Learning with Model Misspecification: Characterization and Robustness

Online Appendix

J. Aislinn Bohren Daniel N. Hauser

February 3, 2021

First version: May 15, 2017

B Derivation of Examples from Section 3

B.1 Example 1: Partisan Bias

Trivially, 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 that satisfies Assumption 3, 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 distribution is 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 $\hat{\psi}_1(L|\omega,\lambda_1) = \pi(\theta_1)F^\omega(1/(1+\lambda_1)) + \pi(\theta_2)F^\omega(.5)$, whereas the true probability of action L is $\psi(L|\omega,\lambda_1) = \pi(\theta_1)F^\omega((1/(1+\lambda_1))^{1/\nu}) + \pi(\theta_2)F^\omega(.5^{1/\nu})$. The construction of $\gamma_1(L,0)$ in Section 3.3 follows from evaluating these expressions at $\lambda_1 = 0$. Similarly,

$$\begin{aligned}\gamma_1(L,\infty) \quad &= \quad \underbrace{\pi(\theta_2)F^L(.5^{1/\nu})\log\frac{F^R(.5)}{F^L(.5)}}_{L-action} \\ &\quad + \underbrace{(\pi(\theta_1)+\pi(\theta_2)(1-F^L(.5^{1/\nu})))\log\frac{\pi(\theta_1)+\pi(\theta_2)(1-F^R(.5))}{\pi(\theta_1)+\pi(\theta_2)(1-F^L(.5))}}_{R-action}\end{aligned}$$

follows from evaluating these expressions at $\lambda_1 = \infty$.

We next characterize how $\Lambda(\omega)$ depends on ν. We write $\gamma_1(\omega, \boldsymbol{\lambda}; \nu)$ and $\Lambda(\omega; \nu)$ to make these expressions' dependence on ν explicit. To simplify notation, define $\alpha_\nu \equiv F^L(.5^{1/\nu})$ 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, $F^R(.5) = 1 - F^L(.5) = 1 - \alpha_1$ 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, and symmetry implies that $\alpha_1 > 1/2$.

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

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

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 $\nu = 0$, θ_2 chooses action R for all signals, $\alpha_0 = 0$. Therefore, $\gamma_1(L, \infty; 0) = \log \frac{1-\pi_A(1-\alpha_1)}{1-\pi_A\alpha_1} > 0$. This establishes that 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, i.e. whether $0 \in \Lambda(L; \nu)$, we need to determine the sign of

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

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 $\nu = 0$, θ_2 chooses action R for all signals, $\alpha_0 = 0$. Therefore,

$$\begin{aligned} \gamma_1(L, 0; 0) \quad &= \quad (1 - \pi_A) \log \frac{1 - \pi_A\alpha_1}{1 - \pi_A(1-\alpha_1)} + \pi_A \log \frac{\alpha_1}{1-\alpha_1} \\ &\geq \quad (1 - \pi_A\alpha_1) \log \frac{1 - \pi_A\alpha_1}{1 - \pi_A(1-\alpha_1)} + \pi_A\alpha_1 \log \frac{\alpha_1}{1-\alpha_1} \\ &= \quad \gamma_1(R, 0; 1) > 0. \end{aligned}$$

This establishes that 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

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

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 immediately 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, i.e. $u_i(a,\omega) = \mathbb{1}_{a=\omega}$. Assume that an equal and positive share of partisan and nonpartisan 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

[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.

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 private signal distribution is symmetric, $F^R(s) = 1 - F^L(1-s)$ and has support $\mathcal{S} = [0,1]$, and all types have common prior $p_0 = 1/2$. Note that signals are aligned since partisan types order signals in the same way as nonpartisan types, i.e. 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 private signal s updates to belief $\frac{p_i(\lambda,s)}{1-p_i(\lambda,s)} = \lambda\left(\frac{s}{1-s}\right)$. It chooses action L if this belief is less than one, which is equivalent to $s < 1/(1+\lambda) = \overline{s}_{i,1}(\lambda)$. 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 $F^\omega(.5)$. 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 nonpartisan types. It believes action L occurs with probability $\hat{\psi}_1(L|\omega,\boldsymbol{\lambda}) = (1-\pi_A)F^\omega(1/(1+\lambda_1)) + \pi_A F^\omega(.5)$. This type misspecification leads the partisan type to underestimate the range of signals for which other agents choose action L, while its signal misspecification causes it to overestimate the probability of these signals. The latter effect dominates, and θ_1 overestimates the frequency of action L. Nonpartisan type θ_3 has a correctly specified model of the signal distribution and believes that other agents do as well, since it does not account for the partisan types. It believes action L occurs with probability $\hat{\psi}_3(L|\omega,\boldsymbol{\lambda}) = (1-\pi_A)F^\omega(1/(1+\lambda_3)) + \pi_A F^\omega(.5)$. This type misspecification leads the nonpartisan type to believe that other agents are choosing L for a larger range of signals than is actually the case, which leads it to overestimate the frequency of L actions. The true probability of action L is

$$\begin{aligned}\psi(L|\omega,\boldsymbol{\lambda}) &= (1-q)((1-\pi_A)F^\omega((1/(1+\lambda_1))^{1/\nu}) + \pi_A F^\omega(.5^{1/\nu})) \\ &\quad + q((1-\pi_A)F^\omega(1/(1+\lambda_3)) + \pi_A F^\omega(.5)).\end{aligned}$$

Although the partisan and nonpartisan 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 nonpartisan 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 cut-offs $q_1 \in (0,1)$ and $q_2 \in (q_1, 1)$ such that:*

1. *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.*
2. *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.*
3. *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 $\alpha_\nu \equiv F^L(.5^{1/\nu})$ 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, $F^R(.5) = 1 - F^L(.5) = 1 - \alpha_1$, 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 these expressions' dependence on ν and q 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 the two possible stationary limit beliefs $(0,0)$ and (∞, ∞).

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

$$\begin{aligned}\gamma_1(L, (\infty,\infty); \nu, q) = \psi(L|L, (\infty,\infty); \nu, q) \log \frac{1-\alpha_1}{\alpha_1} \\ + \psi(R|L, (\infty,\infty); \nu, q) \log \frac{1-\pi_A(1-\alpha_1)}{1-\pi_A \alpha_1},\end{aligned}$$

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

[2] Partisan and nonpartisan types with the same likelihood ratio may choose different actions following a given signal s, as they have different private signal cut-offs.

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 nonpartisan 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 signals, $\alpha_0=0$. 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

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

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 $(0,0)\in\Lambda(L;\nu,q)$.

Finally we 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

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

Therefore, $(0,0) \notin \Lambda(\omega;\nu,q)$ before $(\infty,\infty) \in \Lambda(\omega;\nu,q)$. This establishes the first part of the proposition.

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. □

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:

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

Then $\tilde{\gamma}(\omega,\boldsymbol{\lambda};\nu) = \nu\tilde{\gamma}(\omega,\boldsymbol{\lambda};1)$ since $\int_{s\in\mathcal{S}} \log\left(\frac{s}{1-s}\right)^{\nu} dF^{\omega}(s) = \nu\int_{s\in\mathcal{S}} \log\left(\frac{s}{1-s}\right) dF^{\omega}(s) = \nu\tilde{\gamma}(\omega,\boldsymbol{\lambda};1)$, where $\tilde{\gamma}(\omega,\boldsymbol{\lambda};1)$ is the expected change in the log likelihood ratio in the correctly specified model. Therefore, $\tilde{\gamma}(\omega,\boldsymbol{\lambda};\nu)$ has the same sign as $\tilde{\gamma}(\omega,\boldsymbol{\lambda};1)$. 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) shows that beliefs converges a.s. in state L to the unique element in $\Lambda(L)$. Therefore, correct learning occurs almost surely, independent of ν. □

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 cut-off to choose action a_1. Note that this reduces to p^* for a correctly specified type, i.e. $\hat{p}^*(1) = p^*$.

We first construct the locally stable set. We write $\gamma_i(\omega,\boldsymbol{\lambda};x,\nu)$ 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 $\gamma_i(\omega,\boldsymbol{\lambda};x,\nu)$,

$$\begin{aligned}\Gamma_0(x,\nu) \equiv{}& (F^L(\hat{p}^*(\nu))+x)\log\frac{F^R(p^*)+x}{F^L(p^*)+x} - F^R(\hat{p}^*(\nu))\log\frac{F^R(p^*)}{F^L(p^*)}\\ &+ (F^L(1/2)-F^R(1/2)+F^R(\hat{p}^*(\nu))-F^L(\hat{p}^*(\nu)))\log\frac{F^R(1/2)-F^R(p^*)}{F^L(1/2)-F^L(p^*)}\end{aligned}$$

$$\Gamma_\infty(x,\nu) \equiv -(F^R(\hat{p}^*(\nu))+x)\log\frac{F^R(p^*)+x}{F^L(p^*)+x} + F^L(\hat{p}^*(\nu))\log\frac{F^R(p^*)}{F^L(p^*)}$$
$$+ (F^L(1/2)-F^R(1/2)+F^R(\hat{p}^*(\nu))-F^L(\hat{p}^*(\nu)))\log\frac{F^R(1/2)-F^R(p^*)}{F^L(1/2)-F^L(p^*)}.$$

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

$$\frac{\partial\Gamma_0}{\partial\nu} = \frac{d\hat{p}^*(\nu)}{d\nu} f^L(\hat{p}^*(\nu)) \times\left(\log\frac{F^R(p^*)+x}{F^L(p^*)+x} - \frac{\hat{p}^*(\nu)}{1-\hat{p}^*(\nu)}\log\frac{F^R(p^*)}{F^L(p^*)} - \left(1-\frac{\hat{p}^*(\nu)}{1-\hat{p}^*(\nu)}\right)\log\frac{F^R(1/2)-F^R(p^*)}{F^L(1/2)-F^L(p^*)}\right),$$

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 signals 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 $\Gamma_0(x,1) \le 0$, $\nu \mapsto \Gamma_0(x,\nu)$ 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

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

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, $\frac{\partial\Gamma_0}{\partial x}(x,1) \ge 0$ for all x. Finally, $\Gamma_0(x,\nu) \ge \Gamma_0(x',\nu)$ for $x > x'$. Therefore, as x increases, $\gamma_1(L,0;x,\nu)$ crosses 0 at a lower ν, i.e. 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

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

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 cut-offs 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. □

C.2 Section 4.2 (Naive Learning)

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 $\gamma_i(\omega, \boldsymbol{\lambda}; \hat{\pi}_i)$ 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

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

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 $\hat{\pi}$, i.e. $(\infty,\infty) \notin \Lambda(L;\hat{\pi},\hat{\pi})$, then it does not arise in any corresponding heterogeneous model with average bias $\hat{\pi}$, i.e. $(\infty,\infty) \notin \Lambda(L;\hat{\pi}_1,\hat{\pi}_2)$ 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 cut-off $\overline{\pi} \in (\pi, 1]$ 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

$$\begin{aligned}\gamma_i(L,(0,\infty);\hat{\pi}_i) &= (\pi\alpha_L + (1-\pi)/2)\log\left(\frac{\hat{\pi}_i\alpha_R + \frac{1}{2}(1-\hat{\pi}_i)}{\hat{\pi}_i\alpha_L + \frac{1}{2}(1-\hat{\pi}_i)}\right) \\ &\quad + (\pi(1-\alpha_L) + (1-\pi)/2)\log\left(\frac{\hat{\pi}_i(1-\alpha_R) + \frac{1}{2}(1-\hat{\pi}_i)}{\hat{\pi}_i(1-\alpha_L) + \frac{1}{2}(1-\hat{\pi}_i)}\right) \\ &= \pi(2\alpha_L - 1)\log\left(\frac{\hat{\pi}_i(1-\alpha_L) + \frac{1}{2}(1-\hat{\pi}_i)}{\hat{\pi}_i\alpha_L + \frac{1}{2}(1-\hat{\pi}_i)}\right),\end{aligned}$$

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

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. □

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.

Construction of $\Lambda(\omega)$. 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 $s \geq 1/(\lambda_i + 1) = \bar{s}_{i,1}(\lambda_i)$. Given $\lambda_1 = 1$, $\bar{s}_{1,1}(1) = 0.5$ and 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,

$$\hat{\psi}_3(L|\omega, \boldsymbol{\lambda}) = (1 - \varepsilon)F^\omega(1/(\lambda_2 + 1)) + \varepsilon F^\omega(.5), \tag{24}$$

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,

$$\psi(L|\omega, \boldsymbol{\lambda}) = \pi(\theta_1)F^\omega(.5) + \pi(\theta_2)F^\omega(1/(\lambda_2 + 1)) + \pi(\theta_3)F^\omega(1/(\lambda_3 + 1)). \tag{25}$$

To simplify the exposition, let $\alpha_L \equiv F^L(.5)$ be the probability a level-1 type plays action L in state L and $\alpha_R \equiv F^R(.5)$ 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 signals. 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. We next characterize the region of the type distribution at which each disagreement outcome is locally stable. To do so, 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.

1. 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 cut-off $c_1 > 0$ such that for $\pi(\theta_3) < c_1$, $(0, \infty) \in \Lambda_2(L)$ for all $\pi(\theta_1)$ and $\pi(\theta_2)$.

2. 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 cut-off $c_2 \in (0, 1)$ such that for $\pi(\theta_3) > c_2$, $(0, \infty) \notin \Lambda_2(L)$ and there exists a cut-off $c_3 < 1$ 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 (1). 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.

3. 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 cut-off $c_4 > 0$ such that for $\pi(\theta_2) < c_4$, $(\infty, 0) \notin \Lambda_2(L)$ for all $\pi(\theta_1)$ and $\pi(\theta_3)$.

4. 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 cut-off $c_5 < 1$ such that for $\pi(\theta_2) > c_5$, $(0, \infty) \in \Lambda_2(L)$ and there

exists an interior cut-off $c_6 \in (0,1)$ 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 part (3). 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.

5. 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 cut-off $c_7 \in (0,1)$ such that for $\pi(\theta_1) > c_7$, $(0,\infty) \in \Lambda_2(L)$ and there exists an interior cut-off $c_8 \in (0,1)$ 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 parts (2) and (4). 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 cut-off

$$c_{10} \equiv \frac{\log \frac{\alpha_L}{\alpha_R}}{\log \frac{\alpha_L}{\alpha_R} - \log \frac{1-\alpha_L}{1-\alpha_R}}$$

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

$$c_{12} \equiv \frac{\log \frac{1-\alpha_L}{1-\alpha_R}}{\log \frac{1-\alpha_L}{1-\alpha_R} - \log \frac{\alpha_L}{\alpha_R}}.$$

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.

Maximal Accessibility. When $\Lambda(\omega)$ contains a disagreement outcome, we need to check whether the disagreement outcome is maximally accessible to determine whether it occurs with positive probability from any initial belief. The following lemma establishes that both disagreement outcomes are maximally accessible at all distributions over types and all $\varepsilon \in (0,1]$. This implies that a disagreement outcome arises with positive probability if and only if it is in $\Lambda(\omega)$.

Claim 1. *For any $\pi \in \Delta((\theta_1,\theta_2,\theta_3))$ and $\varepsilon \in (0,1]$, both disagreement outcomes $(0,\infty)$ 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,

$$\frac{\hat{\psi}_2(R|R,(0,0))}{\hat{\psi}_2(R|L,(0,0))} = \frac{\hat{\psi}_3(R|R,(0,0))}{\hat{\psi}_3(R|L,(0,0))} = \frac{1-\alpha_R}{1-\alpha_L}. \tag{26}$$

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,

$$\frac{\hat{\psi}_2(L|R,(\infty,\infty))}{\hat{\psi}_2(L|L,(\infty,\infty))} = \frac{\hat{\psi}_3(L|R,(\infty,\infty))}{\hat{\psi}_3(L|L,(\infty,\infty))} = \frac{\alpha_R}{\alpha_L}. \tag{27}$$

Therefore, $\theta_2 \succ_{(\infty,\infty)} \theta_3$. From Definition 7, this implies that $(\infty,0)$ is maximally accessible.

Construction of $\Lambda_M(\omega)$. Finally, we need to rule out mixed learning outcomes in which θ_2's beliefs converge and θ_3's beliefs cycle, or vice versa. Suppose $\omega = L$ and consider the four possible mixed outcomes.

1. Consider the mixed outcome $(0,\theta_3)$ in which $\langle\lambda_{2,t}\rangle$ does not converge and $\langle\lambda_{3,t}\rangle \to 0$. 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.

2. Consider the mixed outcome (∞,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.

3. Consider the mixed outcome $(0,\theta_2)$. 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.

4. Consider the mixed outcome (∞,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.

Therefore, $\Lambda_M(L) = \emptyset$ and mixed outcomes almost surely do not arise if the state is L. Similar logic rules out mixed outcomes if the state is R.

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

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 $\boldsymbol{\lambda}_2^*, ..., \boldsymbol{\lambda}_L^* = \boldsymbol{\lambda}^*$ such that for $l = 1, ..., L-1$, 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 of Theorem 7′ Given $\kappa \in \{1, ..., k\}$, 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. Choose an $\varepsilon < e^{-E}$, where E is defined in Eq. (14). 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 $Pr(h = \xi) > 0$, $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) > 0$. 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. □

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 $\boldsymbol{\lambda}_2^*, ..., \boldsymbol{\lambda}_L^* = \boldsymbol{\lambda}^*$ such that for $l = 1, ..., L-1$, 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 is identified at certainty and 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.*

The proof of Theorem 3′ shows that Definition 7′ implies the separability condition.

Proof of Theorem 3′. 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 $l = 1, \dots k-\kappa+1$. For each $l = 1, \dots k-\kappa+1$, $\theta_{k-l+1} \succ_{\boldsymbol{\lambda}_l^*} \theta_{k-l}$ implies that the submatrix $\Psi[\theta_{k-l+1}, \theta_{k-l}; a_1, a_M](\boldsymbol{\lambda}_l^*)$ defined in Eq. (17) 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 $\Psi[\theta_j; a_1, a_M](\boldsymbol{\lambda}_l^*)$ is the submatrix of Eq. (16) for type θ_j and actions a_1 and a_M. Therefore, $\boldsymbol{\lambda}_l^*$ is separable at zero for θ_{k-l+1}, since we can set the elements of c to zero for the remaining actions in $\Psi(\boldsymbol{\lambda}_l^*)$. The case where Part (ii) of Definition 7′ holds is analogous, as is the proof for other disagreement outcomes.
□

D.2 Mixed Learning

Consider the mixed outcome $(\boldsymbol{\lambda}_I^*, I)$ in which beliefs converge to $\boldsymbol{\lambda}_I^* \in \{0, \infty\}^{|I|}$ for some subset of social types $I \subset \Theta_S$ and beliefs do not converge for the remaining social types $N \equiv \Theta_S \setminus I$, where $\boldsymbol{\lambda}_I$ denotes the likelihood ratio vector $\boldsymbol{\lambda}$ restricted to a set of types I. 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. It is relatively straightforward to derive conditions under which this outcome is not locally stable. In particular, if $(0,0,0) \in \Lambda_3(\omega)$ or $(0,0,\infty) \in \Lambda_3(\omega)$, then the beliefs of θ_3 to converge with positive probability when $\langle \lambda_{1,t}, \lambda_{2,t} \rangle \to (0,0)$. From here, we can show that this ensures $((0,0), \{\theta_1, \theta_2\})$ almost surely does not occur. This insight generalizes to other mixed outcomes where the beliefs of all but one social type converge. The argument is more involved when ruling out mixed outcomes in which the beliefs of two or more social types do not converge, as we also need to show that a locally stable outcome for the non-convergent types is accessible from other points in the mixed outcome belief space. 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, it is not enough to show that $(0,0,0) \in \Lambda_2(L) \cap \Lambda_3(L)$. We also need to show that for $\boldsymbol{\lambda} \in \{(0,\infty,0),(0,0,\infty),(0,\infty,\infty)\}$, either (i) beliefs will almost surely enter a neighborhood of $(0,0,0)$ from a neighborhood of $\boldsymbol{\lambda}$, or (ii) $\boldsymbol{\lambda} \in \Lambda_2(L) \cap \Lambda_3(L)$. We formalize this idea in the following paragraphs.

We first define the concept of mixed accessibility. Given a mixed outcome, this concept identifies pairs of stationary points for the non-convergent types where, for each non-convergent type whose beliefs differ between the two points, on average, the type's belief drifts away from the first point and towards the second point.

Definition 12 (Mixed Accessible ($k \geq 2$)). *Given mixed outcome $(\boldsymbol{\lambda}^*_I, I)$ with $\boldsymbol{\lambda}^*_I \in \{0,\infty\}^{|I|}$, $I \subset \Theta_S$ and $N \equiv \Theta_S \setminus I$, and distinct stationary beliefs $\boldsymbol{\lambda}_N, \boldsymbol{\lambda}'_N \in \{0,\infty\}^{|N|}$, belief $(\boldsymbol{\lambda}'_N, N)$ is* mixed accessible *from $(\boldsymbol{\lambda}_N, N)$ if $(\boldsymbol{\lambda}^*_I, \boldsymbol{\lambda}_N) \notin \Lambda_i(\omega)$ for each $i \in N$ such that $\lambda'_i \neq \lambda_i$.*

When considering mixed accessibility, we focus on pairs of stationary points in which all non-convergent types whose beliefs differ between the two points have the same beliefs at the initial point.

Definition 13 (Agreement Adjacent). *Given a set of types $N \subset \Theta_S$ and stationary beliefs $\boldsymbol{\lambda}_N, \boldsymbol{\lambda}'_N \in \{0,\infty\}^{|N|}$, beliefs $\boldsymbol{\lambda}_N$ and $\boldsymbol{\lambda}'_N$ are* agreement adjacent *if $\lambda_i = \lambda_j$ for each distinct $\theta_i, \theta_j \in N$ such that $\lambda_i \neq \lambda'_i$ and $\lambda_j \neq \lambda'_j$.*

Using these two definitions and fixing mixed outcome $(\boldsymbol{\lambda}^*_I, I)$ with non-convergent types $N \equiv \Theta_S \setminus I$, we construct a graph $\mathcal{G}(\boldsymbol{\lambda}^*_I, I)$ with nodes $(\boldsymbol{\lambda}_N, N)$ for $\boldsymbol{\lambda}_N \in \{0,\infty\}^{|N|}$ to represent which stationary points for types in N are mixed accessible from agreement adjacent stationary points for types in N.

Definition 14 (Accessible Graph ($k \geq 2$)). *Given $(\boldsymbol{\lambda}^*_I, I)$ and $N \equiv \Theta_S \setminus I$, define the* directed graph *$\mathcal{G}(\boldsymbol{\lambda}^*_I, I)$ with nodes $(\boldsymbol{\lambda}_N, N)$ for $\boldsymbol{\lambda}_N \in \{0,\infty\}^{|N|}$ as follows: there is an edge from $(\boldsymbol{\lambda}_N, N)$ to $(\boldsymbol{\lambda}'_N, N)$ if and only if $(\boldsymbol{\lambda}'_N, N)$ is mixed accessible from $(\boldsymbol{\lambda}_N, N)$ and $(\boldsymbol{\lambda}'_N, N)$ and $(\boldsymbol{\lambda}_N, N)$ are agreement adjacent. A* terminal node *has no edges leaving it.*

We say $(\boldsymbol{\lambda}^*_I, I)$ is *reducible* if $\mathcal{G}(\boldsymbol{\lambda}^*_I, I)$ has no cycles.

We define $\Lambda_M(\omega)$ in the case of $k > 2$ as the set of mixed learning outcomes that are not reducible,

$$\Lambda_M(\omega) \equiv \{(\boldsymbol{\lambda}^*_I, I) | \boldsymbol{\lambda}^*_I \in \{0,\infty\}^{|I|},\ I \subset \Theta_S,\ (\boldsymbol{\lambda}^*_I, I) \text{ is not reducible}\}. \tag{28}$$

This is equivalent to Eq. (8) when $k = 2$. As in the case of $k = 2$, if a mixed learning outcome arises with positive probability, then it must be in $\Lambda_M(\omega)$. Using Eq. (28), the statement of Lemma 4′ is identical to Lemma 4.[3]

[3] An alternative condition that involves bounding $\gamma_i(\omega, \boldsymbol{\lambda})$ across the belief space for $i \in \Theta_S \setminus I$ can also be used to rule out mixed learning.

Lemma 4′ (Unstable Mixed Outcomes ($k \geq 2$))**.** *Assume Assumptions 1 to 4 and consider a learning environment that is identified at certainty. Given $I \subset \Theta_S$ and $N \equiv \Theta_S \setminus I$, if $(\boldsymbol{\lambda}_I^*, I) \notin \Lambda_M(\omega)$, then $Pr(\boldsymbol{\lambda}_{I,t} \to \boldsymbol{\lambda}_I^*$ and $\boldsymbol{\lambda}_{N,t}$ does not converge$) = 0$.*

This lemma implies that if $(\boldsymbol{\lambda}_I^*, I)$ is reducible, then almost surely it does not arise. Intuitively, if $(\boldsymbol{\lambda}_I^*, I)$ is reducible, then when the beliefs of the convergent types I remain in a neighborhood of $\boldsymbol{\lambda}_I^*$, almost surely the beliefs of the non-convergent types also converge. Therefore, almost surely, $(\boldsymbol{\lambda}_I^*, I)$ does not arise.

Reducibility is always satisfied in some important cases and is relatively straightforward to verify. For instance, it is satisfied in misspecified environments that are close to the correctly specified environment, in which $\gamma_i(L, \boldsymbol{\lambda}) < 0$ at $\boldsymbol{\lambda} \in \{0, \infty\}^k$ for all social types θ_i. In such an environment, each node in the graph where κ non-convergent types have belief $\lambda_i = \infty$ is connected to all nodes in which $\kappa' < \kappa$ non-convergent types have belief $\lambda_i = \infty$ and is not connected to any other nodes. Therefore, for any mixed outcome $(\boldsymbol{\lambda}_I^*, I)$ with non-convergent types $N \equiv \Theta_S \setminus I$, every path in the graph $\mathcal{G}(\boldsymbol{\lambda}_I^*, I)$ terminates at node $(0^{|N|}, N)$. For mixed outcome $(0^{|I|}, I)$, this is a convergent point. For other mixed outcomes, this is a point at which each social type's belief moves towards zero in expectation, and therefore, some convergent type $\theta_i \in I$'s belief must eventually exit a neighborhood of $\boldsymbol{\lambda}_I^*$.

Before proving Lemma 4′, we establish the following intermediate result.

Lemma 8. *Given mixed outcome $(\boldsymbol{\lambda}_I^*, I)$ with non-convergent types $N \equiv \Theta_S \setminus I$ and graph $\mathcal{G}(\boldsymbol{\lambda}_I^*, I)$, if $(\boldsymbol{\lambda}_N, N)$ is a terminal node, then $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{\theta_i \in N} \Lambda_i(\omega)$.*

Proof. Let $(\boldsymbol{\lambda}_N, N)$ be a terminal node in $\mathcal{G}(\boldsymbol{\lambda}_I^*, I)$. By definition of terminal node, no nodes that are agreement adjacent to $(\boldsymbol{\lambda}_N, N)$ are mixed accessible from $(\boldsymbol{\lambda}_N, N)$. If any node is mixed accessible from $(\boldsymbol{\lambda}_N, N)$, then there exists a $\theta_i \in N$ such that $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \notin \Lambda_i(\omega)$. Then the node $(\boldsymbol{\lambda}'_N, N)$ where $\lambda'_j = \lambda_j$ for all $\theta_j \neq \theta_i$ is agreement adjacent and mixed accessible, so $(\boldsymbol{\lambda}_N, N)$ is not a terminal node. This is a contradiction. Therefore, if $(\boldsymbol{\lambda}_N, N)$ is a terminal node, then no nodes $(\boldsymbol{\lambda}'_N, N)$ are mixed accessible from $(\boldsymbol{\lambda}_N, N)$. Therefore, by definition of mixed accessibility, $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{\theta_i \in N} \Lambda_i(\omega)$. □

Proof of Lemma 4′. Suppose mixed outcome $(\boldsymbol{\lambda}_I^*, I)$ is reducible, i.e. $(\boldsymbol{\lambda}_I^*, I) \notin \Lambda_M(\omega)$. We will show that this implies that $(\boldsymbol{\lambda}_I^*, I)$ almost surely does not occur. Fix $\varepsilon < e^{-E}$, where E is defined in Eq. (14), and suppose $\boldsymbol{\lambda}_{I,1} \in B_\varepsilon(\boldsymbol{\lambda}_I^*)$. We will show that almost surely, either (i) there exists a time $\tau < \infty$ such that $\boldsymbol{\lambda}_{I,\tau} \notin B_\varepsilon(\boldsymbol{\lambda}_I^*)$ or (ii) $\langle \boldsymbol{\lambda}_t \rangle$ converges for all social types.

By reducibility, at every $\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}$, either $(\boldsymbol{\lambda}_N, N)$ is a terminal node and therefore, by Lemma 8, $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{\theta_i \in N} \Lambda_i(\omega)$, or there exists an agreement adjacent belief $\boldsymbol{\lambda}'_N \in \{0, \infty\}^{|N|}$ that is mixed accessible from $(\boldsymbol{\lambda}_N, N)$ and is a terminal node, so by Lemma 8, $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}'_N) \in \cap_{\theta_i \in N} \Lambda_i(\omega)$. First consider $\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}$ such that $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{\theta_i \in N} \Lambda_i(\omega)$. By the construction in Theorem 1, if beliefs enter $B_\varepsilon((\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N))$, then $\langle \boldsymbol{\lambda}_{N,t} \rangle$ is bounded above by a process that converges to $\boldsymbol{\lambda}_N$ with positive probability, and this probability is uniformly bounded away from zero for any belief in $B_\varepsilon((\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N))$. If $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{\theta_i \in I} \Lambda_i(\omega)$, then $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N)$ is locally stable, so with positive probably, $\langle \boldsymbol{\lambda}_t \rangle \to (\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N)$. Otherwise, if $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \notin \cap_{\theta_i \in I} \Lambda_i(\omega)$, then for some $\theta_i \in I$, $\langle \lambda_{i,t} \rangle$ is bounded below by a process that almost surely leaves $B_\varepsilon(\boldsymbol{\lambda}_I^*)$. Therefore, in the event that $\langle \boldsymbol{\lambda}_{N,t} \rangle \to \boldsymbol{\lambda}_N$, $\langle \boldsymbol{\lambda}_{I,t} \rangle$ almost surely leaves $B_\varepsilon(\boldsymbol{\lambda}_I^*)$.

Next consider $\boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|}$ such that $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \notin \cap_{\theta_i \in N} \Lambda_i(\omega)$. Fix $0 < \varepsilon' < e^{-E}$. We want to show that there exists a $\varepsilon_2 > 0$ such that if initial belief $\boldsymbol{\lambda}_{N,1} \in B_{\varepsilon_2}(\boldsymbol{\lambda}_N)$, then there exists an agreement adjacent belief $(\boldsymbol{\lambda}'_N, N)$ that is mixed accessible from $(\boldsymbol{\lambda}_N, N)$ such that with probability uniformly bounded away from zero in initial belief $\boldsymbol{\lambda}_{N,1}$, beliefs enter a neighborhood $B_{\varepsilon'}(\boldsymbol{\lambda}'_N)$. Given $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N)$, let λ_i denote the component for type $\theta_i \in N$ and λ_i^* denote the component for type $\theta_i \in I$. By the construction in Theorem 1, there exists a $\theta_i \in N$ such that $\langle \lambda_{i,t} \rangle$ is bounded below by a process that almost surely leaves $B_\varepsilon(\lambda_i)$. Let N_U be the set of types $\theta_i \in N$ such that $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \notin \Lambda_i(\omega)$, with $N_{U,0}$ the set of $\theta_i \in N_U$ such that $\lambda_i = 0$ and $N_{U,\infty}$ the set of $\theta_i \in N_U$ such that $\lambda_i = \infty$. We now argue that starting from a neighborhood $B_{\varepsilon_2}(\boldsymbol{\lambda}_N)$ for $\theta_i \in N$ and $B_\varepsilon(\boldsymbol{\lambda}_I^*)$ for $\theta_i \in I$, with positive probability, either $\langle \boldsymbol{\lambda}_{I,t} \rangle$ leaves $B_\varepsilon(\boldsymbol{\lambda}_I^*)$ or $\langle \boldsymbol{\lambda}_{N,t} \rangle$ reaches $B_{\varepsilon'}(\boldsymbol{\lambda}'_N)$ for some agreement adjacent mixed accessible belief $(\boldsymbol{\lambda}'_N, N)$. For $\theta_i \in N_{U,0}$, let N_i be the minimum number of a_M actions it takes for any $\lambda_{i,t} \in [\varepsilon', 1/\varepsilon']$ to hit $1/\varepsilon'$. Similarly, for $\theta_i \in N_{U,\infty}$, let N_i be the minimum number of a_1 actions it takes for any $\lambda_{i,t} \in [\varepsilon', 1/\varepsilon']$ to hit ε'. By the construction in Theorem 1, there exists an $\varepsilon_2 > 0$ such that if $\boldsymbol{\lambda}_{N,1} \in B_{\varepsilon_2/2}(\boldsymbol{\lambda}_N)$, with positive probability there exists a finite t such that $\boldsymbol{\lambda}_{N \setminus N_U, t} \in B_{\varepsilon_2}(\boldsymbol{\lambda}_{N \setminus N_U})$, and $\boldsymbol{\lambda}_{N_U,t} \notin B_{\varepsilon'}(\boldsymbol{\lambda}_{N_U})$.

Choose $\varepsilon_2 > 0$ such that if $\boldsymbol{\lambda}_{N \setminus N_U, 1} \in B_{\varepsilon_2}(\boldsymbol{\lambda}_{N \setminus N_U})$, then after $\sum_{\theta_i \in N_{U,0}} N_i$ actions a_M, $\lambda_{i,t} \in B_{\varepsilon'}(\lambda_i)$ for all $\theta_i \in N \setminus N_{U,0}$, and after $\sum_{i \in N_{U,\infty}} N_i$ actions a_1, $\lambda_{i,t} \in B_{\varepsilon'}(\lambda_i)$ for all $\theta_i \in N \setminus N_{U,\infty}$. Therefore, if $\boldsymbol{\lambda}_{N,1} \in B_{\varepsilon_2/2}(\boldsymbol{\lambda}_N)$ and $\boldsymbol{\lambda}_{I,1} \in B_\varepsilon(\boldsymbol{\lambda}_I^*)$, then with positive probability either (i) there exists a $t < \infty$ such that $\boldsymbol{\lambda}_{I,t} \notin B_\varepsilon(\boldsymbol{\lambda}_I^*)$, or (ii) there exists $t < \infty$ such that for some $\theta_i \in N_U$, $\lambda_{i,t} \notin B_{\varepsilon'}(\lambda_i)$ and for all $\theta_i \in N \setminus N_U$, $\lambda_{i,t} \in B_{\varepsilon_2}(\lambda_i)$. First consider case (ii) and suppose that a type $\theta_i \in N_{U,0}$ leaves. After N_i actions a_M, if $\boldsymbol{\lambda}_{N,t} \in B_{\varepsilon'}(\boldsymbol{\lambda}'_N)$ for some agreement adjacent belief $(\boldsymbol{\lambda}'_N, N)$ that is mixed accessible from $(\boldsymbol{\lambda}_N, N)$, then stop. Otherwise, there exists an $\theta_{i_2} \in N_{U,0}$ such that $\lambda_{i_2,t} \notin B_{\varepsilon'}(\lambda_{i_2})$. Repeat N_{i_2} realizations of a_M. After these $N_i + N_{i_2}$ realizations of a_M, if $\boldsymbol{\lambda}_{N,t} \in B_{\varepsilon'}(\boldsymbol{\lambda}'_N)$ for some agreement adjacent belief $(\boldsymbol{\lambda}'_N, N)$ that is mixed accessible from $(\boldsymbol{\lambda}_N, N)$, then stop. Otherwise, there is an $\theta_{i_3} \in N_{\boldsymbol{\lambda},0}$ such that $\lambda_{i_3,t} \notin B_{\varepsilon'}(\lambda_{i_3})$. Repeat N_{i_3} realizations of a_M, and so on. Therefore, after at most $\sum_{\theta_i \in N_{U,0}} N_i$ realizations of a_M, beliefs have entered the ε' ball around some other stationary point $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}'_N)$ such that $(\boldsymbol{\lambda}'_N, N)$ is agreement adjacent to and mixed accessible from $(\boldsymbol{\lambda}_N, N)$. Therefore, the probability of either $\langle \boldsymbol{\lambda}_{N,t} \rangle$ reaching a neighborhood $B_{\varepsilon'}(\boldsymbol{\lambda}'_N)$ of some agreement adjacent belief $(\boldsymbol{\lambda}'_N, N)$ that is mixed accessible from $(\boldsymbol{\lambda}_N, N)$ or $\langle \boldsymbol{\lambda}_{I,t} \rangle$ leaving the neighborhood $B_\varepsilon(\boldsymbol{\lambda}_I^*)$ is bounded below by the probability of $\sum_{\theta_i \in N_{U,0}} N_i$ realizations of a_M, which is strictly positive. The argument for a type $\theta_i \in N_{U,\infty}$ is analogous.

Consider the graph $\mathcal{G}(\boldsymbol{\lambda}_I^*, I)$. We will choose an $\varepsilon(\boldsymbol{\lambda}_N)$ to correspond to each node $(\boldsymbol{\lambda}_N, N)$. At any terminal node $(\boldsymbol{\lambda}_N, N)$, define $\varepsilon(\boldsymbol{\lambda}_N) = \varepsilon$. For any node $(\boldsymbol{\lambda}'_N, N)$ that only has edges going to terminal nodes, by the above construction, there exists an $\varepsilon(\boldsymbol{\lambda}'_N)$ such that if $\boldsymbol{\lambda}_{N,t} \in B_{\varepsilon(\boldsymbol{\lambda}'_N)}(\boldsymbol{\lambda}'_N)$, then with positive probability, either $\langle \boldsymbol{\lambda}_{N,t} \rangle$ reaches $B_\varepsilon(\boldsymbol{\lambda}_N)$ for terminal node $(\boldsymbol{\lambda}_N, N)$ or $\langle \boldsymbol{\lambda}_{I,t} \rangle$ exits $B_\varepsilon(\boldsymbol{\lambda}_I^*)$. Repeat this process for each node in the graph.

Let $\tau_1 = \min\{t | \boldsymbol{\lambda}_{I,t} \notin B_\varepsilon(\boldsymbol{\lambda}_I^*)\}$. Then almost surely, $\tau_1 < \infty$ or $\langle \boldsymbol{\lambda}_{N,t} \rangle$ enters the neighborhood of a node on the graph constructed above infinitely often,

$$Pr(\tau_1 < \infty \text{ or for some } \boldsymbol{\lambda}_N \in \{0, \infty\}^{|N|},\ \boldsymbol{\lambda}_{N,t} \in B_{\varepsilon(\boldsymbol{\lambda}_N)}(\boldsymbol{\lambda}_N) \text{ i.o.}) = 1.$$

If $\langle \boldsymbol{\lambda}_{N,t} \rangle$ enters the neighborhood of a terminal node $(\boldsymbol{\lambda}_N, N)$ infinitely often, then by Lemma 8, $\boldsymbol{\lambda}_N \in \cap_{\theta_i \in N} \Lambda_i(\omega)$, so either $\boldsymbol{\lambda}_{N,t} \to \boldsymbol{\lambda}_N$ or $\langle \boldsymbol{\lambda}_{I,t} \rangle$ leaves $B_\varepsilon(\boldsymbol{\lambda}_I^*)$. Otherwise. $\langle \boldsymbol{\lambda}_{N,t} \rangle$ enters the neighborhood of some mixed accessible agreement adjacent belief $(\boldsymbol{\lambda}'_N, N)$ infinitely often. Since any path of this form ends at a terminal node, this implies that almost surely, either $\langle \boldsymbol{\lambda}_{N,t} \rangle$ converges or $\langle \boldsymbol{\lambda}_{I,t} \rangle$ leaves $B_\varepsilon(\boldsymbol{\lambda}_I^*)$. Therefore, the mixed outcome $(\boldsymbol{\lambda}_I^*, I)$ almost surely does not arise. □

D.3 Learning Characterization

Given the definitions of $\Lambda_M(\omega)$ and maximal accessibility for $k > 2$, the statement of Lemma 7′ is identical to Lemma 7.

Lemma 7′ (Belief Convergence ($k \geq 2$))**.** *Assume Assumptions 1 to 4 and consider a learning environment that is identified at certainty. If $\Lambda(\omega)$ contains an agreement outcome or maximally accessible disagreement outcome and $\Lambda_M(\omega) = \emptyset$, then for any initial belief $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, there exists a random variable $\boldsymbol{\lambda}_\infty$ with $\operatorname{supp}(\boldsymbol{\lambda}_\infty) = \Lambda(\omega)$ such that $\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}_\infty$ almost surely.*

Proof. Suppose $\Lambda(\omega)$ contains an agreement vector or maximally accessible disagreement vector and $\Lambda_M(\omega) = \emptyset$. Recall that $\mathcal{B}$ is the set of locally stable neighborhoods and $\mathcal{B}_U$ is the set of locally unstable neighborhoods defined in Eq. (15). Let $\tau_1 \equiv \min\{t | \boldsymbol{\lambda}_t \in \mathcal{B}\}$ be the first time that the likelihood ratio enters the set of locally stable neighborhoods. By Lemma 2 in Appendix A.5, there exists a finite sequence of actions such that starting from any initial belief $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, $\langle \boldsymbol{\lambda}_t \rangle$ enters $\mathcal{B}$. This sequence occurs with positive probability. Therefore, the probability of entering $\mathcal{B}$ in finite time is bounded away from zero, $Pr(\tau_1 < \infty) > 0$. If $\langle \boldsymbol{\lambda}_t \rangle$ enters $\mathcal{B}_U$, then by Theorem 1, $\langle \boldsymbol{\lambda}_t \rangle$ almost surely leaves $\mathcal{B}_U$.

If $\langle \boldsymbol{\lambda}_t \rangle$ enters the neighborhood of a mixed outcome $\boldsymbol{\lambda}_I$, by Lemma 4′, $\langle \boldsymbol{\lambda}_t \rangle$ almost surely leaves this neighborhood or converges to a locally stable point. By Lemma 3, $\langle \boldsymbol{\lambda}_t \rangle$ does not converge to a non-stationary belief. Therefore, almost surely, either $\langle \boldsymbol{\lambda}_t \rangle$ does not converge for all types or $\langle \boldsymbol{\lambda}_t \rangle$ converges to a learning outcome in $\Lambda(\omega)$.

Suppose with positive probability $\langle \boldsymbol{\lambda}_t \rangle$ exits and never re-enters the interior of the belief space, $[e^{-E}, e^{E}]^k$, where E is defined in Eq. (14). Then either $\langle \boldsymbol{\lambda}_t \rangle$ enters the neighborhood of each mixed outcome where $|I| = 1$ infinitely often, in which case with probability one they visit a locally stable set, or there exists some i such that λ_i is constant across all neighborhoods that $\langle \boldsymbol{\lambda}_t \rangle$ enters. But then $\langle \boldsymbol{\lambda}_t \rangle$ is in the neighborhood of the mixed outcome λ_i, and by Lemma 4′, almost surely, $\langle \boldsymbol{\lambda}_t \rangle$ must leave this neighborhood or converge to a locally stable point. So almost surely, beliefs either return to $[e^{-E}, e^{E}]^k$ or converge to a locally stable point.

Let $\tau_2 \equiv \min\{\tau | \boldsymbol{\lambda}_t \in \mathcal{B} \; \forall t > \tau\}$ be the time at which $\langle \boldsymbol{\lambda}_t \rangle$ enters $\mathcal{B}$ and never leaves. From Theorem 1, $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}_\infty | \tau_2 < \infty) = 1$, where $\boldsymbol{\lambda}_\infty$ is a random variable with $\text{supp}(\boldsymbol{\lambda}_\infty) \subset \Lambda(\omega)$. Suppose $\tau_2 = \infty$. Then it must be that $\langle \boldsymbol{\lambda}_t \rangle$ enters $\mathcal{B}$ infinitely often, $Pr(\boldsymbol{\lambda}_t \in \mathcal{B} \text{ i.o.}) = 1$. But if $\langle \boldsymbol{\lambda}_t \rangle$ enters a neighborhood of a locally stable belief infinitely often, then almost surely, $\langle \boldsymbol{\lambda}_t \rangle$ converges. This is a contradiction, as we supposed $\tau_2 = \infty$. Therefore, $Pr(\tau_2 < \infty) = 1$. This implies $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}_\infty) = 1$, where $\boldsymbol{\lambda}_\infty$ is a random variable with $\text{supp}(\boldsymbol{\lambda}_\infty) \subset \Lambda(\omega)$. □

The statement of the learning characterization for $k > 2$ is identical to Theorem 4, using the expanded definitions of maximal accessibility (Definition 7′) and $\Lambda_M(\omega)$ (Eq. (28)). 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′.

E Belief-Dependent Models of Inference

In this section, we extend our framework to allow an agent's model of inference to vary with the belief about the state. We then show that with this extension, our framework nests Rabin and Schrag (1999) and Epstein et al. (2010).

E.1 Framework

We 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 of type θ_i uses likelihood ratio $\boldsymbol{\lambda}_t$ to interpret the realized action $\tilde{a}_t$ at time t.

As in Section 2, we focus on settings in which social types believe that actions are informative. When the model of inference depends on the current belief, we need to modify Assumption 3 to ensure that it holds uniformly across all values of the likelihood ratio. We define a notion of uniform informativeness to describe families of signal distributions that are bounded away from uninformative. This definition requires each possible signal realization to be either perceived as informative at all values of the likelihood ratio or perceived as uninformative at all values of the likelihood ratio. Further, it requires at least some signals to be informative.

Definition 15 (Uniformly Informative Signals)**.** *The family of subjective signal distributions* $\{\hat{F}_i^L(s;\boldsymbol{\lambda}), \hat{F}_i^R(s;\boldsymbol{\lambda})\}_{\lambda\in[0,\infty]}$ *are* uniformly informative *if for all* $s \in \mathcal{S}$*, 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* $\inf_{\boldsymbol{\lambda}\in(0,\infty]^k} \hat{F}_i^L(s;\boldsymbol{\lambda}) - \hat{F}_i^R(s;\boldsymbol{\lambda}) > 0$.

We use this definition to modify Assumption 3 so that it rules out sequences of beliefs along which an autarkic type's actions are perceived to become arbitrarily uninformative.

Assumption 3′ (Informative Actions)**.** *There exists an autarkic type* $\theta_j \in \Theta_A$ *with* $\pi(\theta_j) > 0$ *such that the signal cut-offs that describe* θ_j*'s optimal action choices are uniformly bounded away from the highest and lowest signal, i.e.* $\inf \bar{s}_{j,1}\left(\frac{p_0}{1-p_0};\boldsymbol{\lambda}\right) > \inf \mathcal{S}$ *and* $\sup \bar{s}_{j,m}\left(\frac{p_0}{1-p_0};\boldsymbol{\lambda}\right) < \sup \mathcal{S}$*, and each social type* $\theta_i \in \Theta_S$ *believes this autarkic type uniformly exists,* $\inf_{\boldsymbol{\lambda}\in[0,\infty]^k} \hat{\pi}_i(\theta_j;\boldsymbol{\lambda}) > 0$ *and the social type has a uniformly informative subjective signal distributions.*

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

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* $\boldsymbol{\lambda} \in (0,\infty)^k$*.*

Substituting Assumption 3′ for Assumption 3 and adding Assumption 5, the proofs of Theorems 1 to 6 are unchanged, using the following modified version of Eq. (2) for the probability of each action:

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

The proof of Lemma 2 continues to hold with minor modifications.[4]

[4]Other than minor change to notion and a straightforward application of the continuity assumed in Assumption 5, the only new step required is that the bound constructed in equation (9) is no longer independent of $\boldsymbol{\lambda}$. By Assumption 3′, $\bar{s}^*_{1,A} \equiv \inf_{\theta_j\in\Theta_A,\boldsymbol{\lambda}\in[0,\infty]^k} \bar{s}_{1,j}\left(\frac{p_0}{1-p_0};\boldsymbol{\lambda}\right) > \inf \mathcal{S}$, so

$$\begin{aligned}\frac{\hat{\psi}_i(a_1|R,\boldsymbol{\lambda})}{\hat{\psi}_i(a_1|L,\boldsymbol{\lambda})} &\leq \frac{\sum_{\theta_j\in\Theta_A} \hat{\pi}_i(\theta_j;\boldsymbol{\lambda})\hat{F}_i^R(\bar{s}^*_{A,1};\boldsymbol{\lambda}) + \hat{\pi}_i(\Theta_S;\boldsymbol{\lambda})}{\sum_{\theta_j\in\Theta_A} \hat{\pi}_i(\theta_j;\boldsymbol{\lambda})\hat{F}_i^L(\bar{s}^*_{A,1};\boldsymbol{\lambda}) + \hat{\pi}_i(\Theta_S;\boldsymbol{\lambda})} \\ &\leq \sup_{\boldsymbol{\lambda}\in[0,\infty]^k} \frac{(\inf_{\boldsymbol{\lambda}'\in[0,\infty]^k} \hat{\pi}_i(\Theta_A;\boldsymbol{\lambda}'))F_i^R(\bar{s}^*_{A,1};\boldsymbol{\lambda}) + (1-(\inf_{\boldsymbol{\lambda}'\in[0,\infty]^k} \hat{\pi}_i(\Theta_A;\boldsymbol{\lambda}'))}{(\inf_{\boldsymbol{\lambda}'\in[0,\infty]^k} \hat{\pi}_i(\Theta_A;\boldsymbol{\lambda}'))F_i^L(\bar{s}^*_{A,1};\boldsymbol{\lambda}) + (1-(\inf_{\boldsymbol{\lambda}'\in[0,\infty]^k} \hat{\pi}_i(\Theta_A;\boldsymbol{\lambda}'))} < 1.\end{aligned}$$

E.2 Nesting Under- and Overreaction in Epstein et al. (2010)

Epstein et al. (2010) consider an individual learning model where agents under- or overreact to new information. Consider an individual learning problem where agents observe signals drawn from true signal distributions F_R, F_L with finite support $\mathcal{S} \subseteq [0,1]$, as in the paper normalized so that $s = Pr(\omega = R|s)$. Epstein et al. (2010) introduces the following updating procedure. Consider an agent with prior p who observes signal $s \in \mathcal{S}$. This agent updates their beliefs to

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

for some $\alpha \leq 1$. Underreaction to new information corresponds to $\alpha > 0$, overreaction to new information corresponds to $\alpha < 0$, and the correctly specified model corresponds to $\alpha = 0$.

This parametric form of under- and overreaction can be represented in our framework when we allow the misspecification to vary with λ. At any interior belief, a misspecified agent updates as-if they were an agent from Epstein et al. (2010) if and only if they interpret signal $s \in \mathcal{S}$ as

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

The updating rule in Epstein et al. (2010) does not imply a unique interpretation for signals at $\lambda = 0$ or $\lambda = \infty$, since the prior and the posterior are the same regardless of the signal. Since our characterization utilizes the limit of the update rule as $\lambda \to 0$ or ∞ to characterize asymptotic outcomes, we do need to specify how signals are interpreted at extreme beliefs. At $\lambda = 0$, we use the rule from the above formula and at ∞ where the above formula is not well-defined we use

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

This is the unique perceived posterior that satisfies the continuity property required by theorem Lemma 2.[5] Theorem 1 in Bohren and Hauser (2020) shows that there exists a misspecified distribution that induces this bias and satisfies Assumption 1, Assumption 3′ and Assumption 5.

where the last inequality follows from the uniform informativeness of the subjective signal distributions.

[5]In any individual learning problem, any misspecification that induces the same update rule must satisfy

$$\frac{\hat{\psi}(s|R,\lambda)}{\hat{\psi}(s|L,\lambda)} = \frac{\hat{s}(s,\lambda)}{1-\hat{s}(s,\lambda)},$$

so $\hat{s}(s,p)$ pins 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.

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

Rabin and Schrag (1999) present a model of individual learning with confirmation bias. Agents receive a signal that takes one of two possible values s_L, s_R, where s_ω is more likely in state ω than in the other state. But the agent exhibits confirmation bias, if they receive a signal that goes against the state that their prior belief favors then with probability q they misinterpret that signal as the other signal, further pushing beliefs towards the state favored by their prior. In order to nest this model, a slight extension must be made to the framework we've outlined. In particular, we allow for multiple signals that induce the same posterior belief, which in turn allows the mapping $\hat{s}$ to map two signals that induce the same posterior to different misspecified beliefs. As long as the properties of Lemma 2 still are satisfied, this extension does not change the characterization.

Consider an individual learning setting with there are four signals s_{L_1}, s_{L_2}, s_{R_1}, s_{R_2}. All L signals induce the same posterior and all R signals induce the same posterior. Conditional on seeing an L signal, s_{L_2} is drawn with probability q. Similarly, for s_{R_1} and $Pr(s_{L_1} \text{ or } s_{L_2}|\omega = L) = Pr(s_{R_1} \text{ or } s_{R_2}|\omega = R) = s > 1/2$. When $\lambda < 1$ the agent interprets signals as if $\hat{\psi}(s_{L_1})|L, \lambda) = s$, $\hat{\psi}(s_{L_1}|R, \lambda) = 1 - s$ and if s_{L_1} is not drawn one of the other three signals is drawn uniformly at random. Similarly if $\lambda \geq 1$, the agent interprets signals as if $\hat{\psi}(s_{R_2})|R, \lambda) = s$, $\hat{\psi}(s_{R_2}|L, \lambda) = 1 - s$ and if s_{R_2} is not drawn one of the other three signals is drawn uniformly at random. The conclusions of Lemma 2 are immediate, so we can apply the results in the paper to this setting.