Misinterpreting Others and the Fragility of Social Learning*

Mira Frick Ryota Iijima Yuhta Ishii

Abstract

We study to what extent information aggregation in social learning environments is robust to slight misperceptions of others' characteristics (e.g., tastes or risk attitudes). We consider a population of agents who obtain information about the state of the world both from initial private signals and by observing a random sample of other agents' actions over time, where agents' actions depend not only on their beliefs about the state but also on their idiosyncratic types. When agents are correct about the type distribution in the population, they learn the true state in the long run. By contrast, our first main result shows that even arbitrarily small amounts of misperception can generate extreme breakdowns of information aggregation, where in the long run all agents incorrectly assign probability 1 to some fixed state of the world, *regardless* of the true underlying state. This stark discontinuous departure from the correctly specified benchmark motivates independent analysis of information aggregation under misperception. Our second main result shows that *any* misperception of the type distribution gives rise to a specific failure of information aggregation where agents' long-run beliefs and behavior vary only coarsely with the state, and we provide systematic predictions for how the nature of misperception shapes these coarse long-run outcomes. Finally, we show that how sensitive information aggregation is to misperception depends on how rich agents' payoff-relevant uncertainty is. A design implication is that information aggregation can be improved through interventions aimed at simplifying the agents' learning environment.

*This version: December 24, 2018. Frick: Yale University (mira.frick@yale.edu); Iijima: Yale University (ryota.iijima@yale.edu); Ishii: Centro de Investigación Económica, ITAM (yuhta.ishii@itam.mx). This research was supported by National Science Foundation grant SES-1824324. We thank Larbi Alaoui, Nageeb Ali, Dirk Bergemann, Aislinn Bohren, Krishna Dasaratha, Darrell Duffie, Ignacio Esponda, Erik Eyster, Drew Fudenberg, Simone Galperti, John Geanakoplos, Ben Golub, Andrei Gomberg, Marina Halac, Kevin He, Botond Kőszegi, George Mailath, Stephen Morris, Wojciech Olszewski, Romans Pancs, Antonio Penta, Jacopo Perego, Andrea Prat, Larry Samuelson, Rani Spiegler, Philipp Strack, Bruno Strulovici, Xavier Vives, as well as audiences at Boston University, Caltech Workshop on Information and Social Economics, Columbia, ITAM, Michigan, National University of Singapore, Northwestern, Pompeu Fabra, Princeton, Tokyo, and Yale.

1 Introduction

In many economic and social settings, individuals hold limited private information about a payoff-relevant state of the world and rely on observing the behavior of others as a vital source of additional information. Typically, however, others' behavior reflects not only their own information about the state of the world, but is also influenced by their idiosyncratic characteristics. For example, in assessing the quality of a new product or a political candidate, people may draw inferences from the purchasing decisions or stated opinions of others, but these depend at least in part on others' consumption tastes or political preferences. Likewise, in many decentralized markets (such as over-the-counter markets or privately held auctions), agents learn about market fundamentals by observing other participants' trading behavior, yet the latter may also be driven by idiosyncratic features such as risk attitudes, private values or liquidity constraints.

A classic literature examines the possibility of information aggregation in such settings, deriving conditions under which individuals are able to learn the true state of the world in the long run.[1] Importantly, the standard modeling assumption is that individuals possess a *correct* understanding of their environment, in particular the distribution of relevant population characteristics. This conflicts with growing empirical evidence that people are prone to systematic *misperceptions* about such distributions; from under- or overestimating the heterogeneity of socio-political attitudes, consumption tastes or wealth levels in their societies to misjudging the share of "fake" product recommenders or political supporters on review platforms and social networking sites.[2] Such evidence has motivated a burgeoning theoretical literature to incorporate various forms of misspecification into models of social learning.[3]

At the same time, a widely held view of models is summarized by George Box's saying that "all models are wrong," but "cunningly chosen parsimonious models often do provide remarkably useful approximations" (Box, 1979). This raises the question how severe agents' misperceptions must be to motivate departing from the standard model: Does the correctly specified model perhaps offer a good enough approximation as long as the amount of misperception is sufficiently small?

The first main result in this paper suggests a negative answer to this question. We consider a population of agents who obtain information about the state of the world both from initial private signals and by observing a random sample of other agents' actions over time, where agents' actions depend not only on their beliefs about the state but also on their idiosyncratic types. When agents are correct about the type distribution in the population, they learn the true state in the long run. By contrast, we show that even arbitrarily small amounts of misperception about the type distribution can generate extreme breakdowns of information aggregation, where in the long run all agents incorrectly assign probability 1 to some fixed state of the world, *regardless* of the true underlying state.

This stark discontinuous departure motivates analyzing information aggregation under mis-

[1] For surveys, see Vives (2010); Chamley (2004).

[2] See Section 2.4 for references.

[3] We discuss this literature in Section 7.1.

perception in its own right, without extrapolating from the predictions of the correctly specified benchmark. Our second main result shows that *any* misperception about the type distribution gives rise to a specific failure of information aggregation where agents' long-run beliefs and behavior vary only coarsely with the state. Moreover, we provide systematic predictions for how the nature of misperception shapes these coarse long-run outcomes. Finally, our third main result shows that how sensitive information aggregation is to misperception depends on how rich agents' payoff-relevant uncertainty is. An important design implication is that information aggregation can be improved through interventions aimed at simplifying the agents' learning environment.

In our model, a large population of agents choose actions in each period $t \in \{1, 2, \ldots\}$ to maximize their expected utility given their time t information about a fixed but unknown state of the world $\omega \in \Omega = [\underline{\omega}, \overline{\omega}]$. Each agent's utility to a given action depends not only on the state ω, but also on his idiosyncratic type θ, where types in the population are distributed according to some cdf F. Each agent i's time t information about ω consists of two sources: First, in period 0, i observes a private signal about ω; second, in each period up to time t, i randomly meets some other agent j and observes j's action in that period. If agents are correct about the type distribution F, then under our assumptions on signals and payoffs, information aggregation is successful (Lemma 1); intuitively, by observing a sufficiently large sample of other agents' actions, all agents are able to correctly back out the state in the long run.[4]

Our focus is on the case where agents are misspecified about the type distribution, in the sense that they misperceive the true distribution F to be some other cdf $\hat{F}$. Here agents' amount of misperception can be quantified by standard notions of distance between cdfs $\hat{F}$ and F.[5] Misperceiving the type distribution entails the possibility of misinterpreting other agents, in the sense of drawing incorrect inferences about the state from their actions. Nevertheless, one might expect that as long as the amount of misperception is small, this effect should likewise be small, and information aggregation should be approximately successful. Indeed, a natural analogy is with a *single* agent who receives repeated informative signals about the state of the world, but misperceives the mapping from states to signals. When signals are exogenous, a classic result due to Berk (1966) implies that the agent's long-run belief is approximately correct when the amount of his misperception is small (Proposition 0). Moreover, the same is true in recent models of misspecified *active* learning, where signals depend endogenously on the agent's actions and the agent misperceives this dependence, e.g., due to overconfidence in his ability (Heidhues, Koszegi, and Strack, 2018).[6]

Theorem 1 offers a sharp contrast to these single-agent benchmarks. Under social learning, even vanishingly small amounts of misperception can lead to extreme breakdowns of information

[4]Our model is closely related to Duffie and Manso's (2007) random matching model of social learning without misspecification, which also features successful information aggregation. Like them, we abstract away from important forces that cause failures of information aggregation in other correctly specified models, e.g., herding or informational cascades in Bikhchandani, Hirshleifer, and Welch (1992); Banerjee (1992); Smith and Sørensen (2000). This allows us to isolate the effect of vanishingly small misperceptions, which we show can lead to more extreme failures of information aggregation than in *any* correctly specified model.

[5]We use the total variation distance, but our results go through under other standard norms.

[6]See also Nyarko (1991); Fudenberg, Romanyuk, and Strack (2017).

aggregation, where long-run beliefs are *state-independent point masses*: For any state $\hat{\omega}$, there exists a perception $\hat{F}$ that is arbitrarily close to the true type distribution F, but under which in all states ω, all agents' long-run beliefs incorrectly assign probability 1 to $\hat{\omega}$.

Section 4.2 illustrates the logic behind Theorem 1, which relies on showing that in the long run, agents' belief dynamics exhibit two key features. First, there is *mislearning*, in the sense that agents' misperception of the type distribution generates "discrepancies" between expected and observed behavior, to which agents respond by successively misadjusting their beliefs about the state and correspondingly misadjusting their behavior in the next period, generating yet another discrepancy. Second, belief dynamics become *decoupled* from the true state in the long run; this is because under social learning, agents' information over time is based increasingly on other agents' behavior, which depends on the true state only *indirectly* through other agents' beliefs. We show that the combination of these two features renders belief dynamics extremely sensitive to small misperceptions and can lead to the complete unraveling of information aggregation in Theorem 1.

Section 4.3 discusses how these two features drive the sharp contrast between misspecified social learning and the single-agent benchmarks: Mislearning is absent under single-agent passive learning where information is exogenous, but is closely related to the phenomenon of misguided learning in the single-agent active learning model of Heidhues, Koszegi, and Strack (2018). However, single-agent active learning again differs from social learning in that long-run belief dynamics do *not* become decoupled from the true state, because the agent's only source of information in each period are signals that depend *directly* on the state of the world. As a result, single-agent learning, unlike social learning, is not sensitive to small amounts of misperception. We also emphasize that belief dynamics and behavior in the proof of Theorem 1 have the feature that agents' observations (both in finite time and asymptotically) never contradict their mistaken belief that the type distribution is $\hat{F}$, suggesting a sense in which such misperceptions can be a persistent phenomenon.

Theorem 1 highlights that vanishingly small amounts of misperception *can* generate stark discontinuous departures from the correctly specified benchmark. While not every misperception $\hat{F}$ need give rise to breakdowns of information aggregation that are as extreme as in Theorem 1, this suggests that social learning under misperception must be studied independently, without relying on the predictions of the correctly specified model. Theorem 2 therefore investigates information aggregation under *arbitrary* well-behaved true and perceived type distributions F and $\hat{F}$. We show that information aggregation continues to fail, but in general long-run beliefs need not be fully state-independent and instead display the following weaker form of "coarseness:" F and $\hat{F}$ generate a partition of the state space $\Omega = [\underline{\omega}, \overline{\omega}]$ into *finitely* many intervals, and within each such interval, agents' long-run beliefs incorrectly assign probability 1 to the same fixed state. As a result, long-run behavior also varies only coarsely with the true state, remaining constant within each interval of the partition and changing discretely from one interval to the next. As we discuss, this prediction is broadly in line with the widely studied fact that behavior in many economic settings is not finely attuned to economic fundamentals, suggesting a possible new channel for this phenomenon.

Theorem 2 also provides a starting point for analyzing how long-run beliefs vary across different forms of misperception. As an illustration, Section 4.4 shows that when F and $\hat{F}$ are ranked according to first-order stochastic dominance (e.g., when agents under-/overestimate the share of "fake" recommenders) long-run beliefs exhibit drastic overoptimism/-pessimism, and that under- or overestimating population heterogeneity leads to conservative or extreme long-run beliefs, respectively.

Finally, Theorem 3 highlights a key determinant of the fragility of information aggregation, by showing that information aggregation is more sensitive to misperception the richer the state space. Suppose we approximate our continuous state space $\Omega = [\underline{\omega}, \overline{\omega}]$ by an increasingly fine sequence Ω_n of finite state spaces. Then for each n, there is some threshold ε_n such that information aggregation is successful whenever the amount of misperception is below this threshold. However, as the size of the state space grows, ε_n shrinks to 0, so that information aggregation is more and more sensitive to misperception; indeed, arbitrarily small amounts of misperception can give rise to extreme breakdowns of information analogous to Theorem 1 whenever the state space is large enough. As we discuss, this result offers a complementary perspective to Bohren and Hauser (2018), who show that a version of the sequential social learning model (Bikhchandani, Hirshleifer, and Welch, 1992; Banerjee, 1992; Smith and Sørensen, 2000) is robust to small amounts of misspecification but focus on the *binary* state space setting.

Many settings of economic interest naturally feature rich state spaces, from safety levels of new products to market fundamentals under decentralized trade. From a design perspective, Theorem 3 implies that information aggregation in such settings can be improved by simplifying the agents' learning environment: For instance, in the context of news releases by a central bank or consumer protection agency, we highlight a new trade-off between providing more information for agents to aggregate and rendering information aggregation more sensitive to misperception, and argue that this may call for releasing only "vague" information.

The paper proceeds as follows. Section 2 sets up the model. Section 3 establishes two preliminary benchmarks: Successful information aggregation under the correctly specified model and the robustness of single-agent passive learning to small amounts of missperception. Sections 4 and 5 present our main results, Theorems 1–3. Section 6 discusses more general forms of misperception, in particular the interaction between misspecified and correctly specified agents. Finally, Section 7 reviews related literature and offers some concluding remarks.

2 Model

2.1 Environment

There is a continuum of agents with mass normalized to 1. Each agent is endowed with a fixed (preference) type $\theta \in \mathbb{R}$. Each agent's type is his private information. Types in the population are distributed according to a cdf F that admits a positive and continuous density over $\mathbb{R}$. Let $\mathcal{F}$ denote the space of cdfs with these properties.

At the beginning of period 0, a state of the world ω is drawn once and for all from a cdf Ψ

that admits a positive density over a bounded interval $\Omega := [\underline{\omega}, \overline{\omega}] \subseteq \mathbb{R}$. Agents do not observe the realization of ω. At the beginning of each period $t = 1, 2, \ldots$, each agent i chooses an action $a_{it} \in \{0, 1\}$ to myopically maximize his expected utility given his period t information about ω.[7] We assume binary actions only for simplicity; as we discuss in Supplementary Appendix G, analogous insights obtain under continuous actions.

We specify information in the next subsection. Each agent's utility $u(a, \theta, \omega)$ depends on his action, his type, and the state of the world. The utility $u(0, \theta, \omega)$ to action 0 is normalized to be identically 0 for all types and states. We denote the utility $u(1, \theta, \omega)$ to action 1 by $u(\theta, \omega)$ and assume this to be strictly increasing and continuously differentiable in both θ and ω. Moreover, $\lim_{\theta \to \infty} u(\theta, \underline{\omega}) > 0$ and $\lim_{\theta \to -\infty} u(\theta, \overline{\omega}) < 0$; that is, for high (respectively, low) enough types it is always optimal to choose action 1 (respectively, action 0).[8] For each ω, let $\theta^*(\omega)$ denote the threshold type that is indifferent between both actions in state ω; this is uniquely defined by $u(\theta^*(\omega), \omega) = 0$.

2.2 Information

At the end of period 0, each agent i observes a private signal $s_i \in \mathbb{R}$ about the state of the world. Conditional on realized state ω, private signals are drawn i.i.d. across agents from cdf $\Phi(\cdot|\omega)$ with positive density $\phi(\cdot|\omega)$ over $\mathbb{R}$. Private signal distributions satisfy the monotone likelihood ratio property; that is, for each $\omega > \omega'$, $\frac{\phi(s|\omega)}{\phi(s|\omega')}$ is strictly increasing in s, so that higher signal realizations are more indicative of higher states.

At the end of each period $t = 1, 2, \ldots$, each agent i randomly meets another agent j and observes j's period t action a_{jt}.[9] We assume *independent* random matching, in the sense that j's type θ_j is drawn from the type distribution F in the population, independent of i's own type θ_i; Section 7.2 briefly discusses incorporating assortative random matching.

Thus, at the beginning of each period $t = 1, 2, \ldots$, each agent's information about the state consists of two sources: His private signal in period 0; and, if $t \geq 2$, a random sample of other agents' actions in periods $1, \ldots, t-1$. Note that agents do not observe and draw inferences from their utilities; as point (iv) in Section 2.4 discusses, two natural interpretations of this include settings where payoffs are realized only in the long run or where overlapping generations of agents take one-shot actions.

[7] Myopia is without loss in this setting as players' actions do not affect their information.

[8] Such dominant types play the following technical role. They ensure (i) that information aggregation is successful in the correctly specified model (Lemma 1) and (ii) that both actions are observed with positive probability in each period; (ii) avoids the problem of belief-updating after zero probability events, as regardless of their perceptions $\hat{F} \in \mathcal{F}$, agents never encounter observations they considered impossible. In the context of word-of-mouth learning about the quality of a new product, Example 1 interprets such types as "fake" recommenders.

[9] We assume a law of large numbers over a continuum of i.i.d. random variables (i.e., agents' observations of signals and other agents' actions). See Sun (2006), Duffie and Sun (2012) for rigorous formulations.

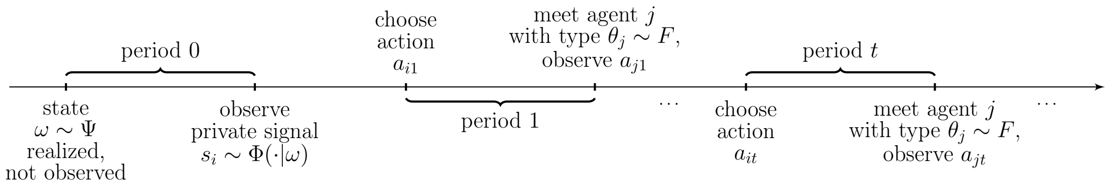

Figure 1: Timeline. At the end of each period, each agent i Bayesian-updates his belief belief about the state given the perception that the type distribution is $\hat{F}$.

2.3 Perceptions and Inferences

In drawing inferences from other agents' actions, we allow for the possibility that agents are misspecified about the type distribution F in the population. Specifically, throughout most of the analysis, we assume that there is some cdf $\hat{F} \in \mathcal{F}$ such that all agents believe the true type distribution to be $\hat{F}$ and believe that $\hat{F}$ is common knowledge. We refer to $\hat{F}$ as agents' ***perceived*** type distribution (***perception*** for short) and focus on the case of misperception, where $\hat{F} \neq F$. This parsimonious departure from the correctly specified model, where $\hat{F} = F$, is enough to convey our main insights, but Section 6 discusses more general misperceptions.

The key implication of misperception is the possibility that agents may draw incorrect inferences about the state from their observations of other agents' actions. We will be particularly interested in the case when the amount of misperception is small. To formalize this, we measure the ***amount of misperception*** by the total variation distance between $\hat{F}$ and F; that is, $\|\hat{F} - F\| := \sup_{B \in \mathcal{B}} |\int \mathbb{1}_B(\theta)\, d\hat{F}(\theta) - \int \mathbb{1}_B(\theta)\, dF(\theta)|$, where $\mathcal{B}$ denotes the Borel σ-algebra on $\mathbb{R}$. Our results go through under other standard norms.[10]

Aside from their potential misperception of the type distribution, agents' inferences are standard. In particular, both the distribution of states Ψ and the private signal distributions $\Phi(\cdot|\omega)$ are common knowledge among agents.[11] Moreover, given perception $\hat{F}$, agents draw inferences from information in each period by Bayesian updating. Figure 1 summarizes the timeline of the model.

In any state ω and at the beginning of each period $t \geq 1$, let $\mu_t^\omega \in \Delta(\Delta(\Omega))$ denote the population distribution over agents' posterior beliefs about the state. To study information aggregation, we consider the distribution of long-run beliefs; that is, the limit $\mu_\infty^\omega := \lim_{t\to\infty} \mu_t^\omega$ with respect to the topology of weak convergence. Whenever μ_∞^ω exists and assigns probability 1 to a Dirac measure $\delta_{\omega'}$ on some state ω', we say that in state ω almost all agents' beliefs converge to a point mass on ω'; if $\omega' = \omega$, then information aggregation in state ω is successful.

[10]See footnote 22.

[11]Incorporating misperceptions about Ψ and/or $\Phi(\cdot|\omega)$ in *addition* to misperceptions about F does not affect our main results. However, if agents are correct about the type distribution F and *only* hold misperceptions about Ψ and/or $\Phi(\cdot|\omega)$, then information aggregation is approximately successful under small enough amounts of misperception. See also footnote 23.

2.4 Examples

The above framework captures numerous economic and social situations:[12]

(i) *Learning from others' behavior.* In assessing the long-term health effects ω of potentially risky behaviors (e.g., recreational drug use) or new products (e.g., GMO foods), individuals may possess only a "fuzzy" understanding $s \sim \Phi(\cdot|\omega)$ of existing research on the subject and obtain additional information by observing other agents' day-to-day behavior and consumption choices a_t.

(ii) *Word-of-mouth communication.* Based on a political candidate's campaign announcement speech or a promotional trailer for an upcoming movie or music album, people may form their own opinions about the expected quality of the candidate or product, but update these opinions after hearing others' assessments.

(iii) *Decentralized markets.* Duffie and Manso (2007) propose a related framework (without type heterogeneity and misperceptions thereof) to capture decentralized markets (e.g., the markets for real estate or over-the-counter securities), where agents who are uncertain about market fundamentals may randomly encounter other participants (e.g., at privately held auctions) and gather additional information by observing their trading behavior (e.g., their bids).

(iv) *Overlapping generations.* An implicit assumption in the previous examples is that agents take actions repeatedly, but do not observe their payoffs to these actions. This fits many settings where states affect payoffs only in the long run (e.g., long-term health effects in (i), the quality of a political candidate once in office in (ii), or an asset conditioned on a distant future event in (iii)). Similar to the literature on sequential social learning (e.g., Bikhchandani, Hirshleifer, and Welch, 1992; Banerjee, 1992; Smith and Sørensen, 2000), our model also fits settings where overlapping generations of agents take one-shot actions whose payoffs they observe privately, and subsequent generations of agents observe a random sample of previous agents' actions.

In each of the above settings, other agents' behavior is influenced not only by their own information, but also by their heterogeneous characteristics θ (e.g., consumption tastes, socio-political preferences, risk attitudes, liquidity constraints). Moreover, as highlighted in the introduction, a growing empirical literature suggests that agents are prone to systematically misperceive the distributions of these characteristics, from under- or overestimating the heterogeneity of socio-political attitudes, consumption tastes or wealth levels in their societies to misjudging the share of "fake" product recommenders or political supporters on review platforms and social networking sites.[13]

[12] In several of these examples, actions are better modeled as continuous rather than binary; as mentioned, we show in Supplementary Appendix G that our analysis generalizes to this setting.

[13] See, e.g., Hauser and Norton (2017), Norton and Ariely (2011) for evidence of under-/overestimation of wealth inequality; Ahler (2014) for overestimation of political attitude polarization; Kunda (1999), Nisbett and Kunda (1985) for misperceptions of numerous taste and attitude distributions in society; and Mayzlin, Dover, and Chevalier (2014) and the references therein for the difficulties of detecting fake reviews.

3 Preliminary Benchmarks

Before turning to analyze information aggregation under misperception about the type distribution in Section 4, we present two preliminary benchmarks that will serve as a helpful contrast to our main results. In Section 3.1, we show that when agents' perception of the type distribution is correct, information aggregation is successful in all states. In Section 3.2, we consider a single agent who observes an exogenous sequence of random actions and misperceives the distribution of actions in each state. We show that in this case, the agent's long-run beliefs are approximately correct as long as the amount of misperception is sufficiently small.

3.1 Information Aggregation under Correct Perceptions

We first show that when there is common knowledge of the correct type distribution F, agents learn the true state in the long run. We will invoke this result in analyzing the case with misperception in Section 4, where we will obtain starkly different conclusions. At the same time, this result highlights the fact that our model does not feature the possibility of herding or related failures of information aggregation that can arise even under correct perceptions. As such, it serves to isolate misperception as the sole source of the breakdown of information aggregation that we will study in Section 4.

Lemma 1 (Information aggregation under correct perceptions)**.** *Suppose that $\hat{F} = F$. Then in any state ω, almost all agents' beliefs converge to a point mass on ω.*

We prove Lemma 1 in Appendix A. Letting $q_t(\omega)$ denote the fraction of agents that take action 0 in state ω and period t, it is sufficient to prove that $\lim_{t\to\infty} q_t(\omega)$ exists and is strictly decreasing in ω. One complication is that agents' action observations, and hence their beliefs about the state, are private, so that calculating $q_t(\omega)$ requires keeping track of the population distribution of agents' beliefs $\mu_t^\omega \in \Delta(\Delta(\Omega))$, which does not admit a tractable expression.[14]

Instead, we first use an inductive argument to show that $q_t(\omega)$ is strictly decreasing in ω for each t. The intuition is quite simple: First, when the realized state ω is low, more agents observe lower private signals in period 0, and consequently, more agents choose action 0 in period 1. As a result, more agents observe action 0 at the end of period 1, and given the first step, action 0 is more indicative of low states than action 1. This in turn leads to fewer agents choosing action 0 in period 2, and so on. To use this to establish that $\lim_{t\to\infty} q_t(\omega)$ exists and is strictly decreasing in ω, we must additionally rule out the possibility that $q_t(\omega)$ becomes very flat in ω in the limit, in which case some states might yield almost the same action frequencies and be impossible for agents to distinguish in the limit. In Appendix A, we establish this through an analysis of the asymptotic properties of the belief distribution μ_t that is based on martingale convergence arguments and the richness types in the population.[15]

[14]In Duffie and Manso's (2007), Duffie, Malamud, and Manso's (2009), and Duffie, Giroux, and Manso's (2010) related models of learning in decentralized markets (see Section 2.4), the population distribution of posteriors can be calculated explicitly. The methods rely on these papers' specific assumptions about state distributions (binary or Gaussian) and homogeneous preferences, and thus do not apply to our setting.

[15]Relatedly, Goeree, Palfrey, and Rogers (2006) establish successful information aggregation in a sequential social

3.2 Single Agent Benchmark with Misperception

In the previous subsection, agents' ability to draw correct inferences about the state from observed actions relied on the fact that they knew the true type distribution. When agents misperceive the type distribution to be $\hat{F} \neq F$, this introduces the possibility of *misinterpreting* observed actions, in the sense that agents might have in mind an incorrect mapping from states to probabilities of observing actions 0 or 1 at any point in time.

Nevertheless, one intuition one might have is that as long as the amount of misperception is small, the effect of such misinterpretation will likewise be small, and agents will "approximately" learn the true state in the long run. Our main results in Section 4 will show that this intuition is not valid. However, to better understand the logic behind these results, it will be helpful to contrast them with the following benchmark, where a small amount of misinterpretation of observed actions does lead to approximately correct long-run beliefs.

Specifically, consider a *single* agent who observes an *exogenous* sequence of binary random variables ("actions") $a_t \in \{0, 1\}$ in all periods $t = 1, 2, \ldots$. Unlike in our original model, where observed actions result from utility-maximizing behavior by other agents, we assume that conditional on realized state ω, a_t is distributed i.i.d. over time: a_t takes value 0 with probability $q(\omega)$ and value 1 with complementary probability, where the mapping $q : \Omega \to (0, 1)$ from states to probabilities of observing action 0 is continuous and strictly decreasing. To capture misinterpretation of observed actions, we consider the possibility that the agent misperceives the mapping q to be $\hat{q} : \Omega \to (0, 1)$, where $\hat{q}$ is again continuous and strictly decreasing.[16]

A classic result due to Berk (1966) (see also Esponda and Pouzo, 2016) characterizes the agent's long-run beliefs in this case: Define the ***Kullback-Leibler (KL) divergence*** between probabilities $p, \hat{p} \in (0, 1)$ to be $\mathrm{KL}(p, \hat{p}) := p \log(\frac{p}{\hat{p}}) + (1 - p) \log(\frac{1-p}{1-\hat{p}})$. Then in any state ω, the agent's long-run belief assigns probability 1 to the state

$$\hat{\omega}(\omega) := \operatorname*{argmin}_{\hat{\omega} \in \Omega} \mathrm{KL}(q(\omega), \hat{q}(\hat{\omega}))$$

that minimizes KL divergence between the true action 0 frequency $q(\omega)$ and the agent's perceived frequency $\hat{q}(\hat{\omega})$. Note that $\hat{\omega}(\omega)$ exists and is unique for each ω, as q and $\hat{q}$ are continuous and strictly decreasing.

An immediate implication is that when the amount of misperception is small, the agent's long-run belief is approximately correct, in the sense that the perceived state $\hat{\omega}(\omega)$ is approximately equal to the true state ω:[17]

learning model with rich types.

[16] In the present model, the agent's learning is passive. In Section 4.3, we discuss models of misspecified single-agent *active* learning and illustrate why long-run beliefs are likewise approximately correct when the amount of misspecification is small.

[17] A similar continuity result holds in the following setting. Suppose a population of agents acquire information and take actions to maximize utility as in Section 2, but only a single agent perceives the type distribution to be $\hat{F}$ (and believes that $\hat{F}$ is common knowledge), while all other agents have common knowledge of the correct type distribution F. Here the action 0 frequency $q_t(\omega)$ in the population is not i.i.d. over time, reflecting correctly specified

Proposition 0. *For any continuous and strictly decreasing* $q, \hat{q} : \Omega \to (0,1)$*, the agent's belief in any state* ω *converges almost surely to a point mass on* $\hat{\omega}(\omega) := \text{argmin}_{\hat{\omega}\in\Omega} \text{KL}(q(\omega), \hat{q}(\hat{\omega}))$*, which is strictly increasing and continuous in* ω*. Moreover, for any* $\delta > 0$*, there exists* $\varepsilon > 0$ *such that if* $\sup_{\omega\in\Omega} |\hat{q}(\omega) - q(\omega)| < \varepsilon$*, then* $\sup_{\omega\in\Omega} |\hat{\omega}(\omega) - \omega| < \delta$*.*

4 Failure of Information Aggregation under Misperception

We now return to analyzing the effect of misperception about the type distribution in the population setting of Section 2. Our main results contrast sharply with the previous two benchmarks.

4.1 Main Results

Our first main result finds that information aggregation is highly non-robust to small amounts of misperception. Whereas Lemma 1 established that under correct perceptions, agents eventually learn the true state, we now show that even arbitrarily small amounts of misperception can lead information aggregation to break down.

The breakdown we derive is very stark: Given any type distribution F, we can find an arbitrarily small amount of misperception under which agents' long-run beliefs are *state-independent point masses*, assigning probability 1 to some fixed state $\hat{\omega}$ regardless of the true state ω. Moreover, long-run beliefs are *arbitrary*, in the sense that *any* state $\hat{\omega}$ can arise as the long-run point-mass belief under some arbitrarily small amount of misperception.

Theorem 1 (Discontinuous breakdown of information aggregation)**.** *Fix any* $F \in \mathcal{F}$ *and* $\hat{\omega} \in \Omega$*. For any* $\varepsilon > 0$*, there exists a perception* $\hat{F} \in \mathcal{F}$ *with* $\|\hat{F} - F\| < \varepsilon$ *under which in any state* ω*, almost all agents' beliefs converge to a point mass on* $\hat{\omega}$*.*

Several prominent social learning models can give rise to unsuccessful information aggregation even when agents are correctly specified; e.g., due to the possibility of herding and/or confounded learning in sequential learning models à la Bikhchandani, Hirshleifer, and Welch (1992), Banerjee (1992), and Smith and Sørensen (2000).[18] However, Theorem 1 generates a more extreme breakdown of information aggregation—long-run beliefs that are state-independent point-masses—that cannot arise under *any* correctly specified Bayesian learning model, because beliefs in such models follow a martingale.

Related instances where agents' long-run beliefs grow confident in an incorrect state have been derived in several recent papers (e.g., Eyster and Rabin, 2010; Bohren, 2016; Gagnon-Bartsch and Rabin, 2017) that incorporate various forms of misspecification into sequential learning models. The

agents' learning. However, since the correctly specified agents learn the true state in the long-run (by Lemma 1), $q_t(\omega)$ converges to $q(\omega) := F(\theta^*(\omega))$ over time. Given this, it can be shown that the single misspecified agent's belief converges to a point mass on $\hat{\omega}(\omega) = \text{argmin}_{\hat{\omega}} \text{KL}(q(\omega), \hat{F}(\theta^*(\hat{\omega})))$, where $\hat{\omega}(\omega)$ is again close to ω when the amount of misperception is small. By contrast, Section 6 considers hybrid populations where some *positive fraction* of agents is misspecified and shows that interactions between correct and misspecified agents are more subtle in this case.

[18] In a random matching setting, see, e.g., Banerjee and Fudenberg (2004), Wolinsky (1990), Blouin and Serrano (2001) for correctly specified models where information aggregation can fail.

key novelty is that Theorem 1 shows that state-independent point-mass beliefs can arise even under *vanishingly small* amounts of misperception. In contrast, the aforementioned papers rely on strong forms of misspecification;[19] indeed, in a general model of misspecified sequential learning that nests several of these types of misspecification, Bohren and Hauser (2018) show that agents learn the true state whenever the amount of misspecification is sufficiently small. One important difference between these models and ours is that they assume a binary state space. In Section 5, we will show that how rich a space of uncertainty agents face has a key impact on how sensitive information aggregation is to small amounts of misperception in our setting.

The fact that an arbitrarily small amount of misperception suffices to bring about this breakdown is also in marked contrast to the single agent benchmark in Section 3.2, where we saw that the agent's long-run beliefs are approximately correct when the amount of misperception is sufficiently small. While Section 3.2 is a model of passive learning, several papers explore single-agent *active* learning and likewise differ from our environment in that learning is robust to small amounts of misperception (Nyarko, 1991; Fudenberg, Romanyuk, and Strack, 2017; Heidhues, Koszegi, and Strack, 2018). Section 4.3 will identify the source of this important difference between social learning and (passive or active) single-agent learning.

As the proof sketch in Section 4.2 will illustrate, for a given type distribution F, the misperceptions $\hat{F}$ that give rise to the extreme breakdown in Theorem 1 are quite specific, though Section 4.4 provides natural examples of such misperceptions (e.g., underestimation of type heterogeneity). Of course, Theorem 1 does not suggest that every misperception leads to state-independent point-mass beliefs. Rather, the key implication is that the correctly specified model need not offer a good approximation of a setting where agents hold even slightly incorrect beliefs about others' characteristics. This suggests that information aggregation under misperception should be studied independently, without relying on the predictions of the correctly specified model.

As a step in this direction, our second main result therefore investigates the effect of *arbitrary* misperceptions $\hat{F}$. While in general long-run beliefs need not be fully state-independent, we show that information aggregation continues to fail, and the failure takes the specific form of "coarse" long-run beliefs and behavior. The result focuses on well-behaved true and perceived type distributions; specifically, we assume that F and $\hat{F}$ are analytic.[20]

Theorem 2 (Coarse information aggregation)**.** *Fix any analytic F, $\hat{F} \in \mathcal{F}$ with $\hat{F} \neq F$. There exists a mapping $\hat{\omega}_\infty : \Omega \to \Omega$ that is weakly increasing and has finite range such that in any state ω, almost all agents' beliefs converge to a point mass on $\hat{\omega}_\infty(\omega)$.*

Information aggregation in Theorem 2 is coarse in the following sense: Since the mapping $\hat{\omega}_\infty$

[19] E.g., in Eyster and Rabin (2010) and Gagnon-Bartsch and Rabin (2017), agents naively believe that each predecessor's action reflects solely that person's private information, *fully* neglecting the fact that predecessors' behavior also reflects their inferences from their own predecessors' behavior.

[20] Function $g : \mathbb{R} \to \mathbb{R}$ is ***analytic*** if it is locally given by a convergent power series; that is, for any $x_0 \in \mathbb{R}$, there is a neighborhood J of x_0 and a sequence of real coefficients $(\alpha_n)_{n=0}^{\infty}$ such that for all $x \in J$, $g(x) = \sum_{n=0}^{\infty} \alpha_n (x - x_0)^n$ and the right-hand side converges. As Section 4.2 illustrates, the only feature of analyticity that Theorem 2 exploits is that analytic $\hat{F} \neq F$ can intersect at most finitely many times on any compact interval; the theorem remains valid even if $\hat{F}, F$ are not analytic but have this feature.

from true states ω to long-run point-mass beliefs $\hat{\omega}_\infty(\omega)$ is weakly increasing and has finite range, it partitions the continuous state space $\Omega = [\underline{\omega}, \overline{\omega}]$ into finitely many intervals, and long-run point-mass beliefs are not necessarily fully state-independent, but are constant within each of these finitely many intervals. This prediction again contrasts with Proposition 0 from the single agent benchmark, where the agent's long-run belief $\hat{\omega}(\omega)$ is a strictly increasing and continuous function of the true state.

As a result of agents' coarse long-run beliefs, their long-run behavior also varies only coarsely with the true state, remaining constant within each interval of the partition generated by $\hat{\omega}_\infty$ and changing discretely from one interval to the next. This prediction is broadly in line with the fact that behavior in many economic settings (e.g., firms' pricing behavior and individuals' consumption-savings decisions) is not finely attuned to economic fundamentals.[21] A rich theory literature provides models of such coarse behavior that are based on the idea that individuals face limitations in their ability to process or acquire information (e.g., Sims, 1998, 2003; Mullainathan, 2002; Jehiel, 2005; Fryer and Jackson, 2008; Gul, Pesendorfer, and Strzalecki, 2017). Theorem 2 highlights a possible complementary channel: Agents in our model do not face any difficulties processing their private information, but coarse behavior emerges because agents' misperceptions of others' characteristics give rise to coarse *aggregation* of this dispersed individual information.

The proofs of Theorems 1 and 2 appear in Appendix B. In the next subsection, we illustrate the basic argument, which relies on a mechanism that we term "decoupled mislearning." As we will see, this mechanism is key in understanding the starkly different effect that small amounts of misperception can have relative to single-agent active or passive learning models.

4.2 Illustration of Theorems 1 and 2

Step 1: Limit model. To illustrate the key ideas, Steps 1 and 2 first present and analyze a heuristic ***limit model***, to which we will refer back throughout the paper. In Step 3 below, we will show that this limit model approximates agents' long-run beliefs and behavior in the original model, allowing us to translate the conclusions obtained in the limit model back into the original model.

We first consider an arbitrary perception $\hat{F}$. The limit model differs from the original model solely in assuming that at the end of each period $t \geq 1$, each agent meets not one, but *infinitely* many other agents and observes their period t actions. By an exact law of large numbers, this means that in any state ω all agents perfectly learn the fraction $q_t(\omega)$ of actions 0 that is played in the population in period t.

In period 1, all agents play the action that maximizes their expected utility given their period 0 private signal. Because of this, if perception $\hat{F} = F$ is correct, then by observing $q_1(\omega)$ at the end of period 1, all agents can correctly back out the true state ω. From period 2 on, agents then follow a threshold strategy with threshold $\theta^*(\omega)$, where types above (respectively, below) $\theta^*(\omega)$ play action 1 (respectively, action 0). Observing this behavior in each period is *consistent* with agents' belief that the state is ω.

[21] See, e.g., Reis (2006a,b) and references therein.

Suppose next that $\hat{F} \neq F$. Since all agents believe $\hat{F}$ to be the true type distribution (and believe this to be common knowledge), observing $q_1(\omega)$ at the end of period 1 leads agents to commonly believe in some state $\hat{\omega}_1(\omega)$. Since $\hat{F} \neq F$, $\hat{\omega}_1(\omega)$ need not equal ω, though by a similar reasoning as in Section 3.2, $|\hat{\omega}_1(\omega) - \omega|$ is negligible when $\|\hat{F} - F\|$ is sufficiently small. Given their belief in $\hat{\omega}_1$, in period 2, agents then follow a threshold strategy with threshold $\theta_1^* := \theta^*(\hat{\omega}_1)$.

A key departure from the correct perceptions case arises at the end of period 2. This takes the form of a possible *inconsistency* between expected behavior and actual observations: Given perception $\hat{F}$, all agents assign probability 1 to observing fraction $\hat{F}(\theta_1^*)$ of actions 0 at the end of period 2. However, since the true type distribution is $F \neq \hat{F}$, agents' actual observation at the end of period 2 is $q_2(\omega) = F(\theta_1^*)$, which typically does not equal $\hat{F}(\theta_1^*)$.

The limit model postulates that agents react to such possible "contradictions" as follows: Upon observing fraction $q_t(\omega)$ of action 0 at the end of any period $t \geq 2$, all agents update beliefs to assign probability 1 to the state $\hat{\omega}_t(\omega) := \operatorname{argmin}_{\hat{\omega} \in \Omega} \mathrm{KL}\left(q_t(\omega), \hat{F}(\theta^*(\hat{\omega}))\right)$. State $\hat{\omega}_t$ is chosen to best explain observation q_t under perception $\hat{F}$, in the sense of minimizing KL-divergence between q_t and the fraction $\hat{F}(\theta^*(\hat{\omega}_t))$ of actions 0 that agents would have expected to observe if $\hat{\omega}_t$ had been common knowledge in period t.

Given the updated belief that the state is $\hat{\omega}_t$, in period $t+1$, agents then follow the threshold strategy with cutoff $\theta_t^* := \theta^*(\hat{\omega}_t)$ and at the end of the period again face a possible discrepancy between expected behavior $\hat{F}(\theta_t^*)$ and actual behavior $F(\theta_t^*)$. Starting with $\hat{\omega}_1$ as derived above, we thus obtain a process of point mass beliefs $\hat{\omega}_t$ in all periods $t \geq 2$ given by

$$\hat{\omega}_t = \operatorname*{argmin}_{\hat{\omega} \in \Omega} \mathrm{KL}\left(F(\theta_{t-1}^*), \hat{F}(\theta^*(\hat{\omega}))\right) \text{ with } \theta_{t-1}^* = \theta^*(\hat{\omega}_{t-1}). \tag{1}$$

Of course, the belief adjustment process (1) is heuristic, as it involves switching from assigning probability 1 to state $\hat{\omega}_{t-1}$ at the beginning of period t to assigning probability 1 to the possibly different state $\hat{\omega}_t$ at the end of period t. However, as we shall see in Step 3, in the long run this adjustment process approximates the Bayesian belief dynamics in the original model: For large t, agents' beliefs in the original model are "close" to point-mass beliefs on $\hat{\omega}_{t-1}$ and $\hat{\omega}_t$ at the beginning and end of period t, but beliefs retain full support throughout, so that belief-updating is well-defined.

Step 2: Long-run beliefs. Turning to long-run beliefs in the limit model, we now show how process (1) leads learning about the state to gradually unravel. To illustrate Theorem 1, we first consider a particular choice of $\hat{F}$ and F, where F is arbitrary and $\hat{F}$ crosses F from below at a single point $\theta^* = \theta^*(\hat{\omega})$ with $\hat{\omega} \in \Omega$, as shown in the left-hand panel of Figure 2. Example 2 in Section 4.4 provides an interpretation in terms of underestimation of population heterogeneity. As explained in Figure 2, starting at any state $\hat{\omega}_1$, the $\hat{\omega}_t$-process in (1) must converge to the limit belief $\hat{\omega}$ that corresponds to the crossing point of $\hat{F}$ and F. Thus, even though, as noted in Step 1 above, the period 1 belief $\hat{\omega}_1$ depends on the true state ω (and indeed is arbitrarily close to ω when the amount of misperception is small), agents' long-run belief assigns probability 1 to the *same* fixed state $\hat{\omega}$ *regardless* of the true state ω. Moreover, observe that by suitably choosing $\hat{F}$, both the

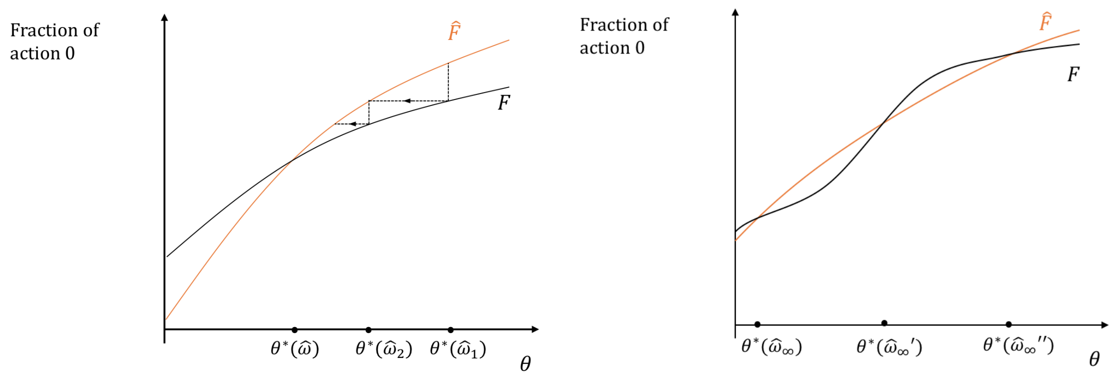

Figure 2: **Left:** Starting with any point-mass belief $\hat{\omega}_1$, agents face a discrepancy between expecting to observe fraction $\hat{F}(\theta^*(\hat{\omega}_1))$ of actions 0 and actually observing $F(\theta^*(\hat{\omega}_1))$. In response, they adjust beliefs to a point mass on $\hat{\omega}_2$, which perfectly explains the latter observation. But following this adjustment, next-period behavior is governed by threshold $\theta^*(\hat{\omega}_2)$, giving rise to another discrepancy between expected and actual behavior $\hat{F}(\theta^*(\hat{\omega}_2))$ and $F(\theta^*(\hat{\omega}_2))$. The process of adjustments continues, converging to the unique belief $\hat{\omega}$ under which expected and actual behavior coincide. **Right:** An example with three steady states $\hat{\omega}_\infty$, $\hat{\omega}'_\infty$, $\hat{\omega}''_\infty$.

amount of misperception $\|\hat{F} - F\|$ and the long-run belief $\hat{\omega} \in \Omega$ can be arbitrary. Together with the justification of the limit model in Step 3, these observations will establish Theorem 1.[22]

We refer to (1) as a process of ***decoupled mislearning***. "Mislearning" refers to the fact that agents' misperception of the type distribution leads them to repeatedly misinterpret others' actions, responding to the discrepancy between expected and observed behavior by misadjusting their beliefs about the state and correspondingly misadjusting their behavior in the next period. Except for boundary cases, the adjusted belief $\hat{\omega}_t$ satisfies $\hat{F}(\theta^*(\hat{\omega}_t)) = F(\theta^*_{t-1})$, so that beliefs at the end of each period t *perfectly* explain the behavior that was observed in the current period. However, each belief adjustment is followed by a corresponding adjustment in *next* period's behavior, which leads to a new discrepancy between expected and actual behavior, triggering yet another belief adjustment. The latter point contrasts with the single agent passive learning benchmark in Section 3.2, where observed action frequencies $q_t(\omega)$ in each period were exogenous and hence were unaffected by changes in the belief about the state.[23]

"Decoupled" refers to a second key feature of our setting: Under social learning, agents' belief-updating over time is based increasingly on other agents' behavior, which depends on the true state only *indirectly* through other agents' beliefs. In the limit model, this takes the especially extreme form that process (1) is decoupled from the true state ω, in the sense that adjustments from $\hat{\omega}_{t-1}$ to $\hat{\omega}_t$ in all periods $t \geq 2$ are completely independent of the realized ω. In Section 4.3,

[22]From this, it is clear that Theorem 1 does not rely on the use of the total variation distance. It remains valid under any norm on $\mathcal{F}$ with the feature that for any $\hat{\omega}$, there are perceptions $\hat{F}$ that are arbitrarily close to F but cross F only once from below at $\theta^*(\hat{\omega})$. Such norms include the sup norm, all L^p norms, the C^1 norm ($\|\hat{F} - F\|_{C^1} := \sup_{\theta \in \mathbb{R}} |\hat{F}(\theta) - F(\theta)| + \sup_{\theta \in \mathbb{R}} |\hat{F}'(\theta) - F'(\theta)|$), etc.

[23]Another contrast is with the case where agents are correct about the type distribution, but hold misperceptions about the distribution Ψ of states or $\Phi(\cdot|\omega)$ of private signals. Such misperceptions only affect agents' inferences at the end of period 1; in all subsequent periods, agents face no discrepancy between actual and expected behavior and do not further adjust beliefs, so that small amounts of misperception have a negligible effect.

we will show that this feature distinguishes social learning from single agent *active* learning models with misspecification (e.g., Heidhues, Koszegi, and Strack, 2018). In the latter, the fact that the agent's actions endogenously affect his information in each period can lead to a similar process of mislearning as above. But the agent's belief-updating in each period remains *directly* tied to the state, and in contrast with the unraveling of beliefs we saw above, this ensures that learning under small amounts of misperception remains approximately successful.

To illustrate Theorem 2, we next consider long-run beliefs for arbitrary F and $\hat{F}$. Let $\mathrm{SS}(F, \hat{F})$ denote the set of steady states of process (1).[24] It is easy to show that any steady state $\hat{\omega}_\infty \in \mathrm{SS}(F, \hat{F})$ satisfies either $F(\theta^*(\hat{\omega}_\infty)) = \hat{F}(\theta^*(\hat{\omega}_\infty))$ or $\hat{\omega}_\infty \in \{\underline{\omega}, \overline{\omega}\}$; thus, steady states either feature no discrepancy between the true and perceived fraction of actions 0 or are boundary points of the state space. Moreover, based on the observation that $\hat{\omega}_{t+1}$ is increasing in $\hat{\omega}_t$ for all t, we can show that in any state ω, the $\hat{\omega}_t$-process converges to some $\hat{\omega}_\infty(\omega) \in \mathrm{SS}(F, \hat{F})$, where $\hat{\omega}_\infty(\omega)$ is weakly increasing in ω.

Observe that when F and $\hat{F}$ are analytic with $\hat{F} \neq F$, then F and $\hat{F}$ coincide in at most finitely many points on the compact interval $[\theta^*(\overline{\omega}), \theta^*(\underline{\omega})]$.[25] As a result, $\mathrm{SS}(F, \hat{F})$ is finite; the right-hand panel of Figure 2 provides an example in which there are three steady states. Hence, the previous paragraph yields a weakly increasing and finite-ranged map $\hat{\omega}_\infty : \Omega \to \Omega$ from realized states ω to limit beliefs $\hat{\omega}_\infty(\omega)$. Together with Step 3, this will establish Theorem 2.

Step 3: Justifying the limit model. Finally, we return to the original model and sketch why in the long run, belief updating and behavior are approximated by the limit model. In the original model, each agent's observations up to period $t+1$ consist of a random sample $(a_1, \ldots, a_t)$ of other agents' actions in periods 1 through t. Belief updating in this model is more complicated than adjustment process (1) for two main reasons: First, agents' inference problem is not time-stationary; second, due to sampling noise, observations (and hence beliefs) differ across agents.

In Appendix B, we overcome both complications by considering the empirical frequency $\bar{a}_t := \frac{1}{t}\sum_{\tau=1}^{t} a_\tau$ of actions that each agent observes. First, since each agent believes that his perception $\hat{F}$ is correct and is shared by everyone, he believes (by Lemma 1) that the population learns the true state in the long run and that behavior converges to the corresponding threshold strategy. Based on this, Lemma B.2 shows that for large enough t, $\bar{a}_t$ provides an approximate sufficient statistic for each agent's inferences; moreover, inferences are approximately time-stationary, in the sense that each agent's time t posterior, while having full support, is close to a point-mass belief on the state $\operatorname{argmin}_{\hat{\omega}} \mathrm{KL}(1 - \bar{a}_t, \hat{F}(\theta^*(\hat{\omega}))$ that best explains $\bar{a}_t$ under perception $\hat{F}$. Second, most agents' observed empirical frequencies $\bar{a}_t$, and hence their beliefs, are very similar in the long run. Specifically, let $q_\tau(\omega)$ denote the true action 0 share in the population at time τ. Then, based on a law of large numbers argument, Lemma B.3 shows that in the long run, an arbitrarily large fraction of agents observes $1-\bar{a}_t$ that is arbitrarily close to the true time average share $\bar{q}_t(\omega) := \frac{1}{t}\sum_{\tau=1}^{t} q_\tau(\omega)$ of action 0, where $\bar{q}_t(\omega)$ converges.[26]

[24] That is, $\mathrm{SS}(F, \hat{F}) := \{\hat{\omega}_\infty \in \Omega : \hat{\omega}_\infty = \operatorname{argmin}_{\hat{\omega} \in \Omega} \mathrm{KL}\left(F(\theta^*(\hat{\omega}_\infty)), \hat{F}(\theta^*(\hat{\omega}))\right)\}$.

[25] This follows from the principle of permanence for analytic functions; see footnote 63.

[26] Here the law of large numbers applies since conditional on each state ω, each agent's action observations

Combining these two observations, Proposition B.1 shows that most agents' long-run belief updating is approximated by the following sequence of commonly held point-mass beliefs $\hat{\omega}_t$ and behavior in the population is close to the corresponding threshold strategy:

$$q_{t+1}(\omega) \approx F(\theta^*(\hat{\omega}_t)) \text{ with } \hat{\omega}_t = \underset{\hat{\omega}\in\Omega}{\operatorname{argmin}} \operatorname{KL}\left(\bar{q}_t(\omega), \hat{F}(\theta^*(\hat{\omega}))\right),$$

where the approximation "$\approx$" becomes arbitrarily precise as $t \to \infty$. Since $\bar{q}_t(\omega)$ converges and hence is close to $q_t(\omega)$ in the long run, this yields

$$\hat{\omega}_t \approx \underset{\hat{\omega}\in\Omega}{\operatorname{argmin}} \operatorname{KL}\left(F(\theta^*(\hat{\omega}_{t-1})), \hat{F}(\theta^*(\hat{\omega}))\right),$$

i.e., an approximate version of the updating process (1) in the limit model. In particular, long-run beliefs in the original model correspond precisely to steady states $\hat{\omega}_\infty(\omega) \in \operatorname{SS}(F, \hat{F})$ of (1).

4.3 Discussion

Persistence of misperceptions. In contrast with the heuristic updating process in the limit model, agents in our model are Bayesian and assign positive probability to every finite action sequence, so they never encounter any "contradictory" information in finite time. Nevertheless, one might wonder whether in the limit as they accumulate infinitely many observations, agents should realize that their perception is incorrect. There is a sense in which this is not the case. This is because (except for boundary cases) almost all agents' observed empirical action frequencies $\lim_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} a_\tau$ converge *precisely* to the prediction $1 - \hat{F}(\theta^*(\hat{\omega}_\infty))$ under their limit belief.[27] This suggests that misperceptions in this environment can be a relatively persistent phenomenon.

Social learning vs. single-agent active learning. Theorem 1 highlights a fundamental distinction between social and single-agent learning: Under the former, but not the latter, small amounts of misperception can have a large negative effect in the long run. Section 3.2 discussed the single-agent *passive* learning benchmark. In addition, several important papers (Nyarko, 1991; Fudenberg, Romanyuk, and Strack, 2017; Heidhues, Koszegi, and Strack, 2018) study the effect of misspecification in single-agent environments with *active* learning, where the agent's actions influence the distribution of signals he observes in each period. Similar to the mislearning process we highlighted in Section 4.2, Heidhues, Koszegi, and Strack (2018) emphasize that this endogeneity of the agent's information can render his belief more and more incorrect over time. However, all aforementioned models again differ from ours in that long-run beliefs under small enough amounts of misspecification are approximately correct.

In Section 4.2, we attributed this difference to the fact that belief updating in our model, but not under active learning, becomes *decoupled* from the true state over time. To see this, consider a single agent who chooses actions from $A = [0, 1]$ in each period to maximize his state-dependent

$(a_1, \ldots, a_t)$ are independently (although not identically) distributed over time.

[27] This observation is an analog of Proposition 6 in Heidhues, Koszegi, and Strack (2018).

utility $u(a,\omega)$. Assume for simplicity that $a^*(\omega) := \operatorname{argmax}_a u(a,\omega)$ is unique for each $\omega \in \Omega$. At the end of each period $t \geq 1$, the agent receives a signal $s_t \in \{0,1\}$ about the state, where the probability $q(a_t,\omega)$ of receiving signal 0 depends both on ω and his period-t action a_t.[28] Assume that the agent perceives mapping $q : A \times \Omega \to (0,1)$ to be $\hat{q} : A \times \Omega \to (0,1)$, where q and $\hat{q}$ are continuous and strictly decreasing in both arguments and are surjective in ω for each a. Analogous to Section 4.2, we consider a limit model where the agent observes not one, but infinitely many draws of signals each period. In period 1 the agent chooses action a_1^* based on his prior and arrives at a point-mass belief in the unique state $\hat{\omega}_1(\omega)$ that explains the observed signal frequency $q(a_1^*,\omega)$ given his perception $\hat{q}$.[29] In period 2, he then chooses action $a_2^* = a^*(\hat{\omega}_1)$, which typically yields a contradiction between the expected and actual signal frequencies $\hat{q}(a_2^*,\hat{\omega}_1)$ and $q(a_2^*,\omega)$. Analogous to (1), we postulate that this gives rise to the following process of point-mass beliefs in periods $t \geq 2$:[30]

$$\hat{\omega}_t = \operatorname*{argmin}_{\hat{\omega}\in\Omega} \mathrm{KL}\left(q(a_{t-1}^*,\omega), \hat{q}(a_{t-1}^*,\hat{\omega})\right) \text{ with } a_{t-1}^* = a^*(\hat{\omega}_{t-1}). \tag{2}$$

Ignoring boundary cases, steady states $\hat{\omega}_\infty$ of (2) satisfy $q(a^*(\hat{\omega}_\infty),\omega) = \hat{q}(a^*(\hat{\omega}_\infty),\hat{\omega}_\infty)$.

Process (2) differs from (1) in one key respect: The true state ω enters *directly* into (2). Because of this, the correctly specified case $\hat{q} = q$ has $\hat{\omega}_\infty = \omega$ as its *unique* steady state. Under mild regularity conditions which ensure that the implicit function theorem applies, a key implication of this uniqueness is that small amounts of misperception continue to yield steady states (and hence long-run beliefs) that are close to the truth; see Supplementary Appendix F. In contrast, as noted in Section 4.2, process (1) is decoupled from ω, because other agents' actions are the only source of new information in each period $t \geq 2$, and these depend on ω only indirectly through others' beliefs. As a result, (1) admits a *continuum* of steady states in the correctly specified case $\hat{F} = F$; indeed, $\mathrm{SS}(F,F) = \Omega$. This renders system (1) fragile against perturbations at $\hat{F} = F$, so that even a slight amount of misperception can have a discontinuous impact on steady states.

Interpretation of discontinuity and countervailing forces. While Theorem 1 exhibits a stark discontinuity of long-run beliefs to small amounts of misperception, our preferred interpretation of this result places less emphasis on the formal discontinuity than on the substantive implication that slight misperceptions can have a large negative impact on social learning. Indeed, below we highlight two countervailing forces that render long-run beliefs continuous in $\hat{F}$, but show that when these forces are weak this substantive implication is unaffected:

Repeated private signals. As in much of the social learning literature, we have assumed that each agent has access to a single private signal about ω, and our results are unaffected if agents receive

[28] The purpose of assuming that action set A is the unit interval and signals s_t are binary is to make this model as analogous as possible to our model, where $a^*(\hat{\omega})$ corresponds to the aggregate action frequency $F(\theta^*(\hat{\omega}))$ and s_t to the observation of a random agent j's action a_{jt}. Neither assumption is essential for robustness of learning in this model.

[29] As in Heidhues, Koszegi, and Strack (2018), our setting has the feature that observing infinitely many signals leads to a point-mass on a unique state: By monotonicity and surjectivity of $\hat{q}$, $\hat{\omega}_1$ is uniquely given by $q(a_1^*,\omega) = \hat{q}(a_1^*,\hat{\omega}_1)$. This abstracts away from possible non-identification problems.

[30] Process (2) is also analogous to the heuristic analysis in Section 3 of Heidhues, Koszegi, and Strack (2018).

private signals in finitely many periods.[31] On the other hand, if agents receive private signals in *all* periods, then agents' long-run beliefs $\hat{\omega}_\infty(\omega)$ depend on the true state ω and vary continuously with $\hat{F}$; intuitively, this setting is a hybrid of social and single-agent learning, and as discussed above, belief dynamics in the latter do not become decoupled from the true state over time.[32] However, whenever agents' repeated private signals are sufficiently uninformative, we obtain an approximate analog of Theorem 1, where for any $\hat{\omega}$ there exist arbitrarily small amounts of misperception $\hat{F}$ such that agents' long-run beliefs $\hat{\omega}_\infty(\omega)$ in each state are *arbitrarily close* to a point mass on $\hat{\omega}$.[33] Thus, whenever agents' dominant source of information is social learning rather than private signals, the basic insight that slight misperceptions of others' characteristics can have a large negative impact on information aggregation remains valid.

Finite horizon. Following much of the literature on information aggregation, our analysis focuses on asymptotic beliefs. For any fixed finite horizon t, it is not difficult to see that the distribution μ_t^ω of agents' posteriors is continuous in $\hat{F}$. Nevertheless, analogous to the previous paragraph, Theorem 1 immediately entails that even under arbitrarily small amounts of misperception, μ_t^ω can be arbitrarily close to a Dirac measure on the state-independent point-mass $\delta_{\hat{\omega}}$ whenever t is large enough. One implication is that in contrast with the correctly specified benchmark, halting agents' interactions after a certain number of periods may improve social learning in the presence of misperception.

4.4 Examples: Nature of Misperception Shapes Long-Run Beliefs

As we have argued, a key implication of Theorem 1 is that information aggregation under misperception should be studied in its own right, without relying on the predictions of the correctly specified benchmark. In characterizing long-run beliefs under any (analytic) F and $\hat{F}$ as steady states $\mathrm{SS}(F,\hat{F})$ of process (1), the proof of Theorem 2 provides a starting point for studying how the *nature* of agents' misperception shapes their long-run beliefs. In the following, we illustrate this for two natural forms of misperception.

For this exercise, it is worth highlighting the following robustness property of such predictions. Even though Theorem 1 shows that the mapping from true and perceived type distributions $(F,\hat{F})$ to steady states $\mathrm{SS}(F,\hat{F})$ exhibits a discontinuity at $\hat{F}=F$ (i.e., in the correct perceptions case), the mapping *is* continuous in both $\hat{F}$ and F for any well-behaved *misperceptions* $\hat{F}\neq F$.[34] Thus,

[31] If private signals are costly, then agents might choose to acquire only finitely many signals under certain classes of cost functions (e.g., Ali, 2018; Burguet and Vives, 2000).

[32] Such hybrid formulations are less typical in the information aggregation literature, as in this case (absent misspecification) each agent can perfectly learn the state on his own without having to draw inferences from other agents' behavior. A recent exception is Harel, Mossel, Strack, and Tamuz (2017) who focus not on information aggregation but on the rate of learning.

[33] More formally, suppose agents receive i.i.d. signal draws from $\Phi(\cdot|\omega)$ in all periods. For any analytic $F,\hat{F}$, similar reasoning as in the proof of Theorem 2 shows that agents' long-run beliefs are given by state-dependent point-mass beliefs $\hat{\omega}_\infty(\omega)\in\mathrm{argmin}_{\hat{\omega}}\,\mathrm{KL}(F(\theta^*(\hat{\omega}_\infty(\omega)),\hat{F}(\theta^*(\hat{\omega})))+\mathrm{KL}(\Phi(\cdot|\omega),\Phi(\cdot|\hat{\omega}))$. Then a similar logic as for Theorem 1 shows that for any F, $\hat{\omega}$, and $\varepsilon>0$, there exists $\hat{F}$ with $\|\hat{F}-F\|<\varepsilon$ and $\delta>0$ such that if $\|\Phi(\cdot|\underline{\omega})-\Phi(\cdot|\overline{\omega})\|<\delta$, then in all states ω, almost all agents' beliefs converge to a point mass on some $\hat{\omega}_\infty(\omega)\in[\hat{\omega}-\varepsilon,\hat{\omega}+\varepsilon]$.

[34] In particular, suppose that $F\neq\hat{F}$ are analytic such that (i) $F'(\theta)\neq\hat{F}'(\theta)$ whenever $F(\theta)=\hat{F}(\theta)$, and (ii) $F(\theta)\neq\hat{F}(\theta)$ at $\theta=\theta^*(\underline{\omega}),\theta^*(\overline{\omega})$. Then one can verify that for any sequences $\hat{F}_n$ and F_n converging to $\hat{F}$ and F,

our framework can provide predictions for long-run beliefs under various well-documented forms of misperception, and these predictions are *robust*, in the sense that they are not sensitive to the exact true and perceived type distributions.

Example 1 (**First-Order Stochastic Dominance and "Fake" Recommendations**). We first consider the possibility that the true type distribution first-order stochastically dominates agents' perceptions or vice versa, so that agents systematically under- or overestimate the share of types above any given level. As discussed in Section 2.4, one natural example of this form of misperception is word-of-mouth communication about a new product where agents underestimate the share of "fake" recommenders. Fake recommenders can be modeled as types $\theta \geq \theta^*(\underline{\omega})$ who take action 1 ("recommend") irrespective of the true quality of the product; underestimating their share then naturally corresponds to $\hat{F}$ being first-order stochastically dominated by F.[35] In the correctly specified model of Section 3.1, the presence of fake recommenders has no effect in the long run, as agents continue to learn the true state.

By contrast, the following result shows that fake recommendations can be a highly effective tool for "manipulating" consumers' beliefs, suggesting a possible rationale for the prevalence of this marketing strategy: As long as consumers slightly underestimate the share of such recommendations, their presence can lead to drastic overoptimism about the quality of the product, in the sense that long-run beliefs are a point mass on the highest quality $\overline{\omega}$, regardless of the true quality ω.

For this result, it is sufficient that F ***strictly first-order stochastically dominates*** $\hat{F}$ ***on the set*** $\Theta^* := (\theta^*(\overline{\omega}), \theta^*(\underline{\omega}))$; that is, $F(\theta) < \hat{F}(\theta)$ for all $\theta \in \Theta^*$, which we denote by $F \succ_{\mathrm{FO}_{\Theta^*}} \hat{F}$. Intuitively, all types below $\theta^*(\overline{\omega})$ (respectively, above $\theta^*(\underline{\omega})$) have the same dominant action 0 (respectively, 1) so that agents' perceptions about the relative type distributions outside Θ^* are irrelevant.

Corollary 1 (Overoptimism/-pessimism). *Fix any $F, \hat{F} \in \mathcal{F}$. If $F \succ_{\mathrm{FO}_{\Theta^*}} \hat{F}$ (respectively, $\hat{F} \succ_{\mathrm{FO}_{\Theta^*}} F$), then in any state ω, almost all agents' beliefs converge to a point mass on $\overline{\omega}$ (respectively, $\underline{\omega}$).*

Figure 3 illustrates the intuition in the limit model. When agents systematically underestimate the share of types above any given level, then under any belief about the state, they are surprised by the lower than expected frequency of action 0. They respond with continual upward adjustments to their beliefs, converging eventually to a point mass on the unique steady state $\overline{\omega}$. □

Example 2 (**Underestimation of Population Heterogeneity**). Another widely documented form of misperception is that in many contexts, individuals tend to underestimate type heterogeneity in society.[36] We capture this by means of the commonly used dispersiveness order (Shaked and

---

$\mathrm{SS}(F_n, \hat{F}_n)$ converges to $\mathrm{SS}(F, \hat{F})$ with respect to the Hausdorff topology. Note that assumption (i) rules out the case in which F and $\hat{F}$ "touch" at some point in $[\theta^*(\underline{\omega}), \theta^*(\overline{\omega})]$.

[35] Concretely, suppose that the true type distribution $F = \beta P + (1-\beta)G$ is a convex combination of a distribution P of "promotional" types whose support is contained in $[\theta^*(\underline{\omega}), +\infty)$ and a full-support distribution G of "genuine" types; and suppose that $\hat{F} = \hat{\beta}P + (1-\hat{\beta})G$ where $\hat{\beta} < \beta$.

[36] For example, several studies (e.g., Norton and Ariely, 2011; Engelhardt and Wagener, 2015) find systematic underestimation of wealth inequality in many countries.

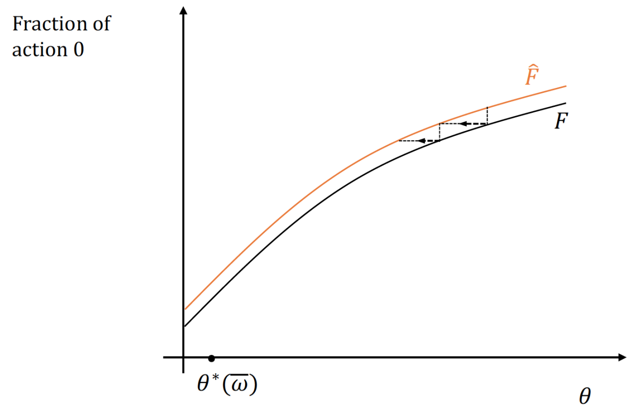

Figure 3: Overoptimism under FOSD. When $F \succ_{\mathrm{FO}_{\Theta^*}} \hat{F}$, beliefs in all states converge to a point mass on $\overline{\omega}$.

Shanthikumar, 2007), whereby F is ***more dispersive*** than $\hat{F}$ if $F^{-1}(x)-F^{-1}(y) \geq \hat{F}^{-1}(x)-\hat{F}^{-1}(y)$ for all type quantiles $x, y \in (0,1)$ with $x > y$; we denote this by $F \succsim_{\mathrm{disp}} \hat{F}$. For example, under Gaussian distributions $F \sim \mathcal{N}(\mu, \sigma^2)$, $\hat{F} \sim \mathcal{N}(\hat{\mu}, \hat{\sigma}^2)$, this takes the simple parametric form that perceived type variance $\hat{\sigma}^2$ is lower than actual variance σ^2. We rule out the possibility that $F \succ_{\mathrm{FO}_{\Theta^*}} \hat{F}$ or $\hat{F} \succ_{\mathrm{FO}_{\Theta^*}} F$, as this is covered by Corollary 1 above.

The following result shows that underestimation of population heterogeneity leads to ***conservative*** long-run beliefs, in the sense that beliefs in all states converge a point mass on an *interior* state $\hat{\omega} \in (\underline{\omega}, \overline{\omega})$:

Corollary 2 (Conservative beliefs)**.** *Fix any analytic $F, \hat{F} \in \mathcal{F}$ with $\hat{F} \neq F$ such that $\hat{F}, F$ are not strictly first-order stochastic dominance ranked on Θ^*. If $F \succsim_{\mathrm{disp}} \hat{F}$, then there exists some $\hat{\omega} \in (\underline{\omega}, \overline{\omega})$ such that in any state ω, almost all agents' beliefs converge to a point mass on $\hat{\omega}$.*

Intuitively, when agents underestimate type heterogeneity, they *overestimate* the sensitivity of the population action distribution against the state, because in any state, they expect different agents to take more similar actions than they actually do. As a result, their belief updates after observing others' actions are more "sluggish" than they should be, leading to conservatism in long-run beliefs. More formally, Figure 4 (left) shows that Corollary 2 corresponds to a setting where $\hat{F}$ crosses F from below in a single point $\theta^*(\hat{\omega})$ that corresponds to an interior state $\hat{\omega} \in (\underline{\omega}, \overline{\omega})$. This is the same setting we analyzed in Figure 2 as part of the proof sketch of Theorem 1, where we saw that $\mathrm{SS}(F, \hat{F}) = \{\hat{\omega}\}$.[37]

In other contexts, people are found to overestimate population heterogeneity.[38] This corresponds to $\hat{F} \succsim_{\mathrm{disp}} F$, where we again assume that $\hat{F}, F$ are not first-order stochastic dominance ranked on Θ^*. As illustrated in Figure 4 (right), in this case $\hat{F}$ crosses F from *above* in a single point $\theta^*(\hat{\omega})$ with $\hat{\omega} \in (\underline{\omega}, \overline{\omega})$, and the limit model predicts convergence to the ***extreme*** beliefs $\underline{\omega}$ and $\overline{\omega}$ in almost

[37]In the proof of Theorem 1, we also allow for the possibility that $\theta^*(\hat{\omega}) \in \{\theta^*(\overline{\omega}), \theta^*(\underline{\omega})\}$, in which case beliefs converge to the corresponding boundary point. In Corollary 2, this is ruled out by the assumption that $\hat{F}, F$ are not first-order stochastic dominance ranked on Θ^*.

[38]See, e.g., Ahler (2014) in the context of perceived political attitudes.

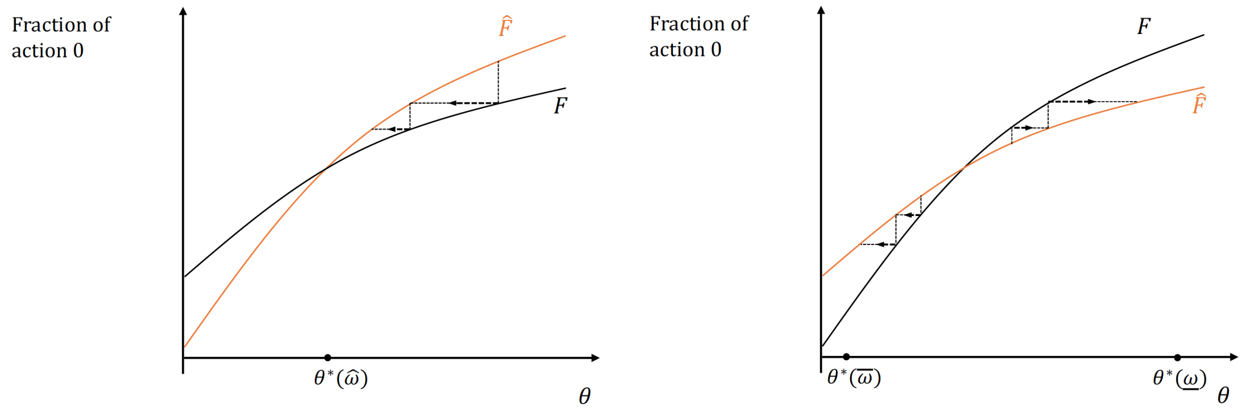

Figure 4: Left: Underestimation of population heterogeneity leads to conservative long-run beliefs. Right: Overestimation leads to extreme beliefs.

all states ω.[39] □

5 Rich vs. Coarse State Spaces

In this section, we show that a key determinant of how fragile information aggregation is under misperception is how rich a space of uncertainty agents face. We also discuss some design implications of this finding. Throughout, we fix some countably infinite set of states $\{\omega^1, \omega^2, ...\}$ that is dense in $\Omega = [\underline{\omega}, \overline{\omega}]$ and let $\Omega_n := \{\omega^1, ..., \omega^n\}$ for each n.

The following result makes two points. First, for any *fixed* finite state space Ω_n, information aggregation is robust, in the sense that if the amount of misperception is small enough agents learn the true state in Ω_n. Second, however, the *larger* the state space Ω_n, the *more sensitive* information aggregation is to small amounts of misperception, and in the limit as $n \to \infty$, we obtain an approximate analog of the extreme breakdown of information aggregation in Theorem 1:

Theorem 3 (Finite state space). *Fix any $F \in \mathcal{F}$.*

1. *Fix any Ω_n. There exists $\varepsilon_n > 0$ such that under any perception $\hat{F} \in \mathcal{F}$ with $\|\hat{F} - F\| < \varepsilon_n$ and in any state $\omega \in \Omega_n$, almost all agents' beliefs converge to a point mass on ω.*

2. *Fix any $\hat{\omega} \in \Omega$. For any $\varepsilon > 0$, there exists N and a perception $\hat{F} \in \mathcal{F}$ with $\|\hat{F} - F\| < \varepsilon$ under which for any state space Ω_n with $n \geq N$ and in any state $\omega \in \Omega_n$, almost all agents' beliefs converge to a point mass on some $\hat{\omega}_\infty(\omega) \in [\hat{\omega} - \varepsilon, \hat{\omega} + \varepsilon]$.*

[39] We state this prediction only for the limit model of Section 4.2, because in addition to $\underline{\omega}$, $\overline{\omega}$, the set of steady states of (1) also includes $\hat{\omega}$, so that the proof of Theorem 2 only implies that agents' long-run point-mass beliefs in the original model are given by a weakly increasing mapping $\hat{\omega}_\infty : \Omega \to \{\underline{\omega}, \hat{\omega}, \overline{\omega}\}$. However, steady state $\hat{\omega}$ is unstable, because in almost all states ω, (1) converges to either $\underline{\omega}$ or $\overline{\omega}$. Given this, we conjecture that in the original model $\hat{\omega}_\infty(\omega)$ likewise takes values $\underline{\omega}$ or $\overline{\omega}$ in almost all states ω, but establishing this requires an argument for non-convergence to unstable steady states of (1), which we do not pursue in the present paper.

The second part of Theorem 3 offers an approximate analog of Theorem 1 in the following sense: For any $\hat{\omega}$, Theorem 1 exhibits arbitrarily small amounts of misperception such that agents' long-run belief is a state-independent point mass on $\hat{\omega}$. In the present setting, agents' beliefs converge to a point mass on $\hat{\omega}_\infty(\omega)$, which may depend on the true state ω. However, Theorem 3 exhibits arbitrarily small amounts of misperception such that in all large enough state spaces Ω_n, agents' long-run belief $\hat{\omega}_\infty(\omega)$ is *arbitrarily close* to a state-independent point mass on $\hat{\omega}$. An immediate implication is that in the first part of Theorem 3, the amount of misperception ε_n below which information aggregation is successful in Ω_n shrinks to 0 as $n \to \infty$.

We prove Theorem 3 in Appendix C. To see the intuition for the first part, suppose that $n = 2$. Consider the limit model from Section 4.2, whose conclusions can again be shown to approximate those of the original model in the long run. Just as in the continuous state setting, after observing the action frequency in the population at the end of period 1, all agents again commonly believe in some state $\hat{\omega}_1$; and as summarized by equation (1), from period 2 on, agents play threshold strategies according to their current point-mass beliefs $\hat{\omega}_t$ and adjust these beliefs at the end of each period to explain the observed action frequency in that period under their misperception $\hat{F}$. Analogous to the continuous state case, $\hat{\omega}_1$ again depends on the true state ω and converges to ω as $\|\hat{F} - F\| \to 0$; in the binary state setting, this means that $\hat{\omega}_1$ in fact equals the true state when $\hat{F}$ is sufficiently close to F.

However, the key difference with the continuous state setting concerns belief adjustments from period 2 on. In the continuous setting, whenever $\hat{F}(\theta^*(\hat{\omega}_1)) \neq F(\theta^*(\hat{\omega}_1))$, then (ignoring boundary cases) agents can find some new point-mass belief $\hat{\omega}_2$ that better explains observed behavior $F(\theta^*(\hat{\omega}_1))$; that is, even small discrepancies between expected and observed behavior trigger a sequence of belief adjustments and corresponding adjustments in behavior, giving rise to the process of decoupled mislearning we highlighted in Section 4.2.

By contrast, in the binary state setting, whenever the amount of misperception is sufficiently small, agents do not further adjust their beliefs in period 2 and beyond. To see this, suppose, say, that $\hat{\omega}_1 = \omega^1$, so that observed behavior in period 2 is $F(\theta^*(\omega^1))$. Then, even though expected behavior $\hat{F}(\theta^*(\omega^1))$ under ω^1 does not perfectly match this observed behavior, there is only one other possible state ω^2, and when $\|\hat{F} - F\|$ is sufficiently small, expected behavior $\hat{F}(\theta^*(\omega^2))$ under ω^2 will be even farther from $F(\theta^*(\omega^1))$ than under ω^2. This is illustrated in Figure 5. Thus, the binary state setting is robust to small amounts of misperception, because small discrepancies between expected and observed behavior do not trigger adjustments to beliefs about the state. Given the coarseness of the state space, believing in the true state ω is sustainable, as the behavior this gives rise to cannot be better explained by any other state.

However, the second part of Theorem 3 implies that as the state space becomes richer, the amount of misperception ε under which information aggregation remains successful becomes smaller and smaller. Intuitively, any discrepancy between observed and expected behavior is more likely to trigger a belief adjustment the more alternative states there are that could explain this discrepancy. In the limit as the number of states approaches infinity, ε shrinks to 0, thus effectively restoring the

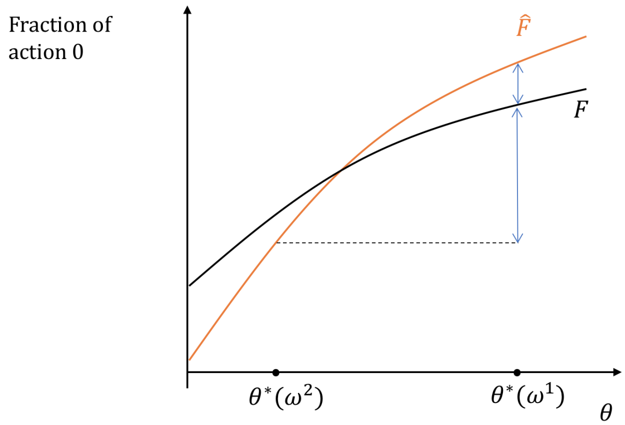

Figure 5: The limit model with a binary state space $\Omega = \{\omega^1, \omega^2\}$. When $\hat{\omega}_t = \omega^1$, then for small enough amount of misperception, actual behavior $F(\theta^*(\omega^1))$ is better explained by expected behavior $\hat{F}(\theta^*(\omega^1))$ under the true state ω^1 than by expected behavior $\hat{F}(\theta^*(\omega^2))$ under the only other state ω^2.

conclusion of Theorem 1.

Many settings of economic interest naturally feature rich state spaces, from safety levels of new products to market fundamentals in decentralized trade settings. In such settings, an important implication of Theorem 3 is that there can be a benefit to "simplifying" agents' learning environment, for instance by making their private information or their payoffs less sensitive to fine details of the fundamentals. We illustrate this with two examples.

Example 3 (Benefits of undetailed public news)**.** Suppose that agents' period 0 private information is obtained in the following manner. Similar to Myatt and Wallace (2014),[40] there is a benevolent sender (e.g., a central bank or consumer protection agency) that acquires and truthfully communicates information about the state of the world (e.g., market fundamentals or the safety of a new product) subject to two frictions: First, there may be limits on the sender's ability to acquire information; second, there is some "receiver noise," in that any given news release might be interpreted differently by different agents.

This is modeled as follows. First, in each state $\omega \in \Omega = [\underline{\omega}, \overline{\omega}]$, the sender observes a signal $\sigma(\omega) \in \mathbb{R}$, where $\sigma : \Omega \to \mathbb{R}$ is weakly increasing; here the partition $\Pi_\sigma := \{\sigma^{-1}(a) : a \in \mathbb{R}\}$ of Ω represents the sender's possibly imperfect information about the state. Second, the sender communicates his signal $\sigma(\omega)$, but each agent observes this signal with some idiosyncratic noise; specifically, agent i observes signal $s_i = \sigma(\omega) + \eta_i$, where η_i is drawn i.i.d. across agents and states from a mean zero distribution with positive log-concave density on $\mathbb{R}$.[41] The induced private signal distributions $\Phi(s_i|\omega) = \Pr(\sigma(\omega) + \eta_i \leq s_i)$ are measurable with respect to the sender's partition Π_σ. Thus, a strictly increasing σ (i.e., a perfectly informed sender) corresponds to the continuous state space setting of Section 2,[42] while if σ has finite range, then the setting is isomorphic to one

[40] See also Morris and Shin (2007); Myatt and Wallace (2011); Pavan (2014); Galperti and Trevino (2018).

[41] We assume for simplicity that idiosyncratic noise η_i does not vary across different sender signal technologies σ. Introducing such variation does not affect our conclusions, as long as η_i always has full support.

[42] Signal distributions satisfy the monotone likelihood ratio property as η_i has log-concave density.

with a finite state space where each state corresponds to a cell of Π_σ.

If in all subsequent period $t \geq 1$ agents draw inferences from each other's behavior as in our model in Section 2, then our analysis implies the following. If agents are correct about the type distribution F, then the better informed the sender (i.e., the finer Π_σ) the better this is for long-run learning, as agents' beliefs converge to a point mass on the correct cell of Π_σ; in particular, if and only if the sender is perfectly informed, agents always learn the exact state. By contrast, if agents' perception $\hat{F}$ is even slightly incorrect, this gives rise to the following new trade-off: On the one hand, the finer Π_σ, the more precise is agents' long-run information about the state *if* aggregation is successful; but on the other hand, information aggregation is more sensitive to misperception and, at worst, may break down completely. As a result, worse informed senders (i.e., coarser partitions Π_σ) can be better. Moreover, even if a benevolent sender has access to precise information about the state, he has a rationale to commit not to fully release it (as might be achieved, e.g., by central banks establishing a reputation for "vague" or "undetailed" announcements).[43] □

Example 4 (Benefits of simple financial assets)**.** Similar to Duffie and Manso (2007) (cf. example (iii) in Section 2.4), suppose that actions $a_t = 1, 0$ correspond to bidding or not bidding for some financial asset that is traded in a decentralized market.[44] Assets are modeled as maps $\rho : \Omega \to \mathbb{R}$ which in each state of the world $\omega \in [\underline{\omega}, \overline{\omega}]$ yield long-run payoff $\rho(\omega)$, where ρ is weakly increasing and continuous; here the partition $\Pi_\rho := \{\rho^{-1}(a) : a \in \mathbb{R}\}$ of Ω reflects how finely the asset conditions on the details of the economic environment. Each agent i's long-run utility to action $a_t = 1$ in state ω is given by $u(\omega, \theta_i) = v_{\theta_i}(\rho(\omega))$, where v_{θ_i} is i's strictly increasing and continuous utility for money.

Note that each $u(\cdot, \theta)$ is measurable with respect to Π_ρ. Thus, analogous to Example 3, a strictly increasing ρ corresponds to the continuous state space setting in Section 2, while if ρ has finite range, then $u(\cdot, \theta)$ is a step function and the setting is isomorphic to one with a finite state space where each state corresponds to a cell of Π_ρ. Hence, as in the previous example, our analysis implies that if agents' perception $\hat{F}$ is incorrect, then coarser partitions (i.e., "simpler" assets) may enable agents to learn more about the state in the long run, while if Π_ρ is too fine (i.e., the asset is too "complex"), agents may come to hold highly incorrect beliefs about the state, with adverse consequences for long-run bidding behavior and payoffs. □

6 More General Perceptions

So far, we have assumed that all agents share the same perception $\hat{F}$ of the the type distribution and that this is common knowledge among agents. Under this assumption, agents' first-order beliefs about F may be incorrect, but their higher-order beliefs (about other agents' beliefs about F and

[43]This rationale for vague communication is complementary to ones based on strategic externalities across agents (e.g. Morris and Shin, 2002) or between sender and receiver (e.g., Crawford and Sobel, 1982).

[44]More realistically, bids might be modeled as continuous, as in Duffie and Manso's (2007) example of a wallet game à la Klemperer (1998). As mentioned, our conclusions are unaffected by continuous action spaces.

others' beliefs about others' beliefs about F etc.) are correct.[45] Thus, Theorem 1 highlights that slightly incorrect first-order beliefs are enough to generate extreme departures from the correctly specified model.

At the same time, an important question in many models featuring agents that are in some way "non-standard" is how such agents interact with standard and sophisticated agents, who are aware of the presence of these non-standard agents.[46] In Section 6.1, we investigate this question in our setting by incorporating a fraction of agents who know the true type distribution. We show that there are learning externalities between the two groups of agents that give rise to a new form of non-robustness: Information aggregation is highly sensitive to sophisticated agents' *second-order* beliefs. In addition, Section 6.2 briefly discusses other generalizations of our baseline model of perceptions.

6.1 Interaction between Correct and Incorrect Agents

Specifically, we now extend our baseline model so that (independently of types) fraction $\alpha \in [0,1]$ of agents (referred to as ***incorrect*** agents) misperceive the type distribution to be $\hat{F}$ and believe that $\hat{F}$ is common knowledge among all agents. The remaining fraction $1-\alpha$ of agents know the true type distribution F and are aware of the presence of incorrect agents and their beliefs. However, we allow for the possibility that they may be (slightly) wrong about the fraction of incorrect agents in the population; specifically, they perceive this fraction to be $\hat{\alpha}$. We refer to this second group of agents as ***correct*** if $\hat{\alpha} = \alpha$ and ***quasi-correct*** if $\hat{\alpha} \neq \alpha$. Our baseline model corresponds to the case where $\alpha = 1$.

We first show that correct agents, who *exactly* know the fraction of incorrect agents, are able to learn the true state in the long run. Moreover, correct agents exert a *positive* externality on the learning of incorrect agents: In contrast with the baseline model, incorrect agents are now able to approximately learn the true state as long as their amount of misperception is sufficiently small.[47]

Proposition 1. *Fix any $\hat{\alpha} = \alpha < 1$.*

1. *Fix any $F \in \mathcal{F}$. Under any $\hat{F} \in \mathcal{F}$ and in any state ω, almost all correct agents' beliefs converge to a point mass on ω.*

2. *Fix any analytic $F \in \mathcal{F}$ and any $\delta > 0$. There exists $\varepsilon > 0$ such that under any analytic $\hat{F} \in \mathcal{F}$ with $\|\hat{F} - F\| < \varepsilon$ and in any state ω, almost all incorrect agents' beliefs converge a point mass on some state $\hat{\omega}_\infty(\omega)$ with $|\hat{\omega}_\infty(\omega) - \omega| < \delta$.*

[45] This setting also has the feature that when $\hat{F}$ is close to F, then agents' hierarchy of beliefs is close (in the product topology) to the correctly specified setting where there is common knowledge of F.

[46] Early explorations of this question include Haltiwanger and Waldman (1985); De Long, Shleifer, Summers, and Waldmann (1990). A recent example is Jehiel (2018) in whose setting the presence of correct agents reduces the welfare of naive agents.

[47] Note that Lemma 1 follows as a special case of Proposition 1 with $\alpha = 0$. Note also that Proposition 1 holds for α arbitrarily close to 1. Thus, even an arbitrarily small fraction of sophisticated agents is enough to enable incorrect agents to approximately learn the state. However, similar to Theorem 3, it can be shown that incorrect agents' learning is more sensitive to misperception the smaller the fraction of sophisticated agents, so that ε in part (2) shrinks to 0 as $\alpha \to 1$.

However, consider next the case of quasi-correct agents, who slightly misperceive the fraction of incorrect agents. Then Proposition 1 breaks down, and now it is the incorrect agents who exert a *negative* externality on quasi-correct agents' learning. Specifically, the following result extends Theorem 1 by showing that the presence of incorrect agents with an arbitrarily small amount of misperception can lead to state-independent and arbitrary long-run beliefs among *both* groups of agents; moreover, for this to occur, the fraction $\alpha > 0$ of incorrect agents can be arbitrarily small and quasi-correct agents' perception $\hat{\alpha}$ of this fraction can be arbitrarily close to the truth, as long as $\hat{\alpha} \neq \alpha$.

Proposition 2. *Fix any $F \in \mathcal{F}$, $\hat{\omega} \in \Omega$, and $\hat{\alpha}, \alpha > 0$ with $\hat{\alpha} \neq \alpha$. For any $\varepsilon > 0$, there exists $\hat{F} \in \mathcal{F}$ with $\|\hat{F} - F\| < \varepsilon$ such that in any state ω, almost all (quasi-correct and incorrect) agents' beliefs converge to a point mass on $\hat{\omega}$.*

The proofs of Propositions 1 and 2 appear in Supplementary Appendix E. The intuition behind Proposition 1 is that since correct agents know the fraction of incorrect agents, their knowledge of F and $\hat{F}$ allows them to back out the true state from observed behavior in the long run. Moreover, similar to footnote 17, the fact that correct agents' long-run behavior depends on the true state ω prevents incorrect agents' beliefs from unraveling too much.

By contrast, in Proposition 2, even if quasi-correct agents are only slightly wrong about α, then despite knowing the true type distribution and incorrect agents' misperception, they too face discrepancies between actual and anticipated behavior. This can give rise to an analogous process of decoupled mislearning as in Section 4.2. Moreover (ignoring boundary cases), incorrect and quasi-correct agents' steady states $(\hat{\omega}^I_\infty, \hat{\omega}^C_\infty)$ under this process must coincide: Indeed, when $\hat{\alpha} \neq \alpha$, this is the only way for quasi-correct agents not to face a discrepancy between the actual action 0 frequency $\alpha F(\theta^*(\hat{\omega}^I_\infty)) + (1-\alpha)F(\theta^*(\hat{\omega}^C_\infty))$ and their expected frequency $\hat{\alpha} F(\theta^*(\hat{\omega}^I_\infty)) + (1-\hat{\alpha})F(\theta^*(\hat{\omega}^C_\infty))$. As a result, for any $\hat{F} \neq F$, long-run outcomes are exactly the same as if there were *only* incorrect agents, yielding the above generalization of Theorem 1.[48]

6.2 Other Perceptions

We briefly comment on further extensions of our baseline model of perceptions.

Heterogeneous perceptions. While our baseline model assumes that agents share the same misperception $\hat{F}$, it might be more realistic to allow each agent i to hold his own misperception $\hat{F}_i$, in the sense that he believes the true type distribution to be $\hat{F}_i$ and (erroneously) believes this to be common knowledge among the population. For example, the false-consensus effect (Ross, Greene, and House, 1977) finds a positive association between people's own characteristics and their perceptions of others' characteristics. Generalizing Theorems 1 and 2 to such heterogeneous perceptions is not difficult, and the only main difference is that this extension naturally gives rise to heterogeneous long-run beliefs (i.e., disagreement).[49] Notably, even when agents' perceptions $\hat{F}_i$

[48] This observation also yields a generalization of Theorem 2.

[49] To see this, consider the limit model from Section 4.2. Let $\lambda \in \Delta(\mathbb{R} \times \mathcal{F})$ denote the joint distribution over

are on average correct (i.e., equal to F), their average long-run belief can be highly incorrect (e.g., state-independent or a coarse function of the state).

Higher-order perceptions. A further generalization of the setting in the previous paragraph is when agents are (partially) aware of the fact that others hold different perceptions.[50] While the general analysis of this case is beyond the scope of our paper, slight misspecifications of higher-order perceptions can also lead to discontinuous break-downs of information aggregation, as we have demonstrated for the hybrid model with correct and incorrect perceptions above.

Nondegenerate perceptions. Finally, as is common in the misspecified learning literature, we have assumed throughout that agents' perceptions are point-mass beliefs on particular distributions $\hat{F}$, so that agents do not update their perceived type distributions. A more general form of misspecification might involve agents holding a (common) prior belief over some set of type distributions, with support $\hat{\mathcal{F}} \subseteq \mathcal{F}$ that does *not* contain the true distribution F, and who update these beliefs (as well as their beliefs about ω) over time. A full analysis of this case is again beyond the scope of the current paper, but our main insights carry over; in particular, it is not difficult to construct arbitrarily small perturbations of the correctly specified model that feature nondegenerate misperceptions but yield the same extreme breakdown of information aggregation as in Theorem 1.

Finally, a substantially different setting is when the support $\hat{\mathcal{F}}$ contains the true type distribution F. In this case, agents are *correctly specified* (under common priors), which falls outside the focus of this paper. In particular, the martingale property of beliefs applies and rules out such extreme breakdowns of information aggregation as the state-independent point-mass beliefs in Theorem 1. However, medium-run predictions (at any fixed period t) under this model approximate those in our baseline model with common $\hat{F}$ whenever agents' prior places sufficiently high probability on a small neighborhood of $\hat{F}$. It is also worth noting that learning in this setting can be subject to an identification problem, as there are typically many combinations of states and type distributions that are consistent with the same observed action distributions, leaving open the possibility of incomplete long-run learning.

7 Concluding Remarks

7.1 Related Literature

Our paper contributes to the burgeoning literature on (Bayesian) learning with misspecified models, which has been studied in a variety of contexts spanning single-agent passive and active learning

types and perceptions in the population, with marginal $\text{marg}_{\mathbb{R}}\lambda$ over types given by F. For each $\hat{F} \in \mathcal{F}$ and fraction $q \in [0, 1]$ of action 0, define $\hat{\omega}(q; \hat{F}) = \text{argmin}_{\hat{\omega}} \text{KL}(q, \hat{F}(\theta^*(\hat{\omega})))$. Then the fraction q_t of action 0 in the population evolves according to $q_{t+1} = \lambda(\{(\theta, \hat{F}) : \theta^*(\hat{\omega}(q_t; \hat{F})) > \theta\}) =: g(q_t)$. Note that transition function g is independent of the realized state; that is, learning is again decoupled. Since g is increasing, q_t converges to one of the steady states of the system. Generalizing Theorem 1, if g crosses the identity function at a single point q^* from above, then $\lim_t q_t = q^*$ regardless of the true state ω, which implies that agents' (heterogeneous) long-run beliefs are state-independent; and this can occur even when all misperceptions in the support of λ are arbitrarily close to F. Likewise, coarse information aggregation arises when g crosses the identity finitely many times on $[0, 1]$.

[50]See Bohren and Hauser (2018) for such a formulation in the context of the sequential learning model.

(e.g., Nyarko, 1991; Rabin, 2002; Rabin and Vayanos, 2010; Ortoleva and Snowberg, 2015; Fudenberg, Romanyuk, and Strack, 2017; Heidhues, Koszegi, and Strack, 2018; He, 2018) and social learning (e.g., Eyster and Rabin, 2010; Guarino and Jehiel, 2013; Bohren, 2016; Gagnon-Bartsch, 2017; Dasaratha and He, 2017; Bohren and Hauser, 2018; Bohren, Imas, and Rosenberg, 2018).[51]

Among single-agent learning models, the most closely related paper is Heidhues, Koszegi, and Strack (2018) who study active learning by an agent who is overconfident in his ability. Similar to our mislearning process in Section 4.2, they highlight a process of "misguided" learning, where the fact that the agent's actions endogenously affect his information in each period can render his beliefs more and more incorrect over time. However, in contrast with our non-robustness result, the agent's long-run belief in this setting is approximately correct when his amount of misperception is small; the same is true in the other aforementioned misspecified single-agent models. We show that this distinction is due to a crucial feature of social learning that is absent under these single-agent learning models, namely the fact that belief updating becomes "decoupled" from the true state over time. Section 4.3 discusses the distinction in more detail.

In the misspecified social learning literature, many aforementioned papers incorporate various specific forms of misspecification into sequential learning models in the style of Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992). As shown by Eyster and Rabin (2010) and subsequent papers, some forms of misspecification can render agents' long-run beliefs confident in incorrect states, similar to our Theorem 1. However, these papers rely on strong forms of misspecification, and in a general model that nests several of these misspecifications Bohren and Hauser (2018) show that information aggregation is robust to small amounts of misspecification.[52] By contrast, Theorem 1 derives extreme breakdowns of information aggregation under arbitrarily small amounts of misperception, highlighting a setting where the correctly specified model is not robust. One important difference between these models and ours is the size of the state space (binary vs. continuous). In Section 5, we show that how rich a state space agents face can play a key role in determining how fragile information aggregation is and discuss design implications of this finding.

Berk-Nash equilibrium developed by Esponda and Pouzo (2016) characterizes possible long-run beliefs in single-agent learning models in terms of minimization of KL-divergence.[53] While long-run beliefs in social learning models (including ours) do not strictly fit into the Berk-Nash framework, our characterization of long-run beliefs is also based on minimization of KL-divergence. Indeed, we show that except for boundary cases, agents' long-run beliefs achieve zero KL-divergence between perceived and actual behavior.

It is worth emphasizing that the discontinuous departure from the correctly specified benchmark in our setting is driven purely by misinterpretation among agents, not by strategic considerations due to direct payoff externalities from others' actions. In the context of screening problems where

[51]See Glaeser and Sunstein (2009); Levy and Razin (2015) for static models of information aggregation with misspecified agents.

[52]In Gagnon-Bartsch (2017), agents' beliefs need not converge under small amounts of misperception, but long-run average beliefs depend continuously on agents' perceptions.

[53]More generally, Esponda and Pouzo (2016) consider strategic games with misspecified agents. See also Esponda and Pouzo (2017) for an analysis of learning by a forward-looking agent.

the principal is misspecified about agents' preferences, Madarász and Prat (2017) highlight the possibility of the latter kind of discontinuity and propose a remedy. Outside the literature on misspecification, Acemoglu, Chernozhukov, and Yildiz (2016) consider agents who observe exogenous public signals and hold non-common full support priors about the signal technology. They focus on higher-order belief disagreement and show that even a small amount of uncertainty can lead to substantial long-run disagreement, due to a non-identification problem in disentangling states and signal technologies. As agents in their model are correctly specified (i.e., their beliefs contain the truth in their support), it is not possible to generate long-run phenomena such as state-independent point-mass beliefs.

While we maintain the classical assumption that agents are Bayesian, a sizeable literature considers models of non-Bayesian social learning where agents update beliefs by employing various exogenous heuristics (e.g., Ellison and Fudenberg, 1993, 1995; DeMarzo, Vayanos, and Zwiebel, 2003; Golub and Jackson, 2010, 2012). Information aggregation can fail in such settings as well, depending on the specifics of the environment and of agents' heuristic rules.[54] Mueller-Frank (2018) considers a model of DeGroot-style learning on a network where an agent can manipulate her updating rule and shows that small amounts of manipulation can have a significant and arbitrary impact on other agents' long-run beliefs. An advantage of non-Bayesian models with simple updating heuristics is that in some situations they may be more realistic, as they do not assume that agents are able to form beliefs about the details of their environment and solve complex inference problems. In contrast, our focus on Bayesian agents is better suited to our goal of of understanding the robustness of the canonical model of rational social learning. This also allows us to explicitly model agents' perceptions and how they affect their process of inference and their interpretation of others, which is usually abstracted away from under heuristic rules.

7.2 Conclusion

We conclude with some remarks on additional directions that seem worthwhile to explore:

First, in order to investigate the robustness of the canonical model of rational social learning, we have maintained the assumption that our agents, while misspecified, are nevertheless Bayesian. One might question the descriptive validity of this assumption, as Bayesian updating requires agents to solve complicated inference problems, and also wonder to what extent the discontinuity results rely on it. While full analysis of this question is beyond the scope of this paper, we note that our results go through under some natural models of simpler inference. For one plausible example, consider the overlapping generations interpretation of our model, and assume that agents perform "coarse inferences" (Guarino and Jehiel, 2013), in the sense that all agents believe the action 0 frequency in all periods to be i.i.d. conditional on each state ω; specifically, they believe this to be equal to the long-run frequency $\hat{F}(\theta^*(\omega))$ under their misperception $\hat{F}$.[55] Under this model, all our results

[54] See, e.g. Molavi, Tahbaz-Salehi, and Jadbabaie (2018) and Sadler (2017) for systematic analyses of this question in the context of learning on networks.

[55] Guarino and Jehiel (2013) analyze a sequential learning model with homogeneous agents, and show that such inference procedures can break information aggregation. They motivate this formulation based on analogy-based

go through, with similar (in fact, somewhat simpler) proofs. Thus, some plausible simplifications of agents' updating processes preserve the fragility of information aggregation. At the same time, an interesting direction for future work is to understand what minimal departures from Bayesian learning might render information aggregation more robust to small amounts of misperception.

Second, while this paper considers a model of social learning through decentralized random interactions that is closest to Duffie and Manso (2007), a natural question is whether analogs of Theorem 1 hold for other social learning models. Our companion note (in preparation) explores this for the following two classes of models. First, a model that is identical to ours except that in each period t, all agents observe a *public* signal s_t of the aggregate action frequency $q_t(\omega)$, instead of (or in addition to) privately observing others' actions; correctly specified versions appear in Vives (1993, 1997); Amador and Weill (2012); Duffie, Malamud, and Manso (2010).[56] Second, a heterogeneous agent version of the sequential learning model of Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992), where agents sequentially choose one-shot actions after observing all previous agents' actions.

Finally, while we have assumed independent random matching, our model can be enriched to allow for *assortativity*, where agents are more likely to encounter other agents with similar types (e.g., income, political attitudes). Introducing assortative interactions allows one to analyze how agents' misperceptions of interaction patterns (e.g., underestimation of assortativity, Frick, Iijima, and Ishii, 2018) influence their beliefs. One new phenomenon that can emerge in this setting is that even when all agents share the same misperception of interaction patterns, agents' long-run beliefs can depend systematically on their types, in line with empirical evidence on correlation between tastes and beliefs.[57] In ongoing work, we pursue this question in more detail.

expectations equilibrium (Jehiel, 2005).

[56] One interpretation of such public signals is as market prices, where agents' actions represent demand/supply.

[57] See, e.g., Bullock, Gerber, Hill, and Huber (2013) in the context of partisan bias in factual beliefs about politics.

Appendix: Main Proofs

A Proof of Lemma 1 and Part 1 of Proposition 1

Consider the hybrid model from Section 6, with true and perceived fraction $\alpha = \hat{\alpha} < 1$ of incorrect agents. In this section, we prove part 1 of Proposition 1, i.e., that in every state ω, almost all correct agents' beliefs converge to δ_ω. Lemma 1 follows immediately as the special case where $\alpha = \hat{\alpha} = 0$.

We make use of the following additional notation. For any belief $H \in \Delta(\Omega)$, let $\theta^*(H)$ denote the type that is indifferent between action 0 and 1; that is, $\int u(\omega, \theta^*(H))dH(\omega) = 0$. By the assumptions on u, such a type exists and is unique and $\theta^*(H)$ is continuous in H under the topology of weak convergence. As in the main text, we let $\theta^*(\omega)$ denote $\theta^*(\delta_\omega)$. Given any private signal s, let $H^s \in \Delta(\Omega)$ denote the Bayesian update of H after observing s. Note that if H is not a Dirac measure, then by the assumptions on signal distributions Φ, H^s strictly first-order stochastically dominates $H^{s'}$ for any $s > s'$, which implies that $\theta^*(H^s) < \theta^*(H^{s'})$.

In each state ω and at the beginning of each period $t \geq 1$, let $\mu_t^{\omega,C}, \mu_t^{\omega,I} \in \Delta(\Delta(\Omega))$ denote the population distributions of posterior beliefs among correct (superscript C) and incorrect (superscript I) agents. Note that the distributions of posteriors are independent across types θ. Let $q_t^C(\omega)$ and $q_t^I(\omega)$ denote the fractions of correct and incorrect agents who choose action 0 in period t and ω; that is,

$$q_t^C(\omega) = \int F(\theta^*(H))d\mu_t^{\omega,C}(H), \quad q_t^I(\omega) = \int F(\theta^*(H))d\mu_t^{\omega,I}(H).$$

Let $q_t(\omega) = \alpha q_t^I(\omega) + (1-\alpha)q_t^C(\omega)$ denote the total fraction of action 0 in the population.

Finally, let $\bar{\mu}_t^{\omega,C}, \bar{\mu}_t^{\omega,I} \in \Delta(\Delta(\Omega))$ denote the *hypothetical* population distributions of posteriors among correct and incorrect agents when agents update beliefs only based on observing actions and do not take into account their private signals in period 0. More precisely, each agent observes actions a_τ in periods $\tau = 1, \ldots, t-1$ that are generated according to $q_\tau(\omega)$ defined above, and correct agents update their beliefs assuming that a_τ is distributed according to q_τ while incorrect agents assume that a_τ is distributed according to q_τ^I. Given this, we can also express q_t^C, q_t^I as

$$q_t^C(\omega) = \int\int F(\theta^*(H^s))d\Phi(s|\omega)d\bar{\mu}_t^{\omega,C}(H), \quad q_t^I(\omega) = \int\int F(\theta^*(H^s))d\Phi(s|\omega)d\bar{\mu}_t^{\omega,I}(H). \quad (3)$$

A.1 Correct Agents' Long-Run Behavior

The following four lemmas establish that (on some measure 1 set of states), $\lim_{t\to\infty} q_t^C(\omega)$ exists and is strictly decreasing in ω. The first lemma proves that q_t^C and q_t^I are strictly decreasing at all finite times.

Lemma A.1. *For each t, $q_t^C(\omega), q_t^I(\omega)$ are strictly decreasing in ω and satisfy $q_t^C(\omega), q_t^I(\omega) \in (0,1)$.*

Proof. The claim that $q_t^C(\omega), q_t^I(\omega) \in (0,1)$ is clear from the fact that types above $\theta^*(\underline{\omega})$ always choose action 1 and types below $\theta^*(\overline{\omega})$ always choose action 0. To show that $q_t^C(\omega), q_t^I(\omega)$ are strictly decreasing in ω, we proceed by induction on t. For $t = 1$, this follows from (3) and the fact that $\bar{\mu}_1^{\omega,C} = \bar{\mu}_1^{\omega,I} = \delta_\Psi$ and $\Phi(\cdot|\omega)$ is strictly increasing in ω with respect to first-order stochastic dominance. Suppose next that the claim holds for all periods up to and including t and consider period $t + 1$.

Let $H_C^{s,a^t} \in \Delta(\Omega)$ denote the posterior belief of a correct agent who observes private signal s and action sequence $a^t = (a_1, ..., a_t) \in \{0, 1\}^t$. Note that H_C^{s,a^t} has full support on Ω, since $q_\tau(\omega) \in (0, 1)$ for all $\tau = 1, .., t$ and all ω (so that a^t occurs with positive probability in each state) and by the full-support assumptions on prior Ψ and on private signals. Consider any $\omega^* > \omega^{**}$. For each $k = 2, ..., t - 2$, we have

$$
\begin{aligned}
&q_{t+1}^C(\omega^*) = \int\int F(\theta^*(H^s))d\Phi(s|\omega^*)d\bar{\mu}_{t+1}^{\omega^*,C}(H) < \int\int F(\theta^*(H^s))d\Phi(s|\omega^{**})d\bar{\mu}_{t+1}^{\omega^*,C}(H) \\
&= \int \sum_{a^t} \prod_{\tau=1}^{t} \left(q_\tau(\omega^*)(1 - a_\tau) + (1 - q_\tau(\omega^*))a_\tau\right) F(\theta^*(H_C^{s,a^t}))d\Phi(s|\omega^{**}) \\
&< \int \sum_{a^t} \prod_{\tau=1}^{t-1} \left(q_\tau(\omega^*)(1 - a_\tau) + (1 - q_\tau(\omega^*))a_\tau\right) \left(q_t(\omega^{**})(1 - a_t) + (1 - q_t(\omega^{**}))a_t\right) F(\theta^*(H_C^{s,a^t}))d\Phi(s|\omega^{**}) \\
&< \int \sum_{a^t} \prod_{\tau=1}^{k-1} \left(q_\tau(\omega^*)(1 - a_\tau) + (1 - q_\tau(\omega^*))a_\tau\right) \prod_{\tau'=k}^{t} \left(q_{\tau'}(\omega^{**})(1 - a_{\tau'}) + (1 - q_{\tau'}(\omega^{**}))a_{\tau'}\right) F(\theta^*(H_C^{s,a^t}))d\Phi(s|\omega^{**}) \\
&< \int\int F(\theta^*(H^s))d\Phi(s|\omega^{**})d\bar{\mu}_{t+1}^{\omega^{**},C}(H) = q_{t+1}^C(\omega^{**}).
\end{aligned}
$$

Here, the first inequality holds since $\theta^*(H^s)$ is strictly decreasing in s for each H in the support of $\bar{\mu}_{t+1}^{\omega^*,C}$ and since $\Phi(\cdot|\omega^*)$ strictly first-order dominates $\Phi(\cdot|\omega^{**})$. For the second inequality, note that since $q_t(\omega)$ is strictly decreasing in ω, $H_C^{s,a^{t-1},1}$ strictly first-order stochastically dominates $H_C^{s,a^{t-1},0}$. Thus, $\theta^*(H_C^{s,a^{t-1},1}) < \theta^*(H_C^{s,a^{t-1},0})$, which together with $q_t(\omega^*) < q_t(\omega^{**})$ yields the second inequality. Iterating this argument yields the remaining inequalities. The fact that q_{t+1}^I is strictly decreasing in ω follows by an analogous argument. □

To prove that q_t^C remains strictly decreasing in ω in the limit as $t \to \infty$, Lemmas A.2 and A.3 first consider the limit of the hypothetical belief distributions $\bar{\mu}_t^{\omega,C}$ that are based only on action observations. Lemma A.2 shows this limit exists almost surely.

Lemma A.2. *There exists a set of states $\Omega^* \subseteq \Omega$ such that $\Psi(\Omega^*) = 1$ and the weak convergent limit $\bar{\mu}_\infty^{\omega,C} := \lim_t \bar{\mu}_t^{\omega,C}$ exists for all $\omega \in \Omega^*$.*

Proof. To formalize correct agents' belief updating based only on observing sequences of actions (not private signals), consider the probability space $(\bar{\Omega}, \mathcal{A}, \mathbb{P})$. Here $\bar{\Omega} := \Omega \times \{0, 1\}^\infty$ is a Polish space (endowed with the product topology) that encodes the realized state $\omega \in \Omega$ and action sequences $a^\infty = (a_1, a_2, ...) \in \{0, 1\}^\infty$, $\mathcal{A}$ denotes the corresponding Borel algebra, and measure $\mathbb{P}$ satisfies

$$
\mathbb{P}\left(\{(\omega, a^\infty) \in \bar{\Omega} : \omega \in E, a_{t_1} = x_1, \ldots, a_{t_k} = x_k\}\right) = \int_E \prod_{j=1}^{k} \left((1 - q_{t_j}(\omega))x_j + q_{t_j}(\omega)(1 - x_j)\right) d\Psi(\omega)
$$

for every Borel set $E \subseteq \Omega$, $t_1, \ldots, t_k \in \mathbb{N}$, and $x_1, \ldots, x_k \in \{0,1\}$. Note that $\mathbb{P}$ exists by the Kolmogorov extension theorem. For each finite action sequence $a^t \in \{0,1\}^t$, Bayesian updating based on $\mathbb{P}$ induces a correct agent's posterior $H(\cdot|a^t) \in \Delta(\Omega)$ over states after observing a^t. Since Ω and $\{0,1\}^\infty$ are both Polish, the posterior belief $H(\cdot|a^\infty) \in \Delta(\Omega)$ conditional on each infinite sequence $a^\infty \in \{0,1\}^\infty$ is also well-defined (see Theorem 9.2.2 in Stroock, 2010).

We claim that $H(\cdot|a^t)$ converges weakly to $H(\cdot|a^\infty)$ $\mathbb{P}$-almost surely. To see this, define the filtration $\mathcal{I}_t := \sigma(a_1, \ldots, a_t) \subseteq \mathcal{A}$ that describes an agent's information at the end of each period t. Let $\mathcal{I}_\infty := \bigcup_{t=0}^\infty \mathcal{I}_t \subseteq \mathcal{A}$. For each Borel set $E \subseteq \Omega$, Levy's upwards theorem applied to the indicator function on E guarantees that as $t \to \infty$, $H(E|a^t) = \mathbb{E}[1_{\omega \in E}|\mathcal{I}_t] \to \mathbb{E}[1_{\omega \in E}|\mathcal{I}_\infty] = H(E|a^\infty)$ holds $\mathbb{P}$-almost surely (see Corollary 5.2.4 in Stroock, 2010). Let $\mathcal{Q}$ denote the set of all rational intervals in Ω, i.e., $\mathcal{Q} = \bigcup_{q_1,q_2 \in \mathbb{Q}\cap\Omega} \{[q_1,q_2],[q_1,q_2),(q_1,q_2],(q_1,q_2)\}$. As $\mathcal{Q}$ is countable, the event that

$$[\lim_{t\to\infty} H(B|a^t) = H(B|a^\infty) \ \ \forall B \in \mathcal{Q}] \tag{4}$$

holds $\mathbb{P}$-almost surely. This implies that $H(\cdot|a^t)$ converges weakly to $H(\cdot|a^\infty)$ $\mathbb{P}$-almost surely.[58]

Thus, there is a set of states $\Omega^* \subseteq \Omega$ with $\mathbb{P}(\Omega^* \times \{0,1\}^\infty) = \Psi(\Omega^*) = 1$ such that conditional on each $\omega \in \Omega^*$, $H(\cdot|a^t)$ converges weakly to $H(\cdot|a^\infty)$ $\mathbb{P}$-almost surely. Consider any $\omega \in \Omega^*$. Since $\bar{\mu}_t^{\omega,C} \in \Delta(\Delta(\Omega))$ is the distribution of $H(\cdot|a^t)$ conditional on ω and since conditional on ω, $H(\cdot|a^t)$ converges to $H(\cdot|a^\infty)$ $\mathbb{P}$-almost surely, this implies that $\bar{\mu}_t^{\omega,C}$ converges weakly, with limit $\bar{\mu}_\infty^{\omega,C} \in \Delta(\Delta(\Omega))$ given by $\bar{\mu}_\infty^{\omega,C}(\mathcal{H}) = \mathbb{P}\left[H(\cdot|a^\infty) \in \mathcal{H} | \, \omega\right]$ for any Borel set $\mathcal{H} \subseteq \Delta(\Omega)$. □

For Ω^* as in Lemma A.2, the following lemma shows that for each $\omega \in \Omega^*$, $\bar{\mu}_\infty^{\omega,C}$ assigns probability 1 to limit posteriors that contain the true state ω in their support.

Lemma A.3. *For any $\omega \in \Omega^*$, $\bar{\mu}_\infty^{\omega,C}(\{H : \mathrm{supp} H \ni \omega\}) = 1$.*

Proof. Fix any $\omega \in \Omega^*$. It suffices to prove the following claim: For any non-empty closed set $E \subseteq \Omega$ and non-empty open set $E' \subseteq \Omega$ such that either (i) $\omega \leq \inf E < \sup E \leq \inf E'$ or (ii) $\sup E' \leq \inf E < \sup E \leq \omega$, we have

$$\bar{\mu}_\infty^{\omega,C}\left(\{H : H(E) = 0 \text{ and } H(E') > 0\}\right) = 0.$$

To see that this claim implies Lemma A.3, note that for each H such that $\omega \notin \mathrm{supp} H$, we have for some n that either $H \in \mathcal{H}_n^+ := \{H' : H'([\omega, \omega + \frac{1}{n}]) = 0 \text{ and } H'((\omega + \frac{1}{n}, \overline{\omega}]) > 0\}$ or $H \in \mathcal{H}_n^- := \{H' : H'([\omega - \frac{1}{n}, \omega]) = 0 \text{ and } H'([\underline{\omega}, \omega - \frac{1}{n})) > 0\}$. But by the above claim, $\bar{\mu}_\infty^{\omega,C}(\mathcal{H}_n^+) = \bar{\mu}_\infty^{\omega,C}(\mathcal{H}_n^-) = 0$ for all n. Hence, by countable additivity of $\bar{\mu}_\infty^{\omega,C}$, $\bar{\mu}_\infty^{\omega,C}(\{H : \mathrm{supp} H \not\ni \omega\}) = 0$.

To prove the claim, we only consider case (i); the proof for case (ii) is analogous. Consider any correct agent and conditional on realized state ω, let (H_t) denote the process of his hypothetical

[58] To see this, note that for any open set U, there exists a countable sequence of pairwise disjoint sets $B_i \in \mathcal{Q}$ such that $\bigcup_{i=1}^\infty B_i = U$. Consider any $K \in \mathbb{N}$ and observe that by (4), $\liminf_{t\to\infty} H(U|a^t) = \liminf_{t\to\infty} \sum_{i=1}^\infty (B_i|a^t) \geq \liminf_{t\to\infty} \sum_{i=1}^K H(B_i|a^t) = \sum_{i=1}^K H(B_i|a^\infty)$ holds $\mathbb{P}$-a.s. Since this is true for any K, it follows that $\liminf_{t\to\infty} H(U|a^t) \geq \sum_{i=1}^\infty H(B_i|a^\infty) = H(U|a^\infty)$ holds $\mathbb{P}$-a.s. Thus, by the Portmanteau theorem, $H(\cdot|a^t)$ converges weakly to $H(\cdot|a^\infty)$ $\mathbb{P}$-a.s.

posterior beliefs that are based only on observing actions (without taking into account his private signal). By the proof of Lemma A.2, (H_t) weakly converges with probability 1, and the limit posterior is distributed according to $\bar{\mu}_\infty^{\omega,C}$.

Consider the process $W_t := H_t(E')/H_t(E)$, which is well-defined at every t since the posterior H_t always has full support. We have

$$\begin{aligned}\mathbb{E}\left[W_{t+1} \mid \mathcal{I}_t, \omega\right] &= \left(q_t(\omega)\frac{q_t(E')}{q_t(E)} + (1-q_t(\omega))\frac{1-q_t(E')}{1-q_t(E)}\right)\frac{H_t(E')}{H_t(E)} \\ &\leq \left(q_t(E)\frac{q_t(E')}{q_t(E)} + (1-q_t(E))\frac{1-q_t(E')}{1-q_t(E)}\right)\frac{H_t(E')}{H_t(E)} = W_t,\end{aligned}$$

where $\mathcal{I}_t$ denotes the filtration generated by the sequence of actions observed by the agent, $q_t(E) := \frac{\int_E q_t(\omega')dH_t(\omega')}{H_t(E)} \in (0,1)$ and $q_t(E') := \frac{\int_{E'} q_t(\omega')dH_t(\omega')}{H_t(E')} \in (0,1)$ denote the probabilities of observing action 0 conditional on events E and E', and the inequality holds because $q_t(\omega) > q_t(E) > q_t(E')$ by the assumption that $\omega \leq \inf E < \sup E \leq \inf E'$ and since $q_t(\omega')$ is strictly decreasing in ω' (Lemma A.1). Thus, W_t is a non-negative supermartingale conditional on ω.

By the martingale convergence theorem, there exists some $W_\infty \in L^1$ such that conditional on ω, $W_t \to W_\infty$ almost surely. Since $W_\infty \in L^1$, $W_\infty < +\infty$ almost surely. Thus, almost surely

$$\liminf_{t\to\infty} H_t(E') > 0 \Rightarrow \limsup_{t\to\infty} H_t(E) > 0.$$

By weak convergence of (H_t) and the Portmanteau theorem, this yields the desired claim. □

Based on Lemma A.3, we now establish that on Ω^*, q_t^C remains strictly decreasing in the limit.

Lemma A.4. *For any $\omega \in \Omega^*$, $q_\infty^C(\omega) := \lim_t q_t^C(\omega)$ exists and is strictly decreasing in ω.*

Proof. Recall from (3) that $q_t^C(\omega) = \int\int F(\theta^*(H^s))d\Phi(s|\omega)d\bar{\mu}_t^{\omega,C}(H)$, where $\int F(\theta^*(H^s))d\Phi(s|\omega)$ is continuous and bounded in H. Thus, since $\bar{\mu}_t^{\omega,C}$ weakly converges to $\bar{\mu}_\infty^{\omega,C}$ on Ω^*, $\lim_t q_t^C(\omega)$ exists for all $\omega \in \Omega^*$ and is given by

$$q_\infty^C(\omega) = \int\int F(\theta^*(H^s))d\Phi(s|\omega)d\bar{\mu}_\infty^{\omega,C}(H).$$

To show that q_∞^C is strictly decreasing, take any $\omega^*, \omega^{**} \in \Omega^*$ such that $\omega^* > \omega^{**}$. If $\bar{\mu}_\infty^{\omega^*,C} = \delta_{\delta_{\omega^*}}$ and $\bar{\mu}_\infty^{\omega^{**},C} = \delta_{\delta_{\omega^{**}}}$, then $q_\infty^C(\omega^*) = F(\theta^*(\omega^*)) < F(\theta^*(\omega^{**})) = q_\infty^C(\omega^{**})$. Thus, suppose that either $\bar{\mu}_\infty^{\omega^*,C} \neq \delta_{\delta_{\omega^*}}$ or $\bar{\mu}_\infty^{\omega^{**},C} \neq \delta_{\delta_{\omega^{**}}}$. We consider the case when $\bar{\mu}_\infty^{\omega^*,C} \neq \delta_{\delta_{\omega^*}}$; the other case is analogous.

We have

$$\begin{aligned}&q_\infty^C(\omega^*) = \int\int F(\theta^*(H^s))d\Phi(s|\omega^*)d\bar{\mu}_\infty^{\omega^*,C}(H) < \int\int F(\theta^*(H^s))d\Phi(s|\omega^{**})d\bar{\mu}_\infty^{\omega^*,C}(H) = \\ &\lim_{t\to\infty}\int\int F(\theta^*(H^s))d\Phi(s|\omega^{**})d\bar{\mu}_t^{\omega^*,C}(H) \leq \lim_{t\to\infty}\int\int F(\theta^*(H^s))d\Phi(s|\omega^{**})d\bar{\mu}_t^{\omega^{**},C}(H) = q_\infty^C(\omega^{**}).\end{aligned}$$

For the first inequality, note that since $\bar{\mu}_\infty^{\omega^*,C} \neq \delta_{\delta_{\omega^*}}$, Lemma A.3 implies that $\bar{\mu}_\infty^{\omega^*,C}$ assigns positive measure to beliefs H that are not Dirac measures. Since for non-Dirac H, $\theta^*(H^s)$ is strictly decreasing in s, the inequality follows from the fact that $\Phi(\cdot|\omega^*)$ strictly first-order stochastically dominates $\Phi(\cdot|\omega^{**})$. The second inequality holds because by the proof of Lemma A.1, $\int\int F(\theta^*(H^s))d\Phi(s|\omega^{**})d\bar{\mu}_t^{\omega^*,C}(H) < \int\int F(\theta^*(H^s))d\Phi(s|\omega^{**})d\bar{\mu}_t^{\omega^{**},C}(H)$ for all t. □

A.2 Completing the Proof

To complete the proof, we show that in all states, almost all correct agents' beliefs converge to a point-mass on the true state. The next lemma first shows this for states $\omega \in \Omega^*$.

Lemma A.5. $\bar{\mu}_\infty^{\omega,C} = \delta_{\delta_\omega}$ *for each* $\omega \in \Omega^*$.

Proof. Fix any $\omega \in \Omega^*$ and any closed interval $E := [\omega_-, \omega_+] \not\ni \omega$. Consider any correct agent and let (H_t) denote the process of his posterior beliefs in state ω when he updates beliefs based only on observing others' actions. It suffices to show that $H_t(E) \to 0$ almost surely.[59]

We focus on the case $\omega_- < \omega_+ < \omega$, as the other case $\omega < \omega_- < \omega_+$ is analogous. Without loss, we assume that $\omega_+ \in \Omega^*$, as otherwise we can expand E by selecting a point from $(\omega_+, \omega) \cap \Omega^*$ as the new upper-bound of E. Pick any $\omega^* \in \Omega^*$ such that $\omega_+ < \omega^* < \omega$ and let $E^* := [\omega^*, \omega]$.

By Lemma A.4, we have $q_\infty^C(\omega_+) > q_\infty^C(\omega^*) > q_\infty^C(\omega)$. Thus, we can pick $\gamma, \delta > 0$ sufficiently small such that for all $q_+, q^*, q \in [F(\theta^*(\overline{\omega})), F(\theta^*(\underline{\omega}))]$ with $q_+ \geq q^* \geq q$, we have that

$$\begin{aligned}&\big(\alpha q + (1-\alpha)q_\infty^C(\omega) + \gamma\big) \log\left(\frac{\alpha q_+ + (1-\alpha)q_\infty^C(\omega_+) + \gamma}{\alpha q^* + (1-\alpha)q_\infty^C(\omega^*) - \gamma}\right) \\ &+ \big(1 - \alpha q - (1-\alpha)q_\infty^C(\omega) + \gamma\big) \log\left(\frac{1 - \alpha q_+ - (1-\alpha)q_\infty^C(\omega_+) + \gamma}{1 - \alpha q^* - (1-\alpha)q_\infty^C(\omega^*) - \gamma}\right)\end{aligned}$$

is less than $-\delta$.[60]

[59] Indeed, since process (H_t) almost surely weakly converges and its limit is distributed according to $\bar{\mu}_\infty^{\omega,C}$, the Portmanteau theorem then implies that $\bar{\mu}_\infty^{\omega,C}(\{H : H(E^o) = 0\}) = 1$, where E^o denotes the interior of E in Ω. In particular, letting $\mathcal{H}_n^+ = \{H : H((\omega + \frac{1}{n}, \overline{\omega}]) > 0\}$ and $\mathcal{H}_n^- = \{H : H([\underline{\omega}, \omega - \frac{1}{n})) > 0\}$, we have $\bar{\mu}_\infty^{\omega,C}(\mathcal{H}_n^+) = \bar{\mu}_\infty^{\omega,C}(\mathcal{H}_n^-) = 0$ for each n. Thus, by countable additivity, $\bar{\mu}_\infty^{\omega,C}(\{H : H \neq \delta_\omega\}) = \bar{\mu}_\infty^{\omega,C}\big(\bigcup_n(\mathcal{H}_n^+ \cup \mathcal{H}_n^-)\big) = 0$, as required.

[60] To see this, note that for all $\beta_+, \beta^*, \beta \in (0,1)$ with $\beta_+ > \beta^* > \beta$, we have that $\beta \log\left(\frac{\beta_+}{\beta^*}\right) + (1-\beta)\log\left(\frac{1-\beta_+}{1-\beta^*}\right) < 0$. Thus, for any compact $B \subseteq (0,1)$, there exist small enough $\gamma, \delta > 0$ such that $(\beta+\gamma)\log\left(\frac{\beta_+ + \gamma}{\beta^* - \gamma}\right) + (1 - \beta + \gamma)\log\left(\frac{1-\beta_+ + \gamma}{1-\beta^* - \gamma}\right) < -\delta$ for all $\beta_+, \beta^*, \beta \in B$ with $\beta_+ > \beta^* > \beta$.

Now consider the process of log likelihood ratios $V_t := \log\left(\frac{H_t(E)}{H_t(E^*)}\right)$. For all t, we have

$$\begin{aligned}\mathbb{E}[V_{t+1} - V_t|\mathcal{I}_t, \omega] &= q_t(\omega)\log\left(\frac{q_t(E)}{q_t(E^*)}\right) + (1-q_t(\omega))\log\left(\frac{1-q_t(E)}{1-q_t(E^*)}\right)\\ &< q_t(\omega)\log\left(\frac{q_t(\omega_+)}{q_t(\omega^*)}\right) + (1-q_t(\omega))\log\left(\frac{1-q_t(\omega_+)}{1-q_t(\omega^*)}\right)\\ &= \left(\alpha q_t^I(\omega) + (1-\alpha)q_t^C(\omega)\right)\log\left(\frac{\alpha q_t^I(\omega_+) + (1-\alpha)q_t^C(\omega_+)}{\alpha q_t^I(\omega^*) + (1-\alpha)q_t^C(\omega^*)}\right)\\ &\quad + (1-\alpha q_t^I(\omega) - (1-\alpha)q_t^C(\omega))\log\left(\frac{1-\alpha q_t^I(\omega_+) - (1-\alpha)q_t^C(\omega_+)}{1-\alpha q_t^I(\omega^*) - (1-\alpha)q_t^C(\omega^*)}\right),\end{aligned}$$

where $\mathcal{I}_t$ denotes the filtration generated by the sequence of actions observed by the agent, and $q_t(E) = \frac{\int_E q_t(\omega')dH_t}{H_t(E)}$ and $q_t(E^*) = \frac{\int_E q_t(\omega')dH_t}{H_t(E^*)}$. Here the inequality holds since q_t^I, q_t^C are strictly decreasing by Lemma A.1 and H_t has full-support over Ω. Moreover, by Lemma A.4, there exists T such that $|q_t^C(\omega) - q_\infty^C(\omega)|, |q_t^C(\omega_+) - q_\infty^C(\omega_+)|, |q_t^C(\omega^*) - q_\infty^C(\omega^*)| < \gamma$ for all $t \geq T$. In addition, $q_t^I(\omega_+), q_t^I(\omega^*), q_t^I(\omega) \in [F(\theta^*(\overline{\omega})), F(\theta^*(\underline{\omega}))]$ with $q_t^I(\omega_+) > q_t^I(\omega^*) > q_t^I(\omega)$ for every t. Thus, by choice of $\gamma, \delta > 0$, the above shows that $\mathbb{E}[V_{t+1} - V_t|\mathcal{I}_t, \omega] < -\delta$ for all $t \geq T$.

Therefore, almost surely, $V_t \to -\infty$ and hence $H_t(E) \to 0$, as required. □

Lemma A.5 implies that $q_\infty^C(\omega) = F(\theta^*(\omega))$ and $\mu_t^{\omega,C}$ weakly converges to δ_{δ_ω} for any $\omega \in \Omega^*$. Now take any $\omega \in \Omega \setminus \Omega^*$. By Lemma A.1, we have the inequalities

$$F(\theta^*(\omega')) = \lim_{t\to\infty} q_t^C(\omega') \geq \limsup_{t\to\infty} q_t^C(\omega) \geq \liminf_{t\to\infty} q_t^C(\omega) \geq \lim_{t\to\infty} q_t^C(\omega'') = F(\theta^*(\omega''))$$

for any $\omega', \omega'' \in \Omega^*$ such that $\omega' < \omega < \omega''$. If $\omega \in (\underline{\omega}, \overline{\omega})$ is interior, then since $\Psi(\Omega^*) = 1$ and Ψ admits a positive density, we can choose $\omega', \omega'' \in \Omega^*$ arbitrarily close to ω. Hence, by continuity of F and θ^*, we have $q_t^C(\omega) \to F(\theta^*(\omega))$. For boundary points $\omega \in \{\underline{\omega}, \overline{\omega}\}$, the same argument shows that $\liminf_{t\to\infty} q_t^C(\underline{\omega}) \geq F(\theta^*(\underline{\omega}))$ and $\limsup_{t\to\infty} q_t^C(\overline{\omega}) \leq F(\theta^*(\overline{\omega}))$. Since $q_t^C(\cdot) \in [F(\theta^*(\overline{\omega})), F(\theta^*(\underline{\omega}))]$, this again implies $q_t^C(\underline{\omega}) \to F(\theta^*(\underline{\omega}))$ and $q_t^C(\overline{\omega}) \to F(\theta^*(\overline{\omega}))$. Thus, $q_\infty^C(\cdot) := \lim_{t\to\infty} q_t^C(\cdot) = F(\theta^*(\cdot))$ exists and is strictly decreasing on the whole of Ω. Given this, for any $\omega \in \Omega \setminus \Omega^*$, the same argument as in the proof of Lemma A.5 shows that $\bar{\mu}_\infty^{\omega;C} = \delta_{\delta_\omega}$. Hence, $\mu_t^{\omega,C}$ weakly converges to δ_{δ_ω}, completing the proof of part 1 of Proposition 1.

B Proofs of Theorems 1 and 2

We prove Theorems 1 and 2 in Sections B.3 and B.4, respectively. Both proofs follow from preliminary results on agents' long-run inferences and beliefs that we establish in Sections B.1 and B.2, in particular Proposition B.1, which shows that long-run beliefs are steady states of the limit model belief adjustment process that we considered in Section 4.2.

B.1 Agents' Long-Run Inferences

In this section, we first consider any agent whose perception is given by $\hat{F}$ and study his inferences from sequences of observed actions $a^{t-1} = (a_1, \dots, a_{t-1})$. The key result is Lemma B.2, which shows that the ***average action*** $\bar{a}^{t-1} = \frac{1}{t-1}\sum_{\tau=1}^{t-1} a_\tau$ that the agent observes up to time $t-1$ provides an "approximate" sufficient statistic for the agent's belief as $t \to \infty$.

Let $\hat{q}_t(\omega)$ denote the agent's *perceived* fraction of action 0 in each period t and state ω. Since the agent believes that $\hat{F}$ is the true type distribution and that $\hat{F}$ is common knowledge among all agents, Lemma 1 and all lemmas used in its proof in Appendix A apply to the agent's *perceptions* of behavior and beliefs in the population. In particular, Lemma 1 implies that the agent believes that in all states ω almost all agents' beliefs converge to a point mass on the true state ω. Thus, the agent believes that behavior in state ω converges to a threshold strategy according to $\theta^*(\omega)$; that is, $\hat{q}_\infty(\omega) := \lim_{t\to\infty} \hat{q}_t(\omega) = \hat{F}(\theta^*(\omega))$ for each ω. Additionally, Lemma A.1 implies that $\hat{q}_t(\cdot)$ is strictly decreasing for each t. Hence, since $\hat{F}(\theta^*(\cdot))$ is continuous on the compact interval $\Omega = [\underline{\omega}, \overline{\omega}]$, it follows that $\hat{q}_t(\cdot)$ converges to $\hat{F}(\theta^*(\cdot))$ uniformly.

Let $H_t(\cdot \mid a^{t-1}, s) \in \Delta(\Omega)$ denote the agent's posterior belief after observing private signal s and action history $a^{t-1} = (a_1, \dots, a_{t-1})$. Because $\hat{q}_t(\omega) \in (0,1)$ for each t and ω (Lemma A.1) and by the full-support assumption on private signals, $H_t(\cdot \mid a^{t-1}, s)$ has full support over Ω with positive density $h_t(\cdot \mid a^{t-1}, s)$ for all a^{t-1} and s . For each pair of states ω', ω'', denote the corresponding ***log likelihood ratio*** by

$$\ell_t(\omega', \omega'' \mid a^{t-1}, s) := \log \frac{h_t(\omega' \mid a^{t-1}, s)}{h_t(\omega'' \mid a^{t-1}, s)}.$$

Lemma B.1 below will provide a lower bound on the log likelihood ratio that depends on histories a^{t-1} only through the average action $\bar{a}^{t-1}$ and holds uniformly across all pairs of states. To state this, we first choose some $\nu^* > 0$ and $\underline{C} < 0 < \overline{C}$ such that $\hat{q}_\infty(\omega) \pm \nu^* \in (0,1)$ for all ω and such that for all $\nu \in (0, \nu^*)$,

$$\begin{aligned}\underline{C} &< \min\left\{\log\frac{1-\hat{F}(\theta^*(\underline{\omega}))-\nu}{1-\hat{F}(\theta^*(\overline{\omega}))+\nu}, \log\frac{\hat{F}(\theta^*(\overline{\omega}))-\nu}{\hat{F}(\theta^*(\underline{\omega}))+\nu}\right\} < 0 \\ &< \max\left\{\log\frac{1-\hat{F}(\theta^*(\overline{\omega}))-\nu}{1-\hat{F}(\theta^*(\underline{\omega}))+\nu}, \log\frac{\hat{F}(\theta^*(\underline{\omega}))-\nu}{\hat{F}(\theta^*(\overline{\omega}))+\nu}\right\} < \overline{C}.\end{aligned}$$

Such values exist since $\hat{q}_\infty(\omega) = \hat{F}(\theta^*(\omega)) \in [\hat{F}(\theta^*(\overline{\omega})), \hat{F}(\theta^*(\underline{\omega}))]$ for each ω and $0 < \hat{F}(\theta^*(\overline{\omega})) < \hat{F}(\theta^*(\underline{\omega})) < 1$.

Moreover, for any $\nu \in [-\nu^*, \nu^*]$, $R \in [0,1]$ and $R', R'' \in [\hat{F}(\theta^*(\overline{\omega})), \hat{F}(\theta^*(\underline{\omega}))]$, we define

$$\Delta^\nu(R, R', R'') := R \log\frac{R'-\nu}{R''+\nu} + (1-R)\log\frac{1-R'-\nu}{1-R''+\nu}.$$

Lemma B.1. *Take any $\nu \in (0, \nu^*)$. There exists some $\hat{t}$ such that for all $t \geq \hat{t}$, all $\omega' \neq \omega''$, all s,* *and all* a^{t-1},

$$\ell_t(\omega', \omega'' \mid a^{t-1}, s) > \log \frac{\phi(s \mid \omega')}{\phi(s \mid \omega'')} + (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t})\Delta^\nu(\bar{a}^{t-1}, \hat{q}_\infty(\omega'), \hat{q}_\infty(\omega'')).$$

Proof. Since $\hat{q}_t$ converges to $\hat{q}_\infty$ uniformly, we can choose $\hat{t}$ such that for all $t \geq \hat{t}$, we have $\sup_{\omega \in \Omega} |\hat{q}_t(\omega) - \hat{q}_\infty(\omega)| < \nu$. Pick any $t \geq \hat{t}$, any $\omega' \neq \omega''$, any private signal s, and any $a^{t-1} = (a_1, \ldots, a_{t-1})$. Then

$$\begin{aligned}
&\ell_t(\omega', \omega'' \mid a^{t-1}, s) - \log \frac{\phi(s \mid \omega')}{\phi(s \mid \omega'')} = \sum_{\tau=1}^{t-1} \left(a_\tau \log \frac{\hat{q}_\tau(\omega')}{\hat{q}_\tau(\omega'')} + (1 - a_\tau) \log \frac{1 - \hat{q}_\tau(\omega')}{1 - \hat{q}_\tau(\omega'')} \right) \\
&> \underline{C}(\hat{t} - 1) + \sum_{\tau=\hat{t}}^{t-1} \left(a_\tau \log \frac{\hat{q}_\tau(\omega')}{\hat{q}_\tau(\omega'')} + (1 - a_\tau) \log \frac{1 - \hat{q}_\tau(\omega')}{1 - \hat{q}_\tau(\omega'')} \right) \\
&\geq \underline{C}(\hat{t} - 1) + \sum_{\tau=\hat{t}}^{t-1} \left(a_\tau \log \frac{\hat{q}_\infty(\omega') - \nu}{\hat{q}_\infty(\omega'') + \nu} + (1 - a_\tau) \log \frac{1 - \hat{q}_\infty(\omega') - \nu}{1 - \hat{q}_\infty(\omega'') + \nu} \right) \\
&> \underline{C}(\hat{t} - 1) + \sum_{\tau=1}^{t-1} \left(a_\tau \log \frac{\hat{q}_\infty(\omega') - \nu}{\hat{q}_\infty(\omega'') + \nu} + (1 - a_\tau) \log \frac{1 - \hat{q}_\infty(\omega') - \nu}{1 - \hat{q}_\infty(\omega'') + \nu} \right) - \overline{C}(\hat{t} - 1) \\
&= (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t}) \left(\bar{a}^{t-1} \log \frac{\hat{q}_\infty(\omega') - \nu}{\hat{q}_\infty(\omega'') + \nu} + (1 - \bar{a}^{t-1}) \log \frac{1 - \hat{q}_\infty(\omega') - \nu}{1 - \hat{q}_\infty(\omega'') + \nu} \right) \\
&= (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t})\Delta^\nu(\bar{a}^{t-1}, \hat{q}_\infty(\omega'), \hat{q}_\infty(\omega'')),
\end{aligned}$$

as required. Here the first inequality holds because by choice of $\underline{C}$, we have for all $\tau = 1, \ldots, \hat{t} - 1$ that

$$a_\tau \log \frac{\hat{q}_\tau(\omega')}{\hat{q}_\tau(\omega'')} + (1 - a_\tau) \log \frac{1 - \hat{q}_\tau(\omega')}{1 - \hat{q}_\tau(\omega'')} \geq \min \left\{ \log \frac{\hat{F}(\theta^*(\overline{\omega}))}{\hat{F}(\theta^*(\underline{\omega}))}, \log \frac{1 - \hat{F}(\theta^*(\underline{\omega}))}{1 - \hat{F}(\theta^*(\overline{\omega}))} \right\} > \underline{C}.$$

The second inequality holds by choice of $\hat{t}$. The third inequality holds because by choice of $\overline{C}$, we have for all $\tau = 1, \ldots, \hat{t} - 1$ that

$$a_\tau \log \frac{\hat{q}_\infty(\omega') - \nu}{\hat{q}_\infty(\omega'') + \nu} + (1 - a_\tau) \log \frac{1 - \hat{q}_\infty(\omega') - \nu}{1 - \hat{q}_\infty(\omega'') + \nu} \leq \max \left\{ \log \frac{1 - \hat{F}(\theta^*(\overline{w})) - \nu}{1 - \hat{F}(\theta^*(\underline{w})) + \nu}, \log \frac{\hat{F}(\theta^*(\underline{w})) - \nu}{\hat{F}(\theta^*(\overline{w})) + \nu} \right\} < \overline{C}.$$

And the final two equalities hold by definition of $\bar{a}^{t-1}$ and Δ^ν. □

Using Lemma B.1, we now show that for large t, the agent's belief is approximately given by a point mass on the state $\hat{\omega} = \min_{\hat{\omega}'} \mathrm{KL}(1 - \bar{a}^{t-1}, \hat{q}_\infty(\hat{\omega}'))$ that minimizes KL-divergence between the observed empirical frequency $1 - \bar{a}^{t-1}$ of action 0 and the agent's perceived long-run fraction of action 0.

Lemma B.2. *Fix any* $\underline{s} \leq \overline{s}$ *and* $R \in (0, 1)$. *Let* $\hat{\omega} := \mathrm{argmin}_{\hat{\omega}'} \mathrm{KL}(R, \hat{q}_\infty(\hat{\omega}'))$. *Then for every* *interval $E \ni \hat{\omega}$ of states with non-empty interior, there exists $\varepsilon > 0$ such that*

$$\lim_{t\to\infty} \inf \left\{ H_t \left(E \mid a^{t-1}, s \right) : s \in [\underline{s}, \overline{s}], 1 - \bar{a}^{t-1} \in [R-\varepsilon, R+\varepsilon] \right\} = 1.$$

Proof. Note that since the agent's posterior admits a positive density, $H_t(E \mid a^{t-1}, s) = H_t(E^\circ \mid a^{t-1}, s)$ for all t, a^{t-1}, and s, where E° is the interior of the interval E. Since E° is an open interval, $E^\circ = (\alpha_1, \alpha_2)$ for some $\alpha_1 < \alpha_2$ with $\underline{\omega} \leq \alpha_1 < \alpha_2 \leq \overline{\omega}$. Let $R_1 := \hat{q}_\infty(\alpha_1)$ and $R_2 := \hat{q}_\infty(\alpha_2)$.

There are three cases to consider:

1. $\hat{\omega} \in (\alpha_1, \alpha_2)$ and $R = \hat{q}_\infty(\hat{\omega})$,

2. $\hat{\omega} = \alpha_2 = \overline{\omega}$ and $R \leq \hat{q}_\infty(\hat{\omega})$,

3. $\hat{\omega} = \alpha_1 = \underline{\omega}$ and $R \geq \hat{q}_\infty(\hat{\omega})$.

We illustrate the argument only for case 1 as it translates easily to the other cases.[61] Moreover, in case 1, we can assume that $\underline{\omega} < \alpha_1 < \hat{\omega} < \alpha_2 < \overline{\omega}$, by restricting to a subset of E if need be. Then we can choose $\xi, \varepsilon, \rho > 0$ such that $R_2 < R - \xi < R + \xi < R_1$ and

$$\rho < \inf \left\{ \Delta^0(R', R'', R''') : R' \in [R-\varepsilon, R+\varepsilon], R'' \in [R-\xi, R], R''' \geq R_1 \right\},$$
$$\rho < \inf \left\{ \Delta^0(R', R'', R''') : R' \in [R-\varepsilon, R+\varepsilon], R'' \in [R, R+\xi], R''' \leq R_2 \right\}.$$

By continuity of KL-divergence, there exists some $\nu \in (0, \nu^*)$ such that

$$\rho < \inf \left\{ \Delta^\nu(R', R'', R''') : R' \in [R-\varepsilon, R+\varepsilon], R'' \in [R-\xi, R], R''' \geq R_1 \right\},$$
$$\rho < \inf \left\{ \Delta^\nu(R', R'', R''') : R' \in [R-\varepsilon, R+\varepsilon], R'' \in [R, R+\xi], R''' \leq R_2 \right\}.$$

Take $M > 0$ such that $-M \leq \log \frac{\phi(s|\omega')}{\phi(s|\omega'')}$ for all $\omega', \omega'' \in [\underline{\omega}, \overline{\omega}]$ and all $s \in [\underline{s}, \overline{s}]$. Let $\hat{t}$ be the cutoff given by Lemma B.1. Then for all $t \geq \hat{t}$, $\omega' \in [\hat{\omega}, \hat{q}_\infty^{-1}(R-\xi)], \omega'' \in \hat{q}_\infty^{-1}([R_1, 1])$, $s \in [\underline{s}, \overline{s}]$, and any a^{t-1} such that $1 - \bar{a}^{t-1} \in [R-\varepsilon, R+\varepsilon]$, we have

$$\begin{aligned} \ell_t(\omega', \omega'' \mid a^{t-1}, s) &> \log \frac{\phi(s \mid \omega')}{\phi(s \mid \omega'')} + (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t})\Delta^\nu(1 - \bar{a}^{t-1}, \hat{q}_\infty(\omega'), \hat{q}_\infty(\omega'')) \\ &\geq -M + (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t})\rho, \end{aligned}$$

where the first inequality holds by Lemma B.1 and the second inequality holds by choice of ν and M above.

[61] In case 2, we choose ξ, ε, $\rho > 0$ such that $\hat{q}_\infty(\overline{\omega}) < R + \xi < R_1$ and

$$\rho < \inf \left\{ \Delta^0(R', R'', R''') : R' \in [R-\varepsilon, R+\varepsilon], R'' \in [R, R+\xi], R''' \geq R_1 \right\}.$$

Finally, in case 3, we choose ξ, ε, $\rho > 0$ such that $R_2 < R - \xi < \hat{q}_\infty(\underline{\omega})$ and

$$\rho < \inf \left\{ \Delta^0(R', R'', R''') : R' \in [R-\varepsilon, R+\varepsilon], R'' \in [R-\xi, R], R''' \leq R_2 \right\}.$$

The remaining steps are analogous to case 1.

Likewise, for all $t \geq \hat{t}$, $\omega' \in [\hat{q}_\infty^{-1}(R+\xi), \hat{\omega}], \omega'' \in \hat{q}_\infty^{-1}([0, R_2])$, $s \in [\underline{s}, \overline{s}]$, and any a^{t-1} such that $1 - \bar{a}^{t-1} \in [R-\varepsilon, R+\varepsilon]$, we have

$$\ell_t(\omega', \omega'' \mid a^{t-1}, s) \geq -M + (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t})\rho.$$

As a result, for all $t \geq \hat{t}$, $s \in [\underline{s}, \overline{s}]$, and any a^{t-1} such that $1 - \bar{a}^{t-1} \in [R-\varepsilon, R+\varepsilon]$, we have

$$\begin{aligned}
&H_t(E \mid a^{t-1}, s) \geq H_t([\hat{\omega}, \hat{q}_\infty^{-1}(R-\xi)] \mid a^{t-1}, s) + H_t([\hat{q}_\infty^{-1}(R+\xi), \hat{\omega}] \mid a^{t-1}, s) \\
&\geq e^{-M+(\underline{C}-\overline{C})(\hat{t}-1)+(t-\hat{t})\rho} \left(\frac{\hat{q}_\infty^{-1}(R-\xi) - \hat{\omega}}{\alpha_1 - \underline{\omega}} H_t\left([\underline{\omega}, \alpha_1] \mid a^{t-1}, s\right) + \frac{\hat{\omega} - \hat{q}_\infty^{-1}(R+\xi)}{\overline{\omega} - \alpha_2} H_t\left([\alpha_2, \overline{\omega}] \mid a^{t-1}, s\right) \right) \\
&\geq K e^{-M+(\underline{C}-\overline{C})(\hat{t}-1)+(t-\hat{t})\rho} H_t\left(\Omega \setminus E \mid a^{t-1}, s\right),
\end{aligned}$$

where the second inequality uses the bounds on log likelihood ratios we obtained above, and in the third line we let $K := \min\left\{\frac{\hat{q}_\infty^{-1}(R-\xi)-\hat{\omega}}{\alpha_1-\underline{\omega}}, \frac{\hat{\omega}-\hat{q}_\infty^{-1}(R+\xi)}{\overline{\omega}-\alpha_2}\right\}$. Since $H_t\left(\Omega \setminus E \mid a^{t-1}, s\right) = 1 - H_t(E \mid a^{t-1}, s)$, this yields

$$H_t(E \mid a^{t-1}, s) \geq \frac{K e^{-M+(\underline{C}-\overline{C})(\hat{t}-1)+(t-\hat{t})\rho}}{K e^{-M+(\underline{C}-\overline{C})(\hat{t}-1)+(t-\hat{t})\rho} + 1},$$

which completes the proof as the right-hand side converges to 1 as $t \to \infty$. □

B.2 Long-Run Beliefs Converge to Steady States

In this section, we fix arbitrary true and perceived type distributions $F, \hat{F} \in \mathcal{F}$ and analyze long-run beliefs and behavior. In each state ω, let $q_t(\omega)$ and $\hat{q}_t(\omega)$ denote the corresponding true and perceived fractions of action 0 in period t, and let $\bar{q}_t(\omega) := \frac{1}{t}\sum_{\tau=1}^{t} q_\tau(\omega)$ denote the true time average of the fraction of action 0 up to period t. Let $\Pr(\cdot \mid \omega)$ denote the probability distribution over observed private signals s and action sequences a^t when signals are distributed according to $\Phi(\cdot \mid \omega)$ and actions in each period τ are distributed according to $q_\tau(\omega)$. Define the set of ***steady states*** $\mathrm{SS}(F, \hat{F}) := \{\hat{\omega}_\infty \in \Omega : \hat{\omega}_\infty = \operatorname{argmin}_{\hat{\omega}\in\Omega} \mathrm{KL}\left(F(\theta^*(\hat{\omega}_\infty)), \hat{F}(\theta^*(\hat{\omega}))\right)\}$.

Focusing on the case when $\mathrm{SS}(F, \hat{F})$ is finite, the key result of this section is Proposition B.1, which shows that agents' long-run beliefs assign probability 1 to steady states. As a preliminary step, the following lemma shows that behavior in each state converges.

Lemma B.3. *Suppose that* $\mathrm{SS}(F, \hat{F})$ *is finite. Then* $R(\omega) := \lim_{t\to\infty} \bar{q}_t(\omega)$ *exists for every* ω *and is weakly decreasing in* ω.

Proof. Fix any $\omega \in \Omega$. Let $\bar{R}(\omega) := \limsup_{t\to\infty} \bar{q}_t(\omega)$ and $\underline{R}(\omega) := \liminf_{t\to\infty} \bar{q}_t(\omega)$. To show that $R(\omega) := \lim_{t\to\infty} \bar{q}_t(\omega)$ exists, suppose for a contradiction that $\overline{R}(\omega) > \underline{R}(\omega)$. Since $\mathrm{SS}(F, \hat{F})$ is finite, we can pick some $R \in (\underline{R}(\omega), \overline{R}(\omega))$ such that $\hat{q}_\infty^{-1}(R) \notin \mathrm{SS}(F, \hat{F})$. Let $\hat{\omega} := \operatorname{argmin}_{\hat{\omega}'\in\Omega} \mathrm{KL}(R, \hat{q}_\infty(\hat{\omega}'))$. Note that $R \neq F(\theta^*(\hat{\omega}))$, as otherwise $\hat{q}_\infty^{-1}(R)$ is a steady state. Below we assume that $R < F(\theta^*(\hat{\omega}))$, as the remaining case, $R > F(\theta^*(\hat{\omega}))$, is analogous.

Pick $\underline{R}, \overline{R}$ such that $\underline{R}(\omega) < \underline{R} < R < \overline{R} < \bar{R}(\omega)$ and $\overline{R} < F(\theta^*(\hat{\omega}))$. Note that we can choose a small enough interval $E \ni \hat{\omega}$ with non-empty interior, a large enough interval $[\underline{s}, \overline{s}]$ of private signals, and a small enough $\nu > 0$, such that for any t, if at least fraction $1-\nu$ of agents with private signals $s \in [\underline{s}, \overline{s}]$ hold beliefs such that $H_t(E \mid a^{t-1}, s) > 1-\nu$, then $q_t(\omega) > \overline{R} + \nu$.[62]

By Lemma B.2, there exists $\varepsilon > 0$ and $\hat{t}_1$ such that for all $t \geq \hat{t}_1$,

$$\inf\left\{H_t\left(E \mid a^{t-1}, s\right) \mid s \in [\underline{s}, \overline{s}], 1-\bar{a}^{t-1} \in [R-\varepsilon, R+\varepsilon]\right\} > 1-\nu.$$

Moreover, we can take ε sufficiently small such that $\underline{R} < R-\varepsilon < R+\varepsilon < \overline{R}$.

Note that each agent observes a sequence of random actions $(a_1, a_2, ...)$ that are independent from each other conditional on the realized state ω. Hence, by the weak law of large numbers for independent (but not necessarily identically distributed) random variables (see, e.g., Theorem 1.2.6 in Stroock, 2010), there exists $\hat{t}_2$ such that for all $t \geq \hat{t}_2$,

$$\Pr\left(1-\bar{a}^{t-1} \in [\bar{q}_{t-1}(\omega) - \frac{\varepsilon}{2}, \bar{q}_{t-1}(\omega) + \frac{\varepsilon}{2}] \mid \omega\right) > 1-\nu.$$

Let $\mathcal{T} := \left\{t \geq \max\{\hat{t}_1, \hat{t}_2\} : \bar{q}_{t-1}(\omega) \in [R-\frac{\varepsilon}{2}, R+\frac{\varepsilon}{2}]\right\}$. Then at all times $t \in \mathcal{T}$, at least fraction $1-\nu$ of agents with private signals $s \in [\underline{s}, \overline{s}]$ hold beliefs such that $H_t(E \mid a^{t-1}, s) > 1-\nu$. Thus, for all $t \in \mathcal{T}$, we have $q_t(\omega) > \bar{R} + \nu$.

Since $\bar{R}(\omega) = \limsup_t \bar{q}_t(\omega) > R + \frac{\varepsilon}{2}$ and $\underline{R}(\omega) = \liminf_t \bar{q}_t(\omega) < R - \frac{\varepsilon}{2}$ and $|\bar{q}_t(\omega) - \bar{q}_{t-1}(\omega)| \leq \frac{1}{t} < \varepsilon$ for all large enough t, we must have an infinite sequence of times $t_k \in \mathcal{T}$ such that

$$\bar{q}_{t_k-1}(\omega) \geq R - \frac{\varepsilon}{2} > \bar{q}_{t_k}(\omega).$$

But then, by definition of $\bar{q}_t$, we have $\bar{q}_{t_k}(\omega) = \frac{t_k-1}{t_k}\bar{q}_{t_k-1}(\omega) + \frac{1}{t_k}q_{t_k}(\omega) > R - \frac{\varepsilon}{2}$, since $q_{t_k}(\omega) > \overline{R} + \nu > R - \varepsilon$ by construction of $\mathcal{T}$. This is a contradiction. Hence, $R(\omega) = \lim_{t\to\infty} \bar{q}_t(\omega)$ exists.

Finally, recall that Lemma A.1 applies to agents' perceived fraction $\hat{q}_t(\omega)$ of action 0 and implies that $\hat{q}_t(\omega)$ is strictly decreasing in ω at each t. Given this, a similar inductive argument as in the proof of Lemma A.1 yields that the true action 0 fraction $q_t(\omega)$ is strictly decreasing in ω for each t. This implies that $R(\omega) = \lim_{t\to\infty} \bar{q}_t(\omega)$ is weakly decreasing in ω, as required. □

We now prove that agents' long-run beliefs assign probability 1 to steady states:

Proposition B.1. *Suppose that* $\mathrm{SS}(F, \hat{F})$ *is finite. Then in all states* ω*, there exists some state* $\hat{\omega}_\infty(\omega) \in \mathrm{SS}(F, \hat{F})$ *such that almost all agents' beliefs converge to a point-mass on* $\hat{\omega}_\infty(\omega)$*. Moreover,* $\hat{\omega}_\infty(\omega)$ *is weakly increasing in* ω*.*

Proof. Fix any $\omega \in \Omega$ and let $R(\omega) := \lim_{t\to\infty} \bar{q}_t(\omega)$, which exists by Lemma B.3. Define

$$\hat{\omega}_\infty(\omega) := \operatorname*{argmin}_{\hat{\omega}\in\Omega} \mathrm{KL}(R(\omega), \hat{q}_\infty(\hat{\omega})).$$

Note that $\hat{\omega}_\infty(\omega)$ is weakly increasing in ω since $R(\omega)$ is weakly decreasing and $\hat{q}_\infty$ is strictly decreasing.

[62] To see this, observe that if all agents' beliefs assigned probability 1 to $\hat{\omega}$ at t, then $q_t(\omega) = F(\theta^*(\hat{\omega})) > \overline{R}$.

Consider any interval $E \ni \hat{\omega}_\infty(\omega)$ with non-empty interior and any $\underline{s} < \overline{s}$. By Lemma B.2, there exists $\varepsilon > 0$ such that

$$\lim_{t\to\infty} \inf \left\{ H_t \left(E \mid a^{t-1}, s \right) \mid s \in [\underline{s}, \overline{s}], 1 - \bar{a}^{t-1} \in [R(\omega) - \varepsilon, R(\omega) + \varepsilon] \right\} = 1.$$

As in Lemma B.3, the weak law of large numbers ensures that

$$\lim_{t\to\infty} \Pr \left(1 - \bar{a}^{t-1} \in [R(\omega) - \varepsilon, R(\omega) + \varepsilon] \mid \omega \right) = 1.$$

Hence, for every $\nu > 0$, we have

$$\lim_{t\to\infty} \Pr \left(H_t \left(E \mid a^{t-1}, s \right) > 1 - \nu \mid \omega \right) \geq \Phi([\underline{s}, \overline{s}] \mid \omega).$$

Since $\underline{s}$ and $\overline{s}$ are arbitrary, the above implies that for every $\nu > 0$ and $E \ni \hat{\omega}_\infty(\omega)$ with non-empty interior,

$$\lim_{t\to\infty} \Pr \left(H_t \left(E \mid a^{t-1}, s \right) > 1 - \nu \mid \omega \right) = 1.$$

Thus, conditional on state ω, almost all agents' beliefs converge to a point-mass on $\hat{\omega}_\infty(\omega)$. But then, $q_t(\omega) \to F(\theta^*(\hat{\omega}_\infty(\omega)))$ as $t \to \infty$, which implies that $R(\omega) = F(\theta^*(\hat{\omega}_\infty(\omega)))$. Since $\hat{q}_\infty(\cdot) = \hat{F}(\theta^*(\cdot))$, this yields $\hat{\omega}_\infty(\omega) = \operatorname{argmin}_{\hat{\omega}\in\Omega} \mathrm{KL}(F(\theta^*(\hat{\omega}_\infty(\omega))), \hat{F}(\theta^*(\hat{\omega})))$; that is, $\hat{\omega}_\infty(\omega) \in \mathrm{SS}(F, \hat{F})$. □

B.3 Proof of Theorem 1

Fix any $F \in \mathcal{F}$, $\hat{\omega} \in \Omega$, and $\varepsilon > 0$. We can pick $\hat{F} \in \mathcal{F}$ such that $\hat{F}$ crosses F from below in the single point $\theta^*(\hat{\omega})$, as shown in Figure 2; clearly, we can also require that $\|F - \hat{F}\| < \varepsilon$. In this case, $\mathrm{SS}(F, \hat{F}) = \{\hat{\omega}\}$. Thus, Proposition B.1 implies that in all states ω, almost all agents' beliefs converge to a point-mass on $\hat{\omega}$. □

B.4 Proof of Theorem 2

Fix any analytic $F, \hat{F} \in \mathcal{F}$ with $\hat{F} \neq F$. Then the set $\{\theta \in [\theta^*(\overline{\omega}), \theta^*(\underline{\omega})] : F(\theta) = \hat{F}(\theta)\}$ is finite (possibly empty).[63] But this implies that $\mathrm{SS}(F, \hat{F})$ is finite, as every $\hat{\omega}_\infty \in \mathrm{SS}(F, \hat{F})$ satisfies either $F(\theta^*(\hat{\omega}_\infty)) = \hat{F}(\theta^*(\hat{\omega}_\infty))$ or $\hat{\omega}_\infty \in \{\underline{\omega}, \overline{\omega}\}$. Thus, Proposition B.1 implies that in every state ω, almost all agents' beliefs converge to a point-mass on a state $\hat{\omega}_\infty(\omega) \in \Omega$, where the mapping $\omega \mapsto \hat{\omega}_\infty(\omega)$ is weakly increasing and has finite range. □

[63] To see this, suppose for a contradiction that $F - \hat{F} = 0$ admits an infinite sequence $\theta_1, \theta_2, \ldots$ of distinct solutions in $\Theta^* := [\theta^*(\overline{\omega}), \theta^*(\underline{\omega})]$. By sequential compactness of Θ^*, restricting to a subsequence if necessary, we can assume that the sequence converges. Then since $F - \hat{F}$ is analytic on $\mathbb{R}$, the principle of permanence implies that $F - \hat{F}$ is identically zero on $\mathbb{R}$, contradicting $\hat{F} \neq F$.

C Proof of Theorem 3

C.1 Theorem 3: Proof of Part 1

Throughout, we fix some $\Omega_n = \{\omega^1, \ldots, \omega^n\} \subseteq \Omega$. Up to relabeling, we can assume that $\omega^1 < \omega^2 < \cdots < \omega^n$. As before, given any true and perceived type distributions $F, \hat{F} \in \mathcal{F}$ and state $\omega \in \Omega_n$, let $q_t(\omega, F, \hat{F})$ denote the true fraction of action 0 in the population at time t, and let $\bar{q}_t(\omega, F, \hat{F}) = \sum_{\tau=1}^{t} q_\tau(\omega, F, \hat{F})$ denote the time average of $q_\tau(\omega; F, \hat{F})$ in periods $\tau = 1, \ldots, t$. Let $\Pr(\cdot \mid \omega, F, \hat{F})$ denote the probability distribution over observed private signals s and action sequences a^t when signals are distributed according to $\Phi(\cdot|\omega)$ and actions in each period τ are distributed according to $q_\tau(\omega, F, \hat{F})$.

In the correctly specified case, $F = \hat{F}$, we omit the dependency on $\hat{F}$ and denote the corresponding quantities by $q_t(\omega, F)$, $\bar{q}_t(\omega, F)$, and $\Pr(\cdot \mid \omega, F)$. Note that under perception $\hat{F}$, $q_t(\omega, \hat{F})$ represents agents' *perceived* fraction of action 0 at time t, as agents believe they are correctly specified; likewise, $\bar{q}_t(\omega, \hat{F})$ denotes agents' perceived time average of action 0 under $\hat{F}$ and $\Pr(\cdot \mid \omega, \hat{F})$ denotes agents' perceived distribution over actions and signals.

Let $H_t(\omega|s, a^{t-1}, \hat{F})$ denote the probability that an agent with perception $\hat{F}$ assigns to state $\omega \in \Omega_n$ following any private signal s and observed action sequence $a^{t-1} = (a_1, \ldots, a_{t-1})$. Finally, for any $F \in \mathcal{F}$ and $\xi > 0$, we let $B(F, \xi) := \{G \in \mathcal{F} : \|G - F\| < \xi\}$ denote the ball of type distributions in $\mathcal{F}$ that are within total variation distance ξ of F.

C.1.1 Preliminary Results

To prove the first part of Theorem 3, we first derive three results that (unlike the results in Appendix B.1–B.2) will enable us to analyze long-run inferences and behavior *uniformly* across all small misperceptions $\hat{F}$.

First, Proposition C.1 shows that whenever $\hat{F}$ is sufficiently close to F, then in every state ω^k agents' perceived behavior $q_t(\omega^k, \hat{F})$ converges at a uniform rate to a small ball around the long-run behavior $F(\theta^*(\omega^k))$ under the correctly specified model.

Proposition C.1. *Fix any $F \in \mathcal{F}$. There exists $\xi^* > 0$ such that for every $\omega^k \in \Omega_n$ and every $\xi \leq \xi^*$,*

$$\lim_{t \to \infty} \sup_{\hat{F} \in B(F, \xi)} |q_t(\omega^k, \hat{F}) - F(\theta^*(\omega^k))| \leq \xi.$$

Proof. See Supplementary Appendix D. □

Second, using Proposition C.1, the following lemma shows that agents' beliefs converge uniformly to a point-mass on the true state whenever their amount of misperception is sufficiently small and their observed empirical action frequency is sufficiently close to long-run behavior under the correctly specified model.

Lemma C.1. *Fix any $F \in \mathcal{F}$ and $\omega^k \in \Omega_n$. There exists $\xi > 0$ and $\varepsilon > 0$ such that for any $\underline{s} \leq \overline{s}$,*

$$\lim_{t\to\infty} \inf \left\{ H_t(\omega^k | s, a^{t-1}, \hat{F}) : s \in [\underline{s}, \overline{s}], 1 - \bar{a}^{t-1} \in [F(\theta^*(\omega^k)) - \varepsilon, F(\theta^*(\omega^k)) + \varepsilon], \hat{F} \in B(F, \xi) \right\} = 1.$$

Proof. For any $\omega^\ell \in \Omega_n$, we write $F^\ell := F(\theta^*(\omega^\ell))$. It suffices to prove that for all $\ell \neq k$, we can find $\xi, \varepsilon > 0$ such that for any $\underline{s} \leq \overline{s}$,

$$\lim_{t\to\infty} \sup \left\{ \log \frac{H_t(\omega^\ell | s, a^{t-1}, \hat{F})}{H_t(\omega^k | s, a^{t-1}, \hat{F})} : s \in [\underline{s}, \overline{s}], 1 - \bar{a}^{t-1} \in [F^k - \varepsilon, F^k + \varepsilon], \hat{F} \in B(F, \xi) \right\} = -\infty.$$

We only consider the case $\ell > k$, as the argument for $\ell < k$ is exactly symmetric. In this case $\omega^\ell > \omega^k$, so that $F^\ell < F^k$ and hence

$$(1 - F^k) \log \frac{1 - F^\ell}{1 - F^k} + F^k \log \frac{F^\ell}{F^k} = -\mathrm{KL}(F^k, F^\ell) < 0.$$

Thus, we can find $\varepsilon > 0$ sufficiently small such that $F^\ell + \varepsilon < F^k - \varepsilon$ and such that for all $R \in [F^k - \varepsilon, F^k + \varepsilon]$, we have

$$(1 - R) \log \frac{1 - F^\ell + \varepsilon}{1 - F^k - \varepsilon} + R \log \frac{F^\ell + \varepsilon}{F^k - \varepsilon} < -\varepsilon. \tag{5}$$

By Proposition C.1, there exists $\xi^* > 0$ such that for all $m = 1, 2, \ldots, n$ and all $\xi \leq \xi^*$,

$$\lim_{t\to\infty} \sup_{\hat{F} \in B(F,\xi)} |q_t(\omega^m, \hat{F}) - F^m| \leq \xi. \tag{6}$$

We pick $\xi \in (0, \xi^*)$ sufficiently small such that $\xi < \varepsilon$ and $\min\left\{F^n, 1 - F^1\right\} - \xi > 0$. The latter implies that for all t and $\hat{F} \in B(F, \xi)$, we have $M := \max\left\{\log \frac{1}{F^n - \xi}, \log \frac{1}{1 - F^1 - \xi}\right\} < \infty$ and

$$\max\left\{ \log \frac{q_t(\omega^\ell, \hat{F})}{q_t(\omega^k, \hat{F})}, \log \frac{1 - q_t(\omega^\ell, \hat{F})}{1 - q_t(\omega^k, \hat{F})} \right\} \leq \max\left\{ \log \frac{1}{\hat{F}(\theta^*(\omega^n))}, \log \frac{1}{1 - \hat{F}(\theta^*(\omega^1))} \right\} \leq M. \tag{7}$$

Moreover, $\xi < \varepsilon$ together with (6) yields a t^* such that for all $t \geq t^*$ and $m = 1, 2, \ldots, n$,

$$\sup_{\hat{F} \in B(F,\xi)} |q_t(\omega^m, \hat{F}) - F^m| \leq \varepsilon. \tag{8}$$

Then for all $t \geq t^*$, $\hat{F} \in B(F, \xi)$, $a^{t-1} = (a_1, \dots, a_{t-1})$ with $1 - \bar{a}^{t-1} \in [F^k - \varepsilon, F^k + \varepsilon]$, and s,

$$\begin{aligned}
&\log \frac{H_t(\omega^\ell | s, a^{t-1}, \hat{F})}{H_t(\omega^k | s, a^{t-1}, \hat{F})} - \log \frac{\phi(s \mid \omega^\ell)}{\phi(s \mid \omega^k)} = \sum_{\tau=1}^{t-1} \left(a_\tau \log \frac{1 - q_\tau(\omega^\ell, \hat{F})}{1 - q_\tau(\omega^k, \hat{F})} + (1 - a_\tau) \log \frac{q_\tau(\omega^\ell, \hat{F})}{q_\tau(\omega^k, \hat{F})} \right) \\
&\leq M t^* + \sum_{\tau=t^*+1}^{t-1} \left(a_\tau \log \frac{1 - q_\tau(\omega^\ell, \hat{F})}{1 - q_\tau(\omega^k, \hat{F})} + (1 - a_\tau) \log \frac{q_\tau(\omega^\ell, \hat{F})}{q_\tau(\omega^k, \hat{F})} \right) \\
&\leq M t^* + (t-1)\bar{a}^{t-1} \log \frac{1 - F^\ell + \varepsilon}{1 - F^k - \varepsilon} + \left((t-1)(1 - \bar{a}^{t-1}) - t^*\right) \log \frac{F^\ell + \varepsilon}{F^k - \varepsilon} \\
&= \left(M - \log \frac{F^\ell + \varepsilon}{F^k - \varepsilon} \right) t^* + (t-1) \left(\bar{a}^{t-1} \log \frac{1 - F^\ell + \varepsilon}{1 - F^k - \varepsilon} + (1 - \bar{a}^{t-1}) \log \frac{F^\ell + \varepsilon}{F^k - \varepsilon} \right) \\
&\leq \left(M - \log \frac{F^\ell + \varepsilon}{F^k - \varepsilon} \right) t^* - (t-1)\varepsilon,
\end{aligned}$$

where the first inequality holds by (7), the second inequality follows from (8) and the fact that $F^\ell + \varepsilon < F^k - \varepsilon$, and the final inequality holds by (5) and the fact that $1 - \bar{a}^{t-1} \in [F^k - \varepsilon, F^k + \varepsilon]$.

Thus, for all $\underline{s} \leq \overline{s}$, we have

$$\begin{aligned}
&\lim_{t \to \infty} \sup \left\{ \log \frac{H_t(\omega^\ell | s, a^{t-1}, \hat{F})}{H_t(\omega^k | s, a^{t-1}, \hat{F})} : s \in [\underline{s}, \overline{s}], 1 - \bar{a}^{t-1} \in [F^k - \varepsilon, F^k + \varepsilon], \hat{F} \in B(F, \xi) \right\} \\
&\leq \sup_{s \in [\underline{s}, \overline{s}]} \log \frac{\phi(s \mid \omega^\ell)}{\phi(s \mid \omega^k)} + \left(M - \log \frac{F^\ell + \varepsilon}{F^k - \varepsilon} \right) t^* - \lim_{t \to \infty} (t-1)\varepsilon = -\infty,
\end{aligned}$$

as required. □

Finally, at any given time t, the following result provides a lower bound on the probability that an agent's empirical frequency of observing action 0 is close to the true time average of action 0, where this bound holds uniformly across all F and $\hat{F}$.

Lemma C.2. *Fix any $F, \hat{F} \in \mathcal{F}$ and $\omega \in \Omega_n$. Then for all $\rho > 0$ and t,*

$$\Pr \left(|1 - \bar{a}^t - \bar{q}_t(\omega, F, \hat{F})| < \rho \mid \omega, F, \hat{F} \right) \geq 1 - \frac{1}{\rho^2 (t-1)}.$$

Proof. Let σ_t^2 be the variance of the random variable $\bar{a}^t$. By Chebyshev's inequality,

$$\Pr \left(|1 - \bar{a}^t - \bar{q}_t(\omega, F, \hat{F})| \geq \rho \mid \omega, F, \hat{F} \right) \leq \frac{\sigma_t^2}{\rho^2}.$$

Note that $\sigma_t^2 \leq \frac{1}{t}$ since $a_1, \dots, a_t \in \{0, 1\}$ are independent conditional on ω. Thus,

$$\Pr \left(|1 - \bar{a}^t - \bar{q}_t(\omega, F, \hat{F})| < \rho \mid \omega, F, \hat{F} \right) = 1 - \Pr \left(|1 - \bar{a}^t - \bar{q}_t(\omega, F, \hat{F})| \geq \rho \mid \omega, F, \hat{F} \right) \geq 1 - \frac{1}{\rho^2 t},$$

as claimed. □

C.1.2 Completing the Proof

We now complete the proof of the first part of Theorem 3. Fix any $F \in \mathcal{F}$ and $\omega^k \in \Omega_n$. We will show that there exists $\xi_k > 0$ such that for all $\hat{F} \in B(F, \xi_k)$ and all $\lambda \in (0,1)$, we have

$$\lim_{t\to\infty} \Pr\left(H_t(\omega^k|s, a^{t-1}; \hat{F}) \geq \lambda|\omega^k, F, \hat{F}\right) \geq \lambda. \tag{9}$$

Then setting $\varepsilon_n := \min_{k=1}^n \xi_k$ yields the desired conclusion. We proceed in four steps.

Step 1. By Lemma C.1, there exists $\xi, \varepsilon > 0$ such that for all $\underline{s} \leq \overline{s}$,

$$\lim_{t\to\infty} \inf \left\{H_t(\omega^k|s, a^{t-1}, \hat{F}) : s \in [\underline{s}, \overline{s}], 1 - \bar{a}^{t-1} \in [\underline{R}, \overline{R}], \hat{F} \in B(F, \xi)\right\} = 1, \tag{10}$$

where $\underline{R} := F(\theta^*(\omega^k)) - \varepsilon$ and $\overline{R} := F(\theta^*(\omega^k)) + \varepsilon$.

Step 2. Next, we pick $\kappa \in (0,1)$ sufficiently large and $\rho \in (0,1)$ sufficiently small such that in any period t, if at least fraction $1-\rho$ of agents assign probability at least κ to state ω^k, then the fraction of agents playing action 0 is in $[\underline{R}+\rho, \overline{R}-\rho]$. To see that such κ and ρ exist, let $\gamma^0(\kappa')$ (respectively, $\gamma^1(\kappa')$) denote the share of types that strictly prefer to play action 0 (respectively, action 1) whenever their posterior assigns probability at least κ' to state ω^k.[64] Then $\gamma^0(\kappa') \leq F(\theta^*(\omega^k)) \leq 1 - \gamma^1(\kappa')$ for all $\kappa' \in [0,1]$, with $\gamma^0(1) = F(\theta^*(\omega^k)) = 1 - \gamma^1(1)$. Thus, since γ^0 and γ^1 are continuous and weakly increasing and $F(\theta^*(\omega^k)) \in [\underline{R}, \overline{R}]$, there exists $\kappa, \rho \in (0,1)$ such that $\underline{R} + \rho < (1-\rho')\gamma^0(\kappa') \leq 1 - \gamma^1(\kappa') + \rho' < \overline{R} - \rho$ for all $\kappa' \geq \kappa$ and $\rho' \leq \rho$. Then κ and ρ are as required, since in any period t, the fraction of agents playing action 0 when fraction at least $(1-\rho')$ of agents assign probability at least κ' to state ω^k is bounded below (resp. above) by $(1-\rho')\gamma^0(\kappa')$ (resp. $1 - \gamma^1(\kappa') + \rho'$).

Step 3. We claim that there exist t_1 and $\xi_k < \xi$ such that for all $t \geq t_1$ and $\hat{F} \in B(F, \xi_k)$, we have $\bar{q}_t(\omega^k, F, \hat{F}) \in [\underline{R}+\rho, \overline{R}-\rho]$. To see this, note first that under the correctly specified model, Proposition C.1 implies that $q_t(\omega^k, F) \to F(\theta^*(\omega^k))$. Thus, since $F(\theta^*(\omega^k)) \in (\underline{R}+\rho, \overline{R}-\rho)$, there exists t_1 such that $q_t(\omega^k, F) \in (\underline{R}+\rho, \overline{R}-\rho)$ for all $t \geq t_1$. Fix any $\underline{s} \leq \overline{s}$ with $\Phi([\underline{s}, \overline{s}] \mid \omega^k) > (1-\rho)$. By picking t_1 sufficiently large, we can additionally assume by (10) that for all $t \geq t_1$, we have

$$\inf \left\{H_t(\omega^k|s, a^{t-1}, \hat{F}) : s \in [\underline{s}, \overline{s}], 1 - \bar{a}^{t-1} \in [\underline{R}, \overline{R}], \hat{F} \in B(F, \xi)\right\} \geq \kappa \tag{11}$$

and $\left(1 - \frac{1}{\rho^2 t}\right) \Phi([\underline{s}, \overline{s}] \mid \omega^k) \geq (1-\rho)$. Since $\bar{q}_{t_1}(\omega^k, F) \in (\underline{R}+\rho, \overline{R}-\rho)$, a continuity argument at t_1 implies that there exists $\xi_k < \xi$ such that for all $\hat{F} \in B(F, \xi_k)$, we also have $\bar{q}_{t_1}(\omega^k, F, \hat{F}) \in (\underline{R}+\rho, \overline{R}-\rho)$. Fix any $\hat{F} \in B(F, \xi_k)$. We prove by induction that for all $t \geq t_1$, $\bar{q}_t(\omega^k, F, \hat{F})$ remains in $[\underline{R}+\rho, \overline{R}-\rho]$. Indeed, suppose this is true for a particular $t' \geq t_1$. We will show that $\bar{q}_{t'+1}(\omega^k, F, \hat{F}) \in [\underline{R}+\rho, \overline{R}-\rho]$.

Since $\bar{q}_{t'+1}(\omega^k, F, \hat{F}) = \frac{t'}{t'+1}\bar{q}_{t'}(\omega^k, F, \hat{F}) + \frac{1}{t'+1}q_{t'+1}(\omega^k, F, \hat{F})$ it suffices to show that $q_{t'+1}(\omega^k, F, \hat{F}) \in [\underline{R}+\rho, \overline{R}-\rho]$. By Step 2, this holds as long as we can show that at $t'+1$, fraction at least $(1-\rho)$

[64] Formally, let $\gamma^0(\kappa) := F(\underline{\theta}_\kappa)$, where $\underline{\theta}_\kappa$ satisfies $\kappa u(\underline{\theta}_\kappa, \omega^k) + (1-\kappa)u(\underline{\theta}_\kappa, \omega^n) = 0$, and let $\gamma^1(\kappa) := 1 - F(\overline{\theta}_\kappa)$, where $\overline{\theta}_\kappa$ satisfies $\kappa u(\overline{\theta}_\kappa, \omega^k) + (1-\kappa)u(\overline{\theta}_\kappa, \omega^1) = 0$.

of agents assign probability at least κ to state ω^k. But note that

$$\Pr\left(H_{t'+1}(\omega^k|s,a^{t'},\hat{F}) \geq \kappa \mid \omega^k, F, \hat{F}\right) \geq \Pr\left(1-\bar{a}^{t'} \in [\underline{R},\overline{R}] \mid \omega^k, F, \hat{F}\right) \Phi([\underline{s},\overline{s}] \mid \omega^k)$$
$$\geq \Pr\left(|1-a^{t'}-\bar{q}_{t'}(\omega^k,F,\hat{F})| < \rho \mid \omega^k, F, \hat{F}\right) \Phi([\underline{s},\overline{s}] \mid \omega^k) \geq \left(1-\frac{1}{\rho^2(t'-1)}\right) \Phi([\underline{s},\overline{s}] \mid \omega^k),$$

where the first inequality holds by (11), the second inequality holds since $\bar{q}_{t'}(\omega^k,F,\hat{F}) \in [\underline{R}+\rho,\overline{R}-\rho]$ by inductive hypothesis, and the third inequality holds by Lemma C.2. Since by choice of t_1, we have $\left(1-\frac{1}{\rho^2 t'}\right)\Phi([\underline{s},\overline{s}] \mid \omega^k) \geq 1-\rho$, this yields the desired conclusion.

Step 4. Finally, fix any $\lambda \in (0,1)$, $\hat{F} \in B(F,\xi_k)$, and $\underline{s} \leq \overline{s}$ such that $\Phi([\underline{s},\overline{s}]|\omega^k) \geq \lambda$. By (10) and since $\xi_k < \xi$, there exists $t_2 \geq t_1$ such that for all $t \geq t_2$, we have

$$\inf\left\{H_t(\omega^k|s,a^{t-1},\hat{F}) : s \in [\underline{s},\overline{s}], 1-\bar{a}^{t-1} \in [\underline{R},\overline{R}]\right\} \geq \lambda.$$

Then for all $t \geq t_2$,

$$\Pr\left(H_t(\omega^k|s,a^{t-1},\hat{F}) \geq \lambda \mid \omega^k\, F, \hat{F}\right) \geq \Pr\left(1-\bar{a}^{t-1} \in [\underline{R},\overline{R}] \mid \omega^k, F, \hat{F}\right) \Phi([\underline{s},\overline{s}]|\omega^k)$$
$$\geq \Pr\left(|1-\bar{a}^t-\bar{q}_t(\omega^k,F,\hat{F})| < \rho \mid \omega^k, F, \hat{F}\right) \Phi([\underline{s},\overline{s}]|\omega^k) \geq \left(1-\frac{1}{\rho^2 t}\right)\lambda,$$

where the first inequality holds by choice of t_2, the second inequality holds since $\bar{q}_t(\omega^k,F,\hat{F}) \in [\underline{R}+\rho,\overline{R}-\rho]$ by Step 3, and the third inequality holds by Lemma C.2 and choice of $[\underline{s},\overline{s}]$. This yields (9), as required. □

C.2 Theorem 3: Proof of Part 2

In each state space Ω_n, define the corresponding set of steady states by

$$\mathrm{SS}_n(F,\hat{F}) := \{\hat{\omega}_\infty \in \Omega_n : \hat{\omega}_\infty = \operatorname*{argmin}_{\hat{\omega}\in\Omega_n} \mathrm{KL}\left(F(\theta^*(\hat{\omega}_\infty)), \hat{F}(\theta^*(\hat{\omega}))\right)\}.$$

The proofs in Appendix B.1-B.2 do not rely on the fact that the state space is continuous and the same arguments go through under finite state spaces. In particular, Proposition B.1 remains valid for each Ω_n and implies that in every state $\omega \in \Omega_n$, there exists some state $\hat{\omega}_\infty(\omega) \in \mathrm{SS}_n(F,\hat{F})$ such that almost all agents' beliefs converge to a point-mass on $\hat{\omega}_\infty(\omega)$.

To prove the second part of Theorem 3, fix any $F \in \mathcal{F}$, $\hat{\omega} \in \Omega$, and $\varepsilon > 0$. Take any perceived type distribution $\hat{F} \in \mathcal{F}$ such that $\|F-\hat{F}\| < \varepsilon$ and $\hat{F}-F$ is strictly increasing with $\hat{F}(\theta^*(\hat{\omega})) - F(\theta^*(\hat{\omega})) = 0$. Then $\kappa := \min_{\omega\in\Omega\setminus[\hat{\omega}-\varepsilon,\hat{\omega}+\varepsilon]} |F(\theta^*(\omega)) - \hat{F}(\theta^*(\omega))|$ satisfies $\kappa > 0$.

Note that $\hat{F}(\theta^*(\omega))$ is uniformly continuous in ω by the compactness of Ω. Thus, there exists $\gamma > 0$ such that for any $\omega',\omega'' \in \Omega$ with $|\omega'-\omega''| < \gamma$, we have $|\hat{F}(\theta^*(\omega')) - \hat{F}(\theta^*(\omega''))| < \kappa$. Moreover, since $\{\omega^1,\omega^2,\ldots\}$ is dense in Ω, we can pick N large enough such that any interval in Ω of length γ contains at least one state from $\Omega_N = \{\omega^1,...,\omega^N\}$.

Consider any state space Ω_n with $n \geq N$. We claim that $\mathrm{SS}_n(F, \hat{F}) \subseteq [\hat{\omega} - \varepsilon, \hat{\omega} + \varepsilon]$. Indeed, consider any $\hat{\omega}_\infty \in \Omega_n \setminus [\hat{\omega} - \varepsilon, \hat{\omega} + \varepsilon]$. We focus on the case $\hat{\omega}_\infty < \hat{\omega} - \varepsilon$, as the case $\hat{\omega}_\infty > \hat{\omega} + \varepsilon$ is analogous. By construction, $F(\theta^*(\omega_\infty)) - \hat{F}(\theta^*(\omega_\infty)) \geq \kappa$. Moreover, since $n \geq N$, there exists some state $\omega' \in (\hat{\omega}_\infty, \hat{\omega}_\infty + \gamma] \cap \Omega_n$. By choice of γ, this yields $F(\theta^*(\hat{\omega}_\infty)) > \hat{F}(\theta^*(\omega')) > \hat{F}(\theta^*(\hat{\omega}_\infty))$, whence $\hat{\omega}_\infty \notin \mathrm{SS}_n(F, \hat{F})$. Thus, Proposition B.1 implies that in any state $\omega \in \Omega_n$, almost all agents' beliefs converge to a point-mass on some state $\hat{\omega}_\infty(\omega) \in \mathrm{SS}_n(F, \hat{F}) \subseteq [\hat{\omega} - \varepsilon, \hat{\omega} + \varepsilon]$, as claimed. □

References

Acemoglu, D., V. Chernozhukov, and M. Yildiz (2016): "Fragility of asymptotic agreement under Bayesian learning," *Theoretical Economics*, 11(1), 187–225.

Ahler, D. J. (2014): "Self-fulfilling misperceptions of public polarization," *The Journal of Politics*, 76(3), 607–620.

Ali, S. N. (2018): "Herding with costly information," *Journal of Economic Theory*, 175, 713–729.

Amador, M., and P.-O. Weill (2012): "Learning from private and public observations of others' actions," *Journal of Economic Theory*, 147(3), 910–940.

Banerjee, A. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107(3), 797–817.

Banerjee, A., and D. Fudenberg (2004): "Word-of-mouth learning," *Games and Economic Behavior*, 46(1), 1–22.

Berk, R. H. (1966): "Limiting behavior of posterior distributions when the model is incorrect," *The Annals of Mathematical Statistics*, pp. 51–58.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 100(5), 992–1026.

Blouin, M. R., and R. Serrano (2001): "A decentralized market with common values uncertainty: Non-steady states," *The Review of Economic Studies*, 68(2), 323–346.

Bohren, J. A. (2016): "Informational herding with model misspecification," *Journal of Economic Theory*, 163, 222–247.

Bohren, J. A., and D. Hauser (2018): "Social Learning with Model Misspecification: A Framework and a Robustness Result," working paper.

Bohren, J. A., A. Imas, and M. Rosenberg (2018): "The Dynamics of Discrimination: Theory and Evidence," working paper.

Box, G. E. (1979): "Robustness in the strategy of scientific model building," in *Robustness in statistics*, pp. 201–236. Elsevier.

Bullock, J. G., A. S. Gerber, S. J. Hill, and G. A. Huber (2013): "Partisan bias in factual beliefs about politics," Discussion paper, National Bureau of Economic Research.

Burguet, R., and X. Vives (2000): "Social learning and costly information acquisition," *Economic Theory*, 15(1), 185–205.

Chamley, C. (2004): *Rational herds: Economic models of social learning*. Cambridge University Press.

Crawford, V. P., and J. Sobel (1982): "Strategic information transmission," *Econometrica: Journal of the Econometric Society*, pp. 1431–1451.

Dasaratha, K., and K. He (2017): "Network Structure and Naive Sequential Learning," .

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990): "Noise trader risk in financial markets," *Journal of Political Economy*, 98(4), 703–738.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003): "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly journal of economics*, 118(3), 909–968.

Duffie, D., G. Giroux, and G. Manso (2010): "Information percolation," *American Economic Journal: Microeconomics*, 2(1), 100–111.

Duffie, D., S. Malamud, and G. Manso (2009): "Information percolation with equilibrium search dynamics," *Econometrica*, 77(5), 1513–1574.

——— (2010): "The Relative Contributions of Private Information Sharing and Public Information Releases to Information Aggregation," *Journal of Economic Theory*, 145, 1574–1601.

Duffie, D., and G. Manso (2007): "Information percolation in large markets," *American Economic Review*, 97(2), 203–209.

Duffie, D., and Y. Sun (2012): "The exact law of large numbers for independent random matching," *Journal of Economic Theory*, 147(3), 1105–1139.

Ellison, G., and D. Fudenberg (1993): "Rules of thumb for social learning," *Journal of political Economy*, 101(4), 612–643.

——— (1995): "Word-of-mouth communication and social learning," *The Quarterly Journal of Economics*, 110(1), 93–125.

Engelhardt, C., and A. Wagener (2015): "Biased perceptions of income inequality and redistribution," mimeo.

Esponda, I., and D. Pouzo (2016): "Berk–Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models," *Econometrica*, 84(3), 1093–1130.

——— (2017): "Equilibrium in misspecified Markov decision processes," working paper.

Eyster, E., and M. Rabin (2010): "Naive herding in rich-information settings," *American economic journal: microeconomics*, 2(4), 221–43.

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

Fryer, R., and M. O. Jackson (2008): "A categorical model of cognition and biased decision making," *The BE Journal of Theoretical Economics*, 8(1).

Fudenberg, D., G. Romanyuk, and P. Strack (2017): "Active learning with a misspecified prior," *Theoretical Economics*, 12(3), 1155–1189.

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

Gagnon-Bartsch, T., and M. Rabin (2017): "Naive social learning, mislearning, and unlearning," working paper.

Galperti, S., and I. Trevino (2018): "Shared Knowledge and Competition for Attention in Information Markets," Discussion paper, Working Paper.

Glaeser, E. L., and C. R. Sunstein (2009): "Extremism and social learning," *Journal of Legal Analysis*, 1(1), 263–324.

Goeree, J. K., T. R. Palfrey, and B. W. Rogers (2006): "Social learning with private and common values," *Economic theory*, 28(2), 245–264.

Golub, B., and M. O. Jackson (2010): "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2(1), 112–49.

——— (2012): "How homophily affects the speed of learning and best-response dynamics," *The Quarterly Journal of Economics*, 127(3), 1287–1338.

Guarino, A., and P. Jehiel (2013): "Social learning with coarse inference," *American Economic Journal: Microeconomics*, 5(1), 147–74.

Gul, F., W. Pesendorfer, and T. Strzalecki (2017): "Coarse competitive equilibrium and extreme prices," *American Economic Review*, 107(1), 109–37.

Haltiwanger, J., and M. Waldman (1985): "Rational expectations and the limits of rationality: An analysis of heterogeneity," *The American Economic Review*, 75(3), 326–340.

Harel, M., E. Mossel, P. Strack, and O. Tamuz (2017): "Groupthink and the Failure of Information Aggregation in Large Groups," working paper.

Hauser, O. P., and M. I. Norton (2017): "(Mis) perceptions of inequality," *Current opinion in psychology*, 18, 21–25.

He, K. (2018): "Mislearning from Censored Data: The Gambler's Fallacy in Optimal-Stopping Problems," working paper.

Heidhues, P., B. Koszegi, and P. Strack (2018): "Unrealistic expectations and misguided learning," *Econometrica*, 86(4), 1159–1214.

Jehiel, P. (2005): "Analogy-based expectation equilibrium," *Journal of Economic theory*, 123(2), 81–104.

——— (2018): "Investment strategy and selection bias: An equilibrium perspective on overoptimism," *American Economic Review*, forthcoming.

Klemperer, P. (1998): "Auctions with almost common values: The Wallet Game'and its applications," *European Economic Review*, 42(3), 757–769.

Kunda, Z. (1999): *Social cognition: Making sense of people*. MIT press.

Levy, G., and R. Razin (2015): "Correlation neglect, voting behavior, and information aggregation," *American Economic Review*, 105(4), 1634–45.

Madarász, K., and A. Prat (2017): "Sellers with misspecified models," *The Review of Economic Studies*, 84(2), 790–815.

Mayzlin, D., Y. Dover, and J. Chevalier (2014): "Promotional reviews: An empirical investigation of online review manipulation," *American Economic Review*, 104(8), 2421–55.

Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie (2018): "A Theory of Non-Bayesian Social Learning," *Econometrica*, 86(2), 445–490.

Morris, S., and H. S. Shin (2002): "Social value of public information," *American Economic Review*, 92(5), 1521–1534.

——— (2007): "Optimal communication," *Journal of the European Economic Association*, 5(2-3), 594–602.

Mueller-Frank, M. (2018): "Manipulating Opinions in Social Networks," working paper.

Mullainathan, S. (2002): "Thinking through categories," working paper.

Myatt, D. P., and C. Wallace (2011): "Endogenous information acquisition in coordination games," *The Review of Economic Studies*, 79(1), 340–374.

——— (2014): "Central bank communication design in a Lucas-Phelps economy," *Journal of Monetary Economics*, 63, 64–79.

Nisbett, R. E., and Z. Kunda (1985): "Perception of social distributions," *Journal of Personality and Social Psychology*, 48(2), 297.

Norton, M. I., and D. Ariely (2011): "Building a better America–One wealth quintile at a time," *Perspectives on Psychological Science*, 6(1), 9–12.

Nyarko, Y. (1991): "Learning in mis-specified models and the possibility of cycles," *Journal of Economic Theory*, 55(2), 416–427.

Ortoleva, P., and E. Snowberg (2015): "Overconfidence in political behavior," *American Economic Review*, 105(2), 504–35.

Pavan, A. (2014): "Attention, coordination and bounded recall," Discussion paper, Discussion Paper, Center for Mathematical Studies in Economics and Management Science.

Rabin, M. (2002): "Inference by believers in the law of small numbers," *The Quarterly Journal of Economics*, 117(3), 775–816.

Rabin, M., and D. Vayanos (2010): "The gambler's and hot-hand fallacies: Theory and applications," *The Review of Economic Studies*, 77(2), 730–778.

Reis, R. (2006a): "Inattentive consumers," *Journal of monetary Economics*, 53(8), 1761–1800.

——— (2006b): "Inattentive producers," *The Review of Economic Studies*, 73(3), 793–821.

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(3), 279–301.

Sadler, E. (2017): "False Information and Disagreement in Social Networks," working paper.

Shaked, M., and G. Shanthikumar (2007): *Stochastic orders*. Springer.

Sims, C. A. (1998): "Stickiness," in *Carnegie-Rochester Conference Series on Public Policy*, vol. 49, pp. 317–356. Elsevier.

——— (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50(3), 665–690.

Smith, L., and P. Sørensen (2000): "Pathological outcomes of observational learning," *Econometrica*, 68(2), 371–398.

Stroock, D. W. (2010): *Probability theory: an analytic view*. Cambridge university press.

Sun, Y. (2006): "The exact law of large numbers via Fubini extension and characterization of insurable risks," *Journal of Economic Theory*, 126(1), 31–69.

Vives, X. (1993): "How fast do rational agents learn?," *The Review of Economic Studies*, 60(2), 329–347.

——— (1997): "Learning from others: a welfare analysis," *Games and Economic Behavior*, 20(2), 177–200.

——— (2010): *Information and learning in markets: the impact of market microstructure*. Princeton University Press.

Wolinsky, A. (1990): "Information revelation in a market with pairwise meetings," *Econometrica*, pp. 1–23.

Supplementary Appendix to
"Misinterpreting Others and the Fragility of Social Learning"

Mira Frick, Ryota Iijima, and Yuhta Ishii

D Proof of Proposition C.1

D.1 Preliminary Results

The proof of Proposition C.1 in the next subsection makes use of the following three lemmas. Given any pair of states ω^k and ω^ℓ and perception $\hat{F}$, let $\delta_t(\omega^\ell, \omega^k, \hat{F}) := q_t(\omega^\ell, \hat{F}) - q_t(\omega^k, \hat{F})$ denote the difference between agents' perceived period t fraction of action 0 in states ω^ℓ and ω^k. By the same inductive argument as in Lemma A.1, $q_t(\omega, \hat{F})$ is strictly decreasing in ω at all t. Hence, $\delta_t(\omega^\ell, \omega^k, \hat{F}) < 0$ whenever $\ell > k$, as we have relabeled states $\omega^i \in \Omega_n$ to be increasing in i.

For all perceptions $\hat{F}$ in a small ball around the true distribution F, our first lemma provides a bound on the perceived fraction of agents that assign high probability to state ω^ℓ when the true state is ω^k, where this bound depends on $\hat{F}$ only through $\delta_t(\omega^\ell, \omega^k, \hat{F})$.

Lemma D.1. *Fix any $F \in \mathcal{F}$ and $0 < \xi < \min\{F(\theta^*(\omega^n)), 1 - F(\theta^*(\omega^1))\}$. There exists a weakly decreasing map $I : (0,1) \to (0,1)$ such that for any $\hat{F} \in B(F, \xi)$, $\omega^\ell \neq \omega^k$, $\underline{s} \leq \overline{s}$ and $\kappa \in (0,1)$, we have*

$$\Pr\left(\max_{s \in [\underline{s}, \overline{s}]} H_t(\omega^\ell | a^{t-1}, s, \hat{F}) \geq \kappa \mid \omega^k, \hat{F}\right) \leq \left(\frac{C}{\kappa}\right)^{\frac{1}{2}} \prod_{\tau=1}^{t-1} I(|\delta_\tau(\omega^\ell, \omega^k, \hat{F})|)$$

for all t, where $C := \max_{s \in [\underline{s}, \overline{s}]} \frac{\phi(s|\omega^\ell)}{\phi(s|\omega^k)}$.

Proof. Let $m := \min\{F(\theta^*(\omega^n)) - \xi, 1 - F(\theta^*(\omega^1)) - \xi\} \in (0,1)$. Note that $1 - m \geq F(\theta^*(\omega^1)) + \xi > F(\theta^*(\omega^n)) - \xi \geq m$. For any $\alpha \in (0, 1-2m]$, define

$$I(\alpha) := \max\left\{\big((1-q)(1-q')\big)^{\frac{1}{2}} + \big(qq'\big)^{\frac{1}{2}} : q, q' \in [m, 1-m] \text{ and } |q - q'| \geq \alpha\right\}$$

and define $I(\alpha) := I(1-2m)$ for all $\alpha \in (1-2m, 1]$. By compactness of the domain of q, q', $I(\cdot)$ is well-defined. Moreover, $I(\cdot)$ is clearly weakly decreasing. Finally, for all $q, q' \in (0,1)$ with $q \neq q'$, Jensen's inequality implies

$$\big((1-q)(1-q')\big)^{\frac{1}{2}} + \big(qq'\big)^{\frac{1}{2}} = q\left(\frac{q'}{q}\right)^{\frac{1}{2}} + (1-q)\left(\frac{1-q'}{1-q}\right)^{\frac{1}{2}} < \left(q\frac{q'}{q} + (1-q)\frac{1-q'}{1-q}\right)^{\frac{1}{2}} = 1.$$

Hence, $I(\alpha) \in (0,1)$ for all α.

Observe that for any $\hat{F} \in B(F,\xi)$ and t, we have

$$
\begin{aligned}
&\Pr\left(\max_{s\in[\underline{s},\overline{s}]} H_t(\omega^\ell|a^{t-1}, s, \hat{F}) \geq \kappa \mid \omega^k, \hat{F}\right) \leq \Pr\left(\max_{s\in[\underline{s},\overline{s}]} \frac{H_t(\omega^\ell|a^{t-1}, s, \hat{F})}{H_t(\omega^k|a^{t-1}, s, \hat{F})} \geq \kappa \mid \omega^k, \hat{F}\right) \\
&= \Pr\left(C \prod_{\tau=1}^{t-1} \left(a_\tau \frac{1-q_\tau(\omega^\ell,\hat{F})}{1-q_\tau(\omega^k,\hat{F})} + (1-a_\tau)\frac{q_\tau(\omega^\ell,\hat{F})}{q_\tau(\omega^k,\hat{F})}\right) \geq \kappa \mid \omega^k, \hat{F}\right) \\
&= \Pr\left(\prod_{\tau=1}^{t-1} \left(a_\tau \frac{1-q_\tau(\omega^\ell,\hat{F})}{1-q_\tau(\omega^k,\hat{F})} + (1-a_\tau)\frac{q_\tau(\omega^\ell,\hat{F})}{q_\tau(\omega^k,\hat{F})}\right)^{\frac{1}{2}} \geq (\kappa/C)^{\frac{1}{2}} \mid \omega^k, \hat{F}\right) \\
&\leq \left(\frac{C}{\kappa}\right)^{\frac{1}{2}} \mathbb{E}\left[\prod_{\tau=1}^{t-1} \left(a_\tau \frac{1-q_\tau(\omega^\ell,\hat{F})}{1-q_\tau(\omega^k,\hat{F})} + (1-a_\tau)\frac{q_\tau(\omega^\ell,\hat{F})}{q_\tau(\omega^k,\hat{F})}\right)^{\frac{1}{2}} \mid \omega^k, \hat{F}\right] \\
&= \left(\frac{C}{\kappa}\right)^{\frac{1}{2}} \prod_{\tau=1}^{t-1} \mathbb{E}\left[\left(a_\tau \frac{1-q_\tau(\omega^\ell,\hat{F})}{1-q_\tau(\omega^k,\hat{F})} + (1-a_\tau)\frac{q_\tau(\omega^\ell,\hat{F})}{q_\tau(\omega^k,\hat{F})}\right)^{\frac{1}{2}} \mid \omega^k, \hat{F}\right] \\
&= \left(\frac{C}{\kappa}\right)^{\frac{1}{2}} \prod_{\tau=1}^{t-1} \left(\left((1-q_\tau(\omega^k,\hat{F})(1-q_\tau(\omega^\ell,\hat{F}))\right)^{\frac{1}{2}} + \left(q_\tau(\omega^k,\hat{F})q_\tau(\omega^\ell,\hat{F})\right)^{\frac{1}{2}}\right),
\end{aligned}
\tag{12}
$$

where the second inequality holds by Markov's inequality, the penultimate equality holds by independence of $(a_1,\ldots,a_t)$, and the final equality holds because conditional on ω^k and $\hat{F}$, a_τ takes value 1 with probability $1-q_\tau(\omega^k,\hat{F})$ and 0 with probability $q_\tau(\omega^k,\hat{F})$.

For every τ and $\hat{F} \in B(F,\xi)$, we have $q_\tau(\omega^k,\hat{F}), q_\tau(\omega^\ell,\hat{F}) \in [\hat{F}(\theta^*(\omega^n)), \hat{F}(\theta^*(\omega^1))] \subseteq [F(\theta^*(\omega^n))-\xi, F(\theta^*(\omega^1))+\xi]$. Hence, by choice of m, $q_\tau(\omega^k,\hat{F}), q_\tau(\omega^\ell,\hat{F}) \in [m, 1-m]$. Moreover, $|q_\tau(\omega^k,\hat{F}) - q_\tau(\omega^\ell,\hat{F})| = |\delta_\tau(\omega^\ell,\omega^k,\hat{F})|$. Thus, combining (12) with the definition of $I(\cdot)$ yields the desired conclusion. □

Second, Lemma D.2 shows that for sufficiently large t, $\delta_t(\omega^{k+1},\omega^k,\hat{F})$ is bounded away from zero uniformly for all perceptions $\hat{F}$ in a small ball around the true distribution F.

Lemma D.2. *Fix any $F \in \mathcal{F}$. There exists $\xi > 0$ such that for all $k = 1, 2, \ldots, n-1$,*

$$\limsup_{t\to\infty} \sup_{\hat{F}\in B(F,\xi)} \delta_t(\omega^{k+1},\omega^k,\hat{F}) < 0.$$

Proof. See Appendix D.3. □

Finally, Lemma D.3 uses Lemmas D.1 and D.2 to show that the perceived fraction of agents who learn the true state converges to 1 at a uniform rate under all perceptions $\hat{F}$ in a small ball around F.

Lemma D.3. *Fix any $F \in \mathcal{F}$. There exists $\xi > 0$ such that for all $k = 1, 2, \ldots, n$, $\lambda \in (0,1)$, and all $\underline{s} \leq \overline{s}$,*

$$\lim_{t\to\infty} \inf_{\hat{F}\in B(F,\xi)} \Pr\left(\min_{s\in[\underline{s},\overline{s}]} H_t(\omega^k|a^{t-1}, s, \hat{F}) \geq \lambda \mid \omega^k, \hat{F}\right) = 1.$$

Proof. By Lemma D.2, there exists $\xi > 0$ such that for every $k = 1, 2, \ldots, n-1$

$$g_{k+1,k}(\xi) := \limsup_{t\to\infty} \sup_{\hat{F}\in B(F,\xi)} \delta_t(\omega^{k+1}, \omega^k, \hat{F}) < 0.$$

Since $g_{k+1,k}(\cdot)$ is weakly increasing, by picking ξ sufficiently small, we can assume additionally that $\xi < \min\{F(\theta^*(\omega^n)), 1 - F(\theta^*(\omega^1))\}$. Let $g := \max\{g_{2,1}(\xi), \ldots, g_{n,n-1}(\xi)\} < 0$.

To prove the result, fix any $\kappa > 0$, $\ell \neq k$, and $\underline{s} \leq \overline{s}$. Note that if $\ell > k$, then $|\delta_t(\omega^\ell, \omega^k, \hat{F})| = |\delta_t(\omega^\ell, \omega^{k+1}, \hat{F})| + |\delta_t(\omega^{k+1}, \omega^k, \hat{F})| \geq |\delta_t(\omega^{k+1}, \omega^k, \hat{F})|$, and likewise if $\ell < k$, then $|\delta_t(\omega^\ell, \omega^k, \hat{F})| \geq |\delta_t(\omega^{\ell+1}, \omega^\ell, \hat{F})|$. Hence, by choice of ξ and m, there exists t^* such that for all $t \geq t^*$ and $\hat{F} \in B(F, \xi)$, we have $|\delta_t(\omega^\ell, \omega^k, \hat{F})| \geq |g| > 0$. Let $C := \max_{s\in[\underline{s},\overline{s}]} \frac{\phi(s|\omega^\ell)}{\phi(s|\omega^k)} > 0$ and let $I(\cdot)$ be as in Lemma D.1. Then

$$\lim_{t\to\infty} \sup_{\hat{F}\in B(F,\xi)} \Pr\left(\max_{s\in[\underline{s},\overline{s}]} H_t(\omega^\ell | a^{t-1}, s, \hat{F}) \geq \kappa \mid \omega^k, \hat{F}\right) \leq \lim_{t\to\infty} \left(\frac{C}{\kappa}\right)^{\frac{1}{2}} (I(|g|))^{t-t^*} = 0,$$

where the inequality holds by Lemma D.1 and since $I(\cdot) \in (0,1)$ is weakly decreasing. □

D.2 Completing the Proof of Proposition C.1

We now prove Proposition C.1. By Lemma D.3, there exists some $\xi^* > 0$ such that for every $\lambda \in (0,1)$, $k = 1, \ldots, n$, and $\underline{s} \leq \overline{s}$

$$\lim_{t\to\infty} \inf_{\hat{F}\in B(F,\xi^*)} \Pr\left(\min_{s\in[\underline{s},\overline{s}]} H_t(\omega^k | a^{t-1}, s, \hat{F}) \geq \lambda \mid \omega^k, \hat{F}\right) = 1. \tag{13}$$

Hence, for every λ, k, and $\underline{s} \leq \overline{s}$, we have

$$\lim_{t\to\infty} \sup_{\hat{F}\in B(F,\xi^*)} \Pr\left(\min_{s\in[\underline{s},\overline{s}]} H_t(\omega^k | a^{t-1}, s, \hat{F}) < \lambda \mid \omega^k, \hat{F}\right) = 0. \tag{14}$$

We will show that ξ^* is as required. Under any (true or perceived) type distribution G, define $\gamma^0(k, \lambda, G)$ to be the (true or perceived) share of types that strictly prefer to play action 0 (respectively, action 1) whenever their posterior assigns probability at least λ to state ω^k.[65] Note that $\gamma^0(k, 1, F) = F(\theta^*(\omega^k)) = 1 - \gamma^1(k, 1, F)$.

Then for every $\hat{F} \in B(F, \xi^*)$, λ, k, $\underline{s} \leq \overline{s}$, and t, we have the following lower and upper bounds

[65] Formally, let $\gamma^0(k, \lambda, G) := G(\underline{\theta}_{k,\lambda})$, where $\underline{\theta}_{k,\lambda}$ satisfies $\lambda u(\underline{\theta}_{k,\lambda}, \omega^k) + (1-\lambda) u(\underline{\theta}_{k,\lambda}, \omega^n) = 0$, and let $\gamma^1(k, \lambda, G) := 1 - G(\overline{\theta}_{k,\lambda})$, where $\overline{\theta}_{k,\lambda}$ satisfies $\lambda u(\overline{\theta}_{k,\lambda}, \omega^k) + (1-\lambda) u(\overline{\theta}_{k,\lambda}, \omega^1) = 0$.

for agents' perceived fraction $q_t(\omega^k, \hat{F})$ of action 0 in state ω^k:

$$\begin{aligned}
&\Phi([\underline{s}, \overline{s}] \mid \omega^k) \Pr\left(\min_{s\in[\underline{s},\overline{s}]} H_t(\omega^k | a^{t-1}, s, \hat{F}) \geq \lambda \mid \omega^k, \hat{F}\right) \gamma^0(k, \lambda, \hat{F}) \\
&\leq q_t(\omega^k, \hat{F}) \\
&\leq (1 - \Phi([\underline{s}, \overline{s}] \mid \omega^k)) + \Phi([\underline{s}, \overline{s}] \mid \omega^k) \Pr\left(\min_{s\in[\underline{s},\overline{s}]} H_t(\omega^k | a^{t-1}, s, \hat{F}) \geq \lambda \mid \omega^k, \hat{F}\right) (1 - \gamma^1(k, \lambda, \hat{F})) \\
&+ \Phi([\underline{s}, \overline{s}] \mid \omega^k) \Pr\left(\min_{s\in[\underline{s},\overline{s}]} H_t(\omega^k | a^{t-1}, s, \hat{F}) < \lambda \mid \omega^k, \hat{F}\right).
\end{aligned}$$

In particular, this holds for $\underline{s} \leq \overline{s}$ such that $\Phi([\underline{s}, \overline{s}] \mid \omega^k)$ is arbitrarily close to 1, so combining this with (13) and (14) yields for every λ and k that

$$\begin{aligned}
&\inf_{\hat{F}\in B(F,\xi^*)} \gamma^0(k, \lambda, \hat{F}) \leq \lim_{t\to\infty} \inf_{\hat{F}\in B(F,\xi^*)} q_t(\omega^k, \hat{F}) \\
&\leq \lim_{t\to\infty} \sup_{\hat{F}\in B(F,\xi^*)} q_t(\omega^k, \hat{F}) \leq \sup_{\hat{F}\in B(F,\xi^*)} (1 - \gamma^1(k, \lambda, \hat{F})).
\end{aligned} \tag{15}$$

Hence, for every λ, k, and every $\xi \leq \xi^*$,

$$\begin{aligned}
&\lim_{t\to\infty} \sup_{\hat{F}\in B(F,\xi)} q_t(\omega^k, \hat{F}) - F(\theta^*(\omega^k)) = \lim_{t\to\infty} \sup_{\hat{F}\in B(F,\xi)} q_t(\omega^k, \hat{F}) - \gamma^0(k, 1, F) \\
&\leq \sup_{\hat{F}\in B(F,\xi)} (1 - \gamma^1(k, \lambda, \hat{F})) - \gamma^0(k, 1, F) \leq 1 - \gamma^1(k, \lambda, F) - \gamma^0(k, 1, F) + \xi,
\end{aligned}$$

where the first inequality holds by (15) and the second inequality holds for any $\hat{F} \in B(F, \xi)$. Since this holds for all $\lambda \in (0, 1)$ and $\lim_{\lambda\to 1} \gamma^1(k, \lambda, F) = 1 - \gamma^0(k, 1, F)$, this yields

$$\lim_{t\to\infty} \sup_{\hat{F}\in B(F,\xi)} \left(q_t(\omega^k, \hat{F}) - F(\theta^*(\omega^k)\right) \leq \xi. \tag{16}$$

Analogously, for every λ, k, and every $\xi \leq \xi^*$,

$$\begin{aligned}
&\lim_{t\to\infty} \inf_{\hat{F}\in B(F,\xi)} q_t(\omega^k, \hat{F}) - F(\theta^*(\omega^k)) = \lim_{t\to\infty} \inf_{\hat{F}\in B(F,\xi)} q_t(\omega^k, \hat{F}) - \gamma^0(k, 1, F) \\
&\geq \gamma^0(k, \lambda, \hat{F}) - \gamma^0(k, 1, F) \geq -\xi + \gamma^0(k, \lambda, F) - \gamma^0(k, 1, F).
\end{aligned}$$

Again, since $\lim_{\lambda\to 1} \gamma^0(k, \lambda, F) = \gamma^0(k, 1, F)$, this yields

$$\lim_{t\to\infty} \inf_{\hat{F}\in B(F,\xi)} \left(q_t(\omega^k, \hat{F}) - F(\theta^*(\omega^k))\right) \geq -\xi. \tag{17}$$

Combining (16) and (17), for every $k = 1, \ldots, n$ and $\xi \leq \xi^*$,

$$\lim_{t\to\infty} \sup_{\hat{F}\in B(F,\xi)} \left|q_t(\omega^k, \hat{F}) - F(\theta^*(\omega^k))\right| \leq \xi,$$

as required. ∎

D.3 Proof of Lemma D.2

It remains to prove Lemma D.2. Fix any $F \in \mathcal{F}$. For any $\ell, k = 1, \ldots, n$ and $\xi > 0$, define

$$g_{\ell,k}(\xi) := \limsup_{t\to\infty} \sup_{\hat{F}\in B(F,\xi)} \delta_t(\omega^\ell, \omega^k, \hat{F}).$$

We will establish the following claim: For any $i = 1, \ldots, n-1$,

$$\exists\, \xi > 0 \text{ such that } g_{k+n-i,k}(\xi) < 0 \text{ for all } k = 1, \ldots, i. \tag{18}$$

Note that Lemma D.2 corresponds to claim (18) when $i = n-1$. Below we prove the claim by induction on i. Our proof repeatedly applies the following lemma:

Lemma D.4. *Let $F \in \mathcal{F}$ and $1 \le \underline{k} < \overline{k} \le n$. Consider a sequence of times t_ℓ and perceptions $\hat{F}_\ell \in \mathcal{F}$ with $\|\hat{F}_\ell - F\| \to 0$. Suppose that for every $\kappa \in (0,1)$ and s,*

$$1 = \lim_{\ell\to\infty} \Pr\left(H_{t_\ell}([\omega^1, \omega^{\underline{k}-1}]|s, a^{t_\ell-1}, \hat{F}_\ell) > \kappa \mid \omega^j, \hat{F}_\ell\right) \qquad \forall j < \underline{k}, \tag{19}$$

$$1 = \lim_{\ell\to\infty} \Pr\left(H_{t_\ell}([\omega^{\overline{k}+1}, \omega^n]|s, a^{t_\ell-1}, \hat{F}_\ell) > \kappa \mid \omega^j, \hat{F}_\ell\right) \qquad \forall j > \overline{k}, \tag{20}$$

$$1 = \lim_{\ell\to\infty} \Pr\left(H_{t_\ell}([\omega^{\underline{k}}, \omega^{\overline{k}}]|s, a^{t_\ell-1}, \hat{F}_\ell) > \kappa \mid \omega^j, \hat{F}_\ell\right) \qquad \forall j \in \left\{\underline{k}, \ldots, \overline{k}\right\}. \tag{21}$$

Then $\limsup_{\ell\to\infty} \delta_{t_\ell}(\omega^{\overline{k}}, \omega^{\underline{k}}, \hat{F}_\ell) < 0$. Moreover, if $\overline{k} < n$, $\limsup_{\ell\to\infty} \delta_{t_\ell}(\omega^{\overline{k}+1}, \omega^{\overline{k}}, \hat{F}_\ell) < 0$, and if $\underline{k} > 1$, $\limsup_{\ell\to\infty} \delta_{t_\ell}(\omega^{\underline{k}}, \omega^{\underline{k}-1}, \hat{F}_\ell) < 0$.

Proof. See Section D.3.1. ∎

We now begin our inductive proof of claim (18). For the base case $i = 1$, we need to show that there exists $\xi > 0$ such that $g_{n,1}(\xi) < 0$. Suppose to the contrary that $g_{n,1}(\xi) = 0$ for all ξ. Then we can find a sequence of times t_ℓ and perceptions $\hat{F}_\ell$ with $\|\hat{F}_\ell - F\| \to 0$ such that $\delta_{t_\ell}(\omega^n, \omega^1, \hat{F}_\ell) \to 0$. But this contradicts Lemma D.4 applied with $\overline{k} = n$ and $\underline{k} = 1$, as in this case (19)–(21) hold trivially.

For the inductive step, suppose claim (18) holds for some $i \in \{1, \ldots, n-2\}$; that is, there exists $\xi^* > 0$ such that $g := \max_{k=1}^{i} g_{k+n-i,k}(\xi^*) < 0$. We will show that (18) holds at $i+1$. For this it suffices to show that for each $k = 1, \ldots, i+1$, there exists $\xi_k > 0$ such that $g_{k+n-(i+1),k}(\xi_k) < 0$, because then setting $\bar{\xi} := \min_{k=1}^{i+1}\{\xi_k\}$, we have $\max_{k=1}^{i+1} g_{k+n-(i+1),k}(\bar{\xi}) < 0$ as required.

Suppose for a contradiction that for some $k \in \{1, \ldots, i+1\}$, we have $g_{k+n-(i+1),k}(\xi_k) = 0$ for all ξ_k. Then we can find some sequence of times t_ℓ and perceptions $\hat{F}_\ell$ such that $\|\hat{F}_\ell - F\| \to 0$ and $\delta_{t_\ell}(\omega^{k+n-(i+1)}, \omega^k, \hat{F}_\ell) \to 0$. Without loss, assume that $\hat{F}_\ell \in B(F, \xi^*)$ for all ℓ. Below we consider separately the case when $k \in \{2, \ldots, i\}$ and when $k \in \{1, i+1\}$.

<u>**Case 1:** $k \in \{2, \dots, i\}$</u>. By choice of ξ^* and g, we have $g_{k+n-i,k}(\xi^*), g_{k+n-i-1,k-1}(\xi^*) \le g < 0$. Since for any τ and $\hat{F}$,

$$\delta_\tau(\omega^{k+n-i}, \omega^k, \hat{F}) = \sum_{j=k}^{k+n-i-1} \delta_\tau(\omega^{j+1}, \omega^j, \hat{F})$$
$$\delta_\tau(\omega^{k+n-i-1}, \omega^{k-1}, \hat{F}) = \sum_{j=k-1}^{k+n-i-2} \delta_\tau(\omega^{j+1}, \omega^j, \hat{F}),$$

we must then have

$$\limsup_{\tau\to\infty} \sup_{\hat{F}\in B(F,\xi^*)} \min_{j=k,\dots,k+n-i-1} \delta_\tau(\omega^{j+1}, \omega^j, \hat{F}) \le \frac{g}{n-i} < \frac{g}{2(n-i)} < 0,$$
$$\limsup_{\tau\to\infty} \sup_{\hat{F}\in B(F,\xi^*)} \min_{j=k-1,\dots,k+n-i-2} \delta_\tau(\omega^{j+1}, \omega^j, \hat{F}) \le \frac{g}{n-i} < \frac{g}{2(n-i)} < 0.$$

Thus, there exists T^* such that for all $\tau \ge T^*$ and ℓ, we have

$$\min_{j=k,\dots,k+n-i-1} \delta_\tau(\omega^{j+1}, \omega^j, \hat{F}_\ell) \le \frac{g}{2(n-i)},$$
$$\min_{j=k-1,\dots,k+n-i-2} \delta_\tau(\omega^{j+1}, \omega^j, \hat{F}_\ell) \le \frac{g}{2(n-i)}.$$

Note that j can take only $n-i$ possible values above. Thus, by the pigeonhole principle, for each ℓ, there must exist $\overline{j}_\ell \in \{k, \dots, k+n-i-1\}$ and $\underline{j}_\ell \in \{k-1, \dots, k+n-i-2\}$ such that

$$\#\left\{\tau \le t_\ell : \delta_\tau(\omega^{\overline{j}_\ell+1}, \omega^{\overline{j}_\ell}, \hat{F}_\ell) \le \frac{g}{2(n-i)}\right\} \ge \frac{t_\ell - T^*}{n-i},$$
$$\#\left\{\tau \le t_\ell : \delta_\tau(\omega^{\underline{j}_\ell+1}, \omega^{\underline{j}_\ell}, \hat{F}_\ell) \le \frac{g}{2(n-i)}\right\} \ge \frac{t_\ell - T^*}{n-i}.$$

Observe that either $\overline{j}_\ell = k+n-i-1$ and $\underline{j}_\ell = k-1$, or else we can assume that $k \le \underline{j}_\ell = \overline{j}_\ell \le k+n-i-2$. Moreover, since both sequences $(\overline{j}_\ell)$ and $(\underline{j}_\ell)$ can take only finitely many values, we can restrict to a subsequence such that $\overline{j}_{\ell_m}$ and $\underline{j}_{\ell_m}$ are constant, say $\overline{j}_{\ell_m} = \overline{j}$ and $\underline{j}_{\ell_m} = \underline{j}$ for all m. Again, we either have $\overline{j} = k+n-i-1$ and $\underline{j} = k-1$, or else we can assume that $k \le \underline{j} = \overline{j} \le k+n-i-2$.

We will derive a contradiction by applying Lemma D.4 along the subsequence ℓ_m. To do so, consider any $\kappa \in (0,1)$ and s. If either $j > \overline{j} \ge \ell$ or $j \le \overline{j} < \ell$, we have

$$\begin{aligned}
&\lim_{m\to\infty} \Pr\left(H_{t_{\ell_m}}(\omega^\ell | a^{t_{\ell_m}-1}, s, \hat{F}_{\ell_m}) \ge \kappa \mid \omega^j, \hat{F}_{\ell_m}\right) \\
&\le \lim_{m\to\infty} \left(\frac{C}{\kappa}\right)^{\frac{1}{2}} \prod_{\tau=1}^{t_{\ell_m}-1} I(|\delta_\tau(\omega^\ell, \omega^j, \hat{F}_{\ell_m})|) \le \lim_{m\to\infty} \left(\frac{C}{\kappa}\right)^{\frac{1}{2}} \left(I\left(\left|\frac{g}{2(n-i)}\right|\right)\right)^{\frac{t_{\ell_m}-T^*}{n-i}} = 0,
\end{aligned} \tag{22}$$

where $C := \frac{\phi(s|\omega^\ell)}{\phi(s|\omega^j)}$. Here the first inequality holds by Lemma D.1, and the second holds because

$I(\cdot) \in (0,1)$ is decreasing, $|\delta_\tau(\omega^\ell, \omega^j, \hat{F}_{\ell_m})| \geq |\delta_\tau(\omega^{\overline{j}+1}, \omega^{\overline{j}}, \hat{F}_{\ell_m})|$ and by choice of $\overline{j}$ the latter exceeds $\left|\frac{g}{2(n-i)}\right|$ at more than $\frac{t_{\ell_m}-T^*}{n-i}$ periods $\tau \leq t_{\ell_m}$.

By an analogous argument using Lemma D.1 and the choice of $\underline{j}$, if either $j > \underline{j} \geq \ell$ or $j \leq \underline{j} < \ell$, we also have

$$\lim_{m\to\infty} \Pr\left(H_{t_{\ell_m}}(\omega^\ell | a^{t_{\ell_m}-1}, s, \hat{F}_{\ell_m}) \geq \kappa \mid \omega^j, \hat{F}_{\ell_m}\right) = 0. \tag{23}$$

To apply Lemma D.4, suppose first that $k \leq \underline{j} = \overline{j} \leq k+n-i-2$. In this case, set $\underline{k} = 1$ and $\overline{k} = \overline{j} > 1$. Then (22) implies that conditions (20) and (21) hold along the subsequence ℓ_m; moreover, (19) holds trivially. Thus, by Lemma D.4, $\limsup_{m\to\infty} \delta_{t_{\ell_m}}(\omega^{\overline{k}+1}, \omega^{\overline{k}}, \hat{F}_{\ell_m}) < 0$. Since $k \leq \overline{k} < \overline{k}+1 \leq k+n-(i+1)$, this implies $\limsup_{m\to\infty} \delta_{t_{\ell_m}}(\omega^{k+n-(i+1)}, \omega^k, \hat{F}_{\ell_m}) < 0$, contradicting the assumption that $\lim_{\ell\to\infty} \delta_{t_\ell}(\omega^{k+n-(i+1)}, \omega^k, \hat{F}_\ell) = 0$.

The remaining possibility is that $\underline{j} = k-1$ and $\overline{j} = k+n-i-1$. In this case, set $\underline{k} = \underline{j}+1$ and $\overline{k} = \overline{j}$. Then (22) and (23) together imply that conditions (19)–(21) hold along the subsequence ℓ_m. Thus, by Lemma D.4, $0 > \limsup_{m\to\infty} \delta_{t_{\ell_m}}(\omega^{\overline{k}}, \omega^{\underline{k}}, \hat{F}_{\ell_m}) = \limsup_{m\to\infty} \delta_{t_{\ell_m}}(\omega^{k+n-(i+1)}, \omega^k, \hat{F}_{\ell_m})$, a contradiction. This completes Case 1.

Case 2: $k = 1, i+1$. The proof when $k = 1$ or $k = i+1$ follows similar lines. We briefly illustrate the case $k = 1$. As in Case 1, the choice of ξ^* and g implies that $g_{k+n-i,k}(\xi^*) \leq g < 0$. Thus, following exactly the same steps as in the first four displayed equations in Case 1 (but ignoring the second line of each equation), we obtain a T^* and for each ℓ some $\overline{j}_\ell \in \{k, \ldots, k+n-i-1\}$ such that

$$\#\left\{\tau \leq t_\ell : \delta_\tau(\omega^{\overline{j}_\ell+1}, \omega^{\overline{j}_\ell}, \hat{F}_\ell) \leq \frac{g}{2(n-i)}\right\} \geq \frac{t_\ell - T^*}{n-i}.$$

As in Case 1, we can restrict to a subsequence such that $\overline{j}_{\ell_m}$ is constant, say $\overline{j}_{\ell_m} = \overline{j}$ for all m.

Consider any $\kappa \in (0,1)$ and s. Just as in Case 1, we can show using Lemma D.1 that whenever either $j > \overline{j} \geq \ell$ or $j \leq \overline{j} < \ell$, we have

$$\lim_{m\to\infty} \Pr\left(H_{t_{\ell_m}}(\omega^\ell | a^{t_{\ell_m}-1}, s, \hat{F}_{\ell_m}) \geq \kappa \mid \omega^j, \hat{F}_{\ell_m}\right) = 0. \tag{24}$$

To apply Lemma D.4, there are two possibilities to consider. Suppose first that $\overline{j} > k = 1$. Then setting $\underline{k} = 1$ and $\overline{k} = \overline{j}$, (24) implies that (20) and (21) hold along the subsequence ℓ_m; moreover, (19) holds trivially. Thus, Lemma D.4 implies $\limsup_{m\to\infty} \delta_{t_{\ell_m}}(\omega^{\overline{k}}, \omega^{\underline{k}}, \hat{F}_{\ell_m}) < 0$, which contradicts $\lim_{\ell\to\infty} \delta_{t_\ell}(\omega^{k+n-(i+1)}, \omega^k, \hat{F}_\ell) = 0$, as $1 = \underline{k} < \overline{k} \leq k+n-(i+1)$.

The only remaining possibility is that $\overline{j} = 1 = k$. In this case, set $\underline{k} = 2$ and $\overline{k} = n$. Then (24) implies that (19) and (21) hold along ℓ_m; moreover, (20) holds trivially. Thus, Lemma D.4 implies $\limsup_{m\to\infty} \delta_{t_{\ell_m}}(\omega^{\underline{k}}, \omega^{\underline{k}-1}, \hat{F}_{\ell_m}) < 0$, which again contradicts $\lim_{\ell\to\infty} \delta_{t_\ell}(\omega^{k+n-(i+1)}, \omega^k, \hat{F}_\ell) = 0$, as $k = \underline{k} - 1 < \underline{k} = 2 \leq k+n-(i+1)$. □

D.3.1 Proof of Lemma D.4

Finally, we prove Lemma D.4. Consider any F, $\underline{k} < \overline{k}$, and sequences (t_ℓ) and $(\hat{F}_\ell)$ as in the statement of the lemma. Let f and $\hat{f}_\ell$ denote the densities of F and $\hat{F}_\ell$ for each ℓ.

We first prove the second and third claims. For the second claim, suppose that $\overline{k} \leq n-1$. Then note that by (20),

$$\limsup_{\ell} q_{t_\ell}(\omega^{\overline{k}+1}, \hat{F}_\ell) \leq \lim_{\kappa\to 1} \sup_{H\in\Delta(\Omega)} \left\{F(\theta^*(H)) : H([\omega^{\overline{k}+1}, \omega^n]) > \kappa\right\} \leq F(\theta^*(\omega^{\overline{k}+1})),$$

while by (21),

$$\liminf_{\ell} q_{t_\ell}(\omega^{\overline{k}}, \hat{F}_\ell) \geq \lim_{\kappa\to 1} \inf_{H\in\Delta(\Omega)} \left\{F(\theta^*(H)) : H([\omega^{\underline{k}}, \omega^{\overline{k}}]) > \kappa\right\} \geq F(\theta^*(\omega^{\overline{k}})).$$

Combining these two observations yields the desired conclusion that $\lim_{\ell\to\infty} \delta_{t_\ell}(\omega^{\overline{k}+1}, \omega^{\overline{k}}, \hat{F}_\ell) \leq F(\theta^*(\omega^{\overline{k}+1})) - F(\theta^*(\omega^{\overline{k}})) < 0$. For the third claim, suppose $\underline{k} > 1$. Then an analogous argument using (19) and (21) yields that $\lim_{\ell\to\infty} \delta_{t_\ell}(\omega^{\underline{k}}, \omega^{\underline{k}-1}, \hat{F}_\ell) \leq F(\theta^*(\omega^{\underline{k}})) - F(\theta^*(\omega^{\underline{k}-1})) < 0$.

It remains to prove the first claim. Suppose for a contradiction that $\limsup_{\ell\to\infty} \delta_{t_\ell}(\omega^{\overline{k}}, \omega^{\underline{k}}, \hat{F}_\ell) = 0$. Restricting to an appropriate subsequence, we can assume that $\lim_{\ell\to\infty} \delta_{t_\ell}(\omega^{\overline{k}}, \omega^{\underline{k}}, \hat{F}_\ell) = 0$.

For any s, θ, and $\kappa \in [0,1)$, consider an agent of type θ with private signal s. Let $\mathcal{H}^{\ell,0}(s,\theta,\kappa)$ (resp. $\mathcal{H}^{\ell,1}(s,\theta,\kappa)$) denote the event that at time t_ℓ, the agent's interim optimal action is 0 (resp. 1) and that he assigns probability greater than κ to states in $[\omega^{\underline{k}}, \omega^{\overline{k}}]$ given perception $\hat{F}_\ell$:

$$\mathcal{H}^{\ell,0}(s,\theta,\kappa) := \left\{a^{t_\ell-1} : \int u(\theta,\omega)dH_{t_\ell}(\omega|s, a^{t_\ell-1}, \hat{F}_\ell) \leq 0, H_{t_\ell}([\omega^{\underline{k}}, \omega^{\overline{k}}]|s, a^{t_\ell-1}, \hat{F}_\ell) > \kappa\right\},$$
$$\mathcal{H}^{\ell,1}(s,\theta,\kappa) := \left\{a^{t_\ell-1} : \int u(\theta,\omega)dH_{t_\ell}(\omega|s, a^{t_\ell-1}, \hat{F}_\ell) \geq 0, H_{t_\ell}([\omega^{\underline{k}}, \omega^{\overline{k}}]|s, a^{t_\ell-1}, \hat{F}_\ell) > \kappa\right\}.$$

Define the probabilities of these events conditional on each state ω^j and perception $\hat{F}_\ell$:

$$q_{t_\ell}(\omega^j, s, \theta, \kappa) = \Pr\left(\mathcal{H}^{\ell,0}(s,\theta,\kappa) \mid \omega^j, \hat{F}_\ell\right), r_{t_\ell}(\omega^j, s, \theta, \kappa) = \Pr\left(\mathcal{H}^{\ell,1}(s,\theta,\kappa) \mid \omega^j, \hat{F}_\ell\right).$$

The remainder of the proof proceeds in two steps. The first step is to establish that these probabilities become "flat" in states ω and signals s as $\ell \to \infty$. The second step obtains a contradiction by considering agents' ex-ante expected payoffs conditional on the events $\mathcal{H}^{\ell,0}(s,\theta,\kappa), \mathcal{H}^{\ell,1}(s,\theta,\kappa)$.

Step 1. For the remainder of the proof, fix any $\kappa \in (0,1)$. Observe that for almost every s and θ,

$$0 = \lim_{\ell\to\infty} r_{t_\ell}(\omega^j, s, \theta, \kappa) = \lim_{\ell\to\infty} q_{t_\ell}(\omega^j, s, \theta, \kappa) \;\forall j \notin \left\{\underline{k}, \dots, \overline{k}\right\}, \tag{25}$$

$$1 = \lim_{\ell\to\infty} r_{t_\ell}(\omega^j, s, \theta, \kappa) + q_{t_\ell}(\omega^j, s, \theta, \kappa) \;\forall j \in \left\{\underline{k}, \dots, \overline{k}\right\}, \tag{26}$$

$$0 = \lim_{\ell\to\infty} \left|q_{t_\ell}(\omega^j, s, \theta, 0) - q_{t_\ell}(\omega^j, s, \theta, \kappa)\right| \;\; \forall j \in \left\{\underline{k} \dots, \overline{k}\right\}. \tag{27}$$

where the first line follows from (19) and (20), and the second and third lines follow from (21).

By restricting to an appropriate subsequence, we can assume that for each $\omega \in \Omega_n$, $q_{t_\ell}(\omega,\cdot,\cdot,\kappa)$ converges in the weak-star topology to some L^∞ function $q_\infty(\omega,\cdot,\cdot)$.[66] Indeed, each $q_{t_\ell}(\omega,\cdot,\cdot,\kappa)$ is an L^∞ function of s and θ, and by the Banach-Alaoglu theorem, the unit ball in L^∞ is compact under the weak-star topology. Likewise, we can assume that $q_{t_\ell}(\omega,\cdot,\cdot,0)$ is weakly convergent.

For all $j \notin \{\underline{k},\dots,\overline{k}\}$, (25) implies that $q_\infty(\omega^j,s,\theta) = 0$ for almost every s and θ. For all $j \in \{\underline{k},\dots,\overline{k}\}$, (27) ensures that $q_\infty(\omega^j,\cdot,\cdot)$ coincides with the weak limit of $q_{t_\ell}(\omega,\cdot,\cdot,0)$. Hence, since $q_{t_\ell}(\omega^j,s,\theta,0)$ is weakly decreasing in s, θ, and j for each ℓ, it follows that $q_\infty(\omega^j,s,\theta)$ is weakly decreasing in s, θ, and $j \in \{\underline{k},\dots,\overline{k}\}$.

We now claim that for almost every θ, $q_\infty(\omega^j,s,\theta)$ is constant in s and $j \in \{\underline{k},\dots,\overline{k}\}$. To see this, note that

$$\begin{aligned}
\lim_{\ell\to\infty} \delta_{t_\ell}(\omega^{\overline{k}},\omega^{\underline{k}},\hat{F}_\ell) &= \lim_{\ell\to\infty} \int\int \left(q_{t_\ell}(\omega^{\overline{k}},s,\theta,0)\phi(s \mid \omega^{\overline{k}}) - q_{t_\ell}(\omega^{\underline{k}},s,\theta,0)\phi(s\mid\omega^{\underline{k}})\right) \hat{f}_\ell(\theta)\,ds\,d\theta \\
&= \int\int \left(q_\infty(\omega^{\overline{k}},s,\theta)\phi(s\mid\omega^{\overline{k}}) - q^0_\infty(\omega^{\underline{k}},s,\theta)\phi(s\mid\omega^{\underline{k}})\right) f(\theta)\,ds\,d\theta \\
&\le \int\int \left(q_\infty(\omega^{\overline{k}},s,\theta) - q_\infty(\omega^{\underline{k}},s,\theta)\right)\phi(s\mid\omega^{\underline{k}}) f(\theta)\,ds\,d\theta \le 0
\end{aligned}$$

where the second equality holds by weak-star convergence of q_{t_ℓ} to q_∞ and because $\|\hat{F}_\ell - F\| \to 0$ by assumption, and the inequalities follow from the fact that $q_\infty(\omega^j,s,\theta)$ is decreasing in s and j. Since we assumed that $\lim_{\ell\to\infty}\delta_{t_\ell}(\omega^{\overline{k}},\omega^{\underline{k}},\hat{F}_\ell) = 0$, the above inequalities hold with equality and thus $q_\infty(\omega^j,s,\theta)$ is constant in $j \in \{\underline{k},\dots,\overline{k}\}$ for almost every (s,θ). As a result,

$$0 = \int\int q_\infty(\omega^{\underline{k}},s,\theta)\left(\phi(s\mid\omega^{\overline{k}}) - \phi(s\mid\omega^{\underline{k}})\right) f(\theta) ds d\theta.$$

Hence, for almost all θ, there exists some $q_\infty(\theta)$ such that $q_\infty(\omega^j,s,\theta)$ for all $j \in \{\underline{k},\dots,\overline{k}\}$ and almost all s.

Step 2. Pick signals $\overline{s} > \underline{s}$ such that there exists some type θ^* that conditional on the event that $\omega \in \{\omega^{\underline{k}},\dots,\omega^{\overline{k}}\}$ prefers action 1 following signal $\overline{s}$, but prefers action 0 following signal $\underline{s}$. Thus, $\int_{\omega^{\underline{k}}}^{\omega^{\overline{k}}} u(\theta^*,\omega)\,dH(\omega|\overline{s}) > 0 > \int_{\omega^{\underline{k}}}^{\omega^{\overline{k}}} u(\theta^*,\omega)\,dH(\omega|\underline{s})$, where $H(\cdot|\overline{s})$ and $H(\cdot|\underline{s})$ denote the Bayesian updates of the prior following $\overline{s}$ and $\underline{s}$. Then picking $\alpha > 0$ sufficiently small, we can assume that the set of types

$$\bar{\Theta} := \left\{\theta : \int_{\underline{k}}^{\overline{k}} u(\theta,\omega)\,dH(\omega|\overline{s}) > \alpha > -\alpha > \int_{\underline{k}}^{\overline{k}} u(\theta,\omega)\,dH(\omega|\underline{s})\right\}$$

satisfies $\int_{\bar{\Theta}} f(\theta)d\theta > 0$. We will derive a contradiction of this with Step 1.

Consider any $\theta \in \bar{\Theta}$, $s \in [\overline{s}, \overline{s}+1]$ and $a^{t_\ell - 1} \in \mathcal{H}^{\ell,0}(s,\theta,\kappa)$. By definition, $0 \ge \int u(\theta,\omega)\,dH(\omega|s,a^{t_\ell-1},\hat{F}_\ell)$; i.e., action 0 is interim optimal for θ following s and $a^{t_\ell-1}$. But then, given signal s, playing action

[66] That is, $\int\int q_{t_\ell}(\omega,s,\theta,\kappa)h(s,\theta)\,ds\,d\theta \to \int\int q_\infty(\omega^j,s,\theta)h(s,\theta)\,ds\,d\theta$ for any L^1 function $h(s,\theta)$.

0 following each $a^{t_\ell-1} \in \mathcal{H}^{\ell,0}(s,\theta,\kappa)$ must also yield a higher ex ante payoff than playing action 1; that is, $0 \geq \int \Pr[a^{t_\ell-1}|\omega,\hat{F}_\ell]u(\theta,\omega)\,dH(\omega|s)$. Summing over all $a^{t_\ell-1}$ in $\mathcal{H}^{\ell,0}(s,\theta,\kappa)$, this implies $0 \geq \int q_{t_\ell}(\omega,s,\theta,\kappa)u(\theta,\omega)\,dH(\omega|s)$. Integrating across all $\theta \in \bar{\Theta}$ and $s \in [\bar{s},\bar{s}+1]$ yields

$$0 \geq \int_{\bar{\Theta}} \int_{\bar{s}}^{\bar{s}+1} \int_{\Omega^n} q_{t_\ell}(\omega,s,\theta,\kappa)u(\theta,\omega)\,dH(\omega|s)f(\theta)\,d\theta\,ds.$$

Taking the limit as $\ell \to \infty$ and using the fact that by Step 1, for almost all θ and s, $q_\infty(\omega^j,s,\theta) = q_\infty(\theta)$ if $j \in \{\underline{k},\ldots,\overline{k}\}$ and $q_\infty(\omega^j,s,\theta) = 0$ if $j \notin \{\underline{k},\ldots,\overline{k}\}$, this implies

$$0 \geq \int\limits_{\bar{\Theta}} q_\infty(\theta) \int\limits_{\bar{s}}^{\bar{s}+1} \int_{\underline{k}}^{\overline{k}} u(\theta,\omega)\,dH(\omega|s)f(\theta)\,d\theta\,ds \geq \alpha \int\limits_{\bar{\Theta}} q_\infty(\theta)f(\theta)\,d\theta, \tag{28}$$

where the second inequality holds by definition of $\bar{\Theta}$.

Considering $\theta \in \bar{\Theta}$, $s \in [\underline{s}-1,\underline{s}]$ and $a^{t_\ell-1} \in \mathcal{H}^{\ell,1}(s,\theta,\kappa)$ and proceeding in an analogous manner to the previous paragraph yields

$$0 \leq \int_{\bar{\Theta}} \int_{\underline{s}-1}^{\underline{s}} \int_{\Omega^n} r_{t_\ell}(\omega,s,\theta,\kappa)u(\theta,\omega)\,dH(\omega|s)f(\theta)\,d\theta\,ds.$$

Taking the limit as $\ell \to \infty$ and using Step 1 along with (25)–(26) and the definition of $\bar{\Theta}$, we obtain

$$0 \leq \int\limits_{\bar{\Theta}} (1-q_\infty(\theta)) \int\limits_{\underline{s}-1}^{\underline{s}} \int_{\underline{k}}^{\overline{k}} u(\theta,\omega)\,dH(\omega|s)f(\theta)\,d\theta\,ds \leq -\alpha \int\limits_{\bar{\Theta}} (1-q_\infty(\theta))f(\theta)\,d\theta. \tag{29}$$

Combining (28) and (29) implies

$$\int\limits_{\bar{\Theta}} f(\theta)\,d\theta \leq \int\limits_{\bar{\Theta}} q_\infty(\theta)f(\theta)\,d\theta \leq 0,$$

which is a contradiction. This concludes the proof of the first claim. □

E Proofs for Section 6

E.1 Proof of Proposition 1

Fix any $\alpha = \hat{\alpha} < 1$. We proved part 1 of Proposition 1 in Appendix A. To show the second part, define for each $F, \hat{F} \in \mathcal{F}$ and $\omega \in \Omega$ the set of steady states

$$\mathrm{SS}(F,\hat{F},\omega) := \left\{\hat{\omega}_\infty \in \Omega : \hat{\omega}_\infty \in \operatorname*{argmin}_{\hat{\omega}\in\Omega} \mathrm{KL}\left(\alpha F(\theta^*(\hat{\omega}_\infty)) + (1-\alpha)F(\theta^*(\omega)), \hat{F}(\theta^*(\hat{\omega}))\right)\right. . \tag{30}$$

The following lemma shows that whenever $\mathrm{SS}(F, \hat{F}, \omega)$ is finite, incorrect agents' long-run beliefs correspond to steady states.

Lemma E.1. *Fix any $F, \hat{F}$ such that $\mathrm{SS}(F, \hat{F}, \omega)$ is finite for each ω. Then in all states ω, there exists some state $\hat{\omega}_\infty(\omega) \in \mathrm{SS}(F, \hat{F}, \omega)$ such that almost all incorrect agents' beliefs converge to a point mass on $\hat{\omega}_\infty(\omega)$.*

Proof. Since Lemma B.2 continues to characterize incorrect agents' inferences from observed actions, the proof proceeds in an analogous manner to that of Proposition B.1. Specifically, let $q_t^C(\omega)$, $q_t^I(\omega) \in [0, 1]$ denote the actual fraction of action 0 among correct and incorrect agents in period t and state ω, and let $\overline{q}_t^C(\omega) := \frac{1}{t}\sum_{\tau=1}^t q_\tau^C(\omega)$ and $\overline{q}_t^I(\omega) := \frac{1}{t}\sum_{\tau=1}^t q_\tau^I(\omega)$ denote the corresponding time averages.

Note that since by the first part of Proposition 1 almost all correct agents learn the true state as $t \to \infty$, it follows that $\lim_{t\to\infty} \overline{q}_t^C(\omega) = \lim_{t\to\infty} q_t^C(\omega) = F(\theta^*(\omega))$ for all ω. Moreover, since $\mathrm{SS}(F, \hat{F}, \omega, \alpha)$ is finite, we can follow the same argument as in the proof of Lemma B.3 to show (using Lemma B.2) that the limit $R^I(\omega) := \lim_{t\to\infty} \overline{q}_t^I(\omega)$ exists for all ω.

For each ω, let

$$\hat{\omega}_\infty(\omega) := \operatorname*{argmin}_{\hat{\omega}\in\Omega} \mathrm{KL}\left(\alpha R^I(\omega) + (1-\alpha)F(\theta^*(\omega)), \hat{F}(\theta^*(\hat{\omega}))\right).$$

Then by the same argument as in the proof of Proposition B.1, we obtain that conditional on each state ω, almost all incorrect agents' beliefs converge to a point mass on $\hat{\omega}_\infty(\omega)$. But then $R^I(\omega) = F(\theta^*(\hat{\omega}_\infty(\omega)))$, whence $\hat{\omega}_\infty(\omega) \in \mathrm{SS}(F, \hat{F}, \omega)$. □

Combined with Lemma E.1, the following lemma completes the proof of the proposition.

Lemma E.2. *Fix any analytic $F \in \mathcal{F}$ and $\delta > 0$. There exists $\varepsilon > 0$ such that for any analytic $\hat{F} \neq F$ with $\|F - \hat{F}\| < \varepsilon$ and every $\omega \in \Omega$:*

1. *$\mathrm{SS}(F, \hat{F}, \omega)$ is finite.*
2. *$|\omega - \hat{\omega}| < \delta$ for every $\hat{\omega} \in \mathrm{SS}(F, \hat{F}, \omega)$.*

Proof. Fix any analytic $F \in \mathcal{F}$ and $\delta > 0$, where we can assume that $\delta < \frac{\overline{\omega}-\underline{\omega}}{2}$. Choose $\varepsilon > 0$ sufficiently small such that $\frac{\varepsilon}{1-\alpha} < |F(\theta^*(\omega)) - F(\theta^*(\omega'))|$ for any pair of states ω, ω' with $|\omega - \omega'| \geq \delta$.

Consider any analytic $\hat{F} \neq F$ with $\|F - \hat{F}\| < \varepsilon$ and any ω. By (30), each $\hat{\omega} \in \mathrm{SS}(F, \hat{F}, \omega)$ satisfies one of the following three cases:

1. $\hat{\omega} \in (\underline{\omega}, \overline{\omega})$ and $\alpha F(\theta^*(\hat{\omega})) + (1-\alpha)F(\theta^*(\omega)) = \hat{F}(\theta^*(\hat{\omega}))$
2. $\hat{\omega} = \overline{\omega}$ and $\alpha F(\theta^*(\overline{\omega})) + (1-\alpha)F(\theta^*(\omega)) \leq \hat{F}(\theta^*(\overline{\omega}))$
3. $\hat{\omega} = \underline{\omega}$ and $\alpha F(\theta^*(\underline{\omega})) + (1-\alpha)F(\theta^*(\omega)) \geq \hat{F}(\theta^*(\underline{\omega}))$.

We first show that $|\omega - \hat{\omega}| < \delta$ for all $\hat{\omega} \in \mathrm{SS}(F, \hat{F}, \omega)$. We consider only the first case, as the remaining cases are analogous. Note that

$$\alpha F(\theta^*(\hat{\omega})) + (1-\alpha)F(\theta^*(\omega)) = \hat{F}(\theta^*(\hat{\omega})) \quad \Leftrightarrow \quad F(\theta^*(\omega)) - \hat{F}(\theta^*(\hat{\omega})) = \frac{\alpha}{1-\alpha}(\hat{F}(\theta^*(\hat{\omega})) - F(\theta^*(\hat{\omega}))),$$

so that $|F(\theta^*(\omega)) - \hat{F}(\theta^*(\hat{\omega}))| \leq \frac{\alpha}{1-\alpha}\varepsilon$. Thus,

$$|F(\theta^*(\omega)) - F(\theta^*(\hat{\omega}))| \leq |F(\theta^*(\omega)) - \hat{F}(\theta^*(\hat{\omega}))| + |\hat{F}(\theta^*(\hat{\omega})) - F(\theta^*(\hat{\omega}))| \leq \frac{\alpha}{1-\alpha}\varepsilon + \varepsilon = \frac{\varepsilon}{1-\alpha}.$$

By choice of ε, this implies $|\omega - \hat{\omega}| < \delta$.

To show that $\mathrm{SS}(F, \hat{F}, \omega)$ is finite, it suffices to show that the equality $\alpha F(\theta^*(\hat{\omega})) + (1 - \alpha)F(\theta^*(\omega)) = \hat{F}(\theta^*(\hat{\omega}))$ admits at most finitely many solutions $\hat{\omega} \in [\underline{\omega}, \overline{\omega}]$. Since F and $\hat{F}$ are analytic and $[\underline{\omega}, \overline{\omega}]$ is compact, if this equality admits infinitely many solutions, then $\alpha F(\theta^*(\hat{\omega})) + (1 - \alpha)F(\theta^*(\omega)) = \hat{F}(\theta^*(\hat{\omega}))$ holds for all $\hat{\omega} \in [\underline{\omega}, \overline{\omega}]$. But the latter is impossible since we have shown that $|\omega - \hat{\omega}| < \delta < \frac{\overline{\omega}-\underline{\omega}}{2}$ holds for any solution $\hat{\omega}$. □

E.2 Proof of Proposition 2

Fix any $F \in \mathcal{F}$, $\hat{\omega} \in \Omega$, $\hat{\alpha}, \alpha > 0$ with $\hat{\alpha} \neq \alpha$ and $\varepsilon > 0$. If $\hat{\alpha} < \alpha$, take $\hat{F} \in \mathcal{F}$ such that $\hat{F} - F$ crosses zero only once at $\theta^*(\hat{\omega})$ from below. If $\hat{\alpha} > \alpha$, take $\hat{F} \in \mathcal{F}$ such that $\hat{F} - F$ crosses zero only once at $\theta^*(\hat{\omega})$ from above. In either case we can additionally require that $\|F - \hat{F}\| < \varepsilon$, as in the proof of Theorem 1. In addition, we can take $\hat{F}$ sufficiently close to F such that the inverse function $F \circ \hat{F}^{-1}$ has a Lipschitz constant less than $\frac{1}{\hat{\alpha}}$.

Let $\hat{q}_t^I(\omega)$ and $\hat{q}_t^C(\omega)$ denote incorrect and quasi-correct agents' perceived population fractions of action 0 in period t and state ω. The proof of Lemma 1 applied to incorrect agents' perceptions implies that $\hat{q}_t^I(\omega)$ is strictly decreasing in ω with $\hat{q}_\infty^I(\omega) := \lim_{t\to\infty} \hat{q}_t^I(\omega) = \hat{F}(\theta^*(\omega))$. Likewise, the proof of Proposition 1 applied to quasi-correct agents' perceptions implies that $\hat{q}_\infty^C(\omega) := \lim_{t\to\infty} \hat{q}_t^C(\omega)$ exists, is strictly decreasing, and satisfies

$$\hat{q}_\infty^C(\omega) = \hat{\alpha}F(\theta^*(\hat{\omega}_\omega)) + (1-\hat{\alpha})F(\theta^*(\omega)) \quad \text{where } \hat{\omega}_\omega = \operatorname*{argmin}_{\hat{\omega}'} \mathrm{KL}\left(\hat{q}_\infty^C(\omega), \hat{F}(\theta^*(\hat{\omega}'))\right). \tag{31}$$

Lemma E.3. *If $\hat{\alpha} < \alpha$ (resp. $\hat{\alpha} > \alpha$), then $\hat{F}(\theta^*(\omega)) - q_\infty^C(\omega)$ crosses zero only once from below (resp. above) at $\omega = \hat{\omega}$.*

Proof. Note that since by construction of $\hat{F}$ the Lipschitz constant of the the RHS of (31) is less than 1, there is a unique solution $\hat{q}_\infty^C(\omega)$ to (31). Given this, we have $\hat{q}_\infty^C(\hat{\omega}) = \hat{F}(\theta^*(\hat{\omega}))$ as $F(\theta^*(\hat{\omega})) = \hat{F}(\theta^*(\hat{\omega}))$. For the remaining claim, we focus on the case $\hat{\alpha} < \alpha$ as the case $\hat{\alpha} > \alpha$ follows a symmetric argument.

Take any $\omega < \hat{\omega}$. Then $\hat{q}_\infty^C(\omega) > \hat{q}_\infty^C(\hat{\omega}) = \hat{F}(\theta^*(\hat{\omega}))$, so that $\hat{\omega}_\omega = \operatorname{argmin}_{\hat{\omega}'} \mathrm{KL}\left(\hat{q}_\infty^C(\omega), \hat{F}(\theta^*(\hat{\omega}'))\right)$ must satisfy $\hat{\omega}_\omega < \omega$ and $\hat{F}(\theta^*(\hat{\omega}_\omega)) \leq \hat{q}_\infty^C(\omega)$. But since $F(\theta) < \hat{F}(\theta)$ for all $\theta > \theta^*(\hat{\omega})$, this implies $F(\theta^*(\hat{\omega}_\omega)) \in (F(\theta^*(\hat{\omega})), \hat{q}_\infty^C(\omega))$. Since by (31), $\hat{q}_\infty^C(\omega) = \hat{\alpha}F(\theta^*(\hat{\omega}_\omega)) + (1-\hat{\alpha})F(\theta^*(\omega))$, this implies

$F(\theta^*(\hat{\omega}_\omega)) < \hat{q}^C_\infty(\omega) < F(\theta^*(\omega)) < \hat{F}(\theta^*(\omega))$, as required. Likewise if $\omega > \hat{\omega}$, then an analogous argument shows $\hat{q}^C_\infty(\omega) > \hat{F}(\theta^*(\omega))$. □

Let $q_t(\omega)$ denote the actual population fraction of action 0 in period t at state ω, and let $\bar{q}_t(\omega) := \frac{1}{t}\sum_{\tau=1}^{t} q_\tau(\omega)$ be its time average. The following lemma uses a similar argument as in Lemma B.3 to show that $\bar{q}_t$ converges to $F(\theta^*(\hat{\omega}))$.

Lemma E.4. *For every ω, $\lim_{t\to\infty} \bar{q}_t(\omega) = F(\theta^*(\hat{\omega}))$.*

Proof. Fix any ω. Let $\overline{R}(\omega) := \limsup_{t\to\infty} \bar{q}_t(\omega)$ and $\underline{R}(\omega) := \liminf_{t\to\infty} \bar{q}_t(\omega)$. Suppose for a contradiction that either $\overline{R}(\omega) > F(\theta^*(\hat{\omega}))$ or $\underline{R}(\omega) < F(\theta^*(\hat{\omega}))$. We consider only the first case, as the second case is analogous.

Consider any $R \in (F(\theta^*(\hat{\omega}), \overline{R}(\omega)]$. We first claim that in state ω and any period t if (i) almost all incorrect agents' beliefs assign probability 1 to $\hat{\omega}^I := \operatorname{argmin}_{\hat{\omega}'} \mathrm{KL}(R, \hat{F}(\theta^*(\hat{\omega}')))$ and (ii) almost all quasi-correct agents' beliefs assign probability 1 to $\hat{\omega}^C := \operatorname{argmin}_{\hat{\omega}'} \mathrm{KL}(R, \hat{q}^C_\infty(\hat{\omega}'))$, then $q_t(\omega) < R$.

To show this claim, we consider only the case $\hat{\alpha} < \alpha$, as the case $\hat{\alpha} > \alpha$ is analogous. By Lemma E.3, $\hat{q}^C_\infty(\omega) > \hat{F}(\theta^*(\hat{\omega}))$ iff $\omega < \hat{\omega}$. Hence, we have $\hat{\omega}^C < \hat{\omega}$ since $R > F(\theta^*(\hat{\omega})) = \hat{F}(\theta^*(\hat{\omega}))$. Likewise, $\hat{\omega}^I < \hat{\omega}$. Thus, since $\hat{F}(\theta^*(\omega)) > \hat{q}^C_\infty(\omega)$ for all $\omega < \hat{\omega}$, it follows that $\hat{\omega} > \hat{\omega}^I > \hat{\omega}^C$.

By definition of $\hat{\omega}^C$, this leaves two cases to consider:

1. $R = \hat{q}^C_\infty(\hat{\omega}^C)$
2. $R > \hat{q}^C_\infty(\hat{\omega}^C)$ and $\hat{\omega}^C = \underline{\omega}$.

In either case, $q_t(\omega) = \alpha F(\theta^*(\hat{\omega}^I)) + (1-\alpha)F(\theta^*(\hat{\omega}^C))$. Moreover, in case 1, (31) implies $R = \hat{\alpha}F(\theta^*(\hat{\omega}^I)) + (1-\hat{\alpha})F(\theta^*(\hat{\omega}^C))$, so that $R > q_t(\omega)$ because $\hat{\alpha} < \alpha$ and $\hat{\omega}^I > \hat{\omega}^C$.

For case 2, we can extend the domain of function $\hat{q}^C_\infty$ from Ω to $\mathbb{R}$ by first extending the domain of function θ^* from Ω to $\mathbb{R}$ (in such a way that θ^* is still continuous, strictly decreasing, and has full range) and then defining $\hat{q}^C_\infty$ by (31) on the whole of $\mathbb{R}$. It is easy to show (using the same argument as above) that the extended $\hat{q}^C_\infty$ continues to satisfy Lemma E.3. Choosing $\tilde{\omega}^C < \overline{\omega}$ such that $R = \hat{q}^C_\infty(\tilde{\omega}^C)$ yields

$$R = \hat{\alpha}F(\theta^*(\hat{\omega}^I)) + (1-\hat{\alpha})F(\theta^*(\tilde{\omega}^C)) > \hat{\alpha}F(\theta^*(\hat{\omega}^I)) + (1-\hat{\alpha})F(\theta^*(\hat{\omega}^C)),$$

where the equality holds by (31) and the inequality holds since $\hat{\omega}^C = \underline{\omega}$. Thus, we again have $R > q_t(\omega)$ because $\hat{\alpha} < \alpha$ and $\hat{\omega}^I > \hat{\omega}^C$.

As a result, by continuity of u and F, there exist signals $\underline{s} < \overline{s}$, intervals of states $E^I \ni \hat{\omega}^I, E^C \ni \hat{\omega}^C$ with non-empty interior, and $\gamma > 0$ such that in state ω and any period t if (i') at least fraction $1-\gamma$ of incorrect agents with private signals $s \in [\underline{s}, \overline{s}]$ hold beliefs such that $H_t(E^I | a^{t-1}, s) \geq 1-\gamma$ and (ii') at least fraction $1-\gamma$ of quasi-correct agents with private signals $s \in [\underline{s}, \overline{s}]$ hold beliefs such that $H_t(E^C | a^{t-1}, s) \geq 1-\gamma$, then $q_t(\omega) < R-\gamma$.

To complete the proof, we consider separately the case where $\overline{R}(\omega) > \underline{R}(\omega)$ and the case where $\overline{R}(\omega) = \underline{R}(\omega)$. In the former case, we can choose $R \in (F(\theta^*(\hat{\omega}), \overline{R}(\omega))$ that additionally satisfies

$R > \underline{R}(\omega)$. Then following a similar argument as in the proof of Lemma B.3 leads to a contradiction. Specifically, for any sufficiently small $\eta > 0$, by definition of $\overline{R}(\omega), \underline{R}(\omega)$ and since $|\bar{q}_t(\omega) - \bar{q}_{t-1}(\omega)| < \eta$ for all large enough t, we can find an infinite sequence of times t_k such that $R - \frac{\eta}{2} \leq \bar{q}_{t_k-1}(\omega) \leq R + \frac{\eta}{2} < \bar{q}_{t_k}(\omega)$. Moreover, by choosing η small enough, the law of large numbers together with Lemma B.2 implies that for all large enough t_k hypotheses (i)' and (ii)' are satisfied. But then $q_{t_k}(\omega) < R - \gamma < R + \frac{\eta}{2}$, so that $\bar{q}_{t_k}(\omega) = \frac{t_k - 1}{t_k}\bar{q}_{t_k-1}(\omega) + \frac{1}{t_k} q_{t_k}(\omega) < R + \frac{\eta}{2}$, a contradiction.

Finally, if $\overline{R}(\omega) = \underline{R}(\omega)$, then we choose $R = \overline{R}(\omega) = \underline{R}(\omega) > F(\theta^*(\hat{\omega}))$. In this case, by the law of large numbers and Lemma B.2, almost all incorrect agents' beliefs converge to a point-mass on $\hat{\omega}^I := \operatorname{argmin}_{\hat{\omega}'} \mathrm{KL}(R, \hat{F}(\theta^*(\hat{\omega}')))$, and almost all quasi-correct agents' beliefs converge to a point-mass on $\hat{\omega}^C := \operatorname{argmin}_{\hat{\omega}'} \mathrm{KL}(R, \hat{q}^C_\infty(\hat{\omega}'))$. Thus, hypotheses (i') and (ii') are satisfied for all large enough t, whence $\lim_{t\to\infty} q_t(\omega) \leq R - \gamma$. This contradicts $\lim_{t\to\infty} \bar{q}_t(\omega) = R$. □

To complete the proof of Proposition 2, let $\hat{\omega}^I := \operatorname{argmin}_{\hat{\omega}'} \mathrm{KL}(F(\theta^*(\hat{\omega})), \hat{q}^I_\infty(\hat{\omega}'))$ and $\hat{\omega}^C := \operatorname{argmin}_{\hat{\omega}'} \mathrm{KL}(F(\theta^*(\hat{\omega})), \hat{q}^C_\infty(\hat{\omega}'))$. Then Lemmas B.2 and E.4 imply that almost all incorrect agents' beliefs converge to a point-mass on $\hat{\omega}^I$ and almost all quasi-correct agents' beliefs converge to a point-mass on $\hat{\omega}^C$. Moreover, since $\hat{q}^I_\infty(\cdot) = \hat{F}(\theta^*(\cdot))$ and $\hat{F}(\theta^*(\hat{\omega})) = F(\theta^*(\hat{\omega}))$ by construction, we must have $\hat{\omega}^I = \hat{\omega}$. Likewise, by Lemma E.3, $\hat{q}^C_\infty(\theta^*(\hat{\omega})) = \hat{F}(\theta^*(\hat{\omega})) = F(\theta^*(\hat{\omega}))$, so that $\hat{\omega}^C = \hat{\omega}$.

F Robustness of Single-Agent Active Learning

Consider the active learning model defined in Section 4.3. We measure the amount of misperception by a "bias" parameter $b \in \mathbb{R}$. Specifically, we write $\hat{q}(a, \omega) = r(a, \omega, b)$ for some C^1 function r that is strictly decreasing in (a, ω) such that $q(a, \omega) = r(a, \omega, 0)$. We also assume that $a^*(\cdot) := \operatorname{argmax}_{a\in A} u(a, \cdot)$ is C^1.

Proposition F.1. *Fix any $\varepsilon > 0$. There exists $\bar{b} > 0$ such that if $|b| < \bar{b}$, then at each $\omega \in \Omega$, process (2) admits a unique steady state $\hat{\omega}_\infty(\omega)$; moreover, $\hat{\omega}_\infty(\omega) \in [\omega - \varepsilon, \omega + \varepsilon]$ and is globally stable.*

Proof. We first show that there exists $\bar{b} > 0$ such that at each $\omega \in \Omega$, process (2) satisfies $\hat{\omega}_t \in [\omega - \varepsilon, \omega + \varepsilon]$ for all $t \geq 2$ whenever $|b| \leq \bar{b}$. To see this, consider identity

$$r(a, \omega, 0) = r(a, \hat{\omega}, b) \tag{32}$$

as a function of $\hat{\omega}$. If $b = 0$, then for any a and ω, (32) admits $\hat{\omega} = \omega$ as the unique solution. Thus, by the implicit function theorem, $\frac{d\hat{\omega}}{db} = \frac{-\frac{\partial}{\partial b} r(a,\hat{\omega},b)}{\frac{\partial}{\partial \hat{\omega}} r(a,\hat{\omega},b)}$ holds at $b = 0$ and $\hat{\omega} = \omega$. But since r is C^1 and $A \times \Omega = [0, 1] \times [\underline{\omega}, \overline{\omega}]$ is compact, $\max_{(a,\omega)\in A\times\Omega} \left| \frac{-\frac{\partial}{\partial b} r(a,\omega,0)}{\frac{\partial}{\partial \omega} r(a,\omega,0)} \right| < \infty$. Hence, there exists $\bar{b} > 0$ such that for every $b \in [-\bar{b}, \bar{b}]$, a, and ω, (32) admits a unique solution $\hat{\omega} \in [\omega - \varepsilon, \omega + \varepsilon]$; that is, process (2) satisfies $\hat{\omega}_t \in [\omega - \varepsilon, \omega + \varepsilon]$ for all $t \geq 2$ from any initial point $\hat{\omega}_1$.

Finally, applying the implicit function theorem to $r(a^*(\hat{\omega}_t), \omega, 0) = r(a^*(\hat{\omega}_t), \hat{\omega}_{t+1}, b)$, we obtain $\frac{d\hat{\omega}_{t+1}}{d\hat{\omega}_t} = -\frac{a^{*\prime}(\hat{\omega}_t)\left(\frac{\partial r(a^*(\hat{\omega}_t),\omega,0)}{\partial a^*} - \frac{\partial r(a^*(\hat{\omega}_t),\hat{\omega}_t,b)}{\partial a^*}\right)}{\frac{\partial r(a^*(\hat{\omega}_t),\hat{\omega}_{t+1},b)}{\partial \hat{\omega}_{t+1}}}$. By uniform continuity of the derivatives (which holds by compactness of the domain $A \times \Omega$), we can choose $\bar{b}$ sufficiently small such that for all $|b| \leq \bar{b}$ and ω, the right hand side is strictly less than 1 in absolute value at all $t \geq 2$. This guarantees that process (2) is a contraction on $[\omega - \varepsilon, \omega + \varepsilon]$. Hence, it admits a unique steady state $\hat{\omega}_\infty(\omega) \in [\omega - \varepsilon, \omega + \varepsilon]$, to which it converges from any initial point. □

G Continuous Actions

This section considers a continuous action space version of our model. We perform steady state analysis (under the limit model) to illustrate why our main insights do not rely on a finite action space. Throughout, we assume that the action space is an interval $A = [\underline{a}, \overline{a}] \subseteq \mathbb{R}$, with $-\infty \leq \underline{a} < \overline{a} \leq \infty$. Let $u(a, \theta, \omega)$ denote type θ's utility to choosing action a in state ω. We assume that for every type $\theta \in \mathbb{R}$ and state $\omega \in \Omega := [\underline{\omega}, \overline{\omega}]$, there exists a unique optimal action $a^*(\theta, \omega) := \text{argmax}_{a \in A} u(a, \theta, \omega)$ which is continuous and strictly increasing in (θ, ω) and such that $a^*(\cdot, \omega)$ has full range for all ω.

Given any true and perceived type distributions $F, \hat{F} \in \mathcal{F}$, we briefly analyze the set of steady states $\text{SS}(F, \hat{F})$ of this model. For each state ω, let $G(\cdot, \omega) \in \Delta(A)$ denote the true cdf over actions in the population when (almost all) agents assign probability 1 to state ω and let $g(\cdot, \omega)$ denote the corresponding density. Likewise, let $\hat{G}(\cdot, \omega)$ and $\hat{g}(\cdot, \omega)$ denote the corresponding perceived action distribution and density when agents assign probability 1 to ω. Note that $G(a, \omega) = F(\theta^*(a, \omega))$ and $\hat{G}(a, \omega) = \hat{F}(\theta^*(a, \omega))$, where $\theta^*(a, \omega)$ satisfies $a = a^*(\theta^*(a, \omega), \omega)$. Let $\text{KL}(H, \hat{H}) := \int \log \left[\frac{h(a)}{\hat{h}(a)}\right] h(a)\, da$ denote the KL divergence between continuous distributions H and $\hat{H}$ with densities h and $\hat{h}$. As in the binary action space setting, we define a steady state $\hat{\omega}^*$ to be a solution to

$$\hat{\omega}^* \in \underset{\hat{\omega}}{\text{argmin}}\, \text{KL}(G(\cdot, \hat{\omega}^*), \hat{G}(\cdot, \hat{\omega})).$$

Thus, as before, in a steady state agents assign probability 1 to a state that minimizes the KL divergence between the corresponding observed action distribution and agents' perceived action distribution. At interior steady states $\hat{\omega}^*$, the first-order condition yields

$$\int \frac{g(a, \hat{\omega}^*)}{\hat{g}(a, \hat{\omega}^*)} \frac{\partial \hat{g}(a, \hat{\omega}^*)}{\partial \hat{\omega}}\, da = 0. \tag{33}$$

Thus, the set of steady states $\text{SS}(F, \hat{F})$ is finite whenever there are at most finitely many $\hat{\omega}^*$ that satisfy (33). A sufficient condition for this is that the left-hand side of (33) is analytic in $\hat{\omega}^*$ and not constantly equal to 0; similar to the logic behind Theorem 2, this is ensured if $F \neq \hat{F}$ are analytic and $\theta^*(a, \cdot)$ is analytic. Moreover, similar to the logic behind Theorem 1, it is easy to construct examples where $\hat{F}$ is arbitrarily close to F but there is only a single (state-independent) steady state, as the following illustrates:

Example 5 (Gaussian type distributions)**.** Consider the quadratic-loss utility $u(a, \theta, \omega) = -(a - \theta - \omega)^2$, which implies that the optimal action takes the form $a^*(\theta, \omega) = \theta + \omega$. Suppose that F and $\hat{F}$ are cdfs of the Gaussian distributions $N(\mu, \sigma^2)$ and $N(\hat{\mu}, \hat{\sigma}^2)$. Then the left-hand side of (33) is given by $\int \frac{\hat{\mu}-\theta}{\hat{\sigma}^2} \frac{\exp[-\frac{(\theta-\mu)^2}{2\sigma^2}]}{\sqrt{2\pi\sigma^2}} d\theta = \frac{\hat{\mu}-\mu}{\hat{\sigma}^2}$. Thus, there is no interior steady state, and whenever $\mu > \hat{\mu}$ (respectively, $\mu < \hat{\mu}$), the unique steady state is given by $\overline{\omega}$ (respectively, $\underline{\omega}$), paralleling Example 1 in the binary action setting. □

Finally, we have focused only on steady state analysis in this section, without considering belief dynamics and establishing convergence to steady states. However, we note that in specific settings such as the Gaussian environment from Example 5 (assuming additionally that states and signals are normally distributed), the evolution of agents' beliefs in every period t admits a simple characterization in terms of difference equations, based on which convergence to the above steady states can be readily established. Details are available on request.