

MISINTERPRETING OTHERS AND THE FRAGILITY OF SOCIAL LEARNING

By

Mira Frick, Ryota Iijima, and Yuhta Ishii

January 2019

Revised March 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2160R



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY

Box 208281

New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Misinterpreting Others and the Fragility of Social Learning*

Mira Frick

Ryota Iijima

Yuhta Ishii

Abstract

We exhibit a natural environment, social learning among heterogeneous agents, where even slight misperceptions can have a large negative impact on long-run learning outcomes. 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 (e.g., tastes or risk attitudes). When agents are correct about the type distribution in the population, they learn the true state in the long run. By contrast, we show, first, 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. Second, any misperception of the type distribution leads long-run beliefs and behavior to vary only coarsely with the state, and we provide systematic predictions for how the nature of misperception shapes these coarse long-run outcomes. Third, we show that how fragile information aggregation is against misperception depends on the richness of agents' payoff-relevant uncertainty; a design implication is that information aggregation can be improved by simplifying agents' learning environment. The key feature behind our findings is that agents' belief-updating becomes "decoupled" from the true state over time. We point to other environments where this feature is present and leads to similar fragility results.

*This version: March 16, 2020. Frick: Yale University (mira.frick@yale.edu); Iijima: Yale University (ryota.ijima@yale.edu); Ishii: Pennsylvania State University (yxi5014@psu.edu). This research was supported by National Science Foundation grant SES-1824324. We thank co-editor Alessandro Lizzeri and four anonymous referees for helpful comments that significantly improved the paper. We also thank Larbi Alaoui, Nageeb Ali, Attila Ambrus, Dirk Bergemann, Aislinn Bohren, Kalyan Chatterjee, Krishna Dasaratha, Darrell Duffie, Ignacio Esponda, Erik Eyster, Drew Fudenberg, Simone Galperti, John Geanakoplos, Ben Golub, Andrei Gomberg, Ed Green, Marina Halac, Kevin He, Shaowei Ke, Igor Kheifets, Botond Köszegi, Vijay Krishna, George Mailath, Stephen Morris, Wojciech Olszewski, Romans Pancs, Antonio Penta, Jacopo Perego, Andrea Prat, Mauricio Romero, Larry Samuelson, Ran Shorrer, Ron Siegel, Rani Spiegler, Philipp Strack, Bruno Strulovici, Xavier Vives, as well as numerous seminar audiences.

1 Introduction

1.1 Motivation and Overview

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 question concerns the possibility of information aggregation in such settings, that is, under what conditions individuals are able to learn the true state of the world in the long run. A large literature has studied this question under the modeling assumption that individuals possess a *correct* understanding of their environment, in particular of the distribution of relevant population characteristics.¹ 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 (see Section 1.2). 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 the latter 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.

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

²See Section 2.4 for references.

This stark discontinuous departure motivates analyzing information aggregation under misperception 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 fragile information aggregation is against misperception depends on the richness of agents’ payoff-relevant uncertainty. A design implication is that information aggregation can be improved through interventions aimed at simplifying the agents’ learning environment. As we discuss, the key feature behind our findings is that agents’ belief-updating becomes “decoupled” from the true state over time, and we point to other learning environments where this feature is present and likewise leads to fragility against even arbitrarily small amounts of misperception.

In our model, a large population of agents choose actions in each period $t \in \{1, 2, \dots\}$ to maximize their expected utility given their time t information about a fixed but unknown state of the world $\omega \in \Omega = [\underline{\omega}, \bar{\omega}]$. 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 cumulative distribution function (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.

Our focus is on the case where agents are misspecified about the type distribution, in the sense that they misperceive the true cdf 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 .³ To isolate the effect of such misperception, we impose appropriate assumptions on signals and payoffs that ensure that if agents are correct about F , then information aggregation is successful (Lemma 1); intuitively, by observing a sufficiently large sample of others’ actions, agents are able to back out the true state in the long run if they interpret observed actions correctly.⁴

By contrast, when agents misperceive the type distribution, this 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 exogenous signals about the state of the world, but misperceives the mapping from states to signals. In this case, a classic result due to Berk (1966) implies that the agent’s long-run belief is approximately correct when the

³We use the total variation distance, but our results go through under other standard norms.

⁴Our correctly specified benchmark is a close variant of Duffie and Manso’s (2007) random matching model of social learning, where information aggregation is likewise successful, as is also the case for sequential social learning models with either rich type heterogeneity, unbounded private signals, or rich actions (e.g., Goeree, Palfrey, and Rogers, 2006; Smith and Sørensen, 2000; Lee, 1993). As we discuss in Section 4.1, information aggregation can fail in other important correctly specified settings, notably due to herding or informational cascades (e.g., Bikhchandani, Hirshleifer, and Welch, 1992; Banerjee, 1992), but the failures of information aggregation we obtain under even vanishingly small amounts of misperception are more extreme and cannot arise in *any* correctly specified model.

amount of his misperception is small (Proposition 0).

Theorem 1 offers a sharp contrast to this single-agent benchmark. In our social learning setting, even vanishingly small amounts of misperception can lead to extreme breakdowns of information 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 cdf F , but under which in all states ω , all agents' long-run beliefs incorrectly assign probability 1 to $\hat{\omega}$.

The logic behind Theorem 1 can be illustrated heuristically as follows. Suppose agents' beliefs converge to a point mass on some state $\hat{\omega}$. Then our payoff assumptions imply that behavior converges to a threshold strategy, where the types taking action 0 are precisely those below some cutoff $\theta^*(\hat{\omega})$ that is monotonic in $\hat{\omega}$. Thus, if agents' perceived type distribution is \hat{F} , they expect fraction $\hat{F}(\theta^*(\hat{\omega}))$ of long-run action observations to be 0. However, the actual fraction is $F(\theta^*(\hat{\omega}))$, as the true cdf is F . Except for boundary cases, we note that there cannot be any discrepancy between expected and actual action observations, i.e., it must be that $\hat{F}(\theta^*(\hat{\omega})) = F(\theta^*(\hat{\omega}))$; intuitively, otherwise long-run beliefs could not concentrate on $\hat{\omega}$, because agents could find some other state $\hat{\omega}' \neq \hat{\omega}$ that better explains their observations. The key feature of the equality $\hat{F}(\theta^*(\hat{\omega})) = F(\theta^*(\hat{\omega}))$ governing potential long-run beliefs is that it is *decoupled* from the true state, in the sense that it does *not* depend on the realized ω . As a result, even for perceptions \hat{F} that are arbitrarily close to F , this equality can admit a unique solution $\hat{\omega}$, and in this case a point mass on $\hat{\omega}$ is the only possible long-run belief, regardless of the true state ω . This heuristic explanation side-steps the question of whether agents' beliefs converge, which is not obvious, since under misperception beliefs do not follow a martingale. Section 4.2 illustrates the argument for convergence, which is based on approximating our original model by a simple "limit model."

The decoupling mechanism highlighted above arises because agents' belief-updating places less and less weight over time on own initial private signals and increasingly more weight on others' behavior, but the latter depends on the true state only *indirectly* through others' beliefs. As we discuss in Section 4.3, this distinguishes our setting from the aforementioned single-agent passive learning benchmark, as well as recent models of misspecified single-agent active learning, such as Heidhues, Koszegi, and Strack (2018). In those models, the agent's belief-updating does not become decoupled from the true state, because it is based on informative signals that depend *directly* on ω even asymptotically; as a result, long-run beliefs are approximately correct under small amounts of misperception. At the same time, Section 7 points to several other natural environments, including single-agent learning under a particular identification failure and other social learning settings, where belief-updating also becomes decoupled, giving rise to similar fragility results as Theorem 1.

While Theorem 1 highlights that vanishingly small amounts of misperception *can* generate stark discontinuous departures from the correctly specified benchmark, not every misperception \hat{F} necessarily gives rise to breakdowns of information aggregation that are as extreme as in Theorem 1. Theorem 2 therefore investigates information aggregation under *arbitrary* well-behaved true and perceived type distributions F and \hat{F} . Based on the decoupling mechanism highlighted above, 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 $\Omega = [\underline{\omega}, \bar{\omega}]$ 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 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 to analyze 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; by contrast, underestimating population heterogeneity leads to conservative long-run beliefs.

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 $\Omega = [\underline{\omega}, \bar{\omega}]$ 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, and in the limit, arbitrarily small amounts of misperception can give rise to extreme breakdowns of information aggregation analogous to Theorem 1. 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, Example 3 highlights a new trade-off between providing more information for agents to aggregate and rendering information aggregation more sensitive to misperception, and we 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 concludes by pointing to other natural learning environments where similar fragility results apply.

1.2 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 (e.g., Nyarko, 1991; Rabin, 2002; Rabin and Vayanos, 2010; Ortoleva and Snowberg, 2015; Fudenberg, Romanyuk, and Strack, 2017; Heidhues, Koszegi, and Strack, 2018, 2019; 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, 2019; Bohren, Imas, and Rosenberg, 2018).⁵ Our main contribution is to point to a natural setting, social learning with misperceptions of others’ characteristics, where even vanishingly small amounts of misspecification can lead to stark breakdowns of learning and to highlight a decoupling mechanism that drives this fragility result.⁶

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. They emphasize that the fact that the agent’s information depends on his actions (and hence his beliefs) can amplify the effect of misperception over time, a force that is present in our setting as well. However, as we discuss in Section 4.3, their setting differs from ours in that learning does not become decoupled from the true state and, as a result, the agent’s long-run belief is approximately correct when his amount of misperception is small.

In the context of social learning, several of the aforementioned papers incorporate various specific forms of misspecification into sequential social learning models and show that long-run beliefs grow confident in incorrect states. However, in contrast with our fragility result, they rely on strong forms of misspecification, and in a general framework that nests several of these misspecifications, Bohren and Hauser (2019) show that information aggregation is robust to small amounts of misspecification.⁷ As we discuss in Section 7.3, this distinction stems from an assumption in these models that rules out decoupled learning; absent this assumption, we show that sequential social learning can likewise be fragile, although the exact fragility mechanism and results differ from our random matching setting.

Esponda and Pouzo (2016) introduce Berk-Nash equilibrium to capture long-run beliefs under repetition of a game with (a single or multiple) misspecified players.⁸ Due to its nonstationarity, our setting does not strictly fit into their framework, but we show that, just as under Berk-Nash, agents’ long-run beliefs are based on minimization of KL-divergence. Indeed, except for boundary cases, long-run beliefs achieve zero KL-divergence between perceived and actual behavior.

Acemoglu, Chernozhukov, and Yildiz (2016) consider agents who observe exogenous public signals and hold non-common full support priors about the signal technology. In contrast with our focus on learning about the state, 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

⁵See Glaeser and Sunstein (2009); Levy and Razin (2015) for static models of information aggregation with misspecified agents. Specific forms of misspecification have also been incorporated into macroeconomic models; for a survey, see Evans and Honkapohja (2012). Hassan and Mertens (2017) consider a general equilibrium model in which agents misperceive signals about economic fundamentals. In their setting, small amounts of misperception have a continuous effect on equilibrium, but they emphasize that the slope of this effect can be arbitrarily large.

⁶We note that our fragility results are driven purely by misinferences from others’ actions, rather than by payoff externalities. The latter channel drives Madarász and Prat’s (2017) finding that screening problems can be highly sensitive to misspecification by the principal about agents’ preferences.

⁷In Gagnon-Bartsch (2017), beliefs can fail to converge under an arbitrarily small amount of misperception, but agents’ long-run average beliefs depend continuously on their misperceptions.

⁸Related approaches include Esponda (2008), Spiegel (2016), and Jehiel (2018). Esponda and Pouzo (2019) formalize Berk-Nash equilibrium for single-agent learning in a Markovian environment.

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.

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. For example, Mueller-Frank (2018) considers DeGroot-style learning on a network, where an agent can manipulate his updating rule, and shows that small amounts of manipulation can have a significant and arbitrary impact on other agents’ long-run beliefs. Given our goal of understanding the robustness of the canonical model of rational social learning, we maintain the classical assumption that agents are Bayesian. This also allows us to explicitly model agents’ perceptions and how they affect their process of inference, which is usually abstracted away from under heuristic rules.

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 density over \mathbb{R} . Let \mathcal{F} denote the space of such cdfs.

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 and continuous density over a bounded interval $\Omega := [\underline{\omega}, \bar{\omega}] \subseteq \mathbb{R}$. Agents do not observe the realization of ω . At the beginning of each period $t = 1, 2, \dots$, each agent i chooses an action $a_{it} \in \{0, 1\}$ to myopically maximize his expected utility given his period t information about ω .⁹ We assume binary actions for simplicity, but as we discuss in Supplementary Appendix E.3, 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. We assume the utility difference $u(1, \theta, \omega) - u(0, \theta, \omega)$ between actions 1 and 0 to be strictly increasing and continuously differentiable in both θ and ω , and denote this difference by $u(\theta, \omega)$. Moreover, $\lim_{\theta \rightarrow \infty} u(\theta, \underline{\omega}) > 0$ and $\lim_{\theta \rightarrow -\infty} u(\theta, \bar{\omega}) < 0$; that is, for high (respectively, low) enough types it is always optimal to choose action 1 (respectively, action 0).¹⁰ For each ω , let $\theta^*(\omega)$ denote the threshold type that is indifferent between both actions in state ω ; $\theta^*(\omega)$ is uniquely defined by $u(\theta^*(\omega), \omega) = 0$ and is strictly decreasing in ω .

⁹Myopia is without loss in this setting as players’ actions do not affect their information.

¹⁰Such 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.

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 any 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} .¹¹ 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 $t = 1, 2, \dots$, each agent i randomly meets another agent j and observes j 's period t action a_{jt} .¹² 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 7.1 briefly discusses incorporating non-independent assortative random matching.

Thus, at the beginning of each period $t = 1, 2, \dots$, 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 $1, \dots, t - 1$. Note that agents do not observe and draw inferences from their utilities; two natural interpretations of this include settings where payoffs are realized only in the long run or where successive generations of agents take one-shot actions (see Section 2.4).

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 certainty. 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 (see footnote 22).

Aside from their potential misperception of the type distribution, agents' inferences are standard. In particular, the true distributions of states, Ψ , and of private signals, $\Phi(\cdot|\omega)$, are common certainty among agents (Section 4.3 briefly discusses incorporating misperceptions about these distributions). Moreover, given perception \hat{F} , agents draw inferences from information in each period by Bayesian

¹¹We assume full-support signal distributions for notational simplicity. Our results remain valid as long as all densities $\phi(\cdot|\omega)$ admit the same support $S \subseteq \mathbb{R}$.

¹²We follow convention in assuming a law of large numbers over a continuum of i.i.d. random variables (i.e., agents' observations of signals and of other agents' actions in a given period). Thus, in aggregate, a deterministic fraction of agents observe each set of signals and action histories. 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 about the state given the perception that the type distribution is \hat{F} .

updating; more precisely, agents’ actions and beliefs in each period follow the perfect Bayesian equilibrium of our environment.¹³ 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.¹⁴ To study information aggregation, we consider the distribution of long-run beliefs; that is, the limit $\mu_\infty^\omega := \lim_{t \rightarrow \infty} \mu_t^\omega$ 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.

2.4 Examples

The above framework captures numerous economic and social situations:

(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).

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

¹³Given common certainty that the type distribution is \hat{F} , equilibrium beliefs about ω are uniquely determined at each history and actions are uniquely determined except for a measure zero set of agents who are indifferent.

¹⁴Given a topological space X , we let $\Delta(X)$ denote the set of Borel measures over X . For $X = \Delta(\Omega)$, we endow X with the topology of weak convergence.

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.¹⁵

Finally, we note that an implicit assumption in the previous examples is that agents take actions repeatedly, but do not observe their payoffs to these actions, as is assumed in a number of social learning models.¹⁶ 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, our model also fits settings where successive generations of agents take one-shot actions whose payoffs they observe privately, and subsequent generations observe a random sample of previous agents’ actions.

3 Preliminary Benchmarks

Before turning to analyze information aggregation under misperception about the type distribution in Section 4, this section presents 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 agents correctly perceive the type distribution F , they 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. This result also highlights that our model does not feature 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 , the key idea is to prove that $\lim_{t \rightarrow \infty} q_t(\omega)$ exists and is strictly decreasing in

¹⁵See, 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.

¹⁶E.g., Duffie and Manso (2007); Mossel, Sly, and Tamuz (2015).

ω , which enables agents to correctly back out the state in the long run by observing sufficiently many actions of others. 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.¹⁷

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 more agents choosing action 0 in period 2, and so on. To use this to establish that $\lim_{t \rightarrow \infty} q_t(\omega)$ 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 the same asymptotic 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 of types in the population.

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 $t = 1, 2, \dots$. 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 $q : \Omega \rightarrow (0, 1)$ 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 $\hat{q} : \Omega \rightarrow (0, 1)$, where \hat{q} is again continuous and strictly decreasing.

¹⁷In 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.

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 $\text{KL}(p, \hat{p}) := p \log(\frac{p}{\hat{p}}) + (1 - p) \log(\frac{1-p}{1-\hat{p}})$. Then in any state ω , the agent’s long-run belief assigns probability 1 to the state

$$\hat{\omega}(\omega) := \underset{\hat{\omega} \in \Omega}{\text{argmin}} \text{KL}(q(\omega), \hat{q}(\hat{\omega}))$$

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. Note also that the minimization problem can equivalently be written as $\underset{\hat{\omega} \in \Omega}{\text{argmin}} |q(\omega) - \hat{q}(\hat{\omega})|$, since $\hat{q}(\Omega)$ is a continuous interval. 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 :

Proposition 0. *For any continuous and strictly decreasing $q, \hat{q} : \Omega \rightarrow (0, 1)$, the agent’s belief in any state ω converges almost surely to a point mass on $\hat{\omega}(\omega) := \underset{\hat{\omega} \in \Omega}{\text{argmin}} \text{KL}(q(\omega), \hat{q}(\hat{\omega})) = \underset{\hat{\omega} \in \Omega}{\text{argmin}} |q(\omega) - \hat{q}(\hat{\omega})|$, 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 con-

founded learning in some sequential learning models (e.g., Bikhchandani, Hirshleifer, and Welch, 1992; Banerjee, 1992).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 that incorporate various forms of misspecification into sequential learning models (e.g., Eyster and Rabin, 2010; Bohren, 2016; Gagnon-Bartsch and Rabin, 2017). The 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 (2019) show that agents learn the true state whenever the amount of misspecification is sufficiently small. We discuss the source of this difference in Section 7.3.

The fact that an arbitrarily small amount of misperception suffices to bring about this breakdown is also in marked contrast to the single-agent passive-learning 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. We discuss the source of this contrast in Section 4.3, where we will also see why more recent models of misspecified single-agent *active* learning likewise lead to approximately successful learning under sufficiently small amounts of misperception.

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

¹⁸Herding is ruled out under either rich type heterogeneity, unbounded private signals, or rich actions (Goeree, Palfrey, and Rogers, 2006; Smith and Sørensen, 2000; Lee, 1993). 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.

¹⁹E.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.

²⁰Function $g : \mathbb{R} \rightarrow \mathbb{R}$ 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

Theorem 2 (Coarse information aggregation). *Fix any analytic F , $\hat{F} \in \mathcal{F}$ with $\hat{F} \neq F$. There exists a mapping $\hat{\omega}_\infty : \Omega \rightarrow \Omega$ 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 $\hat{\omega}_\infty$ 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 $\Omega = [\underline{\omega}, \bar{\omega}]$ 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 ω .

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 the feature that agents’ learning becomes “decoupled” over time.

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

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.

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

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 .

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 certainty), 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 $\hat{\omega}_t(\omega) := \operatorname{argmin}_{\hat{\omega} \in \Omega} \operatorname{KL} \left(q_t(\omega), \hat{F}(\theta^*(\hat{\omega})) \right) = \operatorname{argmin}_{\hat{\omega} \in \Omega} |q_t(\omega) - \hat{F}(\theta^*(\hat{\omega}))|$. 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 certainty 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

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

We briefly note two features of (1). First, 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. But, importantly, 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. This feature 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, but is also present in recent models of single-agent active learning (e.g., [Heidhues, Koszegi, and Strack, 2018](#)). However, Step 2 will highlight a key distinction between our setting and both types of single-agent models.

Second, the belief adjustment process (1) is of course 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




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 $\hat{\omega}_\infty, \hat{\omega}'_\infty, \hat{\omega}''_\infty$.

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 agents do not face any “contradictions” and belief-updating is always well-defined.

Step 2: Long-run beliefs and decoupled learning. We now consider long-run beliefs in the limit model. 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 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

To illustrate Theorem 2, we next consider long-run beliefs for arbitrary F and \hat{F} . Let

$$\text{SS}(F, \hat{F}) := \left\{ \hat{\omega}_\infty \in \Omega : \hat{\omega}_\infty = \underset{\hat{\omega} \in \Omega}{\text{argmin}} \text{KL} \left(F(\theta^*(\hat{\omega}_\infty)), \hat{F}(\theta^*(\hat{\omega})) \right) \right\} \quad (2)$$

denote the set of **steady states** of process (1). It is easy to show that any steady state $\hat{\omega}_\infty$ satisfies

²²From 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.

either $F(\theta^*(\hat{\omega}_\infty)) = \hat{F}(\theta^*(\hat{\omega}_\infty))$ or $\hat{\omega}_\infty \in \{\underline{\omega}, \bar{\omega}\}$; 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 , process (1) converges to some $\hat{\omega}_\infty(\omega) \in \text{SS}(F, \hat{F})$, where $\hat{\omega}_\infty(\omega)$ is weakly increasing in ω . Finally, 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 $[\theta^*(\bar{\omega}), \theta^*(\underline{\omega})]$.²³ As a result, $\text{SS}(F, \hat{F})$ is finite; the right-hand panel of Figure 2 provides an example in which there are three steady states. Thus, process (1) yields a weakly increasing and finite-ranged map $\hat{\omega}_\infty : \Omega \rightarrow \Omega$ from realized states ω to limit beliefs $\hat{\omega}_\infty(\omega) \in \text{SS}(F, \hat{F})$. Together with Step 3, this will establish Theorem 2.

Before proceeding to Step 3, we summarize the key feature driving both Theorems 1 and 2: This is that agents' learning becomes *decoupled* from the true state ω in the long run, as is captured by the fact that the set of steady states in (2) depends on F and \hat{F} , but does *not* depend on ω . Intuitively, this reflects that, over time, agents' belief-updating places less and less weight on own initial private signals and increasingly more weight on observations of others' behavior, but (unlike private signals) others' behavior depends on the true state only *indirectly* through others' beliefs.²⁴ As a result, agents' long-run inferences are driven entirely by the way in which they interpret observed actions. Under correct perceptions, decoupling leads to $\text{SS}(F, F) = \Omega$; that is, *all* states are steady states regardless of ω , but successful information aggregation is nevertheless possible, because initial private signals determine which steady state long-run beliefs converge to. By contrast, when agents even slightly misperceive the type distribution, then as we have seen, the set of steady states can be very limited, because only few states might allow agents to reconcile observed actions with their perceptions. Thus, depending on the nature of misperception, long-run beliefs either depend only coarsely on the true state (as in Theorem 2), or are even fully independent of ω (as in Theorem 1).

As we discuss in Section 4.3, this decoupling mechanism distinguishes our social learning setting from both the single-agent passive learning benchmark in Section 3.2, as well as recent models of misspecified single-agent active learning.

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, \dots, 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.

²³This follows from the principle of permanence for analytic functions; see footnote 53.

²⁴Indeed, in the limit model (though not in our original model), belief-updating becomes decoupled from the true state starting in period 2, as adjustments from $\hat{\omega}_{t-1}$ to $\hat{\omega}_t$ in (1) are independent of the realized ω in all periods $t \geq 2$.

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 $\operatorname{argmin}_{\hat{\omega}} \operatorname{KL}(1 - \bar{a}_t, \hat{F}(\theta^*(\hat{\omega})))$ 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.²⁵

Combining these two observations, Proposition B.1 shows that for large t , most agents’ 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:

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

where the approximation “ \approx ” becomes arbitrarily precise as $t \rightarrow \infty$. Since $\bar{q}_t(\omega)$ converges and hence is close to $q_t(\omega)$ for large t , this yields

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

i.e., an approximate version of the updating process (1) in the limit model. In particular, beliefs in the original model must converge to steady states as given by (2).

4.3 Discussion of Theorems 1 and 2

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; thus, they never encounter any “contradictory” information in finite time. Nevertheless, one might wonder whether in the limit as agents accumulate infinitely many observations, they 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 $\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t a_\tau$ converge *precisely* to the prediction $1 - \hat{F}(\theta^*(\hat{\omega}_\infty))$ under their limit belief.²⁶ This suggests that misperceptions in this environment can be a relatively persistent phenomenon.

Comparison with single-agent learning. Section 3.2 discussed a benchmark model of single-agent *passive* learning. In addition, more recent papers (e.g., Nyarko, 1991; Heidhues, Koszegi, and Strack, 2018, 2019; He, 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. However, just as the passive learning benchmark, all aforementioned models again differ

²⁵Here the law of large numbers applies since conditional on each state ω , each agent’s action observations (a_1, \dots, a_t) are independently (although not identically) distributed over time.

²⁶This observation is an analog of Proposition 6 in Heidhues, Koszegi, and Strack (2018).

from ours in that long-run beliefs under small enough amounts of misspecification are approximately correct. Below we illustrate that this is due to the fact that, as long as the signal technology satisfies an identification assumption, single-agent learning does *not* become decoupled from the true state. At the same time, Section 7.2 points to natural single-agent settings where identification failures in the signal technology give rise to decoupled learning and shows that this can lead to similar fragility results as we found above.

To see the idea, consider an abstract model of single-agent active learning: In each period $t = 1, 2, \dots$, the agent chooses an action $x_t \in X = [0, 1]$ to maximize his expected state-dependent payoff given his belief $\mu_t \in \Delta(\Omega)$ about the state. Assume for simplicity that the optimal action $x^*(\mu)$ is unique for each current belief $\mu \in \Delta(\Omega)$, and write $x^*(\omega) := x^*(\delta_\omega)$. At the end of each period t , the agent observes a signal $a_t \in \{0, 1\}$, where the probability $q(x_t, \omega)$ that $a_t = 0$ depends both on the true state ω and his period- t action x_t .²⁷ Upon observing a_t , the agent updates his belief about the state. However, in so doing, he perceives mapping $q : X \times \Omega \rightarrow (0, 1)$ to be $\hat{q} : X \times \Omega \rightarrow (0, 1)$, where q and \hat{q} are continuous in both arguments. The passive learning benchmark corresponds to the special case where q and \hat{q} do not depend on x .

Assume that $q(x, \omega)$ and $\hat{q}(x, \omega)$ are strictly decreasing in ω for all x , so that both true and perceived signal probabilities *uniquely identify* the state under each action. This is analogous to our assumption under passive learning that $q, \hat{q} : \Omega \rightarrow (0, 1)$ are strictly decreasing in ω , and is also satisfied by the aforementioned active learning models. Then, as in Section 4.2, we can consider a limit model where the agent observes not one, but infinitely many draws of signals each period, and we can set up a process of point-mass beliefs analogous to (1).²⁸ As before, the steady states $\hat{\omega}_\infty$ of this process minimize the discrepancy between the actual and perceived probabilities of signal 0:

$$\text{SS}(q, \hat{q}) = \left\{ \hat{\omega}_\infty \in \Omega : \hat{\omega}_\infty \in \underset{\hat{\omega} \in \Omega}{\text{argmin}} \text{KL}(q(x^*(\hat{\omega}_\infty), \omega), \hat{q}(x^*(\hat{\omega}_\infty), \hat{\omega})) \right\}, \quad (3)$$

with equality $q(x^*(\hat{\omega}_\infty), \omega) = \hat{q}(x^*(\hat{\omega}_\infty), \hat{\omega}_\infty)$ at interior steady states. However, the key difference with (2) is that the set of steady states in (3) *depends* on the true state ω . This reflects that belief-updating in this model does *not* become decoupled from the true state in the long run, because even asymptotically, the actual signal observations $q(x^*(\hat{\omega}_\infty), \omega)$ depend directly on ω , providing an “anchor.” In particular, in contrast with our observation that $\text{SS}(F, F) = \Omega$ in Section 4.2, the correctly specified case $\hat{q} = q$ now has the true state ω as its *unique* steady state. Under mild regularity conditions which ensure that the implicit function theorem applies, this implies that small amounts of misperception continue to yield steady states (and hence long-run beliefs) that are close

²⁷We assume that action set X is the unit interval and signals a_t are binary in order to make this model as analogous as possible to our model, where $x^*(\hat{\omega})$ corresponds to the aggregate action frequency $F(\theta^*(\hat{\omega}))$ and a_t to the observation of a random agent j 's action a_{jt} . However, neither assumption is essential for the points we make.

²⁸Specifically, the agent chooses his first action x_1^* based on his prior and arrives at a point-mass belief in the unique state $\hat{\omega}_1$ that best explains the observed signal frequency $q(x_1^*, \omega)$ given his perception \hat{q} . In all subsequent periods t , given initial belief $\hat{\omega}_{t-1}$, he chooses action $x_t^* = x^*(\hat{\omega}_{t-1})$, but then faces a discrepancy between expected and observed signal frequencies $\hat{q}(x_t^*, \hat{\omega}_{t-1})$ and $q(x_t^*, \omega)$, to which he responds by updating his belief to $\hat{\omega}_t = \underset{\hat{\omega} \in \Omega}{\text{argmin}} \text{KL}(q(x_t^*, \omega), \hat{q}(x_t^*, \hat{\omega}))$.

to the truth; see Supplementary Appendix E.1.

Misperceptions about signal distributions. Theorems 1 and 2 have considered misperceptions about the type distribution F , while assuming that agents are correctly specified about the distributions $\Phi(\cdot|\omega)$ of private signals. Focusing on the limit model in Section 4.2, we briefly illustrate that incorporating misperceptions about $\Phi(\cdot|\omega)$ *in addition* to misperceptions about F would not affect our results; by contrast, if agents correctly perceive the type distribution and *only* (commonly) misperceive $\Phi(\cdot|\omega)$ to be some $\hat{\Phi}(\cdot|\omega)$, then long-run beliefs are approximately correct under small enough amounts of misperception.²⁹ Intuitively, this reflects the fact that long-run inferences are based on observed behavior rather than initial private signals and hence do not depend on $\hat{\Phi}$. Indeed, in the limit model, $\hat{\Phi}$ only affects which belief $\hat{\omega}_1$ agents reach at the end of period 1; in all periods $t \geq 2$, behavior follows a threshold strategy with respect to the current belief $\hat{\omega}_{t-1}$, and hence the perceived type distribution \hat{F} is all that matters for agents' inferences. If $\hat{F} = F$, then agents do not further adjust their period-1 belief $\hat{\omega}_1$, and the latter is close to ω when $\hat{\Phi}$ is close to Φ . If $\hat{F} \neq F$, then regardless of $\hat{\Phi}$, agents' belief process continues to follow (1), yielding the same set of steady states $\text{SS}(F, \hat{F})$ as before and leaving Theorems 1 and 2 unaffected.

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 private signals in finitely many periods.³⁰ On the other hand, if agents receive private signals in *all* periods, then one can show that 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 learning and single-agent passive learning, and as discussed above, belief dynamics in the latter do not become decoupled from the true state over time. 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 $\hat{\omega}$.³¹ 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

²⁹Theorems 1–2 are also unaffected if agents have an incorrect common prior $\hat{\Psi}$ over states, and in this case information aggregation is successful if $\hat{F} = F$.

³⁰If 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).

³¹More 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 $\hat{\omega}_\infty(\omega) \in \text{argmin}_{\hat{\omega}} \text{KL}(F(\theta^*(\hat{\omega}_\infty(\omega))), \hat{F}(\theta^*(\hat{\omega}))) + \text{KL}(\Phi(\cdot|\omega), \Phi(\cdot|\hat{\omega}))$. 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 $\|\Phi(\cdot|\underline{\omega}) - \Phi(\cdot|\bar{\omega})\| < \delta$, then in all states ω , almost all agents' beliefs converge to a point mass on some $\hat{\omega}_\infty(\omega) \in [\hat{\omega} - \varepsilon, \hat{\omega} + \varepsilon]$.

a large negative impact on information aggregation remains valid.³²

Finite horizon. Following much of the literature on information aggregation, our analysis focuses on asymptotic beliefs. Given any 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 $\delta_{\bar{\omega}}$ if t is large enough. One implication is that, under misperception, halting agents’ interactions after a certain number of periods may improve ex-ante expected payoffs; this is not the case for correctly specified social learning models (even in the presence of herding).

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 $\text{SS}(F, \hat{F})$ 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.

Example 1 (First-Order Stochastic Dominance and “Fake” Recommenders). 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 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 .³³ In the correctly specified model of Section 3.1, the presence of fake recommenders has no long-run effect, 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 $\bar{\omega}$, regardless of the true quality ω .

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

³²Likewise, if some fraction of agents ignore others’ actions and act solely based on their initial private signals, the same approximate analog of Theorem 1 holds whenever this fraction is small.

³³Concretely, 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 $\hat{\beta} < \beta$.




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

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

Figure 3 shows the intuition in the limit model. When agents 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 about the state, converging eventually to a point mass on the unique steady state $\bar{\omega}$. \square

Example 2 (Under-/Overestimation of Population Heterogeneity). Another widely documented form of misperception is that in many contexts, individuals tend to underestimate type heterogeneity in society.³⁴ We capture this by means of the commonly used dispersiveness order (Shaked and 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 $F \succ_{\text{disp}} \hat{F}$. 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 $F \succ_{\text{FO}_{\Theta^*}} \hat{F}$ or $\hat{F} \succ_{\text{FO}_{\Theta^*}} F$, 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 to a point mass on an *interior* state $\hat{\omega} \in (\underline{\omega}, \bar{\omega})$:

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 $F \succ_{\text{disp}} \hat{F}$, then there exists some $\hat{\omega} \in (\underline{\omega}, \bar{\omega})$ 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

³⁴For example, several studies (e.g., Norton and Ariely, 2011; Engelhardt and Wagener, 2015) find systematic underestimation of wealth inequality in many countries.




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

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 $\hat{\omega} \in (\underline{\omega}, \bar{\omega})$. As we saw in the proof sketch of Theorem 1, this implies $SS(F, \hat{F}) = \{\hat{\omega}\}$. The same conclusion as in Corollary 2 can also arise in the context of fake product recommendations, if in addition to the fake positive recommenders in Example 1, there are fake negative recommenders (modeled as types $\theta \leq \theta^*(\bar{\omega})$), and agents underestimate the share of both fake types.³⁵

In other contexts, people are found to overestimate population heterogeneity.³⁶ This corresponds to $\hat{F} \succ_{\text{disp}} F$, 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 $\hat{\omega} \in (\underline{\omega}, \bar{\omega})$, and the limit model predicts convergence to the *extreme* beliefs $\underline{\omega}$ and $\bar{\omega}$ in almost all states ω .³⁷ \square

5 Rich vs. Coarse State Spaces

In this section, we show that a key determinant of the fragility of information aggregation to 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 $\{\omega^1, \omega^2, \dots\}$ that is dense in $\Omega = [\underline{\omega}, \bar{\omega}]$ and let $\Omega_n := \{\omega^1, \dots, \omega^n\}$ for each n .

The following result makes two points. First, for any *fixed* finite state space Ω_n (with arbitrary

³⁵Extending footnote 33, suppose the true type distribution is $F = \beta P + \gamma N + (1 - \beta - \gamma)G$, where N is the distribution of “negative promotional” types with $\text{supp}N \subseteq (-\infty, \theta^*(\bar{\omega})]$, while $\hat{F} = \hat{\beta} P + \hat{\gamma} N + (1 - \hat{\beta} - \hat{\gamma})G$ with $\hat{\beta} < \beta, \hat{\gamma} < \gamma$. Then \hat{F} crosses F from below in a single point in Θ^* provided $(\beta + \gamma - \hat{\beta} - \hat{\gamma})(G(\theta^*(\underline{\omega})) - G(\theta^*(\bar{\omega}))) > \gamma - \hat{\gamma}$; this holds whenever the genuine type distribution G puts high enough mass on non-dominant types.

³⁶See, e.g., Ahler (2014) in the context of perceived political attitudes.

³⁷Note that, in addition to $\underline{\omega}, \bar{\omega}$, the set of steady states (2) also includes $\hat{\omega}$. Thus, Theorem 2 only implies that agents’ long-run beliefs in the original model are given by a weakly increasing mapping $\hat{\omega}_\infty : \Omega \rightarrow \{\underline{\omega}, \hat{\omega}, \bar{\omega}\}$. However, steady state $\hat{\omega}$ is unstable under the limit model dynamics, because in almost all states ω , (1) converges to either $\underline{\omega}$ or $\bar{\omega}$. Given this, we conjecture that in the original model $\hat{\omega}_\infty(\omega)$ likewise takes values $\underline{\omega}$ or $\bar{\omega}$ in almost all states ω .

full-support prior $\Psi_n \in \Delta(\Omega_n)$, 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 $n \rightarrow \infty$, 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 $\hat{\omega}_\infty(\omega) \in [\hat{\omega} - \varepsilon, \hat{\omega} + \varepsilon]$.*

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 $n \rightarrow \infty$.

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 $\|\hat{F} - F\| \rightarrow 0$; 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, as 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




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 .

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^1 . 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 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. The following example illustrates this point in the context of public news releases.³⁸

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), 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

³⁸An analogous example can illustrate the benefits of trading “simpler” financial assets; see Example 4 in the previous working paper version, Frick, Iijima, and Ishii (2019b).

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 $\omega \in \Omega = [\underline{\omega}, \bar{\omega}]$, the sender observes a signal $\sigma(\omega) \in \mathbb{R}$, where $\sigma : \Omega \rightarrow \mathbb{R}$ is weakly increasing; here the partition $\Pi_\sigma := \{\sigma^{-1}(m) : m \in \mathbb{R}\}$ 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} .³⁹ 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, while if σ has finite range, then the setting is isomorphic to one with a finite state space where each state corresponds to a cell of Π_σ .

If in all subsequent periods $t \geq 1$ 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).⁴⁰ \square

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 certainty 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 others’ beliefs about others’ beliefs about F etc.) are correct.⁴¹ 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

³⁹Log-concave densities of η_i ensure that signal distributions satisfy the monotone likelihood ratio property. We assume for simplicity that 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.

⁴⁰This 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).

⁴¹This 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 certainty of F .

“non-standard” is how such agents interact with standard and sophisticated agents, who are aware of the presence of these non-standard agents. 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 certainty 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.⁴²

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 $|\hat{\omega}_\infty(\omega) - \omega| < \delta$.*

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 $\hat{\alpha} \neq \alpha$.

⁴²Note 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 $\alpha \rightarrow 1$.

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 D. 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, just as in Lemma 1. Given this, the fact that correct agents' long-run behavior depends on the true state ω provides an anchor for incorrect agents' belief-updating even asymptotically, similar to the discussion of repeated informative private signals in Section 4.3.

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 gives rise to an analogous belief adjustment process as in Section 4.2. Moreover (ignoring boundary cases), incorrect and quasi-correct agents' steady states $(\hat{\omega}_\infty^I, \hat{\omega}_\infty^C)$ 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 $\alpha F(\theta^*(\hat{\omega}_\infty^I)) + (1 - \alpha)F(\theta^*(\hat{\omega}_\infty^C))$ and their expected frequency $\hat{\alpha} F(\theta^*(\hat{\omega}_\infty^I)) + (1 - \hat{\alpha})F(\theta^*(\hat{\omega}_\infty^C))$. 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. Theorem 2 can be generalized analogously.

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 certainty 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).⁴³ Notably, even when agents' perceptions \hat{F}_i 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).

⁴³To see this, consider the limit model from Section 4.2. Let $\lambda \in \Delta(\mathbb{R} \times \mathcal{F})$ denote the joint distribution over 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; as a result, 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 $[0, 1]$.

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.⁴⁴ 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 ; agents then 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 it is not difficult to construct arbitrarily small perturbations of the correctly specified model that feature nondegenerate misperceptions yet yield the same extreme breakdown of information aggregation as in Theorem 1.

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 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, leading to the possibility of incomplete long-run learning.

7 Concluding Discussion

This paper has highlighted a natural learning environment where small amounts of misperception can lead long-run beliefs and behavior to depart significantly from the correctly specified benchmark. While we have made this point in the context of a particular social learning model where agents misperceive others’ characteristics, we view our results as further motivating the study, both empirically and theoretically, of relevant misperceptions in other learning environments. To illustrate that our fragility finding and the decoupling mechanism that drives it are not limited to our specific environment, we briefly discuss some other settings that can give rise to similar results.

7.1 Beyond misperceptions of others’ characteristics

First, remaining within the social learning model of this paper, agents might misperceive features of the environment other than the type distribution F . While information aggregation is robust to some such misperceptions (see the discussion of misperceptions about signal distributions in

⁴⁴See Bohren and Hauser (2019) for such a formulation in the context of a sequential social learning model.

Section 4.3), others lead to analogous fragility results as in the current paper. We briefly discuss two examples, focusing on steady states for simplicity.

Misperceptions about action sampling technology. Consider an extension of our model, in which agents' probability of observing action 0 in state ω and period t is given by $g(q_t(\omega))$, where $q_t(\omega)$ denotes the true share of action 0 in the population and $g : (0, 1) \rightarrow (0, 1)$ is a continuous and strictly increasing function that captures a potentially distorted action sampling technology, for instance, due to greater salience of one of the actions.⁴⁵ In such a setting, agents might naturally misperceive g to be some other mapping $\hat{g} : (0, 1) \rightarrow (0, 1)$, for example, because they fail to fully take into account the sampling distortion. This gives rise to an analogous limit model as in Section 4.2, where (interior) steady states must again yield no discrepancy between the actual and perceived shares of action 0 observations, i.e.,

$$g(F(\theta^*(\hat{\omega}_\infty))) = \hat{g}(\hat{F}(\theta^*(\hat{\omega}_\infty))).$$

Based on this it is easy to see that, even if agents are correct about the type distribution ($\hat{F} = F$), misperceiving the sampling technology leads to analogs of Theorems 1 and 2. For example, if action 0 is less salient than action 1 but agents even slightly underestimate the extent to which this is the case (i.e., $g(q) < \hat{g}(q) < q$ for all $q \in (0, 1)$), then long-run beliefs are a point mass on the highest state $\bar{\omega}$, regardless of the true state ω .

Misperceptions about matching technology. Consider another extension that incorporates assortativity into the random matching technology. Formally, the matching technology $P(\cdot, \cdot)$ is given by a symmetric, continuous and full-support distribution over $\mathbb{R} \times \mathbb{R}$ whose marginal is the type distribution F , and each period, each agent i of type θ_i randomly meets an agent j whose type is drawn from the conditional distribution $P(\cdot|\theta_i)$. We assume that (i) $P(\cdot|\theta) \succ_{\text{FO}} P(\cdot|\theta')$ for all $\theta \geq \theta'$, that is, reflecting assortativity, higher types are more likely than lower types to meet other high types; and (ii) $P(\theta|\theta)$ is strictly increasing in θ .⁴⁶ Our baseline model with independent random matching is the special case where $P = F \times F$.

In such a setting, it may be natural for agents to misperceive the matching technology to be some $\hat{P} \neq P$. Supplementary Appendix E.2 sets up a limit model analogous to Section 4.2. The key novelty is that, even if all types share the same perception \hat{P} , (interior) steady states are now given by *type-dependent* point-mass beliefs $\hat{\omega}_\infty^\theta$: Specifically, agents follow a threshold strategy with a cutoff θ_∞^* that satisfies $P(\theta_\infty^*|\theta_\infty^*) = \hat{P}(\theta_\infty^*|\theta_\infty^*)$, and each type θ 's steady-state belief $\hat{\omega}_\infty^\theta$ ensures that there is no discrepancy between θ 's true and perceived probability of observing action 0, i.e.,

$$P(\theta_\infty^*|\theta) = \hat{P}(\theta_\infty^*|\hat{\omega}_\infty^\theta|\theta).$$

If $P(\theta|\theta)$ and $\hat{P}(\theta|\theta)$ cross only once, then the cutoff type θ_∞^* and steady-state beliefs $\hat{\omega}_\infty^\theta$ are unique; thus, analogous to Theorem 1, long-run beliefs are independent of the true state, regardless

⁴⁵Banerjee and Fudenberg (2004) study such distorted sampling, but assume that agents are aware of the distortion.

⁴⁶This condition is satisfied by many parametric families of P , for example, when P is bivariate Gaussian.

of the amount of agents’ misperception. A natural misperception that can give rise to this is when agents are correct about the type distribution, but underestimate the extent of assortativity.⁴⁷ Supplementary Appendix E.2 formalizes this in a Gaussian setting and shows that it leads the state-independent beliefs $\hat{\omega}_\infty^\theta$ to be increasing in types, in line with empirical evidence on positive association between agents’ preferences and beliefs.⁴⁸

7.2 Single-agent active learning with identification failures

In Section 4.3, we illustrated by looking at steady states, that single-agent learning does not become decoupled when the signal technology satisfies an identification assumption. Our companion paper, [Frick, Iijima, and Ishii \(2019c\)](#), analyzes learning dynamics in a general class of misspecified learning environments that nests single-agent active learning and presents formal belief-convergence results. Just as importantly, however, our companion paper points to natural single-agent settings where, as a result of identification failures in the signal technology, learning *does* become decoupled, and we show that this leads to similar fragility results as in the current paper.

To illustrate, consider the same active learning model as in Section 4.3, but suppose that for some action x_θ , true and perceived signal distributions are completely uninformative, in the sense that $q(x_\theta, \omega) =: q_\theta$ and $\hat{q}(x_\theta, \omega) =: \hat{q}_\theta$ are both constant in ω . Moreover, suppose that whenever the agent’s belief is a point mass on any particular state $\hat{\omega}$, then his optimal action is $x^*(\hat{\omega}) = x_\theta$. In [Frick, Iijima, and Ishii \(2019c\)](#), we highlight that such an identification failure, which we term “non-identification at point-mass beliefs” (NIP), is natural in costly information acquisition settings, where the agent may choose to stop acquiring information whenever he is confident in any given state. Observe that, under this identification failure, the set of steady states in (3) becomes

$$\begin{aligned} \text{SS}(q, \hat{q}) &= \left\{ \hat{\omega}_\infty \in \Omega : \hat{\omega}_\infty \in \underset{\hat{\omega} \in \Omega}{\text{argmin}} \text{KL}(q(x^*(\hat{\omega}_\infty), \omega), \hat{q}(x^*(\hat{\omega}_\infty), \hat{\omega})) \right\} \\ &= \left\{ \hat{\omega}_\infty \in \Omega : \hat{\omega}_\infty \in \underset{\hat{\omega} \in \Omega}{\text{argmin}} \text{KL}(q_\theta, \hat{q}_\theta) \right\} = \Omega. \end{aligned}$$

Hence, in contrast with (3) but as in (2), the set of steady states does *not* depend on the true state ω . This reflects that under non-identification at point-mass beliefs, single-agent learning also becomes decoupled, in that long-run observations $q(x^*(\hat{\omega}_\infty), \omega) = q_\theta$ are no longer tied to ω . However, there is one important difference with the fragility mechanism in the current paper: Whereas Theorem 1 relied on misperceptions for which $\text{SS}(F, \hat{F})$ consists of a *single* state $\hat{\omega}$, the present identification failure leads $\text{SS}(q, \hat{q})$ to always include *all* states in Ω . Thus, determining which steady states beliefs converge to requires substantially different techniques, which we develop in [Frick, Iijima, and Ishii \(2019c\)](#). Using these techniques, we show that, similar to Theorem 1, arbitrarily small amounts of misperception can again lead the agent’s belief to converge to a point mass on the same fixed state $\hat{\omega}$, regardless of the true state ω (see Section 6.2 of [Frick, Iijima, and Ishii, 2019c](#)).

⁴⁷[Frick, Iijima, and Ishii \(2019a\)](#) study the implications of such “assortativity neglect” in static coordination games.

⁴⁸See, e.g., [Bullock, Gerber, Hill, and Huber \(2013\)](#) in the context of partisan bias in factual beliefs about politics.

7.3 Social learning with public action observations

In the present paper, agents learn by *privately* observing random other agents’ actions, analogous to other aforementioned models of social learning through decentralized random interactions. At the same time, other canonical social learning models feature learning based on *public* observations. For example, in sequential social learning models a countable sequence of agents each choose one-shot actions after observing all previous agents’ actions, and in [Vives \(1993\)](#), agents observe a public signal (e.g., a market price) of the aggregate action frequency q_t each period.

The framework in [Frick, Iijima, and Ishii \(2019c\)](#) can also be applied to study the impact of misperceptions in such models. We again show that information aggregation is highly fragile and that vanishingly small amounts of misperception about others’ characteristics can lead beliefs to converge to a state-independent point-mass.⁴⁹ At the same, public observations lead to some differences in both the fragility mechanism and qualitative predictions. Specifically, just as in the current paper, belief-updating under public observations becomes decoupled from the true state, because agents’ long-run inferences are based only on observed actions; however, while in the current paper agents regard long-run action observations as highly indicative of the state, with public observations agents view new actions as uninformative in the long run, because asymptotically agents’ actions are based purely on the public belief rather than additional private information. This gives rise to a similar identification failure as in [Section 7.2](#), as a result of which *all* states are steady states. At a qualitative level, we show that under public observations, information aggregation is fragile even in finite state spaces, in contrast with our findings in [Section 5](#).

Finally, we note that these fragility results contrast with [Bohren \(2016\)](#) and [Bohren and Hauser’s \(2019\)](#) finding that information aggregation is robust to small amounts of misperception in closely related sequential social learning environments. The difference stems from their assumption that agents observe repeated exogenous public signals about the state (or, alternatively, that a fraction of “autarkic” agents act solely based on their private signals). Analogous to the discussion of repeated private signals in [Section 4.3](#), this assumption implies that belief-updating never becomes decoupled from the true state; however, as long as the public signal is sufficiently uninformative (or the fraction of autarkic agents is sufficiently small), one can again establish approximate analogs of our fragility results.

⁴⁹See [Section 6.3 of Frick, Iijima, and Ishii \(2019c\)](#), in the context of sequential social learning. To isolate the effect of misperception, we again consider a setting where due to rich type heterogeneity, information aggregation is successful when agents are correctly specified.

Appendix: Main Proofs

A Proof of Lemma 1

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, $\int (u(1, \omega, \theta^*(H)) - u(0, \omega, \theta^*(H))) dH(\omega) = 0$. 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 $\mu_t^\omega \in \Delta(\Delta(\Omega))$ denote the population distribution of agents' posteriors. Note that the distribution of posteriors is independent across types θ . Let $q_t(\omega)$ denote the fraction of agents choosing action 0 at t and ω , i.e.,

$$q_t(\omega) = \int F(\theta^*(H)) d\mu_t^\omega(H).$$

Finally, let $\bar{\mu}_t \in \Delta(\Delta(\Omega))$ denote the *hypothetical* population distribution of posteriors 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 $\tau = 1, \dots, t-1$ that are generated according to $q_\tau(\omega)$ defined above, and agents update their beliefs assuming that a_τ is distributed according to q_τ . Given this, we can also express q_t as

$$q_t(\omega) = \int \int F(\theta^*(H^s)) d\Phi(s|\omega) d\bar{\mu}_t^\omega(H). \quad (4)$$

A.1 Agents' Long-Run Behavior

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

Lemma A.1. *For each t , $q_t(\omega)$ is strictly decreasing in ω and satisfies $q_t(\omega) \in (0, 1)$.*

Proof. The claim that $q_t(\omega) \in (0, 1)$ is clear from the fact that types above $\theta^*(\underline{\omega})$ always choose action 1 and types below $\theta^*(\bar{\omega})$ always choose action 0. To show that it is strictly decreasing in ω , we proceed by induction on t . For $t = 1$, this follows from (4) and the fact that $\bar{\mu}_1^\omega = \delta_\Psi$ 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 $\Psi^{s, a^t} \in \Delta(\Omega)$ denote the posterior belief of an agent who observes private signal s and action sequence $a^t = (a_1, \dots, a_t) \in \{0, 1\}^t$. Note that Ψ^{s, a^t} has full support on Ω , since $q_\tau(\omega) \in (0, 1)$ for all $\tau = 1, \dots, t$ 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 $k = 2, \dots, t - 2$,

we have

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

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

To prove that q_t remains strictly decreasing in ω in the limit as $t \rightarrow \infty$, Lemmas A.2 and A.3 first consider the limit of the hypothetical belief distribution $\bar{\mu}_t^\omega$ that is based only on action observations. Using standard arguments, 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 $\bar{\mu}_\infty^\omega := \lim_t \bar{\mu}_t^\omega$ exists for all $\omega \in \Omega^*$.*

Proof. To formalize any agent's 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 $a^\infty = (a_1, a_2, \dots) \in \{0, 1\}^\infty$, \mathcal{A} denotes the corresponding Borel algebra, and measure \mathbb{P} satisfies

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

for every Borel set $E \subseteq \Omega$, $t_1, \dots, t_k \in \mathbb{N}$, and $x_1, \dots, x_k \in \{0, 1\}$. For each finite action sequence $a^t \in \{0, 1\}^t$, Bayesian updating based on \mathbb{P} induces an 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](#)).

Define the filtration $\mathcal{I}_t := \sigma(a_1, \dots, a_t) \subseteq \mathcal{A}$ 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 $E \subseteq \Omega$, Lévy's upwards theorem applied to the indicator function on E guarantees that as $t \rightarrow \infty$, $H(E|a^t) = \mathbb{E}[1_{\omega \in E} | \mathcal{I}_t] \rightarrow \mathbb{E}[1_{\omega \in E} | \mathcal{I}_\infty] =$

$H(E|a^\infty)$ holds \mathbb{P} -almost surely (see Corollary 5.2.4 in [Stroock, 2010](#)). By standard arguments, this implies that $H(\cdot|a^t)$ converges weakly to $H(\cdot|a^\infty)$ \mathbb{P} -almost surely.

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 $\bar{\mu}_t^\omega \in \Delta(\Delta(\Omega))$ 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 $\bar{\mu}_t^\omega$ converges weakly, with limit $\bar{\mu}_\infty^\omega \in \Delta(\Delta(\Omega))$ given by $\bar{\mu}_\infty^\omega(\mathcal{H}) = \mathbb{P}[H(\cdot|a^\infty) \in \mathcal{H}|\omega]$ for any Borel set $\mathcal{H} \subseteq \Delta(\Omega)$. \square

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

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

Proof. Fix any $\omega \in \Omega^*$. It suffices to prove the following claim: For any non-empty intervals $E, E' \subseteq \Omega$ with E closed in Ω and E' open in Ω 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

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

To see that this claim implies Lemma A.3, note that for each H such that $\omega \notin \text{supp}H$, we have for some n that either $H \in \mathcal{H}_n^+ := \{H' : H'([\omega, \omega + \frac{1}{n}]) = 0 \text{ and } H'((\omega + \frac{1}{n}, \bar{\omega}]) > 0\}$ 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 $\bar{\mu}_\infty^\omega$, $\bar{\mu}_\infty^\omega(\{H : \text{supp}H \not\ni \omega\}) = 0$.

To prove the claim, we only consider case (i); the proof for case (ii) is analogous. Consider any agent and conditional on realized state ω , let (H_t) denote the process of his hypothetical 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 $\bar{\mu}_\infty^\omega$.

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

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

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 , $W_t \rightarrow W_\infty$ almost surely. Since $W_\infty \in L^1$, $W_\infty < +\infty$ almost surely. Thus, almost surely

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

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

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

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

Proof. Recall from (4) that $q_t(\omega) = \int \int F(\theta^*(H^s)) d\Phi(s|\omega) d\bar{\mu}_t^\omega(H)$, where $\int F(\theta^*(H^s)) d\Phi(s|\omega)$ is continuous and bounded in H . Thus, since $\bar{\mu}_t^\omega$ weakly converges to $\bar{\mu}_\infty^\omega$ on Ω^* , $\lim_t q_t(\omega)$ exists for all $\omega \in \Omega^*$ and is given by

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

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

We have

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

For the first inequality, note that since $\bar{\mu}_\infty^{\omega^*} \neq \delta_{\delta_{\omega^*}}$, Lemma A.3 implies that $\bar{\mu}_\infty^{\omega^*}$ 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, we have $\int \int F(\theta^*(H^s)) d\Phi(s|\omega^{**}) d\bar{\mu}_t^{\omega^*}(H) < \int \int F(\theta^*(H^s)) d\Phi(s|\omega^{**}) d\bar{\mu}_t^{\omega^{**}}(H)$ for all t . \square

A.2 Completing the Proof

To complete the proof, we show that in all states, almost all 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. $\bar{\mu}_\infty^\omega = \delta_{\delta_\omega}$ for each $\omega \in \Omega^*$.

Proof. Fix any $\omega \in \Omega^*$ and any interval $E' := [\omega_-, \omega_+] \not\ni \omega$. Consider any 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 $H_t(E') \rightarrow 0$ almost surely.⁵⁰

⁵⁰Indeed, since process (H_t) almost surely weakly converges and its limit is distributed according to $\bar{\mu}_\infty^\omega$, the Portmanteau theorem then implies that $\bar{\mu}_\infty^\omega(\{H : H(E^o) = 0\}) = 1$, where E^o denotes the interior of E' in Ω . In particular, letting $\mathcal{H}_n^+ = \{H : H((\omega + \frac{1}{n}, \bar{\omega}]) > 0\}$ and $\mathcal{H}_n^- = \{H : H([\underline{\omega}, \omega - \frac{1}{n}]) > 0\}$, we have $\bar{\mu}_\infty^\omega(\mathcal{H}_n^+) = \bar{\mu}_\infty^\omega(\mathcal{H}_n^-) = 0$ for each n . Thus, by countable additivity, $\bar{\mu}_\infty^\omega(\{H : H \neq \delta_\omega\}) = \bar{\mu}_\infty^\omega(\bigcup_n (\mathcal{H}_n^+ \cup \mathcal{H}_n^-)) = 0$, as required.

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 lower-bound of E' . Pick any $\omega^* \in \Omega^*$ such that $\omega < \omega^* < \omega_-$ and let $E := [\omega, \omega^*]$.

Now consider the process of likelihood ratios $W_t := \frac{H_t(E')}{H_t(E)}$. As in the proof of Lemma A.3, this process is a non-negative supermartingale. Thus, there is an L^1 -random variable W_∞ such that $W_t \rightarrow W_\infty$ almost surely. For each t , since $q_t(\cdot)$ is decreasing (Lemma A.1), we have

$$\frac{W_{t+1}}{W_t} = \frac{q_t(E')}{q_t(E)} \leq \frac{\sup_{\omega' \in E'} q_t(\omega')}{\inf_{\omega' \in E} q_t(\omega')} = \frac{q_t(\omega_-)}{q_t(\omega^*)}$$

if $a_t = 0$ and

$$\frac{W_{t+1}}{W_t} = \frac{1 - q_t(E')}{1 - q_t(E)} \geq \frac{\inf_{\omega' \in E'} 1 - q_t(\omega')}{\sup_{\omega' \in E} 1 - q_t(\omega')} = \frac{1 - q_t(\omega_-)}{1 - q_t(\omega^*)}$$

if $a_t = 1$. Since $\lim_{t \rightarrow \infty} q_t(\omega_-) < \lim_{t \rightarrow \infty} q_t(\omega^*)$ by Lemma A.4, there exist $\gamma > 0$ and T such that

$$\left| \frac{W_{t+1}}{W_t} - 1 \right| \geq \gamma$$

for all $t \geq T$ at all histories. This ensures that $W_\infty = 0$ almost surely (as otherwise the convergence $W_t \rightarrow W_\infty$ does not occur). Therefore, $H_t(E') \rightarrow 0$ almost surely, as required. \square

Lemma A.5 implies that $q_\infty(\omega) = F(\theta^*(\omega))$ and μ_t^ω 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

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

for any $\omega', \omega'' \in \Omega^*$ such that $\omega' < \omega < \omega''$. If $\omega \in (\underline{\omega}, \bar{\omega})$ 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 $q_t(\omega) \rightarrow F(\theta^*(\omega))$. For boundary points $\omega \in \{\underline{\omega}, \bar{\omega}\}$, the same argument shows that $\liminf_{t \rightarrow \infty} q_t(\underline{\omega}) \geq F(\theta^*(\underline{\omega}))$ and $\limsup_{t \rightarrow \infty} q_t(\bar{\omega}) \leq F(\theta^*(\bar{\omega}))$. Since $q_t(\cdot) \in [F(\theta^*(\bar{\omega})), F(\theta^*(\underline{\omega}))]$, this again implies $q_t(\underline{\omega}) \rightarrow F(\theta^*(\underline{\omega}))$ and $q_t(\bar{\omega}) \rightarrow F(\theta^*(\bar{\omega}))$. Thus, $q_\infty(\cdot) := \lim_{t \rightarrow \infty} q_t(\cdot) = F(\theta^*(\cdot))$ 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 $\bar{\mu}_\infty^\omega = \delta_{\delta_\omega}$. Hence, μ_t^ω weakly converges to δ_{δ_ω} .

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 $t \rightarrow \infty$.

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 certainty 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, $\hat{q}_\infty(\omega) := \lim_{t \rightarrow \infty} \hat{q}_t(\omega) = \hat{F}(\theta^*(\omega))$ 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 $\Omega = [\underline{\omega}, \bar{\omega}]$, it follows that $\hat{q}_t(\cdot)$ converges to $\hat{F}(\theta^*(\cdot))$ uniformly.

Let $H_t(\cdot | a^{t-1}, s) \in \Delta(\Omega)$ 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, $H_t(\cdot | a^{t-1}, s)$ has full support over Ω with positive density $h_t(\cdot | a^{t-1}, s)$ for all a^{t-1} and s . For each pair of states ω', ω'' , denote the corresponding **log likelihood ratio** by

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

Let $\ell_1(\omega', \omega'')$ denote the log likelihood ratio based on the prior belief.

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 $\underline{C} < 0 < \bar{C}$ such that $\hat{q}_\infty(\omega) \pm \nu^* \in (0, 1)$ for all ω and such that for all $\nu \in (0, \nu^*)$,

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

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

$$\ell_t(\omega', \omega'' \mid a^{t-1}, s) \geq \ell_1(\omega', \omega'') + \log \frac{\phi(s \mid \omega')}{\phi(s \mid \omega'')} + (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t})\Delta^\nu(1 - \bar{a}^{t-1}, \hat{q}_\infty(\omega'), \hat{q}_\infty(\omega'')).$$

Proof. Since \hat{q}_t converges to \hat{q}_∞ 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 $a^{t-1} = (a_1, \dots, a_{t-1})$. Then

$$\begin{aligned} \ell_t(\omega', \omega'' \mid a^{t-1}, s) - \ell_1(\omega', \omega'') - \log \frac{\phi(s \mid \omega')}{\phi(s \mid \omega'')} &= \sum_{\tau=1}^{t-1} \left((1 - a_\tau) \log \frac{\hat{q}_\tau(\omega')}{\hat{q}_\tau(\omega'')} + a_\tau \log \frac{1 - \hat{q}_\tau(\omega')}{1 - \hat{q}_\tau(\omega'')} \right) \\ &\geq \underline{C}(\hat{t} - 1) + \sum_{\tau=\hat{t}}^{t-1} \left((1 - a_\tau) \log \frac{\hat{q}_\tau(\omega')}{\hat{q}_\tau(\omega'')} + a_\tau \log \frac{1 - \hat{q}_\tau(\omega')}{1 - \hat{q}_\tau(\omega'')} \right) \\ &\geq \underline{C}(\hat{t} - 1) + \sum_{\tau=\hat{t}}^{t-1} \left((1 - a_\tau) \log \frac{\hat{q}_\infty(\omega') - \nu}{\hat{q}_\infty(\omega'') + \nu} + a_\tau \log \frac{1 - \hat{q}_\infty(\omega') - \nu}{1 - \hat{q}_\infty(\omega'') + \nu} \right) \\ &> \underline{C}(\hat{t} - 1) + \sum_{\tau=1}^{t-1} \left((1 - a_\tau) \log \frac{\hat{q}_\infty(\omega') - \nu}{\hat{q}_\infty(\omega'') + \nu} + a_\tau \log \frac{1 - \hat{q}_\infty(\omega') - \nu}{1 - \hat{q}_\infty(\omega'') + \nu} \right) - \overline{C}(\hat{t} - 1) \\ &= (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t}) \left((1 - \bar{a}^{t-1}) \log \frac{\hat{q}_\infty(\omega') - \nu}{\hat{q}_\infty(\omega'') + \nu} + \bar{a}^{t-1} \log \frac{1 - \hat{q}_\infty(\omega') - \nu}{1 - \hat{q}_\infty(\omega'') + \nu} \right) \\ &= (\underline{C} - \overline{C})(\hat{t} - 1) + (t - \hat{t})\Delta^\nu(1 - \bar{a}^{t-1}, \hat{q}_\infty(\omega'), \hat{q}_\infty(\omega'')), \end{aligned}$$

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

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

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

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

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

Using Lemma B.1, we now show that for large t , the agent's belief is approximately a point mass on the state $\hat{\omega} = \min_{\hat{\omega}'} \text{KL}(1 - \bar{a}^{t-1}, \hat{q}_\infty(\hat{\omega}'))$ 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 $\underline{s} \leq \bar{s}$ and $R \in (0, 1)$. Let $\hat{\omega} := \operatorname{argmin}_{\hat{\omega}'} \text{KL}(R, \hat{q}_\infty(\hat{\omega}'))$. Then for every

interval $E \ni \hat{\omega}$ of states with non-empty interior, there exists $\varepsilon > 0$ such that

$$\liminf_{t \rightarrow \infty} \{H_t(E | a^{t-1}, s) : s \in [\underline{s}, \bar{s}], 1 - \bar{a}^{t-1} \in [R - \varepsilon, R + \varepsilon]\} = 1.$$

Proof. Note that since the agent's posterior admits a positive density, $H_t(E | a^{t-1}, s) = H_t(E^\circ | a^{t-1}, s)$ 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 $\underline{\omega} \leq \alpha_1 < \alpha_2 \leq \bar{\omega}$. 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. $\hat{\omega} = \alpha_2 = \bar{\omega}$ 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.⁵¹ Moreover, in case 1, we can assume that $\underline{\omega} < \alpha_1 < \hat{\omega} < \alpha_2 < \bar{\omega}$, 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

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

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

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

Take $M > 0$ such that $-M \leq \ell_1(\omega', \omega'') + \log \frac{\phi(s|\omega')}{\phi(s|\omega'')}$ for all $\omega', \omega'' \in [\underline{\omega}, \bar{\omega}]$ and all $s \in [\underline{s}, \bar{s}]$. Let \hat{t} be the cutoff given by Lemma B.1. Then for all $t \geq \hat{t}$, $\omega' \in [\hat{\omega}, \hat{q}_\infty^{-1}(R - \xi)]$, $\omega'' \in \hat{q}_\infty^{-1}([0, R_2])$, $s \in [\underline{s}, \bar{s}]$, and any a^{t-1} such that $1 - \bar{a}^{t-1} \in [R - \varepsilon, R + \varepsilon]$, we have

$$\begin{aligned} \ell_t(\omega', \omega'' | a^{t-1}, s) &\geq \ell_1(\omega', \omega'') + \log \frac{\phi(s | \omega')}{\phi(s | \omega'')} + (\underline{C} - \bar{C})(\hat{t} - 1) + (t - \hat{t})\Delta^\nu(1 - \bar{a}^{t-1}, \hat{q}_\infty(\omega'), \hat{q}_\infty(\omega'')) \\ &\geq -M + (\underline{C} - \bar{C})(\hat{t} - 1) + (t - \hat{t})\rho, \end{aligned}$$

where the first inequality holds by Lemma B.1 and the second inequality holds by choice of ν and M above. Likewise, for all $t \geq \hat{t}$, $\omega' \in [\hat{q}_\infty^{-1}(R + \xi), \hat{\omega}]$, $\omega'' \in \hat{q}_\infty^{-1}([R_1, 1])$, $s \in [\underline{s}, \bar{s}]$, and any a^{t-1}

⁵¹In case 2, we choose $\xi, \varepsilon, \rho > 0$ such that $\hat{q}_\infty(\bar{\omega}) < R + \xi < R_1$ and

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

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

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

The remaining steps are analogous to case 1.

such that $1 - \bar{a}^{t-1} \in [R - \varepsilon, R + \varepsilon]$, we have

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

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

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

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

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

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

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 $\text{Pr}(\cdot \mid \omega)$ 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** $\text{SS}(F, \hat{F}) := \{\hat{\omega}_\infty \in \Omega : \hat{\omega}_\infty = \text{argmin}_{\hat{\omega} \in \Omega} \text{KL}(F(\theta^*(\hat{\omega}_\infty)), \hat{F}(\theta^*(\hat{\omega})))\}$.

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

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

Proof. Fix any $\omega \in \Omega$. Let $\bar{R}(\omega) := \limsup_{t \rightarrow \infty} \bar{q}_t(\omega)$ and $\underline{R}(\omega) := \liminf_{t \rightarrow \infty} \bar{q}_t(\omega)$. To show that $R(\omega) := \lim_{t \rightarrow \infty} \bar{q}_t(\omega)$ exists, suppose for a contradiction that $\bar{R}(\omega) > \underline{R}(\omega)$. Since $\text{SS}(F, \hat{F})$ is finite, we can pick some $R \in (\underline{R}(\omega), \bar{R}(\omega))$ such that $\hat{q}_\infty^{-1}(R) \notin \text{SS}(F, \hat{F})$. Let $\hat{\omega} := \text{argmin}_{\hat{\omega}' \in \Omega} \text{KL}(R, \hat{q}_\infty(\hat{\omega}'))$. 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}, \bar{R} such that $\underline{R}(\omega) < \underline{R} < R < \bar{R} < \bar{R}(\omega)$ and $\bar{R} < F(\theta^*(\hat{\omega}))$. Note that we can choose a small enough interval $E \ni \hat{\omega}$ with non-empty interior, a large enough interval $[\underline{s}, \bar{s}]$ 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

$s \in [\underline{s}, \bar{s}]$ hold beliefs such that $H_t(E \mid a^{t-1}, s) > 1 - \nu$, then $q_t(\omega) > \bar{R} + \nu$.⁵² By Lemma B.2, there exists $\varepsilon > 0$ and \hat{t}_1 such that for all $t \geq \hat{t}_1$,

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

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

Note that each agent observes a sequence of random actions (a_1, a_2, \dots) 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 $\mathcal{T} := \{t \geq \max\{\hat{t}_1, \hat{t}_2\} : \bar{q}_{t-1}(\omega) \in [R - \frac{\varepsilon}{2}, R + \frac{\varepsilon}{2}]\}$. Then at all times $t \in \mathcal{T}$, at least fraction $1 - \nu$ of agents with private signals $s \in [\underline{s}, \bar{s}]$ 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 $q_{t_k}(\omega) > \bar{R} + \nu > R - \varepsilon$ by construction of \mathcal{T} . This is a contradiction. Hence, $R(\omega) = \lim_{t \rightarrow \infty} \bar{q}_t(\omega)$ 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 $R(\omega) = \lim_{t \rightarrow \infty} \bar{q}_t(\omega)$ is weakly decreasing in ω , as required. \square

We now prove that in all states, agents' long-run beliefs assign probability 1 to some steady state:

Proposition B.1. *Suppose that $\text{SS}(F, \hat{F})$ is finite. Then in all states ω , there exists some state $\hat{\omega}_\infty(\omega) \in \text{SS}(F, \hat{F})$ 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 $R(\omega) := \lim_{t \rightarrow \infty} \bar{q}_t(\omega)$, which exists by Lemma B.3. Define

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

Note that $\hat{\omega}_\infty(\omega)$ is weakly increasing in ω since $R(\omega)$ is weakly decreasing and \hat{q}_∞ is strictly decreasing. Consider any interval $E \ni \hat{\omega}_\infty(\omega)$ with non-empty interior and any $\underline{s} < \bar{s}$. By Lemma B.2,

⁵²To see this, observe that if all agents' beliefs assigned probability 1 to $\hat{\omega}$ at t , then $q_t(\omega) = F(\theta^*(\hat{\omega})) > \bar{R}$.

there exists $\varepsilon > 0$ such that

$$\liminf_{t \rightarrow \infty} \{H_t(E \mid a^{t-1}, s) \mid s \in [\underline{s}, \bar{s}], 1 - \bar{a}^{t-1} \in [R(\omega) - \varepsilon, R(\omega) + \varepsilon]\} = 1.$$

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

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

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

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

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

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

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

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, $\operatorname{SS}(F, \hat{F}) = \{\hat{\omega}\}$. Thus, Proposition B.1 implies that in all states ω , almost all agents' beliefs converge to a point-mass on $\hat{\omega}$. \square

B.4 Proof of Theorem 2

Fix any analytic $F, \hat{F} \in \mathcal{F}$ with $\hat{F} \neq F$. Then the set $\{\theta \in [\theta^*(\bar{\omega}), \theta^*(\underline{\omega})] : F(\theta) = \hat{F}(\theta)\}$ is finite (possibly empty).⁵³ But this implies that $\operatorname{SS}(F, \hat{F})$ is finite, as every $\hat{\omega}_\infty \in \operatorname{SS}(F, \hat{F})$ satisfies either $F(\theta^*(\hat{\omega}_\infty)) = \hat{F}(\theta^*(\hat{\omega}_\infty))$ or $\hat{\omega}_\infty \in \{\underline{\omega}, \bar{\omega}\}$. 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. \square

⁵³To see this, suppose for a contradiction that $F - \hat{F} = 0$ admits an infinite sequence $\theta_1, \theta_2, \dots$ of distinct solutions in $\Theta^* := [\theta^*(\bar{\omega}), \theta^*(\underline{\omega})]$. 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

Fix any Ω_n and any true type distribution $F \in \mathcal{F}$. For any perception $\hat{F} \in \mathcal{F}$, let $q_t(\omega; \hat{F})$ and $\hat{q}_t(\omega; \hat{F})$ denote the true and perceived fractions of action in 0 in period t and state $\omega \in \Omega_n$, and let $\bar{q}_t(\omega; \hat{F}) := \frac{1}{t} \sum_{\tau=1}^t q_\tau(\omega; \hat{F})$ and $\hat{\bar{q}}_t(\omega; \hat{F}) := \frac{1}{t} \sum_{\tau=1}^t \hat{q}_\tau(\omega; \hat{F})$ denote the true and perceived time averages up to period t . 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 private signal s and observed action sequence $a^{t-1} = (a_1, \dots, a_{t-1})$.

Take $\nu \in (0, 1)$ sufficiently small that $F(\theta^*(\omega)) \pm 2\nu \in (0, 1)$ for each $\omega \in \Omega_n$, and such that for all pairs of distinct states $\omega, \omega' \in \Omega_n$ and all $R \in [F(\theta^*(\omega)) - 2\nu, F(\theta^*(\omega)) + 2\nu]$,

$$R \log \frac{F(\theta^*(\omega)) - \nu}{F(\theta^*(\omega')) + \nu} + (1 - R) \log \frac{1 - F(\theta^*(\omega)) - \nu}{1 - F(\theta^*(\omega')) + \nu} > \nu \quad (5)$$

Take \underline{C}, \bar{C} such that

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

Our proof uses the following three lemmas:

Lemma C.1. *There exists $\varepsilon_1 \in (0, \nu)$ such that for any $\omega \in \Omega_n$, any t , and any perception \hat{F} with $\|F - \hat{F}\| \leq \varepsilon_1$, if at least fraction $1 - \varepsilon_1$ of agents' posteriors in period t assign probability at least $1 - \varepsilon_1$ to state ω , then $q_{t+1}(\omega; \hat{F}), \hat{q}_{t+1}(\omega; \hat{F}) \in [F(\theta^*(\omega)) - \nu, F(\theta^*(\omega)) + \nu]$.*

Proof. The result follows from a simple continuity argument based on the continuity of u and F . \square

Throughout the rest of the proof, we fix ε_1 as given by Lemma C.1.

Lemma C.2. *Fix any signals $\underline{s} < \bar{s}$, time $\hat{t} > 0$, and $\eta \in (0, 1)$. There exists a time $T_1 = T_1(\hat{t}, \eta, \underline{s}, \bar{s}) > \hat{t}$ such that for any $\omega \in \Omega_n$, private signal $s \in [\underline{s}, \bar{s}]$, and $t \geq T_1$, if perception \hat{F} and observed action history $a^{t-1} = (a_1, \dots, a_{t-1})$ satisfy*

- (i) $\|F - \hat{F}\| \leq \varepsilon_1$ and $q_\tau(\omega'; \hat{F}) \in [F(\theta^*(\omega')) - \nu, F(\theta^*(\omega')) + \nu]$ for each $\tau \in \{\hat{t}, \dots, t\}$ and $\omega' \in \Omega_n$,
- (ii) $\bar{a}^{t-1} \in [F(\theta^*(\omega)) - 2\nu, F(\theta^*(\omega)) + 2\nu]$,

then the posterior probability $H_t(\omega|s, a^{t-1}; \hat{F})$ on state ω is at least η .

Proof. Pick $T_1 = T_1(\hat{t}, \eta, \underline{s}, \bar{s}) > \hat{t}$ sufficiently large that for all $s \in [\underline{s}, \bar{s}]$ and distinct $\omega, \omega' \in \Omega_n$, we have $(\bar{C} - \underline{C})(\hat{t} - 1) + (T_1 - \hat{t})\nu + \ell_1(\omega, \omega') + \log \frac{\phi(s|\omega)}{\phi(s|\omega')} \geq \log \frac{\eta}{1-\eta}$, where $\ell_1(\omega, \omega')$ denotes the log likelihood ratio of ω vs. ω' according to the prior. Now consider any $\omega \in \Omega_n$, $s \in [\underline{s}, \bar{s}]$, $t \geq T_1$, and \hat{F} and a^{t-1} satisfying assumptions (i)–(ii). For any $\omega' \neq \omega$, we will show that the log likelihood

ratio $\ell_t(\omega, \omega' | s, a^{t-1}; \hat{F}) = \log \frac{H_t(\omega | s, a^{t-1}; \hat{F})}{H_t(\omega' | s, a^{t-1}; \hat{F})}$ satisfies $\ell_t(\omega, \omega' | s, a^{t-1}; \hat{F}) \geq \log \frac{\eta}{1-\eta}$. The argument is analogous to Lemma B.1. Indeed, we have

$$\begin{aligned}
& \ell_t(\omega, \omega' | a^{t-1}, s; \hat{F}) - \ell_1(\omega, \omega') - \log \frac{\phi(s|\omega)}{\phi(s|\omega')} = \sum_{\tau=1}^{t-1} \left((1-a_\tau) \log \frac{\hat{q}_\tau(\omega; \hat{F})}{\hat{q}_\tau(\omega'; \hat{F})} + a_\tau \log \frac{1-\hat{q}_\tau(\omega; \hat{F})}{1-\hat{q}_\tau(\omega'; \hat{F})} \right) \\
& \geq \underline{C}(\hat{t}-1) + \sum_{\tau=\hat{t}}^{t-1} \left((1-a_\tau) \log \frac{\hat{q}_\tau(\omega; \hat{F})}{\hat{q}_\tau(\omega'; \hat{F})} + a_\tau \log \frac{1-\hat{q}_\tau(\omega; \hat{F})}{1-\hat{q}_\tau(\omega'; \hat{F})} \right) \\
& \geq \underline{C}(\hat{t}-1) + \sum_{\tau=\hat{t}}^{t-1} \left((1-a_\tau) \log \frac{F(\theta^*(\omega)) - \nu}{F(\theta^*(\omega')) + \nu} + a_\tau \log \frac{1-F(\theta^*(\omega)) - \nu}{1-F(\theta^*(\omega')) + \nu} \right) \\
& \geq \underline{C}(\hat{t}-1) + \sum_{\tau=1}^{t-1} \left((1-a_\tau) \log \frac{F(\theta^*(\omega)) - \nu}{F(\theta^*(\omega')) + \nu} + a_\tau \log \frac{1-F(\theta^*(\omega)) - \nu}{1-F(\theta^*(\omega')) + \nu} \right) - \overline{C}(\hat{t}-1) \\
& = (\overline{C} - \underline{C})(\hat{t}-1) + (t-\hat{t}) \left((1-\bar{a}^t) \log \frac{F(\theta^*(\omega)) - \nu}{F(\theta^*(\omega')) + \nu} + \bar{a}^t \log \frac{1-F(\theta^*(\omega)) - \nu}{1-F(\theta^*(\omega')) + \nu} \right) \\
& \geq (\overline{C} - \underline{C})(\hat{t}-1) + (t-\hat{t})\nu.
\end{aligned}$$

Here, the first inequality holds because for all τ , we have $(1-a_\tau) \log \frac{\hat{q}_\tau(\omega; \hat{F})}{\hat{q}_\tau(\omega'; \hat{F})} + a_\tau \log \frac{1-\hat{q}_\tau(\omega; \hat{F})}{1-\hat{q}_\tau(\omega'; \hat{F})} \geq \min \left\{ \log \frac{F(\theta^*(\bar{\omega})) - \varepsilon_1}{F(\theta^*(\bar{\omega})) + \varepsilon_1}, \log \frac{1-F(\theta^*(\underline{\omega})) - \varepsilon_1}{1-F(\theta^*(\bar{\omega})) + \varepsilon_1} \right\} \geq \underline{C}$. The second inequality follows from assumption (i). The third inequality holds because for all τ , we have $(1-a_\tau) \log \frac{F(\theta^*(\omega)) - \nu}{F(\theta^*(\omega')) + \nu} + a_\tau \log \frac{1-F(\theta^*(\omega)) - \nu}{1-F(\theta^*(\omega')) + \nu} \leq \max \left\{ \log \frac{F(\theta^*(\omega)) - \nu}{F(\theta^*(\bar{\omega})) + \nu}, \log \frac{1-F(\theta^*(\bar{\omega})) - \nu}{1-F(\theta^*(\omega')) + \nu} \right\} \leq \overline{C}$. The final inequality holds by (5) and assumption (ii).

Thus, by choice of T_1 , this implies $\ell_t(\omega, \omega' | s, a^{t-1}; \hat{F}) \geq \log \frac{\eta}{1-\eta}$, as claimed. \square

Lemma C.3. *For any $\eta \in (0, 1)$, there exists $T_2(\eta) > 0$ such that for any $t \geq T_2(\eta)$, any true state $\omega \in \Omega_n$, and any perception \hat{F} , both the true fraction of agents who observe action histories with $\bar{a}^{t-1} \in [\bar{q}_t(\omega; \hat{F}) - \nu, \bar{q}_t(\omega; \hat{F}) + \nu]$ and the perceived fraction of agents who observe action histories with $\bar{a}^{t-1} \in [\hat{q}_t(\omega; \hat{F}) - \nu, \hat{q}_t(\omega; \hat{F}) + \nu]$ is at least η .*

Proof. The result follows from Hoeffding's inequality since true (resp. perceived) action observations a_1, a_2, \dots are drawn independently across t according to $q_t(\omega; \hat{F})$ (resp. $\hat{q}_t(\omega; \hat{F})$). \square

To complete the proof of the first part of Theorem 3, fix any true state $\omega \in \Omega_n$. Take signals $\underline{s}^* < \bar{s}^*$ such that $\Phi(\bar{s}^* | \omega) - \Phi(\underline{s}^* | \omega) \geq \sqrt{1 - \varepsilon_1}$. Let $T_2(\sqrt{1 - \varepsilon_1})$ be as given by Lemma C.3. Note that when $\hat{F} = F$, the proof of Lemma 1 continues to ensure almost all agents' beliefs converge to δ_ω . Thus, there exists $\hat{t} \geq T_2(\sqrt{1 - \varepsilon_1})$ such that for all $t \geq \hat{t}$, we have

$$q_t(\omega; F), \bar{q}_t(\omega; F) \in [F(\theta^*(\omega)) - \frac{\nu}{2}, F(\theta^*(\omega)) + \frac{\nu}{2}].$$

Let $T_1(\hat{t}, 1 - \varepsilon_1, \underline{s}^*, \bar{s}^*)$ be as given by Lemma C.2. Then, by a continuity argument, there exists $\varepsilon_2 \in (0, \varepsilon_1]$ such that for all \hat{F} with $\|F - \hat{F}\| \leq \varepsilon_2$ and all $t \in \{\hat{t}, \dots, T_1(\hat{t}, 1 - \varepsilon_1, \underline{s}^*, \bar{s}^*)\}$,

$$q_t(\omega; \hat{F}), \bar{q}_t(\omega; \hat{F}), \hat{q}_t(\omega; \hat{F}), \hat{\hat{q}}_t(\omega; \hat{F}) \in [F(\theta^*(\omega)) - \nu, F(\theta^*(\omega)) + \nu]. \quad (6)$$

We show by induction that (6) remains valid for all $t \geq T_1(\hat{t}, 1 - \varepsilon_1, \underline{s}^*, \bar{s}^*)$. Fix any \hat{F} with $\|F - \hat{F}\| \leq \varepsilon_2$, and suppose (6) holds up to some $t \geq T_1(\hat{t}, 1 - \varepsilon_1, \underline{s}^*, \bar{s}^*)$. Since $t + 1 \geq T_2(\sqrt{1 - \varepsilon_1})$, Lemma C.3 implies that both the true and perceived fraction of agents with observed action frequency $\bar{a}^t \in [F(\theta^*(\omega)) - 2\nu, F(\theta^*(\omega)) + 2\nu]$ is at least $\sqrt{1 - \varepsilon_1}$. By the inductive hypothesis and Lemma C.2, each such agent's period $t + 1$ belief assigns probability at least $1 - \varepsilon_1$ to ω if his private signal was some $s \in [\underline{s}^*, \bar{s}^*]$. By choice of $\underline{s}^*, \bar{s}^*$, the fraction of agents with $s \in [\underline{s}^*, \bar{s}^*]$ is at least $\sqrt{1 - \varepsilon_1}$. Thus, fraction at least $1 - \varepsilon_1$ of agents assign probability at least $1 - \varepsilon_1$ to ω in period $t + 1$. Hence, by Lemma C.1, we have $q_{t+1}(\omega; \hat{F}), \hat{q}_{t+1}(\omega; \hat{F}) \in [F(\theta^*(\omega)) - \nu, F(\theta^*(\omega)) + \nu]$, which along with (6) at t also ensures $\bar{q}_{t+1}(\omega; \hat{F}), \hat{\bar{q}}_{t+1}(\omega; \hat{F}) \in [F(\theta^*(\omega)) - \nu, F(\theta^*(\omega)) + \nu]$, as claimed.

Finally, consider any perception \hat{F} with $\|F - \hat{F}\| \leq \varepsilon_2$ and any $\eta \in (0, 1)$. Pick signals $\underline{s} < \bar{s}$ such that $\Phi(\bar{s}|\omega) - \Phi(\underline{s}|\omega) \geq \sqrt{\eta}$. Since (6) holds for all $t \geq \hat{t}$, Lemma C.3 implies that for all $t \geq \max\{\hat{t}, T_2(\sqrt{\eta})\}$ at least fraction $\sqrt{\eta}$ of agents observe action frequency $\bar{a}^{t-1} \in [F(\theta^*(\omega)) - 2\nu, F(\theta^*(\omega)) + 2\nu]$. Therefore, by (6) for $t \geq \hat{t}$ and Lemma C.2, any such agent's posterior at all $t \geq \max\{T_1(\hat{t}, \eta, \underline{s}, \bar{s}), \hat{t}, T_2(\sqrt{\eta})\}$ assigns probability at least η on ω if his private signal s was in $[\underline{s}, \bar{s}]$. Thus, by choice of \underline{s}, \bar{s} , fraction at least η of agents assign probability at least η to ω at all $t \geq \max\{T_1(\hat{t}, \eta, \underline{s}, \bar{s}), \hat{t}, T_2(\sqrt{\eta})\}$. Since η was arbitrary, this completes the proof. \square

C.2 Theorem 3: Proof of Part 2

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

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

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 $\hat{\omega}_\infty(\omega) \in \text{SS}_n(F, \hat{F})$ 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 $\{\omega^1, \omega^2, \dots\}$ 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 $\Omega_N = \{\omega^1, \dots, \omega^N\}$.

Consider any state space Ω_n with $n \geq N$. We claim that $\text{SS}_n(F, \hat{F}) \subseteq [\hat{\omega} - \varepsilon, \hat{\omega} + \varepsilon]$. 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 $\hat{\omega}_\infty \notin \text{SS}_n(F, \hat{F})$. 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 $\hat{\omega}_\infty(\omega) \in \text{SS}_n(F, \hat{F}) \subseteq [\hat{\omega} - \varepsilon, \hat{\omega} + \varepsilon]$, as claimed. \square

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.
- 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 (2019): “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,” .
- 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.
- 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. (2008): “Behavioral equilibrium in economies with adverse selection,” *American Economic Review*, 98(4), 1269–91.
- ESPONDA, I., AND D. POUZO (2016): “Berk–Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models,” *Econometrica*, 84(3), 1093–1130.
- (2019): “Equilibrium in misspecified Markov decision processes,” working paper.
- EVANS, G. W., AND S. HONKAPOHJA (2012): *Learning and expectations in macroeconomics*. Princeton University Press.
- EYSTER, E., AND M. RABIN (2010): “Naive herding in rich-information settings,” *American economic journal: microeconomics*, 2(4), 221–43.
- FRICK, M., R. IJIMA, AND Y. ISHII (2019a): “Dispersed Behavior and Perceptions in Assortative Societies,” Cowles Foundation Discussion Paper #2128.
- (2019b): “Misinterpreting Others and the Fragility of Social Learning,” Cowles Foundation Discussion Paper #2160.
- (2019c): “Stability and Robustness in Misspecified Learning Models,” 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.
- 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.
- HASSAN, T. A., AND T. M. MERTENS (2017): “The social cost of near-rational investment,” *American Economic Review*, 107(4), 1059–1103.
- 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.
- (2019): “Convergence in Misspecified Learning Models with Endogenous Actions,” *Available at SSRN 3312968*.
- 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.

- KUNDA, Z. (1999): *Social cognition: Making sense of people*. MIT press.
- LEE, I. H. (1993): "On the convergence of informational cascades," *Journal of Economic theory*, 61(2), 395–411.
- 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.
- MORRIS, S., AND H. S. SHIN (2002): "Social value of public information," *American Economic Review*, 92(5), 1521–1534.
- MOSSEL, E., A. SLY, AND O. TAMUZ (2015): "Strategic learning and the topology of social networks," *Econometrica*, 83(5), 1755–1794.
- 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 (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.
- 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.
- 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.
- SPIEGLER, R. (2016): "Bayesian networks and boundedly rational expectations," *The Quarterly Journal of Economics*, 131(3), 1243–1290.
- 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.
- (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.