SUPPLEMENT TO "LEARNING WITH HETEROGENEOUS MISSPECIFIED MODELS: CHARACTERIZATION AND ROBUSTNESS"
(*Econometrica*, Vol. 89, No. 6, November 2021, 3025–3077)

J. AISLINN BOHREN
Department of Economics, University of Pennsylvania

DANIEL N. HAUSER
Department of Economics, Aalto University School of Business and Helsinki Graduate School of Economics

APPENDIX B: DERIVATION OF EXAMPLES FROM SECTION 3

B.1. *Example 1: Partisan Bias*

SIGNALS AND PREFERENCES ARE ALIGNED (Assumptions 1 and 2) since both types have the same subjective signal distributions and preferences. The autarkic type θ_2 plays both actions with positive probability and the social type θ_1 places positive probability on θ_2, which establishes that Assumption 3 holds. Assumption 4 is redundant in a binary action decision problem, since Assumption 3 guarantees that the social type believes that the autarkic type plays both actions with positive probability. For technical convenience, we assume that the signal distributions are continuous and symmetric, $F^R(s) = 1 - F^L(1-s)$.

From the action probabilities derived in Section 3.1, at likelihood ratio λ_1, type θ_1 believes action L occurs with probability $\hat{\psi}_1(L|\omega, \lambda_1) = \pi(\theta_1)F^\omega(1/(1+\lambda_1)) + \pi(\theta_2)F^\omega(0.5)$, whereas the true probability of action L is $\psi(L|\omega, \lambda_1) = \pi(\theta_1)F^\omega((1/(1+\lambda_1))^{1/\nu}) + \pi(\theta_2)F^\omega(0.5^{1/\nu})$. The construction of $\gamma_1(L, 0)$ in Section 3.3 follows from evaluating these expressions at $\lambda_1 = 0$. Similarly, the construction of $\gamma_1(L, \infty)$ follows from evaluating these expressions at $\lambda_1 = \infty$,

$$\pi(\theta_2)F^L\big(0.5^{1/\nu}\big) \log \frac{F^R(0.5)}{F^L(0.5)}$$

$$+ \big(\pi(\theta_1) + \pi(\theta_2)\big(1 - F^L\big(0.5^{1/\nu}\big)\big)\big) \log \frac{\pi(\theta_1) + \pi(\theta_2)\big(1 - F^R(0.5)\big)}{\pi(\theta_1) + \pi(\theta_2)\big(1 - F^L(0.5)\big)}.$$

We next characterize how $\Lambda(\omega)$ depends on ν. We write $\gamma_1(\omega, \boldsymbol{\lambda}; \nu)$ and $\Lambda(\omega; \nu)$ to make this dependence on ν explicit. To simplify notation, define $\alpha_\nu \equiv F^L(0.5^{1/\nu})$ as the probability that type θ_2 chooses an L action in state L and $\pi_A \equiv \pi(\theta_2)$ as the probability of the autarkic type. By symmetry, $F^R(0.5) = 1 - F^L(0.5) = 1 - \alpha_1$, and by definition of a probability measure, $\pi(\theta_1) = 1 - \pi_A$. Also note that F^L strictly increasing implies that α_ν is strictly increasing in ν, symmetry implies that $\alpha_1 > 1/2$, and F^L continuous implies α_ν is continuous in ν.

First consider $\omega = L$. To determine whether incorrect learning arises, that is, whether $\infty \in \Lambda(L; \nu)$, we need to determine the sign of

$$\gamma_1(L, \infty; \nu) = \pi_A \alpha_\nu \log \frac{1 - \alpha_1}{\alpha_1} + (1 - \pi_A \alpha_\nu) \log \frac{1 - \pi_A(1 - \alpha_1)}{1 - \pi_A \alpha_1}.$$

J. Aislinn Bohren: abohren@sas.upenn.edu
Daniel N. Hauser: daniel.hauser@aalto.fi

Since $\alpha_1 > 1/2$, the update from an L action is negative, $\log \frac{1-\alpha_1}{\alpha_1} < 0$, and the update from an R action is positive, $\log \frac{1-\pi_A(1-\alpha_1)}{1-\pi_A\alpha_1} > 0$. Note both terms are independent of ν. Since α_ν is strictly increasing in ν, the probability of an L action, $\pi_A\alpha_\nu$, is strictly increasing in ν and the probability of an R action, $1 - \pi_A\alpha_\nu$, is strictly decreasing in ν. Therefore, $\gamma_1(L, \infty; \nu)$ is strictly decreasing in ν. At $\nu = 1$, $\gamma_1(L, \infty; 1) < 0$ by the concavity of the log operator. At $\nu = 0$, $\alpha_0 = 0$ and therefore $\gamma_1(L, \infty; 0) = \log \frac{1-\pi_A(1-\alpha_1)}{1-\pi_A\alpha_1} > 0$. Given $\gamma_1(L, \infty; \nu)$ is continuous in ν, there exists a cutoff $\nu_1 \in (0, 1)$ such that for $\nu < \nu_1$, $\gamma_1(L, \infty; \nu) > 0$ and $\infty \in \Lambda(L; \nu)$, and for $\nu > \nu_1$, $\gamma_1(L, \infty; \nu) < 0$ and $\infty \notin \Lambda(L; \nu)$.

To determine whether correct learning arises, that is, whether $0 \in \Lambda(L; \nu)$, we need to determine the sign of

$$\gamma_1(L, 0; \nu) = \left(1 - \pi_A(1 - \alpha_\nu)\right) \log \frac{1 - \pi_A\alpha_1}{1 - \pi_A(1 - \alpha_1)} + \pi_A(1 - \alpha_\nu) \log \frac{\alpha_1}{1 - \alpha_1}.$$

As in the previous case, the update from an L action is negative and the probability of an L action is strictly increasing in ν, while the update from an R action is positive and the probability of an R action is strictly decreasing in ν. Therefore, $\gamma_1(L, 0; \nu)$ is strictly decreasing in ν. At $\nu = 1$, $\gamma_1(L, 0; 1) < 0$ by the concavity of the log operator. At $\nu = 0$, $\alpha_0 = 0$ and therefore,

$$\gamma_1(L, 0; 0) = (1 - \pi_A) \log \frac{1 - \pi_A\alpha_1}{1 - \pi_A(1 - \alpha_1)} + \pi_A \log \frac{\alpha_1}{1 - \alpha_1}$$

$$\geq (1 - \pi_A\alpha_1) \log \frac{1 - \pi_A\alpha_1}{1 - \pi_A(1 - \alpha_1)} + \pi_A\alpha_1 \log \frac{\alpha_1}{1 - \alpha_1}$$

$$= \gamma_1(R, 0; 1) > 0.$$

Given $\gamma_1(L, 0; \nu)$ is continuous in ν, there exists a cutoff $\nu_2 \in (0, 1)$ such that for $\nu < \nu_2$, $\gamma_1(L, 0; \nu) > 0$ and $0 \notin \Lambda(L; \nu)$, and for $\nu > \nu_2$, $\gamma_1(L, 0; \nu) < 0$ and $0 \in \Lambda(L; \nu)$.

Finally we show that $\nu_1 < \nu_2$. Note

$$\gamma_1(L, \infty; \nu) - \gamma_1(L, \infty; 1) = \pi_A(\alpha_\nu - \alpha_1) \left(\log \frac{1 - \alpha_1}{\alpha_1} - \log \frac{1 - \pi_A + \pi_A\alpha_1}{1 - \pi_A\alpha_1} \right),$$

and by the symmetry of the signal distributions, $\gamma_1(L, 0; \nu) - \gamma_1(L, 0; 1) = \gamma_1(L, \infty; \nu) - \gamma_1(L, \infty; 1)$. Moreover, $\gamma_1(L, 0; 1) - \gamma_1(L, \infty; 1)$ is zero at $\pi_A = 0$ and $\pi_A = 1$, and concave in π_A since the second derivative is

$$\frac{(1 - 2\alpha_1)\pi_A}{\left(\pi_A(1 - \alpha_1) + (1 - \pi_A)\right)^2 (\pi_A\alpha_1 + 1 - \pi_A)^2} \leq 0.$$

Therefore, $0 \notin \Lambda(\omega; \nu)$ before $\infty \in \Lambda(\omega; \nu)$. This implies that $\Lambda(L; \nu) = \{\infty\}$ for $\nu \in (0, \nu_1)$, $\Lambda(L; \nu) = \emptyset$ for $\nu \in (\nu_1, \nu_2)$, and $\Lambda(L; \nu) = \{0\}$ for $\nu \in (\nu_2, 1]$.

Next consider $\omega = R$. Then $\gamma(R, \infty; 1) > 0$ and $\gamma(R, 0; 1) > 0$, since only correct learning can occur at $\nu = 1$. The only change in the above expressions is that now the true probabilities of each action are taken with respect to state R rather than state L. Therefore, the comparative statics are similar to the comparative statics in state L: $\gamma_1(R, 0; \nu)$ and $\gamma_1(R, \infty; \nu)$ are decreasing in ν. Therefore, $\gamma_1(R, 0; \nu) > 0$ implies $0 \notin \Lambda(R; \nu)$ for

all $\nu \in (0, 1]$. Similarly, $\gamma_1(R, \infty; \nu) > 0$ implies $\infty \in \Lambda(R; \nu)$ for all $\nu \in (0, 1]$. Therefore, $\Lambda(R; \nu) = \{\infty\}$ for all $\nu \in (0, 1]$.

When there is a single social type, mixed learning and disagreement are trivially not possible. By Theorem 4, the characterization of the locally stable set fully determines asymptotic learning outcomes. This leads to the following proposition, the proof of which follows from the construction of $\Lambda(\omega; \nu)$ above.

PROPOSITION 5—Partisan Bias: *When $\omega = L$, there exist unique cutoffs $0 < \nu_1 < \nu_2 < 1$ such that (i) if $\nu \in (\nu_2, 1]$, then almost surely learning is correct; (ii) if $\nu \in (\nu_1, \nu_2)$, then almost surely learning is cyclical; (iii) if $\nu \in (0, \nu_1)$, then almost surely learning is incorrect. When $\omega = R$, almost surely learning is correct.*

B.2. *Example 2: Partisan Bias and Unawareness*

We construct this variation by adding two types to the setting considered in Example 1. Types θ_1 and θ_2 are partisan types with the same signal misspecification and preferences as in Example 1. Types θ_3 and θ_4 are non-partisan types that correctly interpret signals, $\hat{F}_3^\omega(s) = \hat{F}_4^\omega(s) = F^\omega(s)$; θ_3 is a social type while θ_4 is an autarkic type.[1] Both types have the same preferences as θ_1 and θ_2, that is, $u_i(a, \omega) = \mathbb{1}_{a=\omega}$. Assume that an equal and positive share of partisan and non-partisan types are autarkic, $\pi(\theta_2)/(\pi(\theta_1) + \pi(\theta_2)) = \pi(\theta_4)/(\pi(\theta_3) + \pi(\theta_4)) \in (0, 1)$. Both social types have correct beliefs about the share of autarkic types, but partisan θ_1 believes all agents are partisan, $\hat{\pi}_1(\theta_1) = \pi(\theta_1) + \pi(\theta_3)$ and $\hat{\pi}_1(\theta_2) = \pi(\theta_2) + \pi(\theta_4)$, and analogously, non-partisan θ_3 believes that all agents are non-partisan. Let $q \equiv \pi(\theta_3) + \pi(\theta_4)$ denote the share of non-partisan types and $\pi_A \equiv \pi(\theta_2) + \pi(\theta_4)$ denote the share of autarkic types. To close the model, assume that the signal distributions are continuous and symmetric, $F^R(s) = 1 - F^L(1 - s)$ with support $\mathcal{S} = [0, 1]$, and $p_0 = 1/2$. Signals are aligned since partisan types order signal realizations in the same way as non-partisan types, that is, s^ν is increasing in s (Assumption 1).

The true action probabilities for partisan types θ_1 and θ_2 are identical to those derived in Section 3.1 for Example 1, as are θ_1's subjective action probabilities for each type. A non-partisan type $\theta_i \in \{\theta_3, \theta_4\}$ who has likelihood ratio λ and observes signal realization s updates to belief $\frac{p_i(\lambda,s)}{1-p_i(\lambda,s)} = \lambda(\frac{s}{1-s})$. It chooses action L if this belief is less than 1, which is equivalent to $s < 1/(1 + \lambda) = \bar{s}_{i,1}(\lambda)$. At likelihood ratio λ_3, type θ_3 chooses L with probability $F^\omega(1/(1 + \lambda_3))$. Type θ_4 is autarkic. Therefore, its likelihood ratio is constant at $\lambda_4 = 1$ and it chooses action L with probability $F^\omega(0.5)$. Type θ_3 has correct beliefs about the probability that θ_3 and θ_4 choose action L.

We use these subjective and true action probabilities for each type to construct $\hat{\psi}_1$, $\hat{\psi}_3$, and ψ. Partisan type θ_1 is now also misspecified about the type distribution, since it does not account for the non-partisan types. It believes action L occurs with probability $\hat{\psi}_1(L|\omega, \boldsymbol{\lambda}) = (1 - \pi_A)F^\omega(1/(1 + \lambda_1)) + \pi_A F^\omega(0.5)$. This misspecification about the type distribution leads the partisan type to underestimate the range of signal realizations for which other agents choose action L, while its signal misspecification causes it to overestimate the probability of these signal realizations. The latter effect dominates, and θ_1 overestimates the frequency of action L. Non-partisan type θ_3 has a correctly specified model of the signal distribution and believes that other agents do as well, since

---

[1] In a slight abuse of our previous notation, we maintain θ_2 as the partisan autarkic type for consistency with Example 1, which violates our convention that the first k types are the social types.

it does not account for the partisan types. It believes action L occurs with probability $\hat{\psi}_3(L|\omega, \boldsymbol{\lambda}) = (1 - \pi_A)F^\omega(1/(1+\lambda_3)) + \pi_A F^\omega(0.5)$. This type misspecification leads the non-partisan type to believe that other agents are choosing L for a larger range of signal realizations than is actually the case, which leads it to overestimate the frequency of L actions. The true probability of action L is

$$\psi(L|\omega, \boldsymbol{\lambda}) = (1-q)\big((1 - \pi_A)F^\omega\big((1/(1+\lambda_1))^{1/\nu}\big) + \pi_A F^\omega\big(0.5^{1/\nu}\big)\big)$$
$$+ q\big((1 - \pi_A)F^\omega\big(1/(1+\lambda_3)\big) + \pi_A F^\omega(0.5)\big).$$

Although the partisan and non-partisan social types have different models of the world, their models collapse to the same subjective probability of each action when they have the same current belief: for any $\boldsymbol{\lambda}$ with $\lambda_1 = \lambda_3$, $\hat{\psi}_1(L|\omega, \boldsymbol{\lambda}) = \hat{\psi}_3(L|\omega, \boldsymbol{\lambda})$. Therefore, these types update their likelihood ratios in the same way following each action. For different reasons, their beliefs both move too much towards state R following R actions and too little towards state L following L actions. This implies that when there is a common prior, after any history h_t, beliefs are equal, $\lambda_{1,t} = \lambda_{3,t}$.[2]

Given that the two likelihood ratios move in unison, we can consider the partisan and non-partisan social types as a single type to characterize asymptotic learning outcomes. Disagreement and mixed learning do not arise, since it is not possible to separate beliefs. Global stability immediately follows from local stability for the two agreement outcomes. Therefore, determining the set of parameters (ν, q) for which each agreement outcome is locally stable fully characterizes asymptotic learning outcomes. This leads to the following proposition.

PROPOSITION 6—Partisan Bias: *When $\omega = L$, there exist unique cutoffs $q_1 \in (0, 1)$ and $q_2 \in (q_1, 1)$ such that*:
  (i) *For $q < q_1$, there exist unique cutoffs $0 < \nu_1(q) < \nu_2(q) < 1$ such that if $\nu > \nu_2(q)$, then almost surely learning is correct; if $\nu \in (\nu_1(q), \nu_2(q))$, then almost surely learning is cyclical; and if $\nu < \nu_1(q)$, then almost surely learning is incorrect.*
  (ii) *For $q \in (q_1, q_2)$, there exists a unique cutoff $0 < \nu_2(q) < 1$ such that if $\nu > \nu_2(q)$, then almost surely learning is correct, and if $\nu < \nu_2(q)$, then almost surely learning is cyclical.*
  (iii) *For $q > q_2$, almost surely learning is correct.*
  When $\omega = R$, almost surely learning is correct.

PROOF: The construction of the locally stable set is similar to Example 1. To simplify notation, define $\alpha_\nu \equiv F^L(0.5^{1/\nu})$ as the probability that type θ_2 chooses action L in state L. Given this notation, type θ_4 chooses action L in state L with probability α_1. As in Example 1, $F^R(0.5) = 1 - F^L(0.5) = 1 - \alpha_1$, α_ν is strictly increasing in ν, and $\alpha_1 > 1/2$. We characterize how $\Lambda(\omega)$ depends on ν and q. We write $\gamma_1(\omega, \boldsymbol{\lambda}; \nu, q)$, $\gamma_3(\omega, \boldsymbol{\lambda}; \nu, q)$, and $\Lambda(\omega; \nu, q)$ to make this dependence explicit. Since beliefs move in unison, $\gamma_3(\omega, \boldsymbol{\lambda}; \nu, q) = \gamma_1(\omega, \boldsymbol{\lambda}; \nu, q)$, and therefore, we can focus on characterizing $\gamma_1(\omega, \boldsymbol{\lambda}; \nu, q)$ at $(0, 0)$ and (∞, ∞).

---

[2]Partisan and non-partisan types with the same likelihood ratio may choose different actions following a given signal realization s, as they have different signal cutoffs.

To determine whether $(\infty, \infty) \in \Lambda(L; \nu, q)$, we need to determine the sign of

$$\gamma_1\big(L, (\infty, \infty); \nu, q\big) = \psi\big(L|L, (\infty, \infty); \nu, q\big) \log \frac{1 - \alpha_1}{\alpha_1}$$

$$+ \psi\big(R|L, (\infty, \infty); \nu, q\big) \log \frac{1 - \pi_A(1 - \alpha_1)}{1 - \pi_A \alpha_1},$$

where $\psi(L|L, (\infty, \infty); \nu, q) \equiv \pi_A((1-q)\alpha_\nu + q\alpha_1)$ and $\psi(R|L, (\infty, \infty); \nu, q) \equiv \pi_A((1-q)(1-\alpha_\nu) + q(1-\alpha_1)) + 1 - \pi_A$. Since $\alpha_1 > 1/2$, the update from an L action is negative, $\log \frac{1-\alpha_1}{\alpha_1} < 0$, and the update from an R action is positive, $\log \frac{1-\pi_A(1-\alpha_1)}{1-\pi_A\alpha_1} > 0$. Note both terms are independent of ν and q. Since α_ν is strictly increasing in ν, the probability of an L action, $\psi(L|L, (\infty, \infty); \nu, q)$, is strictly increasing in ν and q, and the probability of an R action, $\psi(R|L, (\infty, \infty); \nu, q)$, is strictly decreasing in ν and q. Therefore, $\gamma_1(L, (\infty, \infty); \nu, q)$ is strictly decreasing in ν and q. At $\nu = 1$, both partisan and non-partisan types are identical, so $\psi(L|L, (\infty, \infty); 1, q) = \pi_A\alpha_1$ and $\psi(R|L, (\infty, \infty); 1, q) = \pi_A(1 - \alpha_1) + 1 - \pi_A$. Therefore, for any $q \in [0, 1]$, $\gamma_1(L, (\infty, \infty); 1, q) < 0$ by the concavity of the log operator. Similarly, at $q = 1$, for any $\nu \in [0, 1]$, $\gamma_1(L, (\infty, \infty); \nu, 1) < 0$ by the concavity of the log operator. At $\nu = 0$, θ_2 chooses action R for all signal realizations, that is, $\alpha_0 = 0$. Therefore, at $q = 0$, $\psi(L|L, (\infty, \infty); 0, 0) = 0$ and $\gamma_1(L, (\infty, \infty); 0, 0) = \log \frac{1-\pi_A(1-\alpha_1)}{1-\pi_A\alpha_1} > 0$. This establishes that there exists a cutoff $q_1 \in (0, 1)$ such that for $q < q_1$, there exists a cutoff $\nu_1(q) \in (0, 1)$ such that for $\nu < \nu_1(q)$, $\gamma_1(L, (\infty, \infty); \nu, q) > 0$ and $(\infty, \infty) \in \Lambda(L; \nu, q)$, and for $\nu > \nu_1(q)$, $\gamma_1(L, (\infty, \infty); \nu, q) < 0$ and $(\infty, \infty) \notin \Lambda(L; \nu, q)$. For $q > q_1$, $\gamma_1(L, (\infty, \infty); \nu, q) < 0$ and $(\infty, \infty) \notin \Lambda(L; \nu, q)$.

To determine whether $(0, 0) \in \Lambda(L; \nu, q)$, we need to determine the sign of

$$\gamma_1\big(L, (0, 0); \nu, q\big) = \psi\big(L|L, (0, 0); \nu, q\big) \log \frac{1 - \pi_A \alpha_1}{\pi_A \alpha_1 + 1 - \pi_A}$$

$$+ \psi\big(R|L, (0, 0); \nu, q\big) \log \frac{\alpha_1}{1 - \alpha_1},$$

where $\psi(L|L, (0, 0); \nu, q) \equiv \pi_A((1-q)\alpha_\nu + q\alpha_1) + 1 - \pi_A$ and $\psi(R|L, (0, 0); \nu, q) \equiv \pi_A((1-q)(1-\alpha_\nu) + q(1-\alpha_1))$. As in the previous case, the update from an L action is negative and the probability of an L action is strictly increasing in ν and q, while the update from an R action is positive and the probability of an R action is strictly decreasing in ν and q. Therefore, $\gamma_1(L, (0, 0); \nu, q)$ is strictly decreasing in ν and q. By similar reasoning to the case of (∞, ∞), at $\nu = 1$, $\gamma_1(L, (0, 0); 1, q) < 0$ for all $q \in [0, 1]$, and at $q = 1$, $\gamma_1(L, (0, 0); \nu, 1) < 0$ for all $\nu \in [0, 1]$ by the concavity of the log operator. At $\nu = 0$ and $q = 0$, $\psi(L|L, (0, 0); 0, 0) = 1 - \pi_A$ since $\alpha_0 = 0$. As in Example 1, $\gamma_1(L, (0, 0); 0, 0) > 0$. This establishes that there exists a cutoff $q_2 \in (0, 1)$ such that for $q < q_2$, there exists a cutoff $\nu_2(q)$ such that for $\nu < \nu_2(q)$, $\gamma_1(L, (0, 0); \nu, q) > 0$ and $(0, 0) \notin \Lambda(L; \nu, q)$, and for $\nu > \nu_2(q)$, $\gamma_1(L, (0, 0); \nu, q) < 0$ and $(0, 0) \in \Lambda(L; \nu, q)$. For $q > q_2$, $\gamma_1(L, (0, 0); \nu, q) < 0$ and $(0, 0) \in \Lambda(L; \nu, q)$.

To show that $q_1 < q_2$ and $\nu_1(q) < \nu_2(q)$ for all $q < q_1$, note $\gamma_1(L, (\infty, \infty); \nu, q) - \gamma_1(L, (\infty, \infty); 1, q)$ is equal to

$$\pi_A(1-q)(\alpha_\nu - \alpha_1)\left(\log \frac{1 - \pi_A\alpha_1}{\pi_A\alpha_1 + 1 - \pi_A} - \log \frac{\alpha_1}{1 - \alpha_1}\right),$$

and by the symmetry of the signal distributions, $\gamma_1(L, (0, 0); \nu, q) - \gamma_1(L, (0, 0); 1, q) = \gamma_1(L, (\infty, \infty); \nu, q) - \gamma_1(L, (\infty, \infty); 1, q)$. Moreover, $\gamma_1(L, (0, 0); 1, q) - \gamma_1(L, (\infty, \infty); 1, q)$ is 0 at $\pi_A = 0$, 0 at $\pi_A = 1$, and concave in π_A since the second derivative is

$$\frac{\pi_A(1 - 4q + 4q^2)(2 - 2\alpha_1 - 1)}{\left(\pi_A(1 - \alpha_1) + 1 - \pi_A\right)^2 (\pi_A \alpha_1 + 1 - \pi_A)^2} \le 0.$$

Therefore, $(0, 0) \notin \Lambda(\omega; \nu, q)$ before $(\infty, \infty) \in \Lambda(\omega; \nu, q)$.

Next consider $\omega = R$. Then $\gamma(R, (\infty, \infty); 1, q) > 0$ and $\gamma(R, (0, 0); 1, q) > 0$ for all $q \in [0, 1]$, since only correct learning can occur at $\nu = 1$. The only change in the above expressions is that now the true probabilities of each action are taken with respect to state R, rather than state L. Therefore, the comparative statics are similar to the comparative statics in state L: $\gamma_1(R, (0, 0); \nu, q)$ and $\gamma_1(R, (\infty, \infty); \nu, q)$ are decreasing in ν and q. Therefore, $\gamma_1(R, (0, 0); \nu, q) > 0$ for all ν and q, which implies $(0, 0) \notin \Lambda(R; \nu, q)$ for all ν and q. Similarly, $\gamma_1(R, (\infty, \infty); \nu, q) > 0$ for all ν and q, which implies $(\infty, \infty) \in \Lambda(R; \nu, q)$ for all ν and q. Therefore, $\Lambda(R; \nu, q) = \{(\infty, \infty)\}$ for all ν and q and learning is almost surely correct. $Q.E.D.$

APPENDIX C: PROOFS FROM SECTION 4

C.1. *Section 4.1 (Overreaction)*

PROOF OF OBSERVATION 1: Suppose agents observe signals directly. Modify the definition of the expected change in the log likelihood ratio to allow for an uncountable signal space (as opposed to a finite action space):

$$\tilde{\gamma}(\omega, \lambda; \nu) \equiv \int_{s \in \mathcal{S}} \log\left(\frac{s}{1 - s}\right)^{\nu} dF^{\omega}(s).$$

Then $\tilde{\gamma}(\omega, \lambda; \nu) = \nu\tilde{\gamma}(\omega, \lambda; 1)$ since $\int_{s \in \mathcal{S}} \log(\frac{s}{1-s})^{\nu} dF^{\omega}(s) = \nu \int_{s \in \mathcal{S}} \log(\frac{s}{1-s}) dF^{\omega}(s) = \nu\tilde{\gamma}(\omega, \lambda; 1)$, where $\tilde{\gamma}(\omega, \lambda; 1)$ is the expected change in the log likelihood ratio in the correctly specified model. Therefore, $\tilde{\gamma}(\omega, \lambda; \nu)$ has the same sign as $\tilde{\gamma}(\omega, \lambda; 1)$. Since correct learning obtains almost surely when agents have a correctly specified model, $\Lambda(L; 1) = \{0\}$ and $\Lambda(R; 1) = \{\infty\}$. This implies that $\Lambda(L; \nu) = \{0\}$ and $\Lambda(R; \nu) = \{\infty\}$ for all $\nu \in [1, \infty)$. Berk (1966) showed that beliefs converges a.s. in state L to the unique element in $\Lambda(L)$. Therefore, correct learning occurs almost surely, independent of ν. Q.E.D.

PROOF OF PROPOSITION 1: Let $x \equiv \pi(\theta_1)/\pi(\theta_2)$ denote the ratio of social to autarkic types. If an agent is an autarkic type with overreaction parameter ν, then $\hat{p}^*(\nu) \equiv \frac{(p^*)^{1/\nu}}{(1 - p^*)^{1/\nu} + (p^*)^{1/\nu}}$ is the signal cutoff to choose action a_1. Note that this reduces to p^* for a correctly specified type, that is, $\hat{p}^*(1) = p^*$.

We first construct the locally stable set. We write $\gamma_i(\omega, \lambda; x, \nu)$ and $\Lambda(\omega; x, \nu)$ to make these expressions' dependence on parameters x and ν explicit. Define $\Gamma_0(x, \nu) \equiv \gamma_1(L, 0; x, \nu)(x + 1)$ and $\Gamma_{\infty}(x, \nu) \equiv \gamma_1(L, \infty; x, \nu)(x + 1)$. Then, from the construction of $\gamma_i(\omega, \lambda; x, \nu)$,

$$\Gamma_0(x, \nu) \equiv \left(F^L(\hat{p}^*(\nu)) + x\right) \log \frac{F^R(p^*) + x}{F^L(p^*) + x} - F^R(\hat{p}^*(\nu)) \log \frac{F^R(p^*)}{F^L(p^*)}$$

$$+ \left(F^L(1/2) - F^R(1/2) + F^R(\hat{p}^*(\nu)) - F^L(\hat{p}^*(\nu))\right) \log \frac{F^R(1/2) - F^R(p^*)}{F^L(1/2) - F^L(p^*)},$$

$$\Gamma_\infty(x, \nu) \equiv - \left(F^R(\hat{p}^*(\nu)) + x\right) \log \frac{F^R(p^*) + x}{F^L(p^*) + x} + F^L(\hat{p}^*(\nu)) \log \frac{F^R(p^*)}{F^L(p^*)}$$

$$+ \left(F^L(1/2) - F^R(1/2) + F^R(\hat{p}^*(\nu)) - F^L(\hat{p}^*(\nu))\right) \log \frac{F^R(1/2) - F^R(p^*)}{F^L(1/2) - F^L(p^*)}.$$

These functions have the same sign as $\gamma_1(L, 0; x, \nu)$ or $\gamma_1(L, \infty; x, \nu)$, respectively. Therefore, the signs of $\Gamma_0(x, \nu)$ and $\Gamma_\infty(x, \nu)$ can be used to characterize the locally stable set $\Lambda(\omega; x, \nu)$. Since there is a single social type, long-run learning is fully determined by $\Lambda(\omega; x, \nu)$.

To show the desired cutoffs exist, we show (i) $\nu \mapsto \Gamma_0(x, \nu)$ crosses zero at most once for a fixed x, (ii) if $0 \notin \Lambda(L; x, \nu)$ for some x', then $0 \notin \Lambda(L; x, \nu)$ for all $x > x'$, (iii) $\infty \notin \Lambda(L; x, \nu)$ for all (x, ν). To show (i), note that the derivative of $\Gamma_0(x, \nu)$ with respect to ν is

$$\frac{\partial \Gamma_0}{\partial \nu} = \frac{d\hat{p}^*(\nu)}{d\nu} f^L(\hat{p}^*(\nu))$$

$$\times \left(\log \frac{F^R(p^*) + x}{F^L(p^*) + x} - \frac{\hat{p}^*(\nu)}{1 - \hat{p}^*(\nu)} \log \frac{F^R(p^*)}{F^L(p^*)} \right.$$

$$\left. - \left(1 - \frac{\hat{p}^*(\nu)}{1 - \hat{p}^*(\nu)}\right) \log \frac{F^R(1/2) - F^R(p^*)}{F^L(1/2) - F^L(p^*)} \right),$$

where we use the property that $f^R(\hat{p}^*(\nu))/f^L(\hat{p}^*(\nu)) = \hat{p}^*(\nu)/(1 - \hat{p}^*(\nu))$, which follows from the normalization that signal realizations are posterior beliefs. The sign of this derivative is the same as the sign of

$$\log \frac{F^R(p^*) + x}{F^L(p^*) + x} + \frac{\hat{p}^*(\nu)}{1 - \hat{p}^*(\nu)} \left(\log \frac{F^R(1/2) - F^R(p^*)}{F^L(1/2) - F^L(p^*)} - \log \frac{F^R(p^*)}{F^L(p^*)} \right)$$

$$- \log \frac{F^R(1/2) - F^R(p^*)}{F^L(1/2) - F^L(p^*)}.$$

This expression is increasing in ν, so $\nu \mapsto \Gamma_0(x, \nu)$ is either decreasing, U-shaped, or increasing. Given $\Gamma_0(x, 1) \leq 0$, $\nu \mapsto \Gamma_0(x, \nu)$ changes signs at most once. Therefore, for a fixed x, there exists a cutoff $\bar{\nu} > 1$ such that $0 \notin \Lambda(L; x, \nu)$ for all $\nu > \bar{\nu}$ and $0 \in \Lambda(L; x, \nu)$ for all $\nu < \bar{\nu}$. For (ii), note that the derivative $\partial \Gamma_0/\partial \nu$ is strictly increasing in x. If we can show that $\Gamma_0(x, 1)$ is increasing in x, then as x increases, $\boldsymbol{\lambda} = 0$ becomes unstable at a lower value of ν. The derivative of $\Gamma_0(x, 1)$ with respect to x is

$$\frac{\partial \Gamma_0}{\partial x} = \log \frac{F^R(p^*) + x}{F^L(p^*) + x} + \frac{F^L(p^*) - F^R(p^*)}{F^R(p^*) + x}.$$

Moreover, the second derivative is

$$\frac{\partial^2 \Gamma_0}{\partial x^2} = -\frac{\left(F^L(p^*) - F^R(p^*)\right)^2}{\left(F^L(p^*) + x\right)\left(F^R(p^*) + x\right)^2} < 0.$$

So $x \mapsto \Gamma_0(x, 1)$ is concave in x and $\lim_{x \to \infty} \frac{\partial \Gamma_0}{\partial x}(x, 1) = 0$. Therefore, $\frac{\partial \Gamma_0}{\partial x}(x, 1) \geq 0$ for all x. Finally, $\Gamma_0(x, \nu) \geq \Gamma_0(x', \nu)$ for $x > x'$. Therefore, as x increases, $\gamma_1(L, 0; x, \nu)$ crosses 0 at a lower ν, that is, if $0 \notin \Lambda(L; x', \nu)$, then $0 \notin \Lambda(L; x, \nu)$. For (iii), the derivative of $\Gamma_\infty(x, \nu)$ with respect to ν is

$$\frac{\partial \Gamma_\infty}{\partial \nu} = \frac{d\hat{p}^*(\nu)}{d\nu} f^R(\hat{p}^*(\nu))$$
$$\times \left(-\log \frac{F^R(p^*) + x}{F^L(p^*) + x} + \frac{1 - \hat{p}^*(\nu)}{\hat{p}^*(\nu)} \log \frac{F^R(p^*)}{F^L(p^*)} \right.$$
$$\left. - \left(\frac{1 - \hat{p}^*(\nu)}{\hat{p}^*(\nu)} - 1 \right) \log \frac{F^R(1/2) - F^R(p^*)}{F^L(1/2) - F^L(p^*)} \right).$$

This derivative is maximized at $x = 0$ for a fixed ν since $\log \frac{F^R(p^*) + x}{F^L(p^*) + x}$ is monotone in x. At $x = 0$, $\frac{\partial \Gamma_\infty}{\partial \nu}(0, \nu) < 0$. Therefore, $\frac{\partial \Gamma_\infty}{\partial \nu}(x, \nu) < 0$ for all (x, ν) and $\infty \notin \Lambda(L; x, \nu)$ for all (x, ν).

The symmetric environment implies identical cutoffs in state R. Therefore, $\bar{\pi}$ and $\bar{\nu}$ exist and satisfy the desired properties. Finally,

$$\lim_{x \to \infty} \lim_{\nu \to \infty} \Gamma_0(x, \nu) = F^R(p^*) - F^L(p^*) - F^R(1/2) \log \frac{F^R(p^*)}{F^L(p^*)} > 0$$

by assumption. Therefore, cyclical learning occurs for some parameters.   *Q.E.D.*

C.2. *Section 4.2 (Naive Learning)*

We first prove Proposition 3 and then Proposition 2, as the latter is based on the former.

PROOF OF PROPOSITION 3: Let $\alpha_L \equiv F^L(1/2)$ be the probability an autarkic type plays action L in state L and $\alpha_R \equiv F^R(1/2)$ be the probability an autarkic type plays action L in state R. Note that $\alpha_L \in (0, 1)$ and $\alpha_R \in (0, 1)$, since private signals are informative. In a slight abuse of notation, let $\hat{\pi}_i$ denote $\hat{\pi}_i(\theta_A)$ and π denote $\pi(\theta_A)$ to abbreviate the following expressions.

We first construct the locally stable set. We write $\gamma_i(\omega, \lambda; \hat{\pi}_i)$ and $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2)$ to make these expressions' dependence on $\hat{\pi}_1$ and $\hat{\pi}_2$ explicit. The local stability of correct learning is determined by the sign of

$$\gamma_i(L, (0, 0); \hat{\pi}_i) = (\pi \alpha_L + 1 - \pi) \log\left(\frac{\hat{\pi}_i \alpha_R + 1 - \hat{\pi}_i}{\hat{\pi}_i \alpha_L + 1 - \hat{\pi}_i} \right) + \pi(1 - \alpha_L) \log\left(\frac{1 - \alpha_R}{1 - \alpha_L} \right).$$

If θ_i has a correctly specified model, $\gamma_i(L, (0, 0); \pi) < 0$. This expression is decreasing in $\hat{\pi}_i$. Therefore, $\gamma_i(L, (0, 0); \hat{\pi}_i) < 0$ for all $\hat{\pi}_i \geq \pi$. This implies that $(0, 0) \in \Lambda(L; \hat{\pi}_1, \hat{\pi}_2)$

for all $\hat{\pi}_1, \hat{\pi}_2$. Therefore, correct learning arises with positive probability at any level of heterogeneity. The local stability of incorrect learning is determined by the sign of

$$\gamma_i\bigl(L, (\infty, \infty); \hat{\pi}_i\bigr) = \pi\alpha_L \log\left(\frac{\alpha_R}{\alpha_L}\right) + \bigl(\pi(1 - \alpha_L) + 1 - \pi\bigr) \log\left(\frac{\hat{\pi}_i(1 - \alpha_R) + 1 - \hat{\pi}_i}{\hat{\pi}_i(1 - \alpha_L) + 1 - \hat{\pi}_i}\right).$$

This expression is increasing in $\hat{\pi}_i$ and is equivalent to the representative agent model at $\hat{\pi}_i = \hat{\pi}$. Therefore, if $\gamma_i(L, (\infty, \infty); \hat{\pi}) < 0$, then $\gamma_1(L, (\infty, \infty); \hat{\pi}_1) < 0$ since $\hat{\pi}_1 \leq \hat{\pi}$ by definition. This implies that if incorrect learning does not arise in the representative agent model with bias $\hat{\pi}$, that is, $(\infty, \infty) \notin \Lambda(L; \hat{\pi}, \hat{\pi})$, then it does not arise in any corresponding heterogeneous model with average bias $\hat{\pi}$, that is, $(\infty, \infty) \notin \Lambda(L; \hat{\pi}_1, \hat{\pi}_2)$ for all $\hat{\pi}_1, \hat{\pi}_2$ such that $(\hat{\pi}_1 + \hat{\pi}_2)/2 = \hat{\pi}$. Further, we know from Bohren (2016) that there exists a cutoff $\overline{\pi} \in (\pi, 1]$ such that for $\hat{\pi}_i > \overline{\pi}$, $\gamma_i(L, (\infty, \infty); \hat{\pi}_i) > 0$, with $\overline{\pi} < 1$ for small enough π. Therefore, $(\infty, \infty) \in \Lambda(L; \hat{\pi}, \hat{\pi})$ for $\hat{\pi} > \overline{\pi}$ and $(\infty, \infty) \in \Lambda(L; \hat{\pi}_1, \hat{\pi}_2)$ for $\hat{\pi}_1 > \overline{\pi}$. The local stability of disagreement is determined by the sign of

$$\gamma_i\bigl(L, (0, \infty); \hat{\pi}_i\bigr) = \bigl(\pi\alpha_L + (1 - \pi)/2\bigr) \log\left(\frac{\hat{\pi}_i\alpha_R + \frac{1}{2}(1 - \hat{\pi}_i)}{\hat{\pi}_i\alpha_L + \frac{1}{2}(1 - \hat{\pi}_i)}\right)$$

$$+ \bigl(\pi(1 - \alpha_L) + (1 - \pi)/2\bigr) \log\left(\frac{\hat{\pi}_i(1 - \alpha_R) + \frac{1}{2}(1 - \hat{\pi}_i)}{\hat{\pi}_i(1 - \alpha_L) + \frac{1}{2}(1 - \hat{\pi}_i)}\right)$$

$$= \pi(2\alpha_L - 1) \log\left(\frac{\hat{\pi}_i(1 - \alpha_L) + \frac{1}{2}(1 - \hat{\pi}_i)}{\hat{\pi}_i\alpha_L + \frac{1}{2}(1 - \hat{\pi}_i)}\right),$$

where the second equality follows from symmetry, $\alpha_R = 1 - \alpha_L$. Given $\alpha_L > 1/2$, $(\hat{\pi}_i(1 - \alpha_L) + \frac{1}{2}(1 - \hat{\pi}_i))/(\hat{\pi}_i\alpha_L + \frac{1}{2}(1 - \hat{\pi}_i)) < 1$ and $2\alpha_L - 1 > 0$. Therefore, $\gamma_i(L, (0, \infty); \hat{\pi}_i) < 0$ for any $\hat{\pi}_i$. This implies that $(0, \infty)$ almost surely does not arise, that is, $(0, \infty) \notin \Lambda(L; \hat{\pi}_1, \hat{\pi}_2)$. Given $\gamma_i(L, (\infty, 0); \hat{\pi}_i) = \gamma_i(L, (0, \infty); \hat{\pi}_i)$, $(\infty, 0)$ almost surely does not arise. Therefore, almost surely disagreement does not arise. The construction of $\Lambda(R; \hat{\pi}_1, \hat{\pi}_2)$ is analogous.

Next, we rule out mixed learning. Since correct learning is always locally stable, the only candidate mixed outcomes are $\lambda_1^* = \infty$ or $\lambda_2^* = \infty$. As argued above, $\gamma_1(L, (0, \infty); \hat{\pi}_1) < 0$ for any $\hat{\pi}_1$ and $\gamma_2(L, (\infty, 0); \hat{\pi}_2) < 0$ for any $\hat{\pi}_2$. This implies $\Lambda_M(L) = \emptyset$. Therefore, mixed learning almost surely does not arise. The construction of $\Lambda_M(R)$ is analogous.

Given $\Lambda_M(\omega) = \emptyset$ and $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2)$ does not contain any disagreement outcomes—and therefore, we do not need to consider maximal accessibility—by Theorem 4, $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2)$ fully characterizes the set of asymptotic learning outcomes. From the above characterization, either $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2) = \{(0, 0)\}$ or $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2) = \{(0, 0), (\infty, \infty)\}$. Therefore, either learning is almost surely correct, or learning is almost surely correct or incorrect with both occurring with positive probability. Further, if $\Lambda(\omega; \hat{\pi}, \hat{\pi}) = \{(0, 0)\}$, then $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2) = \{(0, 0)\}$ for all $\hat{\pi}_1, \hat{\pi}_2$ such that $(\hat{\pi}_1 + \hat{\pi}_2)/2 = \hat{\pi}$, and if $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2) = \{(0, 0), (\infty, \infty)\}$, then $\Lambda(\omega; \hat{\pi}, \hat{\pi}) = \{(0, 0), (\infty, \infty)\}$ at $\hat{\pi} = (\hat{\pi}_1 + \hat{\pi}_2)/2$. $\hspace{2em}$ *Q.E.D.*

PROOF OF PROPOSITION 2: This result follows directly from the constructions of $\gamma_i(\omega, \boldsymbol{\lambda}; \hat{\pi}_i)$ in Proposition 3. Generically, $\gamma_i(\omega, (0, 0); \hat{\pi}_i) \neq 0$ and $\gamma_i(\omega, (\infty, \infty); \hat{\pi}_i) \neq 0$ for $i = 1, 2$. Given an average bias $\hat{\pi}$, consider the case where $\gamma_i(\omega, (0, 0); \hat{\pi}) \neq 0$ and $\gamma_i(\omega, (\infty, \infty); \hat{\pi}) \neq 0$ for $i = 1, 2$. For any $\delta > 0$, there exists an ε such that for $|\hat{\pi}_1 - \hat{\pi}| < \varepsilon/2$ and $|\hat{\pi}_2 - \hat{\pi}| < \varepsilon/2$, $|\gamma_i(\omega, \boldsymbol{\lambda}; \hat{\pi}_i) - \gamma_i(\omega, \boldsymbol{\lambda}; \hat{\pi})| < \delta$ for $\boldsymbol{\lambda} \in \{(0, 0), (\infty, \infty)\}$ and $i = 1, 2$. Choosing δ small enough ensures that $\gamma_i(\omega, \boldsymbol{\lambda}; \hat{\pi}_i)$ and $\gamma_i(\omega, \boldsymbol{\lambda}; \hat{\pi})$ have the same sign. Therefore, $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2) = \Lambda(\omega; \hat{\pi}, \hat{\pi})$ and the heterogeneous set-up has the same set of learning outcomes as the corresponding representative agent set-up. *Q.E.D.*

C.3. *Section 4.3 (Level-k)*

PROOF OF PROPOSITION 4: Let $\boldsymbol{\lambda} = (\lambda_2, \lambda_3)$ denote the vector of likelihood ratios for the social types θ_2 and θ_3. Note $\lambda_{1,t} = 1$ for all t. When type $\theta_i \in \{\theta_1, \theta_2, \theta_3\}$ has current belief λ_i, it chooses action R iff it observes a signal realization $s \geq 1/(\lambda_i + 1)$. Given $\lambda_1 = 1$, type θ_1 chooses action L with probability $F^\omega(0.5)$ and action R with probability $1 - F^\omega(0.5)$, independent of the history. Type θ_2's subjective probability of each L action in the history is the probability that a level-1 type chooses action L, $\hat{\psi}_2(L|\omega, \boldsymbol{\lambda}) = F^\omega(0.5)$ and its subjective probability of each R action is $\hat{\psi}_2(R|\omega, \boldsymbol{\lambda}) = 1 - F^\omega(0.5)$, independent of the history. Given belief λ_2, level-2 chooses an L action with probability $F^\omega(1/(\lambda_2 + 1))$ and an R action with probability $1 - F^\omega(1/(\lambda_2 + 1))$. Type θ_3's subjective probability of each L action is the weighted average of the probability that a level-1 type and a level-2 type choose action L,

$$\hat{\psi}_3(L|\omega, \boldsymbol{\lambda}) = (1 - \varepsilon)F^\omega\big(1/(\lambda_2 + 1)\big) + \varepsilon F^\omega(0.5),$$

which does depend on the history through λ_2. The subjective probability of an R action is analogous. Finally, the *true* probability of an L action depends on the correct distribution over types,

$$\psi(L|\omega, \boldsymbol{\lambda}) = \pi(\theta_1)F^\omega(0.5) + \pi(\theta_2)F^\omega\big(1/(\lambda_2 + 1)\big) + \pi(\theta_3)F^\omega\big(1/(\lambda_3 + 1)\big).$$

To simplify the exposition, let $\alpha_L \equiv F^L(0.5)$ be the probability a level-1 type plays action L in state L and $\alpha_R \equiv F^R(0.5)$ be the probability a level-1 type plays action L in state R. Note that $\alpha_L \in (0, 1)$ and $\alpha_R \in (0, 1)$, since private signals are informative.

Suppose $\omega = L$. We first consider local stability for the level-3 type. At the correct learning outcome, $(0, 0)$, the level-2 type chooses action L for all signal realizations. Therefore, the level-3 type believes that L actions are approximately uninformative for small ε, $\frac{\hat{\psi}_3(L|R,(0,0))}{\hat{\psi}_3(L|L,(0,0))} = \frac{1-\varepsilon+\varepsilon\alpha_R}{1-\varepsilon+\varepsilon\alpha_L} \approx 1$ and R actions are from the level-1 type, $\frac{\hat{\psi}_3(R|R,(0,0))}{\hat{\psi}_3(R|L,(0,0))} = \frac{1-\alpha_R}{1-\alpha_L}$. Since only the level-1 type plays action R, the true probability of an R action is $\pi(\theta_1)(1 - \alpha_L)$. Therefore, for small ε, $\gamma_3(L, (0, 0)) = (\pi(\theta_1)\alpha_L + \pi(\theta_2) + \pi(\theta_3)) \log \frac{1-\varepsilon+\varepsilon\alpha_R}{1-\varepsilon+\varepsilon\alpha_L} + \pi(\theta_1)(1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} \approx \pi(\theta_1)(1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} > 0$ and correct learning is not locally stable for the level-3 type, $(0, 0) \notin \Lambda_3(L)$. Similarly, for small ε, $\gamma_3(L, (\infty, \infty)) \approx \pi(\theta_1)\alpha_L \log \frac{\alpha_R}{\alpha_L} < 0$ and incorrect learning is not locally stable for the level-3 type, $(\infty, \infty) \notin \Lambda_3(L)$. This establishes that correct learning and incorrect learning almost surely do not occur for small ε, as neither outcome is locally stable for level-3 types.

This leaves the disagreement outcomes as candidate learning outcomes. Consider $(0, \infty)$. As in the case of $(0, 0)$, the level-3 type believes that L actions are approximately uninformative and R actions are from the level-1 type. But now, this confirms the level-3

type's belief that the state is R, $\gamma_3(L, (0, \infty)) \approx (\pi(\theta_1)(1 - \alpha_L) + \pi(\theta_3)) \log \frac{1-\alpha_R}{1-\alpha_L} > 0$ and $(0, \infty) \in \Lambda_3(L)$. Similarly, $\gamma_3(L, (\infty, 0)) \approx (\pi(\theta_1)\alpha_L + \pi(\theta_3)) \log \frac{\alpha_R}{\alpha_L} < 0$ and $(\infty, 0) \in \Lambda_3(L)$. Therefore, for small ε, both disagreement outcomes are locally stable for the level-3 type, $\Lambda_3(L) = \{(0, \infty), (\infty, 0)\}$.

Next, we determine whether the disagreement outcomes are locally stable for the level-2 type. The level-2 type believes that all actions are from level-1 types. Therefore, it interprets L and R actions in the same way at both disagreement outcomes. At $(0, \infty)$, the true probability of an L action is $\pi(\theta_1)\alpha_L + \pi(\theta_2)$, while at $(\infty, 0)$, it is $\pi(\theta_1)\alpha_L + \pi(\theta_3)$. Therefore, $\gamma_2(L, (0, \infty)) = (\pi(\theta_1)\alpha_L + \pi(\theta_2)) \log \frac{\alpha_R}{\alpha_L} + (\pi(\theta_1)(1 - \alpha_L) + \pi(\theta_3)) \log \frac{1-\alpha_R}{1-\alpha_L}$ and $\gamma_2(L, (\infty, 0)) = (\pi(\theta_1)\alpha_L + \pi(\theta_3)) \log \frac{\alpha_R}{\alpha_L} + (\pi(\theta_1)(1 - \alpha_L) + \pi(\theta_2)) \log \frac{1-\alpha_R}{1-\alpha_L}$. The signs of these expressions vary with the true distribution of types. To characterize the region of the type distribution at which each disagreement outcome is locally stable, we use the inequalities (a) $\frac{\alpha_R}{\alpha_L} < 1$, (b) $\frac{1-\alpha_R}{1-\alpha_L} > 1$, and (c) from the correctly specified model, $\alpha_L \log \frac{\alpha_R}{\alpha_L} + (1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} < 0$, as well as the property that $\pi \mapsto \gamma_2(L, (0, \infty))$ and $\pi \mapsto \gamma_2(L, (\infty, 0))$ are continuous.

Case (i): As $\pi(\theta_3) \to 0$, $\gamma_2(L, (0, \infty)) \to (\pi(\theta_1)\alpha_L + 1 - \pi(\theta_1)) \log \frac{\alpha_R}{\alpha_L} + \pi(\theta_1)(1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} < 0$ for all $\pi(\theta_1)$, where the negative sign follows from inequalities (a) and (c). Therefore, there exists a cutoff $c_1 > 0$ such that for $\pi(\theta_3) < c_1$, $(0, \infty) \in \Lambda_2(L)$ for all $\pi(\theta_1)$ and $\pi(\theta_2)$.

Case (ii): As $\pi(\theta_3) \to 1$, $\gamma_2(L, (0, \infty)) \to \log \frac{1-\alpha_R}{1-\alpha_L} > 0$ and $\gamma_2(L, (\infty, 0)) \to \log \frac{\alpha_R}{\alpha_L} < 0$. Therefore, there exists an interior cutoff $c_2 \in (0, 1)$ such that for $\pi(\theta_3) > c_2$, $(0, \infty) \notin \Lambda_2(L)$ and there exists a cutoff $c_3 < 1$ such that for $\pi(\theta_3) > c_3$, $(\infty, 0) \notin \Lambda_2(L)$ for all $\pi(\theta_1)$ and $\pi(\theta_2)$, where $c_2 > 0$ follows from part (i). Therefore, there exists an interior cutoff $\bar{\pi}_3 = \max\{c_2, c_3\} \in (0, 1)$ such that if $\pi(\theta_3) > \bar{\pi}_3$, neither disagreement outcome is locally stable for θ_2. Combined with $\Lambda_3(L) = \{(0, \infty), (\infty, 0)\}$, this implies that $\Lambda(L) = \emptyset$ for $\pi(\theta_3) > \bar{\pi}_3$ and small ε.

Case (iii): As $\pi(\theta_2) \to 0$, $\gamma_2(L, (\infty, 0)) \to (\pi(\theta_1)\alpha_L + 1 - \pi(\theta_1)) \log \frac{\alpha_R}{\alpha_L} + \pi(\theta_1)(1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} < 0$ for all $\pi(\theta_1)$, where the negative sign follows from inequalities (a) and (c). Therefore, there exists a cutoff $c_4 > 0$ such that for $\pi(\theta_2) < c_4$, $(\infty, 0) \notin \Lambda_2(L)$ for all $\pi(\theta_1)$ and $\pi(\theta_3)$.

Case (iv): As $\pi(\theta_2) \to 1$, $\gamma_2(L, (0, \infty)) \to \log \frac{\alpha_R}{\alpha_L} < 0$ and $\gamma_2(L, (\infty, 0)) \to \log \frac{1-\alpha_R}{1-\alpha_L} > 0$. Therefore, there exists a cutoff $c_5 < 1$ such that for $\pi(\theta_2) > c_5$, $(0, \infty) \in \Lambda_2(L)$ and there exists an interior cutoff $c_6 \in (0, 1)$ such that for $\pi(\theta_2) > c_6$, $(\infty, 0) \in \Lambda_2(L)$ for all $\pi(\theta_1)$ and $\pi(\theta_3)$, where $c_6 > 0$ follows from case (iii). Therefore, there exists an interior cutoff $\bar{\pi}_2 = \max\{c_5, c_6\} \in (0, 1)$ such that if $\pi(\theta_2) > \bar{\pi}_2$, both disagreement outcomes are locally stable for θ_2. Combined with $\Lambda_3(L) = \{(0, \infty), (\infty, 0)\}$, this implies that $\Lambda(L) = \{(0, \infty), (\infty, 0)\}$ for $\pi(\theta_2) > \bar{\pi}_2$ and small ε.

Case (v): As $\pi(\theta_1) \to 1$, $\gamma_2(L, (0, \infty)) \to \alpha_L \log \frac{\alpha_R}{\alpha_L} + (1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} < 0$ and $\gamma_2(L, (\infty, 0)) \to \alpha_L \log \frac{\alpha_R}{\alpha_L} + (1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} < 0$. Therefore, there exists an interior cutoff $c_7 \in (0, 1)$ such that for $\pi(\theta_1) > c_7$, $(0, \infty) \in \Lambda_2(L)$ and there exists an interior cutoff $c_8 \in (0, 1)$ such that for $\pi(\theta_1) > c_8$, $(\infty, 0) \notin \Lambda_2(L)$ for all $\pi(\theta_2)$ and $\pi(\theta_3)$, where $c_7 > 0$ and $c_8 > 0$ follow from cases (ii) and (iv). Therefore, there exists an interior cutoff $\bar{\pi}_1 = \max\{c_7, c_8\} \in (0, 1)$ such that if $\pi(\theta_1) > \bar{\pi}_1$, $(0, \infty)$ is locally stable for θ_2 and $(\infty, 0)$ is not. Combined with $\Lambda_3(L) = \{(0, \infty), (\infty, 0)\}$, this implies that $\Lambda(L) = \{(0, \infty)\}$ for $\pi(\theta_1) > \bar{\pi}_1$ and small ε.

Fixing $\pi(\theta_2)$, $\gamma_2(L, (0, \infty))$ is increasing in $\pi(\theta_3)$. Given this, we next show that the type distribution can be divided into two connected regions in the simplex such that $(0, \infty) \in \Lambda_2(L)$ or $(0, \infty) \notin \Lambda_2(L)$, and these regions are separated by the unique solution to $\gamma_2(L, (0, \infty)) = 0$. As shown above, at $\pi(\theta_2) = 0$ and $\pi(\theta_3) = 0$, $\gamma_2(L, (0, \infty)) < 0$, and at $\pi(\theta_2) = 0$ and $\pi(\theta_3) = 1$, $\gamma_2(L, (0, \infty)) > 0$. Therefore, there exists a cutoff $c_9 \in (0, 1)$ such that at $\pi(\theta_2) = 0$ and $\pi(\theta_3) = c_9$, $\gamma_2(L, (0, \infty)) = 0$. Similarly, there exists a cutoff $c_{10} \equiv \log \frac{\alpha_L}{\alpha_R} / (\log \frac{\alpha_L}{\alpha_R} - \log \frac{1 - \alpha_L}{1 - \alpha_R})$ such that at $\pi(\theta_1) = 0$ and $\pi(\theta_3) = c_{10}$, $\gamma_2(L, (0, \infty)) = 0$. Given $\gamma_2(L, (0, \infty))$ is linear in $\pi(\theta_2)$ and $\pi(\theta_3)$, the solution to $\gamma_2(L, (0, \infty)) = 0$ is linear in the simplex and represented by the line connecting $(1 - c_9, 0, c_9)$ and $(0, 1 - c_{10}, c_{10})$. This establishes the above statement.

Fixing $\pi(\theta_2)$, $\gamma_2(L, (\infty, 0))$ is decreasing in $\pi(\theta_3)$. Therefore, by similar reasoning, the type distribution can be divided into two connected regions such that $(\infty, 0) \in \Lambda_2(L)$ or $(\infty, 0) \notin \Lambda_2(L)$, and these regions are separated by the unique solution to $\gamma_2(L, (\infty, 0)) = 0$. Given $\gamma_2(L, (\infty, 0))$ is linear in $\pi(\theta_2)$ and $\pi(\theta_3)$, the solution to $\gamma_2(L, (\infty, 0)) = 0$ is linear in the simplex and represented by the line connecting $(1 - c_{11}, c_{11}, 0)$ and $(0, 1 - c_{12}, c_{12})$, where $c_{11} \in (0, 1)$ is the value of $\pi(\theta_2)$ such that $\gamma_2(L, (\infty, 0)) = 0$ when $\pi(\theta_3) = 0$, and $c_{12} \equiv \log \frac{1 - \alpha_L}{1 - \alpha_R} / (\log \frac{1 - \alpha_L}{1 - \alpha_R} - \log \frac{\alpha_L}{\alpha_R})$.

Given the linearity of both solutions, if $c_{10} \geq c_{12}$, then the solution to $\gamma_2(L, (0, \infty)) = 0$ lies above the solution to $\gamma_2(L, (\infty, 0)) = 0$. Therefore, there are three distinct regions such that for small ε, either (i) $\Lambda(L) = \emptyset$, (ii) $\Lambda(L) = \{(0, \infty)\}$, or (iii) $\Lambda(L) = \{(0, \infty), (\infty, 0)\}$. Otherwise, if $c_{10} \leq c_{12}$, the solutions cross exactly once. Therefore, there are four distinct regions such that for small ε, either (i) $\Lambda(L) = \emptyset$, (ii) $\Lambda(L) = \{(0, \infty)\}$, (iii) $\Lambda(L) = \{(\infty, 0)\}$, or (iv) $\Lambda(L) = \{(0, \infty), (\infty, 0)\}$. Note that when the signal distributions are symmetric, $c_{10} \geq c_{12}$. The construction of $\Lambda(R)$ is analogous.

We next show that both disagreement outcomes are maximally accessible at all type distributions. Formally, we show that for any $\pi \in \Delta((\theta_1, \theta_2, \theta_3))$ and $\varepsilon \in (0, 1]$, $(0, \infty)$ and $(\infty, 0)$ are maximally accessible. At $\boldsymbol{\lambda} = (0, 0)$, type θ_2 perceives L actions as stronger evidence of state L than type θ_3,

$$\frac{\hat{\psi}_2(L | R, (0, 0))}{\hat{\psi}_2(L | L, (0, 0))} = \frac{\alpha_R}{\alpha_L} < \frac{\varepsilon + (1 - \varepsilon)\alpha_R}{\varepsilon + (1 - \varepsilon)\alpha_L} = \frac{\hat{\psi}_3(L | R, (0, 0))}{\hat{\psi}_3(L | L, (0, 0))},$$

and both types perceive R actions in the same way,

$$\frac{\hat{\psi}_2(R | R, (0, 0))}{\hat{\psi}_2(R | L, (0, 0))} = \frac{\hat{\psi}_3(R | R, (0, 0))}{\hat{\psi}_3(R | L, (0, 0))} = \frac{1 - \alpha_R}{1 - \alpha_L}.$$

Therefore, $\theta_3 \succ_{(0,0)} \theta_2$. From Definition 7, this implies that $(0, \infty)$ is maximally accessible. At $\boldsymbol{\lambda} = (\infty, \infty)$, type θ_2 perceives R actions as stronger evidence of state R than type θ_3,

$$\frac{\hat{\psi}_2(R | R, (\infty, \infty))}{\hat{\psi}_2(R | L, (\infty, \infty))} = \frac{1 - \alpha_R}{1 - \alpha_L} > \frac{\varepsilon + (1 - \varepsilon)(1 - \alpha_R)}{\varepsilon + (1 - \varepsilon)(1 - \alpha_L)} = \frac{\hat{\psi}_3(R | R, (\infty, \infty))}{\hat{\psi}_3(R | L, (\infty, \infty))},$$

and both types perceive L actions in the same way,

$$\frac{\hat{\psi}_2(L | R, (\infty, \infty))}{\hat{\psi}_2(L | L, (\infty, \infty))} = \frac{\hat{\psi}_3(L | R, (\infty, \infty))}{\hat{\psi}_3(L | L, (\infty, \infty))} = \frac{\alpha_R}{\alpha_L}.$$

Therefore, $\theta_2 \succ_{(\infty,\infty)} \theta_3$. From Definition 7, this implies that $(\infty, 0)$ is maximally accessible. Therefore, a disagreement outcome arises with positive probability if and only if it is in $\Lambda(\omega)$.

Finally, we need to rule out mixed learning outcomes. Suppose $\omega = L$ and consider the four possible mixed outcomes. Consider $(0, \theta_3)$. By the concavity of the log operator, $\alpha_L \log \frac{\alpha_R}{\alpha_L} + (1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} < 0$. Therefore, since $\frac{\alpha_R}{\alpha_L} < 0$, $\gamma_2(L, (0,0)) = (\pi_1 \alpha_L + \pi(\theta_2) + \pi(\theta_3)) \log \frac{\alpha_R}{\alpha_L} + \pi_1(1 - \alpha_L) \log \frac{1-\alpha_R}{1-\alpha_L} < 0$ and $(0, 0) \in \Lambda_2(L)$. By the definition of $\Lambda_M(L)$, this implies that $(0, \theta_3) \notin \Lambda_M(L)$ and this mixed learning outcome almost surely does not arise. Consider (∞, θ_3). This outcome is in $\Lambda_M(L)$ if $(\infty, \infty) \notin \Lambda_2(L)$ and $(0, \infty) \notin \Lambda_2(L)$, which is equivalent to $\gamma_2(L, (\infty, \infty)) < 0$ and $\gamma_2(L, (0, \infty)) > 0$. However, $\gamma_2(L, (\lambda_2, \infty))$ is increasing in λ_2, so this is not possible. Therefore, $(\infty, \theta_3) \notin \Lambda_M(L)$ and this mixed learning outcome almost surely does not arise. Consider $(0, \theta_2)$. This outcome is in $\Lambda_M(L)$ if $(0, 0) \notin \Lambda_3(L)$ and $(0, \infty) \notin \Lambda_3(L)$. From the characterization of $\Lambda(L)$ above, we know that $(0, \infty) \in \Lambda_3(L)$. Therefore, $(0, \theta_2) \notin \Lambda_M(L)$ and this mixed learning outcome almost surely does not arise. Consider (∞, θ_2). This outcome is in $\Lambda_M(L)$ if $(\infty, 0) \notin \Lambda_3(L)$ and $(\infty, \infty) \notin \Lambda_3(L)$. From the characterization of $\Lambda(L)$ above, we know that $(\infty, 0) \in \Lambda_3(L)$. Therefore, $(\infty, \theta_2) \notin \Lambda_M(L)$ and this mixed learning outcome almost surely does not arise. Together, this establishes $\Lambda_M(L) = \emptyset$. Similar logic shows $\Lambda_M(R) = \emptyset$.

Given $\Lambda_M(\omega) = \emptyset$ and both disagreement outcomes are maximally accessible, by Theorem 4, $\Lambda(\omega)$ determines the set of asymptotic learning outcomes. As $\varepsilon \to 1$, $\Lambda(\omega) \subseteq \{(0, \infty), (\infty, 0)\}$. Either $\Lambda(\omega) = \emptyset$, in which case learning is cyclical for both types, or $\Lambda(\omega) \neq \emptyset$, in which case beliefs almost surely converge to a limit random variable with support $\Lambda(\omega)$. The construction of $\Lambda(\omega)$ above establishes the cutoffs on the type distribution such that $\Lambda(\omega) = \emptyset$, $\Lambda(\omega) = \{(0, \infty)\}$, $\Lambda(\omega) = \{(\infty, 0)\}$, or $\Lambda(\omega) = \{(0, \infty), (\infty, 0)\}$. *Q.E.D.*

APPENDIX D: LEARNING CHARACTERIZATION: MORE THAN TWO SOCIAL TYPES

This section proves analogues of the global stability of disagreement, mixed learning, and belief convergence results in Section 3 and Appendix A for any finite number of social types. Together, this establishes a direct analogue of Theorem 4; an analogue of Corollary 2 immediately follows. These results nest the case of $k \leq 2$.

D.1. *Global Stability of Disagreement*

We first prove an analogue of Theorem 7 to show that separability can also be used to establish the global stability of a disagreement outcome when there are more than two social types. We then extend the definition of maximal accessibility and prove that it implies the separability condition, establishing an analogue of Theorem 3.

THEOREM 7′—Global Stability of Disagreement ($k \geq 2$): *Consider a learning environment that is identified at certainty and satisfies Assumptions* 1 *to* 4. *Suppose disagreement outcome* $\boldsymbol{\lambda}^* \in \Lambda(\omega)$ *and, starting from agreement outcome* $\boldsymbol{\lambda}_1^* \in \{0^k, \infty^k\}$, *there exists a finite sequence of adjacent disagreement outcomes* $\boldsymbol{\lambda}_2^*, \dots, \boldsymbol{\lambda}_L^* = \boldsymbol{\lambda}^*$ *such that for* $l = 1, \dots, L-1$, *either* (i) $(\boldsymbol{\lambda}_l^*)_i = 0$, $(\boldsymbol{\lambda}_{l+1}^*)_i = \infty$, *and* $\boldsymbol{\lambda}_l^*$ *is separable at zero for* θ_i, *or* (ii) $(\boldsymbol{\lambda}_l^*)_i = \infty$, $(\boldsymbol{\lambda}_{l+1}^*)_i = 0$, *and* $\boldsymbol{\lambda}_l^*$ *is separable at infinity for* θ_i. *Then* $\boldsymbol{\lambda}^*$ *is globally stable in state* ω.

PROOF: Given $\kappa \in \{1, \ldots, k-1\}$, consider disagreement outcome $\boldsymbol{\lambda}^* = (0^\kappa, \infty^{k-\kappa})$. Suppose $\boldsymbol{\lambda}^* \in \Lambda(\omega)$ and for each $l = 1, \ldots k - \kappa$, $\boldsymbol{\lambda}_l^* = (0^{k-l+1}, \infty^{l-1})$ is separable at zero for type θ_{k-l+1}. Given $\boldsymbol{\lambda}_l^* = (0^{k-l+1}, \infty^{l-1})$ is separable at zero for type θ_{k-l+1}, by Lemma 5, $\boldsymbol{\lambda}_{l+1}^* = (0^{k-l}, \infty^l)$ is adjacently accessible from $\boldsymbol{\lambda}_l^*$. Since this holds for each element of the sequence starting at $\boldsymbol{\lambda}_1^* = 0^k$ and ending at $\boldsymbol{\lambda}_{k-\kappa+1}^* = \boldsymbol{\lambda}^*$, by Lemma 6, $\boldsymbol{\lambda}^*$ is accessible. Fix an initial belief $\boldsymbol{\lambda}_1 \in (0, \infty)^k$ and choose an $\varepsilon < e^{-E}$, where E is defined in Eq. (11). By accessibility, there exists a finite sequence ξ of N actions that occurs with positive probability, such that following ξ, $\boldsymbol{\lambda}_{N+1} \in B_\varepsilon(\boldsymbol{\lambda}^*)$. Given $\boldsymbol{\lambda}^*$ is locally stable, this implies $\Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^* | h = \xi) > 0$. Given $\Pr(h = \xi) > 0$, from any $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, $\Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) > 0$. This establishes that $\boldsymbol{\lambda}^*$ is globally stable. The case in which there is a sequence of stationary beliefs that are separable at infinity is analogous, as is the proof for other disagreement outcomes. *Q.E.D.*

We next use the maximal R-order $\succ_{\boldsymbol{\lambda}}$ to define a sufficient condition for separability, which we refer to as maximally separable. We use this condition to extend the definition of maximal accessibility to the case of more than two social types.

DEFINITION 11—Maximally Separable ($k \geq 2$): Belief $\boldsymbol{\lambda}^* \in \{0, \infty\}^k \setminus \infty^k$ is *maximally separable at zero* for type θ_i with $\lambda_i^* = 0$ if $\theta_j \succeq_{\boldsymbol{\lambda}^*} \theta_i$ for all j with $\lambda_j^* = \infty$ and $\theta_i \succ_{\boldsymbol{\lambda}^*} \theta_j$ for all $j \neq i$ with $\lambda_j^* = 0$. Belief $\boldsymbol{\lambda}^* \in \{0, \infty\}^k \setminus 0^k$ is *maximally separable at infinity* for type θ_i with $\lambda_i^* = \infty$ if $\theta_j \succ_{\boldsymbol{\lambda}^*} \theta_i$ for all $j \neq i$ with $\lambda_j^* = \infty$ and $\theta_i \succeq_{\boldsymbol{\lambda}^*} \theta_j$ for all j with $\lambda_j^* = 0$.

DEFINITION 7′—Maximal Accessibility ($k \geq 2$): Disagreement outcome $\boldsymbol{\lambda}^* \in \{0, \infty\}^k \setminus \{0^k, \infty^k\}$ is *maximally accessible* if, starting from agreement outcome $\boldsymbol{\lambda}_1^* \in \{0^k, \infty^k\}$, there exists a finite sequence of adjacent disagreement outcomes $\boldsymbol{\lambda}_2^*, \ldots, \boldsymbol{\lambda}_L^* = \boldsymbol{\lambda}^*$ such that for $l = 1, \ldots, L-1$, either (i) $(\boldsymbol{\lambda}_l^*)_i = 0$, $(\boldsymbol{\lambda}_{l+1}^*)_i = \infty$, and $\boldsymbol{\lambda}_l^*$ is maximally separable at zero for θ_i, or (ii) $(\boldsymbol{\lambda}_l^*)_i = \infty$, $(\boldsymbol{\lambda}_{l+1}^*)_i = 0$, and $\boldsymbol{\lambda}_l^*$ is maximally separable at infinity for θ_i.

As in the case of $k = 2$, maximal accessibility guarantees that there exists a finite sequence of a_1 and a_M actions that separates beliefs and moves them to a neighborhood of the disagreement outcome. It is straightforward to verify from the primitives of the model and is equivalent to Definition 7 when $k = 2$. Using Definition 7′, the statement of Theorem 3′ is identical to Theorem 3.

THEOREM 3′—Global Stability of Disagreement ($k \geq 2$): *Consider a learning environment that satisfies Assumptions 1 to 4. If disagreement outcome $\boldsymbol{\lambda}^*$ is in $\Lambda(\omega)$ and maximally accessible, then $\boldsymbol{\lambda}^*$ is globally stable in state ω.*

PROOF: We show that Definition 7′ implies the conditions for separability outlined in Theorem 7′. Given $\kappa \in \{1, \ldots, k-1\}$, consider $\boldsymbol{\lambda}^* = (0^\kappa, \infty^{k-\kappa})$. Suppose $\boldsymbol{\lambda}^* \in \Lambda(\omega)$ and $\boldsymbol{\lambda}^*$ is maximally accessible, with $\boldsymbol{\lambda}_l^* = (0^{k-l+1}, \infty^{l-1})$ maximally separable at zero for θ_{k-l+1} for $l = 1, \ldots k - \kappa$. For each $l = 1, \ldots k - \kappa$, $\theta_{k-l+1} \succ_{\boldsymbol{\lambda}_l^*} \theta_{k-l}$ implies that the submatrix $\Psi[\theta_{k-l+1}, \theta_{k-l}; a_1, a_M](\boldsymbol{\lambda}_l^*)$ defined in Eq. (14) has a positive determinant. Therefore, there exists a $c \in \mathbb{R}_+^2$ that solves

$$\Psi[\theta_{k-l+1}, \theta_{k-l}; a_1, a_M](\boldsymbol{\lambda}_l^*) \cdot c = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

By continuity, there exists a perturbation of c to $\tilde{c} \in \mathbb{R}^2_+$ such that

$$\Psi[\theta_{k-l+1}, \theta_{k-l}; a_1, a_M](\boldsymbol{\lambda}^*_l) \cdot \tilde{c} = \begin{pmatrix} G_{k-l+1} \\ G_{k-l} \end{pmatrix},$$

where $G_{k-l+1} > 0$ and $G_{k-l} < 0$. Moreover, $\Psi[\theta_j; a_1, a_M](\boldsymbol{\lambda}^*_l) \cdot \tilde{c} > 0$ for any $j > k - l + 1$ and $\Psi[\theta_j; a_1, a_M](\boldsymbol{\lambda}^*_l) \cdot \tilde{c} < 0$ for any $j < k - l$, where $\Psi[\theta_j; a_1, a_M](\boldsymbol{\lambda})$ is the submatrix of Eq. (13) for type θ_j and actions a_1 and a_M. Therefore, by Definition 8, $\boldsymbol{\lambda}^*_l$ is separable at zero for θ_{k-l+1}, since the definition holds for vector $c' \in (0, \infty)^{|A|}$ with $c'_1 = c_1$, $c'_M = c_2$, and $c'_i = 0$ otherwise. The case of maximal separability at infinity is analogous, as is the proof for the other disagreement outcomes.   *Q.E.D.*

D.2. *Mixed Learning*

Consider the mixed learning outcome $(\boldsymbol{\lambda}^*_C, C)$ in which beliefs converge to $\boldsymbol{\lambda}^*_C \in \{0, \infty\}^{|C|}$ for some subset of social types $C \subset \Theta_S$ with $|C| \in \{1, \ldots, k-1\}$ and beliefs do not converge for the remaining social types $N \equiv \Theta_S \setminus C$, where $\boldsymbol{\lambda}_C$ denotes the likelihood ratio vector $\boldsymbol{\lambda}$ restricted to a set of types C. Without loss of generality, maintain the convention that the first $|C|$ types are the convergent types, that is, $C = \{\theta_1, \ldots, \theta_{|C|}\}$, and the remaining types are the non-convergent types, that is, $N = \{\theta_{|C|+1}, \ldots, \theta_k\}$ (it is always possible to relabel the type space so that this holds).

For example, when $k = 3$, $((0, 0), \{\theta_1, \theta_2\})$ denotes the mixed outcome where θ_1 and θ_2's beliefs converge to zero and θ_3's beliefs do not converge. If $(0, 0, 0) \in \Lambda_3(\omega)$ or $(0, 0, \infty) \in \Lambda_3(\omega)$, then when $\langle \lambda_{1,t}, \lambda_{2,t} \rangle \to (0, 0)$, with positive probability the beliefs of θ_3 also converge in state ω. This is a sufficient condition to establish that $((0, 0), \{\theta_1, \theta_2\})$ almost surely does not occur in state ω. Sufficient conditions to rule out mixed outcomes in which the beliefs of two or more social types do not converge are more involved, as we also need to ensure that the neighborhood of a locally stable outcome for the non-convergent types is reached with positive probability when the beliefs of the convergent types converge. For example, to rule out the mixed outcome $(0, \theta_1)$ in which θ_1's beliefs converge to zero and θ_2 and θ_3 have cyclical learning, in addition to $(0, 0, 0) \in \Lambda_2(\omega) \cap \Lambda_3(\omega)$, we also need to show that from a neighborhood of the other stationary beliefs with $\lambda_1 = 0$, that is, $\boldsymbol{\lambda} \in \{(0, \infty, 0), (0, 0, \infty), (0, \infty, \infty)\}$, either (i) beliefs enter a neighborhood of $(0, 0, 0)$ with positive probability or (ii) $\boldsymbol{\lambda} \in \Lambda_2(\omega) \cap \Lambda_3(\omega)$. The following paragraphs formalize this idea.

We first define the concept of mixed accessibility. The concept applies to pairs of stationary beliefs in which non-convergent types whose components differ between the two belief vectors agree, which we refer to as agreement adjacent beliefs.

DEFINITION 12—Agreement Adjacent: Given a set of types $N \subset \Theta_S$, distinct stationary beliefs $\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}$ and $\boldsymbol{\lambda}'_N \in \{0, \infty\}^{|N|}$ are *agreement adjacent* if $\lambda_i = \lambda_j$ for each $\theta_i, \theta_j \in N$ such that $\lambda'_i \neq \lambda_i$ and $\lambda'_j \neq \lambda_j$.

Trivially, two stationary belief vectors that differ in only one component are agreement adjacent. Given a mixed outcome and a stationary belief for the non-convergent types, the set of stationary beliefs that are mixed accessible from this belief depends on the local stability of this belief for each non-convergent type.

DEFINITION 13—Mixed Accessible ($k \geq 2$): Given mixed outcome $(\boldsymbol{\lambda}^*_C, C)$ with $N \equiv \Theta_S \setminus C$, stationary belief $\boldsymbol{\lambda}'_N \in \{0, \infty\}^{|N|}$ is *mixed accessible* from distinct stationary belief

$\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}$ in state ω if $\boldsymbol{\lambda}'_N$ and $\boldsymbol{\lambda}'_N$ are agreement adjacent and $(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}_N) \notin \Lambda_i(\omega)$ for some $\theta_i \in N$ such that $\lambda'_i \neq \lambda_i$.

As we will show in the proof of Lemma 4′, mixed accessibility is a sufficient condition to establish that with positive probability, the likelihood ratio process either transitions between the neighborhoods of two agreement adjacent stationary beliefs or exits a neighborhood of the mixed outcome. We next define a graph to represent which stationary beliefs are mixed accessible from other stationary beliefs.

DEFINITION 14—Mixed Accessible Graph ($k \geq 2$): Given $(\boldsymbol{\lambda}^*_C, C)$ with $N \equiv \Theta_S \setminus C$, define the *mixed accessible graph* $\mathcal{G}(\boldsymbol{\lambda}^*_C, C; \omega)$ with nodes $\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}$ as follows: there is a directed edge from $\boldsymbol{\lambda}_N$ to $\boldsymbol{\lambda}'_N$ if and only if $\boldsymbol{\lambda}'_N$ is mixed accessible from $\boldsymbol{\lambda}_N$ in state ω.

We say $(\boldsymbol{\lambda}^*_C, C)$ is *reducible* in state ω if $\mathcal{G}(\boldsymbol{\lambda}^*_C, C; \omega)$ has no cycles. We refer to a node with no edges leaving it as a *terminal node*—in other words, a node from which no other nodes are mixed accessible. It follows from the definition of mixed accessibility that $\boldsymbol{\lambda}_N$ is a terminal node in state ω if and only if $(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}_N) \in \bigcap_{\theta_i \in N} \Lambda_i(\omega)$.

We use this graph to define $\Lambda_M(\omega)$ as the set of mixed outcomes that are not reducible,

$$\Lambda_M(\omega) \equiv \left\{ (\boldsymbol{\lambda}^*_C, C) \text{ a mixed outcome } \mid (\boldsymbol{\lambda}^*_C, C) \text{ is not reducible in state } \omega \right\}, \qquad (20)$$

where a mixed outcome corresponds to $\boldsymbol{\lambda}^*_C \in \{0, \infty\}^{|C|}$, $C \subset \Theta_S$, and $|C| \in \{1, \ldots, k-1\}$. This definition is equivalent to Eq. (5) when $k = 2$. Using Eq. (20), the statement of Lemma 4′ is identical to Lemma 4.

LEMMA 4′—Unstable Mixed Outcomes ($k \geq 2$): *Consider a learning environment that is identified at certainty and satisfies Assumptions 1 to 4. If mixed outcome $(\boldsymbol{\lambda}^*_C, C) \notin \Lambda_M(\omega)$, then $\Pr(\boldsymbol{\lambda}_{C,t} \to \boldsymbol{\lambda}^*_C \text{ and } \boldsymbol{\lambda}_{N,t} \text{ does not converge}) = 0$ in state ω, where $N \equiv \Theta_S \setminus C$.*

As in the case of $k = 2$, if a mixed learning outcome arises with positive probability, then it must be in $\Lambda_M(\omega)$. Therefore, if $(\boldsymbol{\lambda}^*_C, C)$ is reducible, then almost surely it does not arise. Intuitively, if $(\boldsymbol{\lambda}^*_C, C)$ is reducible, then when the beliefs of the convergent types are in a neighborhood of $\boldsymbol{\lambda}^*_C$, almost surely either the beliefs of the convergent types leave this neighborhood or the beliefs of the non-convergent types also converge. Reducibility is relatively straightforward to verify and is always satisfied in some important cases. For instance, it holds when $\gamma_i(\omega, \boldsymbol{\lambda}) < 0$ for all $\boldsymbol{\lambda} \in \{0, \infty\}^k$ and $\theta_i \in \Theta_S$ (this includes environments that are close to a correctly specified environment).[3]

PROOF OF LEMMA 4′: Fix state ω and consider mixed outcome $(\boldsymbol{\lambda}^*_C, C)$ with corresponding graph $\mathcal{G}(\boldsymbol{\lambda}^*_C, C; \omega)$ and non-convergent types $N \equiv \Theta_S \setminus C$. Suppose $(\boldsymbol{\lambda}^*_C, C)$ is reducible, that is, $(\boldsymbol{\lambda}^*_C, C) \notin \Lambda_M(\omega)$. Let $\varepsilon \in (0, e^{-E})$, where E is defined in Eq. (11), and suppose $\boldsymbol{\lambda}_1 \in \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C)) \times (0, \infty)^{|N|}$. Let $\tau \equiv \min\{t \mid \boldsymbol{\lambda}_t \notin B_\varepsilon(\boldsymbol{\lambda}^*_C) \times (0, \infty)^{|N|}\}$ be the first time that $\langle \boldsymbol{\lambda}_t \rangle$ leaves a neighborhood of the mixed outcome. We will establish the following claim: almost surely, either (i) $\tau < \infty$ or (ii) $\langle \boldsymbol{\lambda}_t \rangle$ converges for all social types.

---

[3]To see this, consider the graph induced by any mixed outcome $(\boldsymbol{\lambda}^*_C, C)$ with $N \equiv \Theta_S \setminus C$. Each node where κ non-convergent types have belief $\lambda_i = \infty$ has an edge to all agreement adjacent nodes in which $\kappa' < \kappa$ non-convergent types have belief $\lambda_i = \infty$ and does not have an edge to any other nodes. Therefore, every path terminates at node $0^{|N|}$. For any mixed outcome $(0^{|C|}, C)$, this is a convergent point. For other mixed outcomes, this is a point at which some $\theta_i \in C$'s belief eventually exits a neighborhood of $\boldsymbol{\lambda}^*_C$.

By the linearity of the likelihood ratio process, this implies the same holds whenever $\boldsymbol{\lambda}_t \in \text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*)) \times (0, \infty)^{|N|}$, and therefore, $(\boldsymbol{\lambda}_C^*, C)$ almost surely does not occur.

Step 1: Show that for any terminal node $\boldsymbol{\lambda}_N \in \mathcal{G}(\boldsymbol{\lambda}_C^*, C; \omega)$, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle \to (\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N)$ or $\tau < \infty$. Given a terminal node $\boldsymbol{\lambda}_N \in \mathcal{G}(\boldsymbol{\lambda}_C^*, C; \omega)$, as stated above, $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N) \in \bigcap_{\theta_i \in N} \Lambda_i(\omega)$. If $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N) \in \bigcap_{\theta_i \in C} \Lambda_i(\omega)$, then $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N)$ is locally stable, so by Theorem 1, when beliefs are in $B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N)$, then $\langle \boldsymbol{\lambda}_t \rangle \to (\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N)$ with probability uniformly bounded away from zero. Otherwise, if $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N) \notin \bigcap_{\theta_i \in C} \Lambda_i(\omega)$, then there exists a $\theta_i \in C$ such that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N))$, $\langle \lambda_{i,t} \rangle$ is bounded below by a process that almost surely exits $B_\varepsilon(\boldsymbol{\lambda}_C^*)$ (this also follows from the proof of Theorem 1). Therefore, $\tau < \infty$ with probability uniformly bounded away from zero. Together this implies that, starting from the ε-neighborhood of any terminal node $\boldsymbol{\lambda}_N$, with probability uniformly bounded away from zero either $\tau < \infty$ or $\langle \boldsymbol{\lambda}_t \rangle$ converges to $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N)$.

Step 2: Show that for any non-terminal node $\boldsymbol{\lambda}_N \in \mathcal{G}(\boldsymbol{\lambda}_C^*, C; \omega)$, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}_N' \in \mathcal{E}(\boldsymbol{\lambda}_N)} \text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N'))$ or $\tau < \infty$, where $\mathcal{E}(\boldsymbol{\lambda}_N)$ denotes the set of nodes that $\boldsymbol{\lambda}_N$ has edges to. Given a non-terminal node $\boldsymbol{\lambda}_N \in \mathcal{G}(\boldsymbol{\lambda}_C^*, C; \omega)$, let $U(\boldsymbol{\lambda}_N) \subset N$ denote the set of types such that $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N) \notin \Lambda_i(\omega)$ for each $\theta_i \in U(\boldsymbol{\lambda}_N)$ and $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N) \in \Lambda_i(\omega)$ for each $\theta_i \in N \setminus U(\boldsymbol{\lambda}_N)$. As stated above, $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N) \notin \bigcap_{\theta_i \in N} \Lambda_i(\omega)$ for non-terminal nodes, so $U(\boldsymbol{\lambda}_N) \neq \emptyset$.

Step 2a: We first define a space $I(\boldsymbol{\lambda}_N)$ adjacent to $B_\varepsilon(\boldsymbol{\lambda}_N)$ and show that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*) \times I(\boldsymbol{\lambda}_N))$ or $\tau < \infty$. Given a set of types $u \in \mathcal{P}(U(\boldsymbol{\lambda}_N))$, where $\mathcal{P}(\cdot)$ denotes the power set, let $I_{u,i}(\boldsymbol{\lambda}_N) \equiv [\varepsilon, 1/\varepsilon]$ if $\theta_i \in u$ and $I_{u,i}(\boldsymbol{\lambda}_N) \equiv B_\varepsilon((\boldsymbol{\lambda}_N)_i)$ if $\theta_i \in N \setminus u$. Define $I_u(\boldsymbol{\lambda}_N) \equiv \prod_{\theta_i \in N} I_{u,i}(\boldsymbol{\lambda}_N)$ for each $u \in \mathcal{P}(U(\boldsymbol{\lambda}_N))$ and $I(\boldsymbol{\lambda}_N) \equiv \bigcup_{u \in \mathcal{P}(U(\boldsymbol{\lambda}_N)) \setminus \emptyset} I_u(\boldsymbol{\lambda}_N)$. In other words, $I(\boldsymbol{\lambda}_N)$ is the space in which the beliefs of subsets of types in $U(\boldsymbol{\lambda}_N)$ are in $[\varepsilon, 1/\varepsilon]$ and the beliefs of the remaining non-convergent types are in the ε-neighborhood of $\boldsymbol{\lambda}_N$. By the proof of Theorem 1, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, $\langle \lambda_{i,t} \rangle$ exits $B_\varepsilon((\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N)_i)$ for some $\theta_i \in U(\boldsymbol{\lambda}_N) \cup C$. Combined with $\varepsilon < e^{-E}$, which ensures that $\langle \lambda_{i,t} \rangle$ does not enter $B_\varepsilon(\{0, \infty\} \setminus (\boldsymbol{\lambda}_N)_i)$ in the same period it exits $B_\varepsilon((\boldsymbol{\lambda}_N)_i)$ for any $\theta_i \in N$, this implies that with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*) \times I(\boldsymbol{\lambda}_N))$ or $\tau < \infty$.

Step 2b: We next show that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*) \times I(\boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}_N' \in \mathcal{E}(\boldsymbol{\lambda}_N)} \text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N'))$ or $\tau < \infty$. First consider $u \in \mathcal{P}(U(\boldsymbol{\lambda}_N))$ such that $(\boldsymbol{\lambda}_N)_i = 0$ for some $\theta_i \in u$. Suppose $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*) \times I_u(\boldsymbol{\lambda}_N))$. Note $I_{u,i}(\boldsymbol{\lambda}_N) = [\varepsilon, 1/\varepsilon]$ for $\theta_i \in u$ and $I_{u,i}(\boldsymbol{\lambda}_N) \equiv B_\varepsilon((\boldsymbol{\lambda}_N)_i)$ for $\theta_i \in N \setminus u$. Let $V_\infty(\boldsymbol{\lambda}_N, u) \subset \{0, \infty\}^{|N|}$ denote the set of stationary beliefs for non-convergent types in which $\theta_i \in u$ has belief ∞, $\theta_i \in N \setminus u$ such that $(\boldsymbol{\lambda}_N)_i = \infty$ has belief ∞, and $\theta_i \in N \setminus u$ such that $(\boldsymbol{\lambda}_N)_i = 0$ has belief zero or infinity. Then there exists a finite sequence of actions a_M such that, starting from any belief in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*) \times I_u(\boldsymbol{\lambda}_N))$, $\langle \boldsymbol{\lambda}_{N,t} \rangle$ enters $\bigcup_{\boldsymbol{\lambda}_N' \in V_\infty(\boldsymbol{\lambda}_N, u)} \text{int}(B_\varepsilon(\boldsymbol{\lambda}_N'))$.[4] If $\langle \boldsymbol{\lambda}_{C,t} \rangle$ exits $B_\varepsilon(\boldsymbol{\lambda}_C^*)$ during this sequence, then $\tau < \infty$;

---

[4]To see this, note that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}_C^*) \times I_u(\boldsymbol{\lambda}_N))$, then $\langle \lambda_{i,t} \rangle$ is in $[\varepsilon, 1/\varepsilon]$ for $\theta_i \in u$. Therefore, we can construct a finite sequence of a_M actions such that following this sequence, $\langle \lambda_{i,t} \rangle$ is in $\text{int}(B_\varepsilon(\infty))$ for each $\theta_i \in u$. For $\theta_i \in N \setminus u$ such that $(\boldsymbol{\lambda}_N)_i = \infty$, $I_{u,i}(\boldsymbol{\lambda}_N) \equiv B_\varepsilon(\infty)$, and therefore, $\langle \lambda_{i,t} \rangle$ remains in $\text{int}(B_\varepsilon(\infty))$ following this sequence. For $\theta_i \in N \setminus u$ such that $(\boldsymbol{\lambda}_N)_i = 0$, $I_{u,i}(\boldsymbol{\lambda}_N) \equiv B_\varepsilon(0)$. If $\langle \lambda_{i,t} \rangle$ remains in $\text{int}(B_\varepsilon(0))$ or enters $\text{int}(B_\varepsilon(\infty))$ following this sequence for each such $\theta_i \in N \setminus u$, then we are done. Otherwise, continue

otherwise, $\langle \boldsymbol{\lambda}_t \rangle$ is in $\bigcup_{\boldsymbol{\lambda}'_N \in V_\infty(\boldsymbol{\lambda}_N, u)} \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}'_N))$. Each belief $\boldsymbol{\lambda}'_N \in V_\infty(\boldsymbol{\lambda}_N, u)$ is agreement adjacent to $\boldsymbol{\lambda}_N$, as a subset of types in N have belief 0 at $\boldsymbol{\lambda}_N$ and ∞ at $\boldsymbol{\lambda}'_N$, and all other types in N have the same belief at $\boldsymbol{\lambda}_N$ and $\boldsymbol{\lambda}'_N$. By definition of $V_\infty(\boldsymbol{\lambda}_N, u)$, for each $\boldsymbol{\lambda}'_N \in V_\infty(\boldsymbol{\lambda}_N, u)$, $(\boldsymbol{\lambda}_N)_i \neq (\boldsymbol{\lambda}'_N)_i$ for $\theta_i \in u$ such that $(\boldsymbol{\lambda}_N)_i = 0$. Further, $(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}_N) \notin \Lambda_i(\omega)$ for each $\theta_i \in u$. Therefore, each $\boldsymbol{\lambda}'_N \in V_\infty(\boldsymbol{\lambda}_N, u)$ is mixed accessible from $\boldsymbol{\lambda}_N$, which implies $V_\infty(\boldsymbol{\lambda}_N, u) \subset \mathcal{E}(\boldsymbol{\lambda}_N)$. This establishes that, given $u \in \mathcal{P}(U(\boldsymbol{\lambda}_N))$ such that $(\boldsymbol{\lambda}_N)_i = 0$ for some $\theta_i \in u$, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I_u(\boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}'_N \in \mathcal{E}(\boldsymbol{\lambda}_N)} \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}'_N))$ or $\tau < \infty$. Next consider $u \in \mathcal{P}(U(\boldsymbol{\lambda}_N))$ such that $(\boldsymbol{\lambda}_N)_i = \infty$ for some $\theta_i \in u$. Let $V_0(\boldsymbol{\lambda}_N, u)$ denote the set of stationary beliefs for non-convergent types in which $\theta_i \in u$ has belief zero, $\theta_i \in N \setminus u$ such that $(\boldsymbol{\lambda}_N)_i = 0$ has belief zero, and $\theta_i \in N \setminus u$ such that $(\boldsymbol{\lambda}_N)_i = \infty$ has belief zero or infinity. Then substituting a_1 for a_M and $V_0(\boldsymbol{\lambda}_N, u)$ for $V_\infty(\boldsymbol{\lambda}_N, u)$, by similar reasoning, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I_u(\boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}'_N \in \mathcal{E}(\boldsymbol{\lambda}_N)} \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}'_N))$ or $\tau < \infty$. Given that one of these two cases applies to each $u \in \mathcal{P}(U(\boldsymbol{\lambda}_N))$ and $I(\boldsymbol{\lambda}_N) \equiv \bigcup_{u \in \mathcal{P}(U(\boldsymbol{\lambda}_N)) \setminus \emptyset} I_u(\boldsymbol{\lambda}_N)$, this establishes that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I(\boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}'_N \in \mathcal{E}(\boldsymbol{\lambda}_N)} \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}'_N))$ or $\tau < \infty$.

Step 3: Show that for any non-terminal node $\boldsymbol{\lambda}_N \in \mathcal{G}(\boldsymbol{\lambda}^*_C, C; \omega)$, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}'_N \in \mathcal{T}} \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}'_N))$ or $\tau < \infty$, where $\mathcal{T}$ denotes the set of terminal nodes. Given $\Lambda_M(\omega)$ is empty, $(\boldsymbol{\lambda}^*_C, C)$ is reducible and therefore, $\mathcal{G}(\boldsymbol{\lambda}^*_C, C; \omega)$ has no cycles. Therefore, starting from any non-terminal node $\boldsymbol{\lambda}_N \in \mathcal{G}(\boldsymbol{\lambda}^*_C, C; \omega)$ and iterating Step 2 a finite number of times ensures that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}_N))$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}'_N \in \mathcal{T}} \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}'_N))$ or $\tau < \infty$.

Step 4: Finally, we show that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I)$, where $I \equiv (0, \infty)^{|N|} \setminus \bigcup_{\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}} B_\varepsilon(\boldsymbol{\lambda}_N)$, then almost surely, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times \bigcup_{\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}} B_\varepsilon(\boldsymbol{\lambda}_N))$ or $\tau < \infty$. Consider $u \subset N$. Similarly to above, let $I_{u,i} \equiv [\varepsilon, 1/\varepsilon]$ if $\theta_i \in u$, $I_{u,i} \equiv B_\varepsilon(0) \cup B_\varepsilon(\infty)$ if $\theta_i \in N \setminus u$, and $I_u \equiv \prod_{\theta_i \in N} I_{u,i}$. Then by similar reasoning to Footnote 4, there exists a finite sequence of a_1 actions such that, starting from any belief in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I_u)$, $\langle \boldsymbol{\lambda}_{N,t} \rangle$ enters $\bigcup_{\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}} \text{int}(B_\varepsilon(\boldsymbol{\lambda}_N))$ following this sequence. If $\langle \boldsymbol{\lambda}_{C,t} \rangle$ exits $B_\varepsilon(\boldsymbol{\lambda}^*_C)$ during this sequence, then $\tau < \infty$; otherwise, $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times \bigcup_{\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}} B_\varepsilon(\boldsymbol{\lambda}_N))$. Given that this sequence is finite and occurs with probability uniformly bounded away from zero across $B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I_u$, and such a sequence exists for each $u \subset N$, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I)$, then almost surely $\langle \boldsymbol{\lambda}_t \rangle$ enters $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C)) \times \bigcup_{\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}} \text{int}(B_\varepsilon(\boldsymbol{\lambda}_N))$ or $\tau < \infty$.

Taken together, Steps 2–4 establish that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C)) \times (0, \infty)^{|N|}$, then with probability uniformly bounded away from zero, either $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}_N \in \mathcal{T}} \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}_N))$ or $\tau < \infty$. It follows from Theorem 1 that when $\langle \boldsymbol{\lambda}_t \rangle$ enters $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}_N))$ for any $\boldsymbol{\lambda}_N \notin \mathcal{T}$, it almost surely exits $B_\varepsilon(\boldsymbol{\lambda}^*_C, \boldsymbol{\lambda}_N)$, and from Step 4 that when $\langle \boldsymbol{\lambda}_t \rangle$ enters $\text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I)$, it almost surely exits $B_\varepsilon(\boldsymbol{\lambda}^*_C) \times I$. Therefore, when $\boldsymbol{\lambda}_1 \in \text{int}(B_\varepsilon(\boldsymbol{\lambda}^*_C)) \times$

---

repeating a_M until all $\theta_i \in N \setminus u$ with $(\boldsymbol{\lambda}_N)_i = 0$ have beliefs in $\text{int}(B_\varepsilon(0) \cup B_\varepsilon(\infty))$. Given that there are a finite number of such types and, for any such type, a finite number of a_M actions will move its beliefs from $[\varepsilon, 1/\varepsilon]$ to $\text{int}(B_\varepsilon(\infty))$, this will hold following a finite number of additional a_M actions. Following these additional a_M actions, $\langle \boldsymbol{\lambda}_{i,t} \rangle$ remains in $\text{int}(B_\varepsilon(\infty))$ for $\theta_i \in u$, as does $\langle \boldsymbol{\lambda}_{i,t} \rangle$ for $\theta_i \in N \setminus u$ such that $(\boldsymbol{\lambda}_N)_i = \infty$. Therefore, following this sequence, $\langle \boldsymbol{\lambda}_{N,t} \rangle$ is in $\text{int}(B_\varepsilon(\boldsymbol{\lambda}'_N))$ for some $\boldsymbol{\lambda}'_N \in V_\infty(\boldsymbol{\lambda}_N, u)$.

$(0, \infty)^{|N|}$,

$$\Pr\left(\tau < \infty \text{ or } \boldsymbol{\lambda}_t \in \bigcup_{\boldsymbol{\lambda}_N \in \mathcal{T}} B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N) \text{ i.o.} \right) = 1.$$

If $\langle \boldsymbol{\lambda}_t \rangle$ is in $\bigcup_{\boldsymbol{\lambda}_N \in \mathcal{T}} B_\varepsilon(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N)$ infinitely often, then almost surely either $\langle \boldsymbol{\lambda}_t \rangle$ converges to $(\boldsymbol{\lambda}_C^*, \boldsymbol{\lambda}_N)$ for some $\boldsymbol{\lambda}_N \in \mathcal{T}$ or $\tau < \infty$. This establishes the claim.   *Q.E.D.*

D.3. *Learning Characterization*

To establish almost sure convergence, we define an analogous graph to Definition 14 for the case in which all social types are non-convergent types.

DEFINITION 15—Accessible Graph ($k \geq 2$): Define the *accessible graph* $\mathcal{G}(\omega)$ with nodes $\boldsymbol{\lambda} \in \{0, \infty\}^k$ as follows: there is a directed edge from $\boldsymbol{\lambda}$ to $\boldsymbol{\lambda}'$ if and only if $\boldsymbol{\lambda}'$ is mixed accessible from $\boldsymbol{\lambda}$ in state ω.

It follows from the definition of mixed accessibility that $\boldsymbol{\lambda}$ is a terminal node if and only if $\boldsymbol{\lambda} \in \Lambda(\omega)$. Given the definitions of $\Lambda_M(\omega)$ and maximal accessibility for $k > 2$, the statement of Lemma 7$'$ mirrors Lemma 7.

LEMMA 7$'$—Belief Convergence ($k > 2$): *Consider a learning environment that is identified at certainty and satisfies Assumptions 1 to 4. If $\Lambda(\omega)$ contains an agreement outcome or maximally accessible disagreement outcome, $\Lambda_M(\omega) = \emptyset$, and $\mathcal{G}(\omega)$ has no cycles, then for any initial belief $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, there exists a random variable $\boldsymbol{\lambda}_\infty$ with $\mathrm{supp}(\boldsymbol{\lambda}_\infty) = \Lambda(\omega)$ such that $\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}_\infty$ almost surely in state ω.*[5]

PROOF: Fix state ω. Suppose $\Lambda_M(\omega) = \emptyset$ and $\mathcal{G}(\omega)$ has no cycles. Let $\varepsilon \in (0, e^{-E})$, where E is defined in Eq. (11). It follows from the definition of mixed accessibility that $\boldsymbol{\lambda} \in \mathcal{G}(\omega)$ is a terminal node if and only if $\boldsymbol{\lambda} \in \Lambda(\omega)$. Given a terminal node $\boldsymbol{\lambda} \in \Lambda(\omega)$, by Theorem 1, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $B_\varepsilon(\boldsymbol{\lambda})$, then $\langle \boldsymbol{\lambda}_t \rangle \to \boldsymbol{\lambda}$ with probability uniformly bounded away from zero. By analogous reasoning to Step 2 in the proof of Lemma 4$'$, for any non-terminal node $\boldsymbol{\lambda} \in \mathcal{G}(\omega)$, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\mathrm{int}(B_\varepsilon(\boldsymbol{\lambda}))$, then with probability uniformly bounded away from zero, $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}' \in \mathcal{E}(\boldsymbol{\lambda})} \mathrm{int}(B_\varepsilon(\boldsymbol{\lambda}'))$, where $\mathcal{E}(\boldsymbol{\lambda})$ denotes the set of nodes that $\boldsymbol{\lambda}$ has edges to. Given $\mathcal{G}(\omega)$ has no cycles, starting from any non-terminal node $\boldsymbol{\lambda} \in \mathcal{G}(\omega)$, when $\langle \boldsymbol{\lambda}_t \rangle$ is in $\mathrm{int}(B_\varepsilon(\boldsymbol{\lambda}))$, then with probability uniformly bounded away from zero, $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda}' \in \Lambda(\omega)} \mathrm{int}(B_\varepsilon(\boldsymbol{\lambda}'))$. When $\langle \boldsymbol{\lambda}_t \rangle$ is in $I \equiv (0, \infty)^k \setminus \bigcup_{\boldsymbol{\lambda} \in \{0, \infty\}^k} B_\varepsilon(\boldsymbol{\lambda})$, then by similar reasoning to Step 4 in the proof of Lemma 4$'$, almost surely $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda} \in \{0, \infty\}^k} \mathrm{int}(B_\varepsilon(\boldsymbol{\lambda}))$. Taken together, this establishes that when $\langle \boldsymbol{\lambda}_t \rangle$ is in $(0, \infty)^k$, then with probability uniformly bounded away from zero, $\langle \boldsymbol{\lambda}_t \rangle$ enters $\bigcup_{\boldsymbol{\lambda} \in \Lambda(\omega)} \mathrm{int}(B_\varepsilon(\boldsymbol{\lambda}))$. Further, it follows from Theorem 1 that when $\langle \boldsymbol{\lambda}_t \rangle$ enters $\mathrm{int}(B_\varepsilon(\boldsymbol{\lambda}))$ for any $\boldsymbol{\lambda} \in \{0, \infty\}^k \setminus \Lambda(\omega)$, it almost surely exits $B_\varepsilon(\boldsymbol{\lambda})$. Given that $\langle \boldsymbol{\lambda}_t \rangle$ also almost surely exits I, it follows that starting from any $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, $\Pr(\boldsymbol{\lambda}_t \in \bigcup_{\boldsymbol{\lambda} \in \Lambda(\omega)} B_\varepsilon(\boldsymbol{\lambda}) \text{ i.o.}) = 1$. Therefore, almost surely $\langle \boldsymbol{\lambda}_t \rangle$ converges to some $\boldsymbol{\lambda} \in \Lambda(\omega)$.   *Q.E.D.*

---

[5]An alternative condition to $\mathcal{G}(\omega)$ has no cycles is there exists a $\theta_i \in \Theta_S$ such that $\sup_{\boldsymbol{\lambda}_{-i} \in [0, \infty]^{k-1}} \gamma_i(\lambda_i = 0, \boldsymbol{\lambda}_{-i}) < 0$ or $\inf_{\boldsymbol{\lambda}_{-i} \in [0, \infty]^{k-1}} \gamma_i(\lambda_i = \infty, \boldsymbol{\lambda}_{-i}) > 0$. This follows from the observation in Lemma 4$'$ that either beliefs converge or visit each mixed outcome with $|C| = 1$ infinitely often. If the latter occurs with positive probability, then almost surely $\langle \lambda_{i,t} \rangle \to \lambda_\infty \subset \{0, \infty\}$, which contradicts $\Lambda_M(\omega) = \emptyset$.

The statement of the learning characterization for $k > 2$ is analogous to Theorem 4, using the generalized definitions of maximal accessibility (Definition 7′) and $\Lambda_M(\omega)$ (Eq. (20)).

THEOREM 4′—Learning Characterization ($k > 2$): *Consider a learning environment that is identified at certainty and satisfies Assumptions* 1 *to* 4. *When the state is L:*
  (i) **Correct** *learning occurs with positive probability if and only if* $0^k \in \Lambda(L)$.
  (ii) **Incorrect** *learning occurs with positive probability if and only if* $\infty^k \in \Lambda(L)$.
  (iii) **Entrenched Disagreement** *occurs with positive probability if* $\Lambda(L)$ *contains a maximally accessible disagreement outcome and almost surely does not occur if* $\Lambda(L)$ *contains no disagreement outcome. Each maximally accessible* disagreement *outcome in* $\Lambda(L)$ *occurs with positive probability.*
  (iv) **Cyclical Learning** *occurs almost surely for all social types if* $\Lambda(L) = \emptyset$ *and* $\Lambda_M(L) = \emptyset$, *occurs almost surely for at least one social type if* $\Lambda(L) = \emptyset$, *and almost surely does not occur if* $\Lambda(L)$ *contains an agreement outcome or maximally accessible disagreement outcome,* $\Lambda_M(L) = \emptyset$, *and* $\mathcal{G}(L)$ *has no cycles.*
An analogous result holds in state R.

The proof mirrors the case of two social types: it directly follows from Lemma 3, Theorems 1 and 2, Theorem 3′, and Lemmas 4′ and 7′.

APPENDIX E: BELIEF-DEPENDENT MODELS OF INFERENCE

In this section, we extend our framework to allow a type's model of inference to vary with its belief about the state. We show that with this extension, our framework nests Rabin and Schrag (1999) and Epstein, Noor, and Sandroni (2010).

E.1. *Framework*

Modify a type's model of inference as follows. Given likelihood ratio $\boldsymbol{\lambda} \in [0, \infty]^k$, type θ_i has subjective private signal distribution $\hat{F}_i^\omega(s; \boldsymbol{\lambda})$ in state ω and subjective type distribution $\hat{\pi}_i(\theta; \boldsymbol{\lambda})$. An agent uses likelihood ratio $\boldsymbol{\lambda}_t$ to interpret signal $\tilde{s}_t$ or action $\tilde{a}_t$ at time t. Maintain the assumption from Section 2 that $\hat{F}_i^L(\cdot, \boldsymbol{\lambda})$ and $\hat{F}_i^R(\cdot, \boldsymbol{\lambda})$ are mutually absolutely continuous with full support on $\mathcal{S}$ for each $\boldsymbol{\lambda} \in [0, \infty]^k$. Further, social and autarkic types believe that the signal is uniformly informative. When signals are aligned, this can be written as follows: for all $s \in [0, 1]$, either $\hat{F}_i^L(s; \boldsymbol{\lambda}) = \hat{F}_i^R(s; \boldsymbol{\lambda}) = 0$ for all $\boldsymbol{\lambda} \in [0, \infty]^k$, $\hat{F}_i^L(s; \boldsymbol{\lambda}) = \hat{F}_i^R(s; \boldsymbol{\lambda}) = 1$ for all $\boldsymbol{\lambda} \in [0, \infty]^k$, or $\inf_{\boldsymbol{\lambda} \in [0,\infty]^k} \hat{F}_i^L(s; \boldsymbol{\lambda}) - \hat{F}_i^R(s; \boldsymbol{\lambda}) > 0$, with the final case holding for some $s \in [0, 1]$.

As in Section 2, we focus on settings in which social types believe that actions are informative. We need to modify Assumption 3 so that this holds uniformly across $[0, \infty]^k$.

ASSUMPTION 3′—Informative Actions: *For actions* $a \in \{a_1, a_M\}$, *there exists an autarkic type* $\theta_j \in \Theta_A$ *with* $\pi(\theta_j) > 0$ *that plays* a *with probability uniformly bounded away from zero across* $[0, \infty]^k$, *and each social type* $\theta_i \in \Theta_S$ *believes that such an autarkic type exists with probability uniformly bounded away from zero,* $\inf_{\boldsymbol{\lambda} \in [0,\infty]^k} \hat{\pi}_i(\theta_j; \boldsymbol{\lambda}) > 0$.

For technical reasons, we also make the following continuity assumption.

ASSUMPTION 5—Continuity: *For each $\theta_i \in \Theta$, the mapping $\boldsymbol{\lambda} \mapsto (\hat{F}_i^L, \hat{F}_i^R, \hat{\pi}_i)$ is continuous under the total variation norm except at at most a finite number of interior likelihood ratios $\boldsymbol{\lambda} \in (0, \infty)^k$, and $\boldsymbol{\lambda} \to 1/(1 + d\hat{F}_i^L/d\hat{F}_i^R(s; \boldsymbol{\lambda}))$ is continuous at $\boldsymbol{\lambda} \in \{0, \infty\}^k$.*

Substituting Assumption 3′ for Assumption 3 and adding Assumption 5, the modified version of Eq. (1) is

$$\hat{\psi}_i(a_m | \omega, \boldsymbol{\lambda}) \equiv \sum_{j=1}^n \hat{\pi}_i(\theta_j; \boldsymbol{\lambda}) \big(\hat{F}_i^\omega(\bar{s}_{j,m}(\lambda_j; \boldsymbol{\lambda}); \boldsymbol{\lambda}) - \hat{F}_i^\omega(\bar{s}_{j,m-1}(\lambda_j; \boldsymbol{\lambda}); \boldsymbol{\lambda}) \big), \qquad (21)$$

where $\bar{s}_{j,m}(\lambda_j; \boldsymbol{\lambda})$ denotes the signal cutoff for θ_j when it has belief λ_j and social types have belief $\boldsymbol{\lambda}$. Note $\bar{s}_{j,m}$ depends on $(\hat{F}_j^L, \hat{F}_j^R)$, and hence, when these distributions depend on $\boldsymbol{\lambda}$, so does $\bar{s}_{j,m}$. The proof of Lemma 2 continues to hold for Eq. (21) with minor modifications.[6] Theorems 1 to 6 follow.

E.2. *Nesting Under- and Overreaction in Epstein, Noor, and Sandroni (2010)*

Epstein, Noor, and Sandroni (2010) considered an individual learning model where an agent under- or overreacts to signals. They parameterized this bias with the following updating rule: an agent with prior $p \in [0, 1]$ who observes signal realization $s \in \mathcal{S}$ updates her posterior to

$$\Pr(\omega = R | s, p) = (1 - \alpha) \underbrace{\left(\frac{ps}{ps + (1 - p)(1 - s)} \right)}_{\text{correct posterior}} + \alpha p \qquad (22)$$

for some $\alpha \leq 1$. Underreaction corresponds to $\alpha > 0$, overreaction corresponds to $\alpha < 0$, and the correctly specified model corresponds to $\alpha = 0$. This parametric form of under- and overreaction can be represented in the individual learning version of our extended framework as follows. Equation (22) uniquely maps to a type in our framework that forms

---

[6]Aside from minor changes to notation and a straightforward application of the continuity assumed in Assumption 5, there are two main changes. To establish the uniform bound for $a \in \{a_1, a_M\}$ and bounded informativeness for $a \in \mathcal{A}$, it is necessary to account for the subjective type and signal distributions' dependence on $\boldsymbol{\lambda}$. Let $\theta_j \in \Theta_A$ be an autarkic type that $\theta_i \in \Theta_S$ believes satisfies Assumption 3′ for action a_1 and $\bar{s}_{j,1}^* \equiv \inf_{\boldsymbol{\lambda} \in [0,\infty]^k} \bar{s}_{j,1}(\frac{p_0}{1-p_0}; \boldsymbol{\lambda})$. Then the analogue of Eq. (6) is

$$\frac{\hat{\psi}_i(a_1 | R, \boldsymbol{\lambda})}{\hat{\psi}_i(a_1 | L, \boldsymbol{\lambda})} \leq \frac{\hat{\pi}_i(\theta_j; \boldsymbol{\lambda}) \hat{F}_i^R(\bar{s}_{j,1}^*; \boldsymbol{\lambda}) + \hat{\pi}_i(\Theta_S \cup \Theta_A \setminus \{\theta_j\}; \boldsymbol{\lambda})}{\hat{\pi}_i(\theta_j; \boldsymbol{\lambda}) \hat{F}_i^L(\bar{s}_{j,1}^*; \boldsymbol{\lambda}) + \hat{\pi}_i(\Theta_S \cup \Theta_A \setminus \{\theta_j\}; \boldsymbol{\lambda})}$$

$$\leq \sup_{\boldsymbol{\lambda} \in [0,\infty]^k} \frac{\left(\inf_{\boldsymbol{\lambda}' \in [0,\infty]^k} \hat{\pi}_i(\theta_j; \boldsymbol{\lambda}') \right) \hat{F}_i^R(\bar{s}_{j,1}^*; \boldsymbol{\lambda}) + 1 - \inf_{\boldsymbol{\lambda}' \in [0,\infty]^k} \hat{\pi}_i(\theta_j; \boldsymbol{\lambda}')}{\left(\inf_{\boldsymbol{\lambda}' \in [0,\infty]^k} \hat{\pi}_i(\theta_j; \boldsymbol{\lambda}') \right) \hat{F}_i^L(\bar{s}_{j,1}^*; \boldsymbol{\lambda}) + 1 - \inf_{\boldsymbol{\lambda}' \in [0,\infty]^k} \hat{\pi}_i(\theta_j; \boldsymbol{\lambda}')} < 1,$$

where the last line follows from Assumption 3′, which ensures that $\inf_{\boldsymbol{\lambda} \in [0,\infty]^k} \hat{\pi}_i(\theta_j; \boldsymbol{\lambda}) > 0$, and the uniform informativeness of the subjective signal distributions, which ensures that $\inf_{\boldsymbol{\lambda} \in [0,\infty]^k} (\hat{F}_i^L(s_{j,1}^*; \boldsymbol{\lambda}) - \hat{F}_i^R(s_{j,1}^*; \boldsymbol{\lambda})) > 0$. Similar logic establishes that $\hat{\psi}_i(a | \omega, \boldsymbol{\lambda})$ is uniformly bounded away from 0 for all $\boldsymbol{\lambda} \in [0, \infty]^k$, $a \in \mathcal{A}$, and $\omega \in \{L, R\}$, and therefore, a is boundedly informative.

subjective posterior

$$\hat{s}(s, \lambda) = \frac{(1 - \alpha)\left(\dfrac{s}{s\lambda + (1 - s)}\right) + \alpha\left(\dfrac{1}{1 + \lambda}\right)}{1 + (1 - \alpha)\left(\dfrac{(1 - \lambda)s}{s\lambda + (1 - s)}\right) + \alpha\left(\dfrac{1 - \lambda}{1 + \lambda}\right)} \tag{23}$$

following signal realization $s \in \mathcal{S}$ when it has belief $\lambda \in (0, \infty)$. Equation (22) does not map into a unique $\hat{s}(s, \lambda)$ at $\lambda \in \{0, \infty\}$, since the prior and the posterior are the same regardless of the signal realization. Since our learning characterization utilizes the limit of $\hat{s}(s, \lambda)$ as $\lambda \to \{0, \infty\}$, we need to specify how the signal is interpreted at certainty to close the model. At $\lambda = 0$, we use Eq. (23) evaluated at $\lambda = 0$. Equation (23) is not well-defined at $\lambda = \infty$, so we define

$$\hat{s}(s, \infty) \equiv \lim_{\lambda \to \infty} \hat{s}(s, \lambda) = \frac{s}{(1 - \alpha)(1 - s) + (1 + \alpha)s}.$$

This is the unique subjective posterior that satisfies the continuity property required by Lemma 2.[7] This set-up satisfies the properties in Lemma 2, so our learning characterization applies.

E.3. *Nesting Confirmation Bias in Rabin and Schrag (1999)*

Rabin and Schrag (1999) considered an individual learning model where an agent exhibits confirmation bias. The agent observes a signal that takes one of two possible values, s_L or s_R, where s_ω is more likely in state ω than state ω'. Confirmation bias takes the following form: if the agent observes s_ω when she believes ω' is more likely, then with probability $q \in (0, 1)$ she misinterprets the signal realization as $s_{\omega'}$. To represent this model in the individual learning version of our extended framework, we make one additional change to allow multiple signal realizations to induce the same posterior belief. This allows $\hat{s}$ to map two signal realizations that induce the same true posterior to different subjective posteriors. Given this minor extension, this form of confirmation bias can be represented as follows. Suppose $\mathcal{S} = \{l_1, l_2, r_1, r_2\}$. Assume $\Pr(l_1 \text{ or } l_2 | \omega = L) = \Pr(r_1 \text{ or } r_2 | \omega = R) = s > 1/2$, conditional on observing l_1 or l_2, l_2 is realized with probability q, and similarly for r_2. Signal realizations l_1 and l_2 induce the same true posterior, as do r_1 and r_2. When $\lambda > 1$, the agent interprets the signal as if $\hat{\psi}(l_1 | L, \lambda) = s$, $\hat{\psi}(l_1 | R, \lambda) = 1 - s$, $\hat{\psi}(l_2 | L, \lambda) = \hat{\psi}(r_1 | L, \lambda) = \hat{\psi}(r_2 | L, \lambda) = (1 - s)/3$, and $\hat{\psi}(l_2 | R, \lambda) = \hat{\psi}(r_1 | R, \lambda) = \hat{\psi}(r_2 | R, \lambda) = s/3$. Similarly, if $\lambda \leq 1$, the agent interprets the signal as if $\hat{\psi}(r_1 | R, \lambda) = s$, $\hat{\psi}(r_1 | L, \lambda) = 1 - s$, $\hat{\psi}(l_1 | L, \lambda) = \hat{\psi}(l_2 | L, \lambda) = \hat{\psi}(r_2 | L, \lambda) = s/3$, and $\hat{\psi}(l_1 | R, \lambda) = \hat{\psi}(l_2 | R, \lambda) = \hat{\psi}(r_2 | R, \lambda) = (1 - s)/3$. This set-up satisfies the properties in Lemma 2, so our learning characterization applies.

---

[7]In an individual learning setting, any pair of subjective signal distributions that induce the same $\hat{s}$ must satisfy $\frac{\hat{\psi}(s | R, \lambda)}{\hat{\psi}(s | L, \lambda)} = \frac{\hat{s}(s, \lambda)}{1 - \hat{s}(s, \lambda)}$, so $\hat{s}$ determines the properties required by Lemma 2. A consequence of this is that any misspecified distribution that rationalizes $\hat{s}$ will lead to the same behavior. In Bohren and Hauser (2021) we showed that there exist subjective distributions $\hat{F}^L$ and $\hat{F}^R$ that rationalize this $\hat{s}$ and satisfy Assumption 1, Assumption 3′, and Assumption 5.

REFERENCES

BERK, R. H. (1966): "Limiting Behavior of Posterior Distributions When the Model Is Incorrect," *The Annals of Mathematical Statistics*, 37, 51–58. [6]

BOHREN, A. (2016): "Informational Herding With Model Misspecification," *Journal of Economic Theory*, 163, 222–247. [9]

BOHREN, J. A., AND D. N. HAUSER (2021): "Posteriors as Signals in Misspecified Models," Report. [22]

EPSTEIN, L. G., J. NOOR, AND A. SANDRONI (2010): "Non-Bayesian Learning," *The B.E. Journal of Theoretical Economics*, 10 (1), 3. [20,21]

RABIN, M., AND J. L. SCHRAG (1999): "First Impressions Matter: A Model of Confirmatory Bias," *The Quarterly Journal of Economics*, 114, 37–82. [20,22]