PENN INSTITUTE for ECONOMIC RESEARCH
UNIVERSITY of PENNSYLVANIA

The Ronald O. Perelman Center for
Political Science and Economics (PCPSE)
133 South 36th Street
Philadelphia, PA 19104-6297

pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 18-017

Social Learning with Model Misspeciification: A Framework and a Characterization

J. AISLINN BOHREN
University of Pennsylvania
& Carnegie Mellon University

DANIEL N. HAUSER
Aalto University

July 1, 2018
Revised November 7, 2019

https://ssrn.com/abstract=3236842

Social Learning with Model Misspecification: A Framework and a Characterization*

J. Aislinn Bohren[†] Daniel N. Hauser[‡]

November 7, 2019

This paper develops a general framework to study how misinterpreting information impacts learning. We consider sequential social learning and passive individual learning settings in which individuals observe signals and the actions of predecessors. Individuals have incorrect, or *misspecified* models of how to interpret these sources – such as overreaction to signals or misperception of others' preferences. Our main result is a simple criterion to characterize long-run beliefs and behavior based on the underlying form of misspecification. This provides a unified way to compare different forms of misspecification that have been previously studied, as well as generates new insights about forms of misspecification that have not been theoretically explored. It allows for a deeper understanding of how misspecification impacts learning, including exploring whether a given form of misspecification is *conceptually robust*, in that it is not sensitive to parametric specification, whether misspecification has a similar impact in individual and social learning settings, and how model heterogeneity impacts learning. Lastly, it establishes that the correctly specified model is *analytically robust*, in that nearby misspecified models generate similar long-run beliefs.

KEYWORDS: Social learning, model misspecification

*We thank Nageeb Ali, Alex Imas, Shuya Li, George Mailath, Margaret Meyer, Ali Polat, Andrew Postlewaite, Andrea Prat, Yuval Salant, Larry Samuelson, Joel Sobel, Ran Spiegler, Juuso Välimäki and conference and seminar participants for helpful comments and suggestions. Aislinn would like to thank the briq Institute for research support while writing this paper.

†University of Pennsylvania; Email: abohren@sas.upenn.edu

‡Aalto University and Helsinki GSE; Email: daniel.hauser@aalto.fi

1 Introduction

How do individuals learn when they misinterpret information? The literature on misspecified learning typically takes the following approach: fix an incorrect, or *misspecified*, model – such as overreaction to signals or a failure to account for correlated information – and explore how it impacts the long-run beliefs about the state. We know from this literature that asymptotic learning may be *incorrect*, where beliefs converge to the wrong state,[1] *cyclical*, where beliefs do not converge,[2] or *path-dependent*, where multiple learning outcomes arise with positive probability – for example, whether learning is correct or incorrect depends on initial signals.[3]

These insights are usually developed by choosing a parameterized misspecified model that captures a conceptual learning error and assuming that all agents have the same model. But often, there may be multiple parameterizations that capture a given conceptual phenomenon, and these different parameterizations may yield different predictions about asymptotic learning. Alternatively, there may be model heterogeneity, either because agents exhibit varying levels of one form of misspecification or because different agents have conceptually distinct forms of misspecification. Such heterogeneity raises the question of whether the representative agent approach is valid. Finally, the same form of misspecification may have different implications for learning depending on whether the source of information is private, public or social.

This paper develops a general framework to study how misinterpreting information impacts learning. We consider a setting in which individuals learn from signals and the actions of predecessors. Signals can be private, such as past experiences in similar situations, or public, such as the announcements of a government or health agency. Critically, individuals have misspecified models of how to interpret these sources. We develop a simple criterion to characterize long-run beliefs and behavior based on the underlying form of misspecification. This characterization allows for a deeper understanding of how different forms of misspecification impact learning. Specifically, it can be used to explore whether a given form of misspecification is *conceptually robust*, in that its implications are not sensitive to parametric specification, and to determine whether the misspecification has a similar impact in individual and social learning settings. It can also be used to show that misspecification in the settings we consider is *analytically* robust, in that nearby models (in a measure-theoretic sense) generate similar learning outcomes. Therefore, the predictions of a correctly-specified model are a good approximation for

[1] Rabin and Schrag (1999); Epstein, Noor, and Sandroni (2010); Gagnon-Bartsch and Rabin (2017); Frick, Iijima, and Ishii (2019)

[2] Nyarko (1991); Bohren (2016); Fudenberg, Romanyuk, and Strack (2017); Gagnon-Bartsch (2017); Frick, Iijima, and Ishii (2018); Heidhues, Koszegi, and Strack (2018)

[3] Guarino and Jehiel (2013); Eyster and Rabin (2010); Bohren (2016)

settings with small levels of misspecification. In the case of heterogeneity, the characterization can be used to determine how agents with different models influence each others' learning and to evaluate whether a representative agent model is a good approximation. When agents' models are sufficiently distinct, we show that heterogeneity can lead to perpetual *disagreement*, with the beliefs of agents converging to different states despite observing each others' actions. Finally, we demonstrate how our framework can be used to connect conceptually distinct forms of model misspecification that have similar implications for learning.

Our framework captures a rich array of settings in which individuals are systematically biased when processing information and interpreting others' choices. Depending on the context, the empirical literature in psychology and economics documents that individuals overreact or underreact to new information (e.g. overconfidence), slant information towards a preferred state (e.g. motivated reasoning, partisan bias), differentially weight information based on their prior beliefs (e.g. confirmation bias), incorrectly aggregate correlated information (e.g. redundancy neglect), misunderstand strategic interaction (e.g. level-k, cognitive hierarchy), and misperceive others' preferences (e.g. false consensus effect, pluralistic ignorance).[4,5] Our framework can represent these cognitive biases as forms of model misspecification where individuals have incorrect models of the informational environment and how others make decisions. Importantly, the framework is not specific to a given set of biases and can be used to model broad classes of systematic deviations from Bayesian learning with a correctly specified model. For example, it can be used to represent non-Bayesian learning rules, such as the counting heuristic (Ungeheuer and Weber 2017). This representation provides substantial added structure and tractability for analysis.

Our framework encompasses sequential social learning and passive individual learn-

[4]Overconfidence: Edward (1982); Griffin and Tversky (1992); Moore and Healy (2008); Ortoleva and Snowberg (2015); Underconfidence (in social information): Angrisani, Guarino, Jehiel, and Kitagawa (2018); Motivated reasoning / partisan bias: Kunda (1990); Bartels (2002); Brunnermeier and Parker (2005); Köszegi (2006); Bénabou and Tirole (2011); Jerit and Barabas (2012); Confirmation bias: Lord, Ross, and Lepper (1979); Darley and Gross (1983); Plous (1991); Correlation neglect: Kallir and Sonsino (2009); Eyster and Weizsacker (2011); Eyster, Rabin, and Weizsäcker (2018); Enke and Zimmermann (2019); Level-k / cognitive hierarchy: Kübler and Weizsäcker (2004, 2005); Penczynski (2017) Social perception bias: Ross, Greene, and House (1977); Marks and Miller (1987); Miller and McFarland (1987); Gilovich (1990); Miller and McFarland (1991); Grebe, Schmid, and Stiehler (2008). Theories of cognitive limitations provide a foundation for such biases. For example, bounded memory leads to behavior consistent with many documented behavioral biases, including belief polarization, confirmation bias and stickiness (Wilson 2014), while selective awareness leads to confirmation bias and conservatism bias (Gottlieb 2015).

[5]The context of the learning setting will determine which biases are of first order relevance. For example, pluralistic ignorance often arises in contexts where agents believe that a negative trait affects their own behavior, while the false consensus effect arises for non-normative behaviors such as sexual promiscuity or smoking.

ing settings. Each agent is faced with a decision problem in which she selects an action; her payoff depends on this action and an unknown state of the world. Prior to choosing an action, she learns about the state from the actions of her predecessors, a private signal and/or a sequence of public signals. The agent's type specifies her preferences as well as how she interprets signals and prior actions, captured by her subjective beliefs about the signal distribution, others' preferences, and how others interpret signals and prior actions. Model misspecification refers to the case where these subjective distributions *differ* from the true distributions. To maintain structure, we focus on *aligned* environments in which types have a common interpretation of which signals are stronger evidence for each state and the same ordinal, but not necessarily cardinal, preferences over undominated actions. This framework captures all of the information-processing biases cited above and nests several previously developed behavioral models of learning.[6]

Our first main result (Theorem 1) characterizes how asymptotic learning outcomes depend on the form of misspecification. We show that the set of asymptotic learning outcomes that arise with positive probability is determined by two expressions that are straightforward to derive from the primitives of the misspecification: (i) the *expected change in the log likelihood ratio* for each type at each candidate learning outcome; and (ii) an ordering over the type space at an agreement outcome, which we refer to as *maximal accessibility*. The first condition is used to determine whether a learning outcome is locally stable, in that beliefs converge to this limit belief with positive probability from a neighborhood of the limit belief. Building on techniques used in Smith and Sorensen (2000); Bohren (2016), we show that a learning outcome is locally stable if and only if the the log likelihood ratio moves toward this learning outcome in expectation from nearby beliefs.[7] The second condition determines when, starting from a common prior, it is possible to separate the beliefs of different types and push them to the neighborhood of a disagreement outcome (this step follows immediately from the first condition for *agreement* outcomes, where all types have the same limit beliefs).

A challenge in social learning settings is that the informational content of actions depends on the current belief for each type. Therefore, in principle, the asymptotic properties of beliefs could depend on the behavior of beliefs across the infinite belief space. An important feature of our characterization is that the conditions we outline

[6]The applications in Section 5 demonstrate how our framework nests Rabin and Schrag (1999); Epstein et al. (2010); Bohren (2016). More generally, we nest heuristics that reduce to Markovian updating rules. Our framework cannot nest heuristics that reduce to non-Markovian updating rules or are calibrated based on equilibrium objects (e.g. the analogy-based expectation equilibria in Guarino and Jehiel (2013)).

[7]This condition relates to the relative entropy of a type's model in each state. Intuitively, a type's beliefs move towards the state that is more likely to generate the observed pattern of actions and signals. As discussed below, limit beliefs are a Berk-Nash equilibrium (Esponda and Pouzo 2016).

only need to be verified at a *finite* set of beliefs: that is, the set of beliefs in which all types have degenerate beliefs on one of the states. This significantly simplifies the calculations required to use the characterization in specific settings.

Our characterization also establishes a robustness property (Theorem 2): regardless of the form of misspecification, agents almost surely learn the correct state when they have approximately correct models.[8] This shows that agents do not need to know exactly how their peers behave in order to learn from their choices. Complete learning obtains even in the presence of model heterogeneity, as long as all models are sufficiently close to the correct model. This may not seem surprising, since Bayes rule is continuous. But in an infinite horizon setting, a small mistake in each period could sum to a large aggregate mistake. If biases aggregate in this manner, then even arbitrarily small misspecification could interfere with learning, which would in principle limit the applicability of rational learning models. Our results establish when this does not occur.

We close with several applications to illustrate the value of a general framework for studying learning with model misspecification. In the first two applications, we explore whether two forms of signal misspecification – over- and underreaction and confirmation bias – are sensitive to the parametric specifications used to model how agents misinterpret information. We demonstrate a fundamental difference in robustness: confirmation bias is conceptually robust, while over- and underreaction are not. We then demonstrate that the same parameterization of over- and underreaction has a qualitatively different impact in individual and social learning settings. Learning from others' actions interacts with the signal misspecification to lead to cyclical learning in the case of sufficiently severe overreaction and incorrect learning in the case of sufficiently severe underreaction. In contrast, when individuals learn solely from signals, correct learning obtains almost surely for both cases.

In the third application, we examine whether a representative agent model is a good approximation for a setting with model heterogeneity in which agents vary in the degree to which they fail to account for redundant information. When heterogeneity is small, the representative agent model is a good approximation. In contrast, when heterogeneity is large, differing levels of bias actually facilitate learning, in that correct learning obtains for a strictly larger set of parameters compared to the corresponding representative agent model.

In the fourth application, we derive an equivalence in terms of asymptotic learning

[8]The robustness result in Bohren (2016) is a special case of this result. In contrast, Frick et al. (2019) show that robustness can fail in a social learning model with an infinite state space, in that arbitrarily small amounts of misspecification can lead to asymptotic beliefs that are very far from the true state. Madarász and Prat (2016) also find a lack of robustness in a mechanism design setting where a principal's model of an agent's preferences is misspecified.

outcomes between two conceptually distinct forms of misspecification – naive temptation, a form preference misspecification, and partisan bias, a form of signal misspecification. We show that for each level of temptation, there exists a level of partisan bias that yields the same set of asymptotic learning outcomes.

Finally, we study a level-k model of reasoning. This form of model misspecification is prominent in the empirical literature on social learning but relatively unexplored in the theoretical literature, which typically focuses on learning for a single type. It is straightforward to apply our characterization to this setting with multiple types. We show that entrenched disagreement emerges as a robust feature of social learning in a level-k model.

Related Literature. A rich literature explores when model misspecification interferes with learning in both individual and social learning settings. The results are mixed: in some cases, misspecification impedes learning about the state or leads to inefficient behavior, while in other cases, misspecified agents still learn the correct state asymptotically. For example, in a *passive* learning setting, where information arrives independently of past action choices, overweighting information (Epstein et al. 2010; Rabin and Schrag 1999) can lead to incorrect learning while underweighting information (Epstein et al. 2010) does not impede correct learning. In a *social* learning setting, where agents learn from their peers, failing to account for redundant information (Eyster and Rabin 2010; Bohren 2016; Gagnon-Bartsch and Rabin 2017) can lead to incorrect learning, while overestimating redundant information (Bohren 2016) or overestimating the similarity of others' preferences (Gagnon-Bartsch 2017) can lead to non-convergence. In contrast, using coarse reasoning (Guarino and Jehiel 2013) or using a linear updating heuristic that puts sufficient weight on agents' own signals (Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi 2012) leads to correct learning almost surely. In an *active* learning setting, where future signals depend on past action choices, selective attention (Schwartzstein 2014) can lead to incorrect learning, misspecified prior beliefs (Fudenberg et al. 2017; Nyarko 1991) can lead to non-convergence, and overconfidence can lead to inefficiently low effort (Heidhues et al. 2018). In other settings, correlation neglect can lead to inefficient risk-taking (Levy and Razin 2015) and ideological extremeness (Ortoleva and Snowberg 2015).

Our paper is the first to explore a general model of misspecification in social learning settings. By characterizing when misspecification interferes with learning and when it does not, our framework unifies insights about the impact of different forms of misspecification in social and passive learning settings (our characterization does not apply to active learning settings). Molavi, Tahbaz-Salehi, and Jadbabaie (2018) engage in a similar exercise when agents share their beliefs on a network and have imperfect recall.

They nest common learning rules that agents use to aggregate *beliefs* (as opposed to actions), including the canonical Degroot model, and show how this impacts long-run information aggregation. Heidhues, Koszegi, and Strack (2019) develop tools to study convergence in a general class of active learning models with misspecification.

Esponda and Pouzo (2016, 2017) explore the implications of model misspecification for solution concepts. In a Berk-Nash equilibrium, agents have a set of (possibly misspecified) models of the world and play optimally with respect to the model that is the best fit, i.e. the model that minimizes relative entropy with respect to the distribution of outcomes under the equilibrium strategy profile. In our paper, each state corresponds to a model of the world in the Esponda and Pouzo (2016) framework. We show that if beliefs about the state converge, then they converge to a limit belief such that each type's model at that limit belief is the best fit. This set of limit beliefs in our learning characterization is analogous to pure strategy Berk-Nash equilibria. However, there may also be pure strategy Berk-Nash equilibria which are almost surely not limit beliefs in our setting. Further, there may be mixed strategy Berk-Nash equilibria; these equilibria have no analogue in our dynamic setting. In particular, when beliefs do not converge in our setting, the empirical frequency of actions does not correspond to the probability of each action in a mixed strategy Berk-Nash equilibrium.

A related class of papers explore the foundations of non-Bayesian updating and model misspecification. Ortoleva (2012) axiomatizes a non-Bayesian updating rule in which agents switch models when they observe a sufficiently low probability event. Cripps (2018) axiomatizes a class of non-Bayesian updating processes that are independent of how individuals partition information. Frick et al. (2018) show that the false consensus effect can arise when agents' beliefs are derived only from local interactions in an assortative society.

An older statistics literature on model misspecification complements recent work. Berk (1966) and Kleijn and van der Vaart (2006) show that when an agent with a misspecified model is learning from i.i.d. draws of a signal, her beliefs will converge to the distribution that minimizes relative entropy with respect to the true model. Shalizi (2009) extends these result to a class of non-i.i.d. signal processes. Our social learning setting does not fall into this class of processes. In particular, the asymptotic-equipartition property, which describes the long-run behavior of the sample entropy, is generally not satisfied in social learning environments with model misspecification.

The paper proceeds as follows. Section 2 sets up the model. Section 3 outlines the agent's decision problem and evolution of beliefs. Section 4 presents the asymptotic learning characterization and robustness results. Section 5 demonstrates applications of the model, while Section 6 concludes. The Appendix contains the proofs of the main results in the paper, while the Online Appendix takes up extensions and related technical issues.

2 The Common Framework

2.1 The Model

States and Actions. There are two payoff-relevant states of the world, $\omega \in \{L, R\}$. Nature selects one of these states at the beginning of the game according to prior $p_0 \equiv Pr(\omega = R)$. A countably infinite set of agents $t = 1, 2, ...$ act sequentially and choose an action from a finite set $\mathcal{A}$ with $M \equiv |\mathcal{A}| \geq 2$ actions. Let $\tilde{a}_t$ denote agent t's chosen action.[9]

Signals and Histories. Agents learn from private information, public information, and the actions of other agents. Agent t observes a private signal $\tilde{s}_t \in (0, 1)$ governed by conditional c.d.f. F^ω in state ω, the ordered history of past actions $(\tilde{a}_1, ..., \tilde{a}_{t-1})$, and the ordered history of public signals $(\tilde{\sigma}_1, ..., \tilde{\sigma}_{t-1})$, where $\tilde{\sigma}_\tau \in (0, 1)$ is governed by conditional c.d.f. G^ω in state ω.[10] Let $h_t \equiv (\tilde{a}_1, ..., \tilde{a}_{t-1}, \tilde{\sigma}_1, ..., \tilde{\sigma}_{t-1})$ denote the public history.

Signals $\langle \tilde{s}_t \rangle$ and $\langle \tilde{\sigma}_t \rangle$ are i.i.d. across time, conditional on the state and jointly independent. No private or public signal perfectly reveals the state: F^L, F^R and G^L, G^R are mutually absolutely continuous with common supports, denoted by $\mathcal{S}$ and Σ, respectively. Therefore, there exist positive finite Radon-Nikodym derivatives dF^R/dF^L and dG^R/dG^L. Assume that some signals are *informative*, which rules out the trivial case where $dF^R/dF^L = 1$ almost surely and $dG^R/dG^L = 1$ almost surely, and assume that the public signal space Σ is finite. As is standard convention, normalize each signal realization to be the posterior probability that the state is R following a neutral prior, i.e. $s = 1/(1 + dF^L/dF^R(s))$ for all $s \in \mathcal{S}$ and $\sigma = 1/(1 + dG^L/dG^R(\sigma))$ for all $\sigma \in \Sigma$. Private beliefs are bounded if $\inf \mathcal{S} > 0$ and $\sup \mathcal{S} < 1$ and unbounded if the closure of the convex hull of $\mathcal{S}$ is $[0, 1]$. Let $\sigma_R \equiv \sup \Sigma$ denote the public signal that corresponds to the strongest evidence for state R, which we refer to as the *maximal public signal* in state R. Analogously, let $\sigma_L \equiv \inf \Sigma$ denote the maximal public signal in state L.

Types. Each agent has a privately observed type $\tilde{\theta}_t \in \Theta$ drawn from distribution $\pi \in \Delta(\Theta)$, where $\Theta \equiv (\theta_1, ..., \theta_n)$ is a non-empty finite set. An agent's type specifies her model of inference and preferences. A model of inference determines how a type interprets information from signals and actions. Preferences determine which action this type chooses, given its belief about the state. The type space Θ captures both types

[9] We maintain the convention that x_i or x corresponds to an arbitrary element of an ordered set $\mathcal{X}$ and $\tilde{x}_t$ corresponds to a random variable with support $\mathcal{X}$.

[10] Similarly, we maintain the convention that $\tilde{s}_t$ and $\tilde{\sigma}_t$ correspond to the random variable and s and σ denote an arbitrary element $\mathcal{A}$.

that occur with positive probability and types that an agent may mistakenly attribute to other agents; therefore, π may assign probability zero to some types in Θ. Types $\langle\tilde{\theta}_t\rangle$ are i.i.d. across time and independent of both signals.

Models of Inference. For each type θ_i, a model of inference includes (i) a subjective private signal distribution $\hat{F}_i^\omega$ in each state, (ii) a subjective public signal distribution $\hat{G}_i^\omega$ in each state, and (iii) a subjective distribution of types $\hat{\pi}_i \in \Delta(\Theta)$. Assume that each type θ_i believes that no private or public signal perfectly reveals the state and does not observe a signal that is inconsistent with its model of inference: $(\hat{F}_i^L, \hat{F}_i^R)$ and $(\hat{G}_i^L, \hat{G}_i^R)$ are mutually absolutely continuous and have full support on $\mathcal{S}$ and Σ, respectively. Given private signal s and a flat prior, type θ_i's subjective private belief that the state is R is $\hat{s}_i(s) \equiv 1/(1+d\hat{F}_i^L/d\hat{F}_i^R(s))$.[11] Similarly, given public signal σ and a flat prior, type θ_i's subjective belief that the state is R is $\hat{\sigma}_i(\sigma) \equiv 1/(1+d\hat{G}_i^L/d\hat{G}_i^R(\sigma))$. All types share common prior p_0 that the state is R.[12]

We focus on forms of misspecification in which agents have a partial common understanding of how to interpret signals. We assume that each type's subjective private and public signal distributions are *aligned* with the true private and public signal distributions, in that they ordinally rank signals in the same way, in terms of which signals are more or less indicative of state R. In other words, for any two signals s and s', if s is stronger evidence for state R than s' under the true measure, then s is also stronger evidence for state R than s' under the subjective measure. We allow one exception for the possibility that a type believes signals are completely uninformative.

Assumption 1 (Aligned Subjective Signals)**.** *For all $\theta_i \in \Theta$, the subjective private signal distributions are either* aligned *with the true private signal distributions, i.e. for any $s, s' \in \mathcal{S}$ such that $s > s'$, then $\hat{s}_i(s) > \hat{s}_i(s')$ or uninformative, i.e. $\hat{s}_i(s) = 1/2$ for all $s \in \mathcal{S}$. Analogously, the subjective public signal distributions are either aligned with the true public signal distributions or uninformative.*

This assumption allows types to differ in the degree to which a signal influences their belief about the state – both relative to other types and relative to the true distribution.

[11] This set-up implicitly restricts attention to forms of misspecification in which two signals that map into the same true posterior also map into the same subjective posterior. In Online Appendix D, we show that under this restriction, it is without loss of generality to define subjective signal distributions with respect to the private belief space $\mathcal{S} \subset [0,1]$ rather than an arbitrary signal space. We can also define misspecification relative to the relationship between the subjective and true private beliefs, $\hat{s}$ and s. In Online Appendix D, we show that for any strictly increasing function $\hat{s} : \mathcal{S} \to [0,1]$ with $\hat{s}(\inf \mathcal{S}) < 1/2$ and $\hat{s}(\sup \mathcal{S}) > 1/2$, there exists a pair of mutually absolutely continuous probability measures with full support on the signal space that are represented by $\hat{s}$.

[12] It is straightforward to extend the types framework to allow for heterogenous prior beliefs about the state, i.e. type θ_i has a prior belief $p_{i,0}$.

Note that when public signals are aligned, the maximal public signals σ_L and σ_R are also maximal with respect to each agent's subjective public signal distribution.

Preferences. Type θ_i earns payoff $u_i(a, \omega)$ from choosing action a in state ω, where $u_i : \mathcal{A} \times \{L, R\} \to \mathbb{R}$. Given a belief $p \in [0, 1]$ that the state is R, the expected payoff from choosing action a is $(1 - p)u_i(a, L) + pu_i(a, R)$. An agent chooses the action that maximizes her expected payoff. For each type, assume that at least two actions are not weakly dominated, no two actions yield the same payoff in both states, and no action is optimal at a single belief. Without loss of generality, assume that no action is dominated for all types.

We focus on settings in which agents generate information in a common way, in terms of their action choices, by restricting how preferences vary across types. A set of utility functions are aligned if, under complete information, each utility function has the same ordinal ranking over undominated actions.

Assumption 2 (Aligned Preferences)**.** *The set of types Θ have* aligned *preferences, in that there exists a complete order $\succ$ on $\mathcal{A}$ such that if $a \succ a'$, then for all $i = 1, ..., n$, either $u_i(a, R) > u_i(a', R)$ or a is dominated.*[13]

Assumption 2 implies common knowledge that preferences are aligned, since all agents believe that other agents have a type in Θ, and so on. This assumption places no restrictions on how to order actions that are optimal for a single type or how a type ranks its dominated actions. Smith and Sorensen (2000) establish that confounded learning can arise when types have preferences that are not aligned, such as $u_1 = \mathbb{1}_{a=\omega}$ and $u_2 = \mathbb{1}_{a\neq\omega}$. The same is true with misspecification. Therefore, we restrict attention to settings in which confounded learning does not arise in the correctly specified model.

Given Assumption 2, we maintain a complete order over the action space $\mathcal{A}$ by relative preference in state R. Fixing an order $\succ$ that satisfies Assumption 2, index actions to correspond to this order, i.e. $\mathcal{A} \equiv (a_1, ..., a_M)$, where $a_m \succ a_l$ iff $m > l$.[14] Under this order, a_M denotes the *maximal action* in state R, and a_1 denotes the *maximal action* in state L.

Categories of Types. We can broadly group types into four categories based on their models of inference: noise, autarkic, sociable and correct. A *noise* type does not use

[13] For any undominated actions a and a', if $u_i(a, R) > u_i(a', R)$, then $u_i(a, L) < u_i(a', L)$. Therefore, this assumption implies that utility functions also have the same ordinal ranking over undominated actions in state L, i.e. if $a \succ a'$, then for all $i = 1, ..., n$, either $u_i(a', L) > u_i(a, L)$ or a' is dominated.

[14] This order is not necessarily unique. Assumption 2 places no restriction on how actions that are optimal for a single type are ordered. For example, if one type chooses action R_1 when $p \geq 1/2$, and otherwise chooses L_1, and a second type chooses action R_2 when $p \geq 1/2$, and otherwise chooses L_2, then both the orders $R_1 \succ R_2 \succ L_1 \succ L_2$ and $R_2 \succ R_1 \succ L_1 \succ L_2$ satisfy Assumption 2.

its private signal or the history to learn about the state. We can model this using the types framework by defining a noise type to believe that private and public signals are uninformative, $\hat{F}_i^L = \hat{F}_i^R$ and $\hat{G}_i^L = \hat{G}_i^R$. Noise types also believe that actions reflect no information about the state, which is modeled as the belief that all agents are noise types, $\hat{\pi}_i(\Theta_N) = 1$, where Θ_N denotes the set of noise types. An *autarkic* type learns from its private signal, but not the history. It believes that its private signal is informative, $\hat{F}_i^L \neq \hat{F}_i^R$, the public signal is uninformative, $\hat{G}_i^L = \hat{G}_i^R$, and all agents are noise types, $\hat{\pi}_i(\Theta_N) = 1$. To avoid the trivial case in which an autarkic type is observationally equivalent to a noise type, we assume that an autarkic type has preferences such that it has at least two undominated actions on the set of posterior beliefs that arise from its subjective private signal distribution. A type is *sociable* if it uses the history to learn about the state. These types believe that either actions or the public signal are informative. Finally, a *correct* type has a correct model of inference, $\hat{F}_i^\omega = F^\omega$, $\hat{G}_i^\omega = G^\omega$ and $\hat{\pi}_i = \pi$.

Let Θ be ordered such that the first k types are sociable and the remaining $n-k$ types are noise and autarkic types. Let $\Theta_S = (\theta_1, ..., \theta_k)$ denote the set of sociable types, Θ_A denote the set of autarkic types and Θ_N denote the set of noise types.

Adequate Consistent Information. We focus on settings in which adequate information arrives for agents to learn the state in a correctly specified model, and study whether and how misspecification interferes with such learning. We know from Smith and Sorensen (2000) that incomplete learning arises in correctly specified models when there are no public signals, no autarkic types, and private signals are uniformly bounded in strength. The same is true for misspecified models: if actions and public signals cease to reveal information and all types are aware of this, then learning will be incomplete. Assumption 3 rules out such settings by assuming that either public signals are informative or autarkic types occur with positive probability. Since autarkic types do not observe the history, their actions are always informative.

Assumption 3 (Adequate Information)**.** *Either (i) public signals are informative, $dG^R/dG^L \neq 1$, and all sociable types $\theta_i \in \Theta_S$ believe that public signals are informative, $d\hat{G}_i^R/d\hat{G}_i^L \neq 1$, or (ii) there exists an autarkic type $\theta_j \in \Theta_A$ with $\pi(\theta_j) > 0$ that plays actions a_1 and a_M with positive probability, and each sociable type $\theta_i \in \Theta_S$ believes this autarkic type exists, $\hat{\pi}_i(\theta_j) > 0$.*

This assumption ensures that actions or public signals are informative, and sociable types believe that actions or public signals are informative.[15]

[15]While a setting with unbounded private signals would also guarantee complete learning in a cor-

We also focus on settings in which the observed history is consistent with each type's model of inference, in that types do not observe what they believe to be zero-probability histories. Trivially, if there is a single type, $|\Theta| = 1$, then this type has a correct model of the type distribution and consistency is not an issue.[16] With multiple types, a type may have a model of inference that places probability zero on an action that occurs with positive probability. To rule out this possibility, we assume that sociable types believe that there is an autarkic or noise type that plays each action with positive probability (this probability can be arbitrarily small).

Assumption 4 (Consistent Information)**.** *When there are multiple types, $|\Theta| \geq 2$, then for each $a \in \mathcal{A}$ and for each sociable type $\theta_i \in \Theta_S$, there exists an autarkic or noise type $\theta_j \in \Theta_A \cup \Theta_N$ with $\hat{\pi}_i(\theta_j) > 0$ that plays a with positive probability.*

This ensures that each sociable type believes that all histories are on the equilibrium path, and we do not need to model how a type reacts to zero probability events.

Any misspecified model can be slightly perturbed so that it satisfies Assumption 3 and 4 by either (i) perturbing an uninformative public signal distribution so that it is slightly informative, or (ii) perturbing the type distribution to add an autarkic or noise type that occurs with arbitrarily small probability.

Timing. At time t, agent t realizes its type $\tilde{\theta}_t$ and observes the history h_t and private signal $\tilde{s}_t$, then chooses action $\tilde{a}_t$. Then public signal $\tilde{\sigma}_t$ is realized and the history is updated to h_{t+1} to include $(\tilde{a}_t, \tilde{\sigma}_t)$.[17]

2.2 Examples

This framework can capture many information-processing biases, social misperceptions and other models of inference that have been studied theoretically and empirically. The following examples demonstrate how our types framework captures various forms of model misspecification.

Partisan Bias. This captures an information-processing bias that systematically slants signals towards one state (Bartels 2002; Jerit and Barabas 2012). For example, an R-

rectly specified model, our learning characterization requires a stronger notion of *adequate information*, where the minimum and maximum update to beliefs are uniformly bounded away from zero across the belief space. The presence of even a weakly informative public signal or small share of autarkic types guarantees that there is adequate information, whereas unbounded private signals do not. When actions are uninformative at certainty, as is the case with unbounded private signals, and there is no additional source of information (i.e. public signals or autarkic types), learning outcomes will depend on the tail properties of the private signal distribution.

[16] When $|\Theta| = 1$, it must be that $\hat{\pi}^1(\theta_1) = \pi(\theta_1) = 1$, and all observed actions will be consistent. Note that even if only one type actually exists, $\pi(\theta_1) = 1$, if this type believes that there is another type θ_2, i.e. $\hat{\pi}^1(\theta_2) > 0$, then $\Theta = (\theta_1, \theta_2)$ and we are in the case with more than one type.

[17] Allowing agent t to observe $\tilde{\sigma}_t$ before choosing an action does not affect the analysis, but complicates the notation.

partisan type θ_i interprets all signals as being stronger evidence for state R than is actually the case, $\hat{s}_i(s) > s$. A parameterization that captures this is $\hat{F}_i^\omega(s) = F^\omega(s^\nu)$ for $\nu \in (0,1)$ and $F^\omega(s)$ with support $\mathcal{S} = (0,1)$. This leads to private belief $\hat{s}_i(s) = s^\nu$ – after observing signal s, the partisan type has the same belief as a correctly specified type who observes signal s^ν.

Under- and Overreaction. Agents underreact or overreact to signals. For example, $\frac{\hat{s}(s)}{1-\hat{s}(s)} = (\frac{s}{1-s})^\nu$, where $\nu \in [0,1)$ corresponds to underreaction and $\nu \in (1,\infty)$ corresponds to overreaction (Moore and Healy 2008; Epstein et al. 2010; Angrisani et al. 2018).

Correlation Neglect/Naive Learning. Agents underestimate the correlation in the actions of prior agents: the true share of autarkic types is $\pi(\Theta_A)$, but sociable types believe that the share of autarkic types is $\hat{\pi}(\Theta_A) > \pi(\Theta_A)$ (Eyster and Rabin 2010; Bohren 2016; Eyster et al. 2018; Enke and Zimmermann 2019).

Level-k/Cognitive Hierarchy. Level-1 believes all agents are noise types and behaves as an autarkic type. Level-2 believes all agents are level-1 and interprets each prior action as reflecting an independent private signal. Level-3 believes all agents are level-2, and so on. The cognitive hierarchy model is similar, but allows agents to have a richer belief structure over the types of other agents: a level-k type places positive probability on levels 0 through k-1 (Penczynski 2017).

False Consensus Effect. Agents overweight the likelihood that others have similar preferences or models of inference. For example, there are two types with preferences $u_1 \neq u_2$. Both types believe that others have the same preferences as their own, $\hat{\pi}_1(\theta_1) = 1$ and $\hat{\pi}_2(\theta_2) = 1$ (Ross et al. 1977; Marks and Miller 1987; Gagnon-Bartsch 2017; Frick et al. 2018).

Pluralistic Ignorance. Agents underweight the likelihood that others have similar preferences or models of inference. For example, all agents have preferences u_1, but believe that others have preferences u_2, $\pi(\theta_1) = 1$ and $\hat{\pi}_1(\theta_2) = 1$. Alternatively, all agents correctly interpret private signals, but believe that others are overconfident (Miller and McFarland 1987, 1991).

Limited Recall. An alternative interpretation of the social learning setting is that there is a single long-run agent with limited memory: she can recall past actions but not past signals. The agent may also be misspecified in how she recalls these actions (for example, she has selective recall due to motivated reasoning).

The following example contains a complete specification of our framework. We use this example throughout the paper to illustrate our main results.

Example 1 (Partisan Bias). *There are two types of agents: θ_1 is a sociable type and θ_2 is an autarkic type. Both types seek to choose the action that matches the state, $\mathcal{A} = \{L, R\}$ and $u_i(a, \omega) = \mathbb{1}_{a=\omega}$, and have an identical level of partisan bias $\nu \in (0, 1)$, where $\hat{F}_i^\omega(s) = F^\omega(s^\nu)$ as introduced above. To close the model, assume that the public signal is uninformative, the private signal distribution is informative and symmetric, $F^R(s) = 1 - F^L(1 - s)$, there is a positive share $\pi(\theta_2) \in (0, 1)$ of autarkic types, type θ_1 has a correctly specified type distribution, $\hat{\pi}_1 = \pi$, and the prior $p_0 = 1/2$. Trivially, signals and preferences are aligned (Assumption 1 and 2) since both types have the same subjective signal distributions and preferences. We show in Appendix B that Assumption 3.ii holds and Assumption 4 is redundant in a binary action decision problem that satisfies Assumption 3.ii.*

2.3 Discussion of Model

We briefly comment on notable features of the model and possible extensions.

The Type Framework. An agent's type captures her model of the world, which she uses to learn about the environment. We can divide type θ_i's model into a model of the world in state L, which consists of the signal and type distributions in state L, $(\hat{F}_i^L, \hat{G}_i^L, \hat{\pi}_i)$, and analogously, a model of the world in state R, $(\hat{F}_i^R, \hat{G}_i^R, \hat{\pi}_i)$.[18] We take these models as fixed and explore learning about features of the environment that are directly payoff-relevant, i.e. the state.

We focus on *aligned* type spaces in which all agents generate and interpret information in a common way. Aligned signals (Assumption 1) guarantee that agents have a common interpretation of the relative order of signals as evidence for state R, and aligned preferences (Assumption 2) guarantee that the action choices of agents are ordered in a way that reflects the same relative strength of evidence for state R. For example, suppose it is common knowledge that lung cancer is stronger evidence that smoking has a negative impact on the lungs, relative to shortness of breath, but agents differ in their beliefs about the strength of these two signals. Or suppose that agents have the same preferences between a risky and a safe asset when they are certain about the state (whether the risky asset has a high or low expected return), but differ in their risk preferences, and therefore, the threshold belief at which they are willing to start investing in the risky asset. Note that these assumptions do not require all types to choose the same actions at certainty. For example, one type may be systematically more risk averse than another type and prefer less risky actions across all beliefs about the state. A natural economic setting that aligned environments do rule out is some versions of

[18]We implicitly assume that the type distribution is the same in both states. It is a straightforward extension to allow the true and/or subjective type distributions to depend on ω.

horizontally differentiated environments (e.g. the horizontally differentiated preferences studied in Gagnon-Bartsch (2017) in relation to the false consensus effect).

When there are multiple types, agents have heterogenous models and may be misspecified about how other agents learn. This can lead to complicated higher-order beliefs. For example, when an agent believes that other agents have partisan bias, we also need to model what this agent believes that these partisan bias agents believe about others. In our framework, these higher-order beliefs are fully captured by the subjective type distributions. If type θ_i believes that all agents are type θ_j, then type θ_j's subjective distribution $\hat{\pi}_j$ captures θ_i's second order beliefs, the subjective type distributions of the types in the support of $\hat{\pi}_j$ capture third order beliefs, and so on. Therefore, the type space Θ fully determines the hierarchies of beliefs. The assumption that Θ is finite limits the number of models that each type can attribute to other agents. It also rules out infinite chains of models of the form: type θ_i believes all agents are type θ_{i+1}, type θ_{i+1} believes all agents are θ_{i+2}, etc. for $i = 1, 2, ...$.

Individual Learning. Our framework nests an individual learning model in which a long-run agent learns from a sequence of exogenous signals. This is captured by a specification with an informative public signal, an uninformative private signal, and a single type who believes that the public signal is informative and the private signal is uninformative. In this set-up, actions contain no private information and there is no social learning. Therefore, it is isomorphic to a model in which a single long-run agent observes a sequence of signals or multiple long-run agents observe the same sequence of signals.

Extensions. It is straightforward to allow types to receive private signals from different distributions or to believe that other types receive signals from different distributions. To capture this, augment the definition of a type to include both the type's true signal distribution, i.e. type θ_i has private signals drawn from F_i^ω, as well as type's subjective belief about its signal distribution. This extension allows us to model biases that involve interpersonal comparisons related to the quality of information. For instance, a natural way to model an overconfident agent is with a type that believes it draws signals from a more informative distribution than everyone else, i.e. all types observe signals from the same distribution and correctly interpret them, but the overconfident type believes that other agents observe signals from a less informative distribution. The analysis carries through unchanged, albeit with more burdensome notation.

It is also straightforward to allow a type's model of inference to depend on its current belief about the state. This extension allows us to analyze several additional classes of model misspecification, including confirmation bias (nesting (Rabin and Schrag 1999)) and forms of under- and overreaction that vary with the current belief (nesting Epstein et al. (2010)). We explore this extension further in the applications in Section 5 and present the formal set-up in Online Appendix E.4.

In many situations, agents learn from observing the outcomes of others' choices, rather than directly observing their actions. Our learning characterization extends to a setting in which agents learn from outcomes that stochastically depend on actions (see Bohren and Hauser (2019)). A special case of this extension is a setting in which stochastic outcomes are the only source of information (i.e. signals are uninformative) and there is a single sociable type. This case is equivalent to an active individual learning model in which a single myopic long-run agent chooses actions that influence the distribution of information.

For technical convenience, we assume that the action and public signal spaces are finite. Allowing for a continuous action and public signal space would not qualitatively change the analysis. Similar techniques to those we use can be used to analyze a finite state space with more than two states, with the caveats that the definition of an aligned environment is more complicated and the notation is more cumbersome. We use results pertaining to stochastic difference equations in our analysis, which means that generalizing to an infinite state space requires different techniques.

3 Beliefs and Action Choices

Consider an agent of type θ_i who observes history h and private signal s. Assumption 4 ensures that θ_i believes that h occurs with positive probability, and therefore, Bayes rule can be used to update beliefs. The agent uses her model of inference to compute the probability of h in each state, $\hat{P}_i(h|\omega)$, and applies Bayes rule to form the likelihood ratio

$$\lambda_i(h) \equiv \frac{\hat{P}_i(R|h)}{\hat{P}_i(L|h)} = \left(\frac{p_0}{1-p_0}\right)\frac{\hat{P}_i(h|R)}{\hat{P}_i(h|L)} \tag{1}$$

that the state is R versus L. The agent then uses her subjective signal distribution $\hat{s}_i(s)$ to update to

$$\frac{p_i(\lambda_i(h), s)}{1-p_i(\lambda_i(h), s)} \equiv \frac{\hat{P}_i(R|h, s)}{\hat{P}_i(L|h, s)} = \lambda_i(h)\left(\frac{\hat{s}_i(s)}{1-\hat{s}_i(s)}\right). \tag{2}$$

She chooses the action that maximizes her expected payoff with respect to $p_i(\lambda_i(h), s)$.

Let $\boldsymbol{\lambda}(h) \equiv (\lambda_1(h), ..., \lambda_k(h))$ denote the vector of likelihood ratios for sociable types $(\theta_1, ..., \theta_k)$ following history h. For autarkic or noise types, $\lambda_i(h) = p_0/(1-p_0)$ for all h, since these types believe that the history is uninformative. To construct $\boldsymbol{\lambda}(h)$, first consider an agent of type θ_i's decision rule at belief $p_i(\lambda, s)$. Recall that actions $(a_1, ...a_M)$ are ordered by relative preference in state R. Since no two actions yield the same payoff in both states, no action is optimal at a single belief, and preferences are aligned (Assumption 2), there exist belief thresholds $0 = \overline{p}_{i,0} \leq \overline{p}_{i,1} \leq ... \leq \overline{p}_{i,M} = 1$ such that we can partition the belief space into a finite set of closed intervals, with action a_m optimal at $p_i(\lambda, s)$ if $p_i(\lambda, s) \in [\overline{p}_{i,m-1}, \overline{p}_{i,m}]$ and $\overline{p}_{i,m-1} \neq \overline{p}_{i,m}$, and a_m never optimal iff $\overline{p}_{i,m-1} = \overline{p}_{i,m}$. Without loss of generality, assume the tie-breaking rule is to choose the optimal action with the lower index at each interior cut-off $\overline{p}_{i,m} \in (0,1)$, i.e. if $\overline{p}_{i,m-1} \neq \overline{p}_{i,m}$, choose a_m at belief $\overline{p}_{i,m}$. Since there are at least two undominated actions, there are at least two intervals with a non-empty interior. Since signals are aligned (Assumption 1), $p_i(\lambda, s)$ is strictly increasing in s for all $\lambda \in (0, \infty)$. Therefore, for each $\lambda \in (0, \infty)$, we can define the decision rule with respect to signal cut-offs $0 = \overline{s}_{i,0}(\lambda) \leq \overline{s}_{i,1}(\lambda) \leq ... \leq \overline{s}_{i,M}(\lambda) = 1$ such that the agent chooses action a_m at likelihood ratio λ iff she observes private signal $s \in (\overline{s}_{i,m-1}(\lambda), \overline{s}_{i,m}(\lambda)]$ and $\overline{s}_{i,m-1}(\lambda) \neq \overline{s}_{i,m}(\lambda)$.

The signal cut-offs for each type determines how sociable types interpret past action choices. Fix a vector of likelihood ratios $\boldsymbol{\lambda} \in (0, \infty)^k$. An agent of type θ_j chooses action a_m when she observes a private signal in the interval $(\overline{s}_{j,m-1}(\lambda_j), \overline{s}_{j,m}(\lambda_j)]$. An agent of type θ_i believes this occurs with probability $\hat{F}_i^\omega(\overline{s}_{j,m}(\lambda_j)) - \hat{F}_i^\omega(\overline{s}_{j,m-1}(\lambda_j))$ in state ω and believes that type θ_j occurs with probability $\hat{\pi}_i(\theta_j)$. Therefore, she believes that action a_m occurs with probability $\sum_{j=1}^n \hat{\pi}_i(\theta_j)(\hat{F}_i^\omega(\overline{s}_{j,m}(\lambda_j)) - \hat{F}_i^\omega(\overline{s}_{j,m-1}(\lambda_j)))$ in state ω. Let

$$\hat{\psi}_i(a_m, \sigma | \omega, \boldsymbol{\lambda}) \equiv d\hat{G}_i^\omega(\sigma) \sum_{j=1}^n \hat{\pi}_i(\theta_j)(\hat{F}_i^\omega(\overline{s}_{j,m}(\lambda_j)) - \hat{F}_i^\omega(\overline{s}_{j,m-1}(\lambda_j))) \tag{3}$$

denote her belief about the joint probability of action a_m and public signal σ in state ω when the likelihood ratio is $\boldsymbol{\lambda}$, where $\hat{G}_i^\omega$ denotes θ_i's subjective public signal distribution.

From these expressions, we can construct $\boldsymbol{\lambda}(h)$. Each sociable type θ_i initially has likelihood ratio $\lambda_i(h_1) = p_0/(1-p_0)$ at history h_1. At any history h_t with $t > 1$,

$$\lambda_i(h_t) = \left(\frac{p_0}{1-p_0}\right) \prod_{\tau=1}^{t-1} \frac{\hat{\psi}_i(\tilde{a}_\tau, \tilde{\sigma}_\tau | R, \boldsymbol{\lambda}(h_\tau))}{\hat{\psi}_i(\tilde{a}_\tau, \tilde{\sigma}_\tau | L, \boldsymbol{\lambda}(h_\tau))}. \tag{4}$$

The process is recursive: given $\boldsymbol{\lambda}(h_t)$ and $(\tilde{a}_t, \tilde{\sigma}_t)$,

$$\lambda_i(h_{t+1}) = \lambda_i(h_t) \left(\frac{\hat{\psi}_i(\tilde{a}_t, \tilde{\sigma}_t | R, \boldsymbol{\lambda}_t)}{\hat{\psi}_i(\tilde{a}_t, \tilde{\sigma}_t | L, \boldsymbol{\lambda}_t)}\right). \tag{5}$$

Therefore, $\boldsymbol{\lambda}_t \equiv \boldsymbol{\lambda}(h_t)$ is sufficient for the history and we suppress the dependence on h_t going forward.

The behavior of $\langle \boldsymbol{\lambda}_t \rangle_{t=1}^\infty$ determines the learning dynamics for each type. While each type's model of inference determines the *value* of the update to the likelihood ratio following action a_m and signal σ, the true probability

$$\psi(a_m, \sigma|\omega, \boldsymbol{\lambda}) \equiv dG^{\omega}(\sigma) \sum_{i=1}^{n} \pi(\theta_i)(F^{\omega}(\bar{s}_{i,m}(\lambda_i)) - F^{\omega}(\bar{s}_{i,m-1}(\lambda_i))). \tag{6}$$

of (a_m, σ) at a vector of likelihood ratios $\boldsymbol{\lambda} \in (0, \infty)^k$ determines the *probability* that the likelihood ration transitions to this value in state ω. In correctly specified models, $\hat{\psi}_i(a, \sigma|\omega, \boldsymbol{\lambda}) = \psi(a, \sigma|\omega, \boldsymbol{\lambda})$, and the likelihood ratio is a martingale in state L. But misspecification introduces a wedge between the subjective and true probability of observing each action and signal, $\hat{\psi}_i(a, \sigma|\omega, \boldsymbol{\lambda}) \neq \psi(a, \sigma|\omega, \boldsymbol{\lambda})$. This makes characterizing the behavior of $\langle \boldsymbol{\lambda}_t \rangle_{t=1}^{\infty}$ challenging, as it is an equilibrium object with nonlinear transition probabilities that depend on the current value of the process: due to the dependence of equilibrium actions on the current value of $\boldsymbol{\lambda}$, the transition probabilities $\psi(a, \sigma|\omega, \boldsymbol{\lambda})$ and transition values $\frac{\hat{\psi}_i(a,\sigma|R,\boldsymbol{\lambda})}{\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda})}$ also depend on $\boldsymbol{\lambda}$. This presents a technical challenge, as the process fails to satisfy standard conditions from the existing literature on Markov chains.

Example 1 (Partisan Bias, cont.)**.** *Return to the example introduced in Section 2.2. An agent of type θ_i who has likelihood ratio λ and observes private signal s updates to belief $\frac{p_i(\lambda,s)}{1-p_i(\lambda,s)} = \lambda\left(\frac{s^{\nu}}{1-s^{\nu}}\right)$. It chooses action L if this belief is less than one, which is equivalent to $s < (1/(1+\lambda))^{1/\nu} = \bar{s}_{i,1}(\lambda)$. At likelihood ratio λ_1, type θ_1 chooses L with probability $F^{\omega}((1/(1+\lambda_1))^{1/\nu})$. However, it believes that it chooses action L with probability $\hat{F}_1^{\omega}((1/(1+\lambda_1))^{1/\nu}) = F^{\omega}(1/(1+\lambda_1))$. Type θ_2 is autarkic. Therefore, its likelihood ratio is constant at $\lambda_2 = 1$ and it chooses action L with probability $F^{\omega}(.5^{1/\nu})$. But θ_1 believes it chooses L with probability $\hat{F}_1^{\omega}(.5^{1/\nu}) = F^{\omega}(.5)$. This implies that θ_1 overestimates the frequency of action L, since $F^{\omega}(x) > F^{\omega}(x^{1/\nu})$.*

4 Asymptotic Learning

We study the asymptotic learning outcomes – the long-run beliefs about the state – for sociable types. Autarkic and noise types do not learn from the history; therefore, their beliefs following the history are constant across time and their behavior is stationary.

4.1 Asymptotic Learning Outcomes

Without loss of generality, we define asymptotic learning outcomes relative to state L. Let *correct* learning (for type θ_i) denote the event where $\boldsymbol{\lambda}_t \to 0^k$ $(\lambda_{i,t} \to 0)$, *incorrect* learning (for type θ_i) denote the event where $\boldsymbol{\lambda}_t \to \infty^k$ $(\lambda_{i,t} \to \infty)$, and *cyclical* learning (for type θ_i) denote the event where $\boldsymbol{\lambda}_t$ $(\lambda_{i,t})$ does not converge. Learning is *complete* if correct learning occurs almost surely. Agents asymptotically *agree* when all sociable types have the same limit beliefs, $\boldsymbol{\lambda}_t \to \{0^k, \infty^k\}$, and agents *disagree* when some sociable types have incorrect learning and others have correct learning, $\boldsymbol{\lambda}_t \to \{0, \infty\}^k \setminus \{0^k, \infty^k\}$. Learning is *mixed* if some sociable types have correct or incorrect learning and others have cyclical learning, while learning is *stationary* if beliefs converge for all sociable types.[19]

4.2 Asymptotic Learning Characterization.

Our main result characterizes how the limiting behavior of the likelihood ratio depends on two expressions that are straightforward to derive from the primitives of the model, i.e. the type space and the signal distributions. This characterization utilizes results on the stability of nonlinear dynamic systems.[20] First, we characterize the set of stationary beliefs, where the likelihood ratio remains constant. Next, we characterize the behavior of the likelihood ratio when it is in a neighborhood of a stationary belief. Building on techniques developed in Smith and Sorensen (2000), we establish necessary and sufficient conditions for the likelihood ratio to converge to this stationary belief with positive probability, which we refer to as local stability. Third, we determine when the likelihood ratio converges to a locally stable belief with positive probability from any initial belief, which we refer to as global stability. This ensures that our characterization holds independent of the prior belief. Finally, we use these results to determine when the likelihood ratio almost surely converges.

This approach builds on techniques used in Bohren (2016) to characterize asymptotic learning outcomes when there is a single type with a misspecified model of the share of autarkic types. Our key technical innovations are to allow for multiple types, which leads to a vector of likelihood ratios, and to characterize conditions for disagreement. Relative to Bohren (2016), establishing the global stability of disagreement outcomes and belief convergence with multiple types requires novel and different techniques.

It what follows, we summarize the four main steps to derive the characterization formally stated in Theorem 1. See Appendix A.1.1 for a detailed outline of the key lemmas in the proof. To simplify notation, in this section we assume that there are $k \leq 2$ sociable types; we present an analogous derivation for $k > 2$ sociable types in Online Appendix E.2.

[19]We use the term disagreement to refer to the case in which beliefs converge to different limit beliefs. Agents' beliefs will also differ when beliefs do not converge or converge for some types but not others. We do not define *incomplete* learning, where $\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}$ for any $\boldsymbol{\lambda} \notin \{0, \infty\}^k$, as this does not occur in our framework (we show in Lemmas 1 and 2 that Assumption 3 rules out incomplete learning).

[20]In correctly specified models, the likelihood ratio is a martingale, and the Martingale Convergence Theorem can be used to rule out cyclical and incorrect learning. This is not the case in a misspecified model. With even the slightest misspecification, the likelihood ratio is no longer a martingale, as any perturbation of a correctly specified model breaks the equality condition. Therefore, an alternative approach is necessary.

Stationary Beliefs. At a stationary belief, the likelihood ratio remains constant for any action and signal pair that occurs with positive probability.

Definition 1 (Stationary). *$\boldsymbol{\lambda}^* \in [0,\infty]^k$ is* stationary *if for all $(a,\sigma) \in \mathcal{A} \times \Sigma$, either (i) $\psi(a,\sigma|\omega,\boldsymbol{\lambda}^*) = 0$ or (ii) $\boldsymbol{\lambda}^* = \boldsymbol{\lambda}^* \left(\frac{\hat{\psi}_i(a,\sigma|R,\boldsymbol{\lambda}^*)}{\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda}^*)}\right)$ for all $\theta_i \in \Theta_S$.*

Aligned signals and preferences (Assumption 1 and 2) rule out confounded learning, while autarkic types or public signals (Assumption 3) rule out informational herds.[21] These assumptions ensure that actions and/or public signals are informative at any interior belief. Therefore, the set of stationary beliefs correspond to each type placing probability one on either state L ($\lambda = 0$) or state R ($\lambda = \infty$) (Lemma 1 in Appendix A.1). Further, the likelihood ratio almost surely does not converge to non-stationary beliefs (Lemma 2 in Appendix A.1). Therefore, the set of stationary beliefs are the candidate limit points of the likelihood ratio.

Local Stability. A learning outcome is *locally stable* if the likelihood ratio converges to this limit belief with positive probability, from a neighborhood of the belief.

Definition 2 (Local Stability). *$\boldsymbol{\lambda}^*$ is* locally stable *if there exists an $\varepsilon > 0$ and neighborhood $B_\varepsilon(\boldsymbol{\lambda}^*)$ such that $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^* | \boldsymbol{\lambda}_1 \in B_\varepsilon(\boldsymbol{\lambda}^*)) > 0$.*

The first expression for the characterization, the expected change in the log likelihood ratio, determines whether a learning outcome is *locally stable*. For type θ_i, the expected change in the log likelihood ratio at belief $\boldsymbol{\lambda}$ depends on the subjective and true probability of each action,

$$\gamma_i(\boldsymbol{\lambda},\omega) \equiv \sum_{(a,\sigma)\in\mathcal{A}\times\Sigma} \psi(a,\sigma|\omega,\boldsymbol{\lambda}) \log\left(\frac{\hat{\psi}_i(a,\sigma|R,\boldsymbol{\lambda})}{\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda})}\right). \tag{7}$$

Equation (7) has a natural interpretation. Suppose the true state is L and fix a belief $\boldsymbol{\lambda}$. Then $\gamma_i(\boldsymbol{\lambda}, L)$ is the difference between (i) the Kullback-Leibler divergence from type θ_i's subjective model in state L, $\hat{\psi}_i(\cdot|L,\boldsymbol{\lambda})$ to the true model in state L, $\psi(\cdot|L,\boldsymbol{\lambda})$ and (ii) the Kullback-Leibler divergence from θ_i's subjective model in state R, $\hat{\psi}_i(\cdot|R,\boldsymbol{\lambda})$, to the true model in state L, $\psi(\cdot|L,\boldsymbol{\lambda})$. At a given belief $\boldsymbol{\lambda}$, if θ_i's subjective model in state R is closer to the true model than θ_i's subjective model in state L, then this difference is positive, $\gamma_i(\boldsymbol{\lambda}, L) > 0$, and $\log \lambda_i$ moves towards state R in expectation. Otherwise, $\log \lambda_i$ moves towards state L in expectation.

[21] Confounded learning corresponds to convergence to an interior belief at which actions are uninformative even though each type acts based on private information, while informational herding corresponds to convergence to an interior belief at which all types cease to incorporate private information into their action choices (Smith and Sorensen 2000).

The sign of each component of $\boldsymbol{\gamma}(\boldsymbol{\lambda},\omega) \equiv (\gamma_i(\boldsymbol{\lambda},\omega))_{i=1}^k$ determines local stability. Let

$$\Lambda_i(\omega) \equiv \{\boldsymbol{\lambda} \in \{0,\infty\}^k | \gamma_i(\boldsymbol{\lambda},\omega) < 0 \text{ if } \lambda_i = 0 \text{ and } \gamma_i(\boldsymbol{\lambda},\omega) > 0 \text{ if } \lambda_i = \infty\} \tag{8}$$

denote the set of learning outcomes in which the expected change in the log likelihood ratio decreases if $\lambda_i = 0$ and increases if $\lambda_i = \infty$, and let

$$\Lambda(\omega) \equiv \cap_{i=1}^k \Lambda_i(\omega) \tag{9}$$

denote the set that satisfies this property for both sociable types. We restrict attention to misspecified models in which $\gamma_i(\boldsymbol{\lambda},\omega) \neq 0$ for all $\theta_i \in \Theta_S$, $\boldsymbol{\lambda} \in \{0,\infty\}^k$ and $\omega \in \{L,R\}$. This set of misspecified models is *generic* in the set of models that satisfy Assumption 1 - 4.[22]

Building on the results on the local stability of nonlinear stochastic difference equations developed in Appendix C of Smith and Sorensen (2000), we show that a learning outcome $\boldsymbol{\lambda}^*$ in a generic misspecified model is locally stable in state ω if and only if $\boldsymbol{\lambda}^* \in \Lambda(\omega)$ (Lemma 3 in Appendix A.1). In other words, if $\langle \boldsymbol{\lambda}_t \rangle_{t=1}^\infty$ converges for all sociable types, then it must converge to a limit random variable whose support lies in $\Lambda(\omega)$. Intuitively, in order for the likelihood ratio to converge to a given learning outcome with positive probability, in expectation, the log likelihood ratio must move towards this learning outcome from nearby beliefs. This also implies that if $\Lambda(\omega)$ is empty in a generic misspecified model, then almost surely at least one type has cyclical learning. This result significantly simplifies the set of possible limit beliefs. It is straightforward to compute $\Lambda(\omega)$ from the primitives of the model.

Example 1 (Partisan Bias, cont.)**.** *From the characterization of action choices in Section 3,*

$$\begin{aligned}\gamma_1(0,L) &= \underbrace{(\pi(\theta_1) + \pi(\theta_2)F^L(.5^{1/\nu}))\log\frac{\pi(\theta_1)+\pi(\theta_2)F^R(.5)}{\pi(\theta_1)+\pi(\theta_2)F^L(.5)}}_{L-action} \\ &\quad \underbrace{+\,\pi(\theta_2)(1-F^L(.5^{1/\nu}))\log\frac{1-F^R(.5)}{1-F^L(.5)}}_{R-action}\end{aligned}$$

[22] If $\gamma_i(\boldsymbol{\lambda},\omega) = 0$ for $\boldsymbol{\lambda} \in \{0,\infty\}^k$, the stability of $\boldsymbol{\lambda}$ also depends on $\gamma_i(\cdot,\omega)$ in a neighborhood of $\boldsymbol{\lambda}$. We do not consider these non-generic cases, as they significantly complicate the analysis without adding much economic insight.

and

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

in state L. When there is no bias, $\nu = 1$, both expressions are negative and $\Lambda(L) = \{0\}$. As ν decreases, R actions occur more frequently and both expressions increase. When ν is sufficiently small, both expressions are positive and $\Lambda(L) = \{\infty\}$. For intermediate values of ν, $\gamma_1(0, L) > 0$ and $\gamma_1(\infty, L) < 0$, and therefore, $\Lambda(L) = \emptyset$. See Appendix B.1 for this derivation.

Global Stability. We are interested in a characterization of asymptotic learning that is independent of the initial belief. Therefore, we need a stronger notion of stability than local stability. We say a learning outcome is *globally stable* if the likelihood ratio converges to this outcome with positive probability, from *any* initial belief.

Definition 3 (Global Stability). $\boldsymbol{\lambda}^*$ *is* globally stable *if for any initial belief* $\boldsymbol{\lambda}_1 \in (0,\infty)^k$, $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) > 0$.

For an *agreement* outcome, local stability is necessary and sufficient for global stability (Lemma 4 in Appendix A.1). Aligned signals and preferences (Assumption 1 and 2) guarantee that there exist signal and action pairs that move the beliefs of all types in the same direction. Therefore, we can construct sequences of actions and signals that occur with positive probability and move the beliefs of all sociable types to a neighborhood of an agreement outcome. Given this, computing $\Lambda(\omega)$ is the only calculation necessary to determine whether correct or incorrect learning occurs with positive probability in state ω. In a generic misspecified model, these learning outcomes occur with positive probability if and only if the corresponding limit beliefs are in $\Lambda(\omega)$. Note that all stationary learning outcomes are agreement outcomes for the case of $k = 1$.

Global stability does not immediately follow from local stability for *disagreement* outcomes, as it is not always possible to construct a sequence of action and public signal realizations that push the likelihood ratio arbitrarily close to the disagreement outcome. For example, if two types are sufficiently similar, then disagreement may not be possible starting from a common prior, even if it arises when the types' beliefs are far apart. Therefore, while a failure of local stability is sufficient to ensure that a disagreement outcome almost surely does not occur, local stability does not guarantee that the outcome occurs with positive probability.

For a disagreement outcome to be globally stable, it must be possible to separate the beliefs for the type converging to $\lambda_i = 0$ and the type converging to $\lambda_i = \infty$, starting from any initial belief. The second part of our learning characterization is a sufficient condition to separate beliefs using the maximal action and public signal in each state, (a_1, σ_L) and (a_M, σ_R). We first define a partial order on how types update their beliefs following maximal actions and signals. Note that (a_1, σ_L) decrease the likelihood ratio, while (a_M, σ_R) increase the likelihood ratio.

Definition 4 (Maximal R-Order)**.** *The* maximal R-order $\succeq_{\boldsymbol{\lambda}}$ *over* Θ *at likelihood ratio* $\boldsymbol{\lambda}$ *is defined by* $\theta_i \succeq_{\boldsymbol{\lambda}} \theta_j$ *iff* θ_i *interprets both maximal action and public signal pairs as stronger evidence of state* R *than* θ_j*,*

$$\frac{\hat{\psi}_j(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_j(a, \sigma | L, \boldsymbol{\lambda})} \leq \frac{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda})} \tag{10}$$

for $(a, \sigma) \in \{(a_1, \sigma_L), (a_M, \sigma_R)\}$*. Define the corresponding strict order* $\succ_{\boldsymbol{\lambda}}$ *if* (10) *holds with strict inequality for either* (a_1, σ_L) *or* (a_M, σ_R)*.*

We use the maximal R-order to establish conditions under which a neighborhood of a disagreement outcome is reached with positive probability from a neighborhood of an agreement outcome. Suppose $\theta_2 \succ_{(0,0)} \theta_1$. Then in a neighborhood of $(0, 0)$, we can construct a finite sequence of maximal actions and signals that decrease θ_1's beliefs and increase θ_2's beliefs. Such a sequence occurs with positive probability, since it is finite. Therefore, $\theta_2 \succ_{(0,0)} \theta_1$ is sufficient to separate beliefs and move them to a neighborhood of $(0, \infty)$. The intuition is analogous for $\theta_2 \succ_{(\infty,\infty)} \theta_1$ to move beliefs to a neighborhood of $(0, \infty)$, and similarly, $\theta_1 \succ_{(0,0)} \theta_2$ or $\theta_1 \succ_{(\infty,\infty)} \theta_2$ to move beliefs to a neighborhood of $(\infty, 0)$. The following definition formalizes this notion of *maximal accessibility.*

Definition 5 (Maximal Accessibility ($k = 2$))**.** *Disagreement outcome* $(0, \infty)$ *is* maximally accessible *if* $\theta_2 \succ_{(0,0)} \theta_1$ *or* $\theta_2 \succ_{(\infty,\infty)} \theta_1$*, and disagreement outcome* $(\infty, 0)$ *is* maximally accessible *if* $\theta_1 \succ_{(0,0)} \theta_2$ *or* $\theta_1 \succ_{(\infty,\infty)} \theta_2$*.*[23]

As discussed above, the likelihood ratio enters a neighborhood of each agreement outcome with positive probability. Maximal accessibility ensures that the likelihood ratio reaches a neighborhood of the disagreement outcome with positive probability from the neighborhood of an agreement outcome, and local stability establishes convergence from the neighborhood of the disagreement outcome. Therefore, maximal accessibility is a sufficient condition for the global stability of a locally stable disagreement outcome

[23]In Online Appendix E.2, we define an analogous notion of maximal accessibility for $k > 2$ sociable types.

(see Lemma 5 in Appendix A.1).[24] It is straightforward to verify by evaluating (10) at action and signal profiles (a_1, σ_L) and (a_M, σ_R) and either beliefs $(0,0)$ or (∞, ∞). See Section 5.5 for an application that uses maximal accessibility.

Mixed Learning. Next we consider the behavior of the likelihood ratio in the neighborhood of a mixed learning outcome. Consider the outcome in which type θ_1 has correct learning, $\omega = L$ and $\lambda_1^* = 0$, and type θ_2 has cyclical learning. This outcome will almost surely not arise if at $\lambda_1^* = 0$, it is possible for the beliefs of θ_2 to converge, i.e. either $(0,0)$ or $(0,\infty)$ is locally stable for θ_2. Intuitively, if $\langle \lambda_{2,t} \rangle$ converges with positive probability when $\lambda_1^* = 0$, then almost surely $\langle \lambda_{2,t} \rangle$ cannot oscillate infinitely often. Therefore, in order for this mixed outcome to arise with positive probability, it must be that $(0,0) \notin \Lambda_2(\omega)$ and $(0,\infty) \notin \Lambda_2(\omega)$. This ensures that in a neighborhood of $(0,0)$ or $(0,\infty)$, θ_2's beliefs drift away from this outcome.

Generalizing this intuition, let $\Lambda_M(\omega)$ denote the set of mixed learning outcomes in which there are no locally stable beliefs for the non-convergent type. When $k = 2$, this corresponds to

$$\Lambda_M(\omega) \equiv \{(\lambda_i^*, \theta_i) | \lambda_i^* \in \{0, \infty\}, \forall \lambda_{-i} \in \{0, \infty\}, (\lambda_i^*, \lambda_{-i}) \notin \Lambda_{-i}(\omega), i \in \{1, 2\}\}. \quad (11)$$

Trivially, $\Lambda_M(\omega) = \emptyset$ when $k = 1$. We establish that in a generic misspecified model, if a mixed learning outcome is not in $\Lambda_M(\omega)$, then almost surely it does not occur (Lemma 6 in Appendix A.1).[25] Therefore, if $\Lambda_M(\omega)$ is empty, mixed learning almost surely does not arise. It is straightforward to compute $\Lambda_M(\omega)$ from $\Lambda_i(\omega)$.[26] In Section 5, we show that $\Lambda_M(\omega)$ is empty and mixed learning outcomes almost surely do not arise for three commonly studied forms of model misspecification.

Belief Convergence. Finally, we establish that the likelihood ratio converges almost surely for all sociable types when there is at least one globally stable outcome and no locally stable mixed outcomes (Lemma 7 in Appendix A.1).

Theorem 1 combines these steps to characterize how the set of asymptotic learning

---

[24] An alternative sufficient condition for the global stability of $(0, \infty)$ is $(0,0) \in \Lambda_1(\omega) \setminus \Lambda_2(\omega)$ or $(\infty, \infty) \in \Lambda_2(\omega) \setminus \Lambda_1(\omega)$ (i.e. $\gamma_1(\boldsymbol{\lambda}, \omega) < 0$ and $\gamma_2(\boldsymbol{\lambda}, \omega) > 0$ for either agreement outcome $\boldsymbol{\lambda} \in \{(0,0), (\infty, \infty)\}$). This condition can be directly verified from the local stability construction, but it will not be satisfied in applications in which both agreement outcomes are locally stable. An analogous condition holds for $(\infty, 0)$.

[25] When $k > 2$, an analogous condition rules out mixed learning outcomes in which a single type has cyclical learning. We also need to rule out mixed learning outcomes in which more than one type has cyclical learning. This requires joint conditions on $\Lambda_i(\omega)$ for the non-convergent types. See Online Appendix E.2.

[26] An easy sufficient (but not necessary) condition for $\Lambda_M(\omega)$ to be empty is if both agreement outcomes or both disagreement outcomes are locally stable.

outcomes in each state depends on $\Lambda(\omega)$ and $\Lambda_M(\omega)$. The proof follows directly from Lemmas 1 to 7 described above, which are formally stated and proved in Appendix A.1.

Theorem 1. *Consider a generic misspecified model with $k \leq 2$ sociable types that satisfies Assumption 1, 2, 3 and 4. When $\omega = L$:*

1. **Agreement.** Correct *learning occurs with positive probability iff $0^k \in \Lambda(L)$ and* incorrect *learning occurs with positive probability iff $\infty^k \in \Lambda(L)$.*

2. **Disagreement.** *Sociable types* disagree *with positive probability if $\Lambda(L)$ contains a maximally accessible disagreement outcome, and sociable types almost surely do not* disagree *if $\Lambda(L)$ contains no disagreement outcomes. Each maximally accessible* disagreement *outcome in $\Lambda(L)$ occurs with positive probability.*

3. **Cyclical Learning.** Cyclical *learning occurs almost surely for all sociable types if $\Lambda(L)$ and $\Lambda_M(L)$ are empty, and cyclical learning occurs almost surely for at least one sociable type if $\Lambda(L)$ is empty. Cyclical learning almost surely does not occur for any sociable type if $\Lambda(L)$ contains an agreement outcome or maximally accessible disagreement outcome and $\Lambda_M(L)$ is empty.*

An analogous result holds for $\omega = R$.[27]

The conditions for correct and incorrect learning are tight: these learning outcomes arise if and only if the respective limit beliefs are in $\Lambda(\omega)$. For disagreement outcomes, we establish a sufficient condition for the outcome to occur (maximal accessibility) and a sufficient condition for outcome not to occur ($\Lambda(\omega)$ empty). In many applications, all locally stable disagreement outcomes are maximally accessible (see Section 5.5 for an illustration). Therefore, there is no wedge between the sufficient conditions for disagreement to occur and not to occur – a disagreement outcome arises if and only if it is in $\Lambda(\omega)$. However, this is not always the case. When a disagreement outcome is locally stable but not maximally accessible, whether disagreement arises can depend on initial beliefs.

An important feature of Theorem 1 is that the characterization requires calculations at a *finite* set of beliefs. Since action choices depend on beliefs, ψ and $\hat{\psi}_i$ vary with $\boldsymbol{\lambda} \in [0, \infty]^k$. Therefore, in principal, determining the asymptotic properties of the likelihood ratio could require characterizing its behavior across the infinite belief space. However, Theorem 1 establishes that this is not the case. Deriving $\Lambda(\omega)$ and $\Lambda_M(\omega)$ and verifying maximal accessibility only require calculations at the finite set of beliefs

[27] The characterization for more than two sociable types is identical, using the modified definitions of maximal accessibility and $\Lambda_M(\omega)$. See Online Appendix E.2.

$\{0, \infty\}^k$. Thus, Theorem 1 significantly simplifies the characterization of asymptotic beliefs.

Theorem 1 rules out mixed learning when $\Lambda_M(\omega)$ is empty, but does not characterize whether mixed learning arises when $\Lambda_M(\omega)$ is not empty. Mixed learning presents a challenge, as we need to consider the movement of the convergent type's likelihood ratio across all possible beliefs for the non-convergent type to determine whether mixed learning arises with positive probability (in contrast to Theorem 1, where we can restrict attention to stationary beliefs for both types). Theorem 4 in Online Appendix E.3 characterizes sufficient conditions for mixed learning to occur with positive probability.

Example 1 (Partisan Bias, cont.). *Applying Theorem 1 to the characterization of $\Lambda(L)$ above establishes that correct learning occurs a.s. for mild partisan bias (ν high), beliefs fail to converge for intermediate levels, and incorrect learning occurs a.s. for severe partisan bias (ν low). Trivially, mixed learning and disagreement cannot arise when there is a single sociable type. See Proposition 7 in Appendix B.1 for a formal statement and proof of this result.*

Example 2. *To illustrate how Theorem 1 applies to a setting with two sociable types, suppose that there is a partisan type, as in Example 1, and a non-partisan type that has a correctly specified signal distribution. Neither type is aware of others' biases, i.e. the non-partisan type believes that others also have a correctly specified signal distribution and the partisan type believes other have the same misspecified signal distribution. We can use Theorem 1 to show that the learning outcomes are similar to Example 1: for both types, correct learning occurs a.s. for a low share of the partisan type or mild partisan bias, beliefs fail to converge for intermediate levels, and incorrect learning occurs a.s. for a high share of a severely biased partisan type. Even though the non-partisan type correctly interprets signals, its failure to account for the partisan type has an equally severe impact on learning.*

Fig. 1 illustrates these three learning regions as a function of the level of bias ν and share of non-partisan types q. See Appendix B.2 for a formal description and analysis of this variation.

Complete Learning. An immediate consequence of Theorem 1 is that in a correctly specified model, learning is *complete* – correct learning occurs almost surely. To see this, note that when all types are correctly specified, the likelihood ratio is a martingale in state L. Due to the concavity of the log function, this means that the expected change in each type's log likelihood ratio is negative at all beliefs, $\gamma_i(\boldsymbol{\lambda}, L) < 0$ for all $\boldsymbol{\lambda} \in [0, \infty]^k$. Therefore, 0^k is the unique locally stable belief and mixed learning does not arise, i.e.

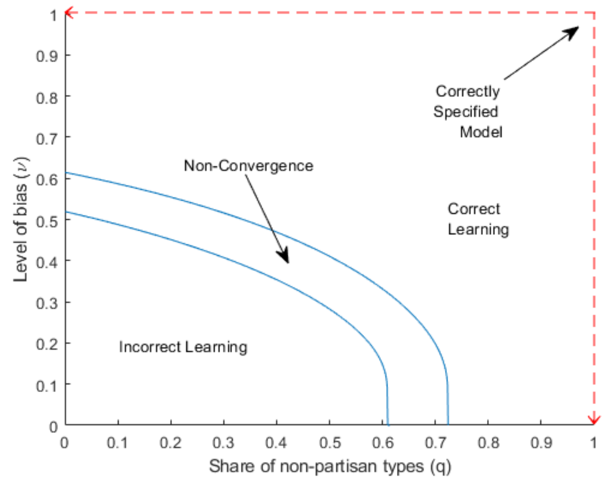

FIGURE 1. Learning Outcomes in Example 2
($\omega = L$, $F^L(s) = 2s - s^2$, $F^R(s) = s^2$, 0.1 share autarkic types)

$\Lambda(L) = \{0^k\}$ and $\Lambda_M(L)$ is empty.[28]

But from Theorem 1, $\gamma_i(\boldsymbol{\lambda}, L) < 0$ at all interior beliefs is not necessary for complete learning: if $\gamma_i(\boldsymbol{\lambda}, L) < 0$ at beliefs $\boldsymbol{\lambda} \in \{0, \infty\}^k$, then learning is complete. Further, complete learning may obtain even if $\gamma_i(\boldsymbol{\lambda}, L) > 0$ for some types at some $\boldsymbol{\lambda} \in \{0, \infty\}^k$. Therefore, Theorem 1 provides much weaker conditions for complete learning.

Corollary 1. *Complete learning obtains in state L if $\Lambda(L) = \{0^k\}$ and $\Lambda_M(L)$ is empty. An analogous result holds for state R.*

These weaker conditions are important for establishing the robustness of complete learning in misspecified models, as with even an arbitrarily small amount of misspecification, the likelihood ratio is no longer a martingale. Section 4.3 uses Corollary 1 to derive several robustness results.

Action Convergence. Belief convergence forces action convergence: each type eventually settles on an action if and only if its beliefs converge. The limit action choice is efficient if learning is correct, and otherwise is inefficient. If learning is cyclical for a type, then that type will choose both efficient and inefficient actions infinitely often.[29]

[28] A similar result holds for all correctly specified types when some types are misspecified. A correctly specified type accurately accounts for other types' misspecification and is able to probabilistically parse out the information conveyed by actions. Therefore, misspecified types do not interfere with the learning of the correctly specified types. See Theorem 3 in Online Appendix E.1.

[29] In the proof of Theorem 1, we show that if the likelihood ratio for a type does not converge, then it enters a neighborhood of each certain belief infinitely often. Therefore, the type chooses both efficient and inefficient actions infinitely often.

4.3 Robustness of Complete Learning.

The next result establishes that correctly specified models are robust to some misspecification, in that learning is complete when sociable types have approximately correct models. This may not seem surprising, since Bayes rule is continuous. But in an infinite horizon setting, nearby models with small per-period differences in belief updating can aggregate to very different limit beliefs. When this is the case, then even arbitrarily small departures from the correctly specified model will interfere with learning. In principle, this would substantially limit the applicability of rational learning models to real-world settings. Theorem 2 establishes that this does not occur in the environment we consider in this paper; complete learning obtains for any form of misspecification in which each sociable type's subjective model is close enough to the true model.

Theorem 2. *Given a generic misspecified model that satisfies Assumption 1, 2, 3 and 4, there exists a $\delta > 0$ such that if $|\hat{\psi}_i(a, \sigma|\omega, \boldsymbol{\lambda}) - \psi(a, \sigma|\omega, \boldsymbol{\lambda})| < \delta$ at $\boldsymbol{\lambda} \in \{0, \infty\}^k$ for all $(a, \sigma) \in \mathcal{A} \times \Sigma$ and $\theta_i \in \Theta_S$, then learning is complete in state ω.*

Robustness follows from the continuity of $\boldsymbol{\gamma}(\boldsymbol{\lambda}, \omega)$ in each type's subjective signal and type distributions. In any correctly-specified model, $\Lambda(L) = \{0^k\}$ and $\Lambda_M(L) = \emptyset$. By continuity, these sets don't change when some misspecification is introduced. The tools in this paper allow for a precise characterization of exactly how robust complete learning is to different forms of misspecification.

Theorem 2 depends on the subjective and true equilibrium probabilities of each action at stationary beliefs. The following corollary presents sufficient conditions on the primitives of the model for these subjective equilibrium probabilities to be close enough to the true equilibrium probabilities: if all sociable types have subjective type and signal distributions close enough to the true distributions, then learning is complete.

Corollary 2. *Given a generic misspecified model that satisfies Assumption 1, 2, 3 and 4, there exists a $\delta > 0$ such that if $||\hat{\pi}_i - \pi|| < \delta$, $||\hat{F}_i^\omega - F^\omega|| < \delta$ and $||\hat{G}_i^\omega - G^\omega|| < \delta$ for all $\theta_i \in \Theta_S$, then learning is complete in state ω, where $|| \cdot ||$ denotes the supremum metric.*

Corollary 2 is more restrictive than Theorem 2, since Theorem 2 can hold when agents are very wrong about the type distribution, as long as the types that they believe exist are "close" to the actual types. For example, suppose that all agents are type θ_1, but believe that all agents are type $\theta_2 \neq \theta_1$. If types θ_1 and θ_2 have similar preferences and subjective signal distributions, then the conditions for Theorem 2 hold, but the conditions for Corollary 2 do not. Similar robustness results hold for an individual type that has a model close to the correctly specified model, regardless of other types.

Example 1 (Partisan Bias, cont.)**.** *For ν sufficiently close to one, $\Lambda(L) = \{0\}$ and complete learning obtains.*

These results contrasts with Frick et al. (2019), who demonstrate a failure of robustness in a social learning setting with misspecification over the distribution of agents' preferences. Their setting differs from ours in that they consider an infinite state space and do not assume adequate information (i.e. Assumption 3) – a continuum of preference types ensures that actions are informative, but the informational content of these actions can become arbitrarily small.[30] With an infinite state space, action choices are much more sensitive to small amounts of misspecification. They show that no matter how small the misspecification in their setting, the limiting belief does not depend on the underlying state when the state space is infinite, but correct learning is robust to misspecification when the state space is finite. Similarly, under an analogue of our adequate information assumption, correct learning is robust to small amounts of misspecification in their setting, even with an infinite state space. Nevertheless, settings with larger state spaces are more sensitive to misspecification and the size of the state space is important to keep in mind when considering robustness in misspecified models.

4.4 Discussion of Results

When agents learn from the action choices of their peers, model misspecification interacts with the endogenous informativeness of actions to give rise to the possibility of cyclical learning or path-dependent learning (for example, both incorrect and correct learning arise with positive probability).[31] These learning outcomes have important economic implications. Cyclical learning is a failure of beliefs (and actions) to settle down, even after an arbitrarily long period of time. This means that in the long-run, action choices oscillate infinitely often between efficient and inefficient actions. When multiple learning outcomes arise, different paths of signal realizations lead to different long-run beliefs. An initial signal that, for instance, a medical technology is dangerous or a new restaurant is low quality when in fact the opposite is true can lead to this misconception becoming entrenched and beliefs converging to the incorrect state. In contrast, if the initial signal had been positive, agents would have learned the correct state. Path-dependent learning can explain why different populations with similar models can come to have very different entrenched views.

[30]Specifically, herds do not arise in the correctly specified model but Lemma 10 in our paper does not hold.

[31]Belief convergence requires that when agents are almost certain of a state, the action and signal frequencies they observe confirm their model in that state. If agents are "surprised" when they are almost certain of either state, this leads to cyclical learning (i.e. $|\Lambda(\omega)| = 0$). In contrast, if there are multiple learning outcomes near which the action frequencies confirm each type's model at that outcome, then multiple learning outcomes occur with positive probability (i.e. $|\Lambda(\omega)| > 1$).

Cyclical and path-dependent learning can also arise in misspecified active individual learning models, where future information depends on the current action choice of the agent, or in misspecified passive individual learning models, where the true signal distribution is independent of the history but an agent's misspecified signal distribution depends on her current belief. For example, cyclical learning emerges in the active individual learning models of Nyarko (1991); Fudenberg et al. (2017), while path-dependent learning arises in the passive individual learning models of Rabin and Schrag (1999); Epstein et al. (2010).

In contrast, cyclical and path-dependent learning do not arise in misspecified passive individual learning models in which an agent's signal misspecification is independent of her current belief. In our framework, this corresponds to settings in which agents learn from the public signal but not from the actions of others. It follows immediately from Theorem 1 that, generically, there is exactly one locally stable learning outcome, $|\Lambda(\omega)| = 1$, and beliefs almost surely converge to this unique outcome.[32]

Path-dependent learning also occurs in correctly specified social learning settings with informational herds (Bikhchandani, Hirshleifer, and Welch 1992; Banerjee 1992; Smith and Sorensen 2000). In contrast to misspecified settings, all but at most one of these limit beliefs must be non-degenerate (i.e. incomplete learning). This difference is economically important. In correctly specified models, informational herds are fragile (Bikhchandani et al. 1992). Even though all agents are playing the same action, they remain uncertain about the state. Therefore, a herd of any length can be overturned by a relatively uninformative public signal or other piece of new information. In contrast, when an incorrect herd persists in our setting, beliefs almost surely converge to the incorrect state. This implies that longer herds will become increasingly difficult to overturn.[33]

Focus on Asymptotic Learning. When incorrect learning, cyclical learning or disagreement arise asymptotically, then we will observe these learning outcomes even in the face of an infinite amount of information. This establishes that the source of these inefficiencies is not a lack of sufficient information to learn the state. Agents are bounded away from efficiency, *irrespective* of the amount of information that they observe. Of course, these asymptotic results also establish that we will observe inefficient choices, belief cycles and disagreement in finite time.

[32]When $\gamma_i(\boldsymbol{\lambda}, \omega)$ is independent of $\boldsymbol{\lambda}$, its sign is constant across the belief space. Therefore, the conditions in Theorem 1 collapse to the standard result that beliefs converge to the state that minimizes the relative entropy from the misspecified model to the correct model (Berk (1966); Shalizi (2009)). When the sign of $\gamma_i(\boldsymbol{\lambda}, \omega)$ is independent of $\boldsymbol{\lambda}$, there is no wedge between local and global stability for disagreement outcomes: if the unique locally stable outcome is a disagreement outcome, then it is also globally stable. Therefore, it is not necessary to check maximal accessibility.

[33]This observation was first made in Eyster and Rabin (2010) in the context of naive social learning.

The expression $\gamma_i(\boldsymbol{\lambda}, \omega)$ that we use to characterize the locally stable set also determines the asymptotic rate of learning. The larger this term is in magnitude, the faster the rate of convergence to (or, depending on the sign, the faster the rate of divergence from) the candidate limit belief from a neighborhood of this belief.

5 Applications

In this section, we demonstrate how our general framework can be used to gain a deeper understanding of how different types of model misspecification affect learning. First, we explore whether two forms of signal misspecification are *conceptually robust*, in that they are not sensitive to the parametric specifications chosen to model how agents misinterpret information. We show that whether the learning results in these settings are robust to alternative specifications depends on the type of misspecification. Second, we compare whether signal misspecification has a similar impact in individual and social learning settings. Third, we examine whether a representative agent model is a good approximation for a setting with model heterogeneity by comparing the learning outcomes for a set of types with heterogenous levels of model misspecification to the learning outcomes for a single type with the average model of the population. Fourth, we demonstrate that our framework can connect conceptually distinct forms of model misspecification that have similar implications for learning and behavior. Finally, we explore entrenched disagreement in a level-k social learning model.

5.1 The Fragility of Underreaction and Overreaction Specifications

We first study conceptual robustness in a setting in which agents underreact or overreact to signals.[34] We show that this form of misspecification is sensitive to the modeling choice, in that different parametric specifications lead to qualitatively different asymptotic learning results. We also compare individual and social learning settings, and show that the interaction between under- or overreaction and learning from others' actions creates long-run inefficiencies that are not present when agents learn solely from signals. This illustrates that the context of a learning environment – specifically, whether agents learn from private or social sources of information – can influence how a bias impacts long-run behavior.

We model overreaction as an agent who forms beliefs as if she has observed the same signal multiple times. Analogously, we model underreaction as an agent who needs to observe multiple realizations of the same signal to reach the posterior induced by the

[34]Both overreaction and underreaction have been observed empirically (Edward 1982; Tversky and Kahneman 1971). Griffin and Tversky (1992) reconcile these disparate findings by showing that context determines whether an individual overreacts or underreacts to information: in updating experiments, subjects underreact to precise estimates of moderate effect sizes, while subjects overreact to noisy estimates of large effect sizes.

true signal distribution. Formally, if a signal induces posterior belief s that the state is R, then the agent updates her likelihood ratio by $\left(\frac{s}{1-s}\right)^{\nu}$, where $\nu \in [0,1)$ captures underreaction and $\nu \in (1,\infty)$ captures overreaction. For example, if $\nu = 2$, an agent double counts signal s, and if $\nu = 1/2$, an agent needs to observe two realizations of signal s in order to arrive at the true posterior induced by one realization of s.

This modeling approach contrasts with Epstein et al. (2010), who represent underreaction as a subjective posterior belief that is a weighted average of the prior belief and the posterior induced by the true signal distribution and overreaction as a subjective posterior belief that places a negative weight on the prior belief and a weight above one on the true posterior. These different modeling choices are of consequence: in an individual learning setting, our forms of underreaction and overreaction lead to fundamentally different asymptotic learning outcomes than the forms considered in Epstein et al. (2010).

Individual Learning. In an individual learning setting (i.e. informative public signals and uninformative private signals), the signal misspecification introduced above corresponds to $\frac{\hat{\sigma}(\sigma)}{1-\hat{\sigma}(\sigma)} = \left(\frac{\sigma}{1-\sigma}\right)^{\nu}$. Our first result establishes that under- and overreaction do not alter the set of asymptotic learning outcomes in this setting, *regardless* of the severity of the bias.

Proposition 1. *In an individual learning setting, correct learning occurs almost surely for any* $\nu \in (0,\infty)$.

To understand the intuition for this result, consider the case of overreaction. Agents overreact symmetrically to signals that favor state L and state R – that is, equally extreme signals for states L and R in the correctly specified model are interpreted as equally extreme signals in the misspecified model. For example, when $\nu = 2$, the symmetric signals $\sigma_1 = 1/3$ and $\sigma_2 = 2/3$ are interpreted as symmetrically more extreme beliefs $\hat{\sigma}(\sigma_1) = 1/5$ and $\hat{\sigma}(\sigma_2) = 4/5$. Further, agents overreact to weak and strong signals in a similar manner. For example, agents double count both weaker signal $\sigma = 3/5$ and stronger signal $\sigma = 4/5$. These properties ensure that overreaction does not affect the sign of the average change in the log likelihood ratio: as in the correctly specified model, it remains negative at all beliefs, $\gamma(\boldsymbol{\lambda}, \omega) < 0$ for all $\boldsymbol{\lambda}$. Therefore, the set of asymptotic learning outcomes is independent of the level of overreaction.[35]

In contrast, Epstein et al. (2010) find that sufficiently severe overreaction leads to

[35] More generally, this result holds for any signal misspecification such that the functions $\sigma/(1-\sigma) \mapsto \hat{\sigma}(\sigma)/(1-\hat{\sigma}(\sigma))$ and $(1-\sigma)/\sigma \mapsto (1-\hat{\sigma}(\sigma))/\hat{\sigma}(\sigma)$ are log-concave. Heidhues et al. (2019) find a similar result for conservatism and base rate neglect in an individual learning model with a continuous state space and a normal prior and noise. Their result also relies on choosing a symmetric parameterization for the misspecification.

the possibility of incorrect learning. In their specification, overreaction alters updating in an asymmetric way. An agent with prior p who overreacts to signal σ that leads to true posterior $\frac{p\sigma}{p\sigma+(1-p)(1-\sigma)}$ updates to posterior belief

$$\tilde{\nu}\left(\frac{p\sigma}{p\sigma+(1-p)(1-\sigma)}\right)+(1-\tilde{\nu})p,$$

where $\tilde{\nu}>1$ is the overreaction parameter. This learning rule is represented by a form of signal misspecification that asymmetrically skews signals: when an agent believes that state R is more likely, she overweights signals in favor of state R more than she overweights signals in favor of state L, and vice versa. For example, when $\tilde{\nu}=1.2$ and $p\approx 1$, $\hat{\sigma}(.25)=.23$ while $\hat{\sigma}(.75)=.83>1-\hat{\sigma}(.25)$. Further, the misspecification depends on the prior: the asymmetry is more pronounced for more extreme priors. Given this asymmetry, sufficiently severe overreaction changes the sign of $\gamma(\boldsymbol{\lambda},\omega)$, and therefore, the set of asymptotic learning outcomes. See Appendix C.1 for the representation of this form of overreaction in our framework.

Relatedly, Bushong and Gagnon-Bartsch (2017) study a setting where an agent is misspecified about her past preferences. The agent has reference dependent utility and underestimates the extent of her reference dependence when recalling past outcomes. This causes her to overreact to past gains and losses. Given that she overreacts symmetrically, as in our setting, she accurately learns the average return. Only if she is also loss averse will she overreact more to losses than gains. As in Epstein et al. (2010), this asymmetry leads to the possibility incorrect learning.

Although complete learning still obtains when under- or overreaction is symmetric, these forms of signal misspecification do impact short-run learning and the volatility of beliefs. An agent who overreacts to signals has beliefs that are more volatile and converge faster than expected, while an agent who underreacts has beliefs that are less volatile and converge slower than expected.

Social Learning. Now consider a social learning setting with informative private signals and uninformative public signals, in which all agents believe private signals are distributed according to $\hat{F}^{\omega}(s)=F^{\omega}(\frac{s^{\nu}}{(1-s)^{\nu}+s^{\nu}})$ in state ω. These misspecified signal distributions lead to the form of under- or overreaction introduced above, i.e. $\frac{\hat{s}(s)}{1-\hat{s}(s)}=\left(\frac{s}{1-s}\right)^{\nu}$. In contrast to individual learning, social learning interacts with under- and overreaction in a way that can interfere with asymptotic learning. Although the signal misspecification is symmetric, these signals are filtered through other agents' action choices in a way that gives rise to asymmetric under- or overreaction to actions. Similar to an asymmetric under- or overreaction to signals, this asymmetric under- or overreaction to actions can alter the set of asymptotic learning outcomes.

To illustrate this possibility, consider a decision problem with four actions, $a \in \{a_1, a_2, a_3, a_4\}$, ordered according to increasing preference in state R. Assume preferences are symmetric: if a_1 is optimal at belief p, then a_4 is optimal at belief $1-p$, and similarly for a_2 and a_3. There are two types of agents: θ_1 is sociable and θ_2 is autarkic. Both types interpret signals in the same way and have the same preferences and type θ_1 has correct beliefs about the share of autarkic types.

Suppose sociable agents are herding on action a_1. Then a_1 actions that confirm the herd are less informative than a_4 actions that contradict the herd. Therefore, overreaction is stronger with respect to the contradictory action a_4. This pulls beliefs away from state L. Similarly, when sociable agents are herding on a_4, they overreact more to contradictory action a_1 than confirmatory action a_4, pulling beliefs away from state R. For sufficiently severe overreaction, this gives rise to cyclical learning. In contrast, when agents underreact to new information, they do not learn enough from contradictory actions. This makes it more difficult to break a herd, and gives rise to the possibility of incorrect learning. Fig. 2 illustrates these learning regions.

Note that in the asymmetric signal misspecification in Epstein et al. (2010), incorrect learning arises when agents overreact to signals, whereas in the asymmetric action misspecification that arises from social learning, incorrect learning arises when agents underreact to signals.

5.2 The Robustness of Confirmation Bias Specifications

We next study conceptual robustness a setting in which an agent misinterprets information to confirm her current beliefs – *confirmation bias*. We use our characterization to show that the asymptotic learning insights derived in Rabin and Schrag (1999) hold for a broad class of specifications of confirmation bias. Therefore, in contrast to the previous section, this form of signal misspecification is less sensitive to modeling choice.[36]

Rabin and Schrag (1999) study an individual learning model in which an agent receives a binary signal $y \in \{l, r\}$ that matches the state with probability $\beta > 1/2$. If the signal contradicts her prior belief (i.e. her belief favors state L and she observes signal r, or vice versa), then with probability q she misreads the signal as a confirmatory signal (i.e. believes she observed l when she observed r). They establish that for sufficiently high q, confirmation bias leads to incorrect learning with positive probability. This set-up makes several strong assumptions to simplify the analysis. First, the severity of the confirmation bias, i.e. the frequency that a signal is misread, is independent of the current belief: an agent exhibits the same level of bias when she believes that state L is

[36] Fryer, Harms, and Jackson (2018) derive similar insights in a model of ambiguous information with limited memory. The interpretation strategies they study map into a misspecified model that mirrors confirmation bias.

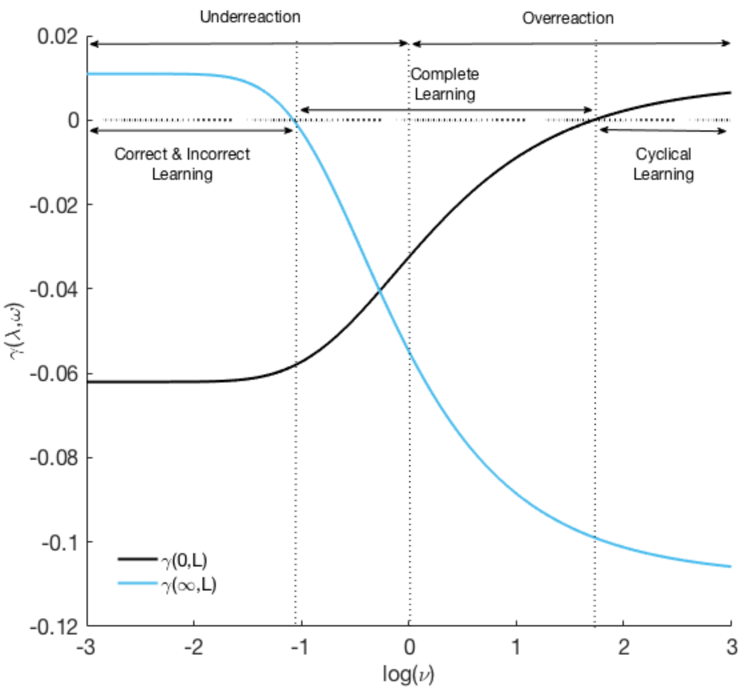

Figure 2. Under/overreaction and social learning
($\omega = L$, $F^L = 2s - s^2$, $F^R = s^2$, $\pi(\theta_1) = .9$, $p_0 = 1/2$, $(p_{i,1}, p_{i,2}, p_{i,3}) = (.25, .5, .75)$ for $i = 1, 2$.)

very likely as when she believes that state L is only slightly more likely than state R. Second, misread contradictory signals are interpreted as having the same informational content as confirmatory signals.

We use our framework to show that the insights gleaned from this stylized form of confirmation bias are robust to more general specifications. Consider the case in which contradictory signals may be viewed as weaker evidence of the more likely state than confirmatory signals and the degree of the slant can depend on the current belief about the state.[37] This corresponds to a form of signal misspecification in which with probability q, a contradictory signal y is interpreted as a weighted average of the true posterior $\sigma(y)$ induced by signal y and the true posterior $\sigma(y')$ induced by the confirmatory signal y'. Formally, when belief $p > 1/2$ that the state is R, with probability q an l signal is slanted towards state R and interpreted as subjective posterior $(1-\nu(p))\sigma(l)+\nu(p)\sigma(r)$ and when $p < 1/2$, with probability q an r signal is slanted towards state L and interpreted as subjective posterior $(1-\nu(p))\sigma(r)+\nu(p)\sigma(l)$, where $\nu : [0,1] \to [0,1]$ captures the slant of the contradictory signal. To simplify exposition, we assume that ν is continuous and symmetric at certainty, $\nu(1) = \nu(0)$. The dependence of $\nu(p)$ on the current

[37]As discussed in Section 2.3, it is a straightforward extension to allow the signal misspecification to depend on the belief about the state (see Online Appendix E.4).

belief about the state allows this slant to vary with the belief about the state. If ν is strictly decreasing on $[0, 1/2]$ and strictly increasing on $[1/2, 1]$, then confirmation bias becomes more severe as the agent's belief becomes more extreme. The model of Rabin and Schrag (1999) corresponds to the case where $\nu(p) = 1$, i.e. contradictory signals are fully slanted, independent of the current belief.

We can apply Theorem 1 to show that the possibility of incorrect learning depends on the probability of misreading signals and the level of slant at certainty, but is independent of how the slant depends on the belief about the state.

Proposition 2. *There exists cut-offs $\overline{q} \in (0, 1)$ and $\overline{\nu}(q) \in (0, 1]$, with $\overline{\nu}(q) < 1$ for $q > \overline{q}$, such that for $q > \overline{q}$ and $\nu(1) > \overline{\nu}(q)$, both incorrect and correct learning arise with positive probability, and for $q < \overline{q}$ or $\nu(1) < \overline{\nu}(q)$, then almost surely learning is correct. If signals are symmetric, $\sigma(r) = 1 - \sigma(l)$, and fully slanted at certainty, $\nu(1) = 1$, then $\overline{q} = 1 - 1/2\sigma(r)$, which is identical to the cut-off when signals are fully slanted across the belief space, $\nu(p) = 1$, as in Rabin and Schrag (1999).*

This establishes that the stark form of confirmation bias in Rabin and Schrag (1999) does not drive the incorrect learning result: when confirmation bias is sufficiently severe, beliefs can become entrenched on both the correct and incorrect state, regardless of its exact form. In fact, when signals are fully slanted at certainty, the cut-off for incorrect learning is identical to that when signals are fully slanted at all beliefs, as in Rabin and Schrag (1999). We obtain a similar result to Proposition 2 when we also allow the probability q of misreading a signal to depend on the belief about the state.

While the possibility of incorrect learning is independent of the shape of $\nu(p)$, the probability of incorrect learning depends crucially on it. Fixing the slant at certainty, an $\nu(p)$ closer to one at interior beliefs amplifies the impact of early signals and makes it more difficult to overturn beliefs that even slightly favor the incorrect state. Fig. 3 illustrates that the probability of incorrect learning is increasing in the severity of the slant using three parameterizations of $\nu(p)$.

5.3 Naive Learning with Model Heterogeneity

A standard assumption in papers that study model misspecification is that all agents have the same misspecified model. This can be viewed as a representative agent approach, and it significantly simplifies the analysis. More realistically, agents will have heterogeneous levels of misspecification. The representative agent approach is valid in the face of heterogeneity if agents process information in a way such that their long-run behavior is approximated by the long-run behavior of the representative agent.

In this section, we use our framework to explore the validity of the representative agent approach when agents exhibit a form of naive learning in which they overestimate

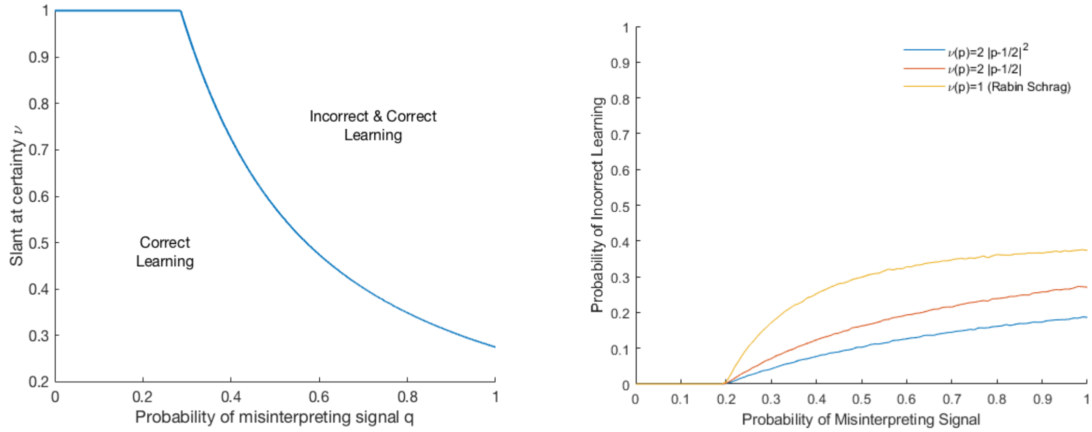

FIGURE 3. Confirmation Bias
$(\sigma(l) = 3/8,\, \sigma(r) = 5/8)$

the private information reflected in actions. Bohren (2016); Eyster and Rabin (2010) study naive learning when all agents have the same misspecified model. We compare learning in a setting where agents have different levels of naivety to a representative agent setting in which a single type has a level of naivety equal to the average naivety of the population. We show that when heterogeneity is small, this representative agent model is a good approximation of the underlying model with heterogeneity. When heterogeneity is large, the models have qualitatively different learning outcomes. Specifically, heterogeneity facilitates learning in that it leads to correct learning for a strictly larger set of parameters than the representative agent model.

As in Bohren (2016), we model naive learning as a misspecified belief about the share of autarkic types. Let θ_A denote the autarkic type and assume $\pi(\theta_A) \in (0, 1)$. To capture model heterogeneity, suppose there are two sociable types, θ_1 and θ_2, that occur with equal probability, $\pi(\theta_1) = \pi(\theta_2)$. Both sociable types overestimate the share of autarkic types, with type θ_2 having a more severe bias, $\pi(\theta_A) < \hat{\pi}_1(\theta_A) \leq \hat{\pi}_2(\theta_A) \leq 1$. This form of misspecification leads agents to underestimate the correlation between prior actions. We compare this setting to a representative agent setting in which a single sociable type believes that the autarkic type occurs with probability $\hat{\pi} \equiv (\hat{\pi}_1(\theta_A)+\hat{\pi}_2(\theta_A))/2$. In other words, the representative agent has a bias equal to the average bias in the heterogenous setting. To close the model, assume that each agent faces a binary decision problem in which she earns a payoff of one from choosing the action that matches the state, $\mathcal{A} = \{L, R\}$ and $u(a, \omega) = \mathbb{1}_{a=\omega}$, all types correctly interpret private signals, sociable types have correct beliefs about the relative frequency of each sociable type, all types have common prior $p_0 = 1/2$, and public signals are uninformative and believed to be uninformative.[38]

We first show that the representative agent model is a good approximation when heterogeneity is sufficiently small.

Proposition 3. *Generically, for any average bias $\hat{\pi}$, there exists an $\varepsilon > 0$ such that if heterogeneity is sufficiently small, $|\hat{\pi}_1(\theta_A) - \hat{\pi}_2(\theta_A)| < \varepsilon$, then the heterogeneous and representative agent settings have the same set of long-run learning outcomes.*

Next, we explore how heterogeneity affects learning. It is a priori unclear whether heterogeneity will facilitate or hinder learning, compared to the representative agent model. The type with milder misspecification may facilitate learning by counteracting the type with more severe misspecification, or the type with the more severe misspecification may distort information in a way that hinders learning for both types. The following result establishes that the first effect dominates and heterogeneity facilitates learning.

Proposition 4. *Suppose the signal distribution is symmetric, $F^L(s) = 1 - F^R(1-s)$. For all $\hat{\pi}_1(\theta_A), \hat{\pi}_2(\theta_A) \in (\pi(\theta_A), 1]$, almost surely learning is either correct or incorrect. If learning is almost surely correct in the representative agent model at $\hat{\pi}$, then learning is almost surely correct in the heterogeneous model for all $\hat{\pi}_1(\theta_A)$ and $\hat{\pi}_2(\theta_A)$ such that $(\hat{\pi}_1(\theta_A) + \hat{\pi}_2(\theta_A))/2 = \hat{\pi}$, and if incorrect learning occurs with positive probability in the heterogeneous model at $\hat{\pi}_1(\theta_A)$ and $\hat{\pi}_2(\theta_A)$, then incorrect learning occurs with positive probability in the representative agent model at $\hat{\pi} \equiv (\hat{\pi}_1(\theta_A) + \hat{\pi}_2(\theta_A))/2$.*

Type θ_1 is more adept at correcting for correlated information, and as a result, asymptotically adopts the inefficient action with lower probability than θ_2. In turn, this helps θ_2 learn the true state. Actions from θ_1 confirm the state and θ_2 overestimates the private information reflected in these actions. This reduces the probability that θ_2 herds on an inefficient action.[39]

Proposition 4 has important implications for policy interventions aimed at mitigating inefficient choices. Suppose a social planner wishes to intervene if and only if agents face the possibility of incorrect learning. The planner measures the average level of bias in the population and uses a representative agent approach to determine whether to intervene. Given Proposition 4, this method will result in *over-intervention*, in that there are levels of bias at which incorrect learning arises in the representative agent model but not in the

[38]The representative agent model is a special case of Bohren (2016).

[39]Heterogeneity does not always improve learning. If heterogeneity leads to fundamentally different biases – for example, if one type overestimates the correlation in prior actions and the other type underestimates it – then sufficient heterogeneity will interfere with long-run learning, even when the average bias is close to the truth (e.g. $\hat{\pi} \approx \pi(\theta_A)$) and learning is complete in the representative agent model.

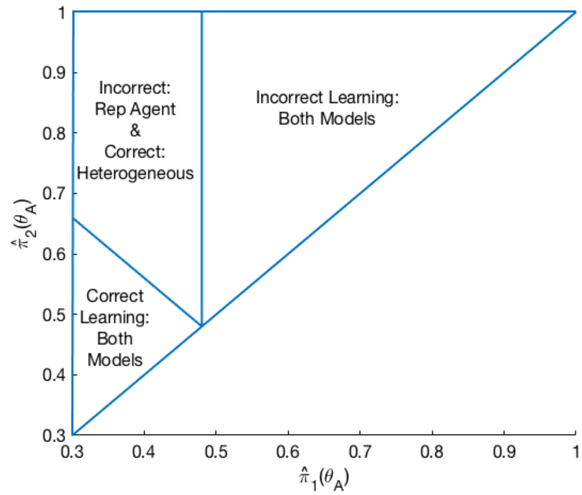

FIGURE 4. Learning in Representative Agent and Heterogeneous Models
($\pi = .3$, $F^L(s) = 2s - s^2$, $F^R(s) = s^2$, $p_0 = 0.5$)

heterogeneous model. However, whenever incorrect learning arises in the heterogeneous model, it also arises in the representative agent model – and therefore, under-intervention will not be an issue. Fig. 4 illustrates these insights.

These results provide insight into the use of the representative agent approach to study how model misspecification impacts learning. The range of heterogeneity plays a key role in determining whether such an approach is appropriate. If heterogeneity is large, then assuming all agents are a single type with the average bias of the population can lead to very different predictions about long-run outcomes. Using these predictions to determine appropriate policy interventions may lead to interventions that are unnecessary or even harmful. In contrast, if the level of heterogeneity is small, then the predictions of the representative agent model are a good approximation. The tools developed in this paper can be used to quickly assess whether a reasonable amount of heterogeneity will lead to substantially different learning predictions.

5.4 Using the Framework to Demonstrate Behavioral Equivalence

We can use our general framework to connect conceptually distinct forms of model misspecification that have similar implications for learning and behavior. To demonstrate this, we show that a model with naive temptation and a model with partisan bias are asymptotically equivalent in terms of the set of learning outcomes that arise.

The temptation setting is as follows. Consider a binary action decision problem in which agents seek to choose the action that matches the state, but differ in whether they are tempted to choose action R when there is uncertainty. The tempted types place higher weight on matching the state in state R, $u_i(a, \omega) = (2 - \beta)\mathbb{1}_{a=\omega=L} + \beta\mathbb{1}_{a=\omega=R}$ for some $\beta \in (1,2)$, while the non-tempted types have symmetric preferences across states, $u_i(a,\omega) = \mathbb{1}_{a=\omega}$. Agents are unaware of others' temptation. We establish an equivalence between this setting and the partisan bias setting in Example 2.

Proposition 5. *For every level of temptation $\beta \in [1,2)$ and share of non-tempted types $q \in [0,1]$, the model in Example 2 with level of partisan bias $\nu(\beta) \equiv \log .5 / \log(1-\beta/2)$ and share of non-partisan types q has an identical set of asymptotic learning outcomes.*

Proposition 5 shows that fundamentally different biases – the temptation model is a form of preference misspecification, while the partisan bias model is a form of signal misspecification – interfere with learning in similar ways. See Appendix C.4 for the full description and analysis of this setting.

5.5 Entrenched Disagreement in a Level-k Learning Model

This application shows that entrenched disagreement emerges as a robust feature of social learning in a level-k model. Level-k models describe how boundedly rational agents draw inference in strategic settings (Costa-Gomes, Crawford, and Iriberri 2009). Agents are characterized by their "depth" of reasoning, where higher levels use progressively more sophisticated reasoning. This misspecified model features prominently in the empirical literature on social learning, but has been relatively unexplored in the corresponding theoretical literature, as characterizing learning outcomes is significantly more complex when agents learn in different ways. Using our framework, each level can be modeled as a type that has a misspecified model of the strategic link between prior actions. When agents have heterogenous models, observing the same information does not ensure that they converge to the same beliefs. In fact, we show that agents can become very certain of different states, despite observing the same patterns of action choices.

Consider a level-k model with four levels. Each level corresponds to a type, $\Theta = \{\theta_0, \theta_1, \theta_2, \theta_3\}$, where the level-0 type is used to anchor the model of level-1.[40] Level-0 chooses an action without learning from signals or the actions of others, i.e. it is a noise type who believes that private signals and prior actions are uninformative, $\hat{s}_0(s) = 1/2$ and $\hat{\pi}_0(\theta_0) = 1$. The level-1, 2 and 3 types accurately learn from private signals but have a misspecified model of how to interpret actions, which is captured by a misspecified type distribution. Level-1 chooses an action solely based on its private signal, i.e. it is an autarkic type who believes that prior actions are uninformative, $\hat{\pi}_1(\theta_0) = 1$. Level-2 believes that prior actions solely reflect private information, i.e. $\hat{\pi}_2(\theta_1) = 1$. It fails to account for redundant information from the prior actions of others.[41] Level-3 is the most

[40] Our framework can allow for higher levels. However, empirical and experimental studies rarely find evidence of levels above level-3.

[41] A level-2 type is analogous to the "BRTNI" agents in Eyster and Rabin (2010) and the "naive

sophisticated type: it understands that prior actions contain redundant information, but does not allow for the possibility that other agents also account for this – it believes that most other agents are level-2, $\hat{\pi}_3(\theta_2) = 1 - \varepsilon$ for some small $\varepsilon > 0$, and for technical reasons, places arbitrarily small probability on the level-1 type, $\hat{\pi}_3(\theta_1) = \varepsilon$.[42] The correctly specified model is not a special case of this set-up for any type, as no type allows for the existence of its own type. To close the model, assume that each agent faces a binary decision problem in which she earns a payoff of one from choosing the action that matches the state, $\mathcal{A} = \{L, R\}$ and $u(a, \omega) = \mathbb{1}_{a=\omega}$, the level-0 type does not actually exist in the population, $\pi(\theta_0) = 0$, the level-1 type occurs with positive probability, $\pi(\theta_1) \in (0, 1)$, all types have common prior $p_0 = 1/2$, and public signals are uninformative and believed to be uninformative.[43]

A level-1 type incorporates solely its private information into its decision, and its likelihood ratio is constant across time, $\lambda_{1,t} = 1$ for all t. Therefore, a level-1 type chooses action R iff it observes a signal $s \geq 1/2$, independent of the history. A level-2 type believes all past actions are from level-1 types. Therefore, it believes that the informational content of actions are independent of the history, and the number of L and R actions are sufficient statistics for its likelihood ratio. One interpretation is the level-2 type uses a simple heuristic that counts the number of L and R actions in the history and uses this number to form beliefs. A level-3 type believes that almost all past actions are from level-2 types. Its subjective probability of an L action depends on the probability that level-2 types choose action L, which *does* depend on the history through λ_2. Therefore, how the level-3 type updates its beliefs depends on the history.

Proposition 6 establishes that there are two distinct regions of learning, which feature cyclical learning or disagreement depending on the true distribution over types.[44]

Proposition 6. *There exists an $\bar{\varepsilon} > 0$ such that if $\varepsilon \in (0, \bar{\varepsilon})$, then either learning is cyclical almost surely or disagreement occurs almost surely. For $\varepsilon \in (0, \bar{\varepsilon})$, there exists a cutoff $\bar{\pi}_3 \in (0, 1)$ such that if $\pi(\theta_3) > \bar{\pi}_3$, then almost surely learning is cyclical, there*

Bayesians" in Hung and Plott (2001). The naive learners in Bohren (2016) can be interpreted as a modified level-2 type that allows for the possibility that other agents are also level-2. In Eyster and Rabin (2010), all agents have the same model – they are all level-2 – while Bohren (2016), level-1 and level-2 types both occur with positive probability.

[42]The exact parameterization of the level-k model, i.e. $\varepsilon = 0$, violates Assumption 3 and 4. In a cognitive hierarchy model (Camerer, Ho, and Chong 2004), level-3 places non-trivial probability on level-1 types. We explore this alternative parameterization in Online Appendix F.2.

[43]These assumptions are made for expositional simplicity. The results from Section 4 apply to any level-k model in which the level-1 type occurs with positive probability, $\pi(\theta_1) > 0$ (to satisfy Assumption 3).

[44]Theorem 1 characterizes the asymptotic learning outcomes as follows: (i) construct the set of locally stable learning outcomes $\Lambda(\omega)$, (ii) show both disagreement outcomes are maximally accessible, and (iii) show $\Lambda_M(\omega)$ is empty. See Appendix C.5 for a detailed construction of $\Lambda(\omega)$ and $\Lambda_M(\omega)$. It follows from (ii) and (iii) that $\Lambda(\omega)$ fully characterizes the set of asymptotic learning outcomes.

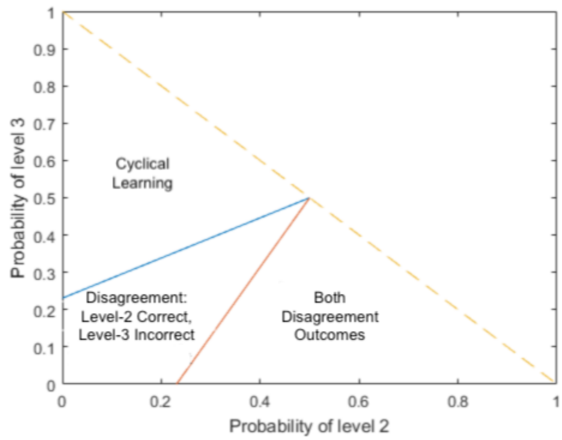

FIGURE 5. Level-k Learning Outcomes
($\omega = L$, $F^L = \frac{10}{3}(s - .5s^2 - .6)$, $F^R(s) = \frac{5}{3}(s^2 - .04)$)

exists a cutoff $\bar{\pi}_2 \in (0,1)$ such that if $\pi(\theta_2) > \bar{\pi}_2$, then both disagreement outcomes arise with positive probability, and there exists a cutoff $\bar{\pi}_1 \in (0,1)$ such that if $\pi(\theta_1) > \bar{\pi}_1$, then disagreement outcome in which level-2 learns the correct state and level-3 learns the incorrect state arises almost surely.

To understand why correct learning almost surely does not arise, suppose that the state is L and the beliefs of both level-2 and level-3 are near zero. When this is the case, a level-3 agent underreacts to L actions since she believes these actions are from herding level-2 agents. However, some L actions are from level-1 agents who reveal private information. In contrast, she accurately interprets R actions, since she correctly attributes these actions to the level-1 agents who reveal private information. Therefore, the level-3 type's belief increases in expectation, moving away from zero, and correct learning is not locally stable. The intuition is similar for incorrect learning.

This leaves the disagreement outcomes and cyclical learning as candidate learning outcomes. Consider disagreement outcome near which level-2 agents choose action L and level-3 agents choose action R. If a large share of agents are level-3, then most agents choose action R. A level-2 agent overreacts to these R actions, since she believes they are from level-1 agents who reveal private information. This pulls the level-2 type's belief away from zero and the disagreement outcome is not locally stable. The intuition for the other disagreement outcome is similar. Therefore, disagreement almost surely does not arise and learning is cyclical.

As the share of level-1 and level-2 types increase, convergence becomes possible and disagreement emerges. Disagreement is driven by level-2 agents' imitation of the more frequent action and level-3 agents' anti-imitation in order to correct for the level-2 agents' overreaction. If a large share of agents are level-1, then most actions do indeed reveal private information. Therefore, the level-2 type's model is close to correct and almost surely level-2 agents learn the correct state. This corresponds to the disagreement outcome in which level-2 agents learn the correct state and level-3 agents learn the incorrect state arising almost surely. In this case, a higher level of reasoning performs strictly *worse* than a lower level of reasoning, since level-3 agents almost surely learn the incorrect state.

If a large share of agents are level-2, then both disagreement outcomes arise with positive probability. Consider the disagreement outcome at which level-2 agents choose L and level-3 agents choose R. Due to the high share of level-2 agents, most agents choose L. A level-2 agent overreacts to these L actions, confirming her belief that the state is L. In contrast, a level-3 agent believes L actions are from herding level-2 agents and R actions are from informative level-1 agents. She does not account for informative level-1 agents playing L or herding level-3 agents playing R. Therefore, she underreacts to L actions and overreacts to R actions. This confirms her belief that the state is R, despite the high frequency of L actions. Therefore, this disagreement outcome is locally stable. The intuition for the other disagreement outcome is similar. Given that both disagreement outcomes arise, learning is path dependent. Two similar populations who learn about the same state from different action histories may converge to very different long-run beliefs.

Fig. 5 illustrates these learning regions. In a social learning experiment, Penczynski (2017) finds that most agents' behavior is consistent with level-1, 2 or 3 types, with a modal type of level-2. His estimate of the type distribution lies in the region where both disagreement outcomes arise with positive probability.

6 Conclusion

We develop a general framework to study sequential social learning and passive individual learning with model misspecification. The framework can capture many biases and heuristics in interpreting information, including those discussed in Section 2.2. Our main result characterizes how asymptotic learning outcomes depend on two expressions that are straightforward to derive from the underlying form of misspecification. This provides a unified way to compare different forms of misspecification that have been previously studied, as well as provides new insights about forms of misspecification that have not been theoretically explored. We use the characterization to show that some forms of misspecification are robust to different parameterizations (i.e. confirmation bias) while others are not (i.e. under and overreaction to signals), and explore whether a representative agent approach yields accurate conclusions about long-run learning in the presence of model heterogeneity. The characterization also yields new insights into how misspecification impacts social learning in a level-k framework and provides a rational for entrenched disagreement, in which agents' beliefs converge to certainty about different states despite observing each others' actions.

References

Angrisani, M., A. Guarino, P. Jehiel, and T. Kitagawa (2018): "Information redundancy neglect versus overconfidence: a social learning experiment," *mimeo*.

Banerjee, A. V. (1992): "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, 107, 797–817.

Bartels, L. M. (2002): "Beyond the running tally: Partisan bias in political perceptions," *Political Behavior*, 24, 117–150.

Bénabou, R. and J. Tirole (2011): "Identity, morals, and taboos: Beliefs as assets," *The Quarterly Journal of Economics*, 126, 805–855.

Berk, R. H. (1966): "Limiting Behavior of Posterior Distributions When the Model is Incorrect," *The Annals of Mathematical Statistics*, 37, 51–58.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *The Journal of Political Economy*, 100, 992–1026.

Bohren, A. (2016): "Informational Herding with Model Misspecification," *Journal of Economic Theory*, 222–247.

Bohren, J. A. and D. Hauser (2019): "Misinterpreting Social Outcomes and Information Campaigns," .

Brunnermeier, M. K. and J. A. Parker (2005): "Optimal expectations," *The American Economic Review*, 95, 1092–1118.

Bushong, B. and T. Gagnon-Bartsch (2017): "Learning with Misattribution of Reference Dependence," *Mimeo*.

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004): "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, 119, 861–898.

Costa-Gomes, M. A., V. P. Crawford, and N. Iriberri (2009): "Comparing Models of Strategic Thinking in Van Huyck, Battalio, and Beil's Coordination Games," *Journal of the European Economic Association*, 7, 365–376.

Cripps, M. W. (2018): "Divisible Updating," *mimeo*.

Darley, J. M. and P. H. Gross (1983): "A hypothesis-confirming bias in labeling effects." *Journal of Personality and Social Psychology*, 44, 20.

Edward, W. (1982): "Conservatism in human information processing," in *Judgment under Uncertainty: Heuristics and Biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky, Cambridge: Cambridge University Press, 359–369.

Enke, B. and F. Zimmermann (2019): "Correlation Neglect in Belief Formation," *The Review of Economic Studies*, 86, 313–332.

Epstein, L. G., J. Noor, and A. Sandroni (2010): "Non-Bayesian Learning," *The B.E. Journal of Theoretical Economics*, 10.

Esponda, I. and D. Pouzo (2016): "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," *Econometrica*, 84, 1093–1130.

——— (2017): "Equilibrium in Misspecified Markov Decision Processes," *Mimeo.*

Eyster, E. and M. Rabin (2010): "Naive Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, 2, 221–243.

Eyster, E., M. Rabin, and G. Weizsäcker (2018): "An Experiment on Social Mislearning," .

Eyster, E. and G. Weizsacker (2011): "Correlation neglect in financial decision-making," .

Frick, M., R. Iijima, and Y. Ishii (2018): "Dispersed Behavior and Perceptions in Assortative Societies," *mimeo.*

——— (2019): "Misinterpreting Others and the Fragility of Social Learning," *Working Paper*, 1–53.

Fryer, R. G., P. Harms, and M. O. Jackson (2018): "Updating Beliefs When Evidence Is Open to Interpretation: Implications for Bias and Polarization," *Journal of the European Economic Association.*

Fudenberg, D., G. Romanyuk, and P. Strack (2017): "Active learning with a misspecified prior," *Theoretical Economics*, 12, 1155–1189.

Gagnon-Bartsch, T. (2017): "Taste Projection in Models of Social Learning," *Mimeo.*

Gagnon-Bartsch, T. and M. Rabin (2017): "Naive Social Learning, Mislearning, and Unlearning," *Mimeo.*

Gilovich, T. (1990): "Differential construal and the false consensus effect," *Journal of personality and social psychology*, 59, 623.

Gottlieb, D. (2015): "Will You Never Learn? Self Deception and Biases in Information Processing," *Mimeo.*

Grebe, T., J. Schmid, and A. Stiehler (2008): "Do individuals recognize cascade behavior of others? – An experimental study," *Journal of Economic Psychology*, 29, 197 – 209.

Griffin, D. and A. Tversky (1992): "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology*, 24, 411–435.

Guarino, A. and P. Jehiel (2013): "Social Learning with Coarse Inference," *American Economic Journal: Microeconomics*, 5, 147–74.

Heidhues, P., B. Koszegi, and P. Strack (2018): "Unrealistic Expectations and Misguided Learning," *Econometrica*, forthcoming.

——— (2019): "Convergence in Misspecified Learning Models with Endogenous Actions," .

Hung, A. A. and C. R. Plott (2001): "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions," *American Economic Review*, 91, 1508–1520.

Jadbabaie, A., P. Molavi, A. Sandroni, and A. Tahbaz-Salehi (2012): "Non-Bayesian social learning," *Games and Economic Behavior*, 76, 210 – 225.

Jerit, J. and J. Barabas (2012): "Partisan perceptual bias and the information environment," *The Journal of Politics*, 74, 672–684.

Kallir, I. and D. Sonsino (2009): "The Neglect of Correlation in Allocation Decisions," *Southern Economic Journal*, 75, 1045–1066.

Kleijn, B. J. and A. W. van der Vaart (2006): "Misspecification in infinite-dimensional Bayesian statistics," *The Annals of Statistics*, 837–877.

Köszegi, B. (2006): "Ego utility, overconfidence, and task choice," *Journal of the European Economic Association*, 4, 673–707.

Kübler, D. and G. Weizsäcker (2004): "Limited depth of reasoning and failure of cascade formation in the laboratory," *The Review of Economic Studies*, 71, 425–441.

——— (2005): "Are longer cascades more stable?" *Journal of the European Economic Association*, 3, 330–339.

Kunda, Z. (1990): "The case for motivated reasoning." *Psychological bulletin*, 108, 480.

Levy, G. and R. Razin (2015): "Correlation neglect, voting behavior, and information aggregation," *American Economic Review*, 105, 1634–45.

Lord, C. G., L. Ross, and M. R. Lepper (1979): "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence." *Journal of personality and social psychology*, 37, 2098.

Madarász, K. and A. Prat (2016): "Sellers with misspecified models," *The Review of Economic Studies*.

Marks, G. and N. Miller (1987): "Ten years of research on the false-consensus effect: An empirical and theoretical review." *Psychological Bulletin*, 102, 72.

Miller, D. T. and C. McFarland (1987): "Pluralistic ignorance: When similarity is interpreted as dissimilarity." *Journal of Personality and Social Psychology*, 53, 298.

——— (1991): "When social comparison goes awry: The case of pluralistic ignorance," in *Social Comparison: Contemporary Theory and Research*, ed. by J. Suls and T. A. Wills, Lawrence Erlbaum Associates, Inc, 287–313.

Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie (2018): "A Theory of Non-Bayesian Social Learning," *Econometrica*, 86, 445–490.

Moore, D. A. and P. J. Healy (2008): "The trouble with overconfidence." *Psychological review*, 115, 502.

Nyarko, Y. (1991): "Learning in mis-specified models and the possibility of cycles," *Journal of Economic Theory*, 55, 416 – 427.

Ortoleva, P. (2012): "Modeling the change of paradigm: Non-Bayesian reactions to unexpected news," *American Economic Review*, 102, 2410–36.

Ortoleva, P. and E. Snowberg (2015): "Overconfidence in political behavior," *American Economic Review*, 105, 504–35.

Penczynski, S. P. (2017): "The nature of social learning: Experimental evidence," *European Economic Review*, 94, 148–165.

Plous, S. (1991): "Biases in the assimilation of technological breakdowns: Do accidents make us safer?" *Journal of Applied Social Psychology*, 21, 1058–1082.

Rabin, M. and J. L. Schrag (1999): "First Impressions Matter: A Model of Confirmatory Bias," *The Quarterly Journal of Economics*, 114, 37–82.

Ross, L., D. Greene, and P. House (1977): "The "false consensus effect": An egocentric bias in social perception and attribution processes," *Journal of experimental social psychology*, 13, 279–301.

Schwartzstein, J. (2014): "Selective Attention and Learning," *Journal of the European Economic Association*, 12, 1423–1452.

Shalizi, C. R. (2009): "Dynamics of Bayesian updating with dependent data and misspecified models," *Electronic Journal of Statistics*, 3, 1039–1074.

Smith, L. and P. Sorensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.

Tversky, A. and D. Kahneman (1971): "Belief in the Law of Small Numbers," *Psychological Bulletin*, 76, 105–110.

Ungeheuer, M. and M. Weber (2017): "The Perception of Dependence, Investment Decisions, and Stock Prices," *Mimeo.*

Wilson, A. (2014): "Bounded Memory and Biases in Information Processing," *Econometrica*, 82, 2257–2294.

A Proofs from Section 4

A.1 Proof of Theorem 1

We establish Theorem 1 through a series of lemmas. In Lemma 1, we characterize the set of stationary beliefs, which are candidate limit points of $\langle \boldsymbol{\lambda}_t \rangle$. Lemma 2 rules out convergence to non-stationary beliefs. Next, Lemma 3 establishes when a stationary belief is locally stable. Lemma 4 establishes that global stability immediately follows from local stability for agreement outcomes, while Lemma 5 establishes that maximal accessibility is a sufficient condition for global stability of disagreement outcomes. In Lemma 6, we show that locally stable mixed learning outcomes must be in $\Lambda_M(\omega)$. Therefore, if $\Lambda_M(\omega)$ is empty, almost surely mixed learning does not arise. Finally, Lemma 7 establishes that the likelihood ratio converges almost surely for all sociable types when there is at least one globally stable stationary outcome and no locally stable mixed outcomes.

We present Lemmas 1 to 4 for any finite number of sociable types $k \geq 1$, as the constructions of local stability and the global stability of agreement outcomes are identical for all finite k. Establishing the global stability of disagreement outcomes and ruling out mixed learning is more involved for more than two sociable types, as the number of possible outcomes increases with k. Therefore, we present Lemmas 5 and 6 for $k = 2$ sociable types (trivially, disagreement and mixed learning are not possible with a single sociable type, i.e. $k = 1$). Finally, we use Lemmas 1 to 6 to establish belief convergence in Lemma 7 for $k \leq 2$ sociable types. The analogues of Lemmas 5 to 7 for $k > 2$ are in Online Appendix E.2.

Throughout this section, assume Assumption 1, 2, 3 and 4. Given $\varepsilon > 0$, define a neighborhood $B_\varepsilon(\boldsymbol{\lambda})$ of $\boldsymbol{\lambda} \in \{0, \infty\}^k$ as $\lambda_i \in [0, \varepsilon)$ if $\lambda_i = 0$ and $\lambda_i \in (1/\varepsilon, \infty]$ if $\lambda_i = \infty$.

A.1.1 Statement of Lemmas

In this section, we state Lemmas 1 to 7 outlined above. The proofs follow in Appendix A.1.2.

By Assumption 3, the set of stationary beliefs correspond to each type placing probability one on either state L ($\lambda = 0$) or state R ($\lambda = \infty$).

Lemma 1 (Stationary Beliefs)**.** *The set of stationary beliefs are* $\{0, \infty\}^k$.

Further, the likelihood ratio almost surely does not converge to non-stationary beliefs.

Lemma 2 (Non-Stationary Beliefs)**.** *Given* $\boldsymbol{\lambda}^*$, *if* $\lambda_i^* \in (0, 1)$ *for some* θ_i, *then* $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) = 0$.

Therefore, if the likelihood ratio converges for all types, then it must converge to a stationary belief $\boldsymbol{\lambda}^* \in \{0, \infty\}^k$.

Next, we determine when the likelihood ratio converges with positive probability. Recall that $\boldsymbol{\lambda}^*$ is *locally stable* if the process $\langle \boldsymbol{\lambda}_t \rangle$ converges to $\boldsymbol{\lambda}^*$ with positive probability from a neighborhood of $\boldsymbol{\lambda}^*$, and that $\gamma_i(\boldsymbol{\lambda}, \omega)$ is the expected change in the log likelihood ratio for type θ_i at belief $\boldsymbol{\lambda}$. Lemma 3 establishes the relationship between the local stability of stationary belief $\boldsymbol{\lambda}^*$ and the sign of $\gamma_i(\boldsymbol{\lambda}^*, \omega)$.

Lemma 3 (Locally Stable Beliefs). *Let $\boldsymbol{\lambda}^* \in \{0, \infty\}^k$ be a stationary belief.*

1. *If $\gamma_i(\boldsymbol{\lambda}^*, \omega) < 0$ for all $\theta_i \in \Theta_S$ such that $\lambda_i^* = 0$ and $\gamma_i(\boldsymbol{\lambda}^*, \omega) > 0$ for all $\theta_i \in \Theta_S$ such that $\lambda_i^* = \infty$, then $\boldsymbol{\lambda}^*$ is locally stable.*

2. *If there exists a $\theta_i \in \Theta_S$ such that $\lambda_i^* = 0$ and $\gamma_i(\boldsymbol{\lambda}^*, \omega) > 0$ or $\lambda_i^* = \infty$ and $\gamma_i(\boldsymbol{\lambda}^*, \omega) < 0$, then $\boldsymbol{\lambda}^*$ is not locally stable and $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) = 0$.*

Lemma 3 uses results on the local stability of nonlinear equations developed in Smith and Sorensen (2000) (Theorems C.1 and C.2). Given Lemma 3, the set $\Lambda(\omega)$ defined in (9) is the set of locally stable beliefs for a generic misspecified model. If there are no locally stable beliefs, i.e. $\Lambda(\omega)$ is empty in a generic misspecified model, then the likelihood ratio almost surely does not converge for at least one type, as Lemma 3 rules out convergence to stationary beliefs that are not locally stable and Lemma 2 rules out convergence to non-stationary beliefs.

We are interested in determining whether convergence occurs with positive probability from any initial value of the likelihood ratio, i.e. global stability. Clearly, the set of globally stable learning outcomes is a subset of the set of locally stable learning outcomes. Therefore, it remains to establish when local stability implies global stability. For agreement outcomes, $\boldsymbol{\lambda}^* \in \{0^k, \infty^k\}$, global stability immediately follows from local stability.

Lemma 4 (Global Stability of Agreement). *For $\boldsymbol{\lambda}^* \in \{0^k, \infty^k\}$, if $\boldsymbol{\lambda}^*$ is locally stable, then $\boldsymbol{\lambda}^*$ is globally stable, i.e. for any initial belief $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) > 0$.*

All types update their beliefs in the same direction following either the maximal action and signal in favor of state L, (a_1, σ_L), or the maximal action and signal in favor of state R, (a_M, σ_R). Therefore, from any initial belief, it is possible construct a finite sequence of action and public signal pairs that occurs with positive probability and pushes the likelihood ratio arbitrarily close to an agreement outcome. Once the likelihood ratio is in a neighborhood of the agreement outcome, local stability establishes positive probability of convergence.

Lemma 5 establishes that maximal accessibility is a sufficient condition for the global stability of a disagreement outcome in the case of two sociable types, $k = 2$ (recall that disagreement is not possible with a single sociable type).

Lemma 5 (Global Stability of Disagreement)**.** *Suppose $k = 2$. If disagreement outcome $\boldsymbol{\lambda}^* \in \{(0, \infty), (\infty, 0)\}$ is locally stable and maximally accessible, then $\boldsymbol{\lambda}^*$ is globally stable.*

Maximal accessibility orders the way each type interprets maximal actions and public signals, which guarantees that there exists a finite sequence of maximal actions and public signals that separates the beliefs of each type in the direction of the disagreement outcome.[45] As before, once the likelihood ratio is sufficiently close to the disagreement outcome, local stability establishes convergence.

As discussed in Section 4.2, a sufficient condition for ruling out mixed outcomes when $k = 2$ is that $\Lambda_M(\omega)$ is empty (recall that mixed learning is not possible with a single sociable type).

Lemma 6 (Unstable Mixed Outcomes)**.** *Suppose $k = 2$. If mixed learning outcome $(\lambda_i^*, \theta_i) \notin \Lambda_M(\omega)$, then $Pr(\lambda_{i,t} \to \lambda_i^*$ and $\lambda_{-i,t}$ does not converge$) = 0$.*

Finally, if there is at least one locally stable agreement outcome or locally stable and maximally accessible disagreement outcome and no locally stable mixed outcomes, then the likelihood ratio converges almost surely for all sociable types (recall that for $k = 1$, both stationary learning outcomes are agreement outcomes and $\Lambda_M(\omega) = \emptyset$).

Lemma 7 (Belief Convergence)**.** *Suppose $k \leq 2$. If $\Lambda(\omega)$ contains an agreement outcome or maximally accessible disagreement outcome and $\Lambda_M(\omega)$ is empty, then for any initial belief $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, there exists a random variable $\boldsymbol{\lambda}_\infty$ with $\text{supp}(\boldsymbol{\lambda}_\infty) = \Lambda(\omega)$ such that $\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}_\infty$ almost surely.*

Theorem 1 immediately follows. Part (1) follows from the local and global stability of agreement outcomes (Lemmas 3 and 4). Part (2) follows from the local and global stability of disagreement outcomes (Lemmas 3 and 5). For part (3), Lemmas 1 and 2 rule out convergence to non-stationary beliefs, Lemma 3 rules out convergence to stationary outcomes that are not locally stable, and Lemma 6 rules out convergence to a mixed learning outcome when $\Lambda_M(\omega)$ is empty. Therefore, if $\Lambda(\omega)$ is empty in a generic misspecified model, there are no locally stable learning outcomes and almost surely the likelihood ratio does not converge for at least one sociable type, establishing the second

[45] While maximal accessibility is simple and easy to verify, it can be restrictive, especially in models with large action or public signal spaces. In Lemma 9 (Appendix A.1.2), we establish a more general sufficient condition to separate beliefs, which we call separability (Definition 6).

statement in part (3).[46] If $\Lambda_M(\omega)$ is also empty, then almost surely the likelihood ratio does not converge for any sociable type, establishing the first statement in part (3). The final statement in part (3) follows from Lemma 7, which establishes when the likelihood ratio converges.

A.1.2 Proofs of Lemmas 1 to 7

Proof of Lemma 1 (Stationary Beliefs). At a stationary belief $\boldsymbol{\lambda}^* \in [0,\infty]^k$,

$$\boldsymbol{\lambda}^* = \boldsymbol{\lambda}^* \left(\frac{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda}^*)}{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda}^*)} \right) \tag{12}$$

for all (a,σ) such that $\psi(a,\sigma|\omega,\boldsymbol{\lambda}^*) > 0$. Trivially, (12) is satisfied for all $\boldsymbol{\lambda}^* \in \{0,\infty\}^k$, independent of $\psi(a,\sigma|\omega,\boldsymbol{\lambda}^*)$. Therefore, all $\boldsymbol{\lambda}^* \in \{0,\infty\}^k$ are stationary. It remains to be determined whether there exist any interior stationary beliefs $\boldsymbol{\lambda}^* \in (0,\infty)^k$.

Suppose $\boldsymbol{\lambda}^* \in (0,\infty)^k$ and Assumption 3.ii holds, i.e. there exists an autarkic θ_j with $\pi(\theta_j) > 0$ that plays a_1 with probability in $(0,1)$, and each sociable type θ_i believes this autarkic type occurs with positive probability, $\hat{\pi}_i(\theta_j) > 0$. Then the true probability of action a_1 and each type's subjective probability of action a_1 at $\boldsymbol{\lambda}^*$ are in $(0,1)$ for each state $\omega \in \{L,R\}$. Further, each type's subjective probability of a_1 in state R is less than its subjective probability of a_1 in state L, since $\hat{F}_i^R < \hat{F}_i^L$. Given $\hat{\sigma}_i(\sigma_L) \le 1/2$, this implies $\hat{\psi}_i(a_1,\sigma_L|R,\boldsymbol{\lambda}^*) < \hat{\psi}_i(a_1,\sigma_L|L,\boldsymbol{\lambda}^*)$ and (12) does not hold for (a_1,σ_L). But (a_1,σ_L) occurs with positive probability in either state, $\psi(a_1,\sigma_L|\omega,\boldsymbol{\lambda}^*) > 0$. Therefore, $\boldsymbol{\lambda}^*$ cannot be stationary.

Suppose $\boldsymbol{\lambda}^* \in (0,\infty)^k$ and Assumption 3.i holds. Then $\hat{\sigma}_i(\sigma_L) < 1/2$ and $\hat{\sigma}_i(\sigma_R) > 1/2$ for all sociable types θ_i. Further, σ_L and σ_R occur with positive probability, independent of $\boldsymbol{\lambda}$. At least one action a occurs with positive probability at $\boldsymbol{\lambda}^*$. Since public signals are informative, it must be that $\frac{\hat{\psi}_i(a,\sigma_L|R,\boldsymbol{\lambda}^*)}{\hat{\psi}_i(a,\sigma_L|L,\boldsymbol{\lambda}^*)} \neq \frac{\hat{\psi}_i(a,\sigma_R|R,\boldsymbol{\lambda}^*)}{\hat{\psi}_i(a,\sigma_R|L,\boldsymbol{\lambda}^*)}$. Therefore, (12) cannot hold for both (a,σ_L) and (a,σ_R). But both action-signal pairs occur with positive probability in either state, $\psi(a,\sigma_L|\omega,\boldsymbol{\lambda}^*) > 0$ and $\psi(a,\sigma_R|\omega,\boldsymbol{\lambda}^*) > 0$. Therefore, $\boldsymbol{\lambda}^*$ cannot be stationary. $\square$

Proof of Lemma 2 (Non-Stationary Beliefs). Suppose beliefs converge to a non-stationary belief $\boldsymbol{\lambda}^* \in [0,\infty]^k \setminus \{0,\infty\}^k$ with positive probability. After action and public signal $(\tilde{a}_t, \tilde{\sigma}_t) = (a_M, \sigma_R)$, by Lemma 10, $\lambda_{i,t+1} - \lambda_{i,t}$ is bounded uniformly away from zero for all sociable types $\theta_i \in \Theta_S$. For sufficiently small $\varepsilon > 0$, if $\boldsymbol{\lambda}_t \in B_\varepsilon(\boldsymbol{\lambda}^*)$,

[46] In the non-generic set of models where $\gamma_i(\boldsymbol{\lambda},\omega) = 0$ for some $\theta_i \in \Theta_S$, the values of $\boldsymbol{\lambda}$ that satisfy $\gamma_i(\boldsymbol{\lambda},\omega) = 0$ may or may not be locally stable. If $\boldsymbol{\lambda}$ is not locally stable, then $\langle \boldsymbol{\lambda}_t \rangle$ almost surely does not converge for at least one type, while if $\boldsymbol{\lambda}$ is locally stable (and maximally accessible), then Lemmas 4 and 5 establish that it is also globally stable.

then after observing $(\tilde{a}_t, \tilde{\sigma}_t) = (a_M, \sigma_R)$, $\lambda_{i,t+1} \notin B_\varepsilon(\boldsymbol{\lambda}^*)$ for any type with an interior belief $\lambda_{i,t} \in (0, \infty)$. The probability $Pr(\exists t < T | (\tilde{a}_t, \tilde{\sigma}_t) = (a_M, \sigma_R))$ converges to one as $T \to \infty$. Therefore, the likelihood ratio almost surely leaves $B_\varepsilon(\boldsymbol{\lambda}^*)$. □

Proof of Lemma 3 (Locally Stable Beliefs). Suppose $\omega = L$. The proof for $\omega = R$ is analogous.

Part 1. Consider $\boldsymbol{\lambda}^* = 0^k$ and suppose $\gamma_i(0^k, L) < 0$ for all sociable types $\theta_i \in \Theta_S$. Then there exists a $\varepsilon > 0$ such that in the neighborhood $B_\varepsilon(0^k) \equiv [0, \varepsilon]^k$ of 0^k,

$$\sum_{(a,\sigma) \in \mathcal{A} \times \Sigma} \psi(a, \sigma | L, 0^k) \sup_{\boldsymbol{\lambda} \in [0,\varepsilon]^k} \log \frac{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda})} < 0. \tag{13}$$

for all $\theta_i \in \Theta_S$. Let

$$g_i(a, \sigma) \equiv \sup_{\boldsymbol{\lambda} \in [0,\varepsilon]^k} \log \frac{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda})}$$

denote the maximal update from action and signal (a, σ) in the neighborhood $[0, \varepsilon]^k$, with $\mathbf{g}(a, \sigma) \equiv (g_1(a, \sigma), ..., g_k(a, \sigma))$. Let

$$\bar{g}_i \equiv \max_{(a,\sigma) \in \mathcal{A} \times \Sigma} g_i(a, \sigma)$$

denote the maximal update across all action and signal pairs in the neighborhood $[0, \varepsilon]^k$, with $\bar{\mathbf{g}} \equiv (\bar{g}_1, ..., \bar{g}_k)$.

For $\delta > 0$, choose a neighborhood $[0, \varepsilon_\delta]^k \subseteq [0, \varepsilon]^k$ with

$$\sup_{\boldsymbol{\lambda} \in [0,\varepsilon_\delta]^k} |\psi(a, \sigma | L, \boldsymbol{\lambda}) - \psi(a, \sigma | L, 0^k)| < \delta.$$

By Lemma 11, $\psi(a, \sigma | L, \boldsymbol{\lambda})$ is continuous at $\boldsymbol{\lambda} = 0^k$, so such a neighborhood exists. Suppose $\boldsymbol{\lambda}_1 \in [0, \varepsilon_\delta]^k$. Let $a(\theta, s, \boldsymbol{\lambda})$ be the optimal action for type θ at beliefs $\boldsymbol{\lambda}$ after observing private signal s. Define the linear system $\langle \boldsymbol{\lambda}_{\delta,t} \rangle_{t=1}^\infty$ as follows: $\boldsymbol{\lambda}_{\delta,1} = \boldsymbol{\lambda}_1$,

$$\log \boldsymbol{\lambda}_{\delta,t+1} = \log \boldsymbol{\lambda}_{\delta,t} + \mathbf{g}(a(\tilde{\theta}_t, \tilde{s}_t, 0^k), \tilde{\sigma}_t),$$

when $(\tilde{\theta}_t, \tilde{s}_t)$ is such that $a(\tilde{\theta}_t, \tilde{s}_t, \boldsymbol{\lambda}) = a(\tilde{\theta}_t, \tilde{s}_t, 0^k)$ for all beliefs $\boldsymbol{\lambda} \in [0, \varepsilon_\delta]$ (note this includes all autarkic types), and

$$\log \boldsymbol{\lambda}_{\delta,t+1} = \log \boldsymbol{\lambda}_{\delta,t} + \bar{\mathbf{g}}$$

otherwise. When $\omega = L$, let $\psi_\delta(a, \sigma)$ be the probability of (a, σ) in the former event and let $\bar{\psi}_\delta$ be the probability of the latter event. Note $\psi_\delta(a, \sigma) \leq \inf_{\boldsymbol{\lambda} \in [0,\varepsilon_\delta]^k} \psi(a, \sigma | L, \boldsymbol{\lambda})$ and

$\bar{\psi}_\delta + \sum_{(a,\sigma)\in\mathcal{A}\times\Sigma} \psi_\delta(a,\sigma|L) = 1$. By Lemma C.1 of Smith and Sorensen (2000), if

$$\bar{\psi}_\delta \bar{g}_i + \sum_{(a,\sigma)\in\mathcal{A}\times\Sigma} \psi_\delta(a,\sigma) g_i(a,\sigma) < 0 \tag{14}$$

for all $\theta_i \in \Theta_S$, then almost surely $\lim_{t\to\infty} \boldsymbol{\lambda}_{\delta,t} = 0^k$. Equation (14) holds for sufficiently small δ, since by (13), it is strictly less than zero at $\delta = 0$.

Let $\delta_1 > 0$ denote an upper bound such that (14) holds for all $\delta < \delta_1$. Whenever $(\tilde{\theta}_t, \tilde{s}_t)$ is such that $a(\tilde{\theta}_t, \tilde{s}_t, \boldsymbol{\lambda}) = a(\tilde{\theta}_t, \tilde{s}_t, 0^k)$ for all $\boldsymbol{\lambda} \in [0, \varepsilon_\delta]$, the process $\langle \log \boldsymbol{\lambda}_{\delta,t}\rangle$ updates by $\mathbf{g}(a,\sigma)$. When $\boldsymbol{\lambda}_t \in [0,\varepsilon_\delta]^k$, by construction this is larger than the update to the process $\langle \log \boldsymbol{\lambda}_t \rangle$, which is $\log \frac{\hat{\psi}_i(a,\sigma|R,\boldsymbol{\lambda}_t)}{\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda}_t)}$ for each type $\theta_i \in \Theta_S$. Otherwise, $\langle \log \boldsymbol{\lambda}_{\delta,t}\rangle$ updates by $\bar{\mathbf{g}}$, which is also larger than the update to $\langle \log \boldsymbol{\lambda}_t\rangle$ when $\boldsymbol{\lambda}_t \in [0,\varepsilon_\delta]^k$. Therefore, for $\delta < \delta_1$, if $\boldsymbol{\lambda}_{\delta,t} \geq \boldsymbol{\lambda}_t$ and $\boldsymbol{\lambda}_{\delta,t} \in [0,\varepsilon_\delta]^k$, then $\boldsymbol{\lambda}_{\delta,t+1} \geq \boldsymbol{\lambda}_{t+1}$. Since $\boldsymbol{\lambda}_{\delta,1} \in [0,\varepsilon_\delta]^k$, as long as it remains in $[0,\varepsilon_\delta]^k$, $\langle \boldsymbol{\lambda}_t \rangle$ is bounded above by a stochastic process that converges to zero almost surely.

Since $\lim_{t\to\infty} \boldsymbol{\lambda}_{\delta,t} = 0^k$ almost surely for $\delta < \delta_1$,

$$Pr(\cup_t \cap_{s\geq t} \{\boldsymbol{\lambda}_{\delta,s} \in [0,\varepsilon_\delta]^k\}) = 1.$$

Therefore, there exists a $t \geq 1$ such that $Pr(\forall s \geq t, \boldsymbol{\lambda}_{\delta,s} \in [0,\varepsilon_\delta]^k) > 0$. Since the system is linear, if this holds at some $t > 1$, it must hold at $t = 1$. Therefore, there exists some $\boldsymbol{\lambda}_{\delta,1} \in [0,\varepsilon_\delta]^k$, with positive probability, $\boldsymbol{\lambda}_{\delta,t}$ remains in $[0,\varepsilon_\delta]^k$ for all $t > 1$ and $\boldsymbol{\lambda}_t \leq \boldsymbol{\lambda}_{\delta,t}$. Moreover, this holds for all $\boldsymbol{\lambda} \leq \boldsymbol{\lambda}_{\delta,1}$. When this happens, since $\lim_{t\to\infty} \boldsymbol{\lambda}_{\delta,t} = 0^k$, it must also be that $\lim_{t\to\infty} \boldsymbol{\lambda}_t = 0^k$. Let $\varepsilon^* = \inf \lambda_{\delta,i,1}$ This establishes that when $\boldsymbol{\lambda}_1 \in [0,\varepsilon^*]^k$, with positive probability, $\lim_{t\to\infty} \boldsymbol{\lambda}_t = 0^k$ i.e. $\boldsymbol{\lambda}^* = 0^k$ is locally stable.

The proofs for the other stationary beliefs are analogous. If $\lambda_i^* = \infty$, substitute λ_i^{-1} for type θ_i and modify the transition rules accordingly.

Part 2. Let $\boldsymbol{\lambda}^* \in \{0,\infty\}^k$ be a stationary belief and suppose that there exists a type, which without loss of generality we denote θ_1, such that $\lambda_1^* = 0$ but $\gamma_1(\boldsymbol{\lambda}^*, L) > 0$. Without loss of generality, suppose the types are ordered so that the first κ types correspond to $\lambda_i^* = 0$ and the latter $k-\kappa$ types correspond to $\lambda_i^* = \infty$. Since $\gamma_1(\boldsymbol{\lambda}^*, L) > 0$, there exists a $\varepsilon > 0$ such that for neighborhood $B_\varepsilon(\boldsymbol{\lambda}^*) \equiv [0,\varepsilon]^\kappa \times [1/\varepsilon, \infty]^{k-\kappa}$ of $\boldsymbol{\lambda}^*$,

$$\sum_{(a,\sigma)\in\mathcal{A}\times\Sigma} \psi(a,\sigma|L,\boldsymbol{\lambda}^*) \inf_{\boldsymbol{\lambda}\in B_\varepsilon(\boldsymbol{\lambda}^*)} \log \frac{\hat{\psi}_1(a,\sigma|R,\boldsymbol{\lambda})}{\hat{\psi}_1(a,\sigma|L,\boldsymbol{\lambda})} > 0. \tag{15}$$

Let $\tau_\varepsilon \equiv \min\{\tau | \boldsymbol{\lambda}_t \in B_\varepsilon(\boldsymbol{\lambda}^*)\ \forall t \geq \tau\}$ be the first time at which beliefs enter $B_\varepsilon(\boldsymbol{\lambda}^*)$ and never exit. Suppose $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) > 0$. Then for all $\varepsilon > 0$, $\tau_\varepsilon < \infty$ with positive probability. We will reach a contradiction by showing that for small enough ε, $\tau_\varepsilon = \infty$ almost surely. Let

$$g_1(a,\sigma) \equiv \inf_{\boldsymbol{\lambda}\in B_\varepsilon(\boldsymbol{\lambda}^*)} \log \frac{\hat{\psi}_1(a,\sigma|R,\boldsymbol{\lambda})}{\hat{\psi}_1(a,\sigma|L,\boldsymbol{\lambda})}$$

denote the minimal update for type θ_1 following action and signal (a,σ) in the neighborhood $B_\varepsilon(\boldsymbol{\lambda}^*)$ and let

$$\underline{g}_1 \equiv \min_{(a,\sigma)\in\mathcal{A}\times\Sigma} g_1(a,\sigma)$$

denote the minimal update across all action and signal pairs in the neighborhood $B_\varepsilon(\boldsymbol{\lambda}^*)$. Suppose $\boldsymbol{\lambda}_\tau \in B_\varepsilon(\boldsymbol{\lambda}^*)$ for some time τ (if such a τ doesn't exist, then clearly $\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*$ is not possible along such a sample path). As above, let $a(\theta,s,\boldsymbol{\lambda})$ be the optimal action for type θ at beliefs $\boldsymbol{\lambda}$ after observing private signal s. Define a linear system $\langle\tilde{\lambda}_t\rangle$ as follows: let $\tilde{\lambda}_\tau = \lambda_{1,\tau}$ and for $t>\tau$,

$$\log\tilde{\lambda}_{t+1} = \log\tilde{\lambda}_t + g_1(a(\tilde{\theta}_t,\tilde{s}_t,\boldsymbol{\lambda}^*),\tilde{\sigma}_t)$$

when $(\tilde{\theta}_t,\tilde{s}_t)$ is such that $a(\tilde{\theta}_t,\tilde{s}_t,\boldsymbol{\lambda}) = a(\tilde{\theta}_t,\tilde{s}_t,\boldsymbol{\lambda}^*)$ for all beliefs $\boldsymbol{\lambda}\in B_\varepsilon(\boldsymbol{\lambda}^*)$ (note this includes all autarkic types), and

$$\log\tilde{\lambda}_{t+1} = \log\tilde{\lambda}_t + \underline{g}_1$$

otherwise. When $\omega = L$, let $\psi(a,\sigma)$ be the probability of (a,σ) in the former event and let $\underline{\psi}$ be the probability of the latter event. Note $\underline{\psi} + \sum_{(a,\sigma)\in\mathcal{A}\times\Sigma}\psi(a,\sigma) = 1$. Choose ε sufficiently small so that

$$\underline{\psi}\,\underline{g}_1 + \sum_{(a,\sigma)\in\mathcal{A}\times\Sigma}\psi(a,\sigma)g_1(a,\sigma) > 0. \tag{16}$$

Given (15), (16) is strictly greater than zero at $\varepsilon = 0$, so such an ε exists. Moreover, $(\log\tilde{\lambda}_{t+1} - \log\tilde{\lambda}_t)_{t=\tau}^\infty$ is an i.i.d. process with expectation equal to (16). By the Law of Large Numbers, almost surely, $\frac{1}{t}(\log\tilde{\lambda}_{t+1} - \log\tilde{\lambda}_t)$ converges to (16), which is positive. Therefore,

$$\lim_{t\to\infty}\log\tilde{\lambda}_t = \lim_{t\to\infty}\left(\log\lambda_{1,\tau} + \sum_{s=\tau}^{t}(\log\tilde{\lambda}_{s+1} - \log\tilde{\lambda}_s)\right) \to \infty.$$

By definition of $\langle\tilde{\lambda}_t\rangle$, if $\lambda_{1,t} \geq \tilde{\lambda}_t$ and $\boldsymbol{\lambda}_t \in B_\varepsilon(\boldsymbol{\lambda}^*)$, then $\lambda_{1,t+1} \geq \tilde{\lambda}_{1,t+1}$. Since $\boldsymbol{\lambda}_\tau \in B_\varepsilon(\boldsymbol{\lambda}^*)$, as long as $\langle\boldsymbol{\lambda}_t\rangle$ remains in $B_\varepsilon(\boldsymbol{\lambda}^*)$ for $t>\tau$, $\langle\lambda_{1,t}\rangle$ is bounded below by the stochastic process $\langle\tilde{\lambda}_t\rangle$. Therefore, if $\langle\boldsymbol{\lambda}_t\rangle$ remains in $B_\varepsilon(\boldsymbol{\lambda}^*)$ for all $t > \tau$

$$\lim_{t\to\infty} \log \lambda_{1,t} \geq \lim_{t\to\infty} \log \tilde{\lambda}_t \to \infty.$$

But this implies that for small enough ε, $\boldsymbol{\lambda_t} \notin B_\varepsilon(\boldsymbol{\lambda}^*)$ for some $t > \tau$. This is a contradiction. So it must be that for small enough ε, $\tau_\varepsilon = \infty$ almost surely. Therefore, $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) = 0$.

Similar logic establishes that for stationary $\boldsymbol{\lambda}^*$ such that $\lambda_1^* = \infty$ and $\gamma_1(\boldsymbol{\lambda}^*, L) < 0$, $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) = 0$. □

Locally Stable Neighborhoods. From Lemma 3, if $\boldsymbol{\lambda}^* \in \Lambda(\omega)$, then $\boldsymbol{\lambda}^*$ is locally stable, i.e. there exists an $\varepsilon > 0$ and a *stable* neighborhood $B_\varepsilon(\boldsymbol{\lambda}^*)$ such that when $\boldsymbol{\lambda}_1 \in B_\varepsilon(\boldsymbol{\lambda}^*)$, $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) > 0$. Also, generically, for each stationary belief $\boldsymbol{\lambda}^* \notin \Lambda(\omega)$, there exists an $\varepsilon > 0$ and an *unstable* neighborhood $B_\varepsilon(\boldsymbol{\lambda}^*)$ such that when $\boldsymbol{\lambda}_1 \in B_\varepsilon(\boldsymbol{\lambda}^*)$, $\langle\boldsymbol{\lambda}_t\rangle$ almost surely leaves this neighborhood.

Fix state ω and define $E > 0$ as the smallest constant such that if $\log \lambda_i \in \mathbb{R}\backslash[-E, E]$ for each $\theta_i \in \Theta_S$, then $\boldsymbol{\lambda}$ is contained in one of these stable or unstable neighborhoods, and let $B_E(\boldsymbol{\lambda}^*)$ denote the corresponding neighborhood for each stationary $\boldsymbol{\lambda}^*$.[47] Let $\mathcal{B}$ denote the union of the stable neighborhoods, $\mathcal{B} \equiv \cup_{\boldsymbol{\lambda}^* \in \Lambda(\omega)} B_E(\boldsymbol{\lambda}^*)$, and let $\mathcal{B}_U$ denote the union of the unstable neighborhoods, $\mathcal{B}_U = \cup_{\boldsymbol{\lambda}^* \in \{0,\infty\}^k \backslash \Lambda(\omega)} B_E(\boldsymbol{\lambda}^*)$. We will use these neighborhoods in the proofs of Lemmas 4 to 7.

Proof of Lemma 4 (Global Stability of Agreement). Suppose the agreement outcome is locally stable, $0^k \in \Lambda(\omega)$, and there are at least two types, $|\Theta| \geq 2$. By Assumption 4, a_1 occurs with positive probability, and by Lemma 10, observing (a_1, σ_L) decreases the likelihood ratio. Given initial likelihood ratio $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, let N be the minimum number of consecutive (a_1, σ_L) actions and signals required for the likelihood ratio to reach the stable neighborhood, $\boldsymbol{\lambda}_{N+1} \in B_E(0^k)$. By Lemma 10, the change in the likelihood ratio following (a_1, σ_L) is bounded away from zero. Therefore, $N < \infty$. Let τ_1 be the first time that $\langle\boldsymbol{\lambda}_t\rangle$ enters $B_E(0^k)$, $\tau_1 \equiv \min\{t | \boldsymbol{\lambda}_t \in B_E(0^k)\}$, let τ_2 be the first that $\langle\boldsymbol{\lambda}_t\rangle$ leaves $B_E(0^k)$ after entering, $\tau_2 \equiv \min\{t > \tau_1 | \boldsymbol{\lambda}_t \notin B_E(0^k)\}$, and let τ_3 be the first time the likelihood ratio enters $B_E(0^k)$ and never leaves, $\tau_3 \equiv \min\{\tau | \boldsymbol{\lambda}_t \in B_E(0^k)\ \forall t \geq \tau\}$. We know that $Pr(\tau_1 < \infty) > 0$, since the probability of transitioning from $\boldsymbol{\lambda}_1$ to $B_E(0^k)$ is bounded below by the probability of initially observing N consecutive (a_1, σ_L) action and signal pairs. Also, $Pr(\tau_2 = \infty) > 0$, since by local stability, when the likelihood ratio is in $B_E(0^k)$, with positive probability, it never leaves. Therefore,

[47] In a slight abuse of notation, we switch from the neighborhood subscript denoting the bound for the likelihood ratio to denoting the bound for the log likelihood ratio. This simplifies notation in subsequent lemmas.

$Pr(\tau_3 < \infty) > Pr(\tau_1 < \infty \wedge \tau_2 = \infty) > 0$. Therefore, with positive probability, the likelihood ratio eventually enters and remains in $B_E(0^k)$. By Lemma 3, if the likelihood ratio remains in $B_E(0^k)$ for all t, beliefs almost surely converge to 0^k. Therefore, if $0^k \in \Lambda(\omega)$, then from any initial belief $\boldsymbol{\lambda}_1 \in (0,\infty)^k$, $Pr(\boldsymbol{\lambda}_t \to 0^k) > 0$.

Suppose the agreement outcome is locally stable, $0^k \in \Lambda(\omega)$, and there is a single type, $|\Theta| = 1$. Then Assumption 3.i must hold and public signals are informative. With a single type, action a_1 may occur with probability zero at some beliefs, and we need to adapt the proof for multiple types. Let $\underline{a}(\lambda)$ be the lowest action that type θ_1 plays at belief λ. When there is a single type θ_1, this type has a correctly specified model of the type distribution (this must be the case when $|\Theta| = 1$, as trivially, $\hat{\pi}^1(\theta_1) = 1$), and therefore, observing $\underline{a}(\lambda)$ at belief λ weakly decreases the likelihood ratio (by similar reasoning to Lemma 10). Therefore, observing $(\underline{a}(\lambda), \sigma_L)$ strictly decreases the likelihood ratio, since public signals are informative. Substituting the sequence $(a(\lambda_t), \sigma_L)_{t=1}^N$ for the sequence of N consecutive (a_1, σ_L) actions and signals, where λ_t is the updated belief following $(a(\lambda_{t-1}), \sigma_L)$, the remainder of the proof is the same as in the multiple types case.

The proof for agreement outcome ∞^k is analogous. $\square$

Intermediate Results for Lemma 5. The following definitions and Lemmas 8 and 9 are used in the proof of Lemma 5. Order the public signals by relative likelihood of state R, $(\sigma_1, \sigma_2, ..., \sigma_{|\Sigma|})$ (note $\sigma_1 = \sigma_L$ and $\sigma_{|\Sigma|} = \sigma_R$). Order the action and signal pairs $((a_1, \sigma_1), (a_2, \sigma_1), ..., (a_M, \sigma_{|\Sigma|}))$ so that pair $M(l-1) + m$ corresponds to action a_m and signal σ_l. Define $A(\boldsymbol{\lambda})$ as the matrix of updates to the log likelihood ratio at beliefs $\boldsymbol{\lambda}$, where each row corresponds to the updates for sociable type θ_i, and each column corresponds to the update following action and signal pair j,

$$(A(\boldsymbol{\lambda}))_{ij} \equiv \log \frac{\hat{\psi}_i((a,\sigma)_j | R, \boldsymbol{\lambda})}{\hat{\psi}_i((a,\sigma)_j | L, \boldsymbol{\lambda})}. \tag{17}$$

Without loss of generality, we consider disagreement outcomes that are ordered so that the first $\kappa \in \{1, ..., k-1\}$ types have belief 0 and the remaining $k - \kappa$ types have belief ∞, i.e. $\boldsymbol{\lambda}^* = (0^\kappa, \infty^{k-\kappa})$. To consider other disagreement outcomes, simply reorder the types so that this holds.

Definition 6 (Separability)**.**

1. *Given $\kappa \in \{1, ..., k\}$, $\boldsymbol{\lambda}^* = (0^\kappa, \infty^{k-\kappa})$ is* separable at zero *if there exist vectors $c \in [0,\infty)^{|\mathcal{A}\times\Sigma|}$ and $G \in \mathbb{R}^k$ with $G_i > 0$ for all $i \geq \kappa$ and $G_i < 0$ for all $i < \kappa$, such that $A(\boldsymbol{\lambda}^*)c = G$.*

2. *Given* $\kappa \in \{0, ..., k-1\}$, $\boldsymbol{\lambda}^* = (0^\kappa, \infty^{k-\kappa})$ *is* separable at infinity *if there exist vectors* $c \in (0, \infty)^{|\mathcal{A} \times \Sigma|}$ *and* $G \in \mathbb{R}^k$ *with* $G_i > 0$ *for all* $i > \kappa + 1$ *and* $G_i < 0$ *for all* $i \leq \kappa + 1$*, such that* $A(\boldsymbol{\lambda}^*)c = G$*.*

Definition 7 (Adjacently Accessible)**.** *Given* $\kappa \in \{1, ..., k\}$*,* $\boldsymbol{\lambda}_2^* = (0^{\kappa-1}, \infty^{k-\kappa+1})$ *is* adjacently accessible *from* $\boldsymbol{\lambda}_1^* = (0^\kappa, \infty^{k-\kappa})$ *if for any* $\varepsilon_2 > 0$*, there exists an* $\varepsilon_1 > 0$ *such that for any* $\boldsymbol{\lambda} \in B_{\varepsilon_1}(\boldsymbol{\lambda}_1^*)$*, there exists a* $\tau(\boldsymbol{\lambda}) < \infty$ *such that if* $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}$*, then* $Pr(\boldsymbol{\lambda}_{t+\tau(\boldsymbol{\lambda})} \in B_{\varepsilon_2}(\boldsymbol{\lambda}_2^*)) > 0$*. The definition is analogous for* $\kappa \in \{0, ..., k-1\}$ *and* $\boldsymbol{\lambda}_2^* = (0^{\kappa+1}, \infty^{k-\kappa-1})$ *adjacently accessible from* $\boldsymbol{\lambda}_1^* = (0^\kappa, \infty^{k-\kappa})$*.*

Lemma 8 (Adjacently Accessible)**.** *Given* $\kappa \in \{1, ..., k\}$*, if* $\boldsymbol{\lambda}_1^* = (0^\kappa, \infty^{k-\kappa})$ *is separable at zero, then* $\boldsymbol{\lambda}_2^* = (0^{\kappa-1}, \infty^{k-\kappa+1})$ *is adjacently accessible from* $\boldsymbol{\lambda}_1^*$*. Given* $\kappa \in \{0, ..., k-1\}$*, if* $\boldsymbol{\lambda}_1^*$ *is separable at infinity, then* $\boldsymbol{\lambda}_2^* = (0^{\kappa+1}, \infty^{k-\kappa-1})$ *is adjacently accessible from* $\boldsymbol{\lambda}_1^*$*.*

Proof. Let $\boldsymbol{\lambda}_1^* = (0^\kappa, \infty^{k-\kappa})$, $\boldsymbol{\lambda}_2^* = (0^{\kappa-1}, \infty^{k-\kappa+1})$ and suppose $\boldsymbol{\lambda}_1^*$ is separable at zero. We will show that for any $\varepsilon_2 > 0$, there exists an $\varepsilon_1 > 0$ such that for any $\boldsymbol{\lambda} \in B_{\varepsilon_1}(\boldsymbol{\lambda}_1^*)$, there exists a $\tau(\boldsymbol{\lambda}) < \infty$ such that if $\boldsymbol{\lambda}_1 = \boldsymbol{\lambda}$, then $Pr(\boldsymbol{\lambda}_{1+\tau(\boldsymbol{\lambda})} \in B_{\varepsilon_2}(\boldsymbol{\lambda}_2^*)) > 0$. Since the log likelihood ratio process is linear, this also holds for any $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}$.

For $\varepsilon > 0$, recall $B_\varepsilon(\boldsymbol{\lambda}_1^*) \equiv [0, \varepsilon)^\kappa \times (1/\varepsilon, \infty]^{k-\kappa}$ denotes a neighborhood of $\boldsymbol{\lambda}_1^*$. Define $K(\varepsilon) \equiv -\log \varepsilon$, and let $[-\infty, -K(\varepsilon))^\kappa \times (K(\varepsilon), \infty]^{k-\kappa}$ denote the corresponding neighborhood of $\log \boldsymbol{\lambda}_1^*$. Define

$$g_{\varepsilon,i}(a, \sigma) \equiv \inf_{\boldsymbol{\lambda} \in B_\varepsilon(\boldsymbol{\lambda}_1^*)} \log \frac{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda})},$$

as the smallest update to the log likelihood ratio when type $i \geq \kappa$ observes (a, σ) and has likelihood ratio in the neighborhood $B_\varepsilon(\boldsymbol{\lambda}_1^*)$, and

$$g_{\varepsilon,i}(a, \sigma) \equiv \sup_{\boldsymbol{\lambda} \in B_\varepsilon(\boldsymbol{\lambda}_1^*)} \log \frac{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda})}$$

as the largest update to the log likelihood ratio when type $i < \kappa$ observes (a, σ) and has likelihood ratio in the neighborhood $B_\varepsilon(\boldsymbol{\lambda}_1^*)$. Finally, define

$$\bar{g}_{\varepsilon,\kappa}(a, \sigma) \equiv \sup_{\boldsymbol{\lambda} \in B_\varepsilon(\boldsymbol{\lambda}_1^*)} \log \frac{\hat{\psi}_\kappa(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_\kappa(a, \sigma | L, \boldsymbol{\lambda})}$$

as the largest update to the log likelihood ratio when type κ observes (a, σ) and has likelihood ratio in the neighborhood $B_\varepsilon(\boldsymbol{\lambda}_1^*)$.

We construct a process that bounds $\langle \boldsymbol{\lambda}_t \rangle$ as long as it remains close to $\boldsymbol{\lambda}_1^*$, and use this process to show that we can separate the log likelihood ratios of types $1, ..., \kappa - 1$ and type κ by an arbitrary amount K while the beliefs of all types remain close to $\boldsymbol{\lambda}_1^*$. By separability at zero, there exist vectors $c \in [0, \infty)^k$ and $G \in \mathbb{R}^k$ that satisfy the separability condition. Moreover, since the rationals are dense in the reals, there exists vector $c \in [0, \infty)^k$ of rational numbers and vector $G \in \mathbb{R}^k$ that satisfies the separability condition.

Therefore, there exists an $\varepsilon_3 > 0$ and integers $c_{a,\sigma} \geq 0$ for each $(a, \sigma) \in \mathcal{A} \times \Sigma$ such that

$$G_i \equiv \sum_{(a,\sigma) \in \mathcal{A} \times \Sigma} c_{a,\sigma} g_{\varepsilon_3, i}(a, \sigma), \tag{18}$$

with $G_i > 0$ for all $i \geq \kappa$ and $G_i < 0$ for all $i < \kappa$. Let

$$\overline{G}_\kappa \equiv \sum_{(a,\sigma) \in \mathcal{A} \times \Sigma} c_{a,\sigma} \bar{g}_{\varepsilon_3, \kappa}(a, \sigma). \tag{19}$$

Next we define processes $\xi_{i,t} \equiv \sum_{s=1}^{t-1} g_{\varepsilon_3,i}(a_s, \sigma_s)$ and $\bar{\xi}_{\kappa,t} \equiv \sum_{s=1}^{t-1} \bar{g}_{\varepsilon_3,\kappa}(a_s, \sigma_s)$. Given a sequence with $c_{a,\sigma}$ realizations of each (a, σ), at time $\tau_1 \equiv \sum_{\mathcal{A} \times \Sigma} c_{a,\sigma} + 1$, the process $\xi_{i,\tau_1} = G_i$ by (18) and $\bar{\xi}_{\kappa,\tau_1} = \overline{G}_\kappa$ by (19). For $i \geq \kappa$, $G_i > 0$, and therefore, $\xi_{i,\tau_1} > 0$, while for $i < \kappa$, $G_i < 0$, and therefore, $\xi_{i,\tau_1} < 0$. Moreover, there exists an $\underline{\mathrm{K}} > 0$ such that for all $i > \kappa$, $\xi_{i,t} \geq -\underline{\mathrm{K}}$ for all $t < \tau_1$, and there exists a $\bar{K} > 0$ such that for all $i < \kappa$, $\xi_{i,t} < \bar{K}$ for all $t < \tau_1$. Therefore, for any $K > 0$, there exists an N_K such that following N_K repetitions of the sequence of $c_{a,\sigma}$ realizations of each (a, σ), at time $\tau_K \equiv N_K \sum_{\mathcal{A} \times \Sigma} c_{a,\sigma} + 1$,

1. $\xi_{i,\tau_K} < -K$ for all $i < \kappa$,
2. $\xi_{i,\tau_K} > 0$ for all $i \geq \kappa$,
3. For all $t < \tau_K$, $\xi_{i,t} \leq \bar{K}$ for all $i < \kappa$ and $\xi_{i,t} \geq -\underline{\mathrm{K}}$ for all $i > \kappa$,
4. $\bar{\xi}_{\kappa,t} \leq N_K \overline{G}_\kappa$ for all $t \leq \tau_K$, with equality at $t = \tau_K$.

In summary, following N_K repetitions of the sequence, the processes $\langle \xi_{i,t} \rangle$ of types $i < \kappa$ and type κ are separated by at least K, and at all t during the repetitions, the process of type $i < \kappa$ is bounded above by $\bar{K}$ and the process of type $i > \kappa$ is bounded below by $-\underline{\mathrm{K}}$. As long as $\boldsymbol{\lambda}_s \in B_{\varepsilon_3}(\boldsymbol{\lambda}_1^*)$ for all $s \leq t$, the change in the log likelihood ratio of $i < \kappa$ is bounded above by $\xi_{i,t}$,

$$\log \lambda_{i,t} - \log \lambda_{i,1} \leq \xi_{i,t} \leq \bar{K},$$

the change in the log likelihood ratio of $i = \kappa$ is bounded above by $\bar{\xi}_{\kappa,t}$,

$$\log \lambda_{\kappa,t} - \log \lambda_{\kappa,1} \leq \bar{\xi}_{\kappa,t},$$

and the change in the log likelihood ratio of $i > \kappa$ is bounded below by $\xi_{i,t}$,

$$\log \lambda_{i,t} - \log \lambda_{i,1} \geq \xi_{i,t} \geq -\underline{\mathrm{K}}.$$

Fix $\varepsilon_2 \in (0, \varepsilon_3)$ and $K > \bar{K}$. Choose an ε_1-neighborhood of $\boldsymbol{\lambda}_1^*$ such that $\log \lambda_{i,1} < -K(\varepsilon_2) - \max(\bar{K}, N_K \overline{G}_\kappa)$ for $i \leq \kappa$ and $\log \lambda_{i,1} > K(\varepsilon_2) + \underline{\mathrm{K}}$ for $i > \kappa$. Note $\varepsilon_1 < \varepsilon_2$. Suppose the initial likelihood ratio $\boldsymbol{\lambda}_1 \in B_{\varepsilon_1}(\boldsymbol{\lambda}_1^*)$. We establish local accessibility in three steps.

Step 1. Repeat N_K realizations of the sequence of $c_{a,\sigma}$ realizations of each (a, σ) to separate the log likelihood ratio of types $i < \kappa$ and κ by K. It follows from items (3) and (4) that $\boldsymbol{\lambda}_t$ remains in $B_{\varepsilon_2}(\boldsymbol{\lambda}_1^*)$ for all $t \leq \tau_K$. Therefore, for each i and at all $t \leq \tau_K$, the process $\xi_{i,t}$ bounds the change in the log likelihood ratio, $\log \lambda_{i,t} - \log \lambda_{i,1} \leq \xi_{i,t}$. After N_K realizations of the sequence, $\log \lambda_{i,\tau_K} < -K(\varepsilon_2) - K$ for $i < \kappa$, and $\log \lambda_{i,\tau_K} > K(\varepsilon_2) + \underline{\mathrm{K}}$ for $i > \kappa$.

Step 2. Next, push type κ's log likelihood ratio to $-K(\varepsilon_3)$ as follows. Continue repeating the sequence of $c_{a,\sigma}$ realizations of each (a, σ) until $\log \lambda_{\kappa,t} > -K(\varepsilon_3)$. By construction, the likelihood ratios of all types $i \neq \kappa$ remain in $B_{\varepsilon_2}(\boldsymbol{\lambda}_1^*)$ after every (a, σ) in this sequence, since at any point in the sequence, $\log \lambda_{i,t} < -K(\varepsilon_2) - K + \bar{K}$ for all $i < \kappa$, and $\log \lambda_{i,t} > K(\varepsilon_2)$ for all $i > \kappa$.

Step 3. Finally, push type κ's log likelihood ratio from $-K(\varepsilon_3)$ to $K(\varepsilon_2)$, while keeping the log likelihood ratio of type $i < \kappa$ less than $-K(\varepsilon_2)$. Given ε_2, there exists an $N_2 < \infty$ such that if $\log \lambda_{\kappa,t} \in [-K(\varepsilon_3), K(\varepsilon_2)]$, then following N_2 realizations of (a_M, σ_R), $\log \lambda_{\kappa,t+N_2} > K(\varepsilon_2)$. Let K_2 be the most any type $i < \kappa$'s log likelihood ratio increases after N_2 realizations of (a_M, σ_R) across all beliefs $\boldsymbol{\lambda} \in B_{\varepsilon_2}(\boldsymbol{\lambda}_1^*)$. Recall that when type κ hit the boundary of $-K(\varepsilon_3)$, $\log \lambda_{i,t} < -K(\varepsilon_2) - K + \bar{K}$ for all $i < \kappa$ and $\log \lambda_{i,t} > K(\varepsilon_2)$ for $i > \kappa$. Therefore, after N_2 realizations of (a_M, σ_R), $\log \lambda_{i,t} < -K(\varepsilon_2) - K + \bar{K} + K_2$ for all $i < \kappa$ and $\log \lambda_{i,t} > K(\varepsilon_2)$ for $i > \kappa$. In order to keep $i < \kappa$ in an ε_2-neighborhood of zero after N_2 realizations of (a_M, σ_R), we need to separate beliefs by at least $K = \bar{K} + K_2$. This determines the K we need to use in step 1.

Following these three steps with $K = \bar{K} + K_2$, the likelihood ratio is in neighborhood $B_{\varepsilon_2}(\boldsymbol{\lambda}_2^*)$. Each step required a finite number of actions and signals that occur with positive probability. Therefore, given ε_1 and ε_2 defined above, for any $\boldsymbol{\lambda} \in B_{\varepsilon_1}(\boldsymbol{\lambda}_1^*)$, there exists a $\tau(\boldsymbol{\lambda}) < \infty$ such that if $\boldsymbol{\lambda}_1 = \boldsymbol{\lambda}$, then $Pr(\boldsymbol{\lambda}_{1+\tau(\boldsymbol{\lambda})} \in B_{\varepsilon_2}(\boldsymbol{\lambda}_2^*)) > 0$. The case of $\boldsymbol{\lambda}_1^*$ separable at infinity is analogous. □

Definition 8 (Accessible)**.** *A belief* $\boldsymbol{\lambda}^*$ *is* accessible *if for any initial belief* $\boldsymbol{\lambda}_1 \in (0,\infty)^k$ *and any* $\varepsilon > 0$*, there exists a* $\tau < \infty$ *such that* $Pr(\boldsymbol{\lambda}_\tau \in B_\varepsilon(\boldsymbol{\lambda}^*)) > 0$*.*

Lemma 9 (Accessible Disagreement)**.** *Suppose* $k = 2$*. If* $(0,0)$ *is separable at zero or* (∞,∞) *is separable at infinity, then* $(0,\infty)$ *is accessible.*

Proof. By Lemma 8, if $(0,0)$ is separable at zero, then $(0,\infty)$ is adjacently accessible from $(0,0)$. Fix initial belief $\boldsymbol{\lambda}_1 \in (0,\infty)^2$ and choose $\varepsilon_2 > 0$. Choose $\varepsilon_1 > 0$ such that for any $\boldsymbol{\lambda} \in B_{\varepsilon_1}((0,0))$, there exists a $\tau_2(\boldsymbol{\lambda}) < \infty$ such that if $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}$, then $Pr(\boldsymbol{\lambda}_{t+\tau_2(\boldsymbol{\lambda})} \in B_{\varepsilon_2}((0,\infty))) > 0$. By adjacent accessibility, such an ε_1 exists. By Lemma 4, there exists a finite sequence ξ_1 of N_1 action and signal pairs that occurs with positive probability, such that following ξ_1, $\boldsymbol{\lambda}_{N_1+1} \in B_{\varepsilon_1}((0,0))$. By adjacent accessibility, there exists a finite sequence ξ_2 of N_2 action and signal pairs that occurs with positive probability, such that following sequences ξ_1 and ξ_2, $\boldsymbol{\lambda}_{N_1+N_2+1} \in B_{\varepsilon_2}((0,\infty))$. Since these sequences occur with positive probability, $Pr(\boldsymbol{\lambda}_{N_1+N_2+1} \in B_{\varepsilon_2}((0,\infty))) > 0$, which is the definition of accessible. The case where (∞,∞) is separable at infinity is analogous. □

Proof of Lemma 5 (Global Stability of Disagreement). Suppose $k = 2$, $(0,\infty) \in \Lambda(\omega)$ and $\theta_2 \succ_{(0,0)} \theta_1$. We first show that $\theta_2 \succ_{(0,0)} \theta_1$ implies that $(0,0)$ is separable at zero. Define the submatrix

$$A_{max} \equiv \begin{pmatrix} \log \frac{\hat{\psi}_2(a_1,\sigma_L|R,(0,0))}{\hat{\psi}_2(a_1,\sigma_L|L,(0,0))} & \log \frac{\hat{\psi}_2(a_M,\sigma_R|R,(0,0))}{\hat{\psi}_2(a_M,\sigma_R|L,(0,0))} \\ \log \frac{\hat{\psi}_1(a_1,\sigma_L|R,(0,0))}{\hat{\psi}_1(a_1,\sigma_L|L,(0,0))} & \log \frac{\hat{\psi}_1(a_M,\sigma_R|R,(0,0))}{\hat{\psi}_1(a_M,\sigma_R|L,(0,0))} \end{pmatrix}.$$

Since $\theta_2 \succ_{(0,0)} \theta_1$, this has a positive determinant. Therefore, there exists a $c \in \mathbb{R}^2_+$ that solves

$$A_{max} c = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

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

$$A_{max} \tilde{c} = \begin{pmatrix} G_2 \\ G_1 \end{pmatrix},$$

where $G_1 < 0$ and $G_2 > 0$. Therefore, by Definition 6, $(0,0)$ is separable at zero, since we can set values of c_j to zero for the remaining action and signal pairs in matrix (17). Therefore, by Lemma 9, $(0,\infty)$ is accessible.

We will next show that for any initial belief, $Pr(\boldsymbol{\lambda}_t \to (0,\infty)) > 0$. Fix initial belief $\boldsymbol{\lambda}_1 \in (0,\infty)^2$ and choose $\varepsilon < e^{-E}$. By accessibility, there exists a finite sequence ξ

of N action and signal pairs that occurs with positive probability, such that following sequences ξ, $\boldsymbol{\lambda}_{N+1} \in B_\varepsilon((0,\infty))$. From $(0,\infty) \in \Lambda(\omega)$, $Pr(\boldsymbol{\lambda}_t \to (0,\infty)|\xi) > 0$. Therefore, from any initial belief $\boldsymbol{\lambda}_1 \in (0,\infty)^2$, $Pr(\boldsymbol{\lambda}_t \to (0,\infty)) > 0$, which implies that $(0,\infty)$ is globally stable. The case where $\theta_2 \succ_{(\infty,\infty)} \theta_1$ is analogous, as is the proof for $(\infty, 0)$. $\square$

Proof of Lemma 6 (Unstable Mixed Outcomes). Suppose $k = 2$ and consider a generic misspecified model in which the mixed learning outcome $(0, \theta_1)$ in which θ_1's belief converges to zero and θ_2's belief doesn't converge. Suppose $(0, \theta_1) \notin \Lambda_M(\omega)$, i.e. $(0,0) \in \Lambda_2(\omega)$ or $(0,\infty) \in \Lambda_2(\omega)$. Without loss of generality, consider the case where $(0,0) \in \Lambda_2(\omega)$. Suppose the initial belief for type θ_1 is near zero, $\lambda_{1,1} \in B_\varepsilon(0)$ for any $\varepsilon < e^{-E}$. We want to show that almost surely, either (i) there exists a $\tau < \infty$ such that $\lambda_{1,\tau} \notin B_\varepsilon(0)$; or (ii) $\langle \boldsymbol{\lambda}_t \rangle$ converges for both types. This will establish that almost surely, the mixed outcome does not occur.

We first characterize how the behavior of $\langle \boldsymbol{\lambda}_t \rangle$ near $(0,0)$ and $(0,\infty)$ depends on $\Lambda_1(\omega)$ and $\Lambda_2(\omega)$. Suppose $(0,0) \in \Lambda_1(\omega)$ (recall by assumption, $(0,0) \in \Lambda_2(\omega)$). By the construction in Lemma 3, for $\varepsilon < e^{-E}$, if $\langle \boldsymbol{\lambda}_t \rangle$ enters $B_\varepsilon(0,0)$, with positive probability, $\langle \boldsymbol{\lambda}_t \rangle$ converges to $(0,0)$. If $(0,0) \notin \Lambda_1(\omega)$, then by the construction in Lemma 3, for $\varepsilon < e^{-E}$, if $\langle \boldsymbol{\lambda}_t \rangle$ enters $B_\varepsilon((0,0))$, then from any belief in $B_\varepsilon((0,0))$, (i) with positive probability uniformly bounded away from zero in the starting belief, $\langle \lambda_{1,t} \rangle$ exits $B_\varepsilon(0)$, and (ii) almost surely, $\langle \boldsymbol{\lambda}_t \rangle$ exits $B_\varepsilon((0,0))$. If $(0,\infty) \in \Lambda_2(\omega)$, the behavior of $\langle \boldsymbol{\lambda}_t \rangle$ in a neighborhood of $(0,\infty)$ is similar. If $(0,\infty) \notin \Lambda_2(\omega)$, then by the construction in Lemma 3, for $\varepsilon < e^{-E}$, if $\langle \boldsymbol{\lambda}_t \rangle$ enters $B_\varepsilon((0,\infty))$, then almost surely, $\langle \boldsymbol{\lambda}_t \rangle$ exits $B_\varepsilon((0,\infty))$.

Let $\tau_1 \equiv \min\{t | \lambda_{1,t} \notin B_\varepsilon(0)\}$ be the first time that θ_1's belief leaves a neighborhood of zero. Then it must be that almost surely, $\tau_1 < \infty$ or $\langle \boldsymbol{\lambda}_t \rangle$ visits a neighborhood of $(0,0)$ or $(0,\infty)$ infinitely often,

$$Pr(\tau_1 < \infty \text{ or } \boldsymbol{\lambda}_t \in B_\varepsilon((0,0)) \cup B_\varepsilon((0,\infty)) \text{ i.o.}) = 1. \tag{20}$$

If $(0,0) \notin \Lambda_1(\omega)$, so $(0,0)$ is not locally stable, then λ_2 almost surely leaves $B_\varepsilon((0,0))$, and $Pr(\tau_1 < \infty \text{ or } \boldsymbol{\lambda}_t \in B_\varepsilon((0,\infty)) \text{ i.o.}) = 1$. Similarly, if $(0,\infty) \notin \Lambda(\omega)$, then λ_2 almost surely leaves $B_\varepsilon((0,\infty))$, and $Pr(\tau_1 < \infty \text{ or } \boldsymbol{\lambda}_t \in B_\varepsilon((0,0)) \text{ i.o.}) = 1$.

Case (i): Suppose $(0,0) \in \Lambda_1(\omega)$ or $(0,\infty) \in \Lambda_1(\omega) \cap \Lambda_2(\omega)$. If $\langle \boldsymbol{\lambda}_t \rangle$ enters a neighborhood of a locally stable belief infinitely often, then $\langle \boldsymbol{\lambda}_t \rangle$ almost surely converges for both types. Therefore, almost surely, $\tau_1 < \infty$ or $\langle \boldsymbol{\lambda}_t \rangle$ converges.

Case (ii): Suppose $(0,0) \notin \Lambda_1(\omega)$ and $(0,\infty) \in \Lambda_2(\omega) \setminus \Lambda_1(\omega)$. Each time $\langle \boldsymbol{\lambda}_t \rangle$ enters $B_\varepsilon((0,0)) \cup B_\varepsilon((0,\infty))$, with positive probability uniformly bounded away from zero in the starting belief, $\langle \lambda_{1,t} \rangle$ exits $B_\varepsilon(0)$. Therefore, if $\langle \boldsymbol{\lambda}_t \rangle$ enters $B_\varepsilon((0,0)) \cup B_\varepsilon((0,\infty))$ infinitely often, $\langle \lambda_{1,t} \rangle$ almost surely exits $B_\varepsilon(0)$. Therefore, almost surely $\tau_1 < \infty$.

Case (iii): Suppose $(0,0) \notin \Lambda_1(\omega)$, $(0,\infty) \notin \Lambda_2(\omega)$. Then $Pr(\tau_1 < \infty$ or $\boldsymbol{\lambda}_t \in B_\varepsilon((0,0))$ i.o.$) = 1$. Each time $\langle \boldsymbol{\lambda}_t \rangle$ enters $B_\varepsilon((0,0))$, with positive probability uniformly bounded away from zero in the starting belief, $\langle \lambda_{1,t} \rangle$ exits $B_\varepsilon(0)$. Therefore, if $\langle \boldsymbol{\lambda}_t \rangle$ enters $B_\varepsilon((0,0))$ infinitely often, $\langle \lambda_{1,t} \rangle$ almost surely exits $B_\varepsilon(0)$. Therefore, almost surely $\tau_1 < \infty$.

The proofs for the other mixed outcomes are analogous. □

Proof of Lemma 7 (Belief Convergence). Suppose $k \leq 2$, $\Lambda(\omega)$ contains an agreement outcome or maximally accessible disagreement outcome, and $\Lambda_M(\omega)$ is empty. Recall that $\mathcal{B}$ is the set of locally stable neighborhoods and $\mathcal{B}_U$ is the set of locally unstable neighborhoods. 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 10, there exists a finite sequence of actions and signals 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 Lemma 3, $\langle \boldsymbol{\lambda}_t \rangle$ almost surely leaves $\mathcal{B}_U$. Therefore, $\langle \boldsymbol{\lambda}_t \rangle$ does not converge to a stationary belief that is not locally stable. If $\langle \boldsymbol{\lambda}_t \rangle$ enters the neighborhood of a mixed outcome, by Lemma 6, $\langle \boldsymbol{\lambda}_t \rangle$ almost surely leaves this neighborhood or converges to a locally stable belief. Therefore, mixed learning outcomes almost surely do not arise. By Lemma 2, $\langle \boldsymbol{\lambda}_t \rangle$ does not converge to a non-stationary belief. Therefore, almost surely, either $\langle \boldsymbol{\lambda}_t \rangle$ does not converge for either type or $\langle \boldsymbol{\lambda}_t \rangle$ converges to a learning outcome in $\Lambda(\omega)$. Since $\langle \boldsymbol{\lambda}_t \rangle$ almost surely leaves the neighborhood of any mixed or unstable outcome, it must be that $\langle \boldsymbol{\lambda}_t \rangle$ enters $\mathcal{B}$ infinitely often, $Pr(\boldsymbol{\lambda}_t \in \mathcal{B}$ 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. □

A.1.3 Intermediate Results

The following two lemmas are intermediate results used to prove Lemmas 1 to 7. They hold for any $k \geq 1$. We first define some additional notation. Fix $\boldsymbol{\lambda} \in [0,\infty]^k$. Let $\underline{a}(\boldsymbol{\lambda})$ denote the $\succ$-minimal action played with positive probability at $\boldsymbol{\lambda}$ and $\overline{a}(\boldsymbol{\lambda})$ denote the $\succ$-maximal action played with positive probability at $\boldsymbol{\lambda}$, where $\succ$ is the order from Assumption 2. If $|\Theta| > 1$, by Assumption 4, these actions are constant and equal to a_1 and a_M respectively, i.e. $\underline{a}(\boldsymbol{\lambda}) = a_1$ and $\overline{a}(\boldsymbol{\lambda}) = a_M$.

Lemma 10 (Minimum Informativeness)**.** *For all sociable types θ_i, the minimal update to the likelihood ratio is uniformly bounded above by one and the maximal update to the* *likelihood ratio is uniformly bounded below by one,*

$$\sup_{\boldsymbol{\lambda}\in[0,\infty]^k} \frac{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}),\sigma_L|R,\boldsymbol{\lambda})}{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}),\sigma_L|L,\boldsymbol{\lambda})} < 1 \textit{ and } \inf_{\boldsymbol{\lambda}\in[0,\infty]^k} \frac{\hat{\psi}_i(\overline{a}(\boldsymbol{\lambda}),\sigma_R|R,\boldsymbol{\lambda})}{\hat{\psi}_i(\overline{a}(\boldsymbol{\lambda}),\sigma_R|L,\boldsymbol{\lambda})} > 1.$$

Lemma 10 implies that for any $\boldsymbol{\lambda}_t \in (0,\infty)^k$, when $(\tilde{a}_t,\tilde{\sigma}_t) = (\underline{a}(\boldsymbol{\lambda}_t),\sigma_L)$, then beliefs update toward state L,

$$\lambda_{i,t+1} = \lambda_{i,t}\left(\frac{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}_t),\sigma_L|R,\boldsymbol{\lambda}_t)}{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}_t),\sigma_L|L,\boldsymbol{\lambda}_t)}\right) < \lambda_{i,t}$$

for all sociable types θ_i. Similarly, when $(\tilde{a}_t,\tilde{\sigma}_t) = (\overline{a}(\boldsymbol{\lambda}_t),\sigma_R)$, then beliefs update toward state R, $\boldsymbol{\lambda}_{t+1} > \boldsymbol{\lambda}_t$.

Proof. We first show that for all $\boldsymbol{\lambda} \in [0,\infty]^k$ and sociable types $\theta_i \in \Theta_S$, $\frac{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}),\sigma_L|R,\boldsymbol{\lambda})}{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}),\sigma_L|L,\boldsymbol{\lambda})} \leq 1$. Fix $\boldsymbol{\lambda} \in [0,\infty]^k$ and let $m(\boldsymbol{\lambda})$ denote the index of action $\underline{a}(\boldsymbol{\lambda})$, e.g. if $\underline{a}(\boldsymbol{\lambda}) = a_1$, then $m(\boldsymbol{\lambda}) = 1$. Consider how sociable type θ_i updates its beliefs following $\underline{a}(\boldsymbol{\lambda})$. Since preferences are aligned, any type $\theta_j \in \Theta$ who chooses $\underline{a}(\boldsymbol{\lambda})$ with positive probability at $\boldsymbol{\lambda}$ chooses this action for any signal $s \in \mathcal{S}$ such that $s \leq \overline{s}_{j,m(\boldsymbol{\lambda})}(\lambda_j)$. Type θ_i believes that θ_j plays $\underline{a}(\boldsymbol{\lambda})$ with probability $\hat{F}_i^\omega(\bar{s}_{j,m(\boldsymbol{\lambda})}(\lambda_j))$. By Lemma A.1 in Smith and Sorensen (2000), $F^R(s) \leq F^L(s)$, with strict equality for $s \in \operatorname{int}(\mathcal{S})$. Since signals are aligned, this is also true for the subjective beliefs. Therefore, $\hat{F}_i^R(\bar{s}_{j,m(\boldsymbol{\lambda})}(\lambda_j)) \leq \hat{F}_i^L(\bar{s}_{j,m(\boldsymbol{\lambda})}(\lambda_j))$. This implies $\hat{P}_i(\underline{a}(\boldsymbol{\lambda})|R,\boldsymbol{\lambda}) \leq \hat{P}_i(\underline{a}(\boldsymbol{\lambda})|L,\boldsymbol{\lambda})$, where $\hat{P}_i(a|\omega,\boldsymbol{\lambda})$ denotes type θ_i's subjective probability of action a in state ω at likelihood ratio $\boldsymbol{\lambda}$, since $\hat{P}_i(a|\omega,\boldsymbol{\lambda})$ is a convex combination of the probability each type chooses a in state ω. Public signals are aligned, so it must be that $\hat{\sigma}_i(\sigma_L) \leq 1/2$, as the maximal public signal in state L is either uninformative or indicative of state L. Therefore, $\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}),\sigma_L|R,\boldsymbol{\lambda}) \leq \hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}),\sigma_L|L,\boldsymbol{\lambda})$.

We next establish the uniform bound. Recall that autarkic types have a likelihood ratio that is constant and equal to $\frac{p_0}{1-p_0}$. Sociable type θ_i's update to the likelihood ratio following $\underline{a}(\boldsymbol{\lambda})$ is bounded by

$$\begin{aligned}
\frac{\hat{P}_i(\underline{a}(\boldsymbol{\lambda})|R,\boldsymbol{\lambda})}{\hat{P}_i(\underline{a}(\boldsymbol{\lambda})|L,\boldsymbol{\lambda})} &= \frac{\sum_{\theta_j\in\Theta_A}\hat{\pi}_i(\theta_j)\hat{F}_i^R(\overline{s}_{j,m(\boldsymbol{\lambda})}(\frac{p_0}{1-p_0})) + \sum_{\theta_j\in\Theta_S}\hat{\pi}_i(\theta_j)\hat{F}_i^R(\overline{s}_{j,m(\boldsymbol{\lambda})}(\lambda_j))}{\sum_{\theta_j\in\Theta_A}\hat{\pi}_i(\theta_j)\hat{F}_i^L(\overline{s}_{j,m(\boldsymbol{\lambda})}(\frac{p_0}{1-p_0})) + \sum_{\theta_j\in\Theta_A}\hat{\pi}_i(\theta_j)\hat{F}_i^L(\overline{s}_{j,m(\boldsymbol{\lambda})}(\lambda_j))} \\
&\leq \frac{\sum_{\theta_j\in\Theta_A}\hat{\pi}_i(\theta_j)\hat{F}_i^R(\overline{s}_{j,m(\boldsymbol{\lambda})}(\frac{p_0}{1-p_0})) + \sum_{\theta_j\in\Theta_S}\hat{\pi}_i(\theta_j)\hat{F}_i^L(\overline{s}_{j,m(\boldsymbol{\lambda})}(\lambda_j))}{\sum_{\theta_j\in\Theta_A}\hat{\pi}_i(\theta_j)\hat{F}_i^L(\overline{s}_{j,m(\boldsymbol{\lambda})}(\frac{p_0}{1-p_0})) + \sum_{\theta_j\in\Theta_S}\hat{\pi}_i(\theta_j)\hat{F}_i^L(\overline{s}_{j,m(\boldsymbol{\lambda})}(\lambda_j))} \\
&\leq \frac{\sum_{\theta_j\in\Theta_A}\hat{\pi}_i(\theta_j)\hat{F}_i^R(\overline{s}_{j,m(\boldsymbol{\lambda})}(\frac{p_0}{1-p_0})) + \hat{\pi}_i(\Theta_S)}{\sum_{\theta_j\in\Theta_A}\hat{\pi}_i(\theta_j)\hat{F}_i^L(\overline{s}_{j,m(\boldsymbol{\lambda})}(\frac{p_0}{1-p_0})) + \hat{\pi}_i(\Theta_S)},
\end{aligned}$$

where the first line follows by definition, the second line follows from $\hat{F}_i^R(s) \leq \hat{F}_i^L(s)$, and the third line follows from $\sum_{\theta_j \in \Theta_S} \hat{\pi}_i(\theta_j)\hat{F}_i^L(s) \leq \hat{\pi}_i(\Theta_S)$ and $\hat{F}_i^R(s) \leq \hat{F}_i^L(s)$. Therefore,

$$\frac{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}), \sigma_L|R, \boldsymbol{\lambda})}{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}), \sigma_L|L, \boldsymbol{\lambda})} \leq \left(\frac{\sum_{\theta_j \in \Theta_A} \hat{\pi}_i(\theta_j)\hat{F}_i^R(\overline{s}_{j,m(\boldsymbol{\lambda})}(\frac{p_0}{1-p_0})) + \hat{\pi}_i(\Theta_S)}{\sum_{\theta_j \in \Theta_A} \hat{\pi}_i(\theta_j)\hat{F}_i^L(\overline{s}_{j,m(\boldsymbol{\lambda})}(\frac{p_0}{1-p_0})) + \hat{\pi}_i(\Theta_S)}\right)\left(\frac{\hat{\sigma}_i(\sigma_L)}{1-\hat{\sigma}_i(\sigma_L)}\right). \quad (21)$$

Note that both terms on the right hand side of (21) are less than or equal to one for all $\boldsymbol{\lambda} \in [0, \infty]^k$. First suppose Assumption 3.i holds. Then $\frac{\hat{\sigma}_i(\sigma_L)}{1-\hat{\sigma}_i(\sigma_L)} < 1$, independent of $\boldsymbol{\lambda}$. Next suppose Assumption 3.ii holds. This implies that $\underline{a}(\boldsymbol{\lambda}) = a_1$, $m(\boldsymbol{\lambda}) = 1$ and there exists at least one autarkic type $\theta_j \in \Theta_A$ with $\hat{\pi}_i(\theta_j) > 0$ and $\overline{s}_{j,1}(\frac{p_0}{1-p_0}) \in \text{int}(\mathcal{S})$, which implies $\hat{F}_i^R(\overline{s}_{j,1}(\frac{p_0}{1-p_0})) < \hat{F}_i^L(\overline{s}_{j,1}(\frac{p_0}{1-p_0}))$. Therefore, the first term on the right hand side of (21) is less than one, independent of $\boldsymbol{\lambda}$. Itn both cases, the right hand side of (21) is uniformly bounded away from one and $\sup_{\boldsymbol{\lambda} \in [0,\infty]^k} \frac{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}),\sigma_L|R,\boldsymbol{\lambda})}{\hat{\psi}_i(\underline{a}(\boldsymbol{\lambda}),\sigma_L|L,\boldsymbol{\lambda})} < 1$.

Similar logic holds for the case of $(\overline{a}(\boldsymbol{\lambda}), \sigma_R)$. $\square$

Lemma 11 (Continuity). *$\boldsymbol{\lambda} \mapsto \psi(a, \sigma|\omega, \boldsymbol{\lambda})$ and $\boldsymbol{\lambda} \mapsto \hat{\psi}_i(a, \sigma|\omega, \boldsymbol{\lambda})$ are continuous at each stationary $\boldsymbol{\lambda}^* \in \{0, \infty\}^k$ for all $(a, \sigma) \in \mathcal{A} \times \Sigma$ and $\omega \in \{L, R\}$.*

Proof. Consider $\boldsymbol{\lambda}^* = 0^k$. Each type $\theta_i \in \Theta_S$ has a unique optimal action at 0^k, independent of the realization of the private signal. Moreover, since no action is optimal at a single belief, there exists an $\varepsilon_1 > 0$ such that if the posterior belief following the private signal is in $[0, \varepsilon_1)^k$, each type plays this action. Let Θ_a denote the set of sociable types who play a at 0^k. Fix $\varepsilon > 0$. Let

$$\delta_1 \equiv \min_{a \in \mathcal{A}} \frac{\varepsilon}{\max\{\pi(\Theta_S \setminus \Theta_a), \pi(\Theta_a)\}}$$

and

$$\delta_2 \equiv \min_{a \in \mathcal{A}, \theta_i \in \Theta_S} \frac{\varepsilon}{\max\{\hat{\pi}^i(\Theta_S \setminus \Theta_a), \hat{\pi}^i(\Theta_a)\}}.$$

and $\delta \equiv \min\{\delta_1, \delta_2\}$. Signals are not perfectly informative, so there exists a $\bar{s}$ such that $1 - \hat{F}_i^\omega(\bar{s}) < \delta$ and $1 - \hat{F}_i^\omega(\bar{s}) < \delta$ for all $\theta_i \in \Theta_S$ and $\omega \in \{L, R\}$. Define

$$\varepsilon_1(\delta) \equiv \frac{\varepsilon_1}{\max_{\theta_i \in \Theta_s} \hat{s}_i(\bar{s})/(1 - \hat{s}_i(\bar{s}))}.$$

Fix an action $a \in \mathcal{A}$ and let q_a denote the probability that a type is autarkic and plays action a. If $\boldsymbol{\lambda} \in [0, \varepsilon_1(\delta))$, then the probability of action a in state ω, denoted $P(a|\omega, \boldsymbol{\lambda})$, is bounded above by $\pi(\Theta_a) + \delta\pi(\Theta_S \setminus \Theta_a) + q_a$ and bounded below by $\pi(\Theta_a)(1-\delta) + q_a$. So $|P(a|\omega, \boldsymbol{\lambda}) - P(a|\omega, 0^k)| \leq \varepsilon$ for all $\boldsymbol{\lambda} \in [0, \varepsilon_1(\delta))^k$. Similarly $|\hat{P}_i(a|\omega, \boldsymbol{\lambda}) - \hat{P}_i(a|\omega, 0^k)| \leq \varepsilon$ for all $\boldsymbol{\lambda} \in [0, \varepsilon_1(\delta))^k$ and $\theta_i \in \Theta_S$, where $\hat{P}_i(a|\omega, \boldsymbol{\lambda})$ denotes θ_i's subjective probability of observing action a in state ω at likelihood ratio $\boldsymbol{\lambda}$. The public signal distribution is independent of $\boldsymbol{\lambda}$. Therefore, this continuity extends to $\psi(a,\sigma|\omega,\boldsymbol{\lambda})$ and $\hat{\psi}_i(a,\sigma|\omega,\boldsymbol{\lambda})$ for all $\theta_i \in \Theta_S$. The proof for other stationary beliefs is identical. $\square$

A.2 Proofs of Theorem 2 and Corollary 2

Proof of Theorem 2. Assume Assumption 1, 2, 3 and 4 and consider a generic misspecified model. Suppose $\omega = L$. For any sociable type $\theta_i \in \Theta_S$, the mapping $\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda}) \mapsto \gamma_i(\boldsymbol{\lambda}, L)$ is continuous. By the concavity of the log operator, $\gamma_i(\boldsymbol{\lambda}, L)$ is negative when $||\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda}) - \psi(a,\sigma|L,\boldsymbol{\lambda})|| = 0$. Therefore, there exists a $\delta > 0$ such that if $||\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda}) - \psi(a,\sigma|L,\boldsymbol{\lambda})|| < \delta$ for $(a,\sigma,\boldsymbol{\lambda}) \in \mathcal{A} \times \Sigma \times \{0,\infty\}^k$ and $\theta_i \in \Theta_S$, then $\gamma_i(\boldsymbol{\lambda}, L) < 0$ for all $\boldsymbol{\lambda} \in \{0,\infty\}^k$ and $\theta_i \in \Theta_S$. Therefore, any locally stable point must have $\lambda_i = 0$ for each sociable type. Therefore, 0^k is the unique locally stable point.

We also need to show that $\Lambda_M(L)$ is empty, i.e. all mixed outcomes are reducible. Consider the mixed outcome $\boldsymbol{\lambda}_I^*$ with convergent types I and non-convergent types $N \equiv \Theta_S \setminus I$. For any node $\boldsymbol{\lambda}_N$ in the graph $\mathcal{G}(\boldsymbol{\lambda}_I^*)$ (as defined in Definition 12), it follows from the choice of δ that for each $i \in N$, $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \Lambda_i(L)$ iff $\lambda_i = 0$. Therefore, each $\boldsymbol{\lambda}'_N$ that is mixed accessible from $\boldsymbol{\lambda}_N$ has fewer $i \in N$ with $\lambda'_i = 0$. Therefore, each path terminates at $0^{|N|}$ and the graph has no cycles, i.e. $\boldsymbol{\lambda}_I^*$ is reducible. Therefore, $\Lambda_M(L)$ is empty. By Theorem 1, if $\Lambda(L) = \{0^k\}$ and $\Lambda_M(L)$ is empty, then the likelihood ratio almost surely converges to 0^k and learning is complete.

Similar logic holds for $\omega = R$. $\square$

Proof of Corollary 2. Assume Assumption 1, 2, 3 and 4, and consider a generic misspecified model. Fix state ω. For any sociable type θ_i, the mapping $(\hat{\pi}_i, \hat{F}_i^\omega, \hat{G}_i^\omega) \mapsto \hat{\psi}_i(a,\sigma|\omega,\boldsymbol{\lambda})$ is continuous. By continuity, for any $\delta_2 > 0$, there exists a $\delta > 0$ such that if $||\hat{\pi}_i - \pi|| < \delta$, $||\hat{F}_i^\omega - F^\omega|| < \delta$ and $||\hat{G}_i^\omega - G^\omega|| < \delta$ for all $\theta_i \in \Theta_S$, then $|\hat{\psi}_i(a,\sigma|\omega,\boldsymbol{\lambda}) - \psi(a,\sigma|\omega,\boldsymbol{\lambda})| < \delta_2$ for all $(a,\sigma) \in \mathcal{A} \times \Sigma$, $\boldsymbol{\lambda} \in \{0,\infty\}^k$ and $\theta_i \in \Theta_S$. Choose δ_2 sufficiently small so that Theorem 2 holds. $\square$

B Derivation of Examples 1 and 2

In this section, we present the analysis for Examples 1 and 2.

B.1 Example 1

From Section 3, the autarkic type θ_2 plays both actions with positive probability and the sociable type θ_1 places positive probability on θ_2, which establishes that Assumption 3.ii holds. Assumption 4 is redundant in a binary action decision problem that satisfies Assumption 3.ii, since Assumption 3.ii guarantees that the sociable type believes that the autarkic type plays both actions with positive probability.

From the action probabilities derived in Section 3, 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

$$\psi(L|\omega, \lambda_1) = \pi(\theta_1)F^\omega((1/(1+\lambda_1))^{1/\nu}) + \pi(\theta_2)F^\omega(.5^{1/\nu}),$$

where in a slight abuse of notation we suppress the dependence of $\hat{\psi}_1$ and ψ on σ since the public signal is uninformative in this example. The construction of $\gamma_1(0, L)$ and $\gamma_1(\infty, L)$ in Section 4.2 follows from evaluating these expressions at $\lambda_1 = 0$ and $\lambda_1 = \infty$, respectively.

We next characterize how $\Lambda(\omega)$ depends on ν. To capture its explicit dependence on ν, let $\gamma_1^\nu(\boldsymbol{\lambda}, \omega)$ correspond to the function $\gamma_1(\boldsymbol{\lambda}, \omega)$ and $\Lambda^\nu(\omega)$ correspond to the set $\Lambda(\omega)$ in the model with partisan bias level ν. 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 $\infty \in \Lambda^\nu(L)$, we need to determine the sign of

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

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, $\gamma_1^\nu(\infty, L)$ is strictly decreasing in ν. At $\nu = 1$, $\gamma_1^1(\infty, L) < 0$ by the concavity of the log operator. At $\nu = 0$, θ_2 chooses action R for all signals, $\alpha_0 = 0$. Therefore, $\gamma_1^0(\infty, L) = \log\frac{1-\pi_A(1-\alpha_1)}{1-\pi_A\alpha_1} > 0$. This establishes that there exists a cutoff $\nu_1 \in (0, 1)$ such that for $\nu < \nu_1$, $\gamma_1^\nu(\infty, L) > 0$ and $\infty \in \Lambda^\nu(L)$ and for $\nu > \nu_1$, $\gamma_1^\nu(\infty, L) < 0$ and $\infty \notin \Lambda^\nu(L)$.

To determine whether correct learning arises, i.e. whether $0 \in \Lambda^\nu(L)$, we need to determine the sign of

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

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, $\gamma_1^\nu(0,L)$ is strictly decreasing in ν. At $\nu = 1$, $\gamma_1^1(0,L) < 0$ by the concavity of the log operator. At $\nu = 0$, θ_2 chooses action R for all signals, $\alpha_0 = 0$. Therefore,

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

This establishes that there exists a cutoff $\nu_2 \in (0,1)$ such that for $\nu < \nu_2$, $\gamma_1^\nu(0,L) > 0$ and $0 \notin \Lambda^\nu(L)$ and for $\nu > \nu_2$, $\gamma_1^\nu(0,L) < 0$ and $0 \in \Lambda^\nu(L)$.

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

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

and by the symmetry of the signal distributions, $\gamma_1^\nu(0,L) - \gamma_1^1(0,L) = \gamma_1^\nu(\infty,L) - \gamma_1^1(\infty,L)$. Moreover $\gamma_1^1(0,L) - \gamma_1^1(\infty,L)$ 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, $0 \notin \Lambda^\nu(\omega)$ before $\infty \in \Lambda^\nu(\omega)$. This implies that $\Lambda^\nu(L) = \{\infty\}$ for $\nu \in (0,\nu_1)$, $\Lambda^\nu(L) = \emptyset$ for $\nu \in (\nu_1,\nu_2)$, and $\Lambda^\nu(L) = \{0\}$ for $\nu \in (\nu_2, 1]$.

Next consider $\omega = R$. Then $\gamma_1^1(\infty,R) > 0$ and $\gamma_1^1(0,R) > 0$, 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: $\gamma_1^\nu(0,R)$ and $\gamma_1^\nu(\infty,R)$ are decreasing in ν. Therefore, $\gamma_1^\nu(0,R) > 0$ implies $0 \notin \Lambda^\nu(R)$ for all $\nu \in (0,1]$. Similarly, $\gamma_1^\nu(\infty,R) > 0$ implies $\infty \in \Lambda^\nu(R)$ for all $\nu \in (0,1]$. Therefore, $\Lambda^\nu(R) = \{\infty\}$ for all $\nu \in (0,1]$.

When there is a single sociable type, mixed learning and disagreement are trivially not possible. By Theorem 1, the characterization of the locally stable set fully determines asymptotic learning outcomes. This leads to the following proposition.

Proposition 7 (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.*

The proof follows immediately from the construction of $\Lambda^\nu(\omega)$ above.

B.2 Example 2

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 sociable type while θ_4 is an autarkic type.[48] 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 sociable 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 public signal is uninformative, the private signal distribution is informative and symmetric, $F^R(s) = 1 - F^L(1 - s)$, and the 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 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 $s < 1/(1 + \lambda) = \bar{s}_{i,1}(\lambda)$. 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 $F^\omega(.5)$. 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

$$\hat{\psi}_1(L|\omega, \boldsymbol{\lambda}) = (1 - \pi_A)F^\omega(1/(1 + \lambda_1)) + \pi_A F^\omega(.5),$$

where again we suppress the dependence of $\hat{\psi}_i$ and ψ on σ since the public signal is uninformative in this example. This type misspecification leads the partisan type to

[48] In a slight abuse of 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 sociable types.

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 sociable 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, $\lambda_{1,t} = \lambda_{3,t}$.[49]

Given that the two likelihood ratios move in unison, we can consider the partisan and nonpartisan sociable 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 8 (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.*

[49] 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.

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.

The construction of the locally stable set is similar to Example 1. We present it in Online Appendix F.1.

C Additional Analysis and Proofs from Section 5

C.1 Section 5.1 (Underreaction and Overreaction)

Proof of Proposition 1. Suppose public signals are informative, private signals are uninformative, and there is a single type θ_1 that interprets public signals according to $\frac{\hat{\sigma}(\sigma)}{1-\hat{\sigma}(\sigma)} = \left(\frac{\sigma}{1-\sigma}\right)^{\nu}$, correctly interprets private signals as uninformative and correctly believes that all agents are type θ_1. Let $\gamma_C(\boldsymbol{\lambda}, \omega) \equiv \sum_{\sigma \in \Sigma} dG^{\omega}(\sigma) \log\left(\frac{\sigma}{1-\sigma}\right)$ denote the expected change in the log likelihood ratio in a correctly specified model. Since correct learning occurs almost surely in the correctly specified model, $\gamma_C(0, L) < 0$ and $\gamma_C(\infty, L) < 0$ in state L and $\gamma_C(0, R) > 0$ and $\gamma_C(\infty, R) > 0$ in state R. Then

$$\gamma(\boldsymbol{\lambda}, \omega) = \sum_{\sigma \in \Sigma} dG^{\omega}(\sigma) \log\left(\frac{\sigma}{1-\sigma}\right)^{\nu} = \nu \sum_{\sigma \in \Sigma} dG^{\omega}(\sigma) \log\left(\frac{\sigma}{1-\sigma}\right) = \nu \gamma_C(\boldsymbol{\lambda}, \omega).$$

Therefore, $\gamma(\boldsymbol{\lambda}, \omega)$ has the same sign as $\gamma_C(\boldsymbol{\lambda}, \omega)$, which implies $\Lambda(L) = \{0\}$ and $\Lambda(R) = \{\infty\}$. Trivially, $\Lambda_M(\omega) = \emptyset$ since there is a single type. By Theorem 1, correct learning occurs almost surely, independent of ν. □

Representation of Epstein et al. (2010). Consider the form of under- and over-reaction specified in Epstein et al. (2010) with constant weights given to the Bayesian update and the prior belief. That is, an agent with current belief p who observes signal σ that leads to true posterior $p' \equiv \frac{p\sigma}{p\sigma + (1-p)(1-\sigma)}$ updates to posterior belief

$$\hat{p}' \equiv \tilde{\nu}\left(\frac{p\sigma}{p\sigma + (1-p)(1-\sigma)}\right) + (1-\tilde{\nu})p, \tag{22}$$

where $\tilde{\nu} \in [0, 1)$ corresponds to underreaction, $\tilde{\nu} > 1$ corresponds to overreaction, and $\tilde{\nu} = 1$ corresponds to the correctly specified model.

This form of signal misspecification can be represented in our framework. Let $\phi(\sigma, p) \equiv p\sigma + (1-p)(1-\sigma)$. Then $\hat{p}' = \frac{\tilde{\nu} p\sigma}{\phi(\sigma,p)} + (1-\tilde{\nu})p$ and $1 - \hat{p}' = \frac{\tilde{\nu}(1-p)(1-\sigma)}{\phi(\sigma,p)} + (1 -$

$\tilde{\nu})(1-p)$. Therefore,

$$\frac{\hat{p}'}{1-\hat{p}'} \quad = \quad \left(\frac{\frac{\tilde{\nu}\sigma}{\phi(\sigma,p)}+1-\tilde{\nu}}{\frac{\tilde{\nu}(1-\sigma)}{\phi(\sigma,p)}+1-\tilde{\nu}}\right)\left(\frac{p}{1-p}\right).$$

To capture this form of under- and overreaction, the signal misspecification needs to depend on the current belief about the state. As discussed in Section 2.3, this is a straightforward extension (see Online Appendix E.4). Let $\hat{\sigma}(\sigma,p)$ denote the extension of the signal misspecification $\hat{\sigma}$ to depend on current belief p. Then

$$\frac{\hat{\sigma}(\sigma,p)}{1-\hat{\sigma}(\sigma,p)} = \frac{\tilde{\nu}\sigma+(1-\tilde{\nu})\phi(\sigma,p)}{\tilde{\nu}(1-\sigma)+(1-\tilde{\nu})\phi(\sigma,p)}$$

yields the representation. This form of signal misspecification satisfies the assumptions of our framework. Therefore, we can use Theorem 1 to determine how the set of learning outcomes depends on misspecification parameter $\tilde{\nu}$.

C.2 Section 5.2 (Confirmation Bias)

Set-up. In Rabin and Schrag (1999), signals are probabilistically misinterpreted. To allow for this, we extend our framework to allow the signal misspecification to map two signals that induce the same true posterior to different misspecified posteriors. It is straightforward to extend all results in our paper to this case. We also allow the severity of confirmation bias to vary with the current belief about the state.

To model probabilistic misinterpretation in our framework, suppose there are four public signals, $y \in \{l_1, l_2, r_1, r_2\}$, where signals $y \in \{l_1, l_2\}$ induce the same posterior, $\sigma_l \equiv \sigma(l_1) = \sigma(l_2)$ and signals $y \in \{r_1, r_2\}$ induce the same posterior, $\sigma_r \equiv \sigma(r_1) = \sigma(r_2)$. Conditional on an l_1 or l_2 signal, the signal is l_2 with probability q and conditional on an r_1 or r_2 signal, the signal is r_2 with probability q. Signals l_2 and r_2 correspond to the signals that are misread, and signals l_1 and r_1 correspond to the signals that are correctly read. Let $\hat{\sigma}(y,p)$ denote the subjective posterior belief for signal y at belief p. The confirmation bias we outlined in Section 5.2 corresponds to

$$\hat{\sigma}(y,p) = \begin{cases} (1-\nu(p))\sigma_l + \nu(p)\sigma_r & \text{if } y = l_2 \text{ and } p > 1/2 \\ (1-\nu(p))\sigma_r + \nu(p)\sigma_l & \text{if } y = r_2 \text{ and } p < 1/2 \\ \sigma(y) & \text{otherwise,} \end{cases} \tag{23}$$

where $\nu : [0,1] \to [0,1]$ is a continuous function.

To complete the model, assume that public signals are informative, $\sigma(r) > 1/2$, and private signals are uninformative and believed to be uninformative.

Proof of Proposition 2. Suppose public signals are informative, private signals are uninformative, and there is a single type θ_1 that interprets public signals according to (23), correctly interprets private signals as uninformative and correctly believes that all agents are type θ_1. Let $\tilde{\nu} \equiv \nu(0)$ denote the slant at certainty (recall $\nu(1) = \nu(0)$ by assumption).

We first characterize the locally stable set $\Lambda(L)$. The local stability of correct learning is determined by the sign of

$$\begin{aligned}\gamma(0,L) \quad &= \quad Pr(\{r_1,r_2\}|L)\left((1-q)\log\left(\frac{\sigma_r}{1-\sigma_r}\right)+q\log\left(\frac{(1-\tilde{\nu})\sigma_r+\tilde{\nu}\sigma_l}{(1-\tilde{\nu})(1-\sigma_r)+\tilde{\nu}(1-\sigma_l)}\right)\right)\\ &\quad +Pr(\{l_1,l_2\}|L)\log\left(\frac{\sigma_l}{1-\sigma_l}\right),\end{aligned}$$

where $P(\{l_1,l_2\}|L) = \frac{1-\sigma_r/(1-\sigma_r)}{\sigma_l/(1-\sigma_l)-\sigma_r/(1-\sigma_r)}$ and $P(\{r_1,r_2\}|L) = 1-P(\{l_1,l_2\}|L)$. At $q=0$, agents have a correctly specified model, so $\gamma(0,L)<0$. As q increases, more weight is placed on the second term and less weight is placed on the first term in the above equation. The second term is less than the first term; therefore, $\gamma(0,L)$ is decreasing in q. Therefore, for all q and $\nu(0)$, $\gamma(0,L)<0$ and correct learning is locally stable, $0 \in \Lambda(L)$.

The local stability of incorrect learning is determined by the sign of

$$\begin{aligned}\gamma(\infty,L) \quad &= \quad Pr(\{l_1,l_2\}|L)\left((1-q)\log\left(\frac{\sigma_l}{1-\sigma_l}\right)+q\log\left(\frac{(1-\tilde{\nu})\sigma_l+\tilde{\nu}\sigma_r}{(1-\tilde{\nu})(1-\sigma_l)+\tilde{\nu}(1-\sigma_r)}\right)\right)\\ &\quad +Pr(\{r_1,r_2\}|L)\log\left(\frac{\sigma_r}{1-\sigma_r}\right).\end{aligned}$$

At $q=0$ or $\tilde{\nu}=0$, agents have a correctly specified model, so $\gamma(\infty,L)<0$. As q increases, more weight is placed on the second term and less weight is placed on the first term in the above equation. The second term is greater than the first term; therefore, $\gamma(\infty,L)$ is increasing in q. Similarly, the second term is increasing in $\tilde{\nu}$, and therefore, so is $\gamma(\infty,L)$. At $q=1$ and $\tilde{\nu}=1$,

$$\gamma(\infty,L) = \log\left(\frac{\sigma_r}{1-\sigma_r}\right) > 0. \tag{24}$$

Therefore, the desired cutoffs $\overline{q} \in (0,1)$ and $\overline{\nu}(q) \in (0,1)$ exist such that incorrect learning is locally stable for $q > \overline{q}$ and $\tilde{\nu} > \overline{\nu}(q)$.

When signals are symmetric and fully slanted at certainty, i.e. $\sigma \equiv \sigma_r = 1-\sigma_l$ and

$\nu(1) = 1$,

$$\gamma(\infty, L) = (2\sigma q + 1 - 2\sigma)) \log\left(\frac{\sigma}{1-\sigma}\right). \tag{25}$$

Given $\log\left(\frac{\sigma}{1-\sigma}\right) > 0$, this expression is positive when $q > 1 - 1/2\sigma$, which is the cut-off in Rabin and Schrag (1999).

The construction of $\Lambda(R)$ is analogous. Given that there is a single type, mixed learning does not arise. Therefore, $\Lambda(\omega)$ fully characterizes the set of asymptotic learning outcomes. □

C.3 Section 5.3 (Naive Learning)

Set-up.

Proof of Proposition 4. 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 $\gamma_i(\boldsymbol{\lambda}, \omega; \hat{\pi}_i)$ 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

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

In the correctly specified model, $\gamma_i((0,0), L; \pi) < 0$, and this expression is decreasing in $\hat{\pi}_i$. Therefore, $\gamma_i((0,0), L; \hat{\pi}_i) < 0$ 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

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

This expression is increasing in $\hat{\pi}_i$ and is equivalent to the representative agent model at $\hat{\pi}_i = \hat{\pi}$. Therefore, if $\gamma_i((\infty,\infty), L; \hat{\pi}) < 0$, then $\gamma_1((\infty,\infty), L; \hat{\pi}_1) < 0$ 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 $\hat{\pi}_i > \overline{\pi}$, $\gamma_i((\infty,\infty), L; \hat{\pi}_i) > 0$, 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

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

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, $\gamma_i((0,\infty), L; \hat{\pi}_i) < 0$ 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 $\gamma_i((\infty, 0), L; \hat{\pi}_i) = \gamma_i((0,\infty), L; \hat{\pi}_i)$, 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 $\gamma_1((0,\infty), L; \hat{\pi}_1) < 0$ for any $\hat{\pi}_1$ and $\gamma_2((\infty, 0), L; \hat{\pi}_2) < 0$ 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 1, $\Lambda(\omega; \hat{\pi}_1, \hat{\pi}_2)$ 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 3. This result follows directly from the constructions of $\gamma_i(\boldsymbol{\lambda}, \omega; \hat{\pi}_i)$ in Proposition 4. Generically, $\gamma_i((0,0), \omega; \hat{\pi}_i) \neq 0$ and $\gamma_i((\infty,\infty), \omega; \hat{\pi}_i) \neq 0$ for $i = 1, 2$. Given an average bias $\hat{\pi}$, consider the case where $\gamma_i((0,0), \omega; \hat{\pi}) \neq 0$ and $\gamma_i((\infty,\infty), \omega; \hat{\pi}) \neq 0$ for $i = 1, 2$. For any $\delta > 0$, there exists an ε such that for $|\hat{\pi}_1 - \hat{\pi}| < \varepsilon/2$ and $|\hat{\pi}_2 - \hat{\pi}| < \varepsilon/2$, $|\gamma_i(\boldsymbol{\lambda}, \omega; \hat{\pi}_i) - \gamma_i(\boldsymbol{\lambda}, \omega; \hat{\pi})| < \delta$ for $\boldsymbol{\lambda} \in \{(0,0), (\infty,\infty)\}$ and $i = 1, 2$. Choosing δ small enough ensures that $\gamma_i(\boldsymbol{\lambda}, \omega; \hat{\pi}_i)$ and $\gamma_i(\boldsymbol{\lambda}, \omega; \hat{\pi})$ 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. $\square$

C.4 Section 5.4 (Using the Framework to Demonstrate Behavioral Equivalence)

Set-up. As in Example 2, consider a setting with a binary action space, $\mathcal{A} = \{L, R\}$ and four types of agents: θ_1 and θ_3 are sociable types, and θ_2 and θ_4 are autarkic types.[50] All types seek to choose the action that matches the state, but differ in whether they are tempted to choose action R when there is uncertainty. Types θ_1 and θ_2 are tempted: their preferences place higher weight on matching the state in state R, $u_i(a, \omega) = (2 - \beta)\mathbb{1}_{a=\omega=L} + \beta\mathbb{1}_{a=\omega=R}$ for some $\beta \in [1, 2)$ and $\theta_i \in \{\theta_1, \theta_2\}$. Types θ_3 and θ_4 are not tempted: their preferences are symmetric across states, $u_i(a, \omega) = \mathbb{1}_{a=\omega}$. Assume that an equal and positive share of tempted and non-tempted 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 sociable types have correct beliefs about the share of autarkic types, but believe that no agents are tempted, $\hat{\pi}_1(\theta_3) = \pi(\theta_1) + \pi(\theta_3)$ and $\hat{\pi}_1(\theta_4) = \pi(\theta_2) + \pi(\theta_4)$ and similarly for $\hat{\pi}_3$. Let $q \equiv \pi(\theta_3) + \pi(\theta_4)$ denote the share of non-tempted types and $\pi_A \equiv \pi(\theta_2) + \pi(\theta_4)$ denote the share of autarkic types. To close the model, assume that the public signal is uninformative, the private signal distribution is informative and symmetric, $F^R(s) = 1 - F^L(1 - s)$, all types have a correct subjective signal distribution, and the prior $p_0 = 1/2$. Note that preferences are aligned since tempted and non-tempted types order actions in the same way (Assumption 2).

Proposition 5 in Section 5.4 establishes that there is an equivalence between this setting with preference misspecification and the signal misspecification in Example 2.

Proof of Proposition 5 Each type θ_i updates to belief $\frac{p_i(\lambda,s)}{1-p_i(\lambda,s)} = \lambda\left(\frac{s}{1-s}\right)$ following likelihood ratio λ and private signal s. Types θ_1 and θ_2 choose action L if this belief is less than $(2-\beta)/\beta$, which is equivalent to $s < (2-\beta)/(2+\beta\lambda-\beta)$. Therefore, $\overline{s}_{1,1}(\lambda) = \overline{s}_{2,1}(\lambda) = (2-\beta)/(2+\beta\lambda-\beta)$. Types θ_3 and θ_4 choose action L if this belief is less than one, which is equivalent to $s < 1/(1+\lambda)$. Therefore, $\overline{s}_{3,1}(\lambda) = \overline{s}_{4,1}(\lambda) = 1/(1+\lambda)$.

At likelihood ratio λ_1, type θ_1 chooses L with probability $F^\omega((2-\beta)/(2+\beta\lambda_3-\beta))$. Type θ_2 is autarkic. Therefore, its likelihood ratio is constant at $\lambda_2 = 1$ and it chooses action L with probability $F^\omega(1-\beta/2)$. 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 $F^\omega(.5)$. Both sociable types θ_1 and θ_3 have correct beliefs about the probability that θ_3 and θ_4 choose action L, but neither accounts for the presence of θ_1 and θ_2. Therefore, at likelihood ratio $\boldsymbol{\lambda}$, type θ_1 believes action L occurs with probability

$$\hat{\psi}_1(L|\omega, \boldsymbol{\lambda}) = (1-\pi_A)F^\omega(1/(1+\lambda_3)) + \pi_A F^\omega(.5),$$

[50]For ease of comparison with Example 2, in a slight abuse of notation we continue to let θ_2 denote an autarkic type, which violates our convention that the first k types are the sociable types.

where again in a slight abuse of notation we suppress the dependence of $\hat{\psi}_i$ and ψ on σ since the public signal is uninformative. Since θ_3 has an identical model as θ_1, its subjective action probability is equivalent, $\hat{\psi}_3(L|\omega, \boldsymbol{\lambda}) = \hat{\psi}_1(L|\omega, \boldsymbol{\lambda})$. The true probability of action L is

$$\begin{aligned}\psi(L|\omega, \boldsymbol{\lambda}) &= (1-q)((1-\pi_A)F^\omega((2-\beta)/(2+\beta\lambda_1-\beta)) + \pi_A F^\omega(1-\beta/2)) \\ &\quad + q((1-\pi_A)F^\omega(1/(1+\lambda_3)) + \pi_A F^\omega(.5)).\end{aligned}$$

From these subjective probabilities, the construction of $\gamma_1^{\beta,q}((0,0),\omega)$ and $\gamma_1^{\beta,q}((\infty,\infty),\omega)$ are similar to Proposition 8, replacing α_ν with $\alpha_\beta \equiv F^L(1-\beta/2)$ and noting that α_β is strictly decreasing in β.

Define $\nu(\beta) \equiv \log .5/\log(1-\beta/2)$. Then α_β in the model with level of temptation β and share of non-tempted types q is equal to $\alpha_{\nu(\beta)}$ in the model in Example 2 with level of partisan bias $\nu(\beta)$ and share of non-partisan types q. Therefore, the models have identical locally stable sets, and hence, an identical set of asymptotic learning outcomes. $\square$

C.5 Section 5.5 (Level-k)

Let $\boldsymbol{\lambda} = (\lambda_2, \lambda_3)$ denote the vector of likelihood ratios for the sociable types θ_2 and θ_3.

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 $\bar{s}_{1,1}(\lambda_1) = 0.5$, type θ_1 chooses action L with probability $F^\omega(0.5)$ and action R with probability $1 - F^\omega(0.5)$. 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 $\hat{\psi}_2(R|\omega, \boldsymbol{\lambda}) = 1 - F^\omega(0.5)$, where in a slight abuse of notation, we suppress the dependence of $\hat{\psi}_i$ on the public signal since it is uninformative. 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,

$$\hat{\psi}_3(L|\omega, \boldsymbol{\lambda}) = (1-\varepsilon)F^\omega(1/(\lambda_2+1)) + \varepsilon F^\omega(.5). \tag{26}$$

The subjective probability of an R action is analogous. Finally, the *true* probability of an L action depends on the correct distribution over types,

$$\psi(L|\omega, \boldsymbol{\lambda}) = \pi(\theta_1)F^\omega(.5) + \pi(\theta_2)F^\omega(1/(\lambda_2+1)) + \pi(\theta_3)F^\omega(1/(\lambda_3+1)). \tag{27}$$

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,

$$\begin{aligned}\gamma_3((0,0),L) &= (\pi(\theta_1)\alpha_L + \pi(\theta_2) + \pi(\theta_3)) \log\left(\frac{1-\varepsilon+\varepsilon\alpha_R}{1-\varepsilon+\varepsilon\alpha_L}\right) \\ &+ \pi(\theta_1)(1-\alpha_L)\log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) \\ &\approx \pi(\theta_1)(1-\alpha_L)\log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) > 0\end{aligned}$$

and correct learning is not locally stable for the level-3 type, $(0,0) \notin \Lambda_3(L)$. Similarly, for small ε,

$$\gamma_3((\infty,\infty),L) \approx \pi(\theta_1)\alpha_L \log\left(\frac{\alpha_R}{\alpha_L}\right) < 0$$

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,

$$\gamma_3((0,\infty),L) \approx (\pi(\theta_1)(1-\alpha_L) + \pi(\theta_3))\log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) > 0.$$

and $(0,\infty) \in \Lambda_3(L)$. Similarly,

$$\gamma_3((\infty,0),L) \approx (\pi(\theta_1)\alpha_L + \pi(\theta_3))\log\left(\frac{\alpha_R}{\alpha_L}\right) < 0$$

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,

$$\begin{aligned}\gamma_2((0,\infty), L) =&(\pi(\theta_1)\alpha_L + \pi(\theta_2)) \log\left(\frac{\alpha_R}{\alpha_L}\right) \\ &+ (\pi(\theta_1)(1-\alpha_L) + \pi(\theta_3)) \log\left(\frac{1-\alpha_R}{1-\alpha_L}\right)\end{aligned}$$

and

$$\begin{aligned}\gamma_2((\infty,0), L) =&(\pi(\theta_1)\alpha_L + \pi(\theta_3)) \log\left(\frac{\alpha_R}{\alpha_L}\right) \\ &+ (\pi(\theta_1)(1-\alpha_L) + \pi(\theta_2)) \log\left(\frac{1-\alpha_R}{1-\alpha_L}\right).\end{aligned}$$

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 $\pi \mapsto \gamma_2((0,\infty), L)$ and $\pi \mapsto \gamma_2((\infty, 0), L)$ are continuous.

1. As $\pi(\theta_3) \to 0$,

$$\begin{aligned}\gamma_2((0,\infty), L) \to \quad & (\pi(\theta_1)\alpha_L + 1 - \pi(\theta_1)) \log\left(\frac{\alpha_R}{\alpha_L}\right) \\ + \quad & \pi(\theta_1)(1-\alpha_L) \log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) < 0\end{aligned}$$

   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 $\pi(\theta_3) \to 1$,

$$\begin{aligned}\gamma_2((0,\infty), L) &\to \log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) > 0 \\ \gamma_2((\infty,0), L) &\to \log\left(\frac{\alpha_R}{\alpha_L}\right) < 0.\end{aligned}$$

   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 $\pi(\theta_2) \to 0$,

$$\begin{aligned} \gamma_2((\infty,0),L) \to \quad & (\pi(\theta_1)\alpha_L + 1 - \pi(\theta_1)) \log\left(\frac{\alpha_R}{\alpha_L}\right) \\ + \quad & \pi(\theta_1)(1-\alpha_L)\log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) < 0, \end{aligned}$$

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 $\pi(\theta_2) \to 1$,

$$\begin{aligned} \gamma_2((0,\infty),L) &\to \log\left(\frac{\alpha_R}{\alpha_L}\right) < 0 \\ \gamma_2((\infty,0),L) &\to \log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) > 0. \end{aligned}$$

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 $\pi(\theta_1) \to 1$,

$$\begin{aligned} \gamma_2((0,\infty),L) &\to \alpha_L \log\frac{\alpha_R}{\alpha_L} + (1-\alpha_L)\log\frac{1-\alpha_R}{1-\alpha_L} < 0 \\ \gamma_2((\infty,0),L) &\to \alpha_L \log\frac{\alpha_R}{\alpha_L} + (1-\alpha_L)\log\frac{1-\alpha_R}{1-\alpha_L} < 0. \end{aligned}$$

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 $\pi(\theta_2)$, $\gamma_2((0,\infty),L)$ 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 $\gamma_2((0,\infty),L) = 0$. As shown above, at $\pi(\theta_2) = 0$ and $\pi(\theta_3) = 0$, $\gamma_2((0,\infty),L) < 0$ and at $\pi(\theta_2) = 0$ and $\pi(\theta_3) = 1$, $\gamma_2((0,\infty),L) > 0$. Therefore, there exists a cutoff $c_9 \in (0,1)$ such that at $\pi(\theta_2) = 0$ and $\pi(\theta_3) = c_9$, $\gamma_2((0,\infty),L) = 0$. Similarly, there exists 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 $\pi(\theta_3) = c_{10}$, $\gamma_2((0,\infty),L) = 0$. Given $\gamma_2((0,\infty),L)$ is linear in $\pi(\theta_2)$ and $\pi(\theta_3)$, the solution to $\gamma_2((0,\infty),L) = 0$ 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 $\pi(\theta_2)$, $\gamma_2((\infty,0),L)$ 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 $\gamma_2((\infty,0),L) = 0$. Given $\gamma_2((\infty,0),L)$ is linear in $\pi(\theta_2)$ and $\pi(\theta_3)$, the solution to $\gamma_2((\infty,0),L) = 0$ 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 $\gamma_2((\infty,0),L) = 0$ 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 $\gamma_2((0,\infty),L) = 0$ lies above the solution to $\gamma_2((\infty,0),L) = 0$. 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 $\Lambda(\omega)$.

Lemma 12. *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.*

Proof. 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,

$$\frac{\hat{\psi}_2(R|R,(0,0))}{\hat{\psi}_2(R|L,(0,0))} = \frac{\hat{\psi}_3(R|R,(0,0))}{\hat{\psi}_3(R|L,(0,0))} = \frac{1-\alpha_R}{1-\alpha_L}. \tag{28}$$

Therefore, $\theta_3 \succeq_{(0,0)} \theta_2$. From Definition 4, 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,

$$\frac{\hat{\psi}_2(L|R,(\infty,\infty))}{\hat{\psi}_2(L|L,(\infty,\infty))} = \frac{\hat{\psi}_3(L|R,(\infty,\infty))}{\hat{\psi}_3(L|L,(\infty,\infty))} = \frac{\alpha_R}{\alpha_L}. \tag{29}$$

Therefore, $\theta_2 \succeq_{(\infty,\infty)} \theta_3$. From Definition 4, 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,

$$\alpha_L \log\left(\frac{\alpha_R}{\alpha_L}\right) + (1-\alpha_L)\log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) < 0.$$

Therefore, since $\frac{\alpha_R}{\alpha_L} < 0$,

$$\begin{aligned}\gamma_2((0,0),L) &= (\pi_1\alpha_L + \pi(\theta_2) + \pi(\theta_3))\log\left(\frac{\alpha_R}{\alpha_L}\right) \\ &+ \pi_1(1-\alpha_L)\log\left(\frac{1-\alpha_R}{1-\alpha_L}\right) < 0.\end{aligned}$$

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 $\gamma_2((\infty,\infty),L) < 0$ and $\gamma_2((0,\infty),L) > 0$. However, $\gamma_2((\lambda_2,\infty),L)$ 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.

Proof of Proposition 6. As $\varepsilon \to 1$, $\Lambda(\omega) \subseteq \{(0,\infty),(\infty,0)\}$. Given $\Lambda_M(\omega) = \emptyset$ and by Lemma 12, both disagreement outcomes are maximally accessible, by Theorem 1, $\Lambda(\omega)$ 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 $\Lambda(\omega) = \{(0,\infty),(\infty,0)\}$.

□

Online Appendix for "Social Learning with Model Misspecification: A Framework and a Characterization"

J. Aislinn Bohren and Daniel N. Hauser

D Posterior Representation.

Let $\mathcal{Z}$ be a signal space. Let $\mu^\omega \in \Delta(\mathcal{Z})$ and $\hat{\mu}^\omega \in \Delta(\mathcal{Z})$ be probability measures on $\mathcal{Z}$ in state ω. Assume μ^L, μ^R and $\hat{\mu}^L, \hat{\mu}^R$ are mutually absolutely continuous. Let $r(z) \equiv 1/(1 + \frac{d\mu^L}{d\mu^R}(z))$ and $\hat{r}(z) \equiv 1/(1 + \frac{d\hat{\mu}^L}{d\hat{\mu}^R}(z))$ denote the posterior belief that the state is L under each set of measures. The c.d.f.s $F_r^\omega(x) \equiv \mu^\omega(z|r(z) \leq x)$ and $\hat{F}_{\hat{r}}^\omega(x) \equiv \hat{\mu}^\omega(z|\hat{r}(z) \leq x)$ are the distributions of the posterior beliefs $r(z)$ and $\hat{r}(z)$ under measures μ^ω and $\hat{\mu}^\omega$, respectively. Given these two measures, we can also define the distribution of $\hat{r}(z)$ under measure μ^ω as $F_{\hat{r}}^\omega(x) \equiv \mu^\omega(z|\hat{r}(z) \leq x)$, and the distribution of $r(z)$ under measure $\hat{\mu}^\omega$ as $\hat{F}_r^\omega(x) \equiv \hat{\mu}^\omega(z|r(z) \leq x)$.

As in the text, we say that two pairs of distributions are aligned if they induce the same ordinal ranking over signals, in terms of the posterior belief. Pairs of distributions are *equality-preserving* if, for all sets of signals that map into a given posterior belief r under one pair of distributions, these signals also map into the same posterior belief $\hat{r}$ under the other pair of distributions (which may differ from r).[51]

Definition 9. (Equality-Preserving Signals). *Mutually absolutely continuous probability measures $\mu^L, \mu^R \in \Delta(\mathcal{Z})^2$ and $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})^2$ are* equality-preserving *if for any $z, z' \in \operatorname{supp} \mu \cap \operatorname{supp} \hat{\mu}$, $\frac{d\mu^L}{d\mu^R}(z) = \frac{d\mu^L}{d\mu^R}(z')$ if and only if $\frac{d\mu^L}{d\mu^R}(z) = \frac{d\mu^L}{d\mu^R}(z')$.*

Multiple signals z and z' can lead to the same posterior belief. Therefore, two pairs of distributions can map to the same distribution over posterior beliefs. The following property describes an equivalence class of probability measures. These measures have the same ordinal ranking of signals and the same distribution over posterior beliefs.

Definition 10 (Equivalent Measures). *Measures μ^L, μ^R and $\hat{\mu}^L, \hat{\mu}^R$ are* equivalent *iff they are aligned, equality-preserving, $\operatorname{supp} \mu = \operatorname{supp} \hat{\mu}$, and $\mu^\omega(z|1/(1 + \frac{d\mu^L}{d\mu^R}(z)) \leq x) = \hat{\mu}^\omega(z|1/(1 + \frac{d\hat{\mu}^L}{d\hat{\mu}^R}(z)) \leq x)$ for all $x \in [0, 1]$ and $\omega \in \{L, R\}$.*

Lemma 13 establishes that when pairs of probability measures $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})^2$ and $\mu^L, \mu^R \in \Delta(\mathcal{Z})^2$ are aligned and equality-preserving, there is a unique representation of $\hat{\mu}^L, \hat{\mu}^R$ as $(\hat{s}, \hat{F}_r^R)$, where $\hat{s} : \operatorname{supp} F_r \to [0, 1]$ is a strictly increasing function mapping the posterior $r(z)$ to the posterior $\hat{r}(z)$ and $\hat{F}_r^R$ is the distribution of $r(z)$ under measure $\hat{\mu}^R$.

[51] If the signal distributions are aligned and equality-preserving, then the mapping from $z \mapsto \frac{\hat{\mu}^L}{\hat{\mu}^R}(z)$ is an order preserving mapping from $(\operatorname{supp} \mu \cap \operatorname{supp} \hat{\mu}, \succeq)$ to $(\mathbb{R}, \geq)$, where $z \simeq z'$ if $\frac{\mu^L}{\mu^R}(z) \geq \frac{\mu^L}{\mu^R}(z')$.

Lemma 13. *Suppose μ^L, μ^R have full support and signals are informative, $\frac{d\mu^R}{d\mu^L}(z) \neq 1$.*

1. *For any mutually absolutely continuous probability measures $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})^2$ that have full support and are equality-preserving and aligned with μ^L, μ^R, there exists a unique $(\hat{s}, \hat{F}_r^R)$, where $\hat{s} : \operatorname{supp} F_r \to [0,1]$ is a strictly increasing function with $\hat{s}(\inf \operatorname{supp} F_r) < 1/2$ and $\hat{s}(\sup \operatorname{supp} F_r) > 1/2$, such that $\hat{s}(r(z)) = 1/(1 + \frac{d\hat{\mu}^L}{d\hat{\mu}^R}(z))$ for all $z \in \mathcal{Z}$ and $\hat{F}_r^R$ is the distribution of $r(z)$ under measure $\hat{\mu}^R$.*

2. *For any strictly increasing function $\hat{s} : \operatorname{supp} F_r \to [0,1]$ and any c.d.f. $\hat{F}_r^R$ with $\operatorname{supp} \hat{F}_r^R = \operatorname{supp} F_r$ and $\int_0^1 \left(\frac{1-\hat{s}(s)}{\hat{s}(s)}\right) d\hat{F}_r^R = 1$, there exist unique (up to an equivalent pair of measures) mutually absolutely continuous probability measures $(\hat{\mu}^L, \hat{\mu}^R) \in \Delta(\mathcal{Z})^2$ that have full support, are equality-preserving and aligned with μ^L, μ^R, and satisfy $\hat{s}(r(z)) = 1/(1 + \frac{d\hat{\mu}^L}{d\hat{\mu}^R}(z))$ for all $z \in \mathcal{Z}$. The measures $\hat{\mu}^L, \hat{\mu}^R$ are equality-preserving and aligned with μ^L, μ^R.*[52]

3. *For any strictly increasing function $\hat{s} : \operatorname{supp} F_r \to [0,1]$ such that $\hat{s}(\inf \operatorname{supp} F_r) < 1/2$ and $\hat{s}(\sup \operatorname{supp} F_r) > 1/2$, there exists mutually absolutely continuous probability measures $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})^2$ that have full support, are equality-preserving and aligned with μ^L, μ^R, and satisfy $\hat{s}(r(z)) = 1/(1 + \frac{d\hat{\mu}^L}{d\hat{\mu}^R}(z))$ for all $z \in \mathcal{Z}$.*

The first part of Lemma 13 implies that $F_{\hat{r}}^{\omega}(\hat{s}(s)) = F_r^{\omega}(s)$ for all $s \in \operatorname{supp}(F_r)$ and $\operatorname{supp} F_{\hat{r}} = \hat{s}(\operatorname{supp} F_r)$. Similarly, $\hat{F}_{\hat{r}}^{\omega}(\hat{s}(s)) = \hat{F}_r^{\omega}(s)$ for all $s \in \operatorname{supp} \hat{F}_r$ and $\operatorname{supp} \hat{F}_{\hat{r}} = \hat{s}(\operatorname{supp} \hat{F}_r)$.

Proof. First establish part (i). Let $(\hat{\mu}^L, \hat{\mu}^R) \in \Delta(\mathcal{Z})^2$ be mutually absolutely continuous probability measures that have full support and are equality-preserving and aligned with (μ^L, μ^R). Define the mapping $\hat{s} : \operatorname{supp} F_r \to [0,1]$ as $\hat{s}(r(z)) = \hat{r}(z)$, where $\hat{r}(z) \equiv 1/(1 + \frac{d\hat{\mu}^L}{d\hat{\mu}^R}(z))$. This is a function since if $r(z) = r(z')$, then $\hat{r}(z) = \hat{r}(z')$, which establishes existence. For any z such that $r(z) > r(z')$, $\hat{r}(z) = \hat{s}(r(z)) > \hat{r}(z') = \hat{s}(r(z'))$ since $(\hat{\mu}^L, \hat{\mu}^R)$ is equality-preserving and aligned. Therefore, $\hat{s}$ is strictly increasing on $\operatorname{supp} F_r$.

By the Bayesian constraint, it must be that $E_{\hat{\mu}}[\hat{r}(z)] = 1/2$, where the expectation is taken with respect to $(\hat{\mu}^L, \hat{\mu}^R)$. Given that (μ^L, μ^R) are informative, equality-preserving and aligned with $(\hat{\mu}^L, \hat{\mu}^R)$, it cannot be that $\hat{r}(z) = 1/2$ for all $z \in \mathcal{Z}$. Therefore, there exist $z, z' \in \mathcal{Z}$ such that $\hat{r}(z) > 1/2$ and $\hat{r}(z') < 1/2$, which implies that there exist $s, s' \in \operatorname{supp} F_r$ such that $\hat{s}(s) > 1/2$ and $\hat{s}(s') < 1/2$. Given that $\hat{s}$ is strictly increasing in s, it immediately follows that $\hat{s}(\inf \operatorname{supp} F_r) < 1/2$ and $\hat{s}(\sup \operatorname{supp} F_r) > 1/2$. Define

[52] Note that if $\hat{F}_r^R$ is a c.d.f. and $\int_0^1 \left(\frac{1-\hat{s}(s)}{\hat{s}(s)}\right) d\hat{F}_r^R = 1$, then it must be that $\hat{s}(\sup(\operatorname{supp} F_r)) > 1/2$ and $\hat{s}(\inf \operatorname{supp} F_r) < 1/2$.

$\hat{F}_r^R(x) \equiv \hat{\mu}^L(z|r(z) \leq x)$. Then $\hat{F}_r^R$ is the distribution of s under measure $\hat{\mu}^R$. Given $\{\hat{s}, \hat{F}_r^R\}$, $\hat{F}_r^L$ is uniquely pinned down by

$$\hat{F}_r^L(x) = \int_0^x \left(\frac{1 - \hat{(}s)}{\hat{(}s)} \right) d\hat{F}_r^R(s)$$

for any $x \in \operatorname{supp} F_r$.

Next, show part (ii). Let $\hat{s} : \operatorname{supp} F_r \to [0, 1]$ be a strictly increasing function and let c.d.f. $\hat{F}_r^R$ be the distribution of $r(z)$ under measure $\hat{\mu}^R$, with $\operatorname{supp} \hat{F}_r^R = \operatorname{supp} F_r$ and $\int_0^1 \left(\frac{1-\hat{s}(s)}{\hat{s}(s)} \right) d\hat{F}_r^R = 1$. By Lemma A.1 in Smith and Sorensen (2000), the distribution of $r(z)$ under measure $\hat{\mu}^L$ is uniquely determined by

$$\hat{F}_r^L(x) = \int_0^x \left(\frac{1 - \hat{s}(s)}{\hat{s}(s)} \right) d\hat{F}_r^R(s).$$

Since $\hat{F}_r^L$ has Radon-Nikodym derivative $\frac{1-\hat{s}(s)}{\hat{s}(s)}$, it induces posterior belief $\hat{s}(s)$ after observing signal z from set of signals $Z = \{z|r(z) = s\}$ that lead to correctly specified posterior s, for any $s \in \operatorname{supp} F_r$. If any other distribution induced the same posterior beliefs, then it would also have Radon-Nikodym derivative $\frac{1-\hat{s}(s)}{\hat{s}(s)}$, so it would be equivalent to $\hat{F}_r^L$. Since $\frac{1-\hat{s}(s)}{\hat{s}(s)} > 0$ and $\hat{F}_r^L(1) = 1$, $\hat{F}_r^L$ is a probability distribution.

Define the random variable $S = r(z)$. $\hat{F}_r^\omega$ defines a probability measure over this random variable in state ω. For any measurable set $A \subseteq \mathcal{Z}$, define

$$\hat{\mu}^\omega(A) = \int E(\mathbb{1}_A|S) d\hat{F}_r^\omega,$$

where E is the conditional expectation defined with respect to μ^R. By the uniqueness and additivity of conditional expectation, for any disjoint, measurable sets $A, B \subseteq \mathcal{Z}$,

$$\hat{\mu}^\omega(A \cup B) = \int E(\mathbb{1}_{A\cup B}|S) d\hat{F}_r^\omega = \int (E(\mathbb{1}_A|S) + E(\mathbb{1}_B|S)) d\hat{F}_r^\omega = \hat{\mu}^\omega(A) + \hat{\mu}^\omega(B),$$

so $\hat{\mu}^\omega$ is a measure. For any set A, if $\hat{\mu}^L(A) = 0$, then $\hat{\mu}^R(A) = 0$ and vice versa, since the integrand used to define $\hat{\mu}^R$ is strictly positive. Therefore, the distributions $(\hat{\mu}^L, \hat{\mu}^R)$ are mutually absolutely continuous with common support $\operatorname{supp} \hat{\mu}$. Also, $\operatorname{supp} \hat{\mu} = \operatorname{supp} \mu$ by construction, so the measures have full support on $\mathcal{Z}$. Moreover, since F_r^ω is unique, $\hat{\mu}^\omega$ is unique up to the probability measure that is used to evaluate $E(\cdot|S)$. For any measurable set $A \subseteq \mathcal{Z}$,

$$\hat{\mu}^L(A) = \int E(\mathbb{1}_A|S) \left(\frac{1 - \hat{s}(S)}{\hat{s}(S)} \right) d\hat{F}_r^R = \int_A \left(\frac{1 - \hat{s}(r(z))}{\hat{s}(r(z))} \right) d\hat{\mu}^R(z),$$

where the first equality follows from the definition of $\hat{F}_r^L$ and the second equality follows from the definition of $\hat{\mu}^R$, so these distributions induce the correct posterior beliefs. Finally, $\hat{\mu}^R(\mathcal{Z}) = \int_0^1 d\hat{F}_r^R(s) = 1$ and $\hat{\mu}^L(\mathcal{Z}) = \int_0^1 d\hat{F}_r^L(s) = 1$, so these are indeed probability measures.

Finally, show part (iii). Suppose $\hat{s} : \operatorname{supp} F_r \to [0,1]$ is a strictly increasing function with $\hat{s}(\inf \operatorname{supp} F_r) < 1/2$ and $\hat{s}(\sup \operatorname{supp} F_r) > 1/2$. Fix any distribution $\hat{F}$ with support $\operatorname{supp} F_r \cap \{s|\hat{s}(s) < 1/2\}$. Then $\int_0^1 \left(\frac{1-\hat{s}(s)}{\hat{s}(s)}\right) d\hat{F}_r(s) < 1$. Similarly, fix a distribution $\hat{G}$ with support $\operatorname{supp} F_r \cap \{s|\hat{s}(s) \geq 1/2\}$. Then $\int_0^1 \left(\frac{1-\hat{s}(s)}{\hat{s}(s)}\right) d\hat{G}(s) > 1$. For any $\lambda \in [0,1]$, let G_λ be the distribution of the compound lottery $G_\lambda = \lambda\hat{F} + (1-\lambda)\hat{G}$. This lottery draws signals from $\hat{F}$ with probability λ and $\hat{G}$ with probability $(1-\lambda)$. The function $H(\lambda) \equiv \int \left(\frac{1-\hat{s}(s)}{\hat{s}(s)}\right) dG_\lambda$ is a continuous mapping from $[0,1]$ to $\mathbb{R}$, so by the intermediate value theorem, there exists a $\lambda^* \in (0,1)$ such that $\int \left(\frac{1-\hat{s}(s)}{\hat{s}(s)}\right) dG_{\lambda^*} = 1$. Let $\hat{F}^R = G_{\lambda^*}$. Then $\hat{F}^R$ is a probability distribution, since it is the convex combination of two distributions. By construction, $\operatorname{supp} \hat{F}_r^R = \operatorname{supp} F_r$ and $\int_0^1 \left(\frac{1-\hat{s}(s)}{\hat{s}(s)}\right) d\hat{F}_r^R = 1$. Therefore, from part(ii), it is possible to construct the desired probability measures $(\hat{\mu}^L, \hat{\mu}^R)$. □

E Additional Results from Section 4

E.1 Complete Learning for Correctly Specified Types

The following result establishes that misspecified types do not interfere with the learning of correctly specified types.

Theorem 3. *Given Assumption 1, 2, 3 and 4, learning is complete for all correctly specified types.*

Proof. Assume Assumption 1, 2, 3 and 4, and suppose $\omega = L$. Under these assumptions, if $\langle\lambda_{i,t}\rangle$ converges for any type θ_i, then the support of the limit belief λ_∞ is a subset of $\{0, \infty\}$, i.e. $\operatorname{supp}(\lambda_\infty) \subset \{0, \infty\}$. Let θ_i be a correctly specified type. Then its subjective probability of each action is equal to the true probability, $\hat{\psi}_i = \psi$. Therefore, given $\omega = L$, $\langle\lambda_{i,t}\rangle$ is a martingale for any learning environment $\{\Theta, \pi, F^L, F^R\}$. By the Martingale Convergence Theorem, $\langle\lambda_{i,t}\rangle$ converges almost surely to a limit random variable λ_∞ with $\operatorname{supp}(\lambda_\infty) \subset [0, \infty)$. This rules out incorrect and cyclical learning. Therefore, zero is the only candidate limit belief for the correctly specified type, $\operatorname{supp}(\lambda_\infty) = \{0\}$, and it must be that $\lambda_{i,t} \to 0$ almost surely. □

E.2 Analogue of Theorem 1 for $k > 2$.

This section proves the analogue of Theorem 1 for more than two sociable types, $k > 2$. The statement of the result is identical to Theorem 1, using the modified definitions of $\Lambda(\omega)$ (defined in (8)), $\Lambda_M(\omega)$ and maximal accessibility (defined for $k > 2$ below). Recall that Lemmas 1 to 4 hold for all $k \geq 1$. Therefore, we prove analogues of Lemmas 5 to 7.

Global Stability of Disagreement. As above, without loss of generality, order a disagreement outcome so that the first κ types have belief zero, and the remaining $k-\kappa$ types have belief infinity, i.e. $\boldsymbol{\lambda} = (0^{\kappa}, \infty^{k-\kappa})$. As in the case of two sociable types, we can use the maximal action and signal pairs to define a sufficient condition for global stability, and use this to prove an analogue of Lemma 5.

Definition 5′ (Maximal Accessibility)**.** *Disagreement outcome* $(0^{\kappa}, \infty^{k-\kappa})$ *is* maximally accessible *if either:*

(i) for all $\kappa' = 0, ..., \kappa-1$*, given* $\boldsymbol{\lambda} = (0^{\kappa'}, \infty^{k-\kappa'})$*,* $\theta_i \succ_{\boldsymbol{\lambda}} \theta_{\kappa'+1}$ *for all* $i > \kappa'+1$ *and* $\theta_{\kappa'+1} \succeq_{\boldsymbol{\lambda}} \theta_i$ *for all* $i < \kappa'+1$*;*

(ii) for all $\kappa' = \kappa+1, ..., k$*, given* $\boldsymbol{\lambda} = (0^{\kappa'}, \infty^{k-\kappa'})$*,* $\theta_i \succeq_{\boldsymbol{\lambda}} \theta_{\kappa'}$ *for all* $i > \kappa'$ *and* $\theta_{\kappa'} \succ_{\boldsymbol{\lambda}} \theta_i$ *for all* $i < \kappa'$*,*

where $\succ_{\boldsymbol{\lambda}}$ *is the maximal R-order defined in Definition 4.*

Note that this definition is equivalent to Definition 5 for the case of $k=2$. If the belief of the type with $\lambda_i^* = 0$ that interprets maximal action and signal pairs as the weakest evidence of state L decreases at a faster rate than the belief of the type with $\lambda_j^* = \infty$ that interprets maximal action and signal pairs as the strongest evidence of state L, then it is possible to find a finite sequence of maximal action and public signal pairs that separate beliefs. Once again, this condition is straightforward to verify from the primitives of the model. As in the case of $k=2$, for any disagreement outcome in $\Lambda(\omega)$, maximal accessibility is a sufficient condition for global stability. Given this revised definition of maximal accessibility, the statement of Lemma 5 is identical.[53]

Lemma 5′ (Global Stability of Disagreement)**.** *If disagreement outcome* $\boldsymbol{\lambda}^* = (0^{\kappa}, \infty^{k-\kappa})$ *is locally stable and maximally accessible, then* $\boldsymbol{\lambda}^*$ *is globally stable.*

Mixed Learning Outcomes. Consider the mixed outcome in which beliefs converge to $\boldsymbol{\lambda}_I^* \in \{0, \infty\}^{|I|}$ for some subset of sociable types $I \subset \Theta_S$, where $\boldsymbol{\lambda}_I^*$ denotes the likelihood ratio vector restricted to set I, and beliefs do not converge for the remaining sociable types $N \equiv \Theta_S \setminus I$. As in the case with two types, denote this mixed outcome as the pair $(\boldsymbol{\lambda}_I^*, I)$, where $I \subseteq \Theta$ is the set of sociable types whose beliefs converge and $\boldsymbol{\lambda}_I^*$ is the vector of limit beliefs these types. For example, a mixed outcome where θ_1 and θ_2's beliefs converge to zero and θ_3's beliefs do not converge is denoted $((0,0), \{\theta_1, \theta_2\})$. This outcome is not locally stable if it is possible for the beliefs of the non-convergent types to

[53] The proof of Lemma 5′ shows that Definition 5′ is sufficient for a much weaker, but more complicated to verify, condition called separability (Definition 6) that utilizes the entire set of actions to separate beliefs.

converge. For example, suppose there are three sociable types. If $(0,0,0) \in \Lambda_1(\omega)$, then the mixed learning outcome in which θ_1 has cyclical learning, and θ_2 and θ_3 have correct learning is not locally stable, since if the beliefs of θ_2 and θ_3 converge to zero, then the beliefs of θ_1 will also almost surely converge to zero. For mixed learning outcomes in which two or more types have cyclical learning, the argument is more involved. To rule out mixed learning, 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, if $(0,0,0) \in \Lambda_1(\omega) \cap \Lambda_2(\omega)$, then to rule out the mixed learning outcome in which θ_1 and θ_2 have cyclical learning and θ_3 has correct learning, we need to show that for $\boldsymbol{\lambda} \in \{(0,\infty,0),(\infty,0,0),(\infty,\infty,0)\}$, either (i) beliefs will almost surely enter a neighborhood of $(0,0,0)$ from a neighborhood of $\boldsymbol{\lambda}$, or (ii) $\boldsymbol{\lambda} \in \Lambda_1(\omega) \cap \Lambda_2(\omega)$. The following definition formalizes this notion.

Definition 11 (Mixed Accessible)**.** *Given mixed outcome $(\boldsymbol{\lambda}_I^*, I)$ for $I \subset \Theta_S$ and non-convergent types $N \equiv \Theta_S \setminus I$, $(\boldsymbol{\lambda}'_N, N)$ is* mixed accessible *from $(\boldsymbol{\lambda}_N, N)$ if $(\boldsymbol{\lambda}'_N, N) \neq (\boldsymbol{\lambda}_N, N)$ and $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \notin \Lambda_i(\omega)$ for each $i \in N$ such that $\lambda'_i \neq \lambda_i$, and $(\boldsymbol{\lambda}'_N, N)$ is* strongly mixed accessible *from $(\boldsymbol{\lambda}_N, N)$ if $(\boldsymbol{\lambda}'_N, N)$ is mixed accessible from $(\boldsymbol{\lambda}_N, N)$ and for each distinct $i, j \in N$ with $\lambda_i \neq \lambda'_i$ and $\lambda_j \neq \lambda'_j$, then $\lambda_i = \lambda_j$.*

Given mixed outcome $(\boldsymbol{\lambda}_I^*, I)$, we construct a graph $\mathcal{G}(\boldsymbol{\lambda}_I^*, I)$ with nodes $(\boldsymbol{\lambda}_N, N)$ such that $\boldsymbol{\lambda}_N \in \{0,\infty\}^{|N|}$ and $N \equiv \Theta_S \setminus I$ to represent which nodes are mixed accessible from other nodes for the non-convergent types.

Definition 12 (Accessible Graph)**.** *Given $(\boldsymbol{\lambda}_I^*, I)$ and $N \equiv \Theta_S \setminus I$, define the directed graph $\mathcal{G}(\boldsymbol{\lambda}_I^*, I)$ with nodes $(\boldsymbol{\lambda}_N, N)$ such that $\boldsymbol{\lambda}_N \in \{0,\infty\}^{|N|}$ as follows: there is an edge from $(\boldsymbol{\lambda}_N, N)$ to $(\boldsymbol{\lambda}'_N, N)$ iff $(\boldsymbol{\lambda}'_N, N)$ is strongly mixed accessible from $(\boldsymbol{\lambda}_N, N)$.*

A *terminal node* $(\boldsymbol{\lambda}_N, N)$ is a node with no edges leaving it.

Definition 13 (Reducible)**.** *A mixed outcome $(\boldsymbol{\lambda}_I^*, I)$ is* reducible *if $\mathcal{G}(\boldsymbol{\lambda}_I^*, I)$ has no cycles.*

If a mixed outcome is reducible, then conditional on the convergent types I remaining in a neighborhood of $\boldsymbol{\lambda}_I^*$, almost surely, the beliefs of the non-convergent types converge. This is a contradiction. Therefore, almost surely, this mixed outcome will not arise. Let $\Lambda_M(\omega)$ denote the set of mixed learning outcomes that are not reducible,

$$\Lambda_M(\omega) \equiv \{(\boldsymbol{\lambda}_I^*, I) | \boldsymbol{\lambda}_I^* \in \{0,\infty\}^{|I|},\ I \subset \Theta_S,\ (\boldsymbol{\lambda}_I^*, I) \text{ is not reducible}\}. \tag{30}$$

As in the case of two sociable types, if a mixed learning outcome arises with positive probability, it must be in $\Lambda_M(\omega)$.

Given the modified definition of $\Lambda_M(\omega)$, the statement of Lemma 6 is identical.[54]

Lemma 6′ (Unstable Mixed Outcomes)**.** *Given $I \subset \Theta_S$ and $N \equiv \Theta_S \setminus I$, if mixed learning outcome $(\boldsymbol{\lambda}_I^*, I) \notin \Lambda_M(\omega)$, then $Pr(\boldsymbol{\lambda}_{I,t} \to \boldsymbol{\lambda}_I^*$ and $\boldsymbol{\lambda}_{N,t}$ does not converge$) = 0$.*

Reducibility is always satisfied in some important cases and is relatively straightforward to verify. For instance, it is satisfied in models near the correctly specified model, in which $\gamma_i(\boldsymbol{\lambda}, \omega) < 0$ at all stationary $\boldsymbol{\lambda} \in \{0, \infty\}^k$ for all sociable types θ_i. In this model, each node in the graph is connected to all nodes with fewer ∞'s than it and is connected to no other nodes. Therefore, for any mixed outcome $(\boldsymbol{\lambda}_I^*, I)$ and set of non-convergent types $N = \Theta_S \setminus I$, every path in the graph 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 λ_i moves towards zero in expectation for all $i \in \Theta_S$, and therefore, some $i \in I$'s beliefs must eventually exit a neighborhood of $\boldsymbol{\lambda}_I^*$.

Belief Convergence. Finally, if there is at least one globally stable agreement or maximally accessible disagreement outcome and no locally stable mixed outcomes, then the likelihood ratio converges almost surely for all types. Given the modified definitions of $\Lambda(\omega)$, $\Lambda_M(\omega)$ and maximal accessibility, the statement of Lemma 7 is identical.

Lemma 7′ (Belief Convergence)**.** *Consider a generic misspecified model. Suppose $\Lambda(\omega)$ contains an agreement outcome or maximally accessible disagreement outcome and $\Lambda_M(\omega)$ is empty. 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.*

E.2.1 Proofs of Lemmas 5′ to 7′.

Using the definition of accessible (Definition 8), the analogue of Lemma 9 for $k > 2$ is as follows. We use this result in the proof of Lemma 5′.

Lemma 9′ (Accessible Disagreement)**.** *Disagreement outcome $\boldsymbol{\lambda}_J^* = (0^\kappa, \infty^{k-\kappa})$ is accessible if there exists a sequence of stationary likelihood ratios $\boldsymbol{\lambda}_1^*, \boldsymbol{\lambda}_2^* \ldots \boldsymbol{\lambda}_J^*$, with $\boldsymbol{\lambda}_1^* \in \{0^k, \infty^k\}$ and $\boldsymbol{\lambda}_j^*$ adjacently accessible from $\boldsymbol{\lambda}_{j-1}^*$ for $j = 2, ..., J$.*

Proof. The proof follows almost directly from Lemma 8. Each element of the sequence $\boldsymbol{\lambda}_j^*$ is adjacently accessible from the previous element of the sequence $\boldsymbol{\lambda}_{j-1}^*$. Starting with $\boldsymbol{\lambda}_J^*$ and any $\varepsilon_J > 0$, there exists an $\varepsilon_{J-1} > 0$ and $\tau_J < \infty$ such that if $\boldsymbol{\lambda}_t \in B_{\varepsilon_{J-1}}(\boldsymbol{\lambda}_{J-1}^*)$, then $Pr(\boldsymbol{\lambda}_{t+\tau_J} \in B_{\varepsilon_J}(\boldsymbol{\lambda}_J^*)) > 0$. Iterating back to $\boldsymbol{\lambda}_1^*$, for any $\varepsilon_J > 0$, there exists an $\varepsilon_1 > 0$ and $\tau_2 < \infty$ such that if $\boldsymbol{\lambda}_t \in B_{\varepsilon_1}(\boldsymbol{\lambda}_1^*)$, then $Pr(\boldsymbol{\lambda}_{t+\sum_{j=2}^J \tau_j} \in B_{\varepsilon_J}(\boldsymbol{\lambda}_J^*)) > 0$.

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

Consider agreement outcome $\boldsymbol{\lambda}_1^* \in \{0^k, \infty^k\}$. By Lemma 4, for any initial belief $\boldsymbol{\lambda}_1 \in (0,\infty)^k$ and any $\varepsilon_1 > 0$, there exists a finite sequence of τ_1 actions and public signals such that following this sequence, $\boldsymbol{\lambda}_{\tau_1+1} \in B_{\varepsilon_1}(\boldsymbol{\lambda}_1^*)$. Therefore, from any initial beliefs, $Pr(\boldsymbol{\lambda}_{\tau_1+1} \in B_{\varepsilon_1}(\boldsymbol{\lambda}_1^*)) > 0$. Therefore, for any $\varepsilon_J > 0$ and initial beliefs $\boldsymbol{\lambda}_1 \in (0,\infty)^k$, $Pr(\boldsymbol{\lambda}_\tau \in B_{\varepsilon_J}(\boldsymbol{\lambda}_J^*)) > 0$, where $\tau \equiv \sum_{j=1}^J \tau_j + 1$. Since each $\tau_j < \infty$, $\tau < \infty$. □

Proof of Lemma 5′. Consider $\boldsymbol{\lambda}^* = (0^\kappa, \infty^{k-\kappa})$. Suppose $\boldsymbol{\lambda}^* \in \Lambda(\omega)$ and $\boldsymbol{\lambda}^*$ is maximally accessible. Consider the sequence of stationary likelihood ratios $\boldsymbol{\lambda}_j^* = (0^{k-j+1}, \infty^{j-1})$ for $j = 1, \dots k-\kappa+1$, and suppose part (ii) of Definition 5′ holds. We first show that this implies separability at zero (Definition 6) for each likelihood ratio in the sequence. For each $j = 1, \dots k-\kappa+1$, define the submatrix

$$A_j \equiv \begin{pmatrix} \log \frac{\hat{\psi}_{k-j+1}(a_1,\sigma_L|R,\boldsymbol{\lambda}_j^*)}{\hat{\psi}_{k-j+1}(a_1,\sigma_L|L,\boldsymbol{\lambda}_j^*)} & \log \frac{\hat{\psi}_{k-j+1}(a_M,\sigma_R|R,\boldsymbol{\lambda}_j^*)}{\hat{\psi}_{k-j+1}(a_M,\sigma_R|L,\boldsymbol{\lambda}_j^*)} \\ \log \frac{\hat{\psi}_{k-j}(a_1,\sigma_L|R,\boldsymbol{\lambda}_j^*)}{\hat{\psi}_{k-j}(a_1,\sigma_L|L,\boldsymbol{\lambda}_j^*)} & \log \frac{\hat{\psi}_{k-j}(a_M,\sigma_R|R,\boldsymbol{\lambda}_j^*)}{\hat{\psi}_{k-j}(a_M,\sigma_R|L,\boldsymbol{\lambda}_j^*)} \end{pmatrix}.$$

Since $\theta_{k-j+1} \succ \theta_{k-j}$, this has a positive determinant. Therefore, there exists a $c \in \mathbb{R}_+^2$ that solves

$$A_j c = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

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

$$A_j \tilde{c} = \begin{pmatrix} G_{k-j+1} \\ G_{k-j} \end{pmatrix},$$

where $G_{k-j+1} > 0$ and $G_{k-j} < 0$. Moreover, by maximal accessibility, for any $j' > k-j+1$,

$$\left(\log \frac{\hat{\psi}_{j'}(a_M,\sigma_R|R,\boldsymbol{\lambda}_j^*)}{\hat{\psi}_{j'}(a_M,\sigma_R|L,\boldsymbol{\lambda}_j^*)}, \log \frac{\hat{\psi}_{j'}(a_1,\sigma_L|R,\boldsymbol{\lambda}_j^*)}{\hat{\psi}_{j'}(a_1,\sigma_L|L,\boldsymbol{\lambda}_j^*)} \right) \cdot \tilde{c} > 0$$

and for any $j' < k-j$,

$$\left(\log \frac{\hat{\psi}_{j'}(a_M,\sigma_R|R,\boldsymbol{\lambda}_j^*)}{\hat{\psi}_{j'}(a_M,\sigma_R|L,\boldsymbol{\lambda}_j^*)}, \log \frac{\hat{\psi}_{j'}(a_1,\sigma_L|R,\boldsymbol{\lambda}_j^*)}{\hat{\psi}_{j'}(a_1,\sigma_L|L,\boldsymbol{\lambda}_j^*)} \right) \cdot \tilde{c} < 0.$$

Therefore, $\boldsymbol{\lambda}_j^*$ is separable at zero, since we can set the elements of c to zero for the remaining action and signal pairs in matrix (17). Therefore, by Lemma 8, $\boldsymbol{\lambda}_{j+1}^*$ is adjacently accessible from $\boldsymbol{\lambda}_j^*$. Since this holds for each element of the sequence, starting at $\boldsymbol{\lambda}_1^* = 0^k$ and ending at $\boldsymbol{\lambda}_J^* = \boldsymbol{\lambda}^*$, by Lemma 9′, $\boldsymbol{\lambda}^*$ is accessible. Similar to the proof of Lemma 9, we can choose $\varepsilon < e^{-E}$, so that the likelihood ratio reaches the locally stable neighborhood of $\boldsymbol{\lambda}^*$ with positive probability. From here, local stability implies that $P(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) > 0$. The case where part (i) of Definition 5′ holds is analogous. □

We use the following result in the proof of Lemma 6′.

Lemma 14. *Given mixed outcome $(\boldsymbol{\lambda}_I^*, I)$, 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 $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{i\in N} \Lambda_i(\omega)$.*

Proof. Let $(\boldsymbol{\lambda}_N, N)$ be a terminal node in $\mathcal{G}(\boldsymbol{\lambda}_I^*, I)$. By definition of terminal node, no nodes are strongly mixed accessible from $(\boldsymbol{\lambda}_N, N)$. If any node is mixed accessible from $(\boldsymbol{\lambda}_N, N)$, then there exists an $i \in N$ 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 $j \neq i$ is strongly 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, $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{i\in N} \Lambda_i(\omega)$. □

Proof of Lemma 6′. Consider a generic misspecified model and 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 $\varepsilon < e^{-E}$ 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 sociable types.

By reducibility, at every $\boldsymbol{\lambda}_N \in \{0,\infty\}^{|N|}$, either $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{i\in N} \Lambda_i(\omega)$ or there exists a $\boldsymbol{\lambda}'_N \in \{0,\infty\}^{|N|}$ such that $(\boldsymbol{\lambda}'_N, N)$ is strongly mixed accessible from $(\boldsymbol{\lambda}_N, N)$ and $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}'_N) \in \cap_{i\in N} \Lambda_i(\omega)$. First consider $\boldsymbol{\lambda}_N \in \{0,\infty\}^{|N|}$ such that $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{i\in N} \Lambda_i(\omega)$. By the construction in Lemma 3, 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 $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \in \cap_{i\in I} \Lambda_i(\omega)$, then $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N)$ is locally stable, so with positive probably, $\boldsymbol{\lambda}_t \to (\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N)$. Otherwise, if $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \notin \cap_{i\in I} \Lambda_i(\omega)$, then for some $i \in I$, $\langle \lambda_{i,t} \rangle$ 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 $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \notin \cap_{i\in N} \Lambda_i(\omega)$. 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 a $(\boldsymbol{\lambda}'_N, N)$ that is strongly 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 $i \in N$ and λ_i^* denote the component for type $i \in I$. By the construction in Lemma 3, there exists an $i \in N$ 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 $i \in N$ such that $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}_N) \notin \Lambda_i(\omega)$, with $N_{U,0}$ the set of $i \in N_U$ such that $\lambda_i = 0$ and $N_{U,\infty}$ the set of $i \in N_U$ such that $\lambda_i = \infty$. We now argue that starting from a neighborhood $B_{\varepsilon_2}(\boldsymbol{\lambda}_N)$ for $i \in N$ and $B_{\varepsilon}(\boldsymbol{\lambda}_I^*)$ for $i \in I$, 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 strongly mixed accessible point $(\boldsymbol{\lambda}'_N, N)$. For $i \in N_{U,0}$, let N_i be the minimum number of (a_M, σ_R) actions it takes for any $\lambda_{i,t} \in [\varepsilon', 1/\varepsilon']$ to hit $1/\varepsilon'$. Similarly, for $i \in N_{U,\infty}$, let N_i be the minimum number of (a_1, σ_L) actions it takes for any $\lambda_{i,t} \in [\varepsilon', 1/\varepsilon']$ to hit ε'. By the construction in Lemma 3, 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 ε_2 such that if $\boldsymbol{\lambda}_{N \setminus N_U,1} \in B_{\varepsilon_2}(\boldsymbol{\lambda}_{N \setminus N_U})$, then after $\sum_{i \in N_{U,0}} N_i$ action and signal realizations (a_M, σ_R), $\lambda_{i,t} \in B_{\varepsilon'}(\lambda_i)$ for all $i \in N \setminus N_{U,0}$, and after $\sum_{i \in N_{U,\infty}} N_i$ action and signal realizations (a_1, σ_L), $\lambda_{i,t} \in B_{\varepsilon'}(\lambda_i)$ for all $i \in N \setminus N_{U,\infty}$. 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 $i \in N_U$, $\lambda_{i,t} \notin B_{\varepsilon'}(\lambda_i)$ and for all $i \in N \setminus N_U$, $\lambda_{i,t} \in B_{\varepsilon_2}(\lambda_i)$. First consider case (ii) and suppose that a type $i \in N_{U,0}$ leaves. After N_i actions and signals (a_M, σ_R), if $\boldsymbol{\lambda}_{N,t} \in B_{\varepsilon'}(\boldsymbol{\lambda}'_N)$ for some $(\boldsymbol{\lambda}'_N, N)$ that is strongly mixed accessible from $(\boldsymbol{\lambda}_N, N)$, then stop. Otherwise, there exists an $i_2 \in N_{U,0}$ such that $\lambda_{i_2,t} \notin B_{\varepsilon'}(\lambda_{i_2})$. Repeat N_{i_2} realizations (a_M, σ_R). After these $N_{i_1} + N_{i_2}$ realizations of (a_M, σ_R), if $\boldsymbol{\lambda}_{N,t} \in B_{\varepsilon'}(\boldsymbol{\lambda}'_N)$ for some $(\boldsymbol{\lambda}'_N, N)$ that is strongly mixed accessible from $(\boldsymbol{\lambda}_N, N)$, then stop. Otherwise, there is an $i_3 \in N_{\boldsymbol{\lambda},0}$ such that $\lambda_{i_3,t} \notin B_{\varepsilon'}(\lambda_{i_3})$. Repeat N_{i_3} realizations (a_M, σ_R), and so on. Therefore, after at most $\sum_{i \in N_{U,0}} N_i$ realizations of (a_M, σ_R), beliefs have entered the ε' ball around some other stationary point $(\boldsymbol{\lambda}_I^*, \boldsymbol{\lambda}'_N)$ such that $(\boldsymbol{\lambda}'_N, N)$ is strongly 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 $(\boldsymbol{\lambda}'_N, N)$ that is strongly 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 $\sum_{i \in N_{U,0}} N_i$ realizations of (a_M, σ_R), which is strictly positive. The argument for a type $i \in N_{U,\infty}$ 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 $B_{\varepsilon(\boldsymbol{\lambda}_N)}(\boldsymbol{\lambda}_N)$ 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,

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

If $\langle \boldsymbol{\lambda}_{N,t} \rangle$ enters the neighborhood of a terminal node $(\boldsymbol{\lambda}_N, N)$ infinitely often, then $\boldsymbol{\lambda}_N \in$

$\cap_{i \in N} \Lambda_i(\omega)$, 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 $(\boldsymbol{\lambda}'_N, N)$ that is strongly mixed accessible 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. □

Proof of Lemma 7′. Suppose $\Lambda(\omega)$ contains an agreement vector or maximally accessible disagreement vector and $\Lambda_M(\omega)$ is empty. Recall that $\mathcal{B}$ is the set of locally stable neighborhoods and $\mathcal{B}_U$ is the set of locally unstable neighborhoods. 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 10, there exists a finite sequence of actions and signals 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 Lemma 3, $\langle \boldsymbol{\lambda}_t \rangle$ almost surely leaves $\mathcal{B}_U$. If $\langle \boldsymbol{\lambda}_t \rangle$ enters the neighborhood of a mixed outcome $\boldsymbol{\lambda}_I$, by Lemma 6′, $\langle \boldsymbol{\lambda}_t \rangle$ almost surely leaves this neighborhood or converges to a locally stable point. By Lemma 2, $\langle \boldsymbol{\lambda}_t \rangle$ 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, $[e^{-E}, e^{E}]^k$. 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 6′, 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 Lemma 3, $Pr(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}_\infty | \tau_2 < \infty) = 1$, 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 $\text{supp}(\boldsymbol{\lambda}_\infty) \subset \Lambda(\omega)$. □

E.3 Sufficient Conditions for Mixed Learning

The intuition for when mixed learning outcomes arise is similar to that for convergent learning outcomes. Consider the mixed learning outcome in which $\lambda_1 \to 0$. A sufficient condition for the likelihood ratio of type θ_1 to converge to zero, independently of λ_2, is that the expected change in $\log \lambda_1$ is negative at zero for all possible beliefs of type θ_2, i.e. $\sup_{\lambda_2} \gamma_1((0, \lambda_2), \omega) < 0$. We also need to ensure that type θ_2's beliefs do not converge. By definition of $\Lambda_M(\omega)$, no limit beliefs with $\lambda_1 = 0$ are locally stable for θ_2, i.e. $(0,0) \notin \Lambda_2(\omega)$ and $(0, \infty) \notin \Lambda_2(\omega)$. Therefore, θ_2's beliefs do not converge. The following theorem states sufficient conditions for mixed learning to occur with positive probability.

Theorem 4. *Suppose there are two sociable types, Assumption 1, 2, 3 and 4 hold, and the true and subjective private signal distributions have a finite number of discontinuities. If mixed outcome $(\lambda_i, \theta_i) \in \Lambda_M(\omega)$ and (i) $\sup_{\lambda_{-i}} \gamma_i((0, \lambda_{-i}), \omega) < 0$ if $\lambda_i = 0$, or (ii) $\inf_{\lambda_{-i}} \gamma_i((\infty, \lambda_{-i}), \omega) > 0$ if $\lambda_i = \infty$, then the mixed outcome occurs with positive probability. If $(\lambda_i, \theta_i) \notin \Lambda_M(\omega)$, then the mixed outcome almost surely does not occur.*

Proof. Suppose $\omega = L$ and suppose the mixed outcome $(\lambda_2, \theta_2) \in \Lambda_M(L)$. As in the proof of Lemma 3, we can construct neighborhoods $(0, e^{-E}]^2$ and $[e^E, \infty) \times (0, e^{-E}]$ such that in each of these neighborhoods, there exists an i.i.d. process that bounds θ_1's updates above as long as beliefs remain in the neighborhood and almost surely converges to zero, and a process that bounds θ_2's updates below (above) in the neighborhood of 0 (∞) and eventually leaves the neighborhood.

Consider the interior of the belief space, $[e^{-E}, e^E]^2$. This space can be partitioned into finitely many closed, convex sets $D_1, D_2, \ldots D_N$ where $\gamma_2(\cdot, L)$ is continuous on the interior of these sets. Consider the set D_j and define the function $\hat{\gamma}_{D_j} : D_j \to \mathbb{R}$ as

$$\hat{\gamma}_{i,D_j}(\boldsymbol{\lambda}) \equiv \begin{cases} \gamma_2(\boldsymbol{\lambda}, L) \text{ if } \boldsymbol{\lambda} \in \text{ interior of } D_j \\ \lim_{x \to \boldsymbol{\lambda}} \gamma_2(x, L) \text{ otherwise.} \end{cases}$$

This is a continuous function. So, for each $(\lambda, 0) \in D_j$, we can construct an open, convex set $B(\lambda, 0)$ such that if $\boldsymbol{\lambda}_t$ is in this set, then $\log \lambda_{2,t+1} - \log \lambda_{2,t}$ is bounded above by

$$g_j(a, \sigma) \equiv \sup_{\boldsymbol{\lambda} \in B(\lambda, 0)} \log \frac{\hat{\psi}_2(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_2(a, \sigma | L, \boldsymbol{\lambda})}.$$

Let

$$\bar{g}_j \equiv \max_{(a,\sigma) \in \mathcal{A} \times \Sigma} g_j(a, \sigma).$$

Define the process

$$\xi_{D_j, t+1} = \xi_{D_j, t} + g_j(a(\tilde{\theta}_t, \tilde{s}_t, (\lambda, 0)), \tilde{\sigma}_t),$$

when $(\tilde{\theta}_t, \tilde{s}_t)$ is such that $a(\tilde{\theta}_t, \tilde{s}_t, \boldsymbol{\lambda}) = a(\tilde{\theta}_t, \tilde{s}_t, (\lambda, 0))$ for all beliefs $\boldsymbol{\lambda} \in D_j$ (note this includes all autarkic types) and

$$\xi_{D_j,t+1} = \xi_{D_j,t} + \bar{g}_j$$

otherwise. When $\omega = L$, let $\psi_j(a,\sigma)$ be the probability of (a,σ) in the former event and let $\bar{\psi}_j$ be the probability of the latter event. By construction, $\log \lambda_{2,t+1} - \log \lambda_{2,t} < \xi_{D_j,t+1} - \xi_{D_j,t}$ if $\boldsymbol{\lambda}_t \in D_j$. Moreover, choose D_j sufficiently small so that

$$\bar{\psi}_j \bar{g}_j + \sum_{(a,\sigma)\in\mathcal{A}\times\Sigma} \psi_j(a,\sigma) g_j(a,\sigma) < 0. \tag{31}$$

As in Lemma 3, this sequence converges to $-\infty$ almost surely. But now the process that bounds the updates changes as the likelihood ratio moves across the state space even if λ_2 stays in a neighborhood of 0, so this is insufficient to conclude that λ_2 converges.

By compactness, we can find a finite collection of open sets $B_{D_j,1} \dots B_{D_j,n}$ that contain all $(\lambda, 0) \in D_j$. Since this procedure can be done for each D_j, there exists a disjoint, finite collection of sets $\mathcal{C} = (C_j)_{i=1}^N$ and an $\varepsilon > 0$ such that these sets contain each point $\boldsymbol{\lambda} \in [e^{-E}, e^E] \times [0, \varepsilon]$, each set is contained in exactly one $B_{D_j,j}$ for some i and j, and each discontinuity point $\boldsymbol{\lambda}$ is contained in a subset of the D_j, where $\gamma_2(\boldsymbol{\lambda}, L) = \hat{\gamma}_{D_j}(\boldsymbol{\lambda})$. Within each set in this cover, $\log \lambda_{2,t} - \log \lambda_{2,1}$ is bounded above by an i.i.d. process $\xi_{C_j,t} - \xi_{C_j,t-1}$, and with positive probability, $\sup_{C_j} E(\xi_{C_j,t} - \xi_{C_j,t-1}) < 0$. It remains to show that with positive probability, λ_2 remains below ε and converges. For each C_j, there exists a sequence of actions and public signals such that $\xi_{C_j} \to -\infty$ and $\sup \xi_t < \log \varepsilon / 2N$. Let N_j be the set of realizations of the process $(\xi_{C_j,t} - \xi_{C_j,t-1})_{t=1}^\infty$ for set C_j. Let $\mathit{\Xi}_j^T$ be the set of these sequences truncated after the first T terms.

Let $\tau_{j,k}$ be the kth time the $\boldsymbol{\lambda}$ process enters C_j, and let $\tau_{j,0} = 0$ and $\xi_{C_j,0} = 0$. Let n be the n that satisfies where $\tau_{j,n} \leq t$ and $\tau_{j,n+1} > t$ if it exists. Let A_t be the event that for all C_j, the sequence $N_j^T = ((\xi_{C_j,\tau_{j,k}+1} - \xi_{j,\tau_{j,k}}))_{k=1}^n$ is in $\mathit{\Xi}_j^{\tau_{j,n}}$ for each j. Elements in this sequence bound the change in the log-likelihood ratio above at each time $\tau_{j,k}$ when the likelihood ratio is in set C_j. Finally let P_{j,N_j^T} be the probability that the process realization of the process $\xi_{C_j,s}$ $s \geq T$ satisfies $(N_j^T, (\xi_{C_j,s+1} - \xi_{C_j,s})_{s=T}^\infty)) \in \mathit{\Xi}_j$, and let P_j

be the probability that $(\xi_{C_j,s+1} - \xi_{C_j,s})_{s=T}^{\infty} \in \Xi_j$. Let c_t be the set C that $\boldsymbol{\lambda}_t$ is in. Then

$$\begin{aligned}
Pr(A_2|c_1) &= Pr(A_1|c_1)Pr(A_2|A_1, c_1) \\
&= Pr(A_1|c_1)E(Pr(A_2|\xi_{c_1,1}, A_1, c_1)|A_1, c_1) \\
&= Pr(A_1|c_1)E[Pr(c_1 \neq c_2|c_1, A_1, \xi_{c_1,1})Pr(A_2|c_1, A_1, \xi_{c_1,1}, c_1 \neq c_2) \\
&+ Pr(c_1 = c_2|c_1, A_1, \xi_1)Pr(A_2|c_1, A_1, \xi_{c_1,1}, c_1 = c_2)|A_1, c_1] \\
&\geq Pr(A_1|c_1)E[Pr(c_1 \neq c_2|c_1) \\
&\quad \sum_{c \neq c_1} Pr(c|c_1, A_1, \xi_{c_1}, c \neq c_1)P_c + Pr(c_1 = c_2|c_1, A_1, \xi_{c_1})P_{c_1,N_c^2}|A_1, c_1] \\
&\geq Pr(A_1|c_1)E(P_{c_1,N_c^2} \prod_{c \neq c_1} P_{c,1}|A_1, c_1) = \prod_{j=1}^{N} P_j > 0
\end{aligned}$$

where the first inequality follows from the fact that at time t, given the current neighborhood beliefs are in c, the probability that the next realization is consistent with the desired sequence is at least the probability that all subsequent realizations of ξ are in that sequence. Now suppose we start at time t, and condition on the current set and the past realizations of the sequences. Then

$$\begin{aligned}
Pr(A_{t+1}|(N_j^t)_{j=1}^n, c_t) &= Pr(A_t|(N_j^t)_{j=1}^N, c_t)Pr(A_{t+1}|A_t, (N_j^t)_{j=1}^N, c_t) \\
&= Pr(A_t|N_j^t)_{j=1}^n, c_t)E[Pr(c_t \neq c_{t+1}|(N_j^{t+1})_{j=1}^N, c_t, A_t) \\
&\quad *Pr(A_{t+1}|(N_j^{t+1})_{j=1}^N, c_t, A_t, c_t \neq c_{t+1}) \\
&+ Pr(c_t = c_{t+1}|(N_j^{t+1})_{j=1}^N, c_t, A_t) \\
&\quad *Pr(A_{t+1}|(N_j^{t+1})_{j=1}^N, c_t, A_t, c_t = c_{t+1})|A_t, c_t, (N_j^t)_{j=1}^N] \\
&\geq \mathrm{Pr}(A_t|(N_c)_{j=1}^N, c_t)E(P_{c_t,N_{c_t}^{t+1}} \prod_{j \neq c_t} P_{j,N_j^t}|A_t, c_t, (N_j^t)_{j=1}^N) \\
&= \prod_{j=1}^{N} P_{j,N_j^t},
\end{aligned}$$

where the inequality follows from similar logic to the previous case. Conditional on knowing the current set c_t and the previous realizations of the sequence when $\boldsymbol{\lambda}_t$ was in c_t, the current realization being consistent does not depend on anything else. Moreover, the current realization being consistent is bounded below by the probability that all future realizations are consistent. Finally

$$\begin{aligned} Pr(A_{t+1}) &= E(Pr(A_{t+1}|(N_j^t)_{j=1}^N c_t)) \geq E(\Pi_{j=1}^N P_{j,N_j^t}) \\ &= E(E(\Pi_{j=1}^N P_{j,N_j^t}|(N_j^{t-1})_{j=1}^N, c_{t-1})) \\ &= E(P_{j,N_j^{t-1}}) \dots \prod_{j=1}^n P_j > 0. \end{aligned}$$

Therefore, $\lim_{T\to\infty} P(A_T) > 0$. By the dominated convergence theorem, $\lim_{T\to\infty} Pr(A_T) = Pr(A)$. Moreover, at any time T, if the event A_T has occurred and $\lambda_{2,1} < \varepsilon/2$, the likelihood ratio updates are bounded above by

$$\log \lambda_{2,T} - \log \lambda_{2,1} \leq \sum_{t=1}^{T-1} (\xi_{c_t,t+1} - \xi_{c_t,t}) < N\varepsilon/2N = \varepsilon/2.$$

So λ_2 never leaves the ε-ball.

Finally, since the mixed outcome $(\lambda_2, \theta_2) \in \Lambda_M(\omega)$, beliefs cannot converge to $(0, 0)$ or $(\infty, 0)$. Otherwise, beliefs would eventually enter either $(0, e^{-E}]^2$ or $[e^E, \infty) \times (0, e^{-E}]$ and never leave. But there exists a process that bounds type θ_1's belief updates below and leaves the neighborhood almost surely, which is a contradiction. Therefore, the mixed outcome occurs with positive probability. □

E.4 Belief-Dependent Signal Misspecification

Some updating rules are represented by a model of inference in which the subjective signal and/or type distributions depend on an agent's current belief. For example, one of the parameterizations of over/underweighting in Section 5.1 and the confirmation bias application in Section 5.2 both have subjective signal distributions that vary with the current belief. In this section, we present the formal set-up of this extension. This extension allows our framework to nest Rabin and Schrag (1999); Epstein et al. (2010).

Specifically, given likelihood ratio λ, an agent of type θ_i has subjective private signal distribution $\hat{F}_i^\omega(s; \lambda)$, subjective public signal distribution $\hat{G}_i^\omega(\sigma; \lambda)$, and subjective type distribution $\hat{\pi}_i(\theta; \lambda)$ in state ω. An agent of type θ_i uses likelihood ratio $\lambda_{i,t}$ to interpret the realized action and public signal $(\tilde{a}_t, \tilde{\sigma}_t)$ at time t.

As in Section 2, we focus on settings in which sociable types believe that actions or public signals are informative. When the model of inference depends on the current belief, we need to modify Assumption 3 to ensure that the adequate information property 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 across the belief space. 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 14 (Uniformly Informative). *The family of subjective private signal distributions* $\{\hat{F}_i^L(s;\lambda), \hat{F}_i^R(s;\lambda)\}_{\lambda\in[0,\infty)}$ *are* uniformly informative *if for all* $s \in \mathcal{S}$, *either* $\hat{F}_i^L(s;\lambda) = \hat{F}_i^R(s;\lambda) = 0$ *for all* $\lambda \in [0,\infty)$, $\hat{F}_i^L(s;\lambda) = \hat{F}_i^R(s;\lambda) = 1$ *for all* $\lambda \in [0,\infty)$, *or* $\inf_{\lambda\in(0,\infty)} \hat{F}_i^L(s;\lambda) - \hat{F}_i^R(s;\lambda) > 0$, *and similarly for the family of subjective public signal distributions.*

We use this definition to modify Assumption 3 in a way that ensures adequate information at all values of the likelihood ratio. This rules out sequences of beliefs along which either public signals or an autarkic type's actions are perceived to become arbitrarily uninformative.

Assumption 3′ (Uniform Adequate Information). *Either (i) public signals are informative,* $dG^R/dG^L \neq 1$, *and each sociable type* $\theta_i \in \Theta_S$ *has uniformly informative subjective public signal distributions, or (ii) there exists an autarkic type* $\theta_j \in \Theta_A$ *with* $\pi(\theta_j) > 0$ *that plays actions* a_1 *and* a_M *with positive probability, and each sociable type* $\theta_i \in \Theta_S$ *believes this autarkic type uniformly exists,* $\inf_{\lambda\in[0,\infty)} \hat{\pi}_i(\theta_j;\lambda) > 0$ *and has uniformly informative subjective private 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* $\lambda \mapsto (\hat{F}_i^L, \hat{F}_i^R, \hat{G}_i^L, \hat{G}_i^R, \hat{\pi}_i)$ *is continuous under the total variation norm except at at most a finite number of interior likelihood ratios* $\lambda \in (0,\infty)$.

Substituting Assumption 3′ for Assumption 3 and adding Assumption 5, the proofs of Lemmas 1 to 7 are unchanged, using the following modified version of (3) for the probability of each action and public signal:

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

The proofs of Lemmas 10 and 11 require minor modifications, but continue to hold. Therefore, Theorems 1 and 2 and Corollaries 1 and 2 continue to hold as stated in Section 4.2.

F Additional Material from Example 2 and Section 5.5

F.1 Proof of Proposition 8 from Example 2

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. To capture its explicit dependence on these parameters, let $\gamma_1^{\nu,q}(\boldsymbol{\lambda}, \omega)$ correspond to the function $\gamma_1(\boldsymbol{\lambda}, \omega)$ and $\Lambda^{\nu,q}(\omega)$ correspond to the set $\Lambda(\omega)$ in the model with partisan bias level ν and share of nonpartisan types q, with an analogous definition of $\gamma_3^{\nu,q}(\boldsymbol{\lambda}, \omega)$. Since beliefs move in unison, $\gamma_3^{\nu,q}(\boldsymbol{\lambda}, \omega) = \gamma_1^{\nu,q}(\boldsymbol{\lambda}, \omega)$, and therefore, we can focus on characterizing $\gamma_1^{\nu,q}(\boldsymbol{\lambda}, \omega)$ at the two possible stationary limit beliefs, $(0, 0)$ and (∞, ∞).

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

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

where

$$\begin{aligned}
\psi^{\nu,q}(L|L,(\infty,\infty)) &\equiv \pi_A[(1-q)\alpha_\nu + q\alpha_1] \\
\psi^{\nu,q}(R|L,(\infty,\infty)) &\equiv \pi_A[(1-q)(1-\alpha_\nu) + q(1-\alpha_1)] + 1 - \pi_A.
\end{aligned}$$

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, $\psi^{\nu,q}(L|L,(\infty,\infty))$, is strictly increasing in ν and q, and the probability of an R action, $\psi^{\nu,q}(R|L,(\infty,\infty))$, is strictly decreasing in ν and q. Therefore, $\gamma_1^{\nu,q}((\infty,\infty), L)$ is strictly decreasing in ν and q. At $\nu = 1$, both partisan and nonpartisan types are identical, so $\psi^{1,q}(L|L,(\infty,\infty)) = \pi_A\alpha_1$ and $\psi^{1,q}(R|L,(\infty,\infty)) = \pi_A(1-\alpha_1) + 1 - \pi_A$. Therefore, for any $q \in [0,1]$, $\gamma_1^{1,q}((\infty,\infty), L) < 0$ by the concavity of the log operator. Similarly, at $q = 1$, for any $\nu \in [0,1]$, $\gamma_1^{\nu,1}((\infty,\infty), L) < 0$ by the concavity of the log operator. At $\nu = 0$, θ_2 chooses action R for all signals, $\alpha_0 = 0$. Therefore, at $q = 0$, $\psi^{0,0}(L|L,(\infty,\infty)) = 0$ and $\gamma_1^{0,0}((\infty,\infty), L) = \log \frac{1-\pi_A(1-\alpha_1)}{1-\pi_A\alpha_1} > 0$. 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 $\nu < \nu_1(q)$, $\gamma_1^{\nu,q}((\infty,\infty), L) > 0$ and $(\infty,\infty) \in \Lambda^{\nu,q}(L)$ and for $\nu > \nu_1(q)$, $\gamma_1^{\nu,q}((\infty,\infty), L) < 0$ and $(\infty,\infty) \notin \Lambda^{\nu,q}(L)$. For $q > q_1$, $\gamma_1^{\nu,q}((\infty,\infty), L) < 0$ and $(\infty,\infty) \notin \Lambda^{\nu,q}(L)$.

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

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

where

$$\begin{aligned}
\psi^{\nu,q}(L|L,(0,0)) &\equiv \pi_A[(1-q)\alpha_\nu + q\alpha_1] + 1 - \pi_A \\
\psi^{\nu,q}(R|L,(0,0)) &\equiv \pi_A[(1-q)(1-\alpha_\nu) + q(1-\alpha_1)].
\end{aligned}$$

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, $\gamma_1^{\nu,q}((0,0),L)$ is strictly decreasing in ν and q. By similar reasoning to the case of (∞,∞), at $\nu = 1$, $\gamma_1^{1,q}((0,0),L) < 0$ for all $q \in [0,1]$ and at $q = 1$, $\gamma_1^{1,q}((0,0),L) < 0$ for all $\nu \in [0,1]$ by the concavity of the log operator. At $\nu = 0$ and $q = 0$, $\psi^{0,0}(L|L,(0,0)) = 1 - \pi_A$ since $\alpha_0 = 0$. As in Example 1, $\gamma_1^{0,0}((0,0),L) > 0$. 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 $\nu < \nu_2(q)$, $\gamma_1^{\nu,q}((0,0),L) > 0$ and $(0,0) \notin \Lambda^{\nu,q}(L)$, and for $\nu > \nu_2(q)$, $\gamma_1^{\nu,q}((0,0),L) < 0$ and $(0,0) \in \Lambda^{\nu,q}(L)$. For $q > q_2$, $\gamma_1^{\nu,q}((0,0),L) < 0$ and $(0,0) \in \Lambda^{\nu,q}(L)$.

Finally we show that $q_1 < q_2$ and $\nu_1(q) < \nu_2(q)$ for all $q < q_1$. Note

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

and by the symmetry of the signal distributions, $\gamma_1^{\nu,q}((0,0),L) - \gamma_1^{1,q}((0,0),L) = \gamma_1^{\nu,q}((\infty,\infty),L) - \gamma_1^{1,q}((\infty,\infty),L)$. Moreover, $\gamma_1^{1,q}((0,0),L) - \gamma_1^{1,q}((\infty,\infty),L)$ is 0 at $\pi_A = 0$, 0 at $\pi_A = 1$, and concave in π_A.[55] Therefore, $(0,0) \notin \Lambda^{\nu,q}(\omega)$ before $(\infty,\infty) \in \Lambda^{\nu,q}(\omega)$. This establishes the first part of the proposition.

Next consider $\omega = R$. Then $\gamma^{1,q}((\infty,\infty),R) > 0$ and $\gamma^{1,q}((0,0),R) > 0$ 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: $\gamma_1^{\nu,q}((0,0),R)$ and $\gamma_1^{\nu,q}((\infty,\infty),R)$ are decreasing in ν and q. Therefore, $\gamma^{\nu,q}((0,0),R) > 0$ for all ν and q, which implies $(0,0) \notin \Lambda^{\nu,q}(R)$ for all

[55] The second derivative is

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

ν and q. Similarly, $\gamma_1^{\nu,q}((\infty,\infty),R) > 0$ for all ν and q, which implies $(\infty,\infty) \in \Lambda^{\nu,q}(R)$ for all ν and q. Therefore, $\Lambda^{\nu,q}(R) = \{(\infty,\infty)\}$ for all ν and q and learning is almost surely correct. $\square$

F.2 Cognitive Hierarchy Parameterization.

In the cognitive hierarchy parameterization, level-3 places non-trivial probability on level-1 types. We study how asymptotic learning varies with level-3's belief about the frequency of the level-2 type, denoted $q \equiv \hat{\pi}_3(\theta_2)$. Aside from allowing any $q \in [0,1)$, we maintain the set-up introduced in Section 5.5. As in Appendix C.5, let $\alpha_L \equiv F^L(1/2)$ be the probability a level-1 type plays action L in state L and $\alpha_R \equiv F^R(1/2)$ 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. To simplify exposition, assume that the types are evenly distributed, $\pi(\theta_1) = \pi(\theta_2) = \pi(\theta_3) = 1/3$, and private signals are symmetrically distributed across states, $\alpha_L = 1 - \alpha_R$.

Similar to the level-k parameterization, we construct $\Lambda(\omega)$ for each q and show that $\Lambda_M(\omega)$ is empty. Lemma 12 holds for all $q \in [0,1)$, and therefore, both disagreement outcomes are maximally accessible. Therefore, $\Lambda(\omega)$ fully characterizes the set of asymptotic learning outcomes. Proposition 9 establishes that there are four regions of learning, featuring incorrect learning, correct learning and disagreement.

Proposition 9 (Cognitive Hierarchy)**.** *The likelihood ratio almost surely converges to a limit random variable with support $\Lambda(\omega) \neq \emptyset$. When $\omega = L$, there exist unique cutoffs $0 < q_1 < q_2 < q_3 < 1$ such that:*

1. *If $q < q_1$, then incorrect and correct learning occur with positive probability, $\Lambda(L) = \{(0,0),(\infty,\infty)\}$.*

2. *If $q \in (q_1,q_2)$, then incorrect learning, correct learning and disagreement occur with positive probability, $\Lambda(L) = \{(0,0),(\infty,\infty),(0,\infty)\}$.*

3. *If $q \in (q_2,q_3)$, then correct learning and disagreement occur with positive probability, $\Lambda(L) = \{(0,0),(0,\infty)\}$.*

4. *If $q > q_3$, then disagreement occurs almost surely, $\Lambda(L) = \{(0,\infty)\}$.*

An analogous result holds for $\omega = R$.

When q is low, level-3 types believe most agents are level-1 and they behave similarly to level-2 types. Both types overreact to confirming actions and underreact to contrary actions. Initial actions have an outsize effect on asymptotic beliefs, as the information from these actions is amplified in every subsequent action. Therefore, whether initial

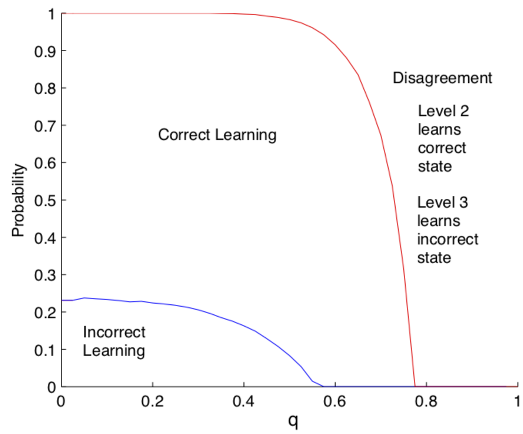

FIGURE 6. Cognitive Hierarchy Learning Outcomes
($\omega = L$, $F^L(s) = \frac{10}{3}(s - \frac{1}{2}s^2 - 3/5)$, $F^R(s) = \frac{5}{3}(s^2 - .04)$)

actions are correct or incorrect will influence whether beliefs build momentum on the correct or incorrect state, leading to either correct or incorrect learning. The models of level-2 and level-3 types are very close, and asymptotic disagreement is not possible.

As q increases, level-2 and level-3 types interpret the action history in an increasingly different way, and disagreement becomes possible. Further, as q increases, level-3 types move closer to the level-k model in which they anti-imitate the more frequent action. Even though level-2's model does not change, the shift in level-3's model leads to behavior that moves level-2's model closer to the correctly specified model. Therefore, disagreement takes a specific form: level-2 learns the correct state, while level-3 learns the incorrect state. Once q is sufficiently large, this disagreement outcome becomes the unique learning outcome, and level-3 almost surely learns the incorrect state, while level-2 almost surely learns the correct state.

Fig. 6 plots the probability of each learning outcome, as a function of q. Increasing q monotonically increases the probability that level-2 learns the correct state, as level-3's behavior mitigates level-2's misspecification. However, increasing q has a non-monotonic effect on the probability that level-3 learns the correct state. At first, raising q moves level-3's model closer to the true model, as it becomes aware of the level-2 type. This increases the probability of complete learning. But above $q = .55$, increasing q moves level-3's model further from the true model, as it begins to overestimate the frequency of the level-2 type. In this specification, $q_1 = .01$, $q_2 = .55$ and $q_3 = .76$.

While this example focuses on a particular distribution of types, $\pi = (0, 1/3, 1/3, 1/3)$, a robustness result that is similar in spirit to Theorem 2 establishes that Proposition 9 holds for nearby type distributions.

Proof of Proposition 9. Suppose $\omega = L$. Consider the level-2 type. Since $\alpha_L > 1/2$,

$$\begin{aligned}
\gamma_2((0,0),L) &= -\left(\frac{1+2\alpha_L}{3}\right)\log\left(\frac{\alpha_L}{1-\alpha_L}\right) < 0 \\
\gamma_2((\infty,0),L) &= \left(\frac{1-2\alpha_L}{3}\right)\log\left(\frac{\alpha_L}{1-\alpha_L}\right) < 0 \\
\gamma_2((0,\infty),L) &= \left(\frac{1-2\alpha_L}{3}\right)\log\left(\frac{\alpha_L}{1-\alpha_L}\right) < 0 \\
\gamma_2((\infty,\infty),L) &= \left(\frac{3-2\alpha_L}{3}\right)\log\left(\frac{\alpha_L}{1-\alpha_L}\right) > 0.
\end{aligned}$$

Therefore, $\Lambda_2(L) = \{(0,0),(0,\infty),(\infty,\infty)\}$. Consider the level-3 type.

$$\begin{aligned}
\gamma_3((\infty,\infty),L) &= \left(\frac{\alpha_L}{3}\right)\log\left(\frac{1-\alpha_L}{\alpha_L}\right) + \left(\frac{3-\alpha_L}{3}\right)\log\left(\frac{q+(1-q)\alpha_L}{q+(1-q)(1-\alpha_L)}\right) \\
\gamma_3((0,\infty),L) &= \left(\frac{1+\alpha_L}{3}\right)\log\left(\frac{q+(1-q)(1-\alpha_L)}{q+(1-q)\alpha_L}\right) + \left(\frac{2-\alpha_L}{3}\right)\log\left(\frac{\alpha_L}{1-\alpha_L}\right) \\
\gamma_3((0,0),L) &= \left(\frac{2+\alpha_L}{3}\right)\log\left(\frac{q+(1-q)(1-\alpha_L)}{q+(1-q)\alpha_L}\right) + \left(\frac{1-\alpha_L}{3}\right)\log\left(\frac{\alpha_L}{1-\alpha_L}\right).
\end{aligned}$$

If $\gamma_3((\infty,\infty),L) > 0$, then $(\infty,\infty) \in \Lambda(L)$. From these expressions, $\gamma_3((\infty,\infty),L)$ is positive at $q = 0$, decreasing in q, and negative at $q = 1$. Therefore, there exists a q_2 such that for $q < q_2$, $(\infty,\infty) \in \Lambda(L)$, and for $q > q_2$, $(\infty,\infty) \notin \Lambda(L)$. If $\gamma_3((0,\infty),L) > 0$, then $(0,\infty) \in \Lambda(L)$ and if $\gamma_3((0,0),L) < 0$, then $(0,0) \in \Lambda(L)$. The expressions $\gamma_3((0,0),L) < \gamma_3((0,\infty),L)$ are both negative at $q = 0$, increasing in q, and positive at $q = 1$. Therefore, there exists $q_1 < q_3$ such that $(0,0) \in \Lambda(L)$ for $q < q_3$ and $(0,0) \notin \Lambda(L)$ for $q > q_3$, while $(0,\infty) \notin \Lambda(L)$ for $q < q_1$ and $(0,\infty) \in \Lambda(L)$ for $q > q_1$. This yields the characterization of $\Lambda(L)$ as a function of q.

It immediately follows from Theorem 1 that the agreement outcomes $(0,0)$ and (∞,∞) are globally stable iff they are in $\Lambda(L)$. By Lemma 12, both disagreement outcomes are maximally accessible. Therefore, $(0,\infty)$ is globally stable iff $(0,\infty) \in \Lambda(L)$.

Finally we have to rule out mixed outcomes. In the region where both correct learning and incorrect learning are locally stable (parts 1 and 2), it immediately follows that $\Lambda_M(L)$ is empty and mixed outcomes almost surely do not arise. Given $\gamma_2((\infty,0),L) < 0$ and

$$\begin{aligned}
\gamma_3((\infty,0),L) &= \left(\frac{1+\alpha_L}{3}\right)\log\left(\frac{1-\alpha_L}{\alpha_L}\right) + \left(\frac{2-\alpha_L}{3}\right)\log\left(\frac{q+(1-q)\alpha_L}{q+(1-q)(1-\alpha_L)}\right) \\
&< \gamma_3((\infty,\infty),L) < 0,
\end{aligned}$$

$(\infty, \theta_2) \notin \Lambda_M(L)$. If disagreement and correct learning are locally stable (part 3), then (∞, θ_2) is the only candidate mixed outcome and therefore, $\Lambda_M(L)$ is empty. If only disagreement is locally stable (part 4), we also have to rule out $\lambda_3^* = 0$. But since $(0,0) \notin \Lambda(L)$, $\gamma_3((0,0), L) > 0$. Also, $\gamma_2((0,0), L) < 0$. Therefore, $(0, \theta_3) \notin \Lambda_M(L)$. Therefore, $\Lambda_M(L)$ is empty for all $q \in [0,1)$.

Given this characterization, by Theorem 1, beliefs almost surely converge to a limit random variable $\boldsymbol{\lambda}_\infty$ with $\operatorname{supp} \boldsymbol{\lambda}_\infty = \Lambda(L)$. □