Bounded Rationality And Learning: A Framework and A Robustness Result*

J. Aislinn Bohren† Daniel N. Hauser‡

May 2017

Abstract

We explore model misspecification in an observational learning framework. Individuals learn from private and public signals and the actions of others. An agent's type specifies her model of the world. Misspecified types have incorrect beliefs about the signal distribution, how other agents draw inference and/or others' payoffs. We establish that the correctly specified model is robust in that agents with approximately correct models almost surely learn the true state asymptotically. We develop a simple criterion to identify the asymptotic learning outcomes that arise when misspecification is more severe. Depending on the nature of the misspecification, learning may be correct, incorrect or beliefs may not converge. Different types may asymptotically disagree, despite observing the same sequence of information. This framework captures behavioral biases such as confirmation bias, false consensus effect, partisan bias and correlation neglect, as well as models of inference such as level-k and cognitive hierarchy.

KEYWORDS: Social learning, model misspecification, bounded rationality

*We thank Nageeb Ali, Alex Imas, Shuya Li, George Mailath, Margaret Meyer, Ali Polat, Andrew Postlewaite, Andrea Prat, Yuval Salant, Larry Samuelson, Ran Spiegler and conference and seminar participants at Carnegie Mellon, ESSET Gerzensee, NASM 2016, Pennsylvania Economic Theory Conference, University of Pennsylvania, University of Pittsburgh and Yale for helpful comments and suggestions.

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

‡Email: dahauser@sas.upenn.edu; University of Pennsylvania

1 Introduction

Faced with a new decision, individuals gather information from many diverse sources before choosing an action. This can include the choices of peers, the announcements of public institutions, such as a government or health agency, and private sources, such as past experiences in similar situations. For example, when deciding whether to enroll in a degree program, an individual may read pamphlets and statistics about the opportunities the program provides, discuss the merits of the program with faculty, and observe the enrollment choices of other students. Learning from these sources requires a model of how to interpret signals, how the choices of other individuals reflect their information, and how to aggregate multiple pieces of information.

A rich literature in psychology and experimental economics documents the myriad of biases that individuals exhibit when processing information and interpreting others' decisions. Individuals have been found to systematically overweight information in favor of their prior beliefs (confirmation bias),[1] overreact or underreact to information (over- and under-confidence),[2] incorrectly aggregate correlated information (correlation neglect),[3] systematically slant information towards a preferred state (motivated reasoning, partisan bias),[4] misunderstand strategic interaction (level-k, cognitive hierarchy),[5] and miscalculate the extent to which others' preferences are similar to their own (false consensus effect, pluralistic ignorance).[6] These biases are forms of model misspecification in which individuals have incorrect models of the informational environment and how others make decisions.

In this paper, we characterize how model misspecification affects long-run learning in a sequential learning framework. Individuals choose between two alternatives. Their payoff depends on their own action choice and an unknown state of the world. Prior to making a decision, an individual learns about the state by observing the actions of her predecessors, a private signal, and a sequence of public signals. An individual's type specifies how she interprets signals, and how she believes others

[1] Darley and Gross (1983); Lord, Ross, and Lepper (1979); Plous (1991).
[2] Moore and Healy (2008).
[3] Enke and Zimmermann (2017); Eyster and Weizsacker (2011); Kallir and Sonsino (2009).
[4] Bartels (2002); Bénabou and Tirole (2011); Brunnermeier and Parker (2005); Jerit and Barabas (2012); Koszegi and Rabin (2006); Kunda (1990).
[5] Kübler and Weizsäcker (2004); Kübler and Weizscker (2005); Penczynski (forthcoming).
[6] Gilovich (1990); Grebe, Schmid, and Stiehler (2008); Marks and Miller (1987); Miller and McFarland (1987, 1991); Ross, Greene, and House (1977).

draw inference and make decisions. Misspecified models of the signal process are modeled as a mapping from the true to the misperceived posterior belief. We provide a foundation for this representation as the reduced form of a misspecified measure over an arbitrary signal space (Appendix A.1). Misspecified models of how others draw inference are captured by a type's perceived distribution over the type space, which can differ from the true distribution. Individuals with different types may coexist, and they can either be aware or unaware of each others' models of the world. Our framework captures the information-processing biases cited above, and nests several previously developed behavioral models of inference.[7]

We study the asymptotic behavior and beliefs of individuals to determine when individuals with misspecified models adopt the desirable action, and whether individuals with different misspecified models conform or disagree. We know from correctly specified observational learning models that individuals asymptotically adopt the desirable action when sufficient information arrives.[8] Misspecification opens the door to learning outcomes – long-run beliefs about the state – that do not occur in the correctly specified model. This includes *incorrect* learning, where beliefs converge to the wrong state with positive probability, *non-stationary incomplete* learning, where beliefs about the state almost surely do not converge, and *disagreement*, where with positive probability, some types learn the correct state and others learn the incorrect state.

Our first main result (Theorem 1) characterizes each type's asymptotic learning outcomes. We show that the set of asymptotic learning outcomes that arise with positive probability depends on two expressions that are straightforward to derive from the primitives of the model – (i) the *expected change in the likelihood ratio* for each type near a candidate limit belief; and (ii) an ordering over the type space, which we refer to as the *total informativeness rank*. The first expression is used to

[7]Appendix B maps Rabin and Schrag (1999) and Epstein, Noor, and Sandroni (2010) into the framework of this paper. It also outlines how our framework can be used to study other non-Bayesian learning rules.

[8]Individuals almost surely adopt the optimal action asymptotically if there are arbitrarily precise private signals (Smith and Sorensen 2000), actions perfectly reveal beliefs (Lee 1993), a subset of individuals who do not observe others' actions (Acemoglu, Dahleh, Lobel, and Ozdaglar 2011), or an infinite sequence of public signals. Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) first studied the sequential observational learning framework with a binary signal space. They demonstrate that incomplete learning may arise when the action space is coarser than the belief space. Ali (2016) shows that incomplete learning can arise even when the action space is isomorphic to the belief space.

determine whether a learning outcome is locally stable, in that beliefs converge to this outcome with positive probability, from a neighborhood of the outcome. We show that a learning outcome is locally stable if and only if the expected change in the log likelihood ratio moves toward this outcome from nearby beliefs. We are interested in a characterization that is independent of the initial belief, and therefore, we need a tighter notion of stability. We say an outcome is globally stable if beliefs converge to this outcome with positive probability, from *any* initial belief. For an *agreement* outcome – all types have the same (possibly incorrect) limit beliefs – we show that local stability is both necessary and sufficient for global stability. However, a *disagreement* outcome – types have different limit beliefs – requires an additional condition to establish global stability. Starting from a common prior, it must be possible to separate the beliefs of different types. The second expression, the total informativeness rank, is a sufficient condition to separate beliefs. Therefore, a disagreement outcome arises with positive probability, from any initial belief, if it is locally stable and total informativeness ranked, while a disagreement outcome almost surely does not arise if it is not locally stable. Given a particular form of misspecification, deriving these two expressions will characterize the set of asymptotic learning outcomes. Correct learning, incorrect learning, non-stationary incomplete learning and disagreement are all possible under certain forms of misspecification.

To establish Theorem 1, we use results from Markov dynamic systems to characterize the limiting behavior of the belief process for each type. The equations of motion for the dynamic system are equilibrium objects that are derived from each agent's optimal choice, as well as their beliefs about the behavior of other agents. An individual's interpretation of others' actions depends on the current belief vector. Therefore, the equations of motion are state-dependent and nonlinear. This presents a technical challenge, as the process fails to satisfy standard conditions from the existing literature on Markov chains (despite the fact that the belief process has a countable state space).

Our second set of main results (Theorems 2 and 3) establish that the correctly specified model is robust to misspecification. As long as individuals have approximately correct models of the signal processes and how others draw inference, then learning is complete in that all types almost surely learn the correct state. Even if there are multiple types of individuals with biases that move in different directions, complete learning obtains, as long as none of these biases are too severe. This may

not seem surprising, since Bayes rule is continuous. But in an infinite horizon setting, a small bias in each period has the potential to sum to a large bias in aggregate. These results establish that this does not occur.

We close with three applications that demonstrate various forms of misspecification – level-k reasoning, partisan bias and confirmation bias. In the level-k application, individuals correctly interpret signals, but have a misspecified model of how others draw inference. Depending on the severity of the misspecification, individuals may learn the correct or incorrect state, and agents with different levels of reasoning may asymptotically disagree.[9] A surprising finding is that a higher level of reasoning may perform strictly *worse* than a lower level of reasoning. Therefore, it may be optimal for an agent to continue reasoning at a lower level, even if he can acquire the ability to use a higher level of reasoning for an arbitrarily small cost.

In the partisan bias application, some individuals systematically slant information towards one of the states. These partisan types believe that all other agents interpret information in the same way as them. Non-partisan types correctly interpret information, but do not account for the slant of the partisan types. We establish that as long as the frequency of partisan types or the level of their bias is not too large, then learning is correct for both types. As the bias and frequency of partisan types increase, both types pass through a region of the parameter space in which the beliefs of neither type converges, before reaching a region in which both types almost surely learn the incorrect state. We also consider the case in which non-partisan types have correct beliefs about the share of partisan types and their level of bias. Disagreement arises almost surely for severe enough levels of partisan bias.

Finally, in the confirmation bias application, a single type systematically slants information towards the state that she believes is more likely. As in Rabin and Schrag (1999), we show that incorrect learning can arise if the degree of confirmation bias is sufficiently high.

Similar to our robustness results for the correctly specified model, our characterizations of asymptotic learning outcomes in misspecified models are robust. Therefore, the insights from these applications are not sensitive to the exact choice of functional form used to pin down each bias. This also establishes that the models nested in our framework are robust to nearby forms of misspecification.

[9] Eyster and Rabin (2010) is a level-k model in which all agents are level-2, but believe that all other agents are level-1.

Related theoretical work shows that information processing biases and incorrect models of inference can lead to incorrect learning and biased beliefs. This includes when agents underweight and overweight new information (Epstein et al. 2010; Rabin and Schrag 1999), selectively pay attention (Schwartzstein 2014), fail to account for redundant information (Bohren 2016; Eyster and Rabin 2010; Gagnon-Bartsch and Rabin 2017), have a coarse model of inference such as the analogy-based expectation equilibrium solution concept (Guarino and Jehiel 2013; Jehiel 2005), overestimate the similarity of others' preferences (Gagnon-Bartsch 2017), or use a non-Bayesian updating heuristic (Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi 2012). Theories of cognitive limitations provide a foundation for certain information-processing biases. Bounded memory can lead to behavior that is consistent with many documented behavioral phenomena, including belief polarization, confirmation bias and stickiness (Wilson 2014). Allowing individuals to selectively interpret signals leads to confirmation bias and conservatism bias (Gottlieb 2015). Bohren (2016) characterizes asymptotic learning outcomes in a model with a single misspecified type who underestimates or overestimates redundant information. Beliefs do not converge when the underestimation is severe, whereas belief may converge to the incorrect state if the overestimation is severe. The robustness result in Bohren (2016) is a special case of the robustness theorems in this paper.

Esponda and Pouzo (2016, 2017) explore the implications of model misspecification for solution concepts. In a Berk-Nash equilibrium, players have a set of (possibly misspecified) models of the world. Individuals play optimally with respect to the model from this set that is the best fit (formally, the model that minimizes relative entropy with respect to the true distribution of outcomes under the equilibrium strategy profile). Nash equilibrium is a special case in which the set of models includes the correctly specified model (which is always the best fit), while our paper corresponds to the case in which each individual has a single (possibly misspecified) model for each state.[10] In our framework, when the belief about the state converges, each type believes that the true state is the state that is the best fit from his set of models, given the frequency of actions and signals that arise when each type is playing optimally with respect to his belief about the state. This is equivalent to a Berk-Nash

[10]In our framework, an agent's type and a state corresponds to a model of the world in the Esponda and Pouzo (2016) framework. Our framework implicitly restricts the feasible models of the world, in that an agent must have the same type in every state.

equilibrium in a dynamic game with infinitely many players.

Madarász and Prat (2016) study optimal mechanism design when the principal's model of the agent's preferences is misspecified, in that it is a finite approximation of the truth. When non-local incentive constraints bind, using the optimal mechanism with respect to a misspecified model can lead to non-vanishing losses, even when the level of misspecification is small. This contrasts with our robustness results, in which the losses from misspecification vanish as the misspecified model approaches the correctly specified model.

An older statistics literature on model misspecification complements recent work. Berk (1966) and Kleijn and van der Vaart (2006) show that when an individual 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. He looks at the limiting distributions of posteriors and establishes conditions for the posterior to converge to the set of distributions that minimize the relative entropy with respect to the true model. These assumptions do not hold in our environment. 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 general model and outlines the individual's decision problem. Section 3 presents the main results, including characterizing the asymptotic learning outcomes under misspecification and establishing robustness. Section 4 develops three applications to explore specific forms of misspecification. Most proofs are in the Appendix.

2 The Common Framework

2.1 The Model

There are two payoff-relevant states of the world, $\omega \in \{L, R\}$, with common prior belief $P(\omega = R) = 1/2$. Nature selects one of these states at the beginning of the game. A countably infinite set of agents $T = \{1, 2, ...\}$ act sequentially and attempt to match the realized state of the world by making a single decision $a_t \in \{L, R\}$.

Information. Agents learn from private information, public information and the actions of other agents. Before choosing an action, each agent t observes the ordered history of past actions $(a_1, ..., a_{t-1})$, a private signal $z_t \in \mathcal{Z}$, where $\mathcal{Z}$ is an arbitrary signal space, and the ordered history of public signals $(y_1, ..., y_t)$, where $y \in \mathcal{Y}$ and $\mathcal{Y}$ is binary. Let $h_t = (a_1, ..., a_{t-1}, y_1, ..., y_{t-1})$ denote the action and public signal history.

Suppose signals $\langle z_t \rangle$ and $\langle y_t \rangle$ are i.i.d. across time, conditional on the state, jointly independent, and drawn according to probability measures $\mu_z^\omega \in \Delta(\mathcal{Z})$ and $\mu_y^\omega \in \Delta(\mathcal{Y})$ in state ω. Assume that no private or public signal perfectly reveals the state, which implies that both μ_z^L, μ_z^R and μ_y^L, μ_y^R are mutually absolutely continuous with common supports, which without loss of generality we assume to be $\mathcal{Z}$ and $\mathcal{Y}$. Finally, assume that some signals are *informative*, which rules out the case where both $d\mu_z^L/d\mu_z^R = 1$ almost surely and $d\mu_y^L/d\mu_y^R = 1$ almost surely.

Given private signal z, the correctly specified private belief that the state is L is $s(z) = 1/(1+d\mu_z^R/d\mu_z^L(z))$. Let c.d.f. $F^\omega(s) \equiv \mu_z^\omega(z|s(z) \leq s)$ denote the distribution of s, and let $[\underline{b}, \bar{b}] \subseteq [0,1]$ denote the convex hull of the common support of private beliefs, $\operatorname{supp} F$. Beliefs are bounded if $0 < \underline{b} < \bar{b} < 1$, and unbounded if $[\underline{b}, \bar{b}] = [0,1]$. Similarly, given public signal y, the correctly specified public belief that that state is L is $\sigma(y) = 1/(1 + d\mu_y^R/d\mu_y^L(y))$, with c.d.f. $G^\omega(\sigma) \equiv \mu_y^\omega(y|\sigma(y) \leq \sigma)$ denoting the distribution of σ. The public signal is binary, so there are at most two public beliefs, $\{\sigma_R, \sigma_L\}$, with $\sigma_R \leq 1/2 \leq \sigma_L$. Let $\operatorname{supp} G$ denote the common support of G^ω.

We will work directly with the correctly specified belief processes $\langle s_t \rangle$ and $\langle \sigma_t \rangle$, where $s_t \equiv s(z_t)$ is referred to as the private signal and $\sigma_t \equiv \sigma(y_t)$ is referred to as the public signal. From Lemma A.1 in Smith and Sorensen (2000) and Lemma 9 in Appendix A.1, $(\operatorname{supp} F, F^L)$ and (σ_R, σ_L) are sufficient for the state signal distributions.

Models of Inference and Payoffs. Agent t has privately observed type $\theta_t \in \Theta$, where Θ is a non-empty finite set and $\pi \in \Delta(\Theta)$ is the distribution over types. Each type θ specifies a payoff structure and a model of inference. In terms of payoffs, all types seek to choose the action that matches the hidden state, but types differ in their costs of errors. Specifically, an agent receives a payoff of 0 if her action matches the realized state. Type θ receives a penalty of $-u_\theta \in (0,1)$ from choosing action L in state R, and a penalty of $-(1-u_\theta)$ from choosing action R in state L.

A agent's model of inference determines how she processes information about the state from signals and prior actions. Type θ's model of inference includes (i) a (possibly misspecified) belief about the likelihood of other types, $\hat{\pi}^\theta \in \Delta(\Theta)$, (ii) a (possibly misspecified) belief about the private signal distribution, $\hat{\mu}_z^{\omega,\theta}(\cdot|p)$, in each state $\omega \in \{L, H\}$ and (iii) a (possibly misspecified) belief about the public signal distribution, $\hat{\mu}_y^{\omega,\theta}(\cdot|p)$, in each state $\omega \in \{L, H\}$, where $p \in [0, 1]$ is the type's belief that the state is L after observing the history but before observing her private signal. This allows an agent's misspecification about the signal distribution to depend on her current belief (for example, to capture confirmation bias). Assume that all distributions are continuous in p under the sup norm.

We place several restrictions on the type of misspecification an agent may have about the state signal distribution. Agents correctly believe that no private or public signal perfectly reveals the state, which implies that both $\hat{\mu}_z^{L,\theta}(\cdot|p), \hat{\mu}_z^{R,\theta}(\cdot|p)$ and $\hat{\mu}_y^{L,\theta}(\cdot|p), \hat{\mu}_y^{R,\theta}(\cdot|p)$ are mutually absolutely continuous for all $p \in [0, 1]$. Agents do not observe signals inconsistent with their models of the world, which implies that both pairs of misspecified measures have full support. Lastly, we say that two pairs of measures have an *equivalent ordinal ranking of signals* if they rank the informativeness of signals in the same order.

Definition 1 (Equivalent Ordinal Ranking of Signals)**.** *Given mutually absolutely continuous probability measures $\mu^L, \mu^R \in \Delta(\mathcal{X})$ and $\nu^L, \nu^R \in \Delta(\mathcal{X})$ on some signal space $\mathcal{X}$, with* $\operatorname{supp}\nu = \operatorname{supp}\mu$*, these pairs of measures have an* equivalent ordinal ranking of signals *if for any $x, x' \in \mathcal{X}$ such that $\frac{d\mu^R}{d\mu^L}(x) \geq \frac{d\mu^R}{d\mu^L}(x')$, then $\frac{d\nu^R}{d\nu^L}(x) \geq \frac{d\nu^R}{d\nu^L}(x')$, with equality iff $\frac{d\mu^R}{d\mu^L}(x) = \frac{d\mu^R}{d\mu^L}(x')$.*

We assume that both the misspecified public and private signal distributions have an equivalent ordinal ranking of signals as the true distributions. This means that if for any two signals $z, z' \in \mathcal{Z}$, if signal z leads to a higher true private belief that the state is L than signal z', then it also leads to a higher misspecified private belief, with an analogous interpretation for the public signal. We make one exception to this assumption to allow for the possibility that a type believes signals are entirely uninformative, $\hat{\mu}_z^{L,\theta} = \hat{\mu}_z^{R,\theta}$ or $\hat{\mu}_y^{L,\theta} = \hat{\mu}_y^{R,\theta}$.

Given private signal z and prior belief $p \in [0, 1]$, the misspecified private belief that the state is L is $\hat{s}^\theta(z, p) = 1/(1 + d\hat{\mu}_z^{R,\theta}/d\hat{\mu}_z^{L,\theta}(z|p))$. By Lemma 8 in Appendix A.1, it is possible to represent the misspecified private belief as a function of the

true private belief, $\hat{s}^\theta(z,p) = r^\theta(s(z),p)$ for a function r^θ that is strictly increasing in its first argument and, when private signals are informative, satisfies $r(\underline{b},\cdot) < 1/2$ and $r(\overline{b},\cdot) > 1/2$. Define the c.d.f. of the perceived distribution of signal s as $\hat{F}^{\omega,\theta}(s) \equiv \hat{\mu}_z^\omega(z|s(z) \leq s)$. Similarly, we can represent the misspecified belief after observing the public signal y and holding prior belief p as $\hat{\sigma}^\theta(y,p) = \rho^\theta(\sigma(y),p)$, where $\rho^\theta(\sigma_R,p) \leq 1/2 \leq \rho^\theta(\sigma_L,p)$ with either both or neither inequalities binding. Therefore, taking (s,σ) as the private and public signals, the tuple $\{r^\theta, \hat{F}^{L,\theta}, \rho^\theta\}$ is sufficient for representing type θ's signal misspecification and we do not need to keep track of the underlying measures on $\mathcal{Z}$ (Lemma 8 in Appendix A.1). The functions $r^\theta(s,\cdot)$ and $\rho^\theta(\sigma,\cdot)$ determine the perceived posterior beliefs following s and σ.[11]

In summary, a type is represented as a tuple $\{u_\theta, \hat{\pi}^\theta, r^\theta, \hat{F}^{L,\theta}, \rho^\theta\}$ that specifies a payoff, belief about other types and model of the state signal distributions.[12] We define several special types. A *rational* type θ_C has a correctly specified model, $\hat{\pi}^C = \pi$, $r^C(s,\cdot) = s$, $\hat{F}^{L,C} = F^L$ and $\rho^C(\sigma,\cdot) = \sigma$. A *noise* type θ_N believes signals and actions are uninformative, $r^N(s,\cdot) = 1/2$, $\rho^N(\sigma,\cdot) = 1/2$ and everyone else is a noise type, $\hat{\pi}^N = \delta_{\theta_N}$. An *autarkic* type θ_A acts solely based on its private signal and does not incorporate the history into its decision-making. It believes everyone else is a noise type, $\hat{\pi}^A = \delta_{\theta_N}$, the public signal is uninformative, $\rho^A(\sigma,\cdot) = 1/2$, and the private signal is informative, $r^A(s,\cdot) \neq 1/2$. We assume $r^\theta(\overline{b},1/2) > u_\theta$ and $r^\theta(\underline{b},1/2) < u_\theta$ to ensure the autarkic type chooses both actions with positive probability (otherwise, it is equivalent to a noise type). There can be multiple autarkic types with different private signal misspecifications and / or an autarkic type with a correctly specified signal distribution. A *sociable* type believes actions are informative and does learn from the history – these are the set of types who are not noise or autarkic types.

Given a set of types Θ, let the vector $(\theta_1,...,\theta_n)$ order Θ such that the first k types are sociable and the remaining $n-k$ types are autarkic or noise types. Let Θ_A denote the set of autarkic types and $\Theta_S = (\theta_1,...,\theta_k)$ denote the set of sociable types.

[11] Further, for any strictly increasing function $r : \operatorname{supp} F_s \to [0,1]$, if $r(\underline{b}) < 1/2$ and $r(\overline{b}) > 1/2$, then there exist a pair of mutually absolutely continuous probability measures with full support on $\Delta(\mathcal{Z})$ that are represented by r.

[12] We can view our restriction to allowing agents to have a single model of signals and inference as the long-run outcome of a learning process across multiple social learning games, where agents begin with multiple types. This is similar to Esponda and Pouzo (2016), who justify Berk-Nash equilibria as the long run outcome of a similar learning process.

We focus on settings where learning is complete in the correctly specified model – that is, an infinite amount of information is revealed through actions or public signals. The following assumption ensures that this is the case by assuming that either there is a positive mass of autarkic types or the public signal is informative.

Assumption 1. *At least one of the following hold: (i)* $\pi(\Theta_A) > 0$*; (ii)* $\sigma_L > 1/2$*.*

We also assume that sociable types have models of inference that believe actions and/or public signals are informative.

Assumption 2. *For each sociable type* θ*, at least one of the following hold: (i)* $\hat{\pi}^\theta(\Theta_A) > 0$*; (ii)* $\rho^\theta(\sigma_L, \cdot) > 1/2$*.*

Finally, we rule out the possibility that an agent observes action choices that are inconsistent with her model of the world.

Assumption 3. *If* $\pi(\Theta_A) > 0$ *or* $[\underline{b}, \bar{b}] = [0, 1]$*, then for each sociable type* θ*, at least one of the following hold: (i)* $\hat{\pi}^\theta(\Theta_A) > 0$*; (ii)* $r^\theta(\underline{b}, \cdot) = 0$ *and* $r^\theta(\bar{b}, \cdot) = 1$*.*

This ensures that when both actions occur with positive probability after any history, every sociable type expects both actions to occur with positive probability.

The timing of the game is as follows. At time t, agent t observes his type θ_t, the history h_t, the private signal s_t, then chooses action a_t. Then public signal y_t is realized, and the history h_{t+1} is updated to include (a_t, y_t).[13]

2.2 The Individual Decision-Problem

Consider an agent of type θ_i who observes history h. Using her model of inference, she computes the probability of this history in each state, $P^i(h|\omega)$, and applies Bayes rule to form the public likelihood ratio $\lambda_i(h)$ that the state is L versus R,

$$\lambda_i(h) = \frac{P^i(h|L)}{P^i(h|R)}.$$

This forms her belief $P^i(L|h) = \lambda_i(h)/(1 + \lambda_i(h))$ for interpreting the private signal when the signal misspecification depends on her current belief, and is also sufficient

[13] Allowing agent t to observe y_t before choosing an action does not change the results, but complicates the notation.

for the history. Next, the agent observes private signal s. Given public belief λ_i, she uses Bayes rule to compute perceived private belief $r^i(s, \lambda_i/(\lambda_i+1))$ that the state is L, where in a slight abuse of notation, we let i index the misspecified posterior belief representation for θ_i. She forms posterior likelihood ratio

$$q^i(\lambda_i, s) = \lambda_i \left(\frac{r^i(s, \lambda_i/(\lambda_i+1))}{1 - r^i(s, \lambda_i/(\lambda_i+1))} \right)$$

that the state is L versus R. The agent maximizes her expected payoff by choosing action L if $q^i(\lambda_i, s) \geq u_i/(1-u_i)$, and action R otherwise. For any public belief λ_i, this decision rule can be represented as a cut-off rule on signal s: choose L if

$$s \geq \overline{s}^i(\lambda_i) \equiv (r^i)^{-1} \left(\frac{u_i}{\lambda_i(1-u_i)+u_i}, \frac{\lambda_i}{1+\lambda_i} \right), \tag{1}$$

and otherwise choose R, where $(r^i)^{-1}$ is the inverse of r^i in the first component. It is common knowledge that each type maximizes payoffs subject to her posterior belief, and therefore, the decision rule of each type is also common knowledge.

2.3 Examples

This framework captures common information-processing biases and models of reasoning about others' action choices. It can be used to study both social and individual learning.[14] The following examples illustrate several types of misspecification.

Level-k and Cognitive Hierarchy. Level-k corresponds to a model in which agents have a misspecified belief about the distribution of types. Level-0 is the noise type. Level-1 believes all other agents are the noise type and behaves as the autarkic type. Level-2 believes all other agents are the autarkic type and interprets all prior actions as independent private signals. Level-3 believes all other 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 agent have a perceived distribution that can place positive probability on types of level-0 through k-1.

[14] Social learning settings are captured by types with informative private signals and non-trivial models of inference about other agents. Individual learning settings are captured by informative public signals, uninformative private signals and agents who do not learn from actions. This is isomorphic to a setting with a single long-run agent of each type.

Partisan Bias. Agents systematically slant signals towards one state. For example, a parameterization that slants signals towards state L is $r(s,p) = s^\nu$, where $\nu < 1$.

Confirmation Bias. Agents overweight information in favor of their prior. That is, they overweight signals in favor of state L when the prior is high, and underweight signals in favor of state L when the prior is low. For example, a symmetric parameterization is $r(s,p) \geq s$ if $p > 1/2$ and $r(s,p) \leq s$ if $p < 1/2$.

Under/Overconfidence. Agents either underweight or overweight signals. For example,

$$\frac{r(s,p)}{1-r(s,p)} = \left(\frac{s}{1-s}\right)^\nu,$$

where $\nu \in [0,1)$ corresponds to underweighting and $\nu \in (1,\infty)$ corresponds to overweighting.

False Consensus Effect. Agents overweight the likelihood that others have similar preferences, when in reality preferences are heterogeneous. For example, there are two types of agents with different costs of choosing the incorrect action, $u_1 \neq u_2$. However, all agents believe that other agents have the same cost, $\hat{\pi}^1(\theta_1) = 1$ and $\hat{\pi}^2(\theta_2) = 1$.

Pluralistic Ignorance. Agents underweight the likelihood that others have similar preferences. For example, all agents have cost of choosing an incorrect action u_1, but believe that others have cost of choosing an incorrect action u_2.

3 Learning Dynamics

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

3.1 The Likelihood Ratio

Let λ_i denote the public likelihood ratio of type θ_i, and define $\boldsymbol{\lambda} \equiv (\lambda_1, ..., \lambda_k)$ as the vector of public likelihood ratios for sociable types (note $\lambda_i = 1$ for all autarkic or noise types θ_i). Recall that the public likelihood ratio for type θ_i after observing history h

depends on how type θ_i perceives the probability of h in each state, $\lambda_i(h) = \frac{P^i(h|L)}{P^i(h|R)}$. In order to calculate $\lambda_i(h)$, we need to determine how $P^i(h|\omega)$ depends on θ_i's model of inference.

Misspecification introduces a wedge between the perceived and true probability of observing each action in h. An agent's type determines how she interprets each action, while the true probability of each action depends on the true signal and type distributions. The true probability that an agent of type θ_i chooses action L when she has public likelihood ratio λ and the state is ω is equal to the probability of observing a private signal below the cutoff $\overline{s}^i(\lambda)$ from decision rule (1). This is determined by the true signal distribution, $F^\omega(\overline{s}^i(\lambda))$. However, type θ_j *believes* that θ_i chooses action L with probability $\hat{F}^{\omega,j}(\overline{s}^i(\lambda))$. This is θ_j's perceived probability of observing a private signal is below θ_i's cutoff. Similarly, the probability of action R is equal to the probability of observing a signal above $\overline{s}^i(\lambda)$, $1 - F^\omega(\overline{s}^i(\lambda))$. The perceived probability is defined analogously.

Given $\boldsymbol{\lambda}$ and state ω, the true probability of action L across all types depends on the true distribution of types,

$$\psi(L|\omega, \boldsymbol{\lambda}) \equiv \sum_{j=1}^{n} F^\omega(\overline{s}^j(\lambda_j))\pi(\theta_j).$$

Similarly, the probability of action R is $\psi(R|\omega, \boldsymbol{\lambda}) = 1 - \psi(L|\omega, \boldsymbol{\lambda})$. Type θ_i's perceived probability of action L depends on her perceived distribution of types $\hat{\pi}^i$ and signals $F^{\omega,i}$,

$$\hat{\psi}_i(L|\omega, \boldsymbol{\lambda}) \equiv \sum_{j=1}^{n} \hat{F}^{\omega,i}(\overline{s}^j(\lambda_j))\hat{\pi}^i(\theta_j).$$

Similarly, her perceived probability of action R is $\hat{\psi}_i(R|\omega, \boldsymbol{\lambda}) = 1 - \hat{\psi}_i(L|\omega, \boldsymbol{\lambda})$.

Each type interprets the history and forms a public likelihood ratio using her perceived probability of actions and public signals. Given a likelihood ratio $\boldsymbol{\lambda}_t$, action a_t and public signal σ_t in period t, the likelihood ratio in the next period is $\boldsymbol{\lambda}_{t+1} = \phi(a_t, \sigma_t, \boldsymbol{\lambda}_t)$, where $\phi : \{L, R\} \times \{\sigma_L, \sigma_R\} \times \mathbb{R}^n_+ \to \mathbb{R}^n_+$, with

$$\phi_i(a, \sigma, \boldsymbol{\lambda}) \equiv \lambda_i \left(\frac{\hat{\psi}_i(a|L, \boldsymbol{\lambda})}{\hat{\psi}_i(a|R, \boldsymbol{\lambda})} \right) \left(\frac{\rho^i(\sigma, \lambda_i/(\lambda_i + 1))}{1 - \rho^i(\sigma, \lambda_i/(\lambda_i + 1))} \right). \tag{2}$$

The transition probability for the likelihood ratio depends on the true probability

of each action and public signal. In a slight abuse of notation, let $\psi(a, \sigma|\omega, \boldsymbol{\lambda}) \equiv \psi(a|\omega, \boldsymbol{\lambda})dG^{\omega}(\sigma)$ denote the probability of action a and public signal σ when the state is ω and the current value of the likelihood ratio is $\boldsymbol{\lambda}$, with analogous notation for $\hat{\psi}(a, \sigma|\omega, \boldsymbol{\lambda})$. Given $\{a_t, \sigma_t, \boldsymbol{\lambda}_t\}$, the process transitions to $\{a_{t+1}, \sigma_{t+1}, \phi(a_t, \sigma_t, \boldsymbol{\lambda}_t)\}$ with probability $\psi(a_{t+1}, \sigma_{t+1}|\omega, \phi(a_t, \sigma_t, \boldsymbol{\lambda}_t))$.[15]

The joint stochastic process $\langle a_t, \sigma_t, \boldsymbol{\lambda}_t\rangle_{t=1}^{\infty}$ is a discrete-time Markov chain starting at $\boldsymbol{\lambda}_1 = \mathbf{1}$. The stochastic properties of this Markov chain determine the learning dynamics for each type. The equations of motion are state-dependent and nonlinear, due to the dependence of equilibrium actions on the current belief vector. This presents a technical challenge, as the process fails to satisfy standard conditions from the existing literature on Markov chains (despite having a countable state space). In the following sections, we use results on the stability of nonlinear stochastic difference equations to characterize the limiting behavior of the likelihood ratio for each type.

3.2 Main Results

Asymptotic Learning Characterization. We first define several asymptotic learning outcomes. Let *incorrect* learning (for type θ_i) denote the event where $\boldsymbol{\lambda}_t \to \infty^k$ ($\lambda_{i,t} \to \infty$), *correct* learning (for type θ_i) denote the event where $\boldsymbol{\lambda}_t \to 0^k$ ($\lambda_{i,t} \to 0$) and *incomplete* learning (for type θ_i) denote the event where $\boldsymbol{\lambda}_t$ ($\lambda_{i,t}$) does not converge or diverge, where 0^k (∞^k) denotes the vector of all zeros (all ∞). Agents asymptotically *agree* when all types have the same limit beliefs, $\boldsymbol{\lambda} \in \{0^k, \infty^k\}$, and agents asymptotically *disagree* when different types have different limit beliefs, $\boldsymbol{\lambda} \in \{0, \infty\}^k \setminus \{0^k, \infty^k\}$. Assumptions 1 and 2 rule out $\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}$ for any $\boldsymbol{\lambda} \notin \{0, \infty\}^k$.

Our main result characterizes the asymptotic learning outcomes in misspecified models. In correctly specified models, the likelihood ratio is a martingale, and the Martingale Convergence Theorem provides a powerful tool to characterize its limit behavior. 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 breaks the equality condition. Therefore, an alternative approach is necessary to characterize limit beliefs. The characterization we develop depends on two expressions

[15] When an agent's interpretation of signals depends on her current belief, this set-up implicitly assumes that the agent uses belief λ_t to interpret both s_t and σ_t. This is for notational simplicity. The results are unchanged if the agent uses λ_t to interpret s_t, and an interim belief that incorporates the information from a_t to interpret σ_t.

that are straightforward to calculate from the primitives of the model – the type space and the signal distributions.

The first expression, the expected change in the log likelihood ratio, determines whether a candidate learning outcome is *locally stable*, in that the likelihood ratio converges to this limit belief with positive probability, from a neighborhood of the belief. Without loss of generality, suppose that the realized state is $\omega = R$. For type θ_i, the expected change in the log likelihood ratio at $\boldsymbol{\lambda} \in \{0, \infty\}^k$ depends on the perceived and true probability of each action,

$$\gamma_i(\boldsymbol{\lambda}) \equiv \sum_{(a,\sigma) \in \{L,R\} \times \{\sigma_L, \sigma_R\}} \psi(a, \sigma | R, \boldsymbol{\lambda}) \log \left(\frac{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda})}{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})} \right). \tag{3}$$

Let $\boldsymbol{\gamma}(\boldsymbol{\lambda}) = (\gamma_1(\boldsymbol{\lambda}), ..., \gamma_k(\boldsymbol{\lambda}))$. The sign of each component of $\boldsymbol{\gamma}(\boldsymbol{\lambda})$ determines local stability. An outcome $\boldsymbol{\lambda}$ is locally stable if and only if $\gamma_i(\boldsymbol{\lambda})$ is negative for types with $\lambda_i = 0$ and positive for types with $\lambda_i = \infty$. Let $\mathit{\Lambda}$ denote the set of learning outcomes that are locally stable,

$$\mathit{\Lambda} \equiv \{\boldsymbol{\lambda} \in \{0, \infty\}^k | \gamma_i(\boldsymbol{\lambda}) < 0 \text{ if } \lambda_i = 0 \text{ and } \gamma_i(\boldsymbol{\lambda}) > 0 \text{ if } \lambda_i = \infty\}. \tag{4}$$

We establish that if $\langle \boldsymbol{\lambda}_t \rangle_{t=1}^{\infty}$ converges, then it must converge to a limit random variable whose support lies in $\mathit{\Lambda}$. Intuitively, in order for the likelihood ratio to converge to a candidate limit point with positive probability, the likelihood ratio must move towards this limit point in expectation from nearby beliefs. It is straightforward to compute $\mathit{\Lambda}$ from the primitives of the model. This result significantly simplifies the set of possible limit beliefs.

We are interested in a characterization of asymptotic learning that is independent of the initial belief, and therefore, we need a tighter notion of stability. A learning outcome is *globally stable* if the likelihood converges to this limit belief with positive probability, from *any* initial belief. For an agreement outcome, we show that local stability is necessary and sufficient for global stability. Therefore, computing $\mathit{\Lambda}$ is the only calculation necessary to determine whether correct or incorrect learning outcomes arise. These learning outcomes arise with positive probability if and only if the corresponding limit beliefs, 0^k or ∞^k, are in $\mathit{\Lambda}$. For disagreement outcomes, a failure of local stability is sufficient to ensure that the outcome almost surely does not arise, but an additional condition is necessary to establish when the outcome arises,

from any initial belief. For a disagreement outcome to be globally stable, it must be possible to separate the beliefs for the types converging to 0 and the types converging to ∞.

The second expression, which we term the total informativeness ranking, is a condition on how the type space is ordered that is sufficient to separate beliefs. To derive the expression, first define a pairwise order that relates how two types interpret actions and signals.

Definition 2 (Pairwise Informativeness Order)**.** *Given* $\boldsymbol{\lambda} \in \{0, \infty\}^k$, $\theta_i \succeq_{\boldsymbol{\lambda}} \theta_j$ *iff*

$$\left|\log\left(\frac{\hat{\psi}_i(R, \sigma_R|L, \boldsymbol{\lambda})}{\hat{\psi}_i(R, \sigma_R|R, \boldsymbol{\lambda})}\right)\right| \geq \left|\log\left(\frac{\hat{\psi}_j(R, \sigma_R|L, \boldsymbol{\lambda})}{\hat{\psi}_j(R, \sigma_R|R, \boldsymbol{\lambda})}\right)\right|$$

and

$$\left|\log\left(\frac{\hat{\psi}_i(L, \sigma_L|L, \boldsymbol{\lambda})}{\hat{\psi}_i(L, \sigma_L|R, \boldsymbol{\lambda})}\right)\right| \leq \left|\log\left(\frac{\hat{\psi}_j(L, \sigma_L|L, \boldsymbol{\lambda})}{\hat{\psi}_j(L, \sigma_L|R, \boldsymbol{\lambda})}\right)\right|.$$

In other words, the most informative action and public signal in favor of state R, (R, σ_R), which unambiguously decreases the likelihood ratio, is more informative for type θ_i than type θ_j. The most informative action and public signal in favor of state L, (L, σ_L), which unambiguously increases the likelihood ratio, is more informative for type θ_j than type θ_i.

We use the pairwise informativeness order to define a ranking over all types. A disagreement outcome is *total informativeness ranked* if the least pairwise informative type in the set of types whose limit beliefs converge to 0 is pairwise more informative than the greatest pairwise informative type whose limit beliefs converge to ∞.

Definition 3 (Total Informativeness Rank)**.** *A disagreement vector* $\boldsymbol{\lambda} = (0^m, \infty^{k-m})$ *is* total informativeness ranked *if for* $i = 1, ..., m$ *and* $j = m+1, ..., k$ *and for either* $\boldsymbol{\lambda}_A \in \{0^k, \infty^k\}$,

1. *There exists an* $i^* \leq m$ *such that* $\theta_i \succeq_{\boldsymbol{\lambda}_A} \theta_{i^*}$ *for all* $i \leq m$ *and there exists an* $j^* > m$ *such that* $\theta_{j^*} \succeq_{\boldsymbol{\lambda}_A} \theta_j$ *for all* $j > m$.

2. $\theta_{i^*} \succ_{\boldsymbol{\lambda}_A} \theta_{j^*}$.

For any disagreement outcome in Λ, the total informativeness rank is a sufficient condition for global stability. If the information from actions and public signals

arrives at a rate such that the least informative type with $\lambda_i = 0$ moves towards 0 at a faster rate than the most informative type with $\lambda_j = \infty$ moves towards 0, then it is possible to find a finite sequence of actions and public signals that sufficiently separate beliefs. Once again, this condition is straightforward to verify from the primitives of the model.

Given Λ and the total informativeness ranking, Theorem 1 characterizes the set of asymptotic learning outcomes.

Theorem 1. *Assume Assumptions 1, 2 and 3 and suppose $\omega = R$.*

- Correct *learning occurs with positive probability if and only if* $0^k \in \Lambda$.
- Incorrect *learning occurs with positive probability if and only if* $\infty^k \in \Lambda$.
- *Agents* disagree *with positive probability if there exists a disagreement vector* $\boldsymbol{\lambda} \in \Lambda$ *that is total informativeness ranked, and agents almost surely do not* disagree *if* Λ *contains no disagreement vectors.*
- Incomplete *learning (non-convergence) occurs almost surely if* Λ *is empty, and beliefs* converge *almost surely if* Λ *is non-empty and either (i)* $0^k \in \Lambda$*, (ii)* $\infty^k \in \Lambda$ *or (iii)* $\exists\, \boldsymbol{\lambda} \in \Lambda$ *that is total informativeness ranked.*

The conditions for correct and incorrect learning are tight. These learning outcomes obtain if and only if the respective limit beliefs are in Λ. Disagreement outcomes are more challenging. We establish a sufficient condition for disagreement to occur, and a sufficient condition for disagreement not to occur. A general necessary and sufficient condition is not possible. In particular, we cannot determine whether the likelihood ratio converges with positive probability to a disagreement vector that is in Λ, but is not total informativeness ranked. In such a case, whether disagreement occurs can depend on initial beliefs. In Section 4, we demonstrate that this will not be an issue for certain forms of misspecification. In these applications, all locally stable disagreement vectors are total informativeness ranked, and there is no wedge between the two sufficient conditions. Therefore, Λ fully characterizes asymptotic learning outcomes.

In Section 3.3, we outline the proof for Theorem 1 through a series of Lemmas. Before proceeding to the proof, we present several additional results.

An immediate consequence of Theorem 1 is that learning is *complete* – correct learning occurs almost surely – in the correctly specified model ($\Lambda = \{0\}$). More generally, even if some types of agents have misspecified models, these misspecified types do not interfere with asymptotic learning for the type that has a correctly specified model. This type has correct beliefs about the distribution of misspecified types, and is able to probabilistically parse out the actual information conveyed by actions. Therefore, learning is complete for the correctly specified type, independent of the other types' outcomes.

Corollary 1. *Assume Assumptions 1, 2 and 3. Learning is* complete *for the correctly specified type θ_C, $\lambda_{C,t} \to 0$ almost surely.*

Robustness of Complete Learning. Our second set of main results establish that the asymptotic learning properties of the correctly specified model are robust to some misspecification, in that learning is complete for sociable types with nearby misspecified models. In correctly specified models, the martingale property of the likelihood ratio, coupled with the Martingale Convergence Theorem, is used to establish complete learning. The likelihood ratio is no longer a martingale with even an arbitrarily small amount of misspecification. However, this is a sufficient, but not necessary, condition for complete learning. From Theorem 1, the behavior of the *log* likelihood ratio yields necessary and sufficient conditions for complete learning. If the log likelihood ratio is a supermartingale, then learning is complete. An even weaker conditions is possible. The log likelihood ratio only needs to decrease in expectation at a *finite* set of vectors in the belief space for learning to be complete. Namely, at each candidate learning outcome $\boldsymbol{\lambda} \in \{0, \infty\}^k$.[16] Due to the concavity of the log operator, if the likelihood ratio is a martingale, then the log likelihood ratio satisfies this condition. Therefore, the correctly specified model is a special case of the set of models in which complete learning obtains.

Theorem 2 presents two sets of sufficient conditions for complete learning to obtain for all sociable types. First, if all sociable types have perceived type and public signal distributions close enough to the true distribution, then learning is complete. This condition places no restrictions on the perceived private signal distributions. Second,

[16] An even weaker condition is possible. The log likelihood ratio needs to decrease in expectation at 0^k, and for any other learning outcome $\boldsymbol{\lambda}$, the log likelihood ratio cannot increase in expectation for all types with $\lambda_i = \infty$ and decrease in expectation for all types with $\lambda_i = 0$. This is precisely the condition that ensures that 0^k is the unique locally stable vector.

if all types have perceived private and public signal distributions close enough to the true distribution, and sociable types' perceived frequency of autarkic types is approximately correct, then learning is complete. This condition holds even if sociable types have very incorrect beliefs over the distribution of different types. For example, all sociable types are type θ, but believe all sociable types are type $\theta' \neq \theta$.

Theorem 2. *Assume Assumptions 1, 2 and 3 and suppose $\omega = R$.*

1. *There exists a $\delta > 0$ such that if $||\hat{\pi}^i - \pi|| < \delta$ and $||\rho^i - I|| < \delta$ for all sociable types θ_i, then learning is* complete.

2. *There exists a $\delta > 0$ such that if $||r^i - I|| < \delta$, $||\hat{F}^{L,i} - F^L|| < \delta$ and $||\rho^i - I|| < \delta$ for all types θ_i, and $|\hat{\pi}^i(\Theta_A) - \pi(\Theta_A)| < \delta$ for all sociable types θ_i, then learning is* complete.

where $||\cdot||$ denotes the supremum metric, and $I : [0,1] \to [0,1]$ denotes an identity function, $I(s) = s$.[17]

This follows from the continuity of $\boldsymbol{\gamma}$ in each type θ_i's belief over the signal and type distributions. Since $\boldsymbol{\gamma}$ is the key expression used to calculate Λ, and $\Lambda = \{0\}$ in the correctly specified model, $\Lambda = \{0^k\}$ is maintained when some misspecification is introduced.

More generally, correct learning obtains for any form of misspecification in which each type's perceived probability of actions and public signals is close enough to the true probability. We present a more general robustness theorem, which depends on the equilibrium objects ψ and $\hat{\psi}_i$. As long as $\hat{\psi}_i$ is close to ψ at all of the candidate limit beliefs $\{0, \infty\}^k$, learning is complete.

Theorem 3. *Assume Assumptions 1, 2 and 3 and suppose $\omega = R$. There exists a $\delta > 0$ such that if for each sociable type θ_i, if $|\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda}) - \psi(a, \sigma | R, \boldsymbol{\lambda})| < \delta$ for all $(a, \sigma, \boldsymbol{\lambda}) \in \{L, R\} \times \{\sigma_L, \sigma_R\} \times \{0, \infty\}^k$, then learning is* complete *for all sociable types.*

In addition to showing that correct learning is robust to small misspecification, the tools in this paper allow for a precise characterization of exactly how robust correct learning is to different types of misspecification. In many interesting examples,

[17] Given set X and metric space Y, the supremum metric between two bounded functions $f : X \to Y$ and $g : X \to Y$ is $||f - g|| = \sup_{x \in X} |f(x) - g(x)|$.

including those developed in Section 4, the region where correct learning is the unique outcome can be quite large. Even agents with misspecified models can still learn the correct state of the world in the long-run.

3.3 Proof of Theorem 1

We establish Theorem 1 through a series of Lemmas. In Lemma 1, we characterize the stationary vectors of the likelihood ratio, which are candidate limit points of $\langle \boldsymbol{\lambda}_t \rangle$. In Lemma 2, we establish when a stationary vector is locally stable. Local stability depends on $\mathbf{fl}$ defined in (4), and Lemma 2 establishes that the set Λ defined in (4) is the set of locally stable vectors. To fully characterize asymptotic learning outcomes, we need to determine when the likelihood ratio converges to a stationary vector, from *any* initial belief. Lemma 3 establishes that global stability immediately follows from local stability for agreement vectors. Lemmas 4 and 5 establish a sufficient condition for a locally stable disagreement vector to also be globally stable. Lemma 6 establishes that when there is at least one globally stable vector, the likelihood ratio converges almost surely. Finally, Lemma 7 rules out convergence to non-stationary vectors.

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

Definition 4. *A vector* $\boldsymbol{\lambda}$ *is* stationary *if for all* $(a, \sigma) \in \{L, R\} \times \{\sigma_R, \sigma_L\}$*, either (i)* $\psi(a, \sigma|\omega, \boldsymbol{\lambda}) = 0$ *or (ii)* $\phi_i(a, \sigma, \boldsymbol{\lambda}) = \boldsymbol{\lambda}$ *for for all* $\theta_i \in \Theta_S$*.*

By Assumptions 1 and 2, actions and/or public signals are informative at any interior belief. Therefore, the set of stationary vectors of the likelihood ratio correspond to each type placing probability 1 on either state L or state R.

Lemma 1. *Assume Assumptions 1 and 2. The set of stationary vectors for* $\boldsymbol{\lambda}$ *are* $\{0, \infty\}^k$*.*

Next, we determine when the likelihood ratio converges to a stationary vector with positive probability. We say stationary vector $\boldsymbol{\lambda}$ is *locally stable* if the process $\langle \boldsymbol{\lambda}_t \rangle$ converges to $\boldsymbol{\lambda}$ with positive probability when $\boldsymbol{\lambda}_1$ is in a neighborhood of $\boldsymbol{\lambda}$.

Definition 5. *A stationary vector* $\boldsymbol{\lambda} \in \{0, \infty\}^k$ *is* locally stable *if there exists an* $\varepsilon > 0$*,* $M > 0$ *and neighborhood* $N = \prod_{i=1}^k N_i$ *with* $N_i = \{\lambda | \lambda < \varepsilon\}$ *if* $\lambda_i = 0$ *and* $N_i = \{\lambda | \lambda > M\}$ *if* $\lambda_i = \infty$*, such that* $P(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda} | \boldsymbol{\lambda}_1 \in N) > 0$*.*

Recall from (3) that $\gamma_i(\boldsymbol{\lambda})$ is the expected change in the log likelihood ratio for type θ_i when $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}$, with $\boldsymbol{\gamma} = (\gamma_1, ..., \gamma_k)$. Lemma 2 establishes the relationship between the local stability of $\boldsymbol{\lambda}$ and the sign of $\gamma_i(\boldsymbol{\lambda})$ for each sociable type.

Lemma 2. *Suppose $\omega = R$ and let $\boldsymbol{\lambda} \in \{0, \infty\}^k$.*

1. *If $\gamma_i(\boldsymbol{\lambda}) < 0$ for all $\theta_i \in \Theta_S$ such that $\lambda_i = 0$ and $\gamma_i(\boldsymbol{\lambda}) > 0$ 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 $\gamma_i(\boldsymbol{\lambda}) > 0$ or $\lambda_i = \infty$ and $\gamma_i(\boldsymbol{\lambda}) < 0$, then $\boldsymbol{\lambda}$ is not locally stable and $P(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}) = 0$.*

Intuitively, if the likelihood ratio moves towards a stationary point in expectation when it is within a neighborhood of the stationary point, then the stationary point is locally stable; otherwise it is not. The likelihood ratio almost surely does not converge to stationary points that are not locally stable. Given Lemma 2, the set Λ defined in (4) is generically the set of locally stable vectors.

Local stability establishes convergence when the likelihood ratio is near a stationary vector. However, we are interested in determining whether convergence to a stationary vector occurs from any initial value of the likelihood ratio. We say a stationary vector is *globally stable* if the likelihood ratio converges to it with positive probability from any initial value.

Definition 6. *A stationary vector $\boldsymbol{\lambda} \in \{0, \infty\}^k$ is* globally stable *if for any initial value $\boldsymbol{\lambda}_1 \in (0, \infty)^k$, $P(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}) > 0$.*

Lemma 2 established that if the likelihood ratio converged to $\boldsymbol{\lambda}$ with positive probability, then $\boldsymbol{\lambda}$ is locally stable. Therefore, the set of globally stable stationary vectors is a subset of the set of locally stable stationary vectors. It remains to establish when local stability implies global stability. For stationary agreement vectors, $\boldsymbol{\lambda} \in \{0^k, \infty^k\}$, global stability immediately follows from local stability.

Lemma 3. *For $\boldsymbol{\lambda} \in \{0^k, \infty^k\}$, if $\boldsymbol{\lambda}$ is locally stable, then $\boldsymbol{\lambda}$ is globally stable.*

All types update their beliefs in the same directly following either an L action and public signal σ_L, or an R action and public signal σ_R. Therefore, it is possible to push the likelihood ratio arbitrarily close to a stationary agreement vector with positive probability by constructing a finite sequence of action and public signal pairs. Once

the likelihood ratio is close enough to the agreement vector, local stability guarantees convergence.

Local stability may not imply global stability for stationary disagreement vectors, $\boldsymbol{\lambda} \in \{0, \infty\}^k \setminus \{0^k, \infty^k\}$. In contrast to agreement vectors, 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 vector. For example, if two types are sufficiently close to each other, then disagreement may arise if their initial beliefs are very far apart, but may not be possible if their initial beliefs are close together. Therefore, there may exist initial values of the likelihood ratio such that a locally stable disagreement vector is reached with probability zero.

Lemma 4 establishes a sufficient condition for the global stability of a stationary disagreement vector when there are two sociable types, $k = 2$. Define the matrix

$$A(\boldsymbol{\lambda}) \equiv \begin{pmatrix} \log \frac{\hat{\psi}_1(L,\sigma_L|L,\boldsymbol{\lambda})}{\hat{\psi}_1(L,\sigma_L|R,\boldsymbol{\lambda})} & \log \frac{\hat{\psi}_1(R,\sigma_R|L,\boldsymbol{\lambda})}{\hat{\psi}_1(R,\sigma_R|R,\boldsymbol{\lambda})} \\ \log \frac{\hat{\psi}_2(L,\sigma_L|L,\boldsymbol{\lambda})}{\hat{\psi}_2(L,\sigma_L|R,\boldsymbol{\lambda})} & \log \frac{\hat{\psi}_2(R,\sigma_R|L,\boldsymbol{\lambda})}{\hat{\psi}_2(R,\sigma_R|R,\boldsymbol{\lambda})} \end{pmatrix}. \tag{5}$$

Lemma 4. *Suppose $k = 2$.*

1. *If $(0, \infty)$ is locally stable and either $\det(A(0,0)) > 0$ or $\det(A(\infty,\infty)) > 0$, then $(0, \infty)$ is globally stable.*

2. *If $(\infty, 0)$ is locally stable and either $\det(A(0,0)) < 0$ or $\det(A(\infty,\infty)) < 0$, then $(\infty, 0)$ is globally stable.*

The determinant conditions in Lemma 4 guarantee that the rate of information arrival is such that it is possible to push beliefs of different types arbitrarily far apart. As before, once the likelihood ratio is sufficiently close to the disagreement vector, then convergence obtains when the disagreement vector is locally stable. Lemma 5 builds on Lemma 4 to establish a sufficient condition for global stability of a stationary disagreement vector when there are more than two sociable types.

Lemma 5. *If disagreement vector $\boldsymbol{\lambda} = (0^m, \infty^{k-m})$ is locally stable and total informativeness ranked, then $(0^m, \infty^{k-m})$ is globally stable.*[18]

Finally, if there is at least one globally stable vector, then the likelihood ratio converges almost surely.

[18]While total informativeness rank is a simple and easy to parse condition, the proof would go through almost unchanged under a more general sufficient condition. Disagreement is globally stable

Lemma 6. *Suppose Λ is non-empty and either (i) $0^k \in \Lambda$, (ii) $\infty^k \in \Lambda$ or (iii) $\exists \boldsymbol{\lambda} \in \Lambda$ that is total informativeness ranked. Then for any initial value $\boldsymbol{\lambda}_1 \in (0,\infty)^k$, there exists a random variable $\boldsymbol{\lambda}$ with $\text{supp}(\boldsymbol{\lambda}) \subset \Lambda$ such that $\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}$ almost surely.*

If there are no locally stable vectors, then the likelihood ratio almost surely does not converge, as Lemma 2 rules out convergence to non-locally stable vectors and the following lemma rules out convergence to non-stationary vectors.

Lemma 7. *If $\boldsymbol{\lambda} \in (0,\infty)^k$, then $P(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}) = 0$.*

Theorem 1 immediately follows. The proofs of Lemmas 1-7 are in Appendix A.2.

4 Applications

We next explore learning in three applications. We illustrate how to calculate the set of asymptotic learning outcomes, Λ, and derive comparative statics for how this set varies with the extent of the misspecification.

4.1 Level-k Model of Inference

Set-up. Suppose that agent types correspond to a level-k model of inference with four levels of reasoning, $\Theta = \{\theta_0, \theta_1, \theta_2, \theta_3\}$.[19,20] Level-0 is a noise type used to model the beliefs of other types, but does not actually exist in the population. The level-1, 2 and 3 types correctly interpret private information, but have misspecified beliefs over the type distribution. Level-1 is an autarkic type – it draws inference solely from its private signal, and is not sophisticated enough to draw inference from the actions of others. This is modeled by specifying that level-1 types believe prior actions are uninformative i.e. all other agents are type θ_0, $\hat{\pi}^1(\theta_0) = 1$.

---

if it is locally stable and the set

$$\{b \in (0,\infty)^4 : A(\boldsymbol{\lambda})b = c \text{ for some } c \text{ such that for any } i \leq m,\ j > m,\ c_i < c_j\},$$

where $(A(\boldsymbol{\lambda}))_{ij} = \log \frac{\hat{\psi}_i(a_j,\sigma_j|L,\boldsymbol{\lambda})}{\hat{\psi}_i(a_j,\sigma_j|L,\boldsymbol{\lambda})}$, is non-empty for some $\lambda \in \{0^k, \infty^k\}$.

[19] Camerer, Ho, and Chong (2004); Costa-Gomes, Crawford, and Iriberri (2009).

[20] It is possible to allow for higher levels, $k > 3$. However, empirical and experimental studies of level-k models rarely find evidence of types above level-3. Penczynski (forthcoming) analyzes experimental data on social learning and finds evidence of level-1, level-2 and level-3 types, with a modal type of level-2, across several learning settings.

Level-2 and level-3 are the sociable types. They believe that actions are informative, but have incorrect models of how others draw inference. Actions reflect both private information and information from the actions of others, but level-2 types do not understand this strategic link. They believe actions solely reflect private information and fail to account for repeated information stemming from prior agents observing a subset of the same action history. This leads level-2 types to overweight the informativeness of actions – their perceived type distribution places probability one on type θ_1, $\hat{\pi}^2(\theta_1) = 1$.

Level-3 types have the most sophisticated reasoning. They understand that some agents act solely based on their private information and some agents misunderstand the strategic link between action choices and the history. However, they do not account for the fact that there are other agents with the same level of reasoning. They believe agents are type θ_2 with probability $p \in [0, 1)$, $\hat{\pi}^3(\theta_2) = p$, and type θ_1 with probability $1 - p$, $\hat{\pi}^3(\theta_1) = 1 - p$. If p is high, they believe most actions are from level-2 types and underweight the informativeness of actions to counteract the overweighting behavior of these level-2 types. If p is low, they believe most actions are from level-1 types and, similar to level-2 types, overweight the informativeness of actions.

To close the model, assume the true distribution of types is equally distributed across levels 1-3, $\pi(\theta_1) = \pi(\theta_2) = \pi(\theta_3) = 1/3$, there are no noise types, $\pi(\theta_0) = 0$, private signals are symmetrically distributed across states, $F^L(1/2) = 1 - F^R(1/2)$, and there are no public signals. All agents have common error penalty $u = 1/2$.[21] Level-1 types occur with positive probability, and level-2 and level-3 types believe that level-1 types occur with positive probability, so Assumptions 1 - 3 are satisfied.

Action Choices and Beliefs. Level-1 types incorporate solely their private information into their decision and their public belief is constant across time, $\lambda_{1,t} = 1$ for all t. When $\theta_t = \theta_1$, the agent chooses $a_t = L$ if $s_t \geq 1/2$ and the informativeness of her action is independent of the history.

Level-2 types believe past actions are from level-1 types, and therefore, are independent and identically distributed. Their perceived probability of each R action in the history is the probability that a level-1 type chooses action R, $\hat{\psi}_2(R|\omega, \boldsymbol{\lambda}) =$

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

$F^\omega(1/2)$, and their perceived probability of each L action is the probability that a level-1 type chooses action L, $\hat{\psi}_2(L|\omega, \boldsymbol{\lambda}) = 1 - F^\omega(1/2)$, which are independent of $\boldsymbol{\lambda} = (\lambda_2, \lambda_3)$. Given the symmetry assumption on F^R and F^L, the difference between the number of R and L actions, $n_t \equiv \sum_{\tau=1}^{t-1} \mathbb{1}_{a_\tau=R} - \mathbb{1}_{a_\tau=L}$, is a sufficient statistic for the public belief of level-2 types,

$$\lambda_{2,t} = \left(\frac{F^L(1/2)}{F^R(1/2)}\right)^{n_t}.$$

When $\theta_t = \theta_2$, she chooses $a_t = L$ if $s_t \geq 1/(\lambda_{2,t} + 1)$. Note the informativeness of level-2 actions *does* depend on the history through n_t.

Level-3 types believe past actions are from either level-1 or level-2 types. Their perceived probability of an R action at time t is a weighted average of the probability that level-1 and level-2 types choose action R,

$$\hat{\psi}_3(R|\omega, \boldsymbol{\lambda}_t) = pF^\omega\left(\frac{1}{\lambda_{2,t}+1}\right) + (1-p)F^\omega(1/2).$$

The perceived probability of an L action is analogous. Both depend on the public belief of the level-2 type, $\lambda_{2,t}$. Therefore, how level-3 types update their public belief following an action depends on the current belief of level-2 types. For example, following an R action,

$$\lambda_{3,t} = \lambda_{3,t-1}\left(\frac{pF^L\left(\frac{1}{\lambda_{2,t}+1}\right) + (1-p)F^L(1/2)}{pF^R\left(\frac{1}{\lambda_{2,t}+1}\right) + (1-p)F^R(1/2)}\right).$$

When $\theta_t = \theta_3$, she chooses $a_t = L$ if $s_t \geq 1/(\lambda_{3,t} + 1)$.

The *actual* probability of an R action at time t depends on the true distribution over types as well as the signal cut-off for each type,

$$\psi(R|\omega, \boldsymbol{\lambda}_t) = \frac{1}{3}F^\omega(1/2) + \frac{1}{3}F^\omega\left(\frac{1}{\lambda_{2,t}+1}\right) + \frac{1}{3}F^\omega\left(\frac{1}{\lambda_{3,t}+1}\right).$$

This is the distribution that governs the transition of $\langle \lambda_{2,t}, \lambda_{3,t} \rangle$. Note that neither level-2 nor level-3 agents have a correctly specified model of inference for any value of p, as neither are aware of level-3 types. Thus, the correctly specified model is not a special case of this level-k model.

Asymptotic Learning. We use Theorem 1 to characterize asymptotic learning outcomes in the level-k model. There are four candidate outcomes: $(0,0)$ corresponds to correct learning for level-2 and level-3 types, (∞,∞) corresponds to incorrect learning for both types, and $(0,\infty)$ and $(\infty,0)$ are the disagreement outcomes in which one type learns the correct state and the other learns the incorrect state. Recall that whether an asymptotic learning outcome $(\lambda_2^*,\lambda_3^*)$ arises depends on the signs of the expected change in the log likelihood ratio for each type, $\gamma_2(\lambda_2^*,\lambda_3^*)$ and $\gamma_3(\lambda_2^*,\lambda_3^*)$. By determining how the sign of (γ_2,γ_3) varies with level-3's belief p about the share of level-2 types, we can characterize the set of candidate learning outcomes Λ (as defined in (4)) for any p.

Suppose the true state is $\omega = R$ and consider the correct learning outcome $(0,0)$. At $(0,0)$, level-2's perceived probability of an R action in state ω is $F^\omega(1/2)$, and its perceived probability of an L action is $1-F^\omega(1/2)$. All level-2 and level-3 types are choosing R actions at these beliefs, so the true probability of an R action is $2/3 + F^R(1/2)/3$. The true probability of an L action is the probability of type θ_1 times the probability this type chooses L, $(1-F^R(1/2))/3$. From (3),

$$\gamma_2(0,0) = \left(\frac{2+F^R(1/2)}{3}\right)\log\left(\frac{F^L(1/2)}{F^R(1/2)}\right) + \left(\frac{1-F^R(1/2)}{3}\right)\log\left(\frac{1-F^L(1/2)}{1-F^R(1/2)}\right),$$

which is negative, since $F^R(1/2) > F^L(1/2)$. The expression $\gamma_3(0,0)$ can be constructed in a similar manner. Whenever $\gamma_3(0,0) < 0$, correct learning is a candidate learning outcome i.e. $(0,0) \in \Lambda$. A similar characterization determines whether other learning outcomes are in Λ.

From Theorem 1, whenever an agreement outcome $(0,0) \in \Lambda$ or $(\infty,\infty) \in \Lambda$, these learning outcomes occur with positive probability. When Λ contains a disagreement outcome, we also need to check whether the disagreement outcome satisfies the total informativeness rank to determine whether it occurs with positive probability. In this example, it turns out that when disagreement arises, learning is correct for level-2 types and incorrect for level-3 types i.e. $(0,\infty)$. At any p such that $(0,\infty) \in \Lambda$, $(0,\infty)$ is also total informativeness ranked, and therefore, occurs with positive probability. The other disagreement outcome, $(\infty,0)$, is never in Λ and almost surely does not arise. Therefore, characterizing Λ fully determines the set of asymptotic learning outcomes. Theorem 4 characterizes how asymptotic learning outcomes depend on p.

Theorem 4. *Suppose $\omega = R$. Then $\boldsymbol{\lambda}_t$ converges almost surely to a limit random*

variable $\boldsymbol{\lambda}_\infty$ *with support* Λ*. There exist unique cutoffs* $0 < p_1 < p_2 < p_3 < 1$ *such that:*

1. *If* $p < p_1$*, then incorrect and correct learning occur with positive probability,* $\Lambda = \{(0,0),(\infty,\infty)\}$*.*

2. *If* $p \in (p_1, p_2)$*, then incorrect learning, correct learning and disagreement occur with positive probability,* $\Lambda = \{(0,0),(\infty,\infty),(0,\infty)\}$*.*

3. *If* $p \in (p_2, p_3)$*, then correct learning and disagreement occur with positive probability,* $\Lambda = \{(0,0),(0,\infty)\}$*.*

4. *If* $p > p_3$*, then disagreement occurs almost surely,* $\Lambda = \{(0,\infty)\}$*.*

Intuition and Discussion. When p is low, level-3 types believe most agents are level-1 and they behave similarly to level-2 types. The models of level-2 and level-3 types are too similar for asymptotic disagreement to occur. Learning is either complete or incorrect. Both types overweight the informativeness of actions, and initial actions have an outsize effect on asymptotic beliefs, as the information from these actions is amplified in every subsequent action. Therefore, whether initial actions are correct or incorrect will influence whether beliefs build momentum on the correct or incorrect state.

As p increases, level-2 and level-3 types interpret the action history in an increasingly different way, and disagreement becomes possible. This disagreement takes a particular form: level-2 types learn the correct state, while the higher order of reasoning level-3 types do not. Although level-2 types have a lower order of reasoning, the impact of their misspecification is mitigated by the behavior of level-3 types. As p increases, level-3 types switch to underweighting the informativeness of actions – they believe a large share of actions are from level-2 types, and therefore, overweighted, so the level-3 types compensate by underweighting these actions. Therefore, the actions of level-3 types do indeed reflect more of their private information, mitigating level-2's bias. Consider $(\infty, 0)$. When the beliefs of level-2 types are near ∞ but the beliefs of level-3 types are near 0, R actions occur frequently enough to pull level-2's beliefs away from incorrect learning, despite level-2's overweighting, and this disagreement outcome does not arise. However, when the beliefs of level-3 types are near ∞ but the beliefs of level-2 types are near 0, even though R actions occur at the same frequency

as near $(\infty, 0)$, level-3 underweights these R actions, and therefore, level-3 beliefs continue moving towards the incorrect state. Thus, beliefs can converge to $(0, \infty)$.

As p increases above p_2, level-3 underweights the action history enough that it completely cancels the overweighting of level-2 types, and level-2 can no longer have incorrect learning. Either both types learn the correct state or they disagree and level-3 learns the incorrect state. Finally, for $p > p_3$, the level-3 types anti-imitate the level-2 types so severely that they almost surely converge to believing the incorrect state.

This characterization yields an interesting take-away on the incentives of an agent to acquire a higher level of reasoning. Suppose that an agent of type k can engage in costly introspection in order to increase his level of reasoning to $k+1$. If a higher level type performs strictly *worse* than a lower level type, then such an agent will not seek to increase his level of reasoning, even when the cost of doing so is arbitrarily small. Further, even if an agent already understands how to reason at a higher level, there is still a higher cognitive cost associated with utilizing this higher level of reasoning, as it involves more complex computations. A level-2 type simply needs to count the number of each action to make a decision, while in addition to this computation, a level-3 type needs to back-out the beliefs of a level-2 type at each previous period to accurately extract level-2's private information. Therefore, it may be optimal for the agent to reason at a lower level even if the cost of switching to higher-level reasoning is arbitrarily small.

Figure 1 plots the probability of each learning outcome, as a function of p. Increasing p monotonically increases the probability that level-2 learns the correct state, as level-3's model mitigates level-2's bias. However, increasing p has a non-monotonic effect on the probability that level-3 learns the correct state. At first, raising p moves level-3's model closer to the true model, which increases the probability of complete learning, but above $p = .55$, increasing p moves level-3's model further from the true model. In this specification, $p_1 = .01$, $p_2 = .55$ and $p_3 = .76$.

While this example focuses on a particular distribution, $\pi = (0, 1/3, 1/3, 1/3)$, a robustness result that is similar in spirit to Theorem 3 establishes that the learning outcomes characterized in Theorem 4 also obtain for nearby distributions π'.

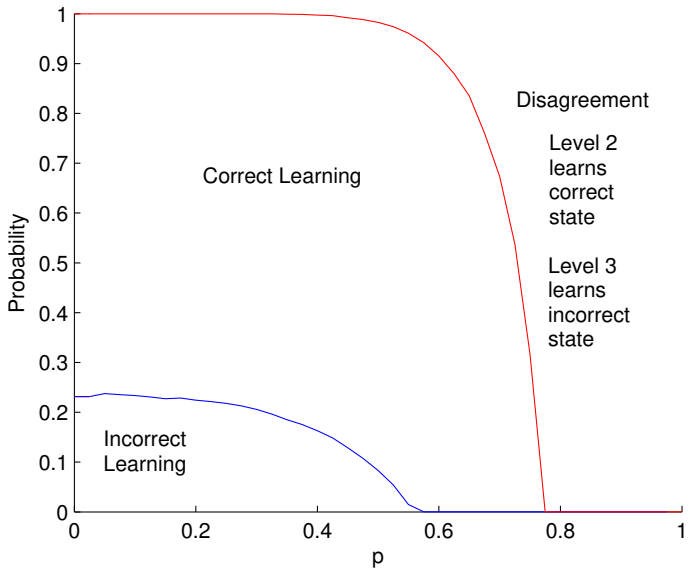

FIGURE 1. Probability of Learning Outcome. ($F_L(s) = \frac{5}{3}(s^2 - .04)$, $F_R = \frac{10}{3}(s - \frac{1}{2}s^2 - 3/5)$, $\operatorname{supp} F = [.2, .8]$)

4.2 Partisan Bias

A large literature in both psychology and economics has provided evidence for biased information processing that systematically slants information towards a particular state. One strand of literature posits that motivated reasoning (Kunda 1990) leads individuals to slant beliefs towards a preferred state due to self-image concerns (Bénabou and Tirole 2011), ego utility (Koszegi and Rabin 2006) or optimism (Brunnermeier and Parker 2005). A related literature in political science explores the impact of party affiliation on information processing. Jerit and Barabas (2012) find that subjects are better at recalling facts that support their political position. Bartels (2002) show that how individuals' evaluations of candidates update in response to new information is consistent with partisan bias. This bias can impede the convergence of beliefs and can even lead to polarization – beliefs moving in opposite directions – after observing the same event. In this application, we seek to model how such a bias affects social learning, but are agnostic as to its source.

Set-up. Suppose that there are two ways in which agents process private information. Some individuals – who we refer to as partisan types – systematically slant

private information in favor of state L. Following any private signal, these partisan types will believe that state L is more likely than it actually is, given the true measure over signals. We model this as a misspecified private signal distribution that slants information in favor of state L, $r^P(s) = s^\nu$ for some $\nu \in (0,1)$. Other individuals are unbiased in that they correctly interpret private information, $r^U(s) = s$. Although partisan and unbiased agents agree on the optimal action choice when the state is known, they will potentially disagree on the optimal action choice following imperfect signals, as the partisan types will believe that signals are more favorable towards state L than unbiased types.

To complete the signal misspecification, we must also specify $\hat{F}^{L,P}$, the perceived distribution of signal s in state L. We assume that the true distribution of signals is unbounded, $\operatorname{supp} F = [0,1]$, and that the perceived distribution satisfies $\hat{F}^{L,P}(s) = F^L(s^\nu)$. This implies that $\hat{F}^{R,P}(s) = F^R(s^\nu)$. Under this specification, whenever a partisan type sees a signal s, they interpret it like an unbiased type would interpret a signal $s' = s^\nu$, which corresponds to stronger evidence for state L. This captures a type who believes signals are manipulated towards state R. For instance, suppose vaccines are dangerous in state L and safe in state R. Then someone who is primed to believe that vaccines are dangerous may look at a study providing evidence that vaccines are safe and believe that the results were falsified to some degree, so results providing a signal of strength s towards state L were actually providing a signal of strength s^ν, more favorable to state L, before they were manipulated.

Suppose that some partisan and unbiased agents observe the history and others do not, so there are four types, $\Theta = \{\theta_P, \theta_U, \theta_{AP}, \theta_{AU}\}$. Types θ_P and θ_{AP} are partisan, with the former a sociable type who learns from the action history and the latter an autarkic type. Types θ_U and θ_{AU} are unbiased sociable and autarkic types, respectively. Let $q = \pi(\theta_{AP}) + \pi(\theta_P)$ denote the share of partisan types. Suppose share $\alpha \in (0,1)$ of both partisan and unbiased types are autarkic, so $\pi(\theta_{AP}) = \alpha q$ and $\pi(\theta_{AU}) = \alpha(1-q)$.

In the presence of partisan types, there is an additional challenge to learn from the actions of others, relative to a model in which all agents correctly interpret the state signal distribution. To accurately interpret actions, an unbiased agent must be aware of the partisan types, and know both the form of their bias (i.e. ν) and their frequency in the population. We assume that agents are not this sophisticated. In particular, unbiased types believe that all agents interpret private information in the

same manner as themselves. Although they have a correct model of the state signal distribution, they incorrectly assume that all other agents do as well.[22] Therefore, they do not invert the bias of the partisan types when learning from actions. This corresponds to believing that no types have partisan bias, $\hat{\pi}^U(\theta_{AP}) = \hat{\pi}^U(\theta_P) = 0$.

Similarly, partisan types believe that all other agents interpret information in the same manner as themselves. In the context of the vaccine example, this means that the partisan types believe that all other types are adjusting for the possibility that information has been manipulated. Although these types have a correct model of how other partisan types interpret information, they have an incorrect model of the state signal distribution driving this process and an incorrect model of how unbiased types interpret information. This corresponds to believing that all types have partisan bias, $\hat{\pi}^P(\theta_{AU}) = \hat{\pi}^P(\theta_U) = 0$, along with perceived state-signal distributions $(\hat{F}^{L,P}, \hat{F}^{R,P})$ that can be represented by $r^P(s) = s^\nu$.

To close the model, assume that both partisan and unbiased types correctly understand how to separate private information from redundant information in actions – that is, they have correct beliefs about the share of autarkic types in the population, $\hat{\pi}^P(\theta_{AP}) = \pi(\theta_{AU}) + \pi(\theta_{AP})$ and $\hat{\pi}^U(\theta_{AU}) = \pi(\theta_{AU}) + \pi(\theta_{AP})$. Assume that there are no public signals, and all agents have common error penalty $u = 1/2$. Autarkic types occur with positive probability, and both sociable types believe autarkic types occur with positive probability, so Assumptions 1-3 are satisfied.

Action Choices and Beliefs. Let $\boldsymbol{\lambda} = (\lambda_P, \lambda_U)$ denote the likelihood ratio vector. At $\lambda \in (0, 1)$, a sociable partisan type plays action R following signals $s \leq \overline{s}^P(\lambda) = 1/(1+\lambda)^{1/\nu}$, while a sociable unbiased type plays action R following signals $s \leq \overline{s}^U(\lambda) = 1/(1+\lambda)$. Similarly, autarkic partisan types play action R following signals $s \leq \overline{s}^P(1) = 0.5^{1/\nu}$, while autarkic unbiased types play action R following signals $s \leq \overline{s}^U(1) = 0.5$. Note that $\overline{s}^P(\lambda) < \overline{s}^U(\lambda)$ – partisan types choose action L for a larger interval of signals, and therefore, with higher frequency.

A partisan type believes that other agents are also partisan. Therefore, she believes that all other agents also use cut-off $\overline{s}^P$, which is lower than the threshold used by unbiased types. The partisan type also has an incorrect belief about the signal distribution – it believes signals are below $\overline{s}^P$ in state ω with probability $\hat{F}^{\omega,P}(\overline{s}^P(\lambda))$,

[22] We relax this assumption and consider the case where unbiased types have correct beliefs about the share of partisan types later in this section.

which is greater than the true probability $F^\omega(\overline{s}^P(\lambda))$. Therefore, she both underestimates the range of signals for which other agents choose action R and overestimates the probability of these signals. The partisan type's perceived probability of an R action is

$$\begin{aligned}\hat{\psi}^P(R|\omega,\lambda_P,\lambda_U) &= (1-\alpha)\hat{F}^{\omega,P}(\overline{s}^P(\lambda_P)) + \alpha\hat{F}^{\omega,P}(\overline{s}^P(1)) \\ &= (1-\alpha)F^\omega(\overline{s}^U(\lambda_P)) + \alpha F^\omega(\overline{s}^U(1)),\end{aligned}$$

where the second equality follows from $\overline{s}^P(\lambda) = \overline{s}^U(\lambda)^{1/\nu}$ and $\hat{F}^{\omega,P}(s) = F^\omega(s^\nu)$.

An unbiased type believes that other agents are also unbiased and use cut-off $\overline{s}^U$, and has a correct belief about the signal distribution. Therefore, she overestimates the range of signals for which other agents choose action R, since some agents are using cut-off $\overline{s}^P < \overline{s}^U$, but correctly estimates the probability of these signals. The unbiased type's perceived probability of an R action is

$$\hat{\psi}^U(R|\omega,\lambda_P,\lambda_U) = (1-\alpha)F^\omega(\overline{s}^U(\lambda_U)) + \alpha F^\omega(\overline{s}^U(1)).$$

This is equal to the true probability that an *unbiased* type plays an R action, and is strictly greater than the true probability of an R action.

Note that if $\lambda_P = \lambda_U$, then $\hat{\psi}^P(R|\omega,\lambda_P,\lambda_U) = \hat{\psi}^U(R|\omega,\lambda_P,\lambda_U)$. Therefore, if the partisan and unbiased type start with the same prior belief, both types update their public likelihood ratio in the same way following an action, and after any history h_t, $\lambda_{P,t} = \lambda_{U,t}$. Although they have different models of the world, their misspecifications collapse to the same misperceived probability of each action in each state. They both overestimate the informativeness of L actions and underestimate the informativeness of R actions. This means that we can consider partisan and unbiased types as a single type to characterize asymptotic learning.[23] It also rules out the possibility of asymptotic disagreement.

Incorrect Learning. When partisan bias is in favor of the incorrect state, then the learning outcome depends on the severity of the partisan bias. If partisan bias is severe, then partisan types choose L for a large range of signals. They believe these signals are less likely than is actually the case, and therefore, they overestimate the

[23]Note that this does not imply that a partisan and unbiased type with public belief λ and private signal s will choose the same action, as they have different cut-offs.

informativeness of L actions. Unbiased types believe that other agents are choosing L for a smaller range of signals than is actually the case, and therefore, they also overestimate the informativeness of these L actions. This leads both partisan and unbiased types to almost surely learn the incorrect state. If partisan bias is not severe (i.e. ν close to one), overweighting the informativeness of L actions is not severe enough to interfere with learning and both types learn the correct state. For intermediate levels, beliefs do not converge. Agents believe L actions are not very informative when beliefs are close to $\lambda = \infty$, as most agents are following the herd and reveal little private information, so beliefs do not converge to ∞. But these agents also underestimate the informativeness of R actions, and therefore, when beliefs are close to 0, L actions pull beliefs away from 0 and prevent correct learning. As discussed above, when partisan bias favors the correct state, learning is complete regardless of the level of bias, as the bias simply speeds up the rate at which beliefs converge to state L. Theorem 5 formalizes these results (the proof is in Appendix A.4).

Theorem 5. *When $\omega = R$, there exists an $\overline{q} \in (0,1)$ such that for $q > \overline{q}$, there exist unique cutoffs $0 < \nu_1(q) < \nu_2(q) < 1$ such that:*

1. *If $\nu > \nu_2(q)$, then learning is correct almost surely, $\Lambda = \{(0,0)\}$.*
2. *If $\nu \in (\nu_1(q), \nu_2(q))$, then learning is incomplete and beliefs do not converge almost surely, $\Lambda = \emptyset$.*
3. *If $\nu < \nu_1(q)$, then learning is incorrect almost surely, $\Lambda = \{(\infty, \infty)\}$.*

and there exists a $\underline{q} < \overline{q}$ such that for $q < \underline{q}$, learning is correct almost surely. When $\omega = L$, learning is correct almost surely, $\Lambda = \{(0,0)\}$.

Figure 2 illustrates the asymptotic learning outcomes as a function of (q, ν) when $\omega = R$. Theorem 5 and Figure 2 also illustrate the robustness of the correctly specified model, in which $q = 0$ and $\nu = 1$. Notice that for (q, ν) close enough to $(0, 1)$, learning is correct almost surely (Theorem 3.1). When ν is close to 1, then correct learning obtains even if all agents have partisan bias ($q = 1$), since the bias is not severe (Theorem 3.2). Similarly, when the share of partisan types is small, q close to 0, then correct learning obtains even if these partisan types have a very severe bias, ν close to 0.

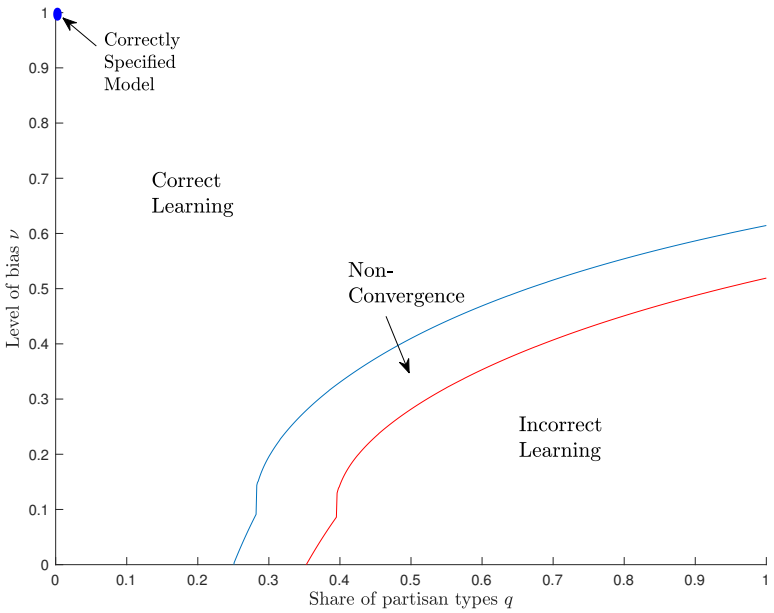

FIGURE 2. Learning Outcomes for (q, ν). ($\omega = R$, $\alpha = .1$, $F_L(s) = s^2$, $F_R(s) = 2s - s^2$, $\operatorname{supp} F = [0, 1]$)

Disagreement. Now suppose there the unbiased type not only has a correct belief about the signal distribution, but also correctly accounts for and parses out the overweighted information in favor of state L from the partisan types. In other words, the unbiased type has a correct belief about the level of partisan bias ν and the share of partisan types q. The next result establishes that disagreement can arise with probability one when the partisan bias favors the incorrect state.

Theorem 6. *Suppose $\omega = R$. There exists a $\overline{q} \in (0, 1)$ such that for $q > \overline{q}$, there exist unique cutoffs $0 < \nu_1(q) < \nu_2(q) < 1$ such that:*

1. *If $\nu > \nu_2(q)$, then learning is correct almost surely, $\Lambda = \{(0, 0)\}$.*
2. *If $\nu \in (\nu_1(q), \nu_2(q))$, then learning is incomplete for the partisan type, $\Lambda = \emptyset$, but the unbiased type still learns the correct state.*
3. *If $\nu < \nu_1(q)$, then disagreement occurs almost surely, $\Lambda = \{(\infty, 0)\}$.*

Therefore, despite observing the same sequence of information, partisan and unbiased types almost surely disagree.

4.3 Confirmation Bias

Confirmation bias is the tendency to interpret information in a way that confirms one's existing beliefs or hypotheses about the world. This bias is well documented in the literature. It has been highlighted as a significant factor in the over-justification of adopted policies in government (Tuchman 1984), continued use of ineffective procedures in medicine (Thomas 1979) and primacy effects in judicial reasoning (Devine and Ostrom 1985). It has also been linked as a cause of overconfidence (Nickerson 1998), which has been implicated as a major reason for the underperformance of individual traders on financial markets (Barber and Odean 2001; Odean 1999).[24]

Set-up. In this application, we show how confirmation bias impacts asymptotic learning. Suppose agents act sequentially and observe a sequence of informative public signals. There is as single type who underweights information that contradicts her prior beliefs (i.e. information that favors the state that she believes to be less likely). The agent correctly interprets information that confirms her prior beliefs in that her perceived posterior is equal to the true posterior following a signal in favor of the more likely state. Given prior belief p that the state is L, an agent interprets public signal σ according to

$$\rho(\sigma, p) = \begin{cases} \sigma & \text{if } \sigma = \sigma_L \text{ and } p \geq \frac{1}{2} \\ \sigma & \text{if } \sigma = \sigma_R \text{ and } p \leq \frac{1}{2} \\ \sigma + \frac{\varrho(p)}{k}(0.5 - \sigma) & \text{otherwise.} \end{cases}$$

where $k > 1$ and $\varrho : [0,1] \to [0,1]$ is a continuous function with $\varrho(0) = \varrho(1) = 1$. Confirmation bias is less severe for higher k, and the correctly specified model corresponds to the limit as $k \to \infty$. If ϱ is strictly decreasing on $[0, 1/2]$ and strictly increasing on $[1/2, 1]$, then the bias becomes more severe as the agent's prior becomes more extreme.

[24] Additionally, Lord et al. (1979) show that when asked to read two studies, one which supports capital punishment and one that does not, proponents of capital punishment place more weight on the former study, while opponents place more weight on the latter; Darley and Gross (1983) found that after being told a child's socioeconomic background, subjects were more likely to rate her performance on a reading test lower when she came from a low socioeconomic background; Plous (1991) documents that, when faced with a non-catastrophic breakdown of a given technology, supporters of the technology become more confident that the safeguard in place will prevent a catastrophic breakdown, while opponents will believe that a catastrophic breakdown is more likely.

To complete the model, assume that public signals are the only source of information and they are informative (i.e. private signals are uninformative, are believed to be uninformative, and $\sigma_L > 1/2$). An agent's beliefs about how other agents interpret public signals is irrelevant, as there is no additional information contained in actions. All agents have common error penalty $u = 1/2$.

Action Choices and Beliefs. There is one type, so we need to keep track of a single likelihood ratio. In period t, given λ_t, agent t chooses action L if $\lambda_t > 1$, and otherwise chooses action R. Actions are uninformative, so following public signal $\sigma_t = \sigma_L$, the likelihood ratio updates to

$$\lambda_{t+1} = \lambda_t \left(\frac{\sigma_L}{1 - \sigma_L} \right)$$

if $\lambda_t \geq 1$ and

$$\lambda_{t+1} = \lambda_t \left(\frac{\sigma_L + \frac{\varrho(\lambda_t/(\lambda_t+1))}{k}(0.5 - \sigma_L)}{1 - \sigma_L - \frac{\varrho(\lambda_t/(\lambda_t+1))}{k}(0.5 - \sigma_L)} \right)$$

if $\lambda_t < 1$. The expressions following $\sigma_t = \sigma_R$ are analogous. The actual probability of signal σ^L is $dF^R(\sigma_L) = (1 - 2\sigma_R)(1 - \sigma_L)/(\sigma_L - \sigma_R)$ in state R and $dF^L(\sigma_L) = (1 - 2\sigma_R)\sigma_L/(\sigma_L - \sigma_R)$ in state L (Lemma 9 in Appendix A.1).

Asymptotic Learning. When agents have confirmation bias, they underweight signals that do not confirm their current belief. Following a *contrary* signal – a signal in favor of the less likely state – beliefs move more slowly away from the favored state, relative to the correctly specified model. When the confirmation bias is severe enough, it is very unlikely that agents will see enough information to overturn their prior misconceptions. Suppose the true state is R but an initial set of signals favor state L. If confirmation bias is severe, it is difficult to recover from this trap, as agents place less and less weight on contrary R signals as their beliefs move towards state L. Therefore, incorrect learning arises with positive probability. But if confirmation bias is relatively low, then agents will eventually see enough signals in favor of state R to overcome their preconceptions and incorrect learning almost surely does not occur. Correct learning always occurs with positive probability, since with positive probability, agents come to believe the correct state is more likely, and once this occurs, the bias only increases the rate at which their beliefs move towards the correct

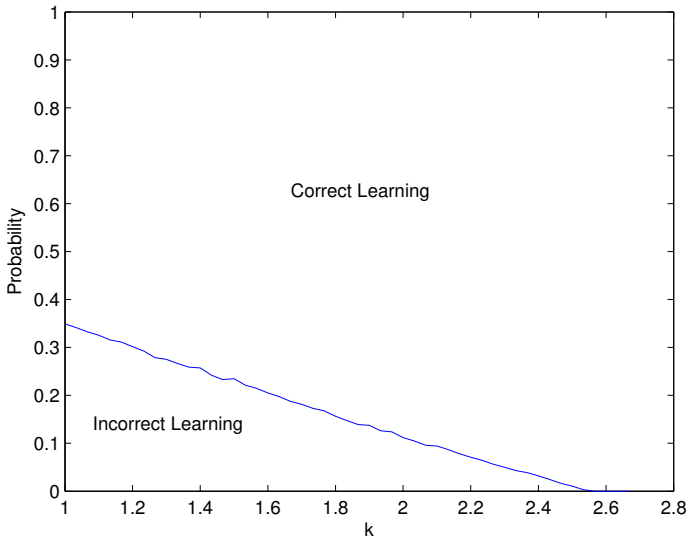

FIGURE 3. Probability of Learning Outcome. $(\varrho(p) = \left(\frac{1}{2} \cdot |p - 1/2|\right)^{\frac{1}{8}}, \sigma_L = 5/8, \sigma_R = 3/8)$

state. This is similar to the results from Rabin and Schrag (1999).[25]

Theorem 7. *Suppose $\omega = R$. There exists a unique cutoff $\bar{k} > 1$ such that*

1. *If $k > \bar{k}$ then learning is correct almost surely.*
2. *If $k < \bar{k}$ then both correct and incorrect learning occur with positive probability, and beliefs converge almost surely.*

Figure 3 plots the probability of correct and incorrect learning as a function of k. Increasing k monotonically decreases the probability of incorrect learning. Correct learning almost surely occurs for $k > 2.55$.

5 Conclusion

Our paper develops a general framework for learning with model misspecification. We characterize how the asymptotic learning outcomes depend on the primitives of the model, including the ways in which agents misinterpret private and public

[25] We can nest the model in Rabin and Schrag (1999) with a minor extension to our framework – allowing for four public signals. All results in this paper easily extend to any finite number of public signals, so this is a straightforward extension. Appendix B outlines the mapping between this paper and Rabin and Schrag (1999).

information, and draw inference from the actions of others. When agents' models of the world are misspecified, correct learning – individuals eventually place probability one on the true state – is no longer guaranteed. Asymptotic learning may be incorrect, individuals may perpetually disagree, or beliefs may not converge at all. We establish that the correctly specified model is robust, in that correct learning is guaranteed for approximately correctly specified models, regardless of the form of misspecification. These results yield insights about new forms of misspecification, as well as unify particular types of misspecification that have already been studied.

A Appendix

A.1 Posterior Representation.

Let $\mathcal{Z}$ be a signal space. Suppose signals $\{z_n\}$ are i.i.d., conditional on the state, and drawn according to probability measure $\mu^\omega \in \Delta(\mathcal{Z})$ in state $\omega \in \{L, R\}$. Assume μ^L, μ^R are mutually absolutely continuous, and therefore have common support, which we assume to be full. Define the posterior belief $s(z) \equiv 1/(1 + \frac{d\mu^R}{d\mu^L}(z))$ that the state is L. The c.d.f. $F_s^\omega(x) \equiv \mu^\omega(z|s(z) \leq x)$ is the distribution of the posterior belief s, with common support $\operatorname{supp} F_s$. Let $[\underline{b}, \overline{b}] \subseteq [0, 1]$ denote the convex hull of $\operatorname{supp} F_s$. Assume signals are informative, which rules out $d\mu^L/d\mu^R = 1$ almost surely.

Suppose an agent has a misspecified probability measure $\hat{\mu}^\omega \in \Delta(\mathcal{Z})$ about the distribution of signals, where $\hat{\mu}^L, \hat{\mu}^R$ are mutually absolutely continuous with common support. Assume the misspecified measures also have full support, so that agents do not observe signals that are inconsistent with their model of the world. When an agent observes signal z, she has misspecified posterior belief $\hat{s}(z) \equiv 1/(1 + \frac{d\hat{\mu}^R}{d\hat{\mu}^L}(z))$. The c.d.f. $\hat{F}_{\hat{s}}^\omega(x) \equiv \hat{\mu}^\omega(z|\hat{s}(z) \leq x)$ is the perceived distribution of $\hat{s}$. We also define the c.d.f. $F_{\hat{s}}^\omega(x) \equiv \mu^\omega(z|\hat{s}(z) \leq x)$ as the true distribution of $\hat{s}$ and the c.d.f. $\hat{F}_s^\omega(x) \equiv \hat{\mu}^\omega(z|s(z) \leq x)$ as the perceived distribution of s.

We define two properties of probability measures. The first describes a property of the relationship between two pairs of measures, which lead to the same ordinal mapping between sets of signals and posterior beliefs.

Definition 7 (Equivalent Ordinal Ranking of Signals)**.** *Given mutually absolutely continuous probability measures $\mu^L, \mu^R \in \Delta(\mathcal{Z})$ with Radon-Nikodym derivative $f(z) = \frac{d\mu^R}{d\mu^L}(z)$, mutually absolutely continuous probability measures $\nu^L, \nu^R \in \Delta(\mathcal{Z})$ with*

$\operatorname{supp}\nu = \operatorname{supp}\mu$ *and Radon-Nikodym derivative* $g(z) = \frac{d\nu^R}{d\nu^L}(z)$ *have an* equivalent ordinal ranking of signals *if for any* $z, z' \in \mathcal{Z}$ *such that* $f(z) \geq f(z')$*, then* $g(z) \geq g(z')$*, with equality iff* $f(z) = f(z')$*.*

The second describes an equivalence class of probability measures, which have the same support of posterior beliefs, distributions over posterior beliefs and ordinal ranking of signals.

Definition 8 (Equivalent Measures)**.** *Mutually absolutely continuous probability measures* $\mu^L, \mu^R \in \Delta(\mathcal{Z})$ *and* $\nu^L, \nu^R \in \Delta(\mathcal{Z})$ *are* equivalent *iff* $\operatorname{supp}\mu = \operatorname{supp}\nu$*,* $\mu^\omega(z|1/(1+\frac{d\mu^R}{d\mu^L}(z)) \leq x) = \nu^\omega(z|1/(1+\frac{d\nu^R}{d\nu^L}(z)) \leq x)$ *for all* $x \in [0,1]$ *and they have an equivalent ordinal ranking of signals.*

Lemma 8 establishes that when a pair of misspecified probability measures has an equivalent ordinal ranking of signals as the true measures, there is a unique mapping between a set of misspecified measures $(\hat{\mu}^L, \hat{\mu}^R) \in \Delta\mathcal{Z}$ and a representation $(r, \hat{F}_s^L)$, where $r : \operatorname{supp} F_s \to [0,1]$ is a strictly increasing function mapping the true posterior s to the misspecified posterior $\hat{s}$ and $\hat{F}_s^L$ is the c.d.f. of the perceived distribution of s in state L.

Lemma 8. *Let* $\mu^L, \mu^R \in \Delta(\mathcal{Z})$ *be a set of mutually absolutely continuous probability measures with full support. Assume signals are informative.*

1. *For any mutually absolutely continuous misspecified probability measures* $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})$ *that have full support and an equivalent ordinal ranking of signals, there exists a unique* $(r, \hat{F}_s^L)$*, where* $r : \operatorname{supp} F_s \to [0,1]$ *is a strictly increasing function with* $r(\bar{b}) > 1/2$ *and* $r(\underline{b}) < 1/2$*, such that* $\hat{s}(z) = r(s(z))$ *for all* $z \in \mathcal{Z}$ *and* $\hat{F}_s^L$ *is the c.d.f. of the perceived distribution of* s *in state* L*.*

2. *For any strictly increasing function* $r : \operatorname{supp} F_s \to (0,1)$ *and any c.d.f.* $\hat{F}_s^L$ *with* $\operatorname{supp}\hat{F}_s^L = \operatorname{supp} F_s$ *and* $\int_0^1 \left(\frac{1-r(s)}{r(s)}\right) d\hat{F}_s^L = 1$*, there exist unique (up to an equivalent measure) mutually absolutely continuous probability measures* $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})$ *that have full support and satisfy* $r(s(z)) = 1/(1+\frac{d\hat{\mu}^R}{d\hat{\mu}^L}(z))$ *for all* $z \in \mathcal{Z}$*. The measures* $\hat{\mu}^L, \hat{\mu}^R$ *have an equivalent ordinal ranking of signals to* μ^L, μ^R*.*[26]

[26]Note that if $\hat{F}_s^L$ is a c.d.f. and $\int_0^1 \left(\frac{1-r(s)}{r(s)}\right) d\hat{F}_s^L = 1$, then it must be that $r(\bar{b}) > 1/2$ and $r(\underline{b}) < 1/2$.

3. *For any strictly increasing function* $r : \operatorname{supp} F_s \to (0,1)$*, if* $r(\underline{b}) < 1/2$ *and* $r(\overline{b}) > 1/2$*, then there exist mutually absolutely continuous probability measures* $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})$ *that have full support and satisfy* $r(s(z)) = 1/(1 + \frac{d\hat{\mu}^R}{d\hat{\mu}^L}(z))$ *for all* $z \in \mathcal{Z}$*.*

Proof. First establish part (i). Let $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})$ be probability measures that are mutually absolutely continuous with full support and strictly preserve the ordinal ranking of signals. Define the mapping $r : \operatorname{supp}(F_s) \to [0,1]$ as $r(s(z)) = \hat{s}(z)$. This is a function since if $s(z) = s(z')$, then $\hat{s}(z) = \hat{s}(z')$, which establishes existence. For any z such that $s(z) > s(z')$, $\hat{s}(z) = r(s(z)) > \hat{s}(z') = r(s(z'))$ since $\hat{\mu}^L, \hat{\mu}^R$ strictly preserve the ordinal ranking of signals. Therefore, r is strictly increasing on $\operatorname{supp} F_s$.

By the Bayesian constraint, it must be that $\hat{E}[\hat{s}(z)] = 1/2$, where the expectation is taken with respect to the misspecified measures. Given that the true measures are informative and the misspecified measures strictly preserve the ordinal ranking of signals, it cannot be that $\hat{s}(z) = 1/2$ for all $z \in \mathcal{Z}$. Therefore, there exist $z, z' \in \mathcal{Z}$ such that $\hat{s}(z) > 1/2$ and $\hat{s}(z') < 1/2$, which implies that there exist $s, s' \in \operatorname{supp} F_s$ such that $r(s) > 1/2$ and $r(s') < 1/2$. Given that r is strictly increasing in s, it immediately follows that $r(\overline{b}) > 1/2$ and $r(\underline{b}) < 1/2$. Define $\hat{F}_s^L(x) \equiv \hat{\mu}^L(z|s(z) \leq x)$. Then $\hat{F}_s^L$ is the perceived c.d.f. of s under measure $\hat{\mu}^L$. Given $\{r, \hat{F}_s^L\}$, $\hat{F}_s^R$ is uniquely pinned down by

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

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

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

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

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

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

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

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

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

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 $\hat{\mu}^R$ and $\hat{\mu}^L$ are mutually absolutely continuous with common support $\operatorname{supp}\mu$. Also, $\operatorname{supp}\hat{\mu} = \operatorname{supp}\mu$ by construction, so the measures have full support on $\mathcal{Z}$. Moveover, since F_s^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 $A \subseteq \mathcal{Z}$,

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

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

Finally, show part (iii). Suppose $r : \operatorname{supp}F_s \to [0,1]$ is a strictly increasing function with $r(\underline{b}) < 1/2$ and $r(\bar{b}) > 1/2$. Fix any distribution $\hat{F}(\cdot)$ with support $\operatorname{supp}F_s \cap \{s|r(s) < 1/2\}$. Then $\int_0^1 \left(\frac{1-r(s)}{r(s)}\right)d\hat{F}_s(s) < 1$. Similarly, fix a distribution $\hat{G}(\cdot)$ with support $\operatorname{supp}F_s \cap \{s|r(s) \geq 1/2\}$. Then $\int_0^1 \left(\frac{1-r(s)}{r(s)}\right)d\hat{G}(s) > 1$. For any $\lambda \in [0,1]$, let $\hat{F}_\lambda$ be the distribution of the compound lottery $\hat{F}_\lambda = \lambda\hat{F} + (1-\lambda)\hat{G}$. This lottery draws signals from $\hat{F}$ with probability λ and $\hat{G}$ with probability $(1-\lambda)$. The function $H(\lambda) \equiv \int \left(\frac{1-r(s)}{r(s)}\right)d\hat{F}_\lambda$ 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 $\int \left(\frac{1-r(s)}{r(s)}\right)d\hat{F}_{\lambda^*} = 1$. Let $\hat{F}^L = \hat{F}_{\lambda^*}$. Then $\hat{F}^L$ is a probability distribution, since it is the convex combination of two distributions. By construction, $\operatorname{supp}\hat{F}_s^L = \operatorname{supp}F_s$

and $\int_0^1 \left(\frac{1-r(s)}{r(s)}\right) d\hat{F}_s^L = 1$. Therefore, from part(ii), it is possible to construct the desired probability measures $\hat{\mu}^L, \hat{\mu}^R$. □

The first part of Lemma 8 implies that $F_{\tilde{s}}^{\omega}(r(s)) = F_s^{\omega}(s)$ for all $s \in \text{supp}(F_s)$ and $\text{supp}(F_{\hat{s}}) = r(\text{supp}(F_s))$. Similarly, $\hat{F}_{\hat{s}}^{\omega}(r(s)) = \hat{F}_s^{\omega}(s)$ for all $s \in \text{supp}(\hat{F}_s)$ and $\text{supp}(\hat{F}_{\hat{s}}) = r(\text{supp}(\hat{F}_s))$.

Lemma 9. *Given mutually absolutely continuous probability measures $\mu^L, \mu^R \in \Delta(\mathcal{Z})$, $\text{supp}\, F_s$ and F_s^L are sufficient for the state-signal distribution. If signals are binary, $|\text{supp}\, F_s| = 2$, then $\text{supp}\, F_s$ is sufficient for the state-signal distribution.*

Proof. The first part follows immediately from Lemma A.1 in Smith and Sorensen (2000). Given $\text{supp}\, F_s$ and F_s^L, F_s^R is uniquely pinned down by

$$F_s^R(x) = \int_0^x \left(\frac{1-s}{s}\right) dF_s^L(s)$$

for any $x \in \text{supp}\, F_s$.

In the case of binary signals, there are two possible posterior beliefs. Without loss of generality, denote these beliefs s_R and s_L, with $s_R \leq s_L$. It must be that $s_R \leq 1/2 \leq s_L$, where the equality either binds for both or neither posteriors, in order to satisfy the Bayesian constraint $E[s] = 1/2$. Then F_s^L and F_s^R are uniquely pinned down by $\{s_R, s_L\}$. To see this, note that by definition, $dF_s^L(s_L)/dF_s^R(s_L) = s_L/(1-s_L)$ and $dF_s^L(s_R)/dF_s^R(s_R) = s_R/(1-s_R)$. Since F_s^L is a c.d.f., $dF_s^L(s_L) + dF_s^L(s_R) = 1$. Therefore,

$$\left(\frac{s_L}{1-s_L}\right) dF^R(s_L) + \left(\frac{s_R}{1-s_R}\right) dF^R(s_R) = 1. \tag{6}$$

Similarly, $dF_s^R(s_L) + dF_s^R(s_R) = 1$. Plugging in $dF_s^R(s_R) = 1 - dF_s^R(s_L)$ to (6) pins down the unique $dF_s^R(s_L) \in (0,1)$, and therefore, $dF_s^R(s_R)$. $F_s^L(s_R)$ is pinned down by $dF_s^L(s_R)/dF_s^R(s_R) = s_R/(1-s_R)$, and similarly for $dF_s^L(s_L)$. □

A.2 Proof of Theorem 1

Throughout this section, assume Assumptions 1, 2 and 3 hold and suppose $\omega = R$.

Proof of Lemma 1. At a stationary vector $\boldsymbol{\lambda}^*$, $\phi_i(a,\sigma,\boldsymbol{\lambda}^*) = \boldsymbol{\lambda}^*$ for all (a,σ) such that $\psi_i(a,\sigma|\omega,\boldsymbol{\lambda}^*) > 0$. When $\pi(\Theta_A) > 0$, both actions occur with positive probability at all $\boldsymbol{\lambda} \in [0,\infty]^k$, since autarkic types play both actions with positive probability independent of the history. Both public signals always occur with positive probability since the distribution is independent of $\boldsymbol{\lambda}$. Therefore, at all $\boldsymbol{\lambda} \in [0,\infty]^k$, either $\psi_i(L,\sigma_L|\omega,\boldsymbol{\lambda}) > 0$ or $\psi_i(R,\sigma_R|\omega,\boldsymbol{\lambda}) > 0$ (or both). By Assumption 2, actions and/or public signals are perceived to be informative by all sociable types θ_i, so at all $\boldsymbol{\lambda} \in [0,\infty]^k$,

$$\frac{\hat{\psi}_i(L,\sigma_L|L,\boldsymbol{\lambda})}{\hat{\psi}_i(L,\sigma_L|R,\boldsymbol{\lambda})} > 1 \qquad \text{and} \qquad \frac{\hat{\psi}_i(R,\sigma_R|L,\boldsymbol{\lambda})}{\hat{\psi}_i(R,\sigma_R|R,\boldsymbol{\lambda})} < 1.$$

Therefore, $\phi(a,\sigma,\boldsymbol{\lambda}) = \boldsymbol{\lambda}$ at all (a,σ) such that $\psi_i(a,\sigma|\omega,\boldsymbol{\lambda}) > 0$ if and only if $\boldsymbol{\lambda} \in \{0,\infty\}^k$. □

Proof of Lemma 2.

Part 1. Consider the stationary vector $\mathbf{0}$. Since $\gamma_i(\mathbf{0}) < 0$ for all i, there exists a neighborhood of $\mathbf{0}$, $[0,M]^k$, such that given any likelihood ratio vector $\boldsymbol{\lambda}_{a,\sigma} \in [0,M]^k$ for each a,σ pair

$$\sum_{a,\sigma} \psi_i(a,\sigma|R,\mathbf{0}) \log \frac{\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda}_{a,\sigma})}{\hat{\psi}_i(a,\sigma|R,\boldsymbol{\lambda}_{a,\sigma})} < 0.$$

Let

$$g_{i,a,\sigma} = \sup_{\boldsymbol{\lambda} \in [0,M]^k} \log \frac{\hat{\psi}_i(a,\sigma|L,\boldsymbol{\lambda}_{a,\sigma})}{\hat{\psi}_i(a,\sigma|R,\boldsymbol{\lambda}_{a,\sigma})}$$

and let

$$\bar{g}_i = \max_{a,\sigma} g_{i,a,\sigma}.$$

Fix an $\varepsilon > 0$ and define a neighborhood $[0,M_\varepsilon]^k \subseteq [0,M]^k$ such that

$$\inf_{\boldsymbol{\lambda} \in [0,M_\varepsilon]^k} |\psi(a,\sigma,\boldsymbol{\lambda}) - \psi(a,\sigma,0)| < \varepsilon/4.$$

Define the linear system $\langle \hat{\boldsymbol{\lambda}}_{\varepsilon,t} \rangle$ as follows.

$$\hat{\boldsymbol{\lambda}}_{\varepsilon,t} = \exp(g_{a,\sigma}) \hat{\boldsymbol{\lambda}}_{t-1},$$

when public signal σ is realized and the type drawn in period t would play a for all $\boldsymbol{\lambda} \in [0, M_\varepsilon]$, and

$$\hat{\boldsymbol{\lambda}}_{\varepsilon,t} = \exp(\bar{g}_i)\hat{\boldsymbol{\lambda}}_{t-1}$$

otherwise (let $\bar{\varepsilon}$ be the probability of this event). This is a linear system in each coordinate, so by lemma C.1 of Smith and Sorensen (2000) if

$$\exp(\bar{g}_i)^{\bar{\varepsilon}} \prod_{a,\sigma} \exp(g_{i,a,\sigma})^{\inf_{\boldsymbol{\lambda}\in[0,M_\varepsilon]^k} \psi_i(a,\sigma|R,\boldsymbol{\lambda})} < 1.$$

This holds for sufficiently small ε, ε_1, since this is strictly less than 1 at $\varepsilon = 0$.

So whenever a private signal is drawn such that a type would play a for any $\boldsymbol{\lambda} \in [0, M_\varepsilon]^k$, $\hat{\boldsymbol{\lambda}}_{\varepsilon_1,t}$ updates by $\exp(g_{a,\sigma})$ which is by construction larger than the actual update. Otherwise, $\boldsymbol{\lambda}_t$ updates by $\bar{g}$, which is larger than all possible updates. Therefore $\hat{\lambda}_{i,\varepsilon_1,t-1} = \lambda_{i,t-1}$ then $\hat{\lambda}_{i,\varepsilon_1,t} \geq \lambda_{i,t}$ for all i. So if $\boldsymbol{\lambda}_0 \in [0, M_{\varepsilon_1}]^k$ then it is bounded above by a RV that converges almost surely as long as it remains in $[0, M_{\varepsilon_1}]^k$.

Since $\hat{\boldsymbol{\lambda}}_{\varepsilon_1,t} \to \mathbf{0}$ almost surely

$$Pr(\cup_t \cap_{s\geq t} \{\hat{\boldsymbol{\lambda}}_s \in [0, M_{\varepsilon_1}]^k\}) = 1$$

So there must exist some $t \geq 0$ such that $Pr(\forall s \geq t, \hat{\boldsymbol{\lambda}}_{\varepsilon_1,s} \in [0, M_{\varepsilon_1}]^k$ and since the system is linear, if this holds at some $t > 0$, it must hold at $t = 0$. So, with positive probability, if $\hat{\boldsymbol{\lambda}}_{\varepsilon_1,0} \in [0, M_{\varepsilon_1}]^k$, it remains in $[0, M_{\varepsilon_1}]^k$ forever and is thus always larger than $\boldsymbol{\lambda}$. When this happens, since $\hat{\boldsymbol{\lambda}}_{\varepsilon_1}$ converges to $\mathbf{0}$, so does $\boldsymbol{\lambda}$.

The proof in the other cases is analogous. If $\lambda_i^* = \infty$, consider the λ_i^{-1} instead for that component and modify the transition rules accordingly. □

Part 2. Suppose $\boldsymbol{\lambda}^*$ is stationary and there exists a $\theta_i \in \Theta_S$ such that $\lambda_i = 0$ and $\gamma_i(\boldsymbol{\lambda}^*) > 0$ or $\lambda_i = \infty$ and $\gamma_i(\boldsymbol{\lambda}^*) < 0$. Suppose $P(\boldsymbol{\lambda}_t \to \boldsymbol{\lambda}^*) > 0$ so that the likelihood ratio converges to this vector with positive probability. Let

$$\Theta_R \equiv \{\theta_i \in \Theta_S | \lambda_i = 0\}$$

be the set of sociable types with limit belief 0 and $\Theta_L \equiv \Theta_S \setminus \Theta_R$ be the set of sociable types with limit belief ∞. Given $\theta_i \in \Theta_S$, define

$$g_i(a, \sigma, \boldsymbol{\lambda}) \equiv \log \frac{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda})}{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})}.$$

for all $i \in \Theta_R$ and

$$g_i(a, \sigma, \boldsymbol{\lambda}) \equiv \log \frac{\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})}{\hat{\psi}_i(a, \sigma | L, \boldsymbol{\lambda})}$$

for all other types. The log-likelihood ratio process $\langle \log \boldsymbol{\lambda}_t \rangle$ follows law of motion $\log \lambda_{i,t+1} = \log \lambda_{i,t} + g_i(a_t, \sigma_t, \boldsymbol{\lambda}_t)$ for each $\theta_i \in \Theta_S$. Fix a nbhd $[0, M]^k$ and define an i.i.d. sequence of random variables

$$\alpha_t^\theta = \begin{cases} L \text{ if } \theta_t \text{ plays } L \text{ at } (\boldsymbol{\lambda}_t, s_t) \text{ for any } \boldsymbol{\lambda} \in [0, M] \\ R \text{ if } \theta_t \text{ plays } R \text{ at } (\boldsymbol{\lambda}_t, s_t) \text{ for any } \boldsymbol{\lambda} \in [0, M] \\ R \text{ if the above doesn't hold and } \theta \in \Theta_R \\ L \text{ otherwise} \end{cases}$$

Fix $\varepsilon > 0$, and choose M such that the probability that either of the first two cases do not occur is at most ε.
By Lemma 2, for small $\varepsilon > 0$

$$\sum_{\alpha, \sigma} \psi_{i,\alpha}(L, \sigma, \boldsymbol{\lambda}^*) g_i(\alpha, \sigma, \boldsymbol{\lambda}^*) > (<) 0.$$

for all $\theta_i \in \Theta_R (\Theta_L)$, where $\psi_\alpha(\alpha, \sigma, \boldsymbol{\lambda}^*)$ is the probability of (α, σ) given the α^θ random variable.

By continuity, there exists an $\bar{M} > 0$ such that

$$\sum_{\alpha, \sigma} \psi_{i,\alpha}(L, \sigma, \boldsymbol{\lambda}^*) g_i(\alpha, \sigma, \boldsymbol{\lambda}_{\alpha,\sigma}) > (<) 0.$$

inequalities holds for any four $\boldsymbol{\lambda}_{\alpha,\sigma} \in [0, \bar{M}]^k$, for any $\theta_i \in \Theta_R (\Theta_L)$

Let

$$g_{i,a,\sigma} = \inf_{\boldsymbol{\lambda} \in [0, \bar{M}]^k} g_i(a, \sigma, \boldsymbol{\lambda}).$$

By construction

$$\sum_{\alpha,\sigma} \psi_{\alpha,i}(\alpha,\sigma|R,\boldsymbol{\lambda}) g_{i,a,\sigma} > 0.$$

In a neighborhood of the non-locally stable vector $\boldsymbol{\lambda}$, $\sum g_i(a_t,\sigma_t,\boldsymbol{\lambda}) \geq \sum g_{i,\alpha_t^i,\sigma_t}$. Since $\langle \alpha_t \rangle$ and $\langle \sigma_t \rangle$ are i.i.d. processes,

$$\lim_{T\to\infty} P\left(\frac{1}{T}\sum_{t=0}^{T} g_{i,\alpha_t^i,\sigma_t} > 0\right) = 1$$

by the Strong Law of Large Numbers. Let τ_1 be the first time beliefs enter the set $[0,M]^k$ and never leave for type θ_i. This implies that

$$\lambda_{i,t} = \lambda_{i,\tau_1} + \sum_{i=\tau_1}^{t-1} g_i(\alpha_t,\sigma_t,0) \to \infty \text{ a.s.},$$

which is a contradiction. □

The following Lemma is an intermediate result used in Lemma 3.

Lemma 10. *For any* $\log \boldsymbol{\lambda} \in \mathbb{R}^k$, $\inf g(L,\sigma_L,\boldsymbol{\lambda}) > 0$ *and* $\sup g(R,\sigma_R,\boldsymbol{\lambda}) < 0$.

Proof. L actions are always perceived to occur (weakly) more frequently in state L and R actions are always perceived to occur more frequently in state R. Similarly, σ_R signals are always perceived to occur more frequently in state R and σ_L signals are perceived to occur more frequently in σ_L. Under Assumption 2, agents either believe there is a positive mass of autarkic types or the public signal is informative. Suppose type θ_i believes there is a positive mass of autarkic types. Following an L action, $\log \lambda_i$ updates to

$$\log \frac{Pr(L|\theta \in \Theta_A, \omega = L)\hat{\pi}^i(\Theta_A) + \hat{\pi}^i(\Theta_S) Pr(L|\theta \notin \Theta_A, \omega = L)}{Pr(L|\theta \in \Theta_A, \omega = R)\hat{\pi}^i(\Theta_A) + \hat{\pi}^i(\Theta_A) Pr(L|\theta \notin \Theta_A, \omega = R)}$$

where Pr is the misperceived probability. This is bounded below by

$$\log \frac{Pr(L|\theta \in \Theta_A, \omega = L)\hat{\pi}^i(\Theta_A)}{Pr(L|\theta \in \Theta_A, \omega = R)\hat{\pi}^i(\Theta_A) + \hat{\pi}^i(\Theta_S)1} > 0.$$

Similar logic holds for R actions.

Suppose type θ_i believes that the public signal is informative. Then the minimal informativeness of σ_L is always positive, so the log-likelihood ratio updates are

bounded below uniformly. □

Proof of Lemma 3. Suppose $\mathbf{0} \in \Lambda_L$. Let $\mathcal{J}$ denote the locally stable neighborhood defined in Lemma 2 and choose $M > 0$ so that if $\log \boldsymbol{\lambda} \in \mathbb{R}^k \setminus [-M, M]^k$ then it is contained in one of the neighborhoods of stationary points constructed in Lemma 2. Let N be the minimal number of consecutive (R, σ_R) action and signal pairs required for the likelihood ratio of all sociable types to reach $\mathcal{J}$, given initial likelihood ratio $\log \boldsymbol{\lambda}_0 \in [-M, M]^k$. N exists by Lemma 10.

Let τ_3 be the first time that λ_i enters $\mathcal{J}$ for all $\theta_i \in \Theta_S$, and let τ_4 be the first time any type's beliefs leave after entering. We know that $P(\tau_3 < \infty) = 1$, since if they did not, $\log \boldsymbol{\lambda} \in [-M, M]^k$ infinitely often, and the probability of transitioning from $[-M, M]^k$ to $\mathcal{J}$ is bounded below by the probability of observing N action and signal pairs (R, σ_R).

Also, $P(\tau_4 < \infty) < 1$, since beliefs enter $\mathcal{J}$ and never leave with positive probability due to local stability. So $P(\boldsymbol{\lambda}_t \notin \mathcal{J}\ i.o.) = 0$. Let τ_5 be the first time the likelihood ratio enters the $\mathcal{J}$ set and stays there forever. $P(\tau_5 < \infty) = 1$, so the likelihood ratio remains in the $\mathcal{J}$ almost surely. By Lemma 2, if the likelihood ratio remains in $\mathcal{J}$ forever, then beliefs must converge. □

Let $\mathcal{J}$ be the neighborhood constructed in Lemma 2 and let $M > 0$ be such that if $\boldsymbol{\lambda} \in \mathbb{R} \setminus [-M, M]$ then it is contained in one of the neighborhoods constructed in Lemma 2 (either the neighborhood where beliefs converge with positive probability or the nbhd where beliefs leave with probability 1).

Proof of Lemma 4. Let $k = 2$ and first suppose signals are bounded. The linear equation

$$A(0,0)\begin{pmatrix} c \\ d \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

has a solution where (c, d) are positive if and only if $det(A(0,0)) > 0$. Therefore, if $\det(A(0,0)) > 0$ then there exist c, d such that

$$c \log \frac{\hat{\psi}_i(L|0,L)}{\hat{\psi}_i(L|0,R)} + d \log \frac{\hat{\psi}_i(R|0,L)}{\hat{\psi}_i(R|0,R)}$$

is negative for θ_1 and positive for θ_2. Moreover, for some $-M' < -M$ if $\log \boldsymbol{\lambda} \in (-\infty, -M']^k$ of 0, this will still hold.

Let

$$\xi_{1,t} = \sum g_{1,a_t,\sigma_t}$$

where $g_{1,a,\sigma} = \sup_{\log \boldsymbol{\lambda} \in (-\infty,-M']^2} g_1(a,\sigma,\boldsymbol{\lambda})$ and

$$\xi_{2,t} = \sum g_{2,a_t,\sigma_t}$$

where $g_{2,a_t,\sigma_t} = \inf_{\lambda \in (-\infty,-M'']^2} g_2(a,\sigma,\boldsymbol{\lambda})$.

For any $K_2 > 0$ there exists a sequence of actions $(a_t,\sigma_t)_{t=1}^T$ and a finite number K_1, where T is some finite number such that

1. $\xi_{1,T} < 0$
2. $\xi_{2,T} > K_2$
3. $\xi_{1,t} < K_1$ for all t.

This sequence exists because there are rational numbers P and Q such that

$$P \log\left(\frac{\hat{\psi}_i(L|L,\lambda)}{\hat{\psi}_i(L|R,\lambda)}\right) + Q \log\left(\frac{\hat{\psi}_i(R|L,\lambda)}{\hat{\psi}_i(R|R,\lambda)}\right)$$

is less than 0 for the first type and is greater than 0 for the second. So there exists a non-zero $N \in \mathbb{N}$ such that NP and NQ are integers. Then after NP $(L,\sigma_L)'s$ and NQ (R,σ_R)'s, λ_1 decreases and λ_2 increases. So a finite sequence of actions that satisfies the three properties exists.

Let $\boldsymbol{\lambda}_0 \in (-\infty,-M]^2$. As long as $\log \boldsymbol{\lambda} \in (-\infty,-M']^2$, ξ_1 bounds the updates to θ_1's beliefs above, $\log \boldsymbol{\lambda}_{1,t} - \log \boldsymbol{\lambda}_{1,0} < \xi_{1,t}$, and ξ_2 bounds θ_2's beliefs below $\log \boldsymbol{\lambda}_{2,t} - \log \boldsymbol{\lambda}_{2,0} > \xi_{2,t}$.

Let $(-\infty,-M']$ be the set of log-likelihood ratios constructed in Lemma 2 around $\lambda^* = 0^k$. The above construction implies that for $K_2 = 1$ if $\log \lambda_1 < -M' - K_1 - K$ for any $K > 0$ where K_1 is the K_1 that corresponds to $K_2 = 1$, then there exists a sequence of actions such that $\lambda_1 < -M' - K$ and λ_2 is outside of $(-\infty,M']$ if $\boldsymbol{\lambda}_1 \in (-\infty,-M']$.

Let N_1 be the smallest number of consecutive (L,σ_L) actions and signals it takes for $\lambda_{2,t}$ to go from a point outside of $(-\infty,-M']$ to $[M,\infty)$. This can at most increase $\log \lambda_{1,t}$ by $K < \infty$ by lemma 10. So, if $\log \lambda_{1,t} < \sup -M' - K - K_1$ for large enough K_1, then there exists a finite sequence of S actions such that

1. $\log \lambda_{1,t} < -M'$ for all t.

2. $\log \lambda_{2,S} > M$.

Since any finite sequence of actions occurs with positive probability, and beliefs converge with positive probability once the $\log \boldsymbol{\lambda}$ enters $(-\infty, -M'] \times [M, \infty)$, beliefs converge with positive probability to $\boldsymbol{\lambda}^* = (0, \infty)$ if this is true.

So, with positive probability, from any initial $\boldsymbol{\lambda}_0 \in (0, \infty)^2$, $\boldsymbol{\lambda}$ enters a neighborhood of $(0, \infty)$ where beliefs converge with positive probability. So, disagreement occurs with positive probability. □

Proof of Lemma 5. By Lemma 4, we can separate θ_{i^*} and θ_{j^*}, since

$$A(\boldsymbol{\lambda}_{(i^*,j^*)}) = \begin{pmatrix} g_{i^*}(L, \sigma_L, \boldsymbol{\lambda}) & g_{i^*}(R, \sigma_R, \boldsymbol{\lambda}) \\ g_{j^*}(L, \sigma_L, \boldsymbol{\lambda}) & g_{j^*}(R, \sigma_R, \boldsymbol{\lambda}) \end{pmatrix}$$

has positive determinant by the assumption that $\theta_{i^*} \succ_{\boldsymbol{\lambda}} \theta_{j^*}$. As in the proof of Lemma 4, let $\xi_{i,t} \equiv \sum g_{i,a_t,\sigma_t}$, where

$$g_{j,a_t,\sigma_t} \equiv \inf_{\log \boldsymbol{\lambda} \in (-\infty, -M']^k} g_j(a, \sigma, \boldsymbol{\lambda})$$

for all $j > m$, and

$$g_{i,a_t,\sigma_t} \equiv \sup_{\log \lambda \in (-\infty, -M']^k} g_i(a, \sigma, \boldsymbol{\lambda})$$

for all other $i \leq m$. By the argument from Lemma 4, for any $K_{j^*} \in \mathbb{R}_+$, there exists a sequence of actions, a $T \in \mathbb{N}$, and a $K_{i^*} \in \mathbb{R}_+$ such that

$$\begin{aligned} \xi_{i^*,T} &< 0 \\ \xi_{j^*,T} &> K_{j^*} \\ \xi_{i^*,t} &< K_{i^*} \text{ for all } t \leq T \end{aligned}$$

Moreover, the $\succeq$ relation implies that for any $\theta_i \succeq \theta_j$, the log likelihood ratio for θ_i has increased more (or decreased less) than the log likelihood ratio for θ_j as long as both likelihood ratios remain in $(-\infty, -M']$. Therefore, there exists a sequence of actions such that $\lambda_i \in (\infty, -M' - K]$ for all types θ_i where $i \leq m$ for any $K > 0$ and $\lambda_j \notin -(\infty, -M']$ for all other types j.

Let N_1 be the minimum number of consecutive (L, σ_L) actions and signals such that for any type $j > m$ such that if $\log \boldsymbol{\lambda}_{j,t} = -M'$ then $\log \boldsymbol{\lambda}_{j,t+N_1} > M$ (denote this by $\mathcal{J}_D^j$ for each type θ_j). This minimum number of (L, σ_L) actions exists by Lemma 10. There exists a maximum amount that these N_1 (L, σ_L)s can increase $\log \lambda_j$ for any $j \leq m$. Let K be the maximum amount that this can increase $\log \lambda_j$ for any $j \leq m$ (i.e. after N_1 (L, σ_L)s, $\log \lambda_{j,N_1} - \log \lambda_{j,0} < K$ for any $\boldsymbol{\lambda}_0$).

Therefore, from any initial $\boldsymbol{\lambda}_0$, there exists a sequence such that $\lambda_j \in (-\infty, -M' - K]$ for all types $i \leq m$ and $\lambda_i \notin (-\infty, -M']$ for all other types. With positive probability, N_1 consecutive (L, σ_L)s occur after the likelihood ratio reaches this set. After N_1 consecutive (L, σ_L)s $\log \boldsymbol{\lambda} \in (-\infty, -M]^m \times [M, \infty)^{k-m}$. Therefore $\boldsymbol{\lambda}$ enters a neighborhood of $(0^m, \infty^{k-m})$ where beliefs converge with positive probability. So, disagreement occurs with positive probability. □

Proof of Lemma 6. Let $\mathcal{J}$ be the neighborhood constructed in Lemma 2 and let $M > 0$ be such that if $\log \boldsymbol{\lambda} \in \mathbb{R} \setminus [-M, M]$ then it is contained in one of the neighborhoods constructed in lemma 2 (either the neighborhood where beliefs converge with positive probability or the nbhd where beliefs leave with probability 1). Let $\tau_1 = \inf\{t : \boldsymbol{\lambda}_t \in \mathcal{J}\}$.

First, we show that $Pr(\tau_1 < \infty) = 1$. Suppose $Pr(\tau_1 < \infty) < 1$. If $\mathbf{0}$ or $\boldsymbol{\infty}$ are stable, then there exists a sequence of actions such that for any point in $\log \boldsymbol{\lambda}_0 \in [-M, M]^k$ enters the part of $\mathcal{J}$ containing a locally stable point. By the proof of Lemma 5 for any disagreement point there exists a finite sequence of actions such that from any point in $\log \boldsymbol{\lambda}_0 \in [-M, M]^k$ beliefs eventually enter the part of $\mathcal{J}$ containing a locally stable point. Therefore, the probability of entering the part of $\mathcal{J}$ from any point in $[-M, M]^k$ in finite time is bounded away from 0. Moreover, if beliefs never entered a part of $\mathcal{J}$ containing a locally stable point, then with probability 1 $\log \boldsymbol{\lambda}_t \in [-M, M]^k$ i.o. Since the probability of entering $\mathcal{J}$ from $[-M, M]^k$ is bounded away from 0, beliefs must eventually enter $\mathcal{J}$ near a locally stable point. Therefore, $Pr(\tau_1 < \infty) = 1$. Let $\tau_2 = \inf\{t > \tau_1 : \boldsymbol{\lambda}_t \notin \mathcal{J}\}$. By Lemma 2, $Pr(\tau_2 < \infty) < 1$. Therefore, $Pr(\boldsymbol{\lambda}_t \notin \mathcal{J} \text{ i.o.}) = 0$, so beliefs converge almost surely. □

Proof of Lemma 7. Suppose beliefs converged to a non-stationary point $\boldsymbol{\lambda}^* \in (0, \infty)^k$ with positive probability. After an L action and a σ_L public signal, the likelihood ratio must increase for all sociable types, by Assumptions 1-3. Moreover,

for any $M > 0$, if $\log \lambda_i \in [-M, M]$, this update is bounded uniformly away from 0. For $\varepsilon > 0$, let $B_\varepsilon(\boldsymbol{\lambda}^*)$ be an open ε-ball around $\boldsymbol{\lambda}^*$. For sufficiently small $\varepsilon > 0$, if $\boldsymbol{\lambda} \in B_\varepsilon(\boldsymbol{\lambda}^*)$, then observing (L, σ_L) causes the likelihood ratio to leave $B_\varepsilon(\boldsymbol{\lambda}^*)$. The probability of (L, σ_L) never occurring converges to 0 as $t \to \infty$. Therefore, the likelihood ratio leaves any ε-ball around a non-stationary point almost surely. □

Proof of Corollary 1. Given Assumption 1, if the correctly specified type θ_C has a stationary limit belief, then the support of the limit belief is a subset of $\{0, \infty\}$. Also, for θ_C, the perceived probability of each action is equal to the true probability, $\hat{\psi}_C = \psi$. Therefore, $\lambda_{C,t}$ is a martingale for any $\{\Theta, \pi\}$. By the Martingale Convergence Theorem, $\lambda_{C,t}$ converges almost surely to a limit random variable λ_∞ with $\text{supp}(\lambda_\infty) \subset [0, \infty)$. This rules out incorrect and non-stationary incomplete learning. Therefore, 0 is the only candidate limit point and it must be that $\lambda_{C,t} \to 0$ almost surely. □

A.3 Proofs of Theorem 2 and 3

Proof of Theorem 2. Assume Assumptions 1, 2 and 3 and suppose $\omega = R$.

Part 1: For any type θ, the function $(\hat{\pi}^\theta, \rho^\theta) \mapsto \hat{\psi}^\theta(a, \sigma|\omega, \boldsymbol{\lambda})$ is continuous. By continuity, given $\delta_2 > 0$, there exists a $\delta > 0$ such that if $||\hat{\pi}^{\theta_i} - \pi|| < \delta$ and $||\rho^{\theta_i} - r|| < \delta$ for all sociable types θ_i, then $||\hat{\psi}_i(a, \sigma|\omega, \cdot) - \psi(a, \sigma|\omega, \cdot)|| < \delta_2$ for $(a, \sigma) \in \{L, R\} \times \{\sigma_L, \sigma_R\}$. Thus, δ_2 can be chosen to be sufficiently small so that at every $\boldsymbol{\lambda} \in \{0, \infty\}$

$$|\gamma(\boldsymbol{\lambda}) - \gamma_C(\boldsymbol{\lambda})| < \min_{i, \boldsymbol{\lambda} \in \{0,\infty\}^k} |\gamma_i(\boldsymbol{\lambda})|/2,$$

where γ_C is the corresponding γ for the model where $\hat{\pi} = \pi$ and $\rho^\theta = r$. So δ can be chosen so that the sign of γ in the misspecified model matches the sign of γ_C at all stationary points. Since

$$\gamma_C(\boldsymbol{\lambda}) = \sum_{a,\sigma} \psi(a, \sigma, R) \log \frac{\psi(a, \sigma, L)}{\psi(a, \sigma, L)} < 0,$$

by Theorem 3, learning is complete.

Part 2: Let $\varepsilon = \min_{i, \boldsymbol{\lambda} \in \{0,\infty\}^k} |\gamma_i(\boldsymbol{\lambda})|/2$. There exists a $\delta > 0$ such that if $||r^\theta - r|| < \delta$ and $||\rho^\theta - r|| < \delta$ for all types θ, $||\hat{F}^{L,i} - F^L|| < \delta$, and $\left|\hat{\pi}^\theta(\Theta_A) - \pi(\Theta_A)\right| < \delta$ for

all sociable types θ such that:

1. The empirical frequency with which autarkic types θ play each action is

$$|F^{\omega}(\frac{1}{2}) - \hat{F}^{\omega}(r^{\theta})^{-1}(\frac{1}{2}, \frac{1}{2})| < \varepsilon,$$

since there always exists a δ sufficiently small such that

$$|| \int \frac{1 - r(p)}{r(p)} d\hat{F}^{L} - \int \frac{1-p}{p} dF^{L} || < \varepsilon.$$

2. For all sociable types, at any stationary $\boldsymbol{\lambda}$, the probability of any action is either 1 or 0 in both the misspecified and correctly specified models.

3. Binary signals imply that the perceived probability of each signal is continuous. At any stationary vector $\boldsymbol{\lambda}$, the perceived probabilities of each public signal satisfies $|G(\sigma_L) - \hat{G}^{\theta}(\sigma_L)| < \varepsilon$.

This implies $\hat{\psi}_i$ can be made sufficiently close to ψ at every stationary vector so that $||\gamma_C(\boldsymbol{\lambda}) - \gamma(\boldsymbol{\lambda})|| < \varepsilon$ where γ_C is γ in the correctly specified model. Therefore, γ has the same sign as γ_C for all stationary $\boldsymbol{\lambda}$. By Theorem 3, learning is complete. □

Proof of Theorem 3. Assume Assumptions 1, 2 and 3 and suppose $\omega = R$. For any sociable type θ_i, the mapping $(\hat{\psi}_i(a, \sigma | R, \boldsymbol{\lambda})) \mapsto \gamma(\boldsymbol{\lambda})$ is continuous, and by the concavity of the log operator, is negative when $||\hat{\psi}_i(a, \sigma | R, \cdot) - \psi(a, \sigma | R, \cdot)|| = 0$. Therefore, there exists a $\delta_i > 0$ such that if $||\hat{\psi}_i(a, \sigma | R, \cdot) - \psi(a, \sigma | R, \cdot)|| < \delta_i$ for $(a, \sigma) \in \{L, R\} \times \{\sigma_L, \sigma_R\}$, then $\gamma_i(\boldsymbol{\lambda}) < 0$ at all stationary vectors. Therefore, any locally stable point must have $\lambda_i = 0$. This holds for all sociable types θ_i, so $\boldsymbol{\lambda} = \mathbf{0}$ is the unique locally stable point. By Theorem 1, the likelihood ratio converges to $\mathbf{0}$ almost surely and learning is complete. □

A.4 Proofs from Section 4

Proof of Theorem 4. Suppose $\omega = R$. Let $x \equiv F^{R}(1/2)$ be the probability a level-1 type plays action R. At a stationary vector (λ_2, λ_3), whether this vector is in

Λ is determined by the sign of

$$\gamma_i(\lambda_2,\lambda_3) = \psi(R|R,\lambda_2,\lambda_3)\log\frac{\hat{\psi}_i(R|L,\lambda_2,\lambda_3)}{\hat{\psi}_i(R|R,\lambda_2,\lambda_3)} + \psi(L|R,\lambda_2,\lambda_3)\log\frac{\hat{\psi}_i(L|L,\lambda_2,\lambda_3)}{\hat{\psi}_i(L|R,\lambda_2,\lambda_3)}$$

for each type. Consider the level-2 type. Since $x > 1/2$,

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

Therefore, $(0,0)$, $(0,\infty)$ and (∞,∞) are locally stable for level-2 and $(\infty,0)$ is not locally stable. Consider the level-3 type.

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

If $\gamma_3(\infty,\infty) > 0$, then $(\infty,\infty) \in \Lambda$. From these expressions, $\gamma_3(\infty,\infty)$ is positive at $p = 0$, decreasing in p and negative at $p = 1$. Therefore, there exists a p_2 such that for $p < p_2$, $(\infty,\infty) \in \Lambda$, and for $p > p_2$, $(\infty,\infty) \notin \Lambda$. If $\gamma_3(0,\infty) > 0$, then $(0,\infty) \in \Lambda$ and if $\gamma_3(0,0) < 0$, then $(0,0) \in \Lambda$. The expressions $\gamma_3(0,0) < \gamma_3(0,\infty)$ are both negative at $p = 0$, increasing in p and positive at $p = 1$. Therefore, there exists $p_1 < p_3$ such that $(0,0) \in \Lambda$ for $p < p_3$ and $(0,0) \notin \Lambda$ for $p > p_3$, while $(0,\infty) \notin \Lambda$ for $p < p_1$ and $(0,\infty) \in \Lambda$ for $p > p_1$.

It immediately follows from Theorem 1 that the agreement outcomes $(0,0)$ and (∞,∞) arise with positive probability if and only if they are in Λ, and when at least one agreement vector is in Λ, beliefs converge (i.e. for $p < p_3$). It remains to show that if $(0,\infty) \in \Lambda$, then $(0,\infty)$ is total informativeness ranked, which establishes that this outcome arises with positive probability if and only if it is in Λ, and also

establishes belief convergence for the case of $p > p_3$. To apply Lemma 5, it must be that for some $\boldsymbol{\lambda} \in \{(0,0),(\infty,\infty)\}$, $\theta_2 \succeq_{\boldsymbol{\lambda}} \theta_3$. This is satisfied for $\succeq_{(0,0)}$. In particular,

$$|g_2(R,(0,0))| = \left|\log\left(\frac{1-x}{x}\right)\right| > \left|\log\left(\frac{p+(1-p)(1-x)}{p+(1-p)x}\right)\right| = |g_3(R,(0,0))|$$
$$g_2(L,(0,0)) = \log\left(\frac{x}{1-x}\right) = g_3(L,(0,0)).$$

Intuitively, both types make the same inference from L actions around $(0,0)$, which they believe must come from a level-1 type. But the level-2 type believes that R actions are stronger evidence of state R than the level-3 type, because the level-3 type underweights the informativeness of these actions to account for the possibility of level-2 types. Therefore, the conditions for the pairwise informativeness order defined in Definition 2 are satisfied, $\theta_2 \succeq_{(0,0)} \theta_3$. Given that these are the only two types, the conditions in Definition 6 for $(0,\infty)$ to be total informativeness ranked are also satisfied.

Proof of Theorem 5. Both partisan and unbiased types believe that share α of agents are autarkic. Partisan types think these autarkic types are also partisan, while unbiased types think these autarkic types are also unbiased. Let $x_P^\omega(\nu) \equiv F^\omega(0.5^{1/\nu})$ be the probability that the partisan autarkic type plays action R and $x_U^\omega \equiv F^\omega(0.5)$ be the probability that the unbiased autarkic type plays action R in state ω. Then $x_P^R(\nu) \leq x_U^R$ and $x_P^L(\nu) \leq x_U^L$ for all $\nu \in (0,1)$, since partisan types slant information in favor of state L. Moreover, action R occurs more often in state R, so $x_U^R > x_U^L$ and $x_P^R(\nu) > x_P^L(\nu)$ for all $\nu \in (0,1)$. Unbiased types believe that autarkic types play action R with probability x_U^ω, and partisan types believe that autarkic types play action R with probability $\hat{F}^{\omega,P}(0.5^{1/\nu}) = F^\omega(0.5) = x_U^\omega$.

Let $\gamma_P^{\nu,q}(\boldsymbol{\lambda})$ be the value of $\gamma_P(\boldsymbol{\lambda})$ in the model with partisan bias level ν and frequency q, with an analogous definition of $\gamma_U^{\nu,q}(\boldsymbol{\lambda})$. Since partisan and unbiased sociable types have the same perceived probability of each action, beliefs can never separate and asymptotic disagreement is not possible. Additionally, $\gamma_P^{\nu,q} = \gamma_U^{\nu,q}$, and therefore, we only need to check the sign of $\gamma_P^{\nu,q}(0,0)$ to determine whether $(0,0)$ is locally stable, and the sign of $\gamma_P^{\nu,q}(\infty,\infty)$ to determine whether (∞,∞) is locally stable. Recall that global stability immediately follows for agreement vectors.

Suppose $\omega = R$. To determine whether $(\infty,\infty) \in \Lambda$ at (ν,q), we need to determine

the sign of

$$\gamma_P^{\nu,q}(\infty,\infty) = \psi^{\nu,q}(R|R,\infty^2)\log\left(\frac{\hat{\psi}_P(R|L,\infty^2)}{\hat{\psi}_P(R|R,\infty^2)}\right) + \psi^{\nu,q}(L|R,\infty^2)\log\left(\frac{\hat{\psi}_P(L|L,\infty^2)}{\hat{\psi}_P(L|R,\infty^2)}\right),$$

where

$$\begin{aligned}
\hat{\psi}_P(R|\omega,\infty^2) &= \alpha x_U^{\omega} \\
\hat{\psi}_P(L|\omega,\infty^2) &= \alpha(1-x_U^{\omega}) + 1 - \alpha \\
\psi^{\nu,q}(R|R,\infty^2) &= \alpha q x_P^R(\nu) + \alpha(1-q)x_U^R \\
\psi^{\nu,q}(L|R,\infty^2) &= \alpha q(1-x_P^R(\nu)) + \alpha(1-q)(1-x_U^R) + 1 - \alpha.
\end{aligned}$$

If $\nu = 1$, then $x_P^R(1) = x_U^R$, so $\psi^{1,q}(R|R,\infty^2) = \hat{\psi}_P(R|R,\infty^2)$ and $\psi^{1,q}(L|R,\infty^2) = \hat{\psi}_P(L|R,\infty^2)$. Therefore, $\gamma_P^{1,q}(\infty,\infty) < 0$ by the concavity of the log operator, for any q. At $\nu = 0$ and $q = 1$, $x_P^R(0) = 0$ and therefore $\psi^{0,1}(R|R,\infty^2) = 0$. Note that R actions decrease the likelihood ratio, $\log\left(\frac{\hat{\psi}_P(R|L,\infty^2)}{\hat{\psi}_P(R|R,\infty^2)}\right) < 0$, while L actions increase the likelihood ratio, $\log\left(\frac{\hat{\psi}_P(L|L,\infty^2)}{\hat{\psi}_P(L|R,\infty^2)}\right) > 0$, independently of q and ν. Therefore, $\gamma_P^{0,1}(\infty,\infty) > 0$. Also, $\psi^{\nu,q}(R|R,\infty^2)$ is strictly decreasing in q and strictly increasing in ν, since $x_P^R(\nu)$ is strictly increasing in ν. Therefore, $\gamma_P^{\nu,q}(\infty,\infty)$ is strictly decreasing in ν and increasing in q. Therefore, there exists a cutoff q_1 such that for $q > q_1$, there exists a cutoff $\nu_1(q) > 0$ such that for $\nu < \nu_1(q)$, $\gamma_P^{\nu,q}(\infty,\infty) > 0$ and (∞,∞) is locally stable, while for $\nu > \nu_1(q)$, $\gamma_P^{\nu,q}(\infty,\infty) < 0$ and (∞,∞) is not locally stable.

To determine whether $(0,0) \in \Lambda$ at (ν,q), we need to determine the sign of

$$\gamma_P^{\nu,q}(0,0) = \psi^{\nu,q}(R|R,0^2)\log\left(\frac{\hat{\psi}_P(R|L,0^2)}{\hat{\psi}_P(R|R,0^2)}\right) + \psi^{\nu,q}(L|R,0^2)\log\left(\frac{\hat{\psi}_P(L|L,0^2)}{\hat{\psi}_P(L|R,0^2)}\right),$$

where

$$\begin{aligned}
\hat{\psi}_P(R|\omega,0^2) &= \alpha x_U^{\omega} + 1 - \alpha \\
\hat{\psi}_P(L|\omega,0^2) &= \alpha(1-x_U^{\omega}) \\
\psi^{\nu,q}(R|R,0^2) &= \alpha q x_P^R(\nu) + \alpha(1-q)x_U^R + 1 - \alpha \\
\psi^{\nu,q}(L|R,0^2) &= \alpha q(1-x_P^R(\nu)) + \alpha(1-q)(1-x_U^R).
\end{aligned}$$

If $\nu = 1$, then $x_P^R(1) = x_U^R$, so $\psi^{1,q}(R|R,0^2) = \hat{\psi}_P(R|R,0^2)$ and $\psi^{1,q}(L|R,0^2) =$

$\hat{\psi}_P(L|R,0^2)$. Therefore, $\gamma_P^{1,q}(0,0) < 0$ by the concavity of the log operator. At $\nu = 0$ and $q = 1$, then $x_P^R(0) = 0$, and therefore $\psi^{0,1}(R|R,0^2) = 1-\alpha$. Therefore, $\gamma_P^{0,1}(0,0) > 0$. Moreover, $\gamma_P^{\nu,q}(0,0)$ is strictly increasing in q and strictly decreasing in ν, since $x_P^R(\nu)$ is strictly increasing in ν. Therefore, there exists a cut-off $q_2 < 1$ such that for any $q > q_2$, there exists a cutoff $\nu_2(q)$ such that for $\nu < \nu_2(q)$, $\gamma_P^{\nu,q}(0,0) > 0$ and $(0,0)$ is not locally stable, and for $\nu > \nu_2(q)$, $\gamma_P^{\nu,q}(0,0) < 0$ and $(0,0)$ is locally stable.

Suppose $\omega = L$. Then $\gamma^{1,q}(\infty,\infty) > 0$ and $\gamma^{1,q}(0,0) > 0$ for all $q \in [0,1]$, since only correct learning can occur for $\nu = 1$. The only change in the above expressions is that now the true measures are taken for state L, rather than state R. Therefore, all of the comparative statics on γ are preserved. As above, for any q, $\gamma^{\nu,q}(0,0)$ is decreasing in ν. Therefore, $\gamma^{\nu,q}(0,0) > 0$ for all ν and q, and incorrect learning is never locally stable. Also, for any q, $\gamma^{\nu,q}(\infty,\infty)$ is decreasing in ν. Therefore, $\gamma^{\nu,q}(\infty,\infty) > 0$ for all ν and q, and correct learning is always locally stable.

Proof of Theorem 6. Theorem 1 establishes that that $\lambda \to 0$ almost surely for a correctly specified type. Moreover, since the beliefs of the correctly specified type are a martingale, $\gamma_U(\lambda_P,\lambda_U) < 0$ for all $(\lambda_P,\lambda_U) \in [0,1]^2$. In order to establish this result, all that remains is to sign $\gamma_P^{\nu,q}(0,0)$ and $\gamma_P^{\nu,q}(\infty,0)$. Since the partisan type believes that all types are partisan, $\hat{\psi}_P$ remains unchanged from the proof of Theorem 5. But now at $(0,0)$,

$$\begin{aligned}
\psi^{\nu,q}(R|R,0^2) &= \alpha q x_P^R(\nu) + \alpha(1-q)x_U^R + (1-\alpha) \\
\psi^{\nu,q}(L|R,0^2) &= \alpha q(1-x_P^R(\nu)) + \alpha(1-q)(1-x_U^R).
\end{aligned}$$

and at $(0,\infty)$,

$$\begin{aligned}
\psi^{\nu}(R|R,(\infty,0)) &= \alpha(1-q)x_U^R + \alpha q x_P^R(\nu) + (1-\alpha)(1-q) \\
\psi^{\nu}(L|R,(\infty,0)) &= \alpha(1-q)(1-x_U^R) + \alpha q(1-x_P^R(\nu)) + \alpha q.
\end{aligned}$$

As before, as ν increases, $\psi^{\nu}(R|R,\boldsymbol{\lambda})$ increases and $\psi^{\nu}(L|R,\boldsymbol{\lambda})$ decreases. So, as long as $\gamma_P^{0,q}(0,0) > 0$ and $\gamma_P^{0,q}(\infty,0) > 0$, both these cutoffs exist.

Proof of Theorem 7. Suppose $\omega = R$. Let $\gamma^k(\lambda)$ be the value of $\gamma(\lambda)$ in the model with parameter k. Specifically,

$$\gamma^k(0) = dF^R(\sigma_L)\log\left(\frac{\sigma_L + \frac{1}{k}(0.5-\sigma_L)}{1-\sigma_L-\frac{1}{k}(0.5-\sigma_L)}\right) + dF^R(\sigma_R)\log\left(\frac{\sigma_R}{1-\sigma_R}\right)$$
$$\gamma^k(\infty) = dF^R(\sigma_L)\log\left(\frac{\sigma_L}{1-\sigma_L}\right) + dF^R(\sigma_R)\log\left(\frac{\sigma_R + \frac{1}{k}(0.5-\sigma_R)}{1-\sigma_R-\frac{1}{k}(0.5-\sigma_R)}\right)$$

since $\varrho(0) = \varrho(1) = 1$. With a single type, global stability follows immediately from local stability. Determining how the sign of $\gamma^k(0)$ and $\gamma^k(\infty)$ varies with k will characterize the local stability set, and therefore, asymptotic learning outcomes.

We know that learning is almost surely correct in the correctly specified model, so it must be that $\gamma^\infty(0) < 0$ and $\gamma^\infty(\infty) < 0$. The bias ρ is continuous in k, so $\gamma^k(0)$ and $\gamma^k(\infty)$ are continuous in k. The informativeness of σ_L when beliefs are near 0 is

$$\log\left(\frac{\sigma_L + \frac{1}{k}(0.5-\sigma_L)}{1-\sigma_L-\frac{1}{k}(0.5-\sigma_L)}\right).$$

This expression is increasing in k since $\sigma_L > 1/2$. Increasing k has no effect on the informativeness of σ_R when beliefs are near 0. Therefore, $\gamma^k(0)$ is increasing in k. Since $\gamma^\infty(0) < 0$, $\gamma^k(0) < 0$ for all k. This means that correct learning occurs with positive probability for all $k > 1$.

In contrast, the informativeness of σ_R when beliefs are near ∞,

$$\left|\log\left(\frac{\sigma_R + \frac{1}{k}(0.5-\sigma_R)}{1-\sigma_R-\frac{1}{k}(0.5-\sigma_R)}\right)\right|,$$

is increasing in k. Therefore, $\gamma^k(\infty)$ is decreasing in k. Since $\gamma^\infty(\infty) < 0$, incorrect learning can only occur if k is low enough. At $k = 1$, σ_R is perceived to be uninformative near ∞, and the likelihood ratio moves towards state L,

$$\gamma^1(\infty) = dF^R(\sigma_L)\log\left(\frac{\sigma_L}{1-\sigma_L}\right) + dF^R(\sigma_R)\log 1 > 0.$$

Therefore, there exists a cut-off $\bar{k} > 1$ such that for $k < \bar{k}$, $\gamma^k(\infty) > 0$ and incorrect learning occurs with positive probability.

B Examples of Nested Models

This paper nests the boundedly rational models of several other papers, including Rabin and Schrag (1999) and Epstein et al. (2010).

B.1 Rabin and Schrag (1999)

Rabin and Schrag (1999) examines individual learning with confirmation bias. Agents receive a binary signal, but if they receive a signal that goes against their prior beliefs then with probability q they misinterpret that signal as the other signal (which agrees with their prior belief). In order to nest this model, a slight extension must be made to the framework we've outlined. In particular, this requires four public signals and the mapping ρ must be able to map two public signals that induce the same posterior to different misspecified beliefs. It is straightforward to extend all arguments made in this paper to this case.

This is a misspecified model with one type θ. There are 4 public signals σ_{L_1}, σ_{L_2}, σ_{R_1}, σ_{R_2}. All L signals induce the same posterior and all R signals induce the same posterior. Conditional on seeing an L signal, σ_{L_2} is draw with probability q. Similarly, for σ_{R_1} and $Pr(\sigma_{L_1} \text{ or } \sigma_{L_2}|\omega = R) = Pr(\sigma_{R_1} \text{ or } \sigma_{R_2}|\omega = L) = \sigma < 1/2$. If $\lambda > 1$, then $\rho(\sigma_{L_2}) = \sigma$ and all other signals are interpreted correctly. If $\lambda < 1$ then $\rho(\sigma_{R_1}) = 1 - \sigma$ and all other signals are interpreted correctly.

The parameter q indexes the degree of confirmation bias. Higher q means it is more likely that agents misinterpret signals that go against their prior. Under this specification,

$$\gamma(0) = (1-q)\left(\sigma \log\left(\frac{1-\sigma}{\sigma}\right) + (1-\sigma)\log\left(\frac{\sigma}{1-\sigma}\right)\right) + q\log\left(\frac{\sigma}{1-\sigma}\right).$$

and

$$\gamma(\infty) = (1-q)\left(\sigma \log\left(\frac{1-\sigma}{\sigma}\right) + (1-\sigma)\log\left(\frac{\sigma}{1-\sigma}\right)\right) + q\log\left(\frac{1-\sigma}{\sigma}\right).$$

As q increases, more weight is placed on the last term, which is negative when $\lambda = 0$ and positive when $\lambda = \infty$.

B.2 Epstein et al. (2010)

Epstein et al. (2010) considers an individual learning model where agents overweight beliefs towards the prior or towards the posterior. Specifically, an agent with prior p who would update her beliefs to $BU(p)$ instead updates to

$$(1-\alpha)BU(p)+\alpha p$$

for some $\alpha \leq 1$. When $\alpha = 0$, this is the correct model, for $\alpha > 0$ agents overweight the prior and for $\alpha < 0$, agents overweight new information. For simplicity of notation, suppose that $Pr(\sigma_L|\omega = R) = Pr(\sigma_R|\omega = L) = \sigma < 0.5$. In our framework, this is a model with a single agent type who only receives public signal σ and maps this signal to

$$\rho(\sigma,p) = \frac{\frac{\sigma(1-\alpha)}{(1-\sigma)(1-p)+ps}+\alpha}{\frac{1}{1-p}+\frac{1-2p}{1-p}\left(\frac{\sigma(1-\alpha)}{(1-\sigma)(1-p)+p\sigma}+\alpha\right)},$$

with

$$\rho(\sigma,1) = \frac{\sigma}{(1-\alpha)(1-\sigma)+(1+\alpha)\sigma},$$

which implies that $\rho(\sigma,1) = \lim_{p\to 1}\rho(\sigma,1)$.[27]

Under this misspecification, whenever an agent with prior p_t updates their beliefs, the likelihood ratio becomes

$$\lambda_{t+1} = \frac{\frac{p_t\sigma(1-\alpha)}{(1-\sigma)(1-p_t)+p_t\sigma}+\alpha p_t}{1-\frac{p_t\sigma(1-\alpha)}{(1-\sigma)(1-p_t)+p_t\sigma}-\alpha p_t}.$$

Therefore, the Bayes update is

$$p_{t+1} = \frac{p_t\sigma(1-\alpha)}{(1-\sigma)(1-p_t)+p_t\sigma}+\alpha p_t.$$

Therefore, the update rule from Epstein et al. (2010) can be represented in our framework.

[27] Epstein et al. (2010) does not identify how signals are interpreted at 0 or 1, since beliefs are stationary at these points. In order to characterize asymptotic outcomes, the tools developed in this paper show ho the limit of the update rule as $p \to 0$ or 1 can be used to characterize asymptotic outcomes of the model in Epstein et al. (2010).

Under this specification, the likelihood ratio update is

$$\lambda_{t+1}/\lambda_t = \frac{\frac{\sigma(1-\alpha)}{(1-\sigma)(1-p_t)+p_t\sigma}+\alpha}{\frac{(1-\alpha)(1-\sigma)}{(1-\sigma)(1-p_t)+p_t\sigma}+\alpha}$$

As $p \to 1$, the likelihood ratio update conveges to

$$\frac{1}{(1-\alpha)\frac{1-\sigma}{\sigma}+\alpha}$$

and as $p \to 0$, the likelihood ratio update converges to

$$\frac{\sigma(1-\alpha)}{1-\sigma}+\alpha$$

In an environment with symmetric binary signals,

$$\gamma(0) = \sigma \log[(1-\alpha)\frac{1-\sigma}{\sigma}+\alpha] + (1-\sigma)\log[(1-\alpha)\frac{\sigma}{1-\sigma}+\alpha],$$

and

$$\gamma(\infty) = \sigma \log \frac{1}{(1-\alpha)\frac{\sigma}{1-\sigma}+\alpha} + (1-\sigma)\log\frac{1}{(1-\alpha)\frac{1-\sigma}{\sigma}+\alpha}.$$

B.3 Overestimating Bayesianism

An interesting class of models that are nested in the framework in this paper are models where agents believe that others correctly interpret their private information and update using Bayes rule, but when faced with their own decision, agents mistakenly interpret their private signal and fail to update correctly. For instance, an agent may overweight her prior when forming new beliefs, as in Epstein et al. (2010), or update beliefs in a way that favors her preferred state, while still believing that everyone else is forming beliefs using the correctly specified model. In such a model, the agent correctly interprets the actions of others but incorrectly combines it with her own information.[28]

[28] In our framework, agents use Bayes rule to update beliefs from a sequence of actions. But the techniques in this paper can easily be generalized to incorporate fully non-Bayesian update rules of the form $\log \lambda_{t+1} = \log \lambda_t + \hat{r}\left(\log \frac{\hat{\psi}(a_t|L,\lambda_t)}{\hat{\psi}(a_t|R,\lambda_t)}, \lambda_t\right)$. Using such an updating rule, $\gamma(\lambda) = \sum \psi(a|R,\lambda)\hat{r}\left(\log \frac{\hat{\psi}(a|L,\lambda)}{\hat{\psi}(a|R,\lambda)}, \lambda\right)$. Proofs remain otherwise unchanged.

If a non-Bayesian update rule can be represented by a misspecified private belief function $r(\cdot, p)$, the analysis of this paper goes through unchanged, and $\gamma(\cdot)$ can be used to characterize the set of globally stable points. Moreover, if signals are unbounded, there exists a misspecified model that represents this learning rule.[29]

Lemma 11. *Suppose $r : [0,1] \to [0,1]$ is a strictly increasing function on $\operatorname{supp} F_s$ and $r(\operatorname{supp} F_s) \subseteq \operatorname{supp} F_s$. Then there exist mutually absolutely continuous measures $\hat{\mu}^L, \hat{\mu}^R \in \Delta(\mathcal{Z})$ such that the perceived posterior distribution at belief $\hat{s}$ is equal to the true posterior distribution at signal $s = \hat{s}$, $\hat{F}^{\omega}_{\hat{s}}(\hat{s}) = F^{\omega}_{s}(\hat{s})$.*

Proof. Let $\hat{F}^L_s(s) \equiv F^L_s(r(s))$. This satisfies $F^L_s(s) = \hat{F}^L_{\hat{s}}(s)$. It remains to show that $\hat{F}^R_s$ also satisfies this identity. By Lemma A.1 in Smith and Sorensen (2000)

$$F^R_s(r(s)) = \int_0^{r(s)} \frac{1-p}{p} dF^L_s(p), \tag{7}$$

and it must be that

$$\hat{F}^R_s(s) = \int_0^{s} \frac{1-r(q)}{r(q)} d\hat{F}^L_s(q).$$

Applying the change of variables formula to (7)

$$\hat{F}^R_s(s) = \int_0^{s} \frac{1-r(q)}{r(q)} d\hat{F}^L_s(q).$$

So $\hat{F}^{\omega}_s(s) = F^{\omega}_s(r(s))$ in both states. □

By Theorem 3, learning is robust if the update rule is not far away from Bayes rule (i.e. if $||r(s)-s||$ is sufficiently small), and correct learning occurs almost surely. Moreover, even with bounded signals, the arguments for Theorem 1 remain unchanged, so the set Λ can be used to characterize the set of globally stable points.

[29] When signals are bounded, there also exists a misspecified model that represents this learning rule under a slight extension to the model that allows types to have heterogeneous signal distributions. If $r : [0,1] \to [0,1]$ is a strictly increasing function on $\operatorname{supp} F_s$, but $r(\operatorname{supp} F_s) \not\subseteq \operatorname{supp} F_s$ (for example, when signals are bounded and $r(s) = s^{\nu}$), allowing for heterogeneous signal distributions will yield an analogous result. Under this extension, type θ's private signal and type misspecification encodes $(\hat{\mu}^{\theta}, \hat{\pi}^{\theta}, \mu^{\theta})$, where μ^{θ} is the true distribution of type θ's signal and $\hat{\mu}^{\theta}$ is the perceived distribution. A misspecified agent has a perceived measure over signals that is represented by $r(s)$, but believes that all other agents are type $(\mu^{\theta}, \hat{\pi}^{\theta}, \mu^{\theta})$, where μ^{θ} is the true signal distribution. In this environment, all agents are misspecified and believe all agents are interpreting information correctly. The main results of the paper easily extend to this setting.

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011): "Bayesian Learning in Social Networks," *The Review of Economic Studies*, 78, 1201–1236.

Ali, N. (2016): "Social Learning with Endogenous Information," *Mimeo*.

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

Barber, B. M. and T. Odean (2001): "Boys will be boys: Gender, overconfidence, and common stock investment," *The quarterly journal of economics*, 116, 261–292.

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.

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

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.

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

Devine, P. G. and T. M. Ostrom (1985): "Cognitive mediation of inconsistency discounting." *Journal of Personality and Social Psychology*, 49, 5.

Enke, B. and F. Zimmermann (2017): "Correlation neglect in belief formation," .

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): "Nave Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, 2, 221–43.

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

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.

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

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

Jehiel, P. (2005): "Analogy-based expectation equilibrium," *Journal of Economic Theory*, 123, 81 – 104.

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.

Koszegi, B. and M. Rabin (2006): "A model of reference-dependent preferences," *The Quarterly Journal of Economics*, 121, 1133–1165.

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.

Kübler, D. and G. Weizscker (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.

Lee, I. H. (1993): "On the Convergence of Infomational Cascades," *Journal of Economic Theory*, 61, 395–411.

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." .

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

Nickerson, R. S. (1998): "Confirmation bias: A ubiquitous phenomenon in many guises." *Review of general psychology*, 2, 175.

Odean, T. (1999): "Do Investors Trade Too Much?" *The American Economic Review*, 89, 1279–1298.

Penczynski, S. (forthcoming): "The Nature of Social Learning: Experimental Evidence," *European Economic Review*.

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.

Thomas, L. (1979): *The Medusa and the Snail: More Notes of a Biology Watcher*, Penguin Books.

Tuchman, B. (1984): *The March of Folly: From Troy to Vietnam*, Knopf.

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