebra Σ, and let $p \in \Delta(\Theta)$ be a (Σ-measurable) probability measure on Θ. As before, we interpret $\theta \sim p$ as an unknown parameter of interest.

The space of signal realizations $(\mathcal{X}, d_X)$ is again a complete separable metric space endowed with its Borel σ-algebra $\mathcal{B}$. There is an infinite sequence of signal realizations $X_1, X_2, \ldots$ taking values in the set $\mathcal{X}^\infty = \mathcal{X}_1 \times \mathcal{X}_2 \times \ldots$ where each $\mathcal{X}_t$ is a copy of $\mathcal{X}$. Conditional on the realized θ, signals $X_1, X_2, \ldots$ are generated iid according to a conditional density f_θ, and we refer to each X_t as the period-t signal.

The full state space is $\Omega = \Theta \times \mathcal{X}^\infty = \Theta \times \mathcal{X}_1 \times \mathcal{X}_2 \times \ldots$ and it is equipped with the product σ-algebra $\Sigma \times \mathcal{B}_1 \times \mathcal{B}_2 \times \ldots$ where each $\mathcal{B}_t$ is a copy of $\mathcal{B}$. Throughout, we use P to denote the measure on Ω induced by p and the family $(f_\theta)_{\theta \in \Theta}$, and we use P_θ to denote the conditional measure on $\mathcal{X}^\infty$ when the parameter is θ.

6.2 Binary Example

First consider a single-agent environment with two possible parameter values $\theta \in \{A, B\}$. Each period $t \in \mathbb{Z}_+$ a signal realization from $\{a, b\}$ is generated iid according to

	a	b
A	q	$1-q$
B	$1-q$	q

where $q > 1/2$. Will an agent who holds a prior belief that the probability of A is $p \in (0,1)$ eventually learn the value of the parameter?

Suppose first that the parameter is $\theta = A$, in which case signals are drawn iid according to $f_A = (q, 1-q)$. For any infinite sequence $\mathbf{x} \in \{a,b\}^\infty$ and any $t \in \mathbb{Z}_+$, let

$$n_t(\mathbf{x}) \equiv \#\{1 \le t' \le t : x_{t'} = a\}$$

denote the number of a-realizations among the first t realizations of $\mathbf{x}$. By the strong law of large numbers, there is a set $\mathcal{X}_0^\infty \subseteq \mathcal{X}^\infty$ of P_A-measure 1 such that

$$\lim_{t\to\infty} \frac{n_t(\mathbf{x})}{t} = q \quad \forall \mathbf{x} \in \mathcal{X}_0^\infty.$$

That is, the limiting fraction of a-realizations is q along each sequence in $\mathcal{X}_0^\infty$.

Since signals are assumed to be conditionally independent, the agent's posterior belief about A following any sequence $(x_1, \ldots, x_t)$ depends only on the count of a and b-realizations. Let n denote the number of a-realizations. Then applying Bayes' rule (Section 2.2.1), the agent's posterior belief is

$$\begin{aligned} P(\theta = A \mid x_1, \ldots, x_t) &= \frac{pq^n(1-q)^{t-n}}{pq^n(1-q)^{t-n} + (1-p)(1-q)^n q^{t-n}} \\ &= \frac{1}{1 + \frac{1-p}{p}\left(\frac{1-q}{q}\right)^{2n-t}} \end{aligned} \tag{6.1}$$

Along any $\mathbf{x} \in \mathcal{X}_0^\infty$ we have

$$\lim_{t\to\infty} P(\theta = A \mid x_1, \ldots, x_t) = \lim_{t\to\infty} \left(1 + \frac{1-p}{p}\left[\left(\frac{1-q}{q}\right)^{2\frac{n_t(\mathbf{x})}{t} - 1}\right]^t\right)^{-1} = 1$$

recalling that $q > 1/2$ by assumption.

So the agent's posterior belief P_A-almost surely converges to certainty of the correct value of the parameter, A. An identical argument shows that when the parameter is B then the agent's posterior belief P_B-almost surely converges to certainty of B. Thus the agent (eventually) learns the parameter.

6.3 Doob's Consistency Theorem

A classic result due to Doob (1949) generalizes the individual learning result from the previous section.[1]

[1] Our presentation of this material follows Miller (2018).

Assumption 3 (Identifiability)**.** *If $\theta \neq \theta'$, then $P_\theta \neq P_{\theta'}$.*

In words, Assumption 3 is satisfied if no pair of parameter values induce the same distribution over signals, meaning the parameter is identifiable from its observable implications.

Proposition 17. *Suppose Assumption 3 is satisfied, and let $g : \Theta \to \mathbb{R}$ be any measurable function satisfying $\mathbb{E}|g(\theta)| < \infty$. Then*

$$\lim_{t\to\infty} \mathbb{E}(g(\theta) \mid X_1, X_2, \ldots, X_t) = g(\theta) \quad P\text{-a.s.}$$

In the special case where $g(\theta) = \theta$, the result implies that the posterior expectation of θ converges to its true value almost surely. The following proposition is a Bayesian analogue of the above result, and says that posterior beliefs converge almost surely to a degenerate measure at the true state.

Proposition 18 (Posterior Consistency)**.** *Suppose Assumption 3 holds. Then, there exists a set $\Theta' \subseteq \Theta$ with $p(\Theta') = 1$ such that for every $\theta_0 \in \Theta'$ and every neighborhood B of θ_0,*

$$\lim_{t\to\infty} \mathbb{P}(\theta \in B \mid X_1, X_2, \ldots, X_t) = 1 \quad P_{\theta_0}\text{-a.s.}$$

That is, for any prior distribution, the posterior belief is guaranteed to concentrate in a neighborhood of the true parameter θ—except possibly on a set of parameter values that has measure zero under the agent's prior.

Remark 6.1. The qualification that learning occurs except on a set of "measure zero under the agent's prior" is less harmless than it might initially seem. Consider $\Theta = \mathbb{R}$ where the agent's prior $p \in \Delta(\Theta)$ is a point mass at $\theta = 0$. Then the posterior is also a point mass at zero, so the agent will fail to learn any parameter which is different from 0. But because the set $\mathbb{R}\backslash\{0\}$ has measure zero under the agent's prior, the statement of the result holds in a trivial sense. See also the subsequent discussion in Section 7.2.1.

Remark 6.2. Proposition 18 implies that the agent's posterior belief converges almost surely to a point mass on the true parameter in the topology of weak convergence, i.e., there is a P_θ-measure 1 set of sequences of signal realizations such that

$$d(P^t, \delta_\theta) \to 0$$

along each of these sequences, where d denotes the Levy-Prokhorov metric and $P^t \in \Delta(\Theta)$ denotes the posterior belief after observing the first t coordinates of the sequence. Since d is a metric, we also have that for any alternative prior $\widetilde{p} \in \Delta(\Theta)$ and corresponding posterior belief $\widetilde{P}^t \in \Delta(\Theta)$ (updating to the same t realizations),

$$d(P^t, \widetilde{P}^t) \leq d(P^t, \delta_\theta) + d(\delta_\theta, \widetilde{P}^t).$$

Since the RHS converges to zero almost surely (by Proposition 18), the two agents' posterior beliefs converge to one another almost surely in the topology of weak convergence. The subsequent section provides an even stronger version of this result.

6.4 Merging of Beliefs

Assume that for each $t \geq 1$, a unique conditional probability distribution $P^t(x_1, \ldots, x_t)(C)$ exists for all realized sequences $x_1, \ldots, x_t \in \mathcal{X}_1 \times \cdots \times \mathcal{X}_t$ and unknown events $C \in \mathcal{B}_{t+1} \times \mathcal{B}_{t+2} \times \ldots$.[2] Blackwell and Dubins (1962) show that even if players start out with different prior beliefs, their conditional beliefs will merge to one another in a strong sense.

To state the result formally, recall that for any two probability measures μ_1, μ_2 defined on the same σ-algebra $\mathcal{F}$, *total variation distance* and *absolute continuity* are defined as follows.

Definition 6.1. *The* total variation distance *between μ_1 and μ_2 is*

$$d_{TV}(\mu_1, \mu_2) = \sup_{D \in \mathcal{F}} |\mu_1(D) - \mu_2(D)|$$

Definition 6.2. *If $\mu_2(D) = 0$ implies $\mu_1(D) = 0$ for every $D \in \mathcal{F}$, then μ_1 is* absolutely continuous *with respect to μ_2, denoted $\mu_1 \ll \mu_2$.*

Now we are ready to state the main result:

Proposition 19. *Suppose $p, \widetilde{p} \in \Delta(\Theta)$ are absolutely continuous with respect to one another, and define P, $\widetilde{P}$ to be the measures on Ω induced by the respective priors $p, \tilde{p}$, and the family $(P_\theta)_{\theta \in \Theta}$. Then*

$$\lim_{t \to \infty} d_{TV}(P^t(x_1, \ldots, x_t), \widetilde{P}^t(x_1, \ldots, x_t)) = 0 \quad P\text{-almost surely}$$

That is, if two agents hold different prior beliefs about the parameter but agree on the set of measure-0 events, then their conditional beliefs merge in a strong sense: For *all* measurable future events, agents eventually assign similar probabilities.

Example 6.1. To clarify the difference between this result and the one examined in the previous section, consider the problem of learning the unknown bias of a coin, which is parametrized to $p \in [0, 1]$. A coin whose bias is p lands on Heads with probability p and lands on Tails with probability $1 - p$. Two agents have different prior beliefs on $[0, 1]$ and each observe t independent flips of this coin.

Proposition 18 says that the two agents will eventually learn the bias of the coin as t grows large. Proposition 19 says instead: Suppose the two agents have observed t independent flips of the coin; then, their beliefs over all events regarding the future—e.g., that over half of the remaining coin flips will turn up Heads, or that the limiting fraction of Heads realizations is $1/2$—must eventually become close (uniformly across such events).

[2] Blackwell and Dubins (1962) work with the more general notion of "predictive probabilities" P where conditional probabilities can be defined.

6.5 (Expected) Disagreement

We now turn to the impact of information on agents' second-order beliefs—i.e., what they think about what others think. Kartik, Lee and Suen (2021) show that when signals satisfy an MLRP condition, then agents with different beliefs expect information to reduce the extent of disagreement.

Here we assume the set of parameters $\Theta \subseteq \mathbb{R}$ is finite and ordered. Two signals X and $\widetilde{X}$ respectively take values in $\mathcal{X}$ and $\widetilde{\mathcal{X}}$, and we assume that X is Blackwell more informative than $\widetilde{X}$. There are two agents, Ann and Bob, who have common knowledge of the conditional distributions $\{f_{X|\theta}(x \mid \theta)\}_{\theta\in\Theta}$ and $\{f_{\widetilde{X}|\theta}(\widetilde{x} \mid \theta)\}_{\theta\in\Theta}$. But Ann and Bob hold different prior beliefs $f^A_\theta, f^B_\theta \in \Delta(\Theta)$ about the parameter. We use F^A and F^B to denote their perceived joint distributions of $(\theta, X, \widetilde{X})$ (induced by the respective priors and the common knowledge signal distributions), and $\mathbb{E}_A$ and $\mathbb{E}_B$ to denote expectations with respect to these distributions.

Assumption 4. *There is an order $\succ$ on $\mathcal{X}$ and an order $\widetilde{\succ}$ on $\widetilde{\mathcal{X}}$ such that the families $\{f_{X|\theta}(\cdot \mid \theta)\}_{\theta\in\Theta}$ and $\{f_{\widetilde{X}|\theta}(\cdot \mid \theta)\}_{\theta\in\Theta}$ each have MLRP (see Definition 3.2).*

Assumption 5. *Bob's prior f^B_θ likelihood-ratio dominates Ann's prior f^A_θ (see Definition 3.1).*

The agents' prior expectations of the parameter are $\mu_A \equiv \mathbb{E}_A(\theta)$ and $\mu_B \equiv \mathbb{E}_B(\theta)$. We are interested in Ann's prior expectation of Bob's posterior expectation (updated to X), and Bob's prior expectation of Ann's posterior expectation (updated to X), respectively denoted by

$$\mu_{AB}(X) \equiv \mathbb{E}_A[\mathbb{E}_B(\theta \mid X)]$$
$$\mu_{BA}(X) \equiv \mathbb{E}_B[\mathbb{E}_A(\theta \mid X)]$$

Proposition 20. *Suppose Assumptions 4 and 5 are satisfied. If X is Blackwell more informative than $\widetilde{X}$, then*

$$\mu_A \le \mu_{AB}(X) \le \mu_{AB}(\widetilde{X}) \le \mu_B$$

$$\mu_A \le \mu_{BA}(\widetilde{X}) \le \mu_{BA}(X) \le \mu_B$$

That is, Ann expects that a more informative experiment will, in expectation, bring Bob's posterior mean closer to Ann's prior, and vice versa. These are both subjective statements, and indeed only one of Ann and Bob can be correct.

We'll prove this proposition using the following relationships, which are left as an exercise.

Exercise 6.1 (G). *Prove the following statements:*

(a) $F^B_{\theta|X}(\theta \mid X = x)$ *first-order stochastically dominates* $F^A_{\theta|X}(\theta \mid X = x)$ *for every signal realization* $x \in \mathcal{X}$

(b) $F^B_{X|\widetilde{X}}(X \mid \widetilde{X} = \tilde{x})$ *first-order stochastically dominates* $F^A_{X|\widetilde{X}}(X \mid \tilde{X} = \tilde{x})$ *for every signal realization* $\tilde{x} \in \widetilde{\mathcal{X}}$

Proof. Part (a) of Exercise 6.1 implies $\int \theta dF^A_{\theta|X}(\theta \mid x) \leq \int \theta dF^B_{\theta|X}(\theta \mid x)$ for every realization x, so also

$$\int\int \theta dF^A_{\theta|X}(\theta \mid x) dF^A_X(x) \leq \int\int \theta dF^B_{\theta|X}(\theta \mid x) dF^A_X(x). \tag{6.2}$$

By assumption that $\{f_{X|\theta}(\cdot \mid \theta)\}_{\theta\in\Theta}$ has MLRP, the integral $\int \theta dF^B_{\theta|X}(\theta \mid x)$ is an increasing function of x. Moreover, Part (b) of Exercise 6.1 says that F^B_X first-order stochastically dominates F^A_X (taking $\widetilde{X}$ to be any constant signal). Thus

$$\int\int \theta dF^B_{\theta|X}(\theta \mid x) dF^A_X(x) \leq \int\int \theta dF^B_{\theta|X}(\theta \mid x) dF^B_X(x). \tag{6.3}$$

Together, (6.2) and (6.3) imply

$$\int\int \theta dF^A_{\theta|X}(\theta \mid x) dF^A_X(x) \leq \int\int \theta dF^B_{\theta|X}(\theta \mid x) dF^A_X(x) \leq \int\int \theta dF^B_{\theta|X}(\theta \mid x) dF^B_X(x)$$

which is precisely the desired inequality $\mu_A \leq \mu_{AB}(X) \leq \mu_B$. It follows by identical arguments that $\mu_A \leq \mu_{BA}(X) \leq \mu_B$.

To show that $\mu_{AB}(\widetilde{X}) \geq \mu_{AB}(X)$, we use the fact that (since X Blackwell-dominates $\widetilde{X}$) we can generate the two variables in such a way that $\widetilde{X}$ is conditionally independent of θ conditional on X.[3] Then on this probability space

$$\begin{aligned}
\mu_{AB}(\widetilde{X}) &= \mathbb{E}_A\left[\mathbb{E}_B\left(\theta \mid \widetilde{X}\right)\right] \\
&= \mathbb{E}_A\left[\mathbb{E}_B\left(\mathbb{E}_B\left(\theta \mid X, \widetilde{X}\right) \mid \widetilde{X}\right)\right] && \text{by L.I.E.} \\
&= \mathbb{E}_A\left[\mathbb{E}_B\left(\mathbb{E}_B\left(\theta \mid X\right) \mid \widetilde{X}\right)\right] && \text{since } \widetilde{X} \perp\!\!\!\perp \theta \mid X \\
&= \int\int \mathbb{E}_B(\theta \mid x) dF^B_{X|\widetilde{X}}(x \mid \tilde{x}) dF_A(\tilde{x}) \\
&\geq \int\int \mathbb{E}_B(\theta \mid x) dF^A_{X|\widetilde{X}}(x \mid \tilde{x}) dF_A(\tilde{x}) \\
&= \mathbb{E}_A\left[\mathbb{E}_A\left(\mathbb{E}_B\left(\theta \mid X\right) \mid \widetilde{X}\right)\right] \\
&= \mathbb{E}_A\left[\mathbb{E}_B\left(\theta \mid X\right)\right] && \text{by L.I.E.} \\
&= \mu_{AB}(X)
\end{aligned}$$

where the crucial inequality follows by observing that $\mathbb{E}_B(\theta \mid x)$ is an increasing function of x (by Assumption 4) while $F^B_{X|\widetilde{X}}(\cdot \mid \tilde{x})$ first-order stochastically dominates $F^A_{X|\widetilde{X}}(\cdot \mid \tilde{x})$ for every realization of $\tilde{x}$ (by Part (b) of Exercise 6.1).

Since the previous arguments apply to show also that $\mu_{AB}(\widetilde{X}) \leq \mu_B$, we are done. ∎

[3]See Remark 4.1 for further detail. Note also that the correlation between X and $\widetilde{X}$ is irrelevant for the comparison of $\mu_{AB}(X)$ and $\mu_{AB}(\widetilde{X})$.

6.6 Common Learning

Suppose Assumption 3 (Identifiability) holds, so that agents eventually learn the true parameter. Does this imply that agents will eventually have *common knowledge* of the true parameter? Cripps et al. (2008) adapt Monderer and Samet (1989)'s definition of common q-belief for the present learning environment, and show that individual learning does imply common learning when the set of signal realizations is finite, but that this implication may otherwise fail.

In what follows recall that each state $\omega \in \Omega = \Theta \times \mathcal{X}^\infty$ describes both the value of the parameter and the infinite sequence of signal profiles. As before, P_θ denotes the measure on $\mathcal{X}^\infty$ conditional on parameter θ, and again assume that Θ is finite. There are two agents $i = 1, 2$, and (different from the previous sections) we decompose $\mathcal{X} = \mathcal{X}^1 \times \mathcal{X}^2$ where $\mathcal{X}^i$ denotes the set of agent i signal realizations. Each agent privately observes their own signal each period. We use $h_{it}(\omega) = (x_1^i(\omega), \ldots, x_t^i(\omega))$ for agent i's history at time t when ω is the realized state, and $\mathcal{H}_{it}$ to denote the filtration induced by agent i's histories.

Definition 6.3. *For any $q \in [0,1]$ and (measurable) event F, agent i* q-believes *in F at time t on*

$$B_{it}^q(F) = \{\omega \in \Omega \mid P(F \mid h_{it}(\omega)) \geq q\}$$

Definition 6.4. *For any $q \in [0,1]$, there is* common q-belief *in F at time t on*

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

where $B_t^q(F) = B_{1t}^q(F) \bigcap B_{2t}^q(F)$.

Definition 6.5 (Individual Learning). *Agent i* learns *θ if for each $q \in (0,1)$ there exists $T < \infty$ such that*

$$P_\theta(B_{it}^q(\{\theta\} \times \mathcal{X}^\infty)) > q \quad \forall t > T$$

Equivalently: $\lim_{t\to\infty} P_\theta(B_{it}^q(\{\theta\} \times \mathcal{X}^\infty)) = 1$ *for all $q \in (0,1)$. Agent i* individually learns *if the agent learns each $\theta \in \Theta$.*

Definition 6.6 (Common Learning). *Agents* commonly learn *θ if for each $q \in (0,1)$ there exists $T < \infty$ such that*

$$P_\theta(C_t^q(\{\theta\} \times \mathcal{X}^\infty)) > q \quad \forall t > T$$

Equivalently: $\lim_{t\to\infty} P_\theta(C_t^q(\{\theta\} \times \mathcal{X}^\infty)) = 1$ *for all $q \in (0,1)$. Agents* commonly learn *if they commonly learn each $\theta \in \Theta$.*

Clearly if signals are perfectly correlated (or public), so that $P(\theta \mid \mathcal{H}_{1t}) = P(\theta \mid \mathcal{H}_{2t})$ for all θ and t, then individual learning implies common learning. This result also holds at the other extreme of independent signals.

Proposition 21. *Suppose agents individually learn, and their signals are conditionally independent given the parameter. That is, there exist families* $(P^i_\theta)_{\theta\in\Theta}$*, with each* $P^i_\theta \in \Delta(\mathcal{X}^i)$*, such that* $P_\theta(A \times B) = P^1_\theta(A)P^2_\theta(B)$ *for each* $\theta \in \Theta$ *and measurable* $A \subseteq \mathcal{X}^1$*,* $B \subseteq \mathcal{X}^2$*. Then, agents commonly learn.*

Cripps et al. (2008) proves this proposition using a result from Monderer and Samet (1989) (adapted to the present learning context).

Lemma 2. *Agents commonly learn if and only if for every* $\theta \in \Theta$ *and* $q \in (0, 1)$*, there is a sequence of events* F_t *and a period* T *such that for all* $t > T$*,*

(a) $F_t \subseteq B^q_t(\theta)$ *("*θ *is q-believed on* F_t *at time t")*

(b) $P_\theta(F_t) > q$ *("probability of* F_t *is sufficiently high")*

(c) $F_t \subseteq B^q_{it}(F_t)$ *for* $i = 1, 2$ *("*F_t *is evident q-belief at time t")*

We'll now prove Proposition 21.

Proof. Henceforth write $\{\theta\}$ for the event $\{\theta\} \times \mathcal{X}^\infty$. Define $F_t = \{\theta\} \cap B^{\sqrt{q}}_t(\theta)$ to be the set of states at which θ is true and both agents $\sqrt{q}$-believe it. We'll verify that the conditions of Lemma 2 hold for the sequence of events $(F_t)^\infty_{t=1}$, from which Proposition 21 follows.

First observe that

$$\begin{aligned} F_t &\subseteq B^{\sqrt{q}}_t(\theta) && \text{by definition of } F_t \\ &\subseteq B^q_t(\theta) && \text{since } q < \sqrt{q} \end{aligned}$$

yielding Part (a) of Lemma 2. Part (b) holds since individual learning implies that there exists $T < \infty$ such that for both agents $i = 1, 2$,

$$P_\theta\left(B^{\sqrt{q}}_{it}(\theta)\right) > \sqrt{q} \quad \forall t > T$$

and thus

$$P_\theta(F_t) = P_\theta\left(B^{\sqrt{q}}_{1t}(\theta)\right) P_\theta\left(B^{\sqrt{q}}_{2t}(\theta)\right) > q \quad \forall t > T$$

from the assumption of conditional independence.

It remains to show Part (c). First rewrite the set $B^q_{1t}(F_t)$ as follows:

$$\begin{aligned} B^q_{1t}(F_t) &= \{\omega \mid \mathbb{E}\left[\mathbb{1}_{F_t} \mid \mathcal{H}_{1t}\right] \geq q)\} && \text{by definition of } B^q_{1t} \\ &= \left\{\omega \mid \mathbb{E}\left[\mathbb{1}_{B^{\sqrt{q}}_{1t}(\theta)} \mathbb{1}_{B^{\sqrt{q}}_{2t}(\theta)\cap\{\theta\}} \mid \mathcal{H}_{1t}\right] \geq q\right\} && \text{by definition of } F_t \\ &= \left\{\omega \mid \mathbb{1}_{B^{\sqrt{q}}_{1t}(\theta)} \mathbb{E}\left[\mathbb{1}_{B^{\sqrt{q}}_{2t}(\theta)\cap\{\theta\}} \mid \mathcal{H}_{1t}\right] \geq q\right\} && \text{since } B^{\sqrt{q}}_{1t}(\theta) \in \mathcal{H}_{1t} \\ &= B^{\sqrt{q}}_{1t}(\theta) \cap B^q_{1t}\left(B^{\sqrt{q}}_{2t}(\theta) \cap \{\theta\}\right) \end{aligned}$$

By definition we have that $F_t \subseteq B^{\sqrt{q}}_{1t}(\theta)$. As above, individual learning implies existence of T sufficiently large that $P_\theta\left(B^{\sqrt{q}}_{2t}(\theta)\right) > \sqrt{q}$ for all $t > T$. Since signals are conditionally independent, agent 1's history is uninformative about agent 2's history, implying that

$$P_\theta\left(B_{2t}^{\sqrt{q}}(\theta) \mid \mathcal{H}_{1t}\right) \geq \sqrt{q} \tag{6.4}$$

holds uniformly across agent 1 histories (for all $t > T$). So on F_t (for $t > T$) we have

$$P(B_{2t}^{\sqrt{q}}(\theta) \cap \{\theta\} \mid \mathcal{H}_{1t}) = \underbrace{P_\theta(B_{2t}^{\sqrt{q}}(\theta) \mid \mathcal{H}_{1t})}_{>\sqrt{q}\text{ by (6.4)}} \underbrace{P(\theta \mid \mathcal{H}_{1t})}_{>\sqrt{q}\text{ since } F_t \subseteq B_{1t}^{\sqrt{q}}(\theta)} > q.$$

Apply Lemma 2 and we are done. ■

Remark 6.3. This proof extends for arbitrary finite numbers of agents, setting $F_t = \{\theta\} \cap B_t^{\sqrt[n]{q}}(\theta)$.

Although common learning is implied by individual learning when agents have either perfect information or no information about the other agent's history, intermediate cases of correlation can break this result.

Example 6.2. (Twist on Rubinstein (1989)'s email game.) The unknown parameter is $\theta \in \{\theta', \theta''\}$, where $0 \leq \theta' < \theta'' \leq 1$. Suppose that every period a signal profile is independently drawn according to:

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

This signal structure generalizes the information structure in the email game from Section 1.3, where $\theta = 1$ corresponds to state a in the email game and $\theta = 0$ corresponds to state b.

Agents observe repeated independent realizations of the signal. Will they commonly learn the game parameter? When θ is restricted to values 0 and 1 (as per Rubinstein (1989)'s email game), the answer is yes.

Exercise 6.2 (G). *Prove that common learning occurs if $\theta \in \{\theta', \theta''\} \equiv \{0, 1\}$.*

But common learning fails whenever $0 < \theta' < \theta'' < 1$ as agents cannot commonly learn θ'', the parameter placing more weight on the lower signal realizations. Intuitively, when 1 sees the signal k, then he believes with some probability (that can be uniformly lower bounded across histories) that 2 has also observed at least k. And if 2 observes k, then he believes with some probability (that again can be uniformly lower bounded) that 1 observed $k+1$. Since the number of signal realizations is infinite, there is unbounded contagion upwards: The agent always believes with some probability that the other agent believes with some probability that he has observed... such a large signal that he believes that the state is (very likely to be) θ'. And thus we cannot establish common q-belief of θ'' for large q.

The main result in Cripps et al. (2008) establishes that infinite signal realizations are critical to the previous counterexample. When the number of signal realizations is finite, then individual learning always implies common learning.

Assumption 6 (Finite Signal Sets)**.** $|\mathcal{X}^1|, |\mathcal{X}^2| < \infty$

Proposition 22. *If Assumption 6 is satisfied, then individual learning implies common learning.*

A brief idea of the proof follows. Define $\pi^\theta(ij)$ to be the probability of realization $(x_t^1, x_t^2) = (i, j)$ when the parameter is θ, and define

$$\phi^\theta(i) = \sum_{j \in \mathcal{X}^2} \pi^\theta(ij)$$

to be the marginal probability of signal i, with $\phi^\theta \equiv (\phi^\theta(i))_{i \in \mathcal{X}^1}$. Likewise define

$$\psi^\theta(j) = \sum_{i \in \mathcal{X}^1} \pi^\theta(ij)$$

to be the marginal probability of signal j, with $\psi^\theta \equiv (\psi^\theta(j))_{j \in \mathcal{X}^2}$. Then (by the results in Section 6.3), individual learning follows whenever $\phi^\theta \neq \phi^{\theta'}$ and $\psi^\theta \neq \psi^{\theta'}$ for every $\theta \neq \theta'$.

Define $\hat{\phi}_t$ to be the empirical frequency of agent 1 signals and $\hat{\psi}_t$ to be the empirical frequency of agent 2 signals. Under the assumption of individual learning, empirical frequencies must converge to the theoretical frequencies, i.e., for each parameter θ, $\hat{\phi}_t \to \phi^\theta$ and $\hat{\psi}_t \to \phi^\theta$ P_θ-almost surely. Thus each agent eventually assigns a high probability to true θ.

The crucial next step is establishing that when agent 1 assigns a high probability to θ, he believes that agent 2 does as well (and vice versa). To see why this might be the case, let M_1^θ be the $|\mathcal{X}^1| \times |\mathcal{X}^2|$ matrix whose (i,j)-th entry is $\frac{\pi^\theta(ij)}{\phi^\theta(i)}$, i.e. the conditional probability (under θ) that agent 2 observes j given that agent 1 observed i, and define M_2^θ analogously. Then $\hat{\phi}_t M_1^\theta$ is agent 1's expectation of agent 2's realized frequencies (conditional on θ), and $\hat{\phi}_t M_1^\theta M_2^\theta$ is agent 1's expectation of agent 2's expectation of agent 1's realized frequencies (again conditional on θ). Observe (by algebra) that

$$\phi^\theta M_1^\theta = \psi^\theta$$
$$\psi^\theta M_2^\theta = \phi^\theta$$

so $\phi^\theta M_1^\theta M_2^\theta = \phi^\theta$. Indeed the matrix $M_{12}^\theta \equiv M_1^\theta M_2^\theta$ is a Markov transition matrix on $\mathcal{X}^1$ with stationary distribution ϕ^θ, and it is moreover a contraction mapping on $\Delta(\mathcal{X}^1)$. These properties together imply that the higher order beliefs cannot run away from the agent's first-order belief as they did in Example 6.2.

6.7 Additional Exercises

EXERCISE 6.3 (G*). *Let $\theta \sim \mathcal{N}(0,1)$ be an unknown parameter. Each agent $i = 1,2$ observes n signals $X_1^i, \ldots, X_n^i$ where each*

$$X_m^i = \theta + \varepsilon_m^i$$

with $\varepsilon_m^i \sim \mathcal{N}(0,1)$ independent of θ, independent across agents, and independent across signals. Suppose that the true value of θ is strictly positive, and let E_p be the event that the two agents have common p-belief that θ is positive, where $p > 1/2$. What is the probability of E_p under the actual data-generating process?