

DISCUSSION PAPER SERIES

DP16788

Belief Convergence under Misspecified Learning: A Martingale Approach

Mira Frick, Ryota Iijima and Yuhta Ishii

ORGANIZATIONAL ECONOMICS

CEPR

Belief Convergence under Misspecified Learning: A Martingale Approach

Mira Frick, Ryota Iijima and Yuhta Ishii

Discussion Paper DP16788
Published 07 December 2021
Submitted 07 December 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Organizational Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Mira Frick, Ryota Iijima and Yuhta Ishii

Belief Convergence under Misspecified Learning: A Martingale Approach

Abstract

We present an approach to analyze learning outcomes in a broad class of misspecified environments, spanning both single-agent and social learning. We introduce a novel “prediction accuracy” order over subjective models, and observe that this makes it possible to partially restore standard martingale convergence arguments that apply under correctly specified learning. Based on this, we derive general conditions to determine when beliefs in a given environment converge to some long-run belief either locally or globally (i.e., from some or all initial beliefs). We show that these conditions can be applied, first, to unify and generalize various convergence results in previously studied settings. Second, they enable us to analyze environments where learning is “slow,” such as costly information acquisition and sequential social learning. In such environments, we illustrate that even if agents learn the truth when they are correctly specified, vanishingly small amounts of misspecification can generate extreme failures of learning.

JEL Classification: N/A

Keywords: N/A

Mira Frick - mira.frick@gmail.com
Yale University and CEPR

Ryota Iijima - ryota.ijima@yale.edu
Yale University

Yuhta Ishii - yishii1986@gmail.com
Pennsylvania State University

Belief Convergence under Misspecified Learning: A Martingale Approach*

Mira Frick Ryota Iijima Yuhta Ishii

First version: October 1, 2019

This version: December 4, 2021

Abstract

We present an approach to analyze learning outcomes in a broad class of misspecified environments, spanning both single-agent and social learning. We introduce a novel “prediction accuracy” order over subjective models, and observe that this makes it possible to partially restore standard martingale convergence arguments that apply under correctly specified learning. Based on this, we derive general conditions to determine when beliefs in a given environment converge to some long-run belief either locally or globally (i.e., from some or all initial beliefs). We show that these conditions can be applied, first, to unify and generalize various convergence results in previously studied settings. Second, they enable us to analyze environments where learning is “slow,” such as costly information acquisition and sequential social learning. In such environments, we illustrate that even if agents learn the truth when they are correctly specified, vanishingly small amounts of misspecification can generate extreme failures of learning.

*Frick: Yale University (mira.frick@yale.edu); Iijima: Yale University (ryota.iijima@yale.edu); Ishii: Pennsylvania State University (yxi5014@psu.edu). This research was supported by National Science Foundation grant SES-1824324. Previous versions of this paper were circulated under the title “Stability and Robustness in Misspecified Learning Models.” We thank the editor, three anonymous referees, Nageeb Ali, Aislinn Bohren, Christophe Chamley, In-Koo Cho, Ignacio Esponda, Erik Eyster, Drew Fudenberg, Marina Halac, Daniel Hauser, Kevin He, Paul Heidhues, Michihiro Kandori, Giacomo Lanzani, Yves Le Yaouanq, Bart Lipman, Weicheng Min, Pietro Ortoleva, Larry Samuelson, Ron Siegel, Philipp Strack, Yuichi Yamamoto, Leeat Yariv, as well as seminar audiences at Caltech, Harvard-MIT, Hitotsubashi, LMU Munich, Oxford, Princeton, QMUL, RINSE, Seminar in Economic Theory, University of Tokyo, Toronto, Western, and Yale.

1 Introduction

1.1 Motivation and overview

Motivated in part by empirical evidence that individuals face numerous systematic cognitive biases and limitations, a growing literature recognizes the need to enrich classic economic models of single-agent and social learning by allowing for the possibility that agents may hold an incorrect, simplified or, for short, *misspecified* view of the data generating process. Many papers have demonstrated how various forms of misspecification alter learning outcomes in a wide range of economic applications, from learning about the return to effort by a worker who is overconfident in her ability, to social learning about the quality of a new product by consumers who are incorrect about others’ preferences.

Learning dynamics of such models tend to be non-trivial to analyze. A primary reason is that when agents are misspecified, their belief (i.e., posterior ratio) process is no longer a martingale (with respect to the true data generating process), so standard convergence arguments do not apply. The analysis is further complicated by the fact that in most aforementioned settings information depends endogenously on agents’ actions, and hence may be influenced by their misspecification.¹ As a result, much existing work has derived learning outcomes using approaches tailored specifically to each application, while only recently the focus has turned to developing general tools to analyze the asymptotics of misspecified learning dynamics (see Section 1.2 for a discussion of related literature).

This paper contributes to the latter goal by presenting an approach to analyze learning outcomes in a broad class of misspecified environments, spanning both single-agent and social learning. We introduce novel “prediction accuracy” orderings over subjective models that allow one to partially restore the standard martingale convergence method. Based on this, we derive general conditions to determine when beliefs in a given environment converge to some long-run belief either locally or globally (i.e., from some or all initial beliefs). We show that these conditions can be applied, first, to unify and generalize various convergence results in previously studied settings. Second, they enable us to analyze a natural class of environments, including costly information acquisition and sequential social learning, where learning is “slow.” In such environments, we illustrate that even if agents learn the truth when they are correctly specified, vanishingly small amounts of misspecification can generate extreme failures of learning.

To nest a wide range of applications and make the logic of belief convergence transparent, Section 2 sets up an abstract environment, where agents, actions, and preferences are not

¹This contrasts with a literature in statistics that studies learning by a passive observer who receives exogenous signals about which he is misspecified (e.g., Berk, 1966).

explicitly modeled. Instead, we consider a belief process μ_t over some set of states of the world, which from any initial belief μ_0 evolves in the following manner. Each period $t = 0, 1, \dots$, a signal z_t is drawn according to a true signal distribution P_{μ_t} that—capturing endogeneity of signals—may depend on the current belief μ_t . Following the realization of z_t , belief μ_t is updated to μ_{t+1} via Bayes’ rule based on the perception that the signal distribution at each state ω and belief μ_t is $\hat{P}_{\mu_t}(\cdot|\omega)$. Capturing potential misspecification, the true signal distribution need not coincide with any of the perceived distributions. Remark 1 illustrates how leading economic models of single-agent and social learning map into this environment.

Section 3 analyzes belief convergence. We begin by introducing an order over states that compares how well they predict the true signal distribution at any given belief: For any $q > 0$, we say that state ω *q-dominates* state ω' at belief μ if the perceived signal distribution $\hat{P}_{\mu}(\cdot|\omega)$ in state ω comes “closer” to the true distribution P_{μ} than does the perceived distribution $\hat{P}_{\mu}(\cdot|\omega')$ in state ω' . Here closeness is measured using the moment-generating function (evaluated at q) of the perceived log-likelihood ratio of states. This order refines the usual comparison based on Kullback-Leibler divergence, which features prominently in existing analyses of misspecified learning. A simple but key observation is that, throughout any range of beliefs where *q*-dominance obtains, the *q*th power of the posterior ratio process becomes a nonnegative supermartingale. This allows one to locally restore standard martingale convergence arguments from the correctly specified setting, providing a useful approach to analyze asymptotic beliefs.

Building on this observation, we derive conditions that ensure that a given point-mass belief δ_{ω} is (i) *locally stable*, (ii) *globally stable*, or (iii) *unstable*, in the sense that the belief process μ_t converges to δ_{ω} either (i) from any initial belief that is sufficiently close to δ_{ω} , or (ii) from all initial full-support beliefs, or (iii) escapes any small enough neighborhood of δ_{ω} .

By applying the above martingale observation, Theorem 1 shows that δ_{ω} is locally stable if state ω strictly *q*-dominates all other states ω' at all beliefs μ in a neighborhood of δ_{ω} , except possibly at the belief $\mu = \delta_{\omega}$. We provide an analogous condition for instability. The fact that these conditions do not impose *q*-dominance at the point-mass belief δ_{ω} is essential for analyzing environments with slow learning, a property we explain below.

Using martingale arguments, we also obtain two conditions for global stability that strengthen the local stability criterion in Theorem 1 in complementary ways. Theorem 2 shows that δ_{ω} is globally stable if state ω uniquely survives the iterated elimination of (globally) strictly dominated states. Proposition 1 restricts the prediction accuracy ranking only near point-mass beliefs, but imposes more structure on how states are ordered.

Section 4 applies the preceding stability results to two classes of economic applications. Section 4.1 considers single-agent active learning in rich one-dimensional state spaces, as

in many important applications in the literature. We show that the iterated elimination criterion in Theorem 2 is straightforward to verify in this setting and can be used to unify and generalize convergence results in applications such as monopoly pricing with a misspecified demand curve (e.g., Esponda and Pouzo, 2016; Heidhues, Kőszegi, and Strack, 2021), effort choice by an overconfident agent (Heidhues, Kőszegi, and Strack, 2018), and optimal stopping under the gambler’s fallacy (He, 2021).

Section 4.2 studies environments that feature *slow learning*: That is, as agents grow confident in any state, their behavior generates less and less informative new signals, so the speed of belief convergence vanishes near point-mass beliefs. This is a well-known property of several important economic applications: For example, under sequential social learning, later agents’ actions reveal less and less about their private information, as they increasingly base their action choices on the information conveyed by earlier agents’ actions; likewise, under costly information acquisition, an agent may acquire increasingly less precise signals the more confident she becomes. Existing approaches to analyze learning outcomes under misspecification (Section 1.2) do not in general apply to such settings, as these approaches measure prediction accuracy using Kullback-Leibler divergence, which can be too coarse to determine stability/instability when signal informativeness vanishes near point-mass beliefs (see Remark 2). In contrast, our stability results based on q -dominance apply to these settings, and we highlight that slow learning can lead to fragility against misspecification: Even if agents learn the true state when they are correctly specified, vanishingly small amounts of misspecification can generate extreme failures of learning. For example, under social learning about the safety of a new product, if agents even slightly underestimate others’ risk tolerance, then, regardless of the product’s actual safety, long-run beliefs always become confident in the highest possible safety level (Section 4.2.2); similarly, if an agent has even a slight tendency to distort feedback about her ability in an “ego-biased” manner and if acquiring feedback is even slightly costly, then her long-run beliefs will display drastic overconfidence in her ability (Section 4.2.1).

1.2 Related literature

Our paper builds on Esponda and Pouzo (2016), who define a general steady-state notion for misspecified learning dynamics, Berk-Nash equilibrium, nesting other influential steady-state concepts that capture more specific forms of misspecification (e.g., Eyster and Rabin, 2005; Jehiel, 2005; Esponda, 2008; Spiegel, 2016). It is known that, while any locally stable belief is a Berk-Nash equilibrium (Lemma 1 establishes this in our setting), the converse is not in general true (e.g., Nyarko, 1991). We provide stability criteria that determine which Berk-

Nash equilibria learning dynamics in a given environment converge to locally or globally. We also point to natural settings where the set of stable equilibria is not robust to the details of agents’ misspecification. Our martingale approach relies on measuring prediction accuracy using q -dominance, which refines the measure based on Kullback-Leibler divergence that underlies Berk-Nash equilibrium.

Several important earlier papers have examined the convergence of misspecified learning dynamics in a variety of single-agent (e.g., Nyarko, 1991; Schwartzstein, 2014; Fudenberg, Romanyuk, and Strack, 2017; Heidhues, Kőszegi, and Strack, 2018; He, 2021; Bushong and Gagnon-Bartsch, 2019; Cong, 2019; Heidhues, Kőszegi, and Strack, 2021) and social learning settings (e.g., Eyster and Rabin, 2010; Bohren, 2016; Gagnon-Bartsch, 2017; Bohren, Imas, and Rosenberg, 2019). The approaches in these papers are either tailored to particular environments and forms of misspecification or apply in more general settings but rely on specific parametric assumptions (e.g., Gaussian signals in Fudenberg, Romanyuk, and Strack, 2017; Heidhues, Kőszegi, and Strack, 2021).

Our paper contributes to a recent focus in the literature on developing more unified approaches to establish convergence under misspecified learning. In binary-state environments, Bohren and Hauser (2021) provide general conditions for local and global stability of beliefs based on Kullback-Leiber divergence. A key challenge they address is to allow for heterogeneous models across different agents (as is natural under social learning), which we do not consider in this paper.² Instead, relying on our martingale approach based on q -dominance, we derive results that apply to rich state spaces (e.g., Section 4.1) and environments with slow learning (e.g., Section 4.2), to which their methods do not apply. In settings that do not feature slow learning, Bohren and Hauser (2021) show that successful learning is robust to small amounts of misspecification; complementary to this, Section 4.2 sheds light on ways in which slow learning can lead to fragility against misspecification.

In general-state environments, Esponda, Pouzo, and Yamamoto (2021) (EPY) and Fudenberg, Lanzani, and Strack (2021a) (FLS) analyze action convergence under single-agent learning. Unlike our paper, the convergence results in EPY and FLS do not apply to social learning settings or environments with infinite actions;³ at the same time, both papers address important settings/questions to which our results do not apply. In particular, EPY develop a methodology to analyze asymptotic action frequencies based on approximating these by a differential inclusion. Unlike our paper, their paper also characterizes asymptotic action frequencies when beliefs/actions do not converge. FLS provide tight conditions that

²Some of our results can be extended to heterogeneous models; see Appendix G of the previous version Frick, Iijima, and Ishii (2020b).

³The more recent paper by Murooka and Yamamoto (2021) extends EPY to settings with infinite actions and/or strategic externalities, but also does not consider social learning.

relate action convergence to the agent’s payoffs, while our conditions for belief convergence do not explicitly involve the agent’s incentives; their convergence proofs build on the martingale approach we introduce in this paper. To analyze the agent’s forward-looking incentives, FLS also derive results on the rate at which beliefs concentrate (see also Fudenberg, Lanzani, and Strack, 2021b). The approaches in EPY and FLS are again based on Kullback-Leibler divergence; as noted, this measure can be too coarse to identify long-run outcomes in settings such as slow learning environments (Remark 2).

Some environments in the literature are not nested by the current framework, notably models with intertemporally correlated signals and social learning settings with private action observations.⁴ The latter includes our previous paper, Frick, Iijima, and Ishii (2020a), which, similar to Section 4.2.2, highlights the fragility of social learning against misspecification about others’ preferences. As we discuss (Section 4.2.3), the logic and nature of this fragility result differs from the current paper, as the setting in Frick, Iijima, and Ishii (2020a) does not display slow learning.

2 Model

2.1 Setup

We conduct our general analysis in the following abstract environment, where agents, actions, and preferences are not explicitly modeled. This allows us to simultaneously nest a variety of single-agent and social learning models and makes the logic of belief convergence transparent. For any topological space X , we endow X with its Borel σ -algebra and let $\Delta(X)$ denote the space of Borel probability measures on X .

There is a set of *states* Ω . For the analysis in the main text, we assume that Ω is finite; Appendix B provides results for infinite state spaces. At the beginning of each period $t = 0, 1, \dots$, there is a *belief* $\mu_t \in \Delta(\Omega)$; we endow $\Delta(\Omega) \subseteq \mathbb{R}^{|\Omega|}$ with the sup norm. The initial belief μ_0 is exogenous and has full support.⁵ The evolution of beliefs is determined as follows: At the end of each period t , a signal z_t from a topological space Z is drawn according to P_{μ_t} , where $P_{\mu} \in \Delta(Z)$ denotes the *true signal distribution* at current belief μ . After signal z_t realizes, belief μ_t is updated to μ_{t+1} via Bayes’ rule according to a collection of conditional *perceived signal distributions*: Specifically, at each current belief μ , the perceived signal distribution conditional on state ω is $\hat{P}_{\mu}(\cdot|\omega) \in \Delta(Z)$. We assume that, for

⁴See, e.g., Rabin (2002); Ortoleva and Snowberg (2015); Esponda and Pouzo (2019); Molavi (2019); Cho and Kasa (2017) for the former, and Dasaratha and He (2020); Levy and Razin (2018) for the latter.

⁵The full-support assumption is without loss; if μ_0 assigns zero probability to some states, the same analysis and results below apply up to eliminating those states from Ω .

each ω and μ , P_μ and $\hat{P}_\mu(\cdot|\omega)$ admit continuous Radon-Nikodym derivatives p_μ and $\hat{p}_\mu(\cdot|\omega)$ with respect to some σ -finite measure ν on Z ; as usual, when Z is finite (resp. $Z = \mathbb{R}$), we take ν to be the counting (resp. Lebesgue) measure. The updated belief following signal z_t satisfies

$$\mu_{t+1}(\omega) = \frac{\mu_t(\omega)\hat{p}_{\mu_t}(z_t|\omega)}{\sum_{\omega' \in \Omega} \mu_t(\omega')\hat{p}_{\mu_t}(z_t|\omega')}, \quad \forall \omega \in \Omega.$$

By allowing the true and perceived signal distributions to depend on the current belief, the model can nest applications where signals depend endogenously on agents' actions, which depend on their current beliefs; see Remark 1. Capturing possible misspecification, the true signal distribution need not coincide with any of the perceived signal distributions. We refer to the case where for some true state ω^* , $P_\mu = \hat{P}_\mu(\cdot|\omega^*)$ for all μ , as the *correctly specified* benchmark. Throughout, we impose the following regularity assumption:

Assumption 1.

1. (Absolute continuity). For each $\omega \in \Omega$ and $\mu \in \Delta(\Omega)$, $\text{supp}P_\mu \subseteq \text{supp}\hat{P}_\mu(\cdot|\omega)$.
2. (Well-behaved likelihood ratios). There exist a ν -integrable function $h : Z \rightarrow \mathbb{R}_+$ and $q^* > 0$ such that $\sup_{\mu, \omega, \omega'} \left(\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')} \right)^{q^*} p_\mu(z) \leq h(z)$ for all $z \in Z$.⁶
3. (Belief continuity near point-mass beliefs). For each $\omega \in \Omega$, there is a neighborhood $B \ni \delta_\omega$ such that for all $\omega', \omega'' \in \Omega$, $\mu \in B$ and $z \in Z$, we have that $p_\mu(z)$, $\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')}$, and $p_\mu(z) \frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega'')}$ are continuous in μ .

Assumption 1.1 is standard in the literature and rules out belief-updating after signals that are perceived to realize with zero probability. The remaining assumptions are technical conditions that are satisfied in most applications in the literature: Assumption 1.2 is a regularity condition on the integrability of perceived likelihood ratios, which will be important for our martingale approach based on moment generating functions in Section 3.1. This rules out that the distribution of perceived log-likelihood ratios $\log \frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')}$, when z is drawn from P_μ , is heavy-tailed (i.e., the moment-generating function is infinite at all non-zero arguments); commonly used parametric distributions (e.g., Gaussian) are not heavy-tailed. Assumption 1.3 imposes continuity with respect to beliefs on signal densities, but is only assumed near point-mass beliefs; this simplifies the statements of our stability results.

Remark 1. We illustrate how two leading classes of applications map into this model.

Single-agent learning: The state space Ω represents an agent's uncertainty about the environment (e.g., a monopolist's uncertainty about market conditions). Each period

⁶Throughout the paper, we use the convention that $\frac{0}{0} = 0$, $\frac{1}{0} = \infty$, $0 \log 0 = 0$, and $\log \infty = \infty$.

$t = 0, 1, \dots$, the agent chooses an action a_t (e.g., a price) from a discrete or continuous space A and observes a signal $z_t \in Z$ (e.g., realized demand). Each action a induces a true signal distribution $G_a \in \Delta(Z)$, but the agent updates her belief $\mu_t \in \Delta(\Omega)$ based on the perceived signal distributions $\hat{G}_a(\cdot|\omega) \in \Delta(Z)$ (e.g., the monopolist may hold a misspecified model of the demand function). The agent's action choice $a_t = a(\mu_t)$ is Markovian in her belief, for example, because she maximizes subjective expected discounted payoffs (e.g., revenue).

Active learning environments of this form map into our model by setting $P_\mu = G_{a(\mu)}$ and $\hat{P}_\mu(\cdot|\omega) = \hat{G}_{a(\mu)}(\cdot|\omega)$. The above assumptions on P, \hat{P} translate into assumptions on G, \hat{G} , and $a(\cdot)$ in a direct manner. For example, Assumption 1.3 holds if $a(\cdot)$ is continuous in μ near point-mass beliefs and $G_a, \hat{G}_a(\cdot|\omega)$ admit densities that satisfy the corresponding continuity conditions with respect to a .⁷

In addition to monopoly pricing, Section 4 will analyze several other concrete active learning problems, including costly information acquisition and effort choice. Beyond active learning, our model can also capture single-agent learning settings where true signal distributions are exogenous, but perceived signal distributions depend on μ due to certain belief-dependent departures from Bayesian updating, such as confirmation bias.

Social learning: Consider a sequential social learning setting à la [Smith and Sørensen \(2000\)](#). There is a fixed and unknown state $\omega^* \in \Omega$ (e.g., the safety of a new product). Each period $t = 0, 1, \dots$, agent t chooses a one-shot action $z_t \in Z = \{0, 1\}$ (e.g., whether or not to adopt the product) after observing the public sequence (z_0, \dots, z_{t-1}) of predecessors' actions and a private signal $s_t \in \mathbb{R}$ that is drawn i.i.d. conditional on state ω^* according to a cdf $\Phi(\cdot|\omega^*)$. Agents have private preference types $\theta_t \in \mathbb{R}$ (e.g., risk attitudes), which are drawn independently across agents, states, and signals according to a cdf F . Starting with some full-support common prior $\mu_0 \in \Delta(\Omega)$, agent t chooses z_t to maximize her expected utility,

$$z_t \in z(\mu_t, \theta_t, s_t) := \underset{z \in Z}{\operatorname{argmax}} \mathbb{E}_{\mu_t}[u(z, \theta_t, \omega)|\theta_t, s_t],$$

where μ_t denotes the *public belief*, i.e., the Bayesian update of μ_0 based solely on the public action sequence (z_0, \dots, z_{t-1}) . In Section 4.2.2, we will analyze this setting when agents are misspecified about others' preferences: In updating beliefs to μ_t , all agents misperceive the type distribution F in the population to be some other cdf \hat{F} .⁸

To map this into our model, consider the public belief process μ_t and identify signals

⁷Note that if A is discrete, then $a(\cdot)$ is not in general globally continuous (unless there is a dominant action), but Assumption 1.3 is satisfied as long as, for each ω , there is a neighborhood $B_\omega \ni \delta_\omega$ such that $a(\mu)$ is constant across all beliefs $\mu \in B_\omega$. The formulation also allows A to be a set of mixed actions; in this case, we treat Z as the product space of realized signals and actions.

⁸Alternatively or additionally, agents might misperceive the private signal distributions $\Phi(\cdot|\omega)$.

with actions z_t . Given μ_t , z_t is stochastic due to the random realization of agent t 's type θ_t and private signal s_t . The true probability of each action z given public belief μ is⁹

$$p_\mu(z) = p_\mu(z|\omega^*) = \Pr_{\theta \sim F, s \sim \Phi(\cdot|\omega^*)} [z \in z(\mu, \theta, s)];$$

however, because agents misperceive the type distribution F to be \hat{F} , public beliefs are updated according to the perceived action probabilities

$$\hat{p}_\mu(z|\omega) = \Pr_{\theta \sim \hat{F}, s \sim \Phi(\cdot|\omega)} [z \in z(\mu, \theta, s)].$$

Beyond this particular setting, our model also nests any other social learning environment in which agents' actions are Markovian in a public belief, including learning from market prices (e.g., [Vives, 1993](#)) or strategic experimentation (e.g., [Bolton and Harris, 1999](#)). \blacktriangle

2.2 Stability notions

Given any true and perceived signal distributions and initial belief μ_0 , our model generates a Markov process over beliefs. Let \mathbb{P}_μ denote the induced probability measure over sequences of beliefs (μ_t) with $\mu_0 = \mu$. We seek to analyze which states ω long-run beliefs can grow confident in, in the sense that process μ_t converges to the point-mass belief δ_ω either locally or globally as a function of initial beliefs. Formally, we consider the following stability notions:¹⁰

Definition 1. Consider any $\omega \in \Omega$. Belief δ_ω is:

1. **locally stable** if for any $\gamma < 1$, there exists a neighborhood $B \ni \delta_\omega$ such that $\mathbb{P}_\mu[\mu_t \rightarrow \delta_\omega] \geq \gamma$ for each initial belief $\mu \in B$;
2. **globally stable** if $\mathbb{P}_\mu[\mu_t \rightarrow \delta_\omega] = 1$ for each initial belief μ ;
3. **unstable** if there exists a neighborhood $B \ni \delta_\omega$ such that $\mathbb{P}_\mu[\exists t, \mu_t \notin B] = 1$ for each initial belief $\mu \in B$.

Local stability requires that beliefs converge with positive probability to δ_ω from any initial belief in some open set B around δ_ω , where the probability of converging to δ_ω can be made arbitrarily close to 1 as long as B is small enough.¹¹ More strongly, global stability

⁹We assume that the true and perceived probability that the set of interim-optimal actions $z(\mu, \theta, s)$ is single-valued is 1. The additional restrictions imposed in Section 4.2.2 will ensure that this is the case.

¹⁰Similar stability notions are considered by [Smith and Sørensen \(2000\)](#); [Bohren and Hauser \(2021\)](#).

¹¹We do not consider a stronger version of local stability that allows for $\gamma = 1$. Unless global stability holds, this notion is too demanding in most settings (due to the possibility of signal realizations that push beliefs outside neighborhood B).

requires that beliefs converge to δ_ω with probability 1 from *any* initial belief (recall that initial beliefs are assumed full-support). By contrast, δ_ω is unstable if starting from any initial belief μ in some small enough neighborhood B of δ_ω , beliefs eventually escape B with probability 1. Clearly, if δ_ω is unstable, it is not locally stable.¹²

By focusing on the stability/instability of point-mass beliefs δ_ω , this paper does not analyze when long-run beliefs are mixed, i.e., assign positive probability to multiple states. Long-run beliefs are never mixed in environments that satisfy an identification condition, whereby at any mixed μ , there is a possible signal realization that leads beliefs to update in favor of one state in the support of μ rather than some other state (see Lemma 10 in Appendix A for the formal statement). This condition is satisfied in most existing settings studied in the misspecified learning literature, including all applications in this paper. At the same time, this rules out some important applications, such as active learning settings where agents stop observing informative signals at some mixed belief (e.g., McLennan, 1984, bandit problems) and social learning settings that feature herding or confounded learning (Bikhchandani, Hirshleifer, and Welch, 1992; Banerjee, 1992; Smith and Sørensen, 2000). Section 5 briefly discusses how our techniques might be extended to such settings, which have thus far been studied mostly without misspecification.

2.3 Berk-Nash equilibrium and slow learning

A necessary condition for stability has been proposed by Esponda and Pouzo (2016). For any $P, \hat{P} \in \Delta(Z)$ with densities p, \hat{p} , define the **Kullback-Leibler (KL) divergence** of \hat{P} relative to P by $\text{KL}(P, \hat{P}) := \int \log \frac{p(z)}{\hat{p}(z)} dP(z)$. When signals are drawn repeatedly according to the distribution P , this measures how close \hat{P} comes to predicting the long-run signal distribution, by considering the expected log-likelihood ratio of signals between P and \hat{P} . Given any true and perceived signal distributions, we call belief δ_ω a **Berk-Nash equilibrium (BeNE)** if

$$\omega \in \underset{\omega' \in \Omega}{\text{argmin}} \text{KL} \left(P_{\delta_\omega}, \hat{P}_{\delta_\omega}(\cdot | \omega') \right). \quad (1)$$

Condition (1) is a fixed-point requirement, which says that at belief δ_ω , the perceived signal distribution that comes closest to the true signal distribution P_{δ_ω} is the distribution $\hat{P}_{\delta_\omega}(\cdot | \omega)$ in state ω . Thus, if beliefs converge to δ_ω , then state ω itself best predicts the induced long-run signal distribution. This is a straightforward adaptation of Esponda and Pouzo (2016) to our setting, focusing only on point-mass beliefs.¹³ Analogous to Esponda and Pouzo (2016),

¹²Note that it is possible that δ_ω is neither unstable nor locally stable, for example, if, for every neighborhood $B \ni \delta_\omega$, whether or not beliefs converge to δ_ω varies across initial beliefs $\mu_0 \in B$.

¹³Esponda and Pouzo (2016) consider settings where multiple agents choose actions given their beliefs about a payoff-relevant parameter and about other agents' behavior. A BeNE requires agents' beliefs to

we show that this is a necessary condition for δ_ω to be locally stable:

Lemma 1. *If δ_ω is not a BeNE, then δ_ω is unstable.*

While condition (1) is necessary for local stability, it is not in general sufficient, as many environments feature multiple BeNE, some of which are stable while others are unstable. Thus, our sufficient conditions for stability will take the form of refinements of BeNE.

A class of environments with a particularly stark multiplicity of BeNE is the following. We say that **slow learning** obtains if, for any $\omega, \omega', \omega'' \in \Omega$ and ν -almost all z ,

$$\lim_{\mu \rightarrow \delta_\omega} \hat{p}_\mu(z|\omega') = \lim_{\mu \rightarrow \delta_\omega} \hat{p}_\mu(z|\omega''). \quad (2)$$

That is, the (perceived) information content of each signal z vanishes as the belief μ grows confident in any particular state ω . Under (2), the expected change in log-posterior ratios, $\mathbb{E}_{P_{\mu_t}} [\log \frac{\mu_{t+1}(\omega')}{\mu_{t+1}(\omega'')} - \log \frac{\mu_t(\omega')}{\mu_t(\omega'')}] = \int \log \frac{\hat{p}_{\mu_t}(z|\omega')}{\hat{p}_{\mu_t}(z|\omega'')} dP_{\mu_t}(z)$, vanishes as beliefs μ_t approach any point-mass belief δ_ω , capturing the sense in which learning is slow. Under Assumption 1, slow learning implies that $\hat{p}_{\delta_\omega}(z|\omega')$ is constant in ω' at each δ_ω . From this it is immediate that *every* point-mass belief δ_ω is a BeNE.

As a large literature highlights (for surveys, see Vives, 2010; Chamley, 2004), slow learning is a central feature of many social learning models (e.g., the sequential social learning environment in Remark 1): In these settings, new action observations convey less and less information as the public belief grows confident, because agents base their action choices increasingly on the public belief rather than their private information.¹⁴ As we illustrate in Section 4.2.1, slow learning also arises naturally in single-agent settings if information acquisition is costly, in which case the agent acquires less and less informative signals as she grows confident in any state. By contrast, if every action chosen by the agent generates non-vanishingly informative signals about the state (as in the applications in Section 4.1), then learning is not slow.

minimize KL-divergence relative to the feedback about the parameter and others' behavior that is generated by their optimal actions at these beliefs. They also allow for mixed BeNE and show that if beliefs converge to μ^* with positive probability, then μ^* must be a BeNE belief (see their Lemma 2 and Theorem 2).

¹⁴Herding is an extreme form of slow learning, where belief-updating ceases completely at some mixed belief. But even absent herding, sequential social learning is generally slow, as quantified by Vives (1993); Hann-Caruthers, Martynov, and Tamuz (2018); Rosenberg and Vieille (2019). These papers employ different quantifications of learning speed, but all the settings studied satisfy (2).

3 Stability analysis

3.1 Prediction accuracy orders and martingale approach

Before presenting our conditions for local stability, instability, and global stability of beliefs, we introduce orders over states that compare how well they predict the true signal distribution at each belief μ . These prediction accuracy orders will play a central role in our stability analysis and the martingale arguments on which it relies.

Given any belief μ , say that state ω **KL-dominates** ω' **at** μ , denoted $\omega \succ_{\mu}^{\text{KL}} \omega'$, if

$$\text{KL}\left(P_{\mu}, \hat{P}_{\mu}(\cdot|\omega)\right) - \text{KL}\left(P_{\mu}, \hat{P}_{\mu}(\cdot|\omega')\right) := \int \log\left(\frac{\hat{p}_{\mu}(z|\omega')}{\hat{p}_{\mu}(z|\omega)}\right) dP_{\mu}(z) \leq 0. \quad (3)$$

That is, at belief μ , the perceived signal distribution in state ω achieves lower KL-divergence relative to the true distribution than does the perceived signal distribution in state ω' . Write $\omega \succ_{\mu}^{\text{KL}} \omega'$ if inequality (3) is strict. Note that δ_{ω} is a BeNE if and only if $\omega \succ_{\delta_{\omega}}^{\text{KL}} \omega'$ for all ω' .

Our analysis relies on the following refinement of \succ_{μ}^{KL} . Given any $q > 0$, say that ω **q -dominates** ω' **at** μ , denoted $\omega \succ_{\mu}^q \omega'$, if

$$\int \left(\frac{\hat{p}_{\mu}(z|\omega')}{\hat{p}_{\mu}(z|\omega)}\right)^q dP_{\mu}(z) \leq 1, \quad (4)$$

and write $\omega \succ_{\mu}^q \omega'$ if inequality (4) is strict. To see the connection between q -dominance and KL-dominance, consider the random variable $X = \log\left(\frac{\hat{p}_{\mu}(z|\omega')}{\hat{p}_{\mu}(z|\omega)}\right)$, i.e., the perceived log-likelihood ratio of states ω' vs. ω , when signals z are drawn according to the true signal distribution P_{μ} . Then the left-hand side of (3) is the expectation of X , while the left-hand side of (4) is the moment-generating function $M_X(q) = \mathbb{E}[e^{qX}]$ of X evaluated at q .

Whereas \succ_{μ}^{KL} is complete (by the representation on the LHS of (3)), \succ_{μ}^q is in general incomplete.¹⁵ However, the q -dominance orders are nested and approximate KL-dominance as $q \rightarrow 0$:

Lemma 2. *Fix any belief μ and states ω, ω' .*

1. *If $\omega \succ_{\mu}^q \omega'$ for some $q > 0$, then $\omega \succ_{\mu}^{\text{KL}} \omega'$ and $\omega \succ_{\mu}^{q'} \omega'$ for all $q' \in (0, q)$.*
2. *If $\omega \succ_{\mu}^{\text{KL}} \omega'$, then there exists $q > 0$ such that $\omega \succ_{\mu}^q \omega'$.*

¹⁵Note that q -dominance bears some formal resemblance to a generalization of KL-divergence known as Rényi divergence. However, whereas KL-dominance amounts to comparing $\text{KL}(P_{\mu}, \hat{P}_{\mu}(\cdot|\omega))$ and $\text{KL}(P_{\mu}, \hat{P}_{\mu}(\cdot|\omega'))$, q -dominance is not equivalent to comparing the corresponding Rényi divergences.

To understand the role that q -dominance will play in our analysis, first consider the correctly specified benchmark, where for some true state ω^* , $P_\mu = \hat{P}_\mu(\cdot|\omega^*)$ for all μ . In this case, $\omega^* \succsim_\mu^1 \omega$ for all μ and ω ; indeed, (4) holds with equality when $q = 1$.¹⁶ This implies a well-known property of correctly specified learning: The posterior ratio process $\frac{\mu_t(\omega)}{\mu_t(\omega^*)}$ is a nonnegative martingale with respect to \mathbb{P}_{μ_0} and the filtration generated by (μ_t) , as

$$\mathbb{E}_{\mathbb{P}_{\mu_0}} \left[\frac{\mu_{t+1}(\omega)}{\mu_{t+1}(\omega^*)} | (\mu_s)_{s \leq t} \right] = \frac{\mu_t(\omega)}{\mu_t(\omega^*)} \int \left(\frac{\hat{p}_{\mu_t}(z|\omega)}{\hat{p}_{\mu_t}(z|\omega^*)} \right) dP_{\mu_t}(z) = \frac{\mu_t(\omega)}{\mu_t(\omega^*)}.$$

The martingale property is central to analyzing long-run beliefs under correctly specified learning. In particular, it implies that, by Doob's convergence theorem, $\frac{\mu_t(\omega)}{\mu_t(\omega^*)}$ converges almost surely (a.s.) to a nonnegative random limit.

Under misspecified learning, there is in general no state that globally 1-dominates all other states. As a result, the martingale property is lost. However, the definition of \succsim_μ^q immediately implies a key observation: Throughout any region of beliefs where q -dominance obtains, the q th power of the posterior ratio process becomes a nonnegative supermartingale.

Lemma 3. *Suppose there exist $q > 0$ and $B \subseteq \Delta(\Omega)$ such that $\omega \succsim_\mu^q \omega'$ for all $\mu \in B$. Then, for any initial belief μ_0 , the process ℓ_t defined by*

$$\ell_t := \left(\frac{\mu_{\min\{t, \tau\}}(\omega')}{\mu_{\min\{t, \tau\}}(\omega)} \right)^q \quad \text{with } \tau := \inf\{s : \mu_s \notin B\} \quad (5)$$

is a nonnegative supermartingale with respect to \mathbb{P}_{μ_0} and the filtration generated by μ_t .

Proof. Observe $\mathbb{E}_{\mathbb{P}_{\mu_0}}[\ell_{t+1} | (\mu_s)_{s \leq t}] = \begin{cases} \ell_t \int \left(\frac{\hat{p}_{\mu_t}(z|\omega')}{\hat{p}_{\mu_t}(z|\omega)} \right)^q dP_{\mu_t}(z) \leq \ell_t & \text{if } \mu_s \in B \forall s \leq t \\ \ell_t & \text{otherwise.} \end{cases} \quad \square$

Under the assumptions in Lemma 3, standard martingale methods from the correctly specified setting, such as Doob's convergence theorem and Markov's inequality, can be applied locally, to the stopped process ℓ_t . Such arguments will play a key role throughout our stability analysis, by providing useful information on the asymptotic behavior of the original belief process μ_t . As we discuss in Remark 2, q -dominance is essential to this approach, as analogous arguments do not apply under KL-dominance.

3.2 Local stability and instability

Based on the preceding observations, our first main result provides sufficient conditions for belief δ_ω to be locally stable or unstable:

¹⁶That is, $\int \left(\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega^*)} \right) dP_\mu(z) = \int \left(\frac{\hat{p}_\mu(z|\omega)}{p_\mu(z)} \right) p_\mu(z) d\nu(z) = 1$.

Theorem 1. Consider any $\omega \in \Omega$. Belief δ_ω is:

1. locally stable if there exists $q > 0$ and a neighborhood $B \ni \delta_\omega$ such that

$$\omega \succ_\mu^q \omega' \text{ for all } \omega' \neq \omega \text{ and } \mu \in B \setminus \{\delta_\omega\}. \quad (6)$$

2. unstable if there exists $q > 0$ and a neighborhood $B \ni \delta_\omega$ such that

$$\text{for some } \omega' \neq \omega, \text{ we have } \omega' \succ_\mu^q \omega \text{ for all } \mu \in B \setminus \{\delta_\omega\}. \quad (7)$$

By the first part, δ_ω is locally stable if for some q , state ω strictly q -dominates all other states at all beliefs in some neighborhood of δ_ω , except possibly at the belief δ_ω , where this dominance need only be weak.¹⁷ Thus, condition (6) strengthens BeNE, which requires that ω weakly KL-dominates all other states at the belief δ_ω , in two ways: First, by comparing the prediction accuracy of ω against other states at beliefs in a neighborhood B of δ_ω ; second, by imposing strict q -dominance rather than weak KL-dominance throughout $B \setminus \{\delta_\omega\}$. The second part provides an analogous condition for instability; combined with Lemma 2, this result also implies Lemma 1.

The proof of Theorem 1 is a simple application of the martingale construction in the previous section. To see the idea, suppose that $\Omega = \{\omega, \omega'\}$ is binary. For the first part, consider the stopped process $\ell_t(\omega') := \left(\frac{\mu_{\min\{t,\tau\}}(\omega')}{\mu_{\min\{t,\tau\}}(\omega)}\right)^q$ with $\tau := \inf\{s : \mu_s \notin B\}$. By Lemma 3, this is a nonnegative supermartingale. Thus, by Doob's convergence theorem, ℓ_t converges a.s. to a nonnegative random limit ℓ_∞ . Based on this, we first show that if the belief process μ_t remains in B forever with positive probability, then conditional on this event, μ_t converges to δ_ω a.s.: Otherwise, the random limit belief $\mu_\infty \in B$ would be mixed with positive probability, which we show is impossible by (6). Second, by applying Markov's inequality to ℓ_∞ , we show that the probability that μ_t remains in B forever can be made arbitrarily close to 1 by restricting to initial beliefs μ_0 in a small enough subneighborhood $B' \subseteq B$ around δ_ω . Combining these observations implies that δ_ω is locally stable. For the second part of Theorem 1, we apply Doob's theorem to the nonnegative supermartingale $\ell_t(\omega') := \left(\frac{\mu_{\min\{t,\tau\}}(\omega)}{\mu_{\min\{t,\tau\}}(\omega')}\right)^q$ with $\tau := \inf\{s : \mu_s \notin B\}$, to show that μ_t a.s. leaves B .

The fact that conditions (6) and (7) do not impose strict dominance at the point-mass belief δ_ω is essential for applying Theorem 1 to environments with slow learning: Indeed, under (2), the difference in prediction accuracy across states vanishes as μ approaches any point-mass belief.¹⁸

¹⁷The weak dominance $\omega \succ_{\delta_\omega}^q \omega'$ follows from (6) and Assumption 1.

¹⁸That is, (3) and (4) hold with equality when $\mu = \delta_\omega$.

Note that conditions (6) and (7) feature existential quantifiers over q and B . The following example illustrates how q and B can be found straightforwardly from the relationship between P and \hat{P} ; we will apply similar observations to analyze the economic applications in Section 4.2.

Example 1. Consider $Z = \{0, 1\}$ and any δ_ω . Under slow learning, perceived signal probabilities $\hat{p}_\mu(1|\omega')$ become independent of ω' as μ approaches δ_ω . Suppose these perceptions understate the truth in any small enough neighborhood B of δ_ω , i.e., $\hat{p}_\mu(1|\omega') \leq p_\mu(1)$ for all ω' and $\mu \in B$ (the opposite case is analogous). Consider two possibilities near δ_ω :

- *Perceived signal probabilities in state ω are closest to the truth:* That is, $\hat{p}_\mu(1|\omega') < \hat{p}_\mu(1|\omega)$ for all $\omega' \neq \omega$ and $\mu \in B \setminus \{\delta_\omega\}$. Then, $\omega \succ_\mu^q \omega'$ for any $q \in (0, 1)$.¹⁹ Thus, δ_ω is locally stable by (6).
- *Perceived signal probabilities in some other state ω' are closer to the truth:* That is, for some $\omega' \neq \omega$, $\hat{p}_\mu(1|\omega') > \hat{p}_\mu(1|\omega)$ for all $\mu \in B \setminus \{\delta_\omega\}$. Then, analogously, $\omega' \succ_\mu^q \omega$ for all $q \in (0, 1)$. Thus, δ_ω is unstable by (7). \blacktriangle

At the same time, an immediate corollary of Theorem 1 is the following more demanding sufficient condition for local stability, which is easy to verify in environments that do not feature slow learning (or other ties in prediction accuracy). Call δ_ω a **strict BeNE** if $\omega \succ_{\delta_\omega}^{\text{KL}} \omega'$ for all $\omega' \neq \omega$. By Lemma 2 and Assumption 1, any strict BeNE satisfies (6).

Corollary 1. *If δ_ω is a strict BeNE, then δ_ω is locally stable.*

Bohren (2016) (extended by Bohren and Hauser (2021) to heterogeneous beliefs) derived an analog of Corollary 1 under binary states $|\Omega| = 2$ and finite Z . Their proofs use a “local approximation” argument that is different from our martingale approach and does not extend to settings that feature slow learning.²⁰

While Corollary 1 is not applicable under slow learning, a convenient feature is that it only involves considering KL-prediction accuracy differences at the single belief δ_ω . Under slow learning, Theorem 1 can be used to derive a condition for local stability with a similar feature: This condition only involves computing the *derivative* of the KL-prediction accuracy differences at the belief δ_ω ; see Online Appendix D.1.

¹⁹Indeed, $\hat{p}_\mu(1|\omega') < \hat{p}_\mu(1|\omega) \leq p_\mu(1)$ implies $\sum_z p_\mu(z) \left(\frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)} \right)^q \leq \sum_z \hat{p}_\mu(z|\omega) \left(\frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)} \right)^q < 1$ for any $q \in (0, 1)$, where the final inequality follows from Jensen’s inequality and the concavity of $f(x) = x^q$.

²⁰Specifically, they locally bound the log-likelihood ratio process under (P, \hat{P}) by the corresponding process under a different environment (Q, \hat{Q}) with the property that Q_μ, \hat{Q}_μ are independent of μ and that beliefs converge to δ_ω a.s. (by the law of large numbers). The construction of (Q, \hat{Q}) requires the log-likelihood ratio process under (P, \hat{P}) to have non-vanishing drift near δ_ω , which implies that $\omega \succ_{\delta_\omega}^{\text{KL}} \omega'$ for $\omega' \neq \omega$.

Remark 2. To understand the importance of refining KL-dominance to q -dominance, suppose (6) is weakened to the assumption that in some neighborhood $B \ni \delta_\omega$,

$$\omega \succ_\mu^{\text{KL}} \omega' \text{ for all } \omega' \neq \omega \text{ and } \mu \in B \setminus \{\delta_\omega\}. \quad (8)$$

Then the stopped processes $\log \left(\frac{\mu_{\min\{t,\tau\}}(\omega')}{\mu_{\min\{t,\tau\}}(\omega)} \right)$ with $\tau := \inf\{s : \mu_s \notin B\}$ are supermartingales. However, since these supermartingales are *unbounded* below as μ_t approaches δ_ω , the above arguments based on Doob's convergence theorem and Markov's inequality no longer apply. Indeed, Online Appendix D.2 provides an example where (8) holds but δ_ω is unstable. This illustrates that KL-dominance conditions are not in general enough to determine local stability/instability. \blacktriangle

Finally, we note that when Ω is infinite, strict BeNE need not be locally stable, as shown by Proposition 1 in Heidhues, Kőszegi, and Strack (2021). Appendix B.2 provides conditions for local stability under infinite states.

3.3 Global stability

Global stability is significantly more demanding than local stability. For instance, even if δ_ω is the unique locally stable belief, it need not be globally stable. In this section, we use our martingale approach to obtain two sufficient conditions for global stability that strengthen the local stability criterion in Theorem 1 in complementary ways. Both conditions place some additional restrictions on the environment, but we illustrate their usefulness with the applications in Section 4.

3.3.1 Iterated elimination of dominated states

Our first approach extends the previous local stability arguments by constructing supermartingales that apply not only near δ_ω but more globally.

We employ a generalization of global stability to sets of beliefs: Call $M \subseteq \Delta(\Omega)$ a **globally stable set** if $\mathbb{P}_\mu[\inf_{\nu \in M} \|\mu_t - \nu\| \rightarrow 0] = 1$ for every initial belief μ . Note that $\Delta(\Omega)$ is trivially globally stable. We show that global stability is preserved under the following process of **iterated elimination of dominated states**, defined similarly to the iterated elimination of dominated strategies in games: For each subset $\Omega' \subseteq \Omega$, let

$$S(\Omega') := \{\omega \in \Omega' : \nexists \omega' \in \Omega' \text{ s.t. } \omega' \succ_\mu^{\text{KL}} \omega \text{ for all } \mu \in \Delta(\Omega')\}.$$

Then recursively define $S^0(\Omega) := \Omega$, $S^{k+1}(\Omega) := S(S^k(\Omega))$ for all $k = 0, 1, \dots$, and $S^\infty(\Omega) :=$

$\bigcap_{k \in \mathbb{N}} S^k(\Omega)$. We say that **belief continuity** holds if Assumption 1.3 is satisfied not only near point-mass beliefs, but at all beliefs $\mu \in \Delta(\Omega)$.²¹

Theorem 2. *Assume belief continuity holds. Then $\Delta(S^\infty(\Omega))$ is globally stable. In particular, if $S^\infty(\Omega) = \{\omega\}$ for some $\omega \in \Omega$, then belief δ_ω is globally stable.*

To prove Theorem 2, we show inductively that $\Delta(S^k(\Omega))$ is globally stable for all k . Since $\Delta(\Omega)$ is globally stable, it suffices to show that whenever $\Delta(\Omega')$ is globally stable for some $\Omega' \subseteq \Omega$, then so is $\Delta(S(\Omega'))$. This can again be established using martingale arguments. To see the idea, suppose that $S(\Omega') = \Omega' \setminus \{\omega'\}$. Then by Lemma 2 and belief continuity, there exist $q > 0$ and $\omega'' \in \Omega'$ such that $\omega'' \succ_\mu^q \omega'$ for all $\mu \in \Delta(\Omega')$, and hence also $\omega'' \succ_\mu^q \omega'$ for all μ in any small enough neighborhood $B \supseteq \Delta(\Omega')$.²² Thus, by Lemma 3,

$$\left(\frac{\mu_{\min\{t,\tau\}}(\omega')}{\mu_{\min\{t,\tau\}}(\omega'')} \right)^q \quad \text{with } \tau = \inf\{s : \mu_s \notin B\} \quad (9)$$

is a nonnegative supermartingale. Similar to Theorem 1, this implies that (i) from any initial $\mu \in B$, μ_t remains forever in B with positive probability; and (ii) $\mu_t(\omega')$ converges to 0 a.s. conditional on remaining in B . We show that combined with the assumption that $\Delta(\Omega')$ (and hence $B \supseteq \Delta(\Omega')$) is globally stable, this yields that $\Delta(\Omega' \setminus \{\omega'\})$ is globally stable.

Note that although the definition of iterated elimination uses strict KL-dominance, q -dominance again plays an essential role in the proof, by allowing us to construct the nonnegative supermartingale (9). Appendix B.1 shows that Theorem 2 remains true unchanged in arbitrary compact metric state spaces, by extending the above martingale arguments.

At a high level, the iterated elimination approach can be seen as an abstract generalization of arguments in Heidhues, Kőszegi, and Strack (2018) and He (2021), who analyze specific single-agent settings with one-dimensional states and actions.²³ Their analysis considers the largest interval of states that is contained in the support of the agent's long-run belief and shows, using iterated contraction arguments, that this must collapse to a singleton. While their proofs involve analyzing the slope of the agent's perceived log-likelihood functions with respect to the one-dimensional state, our proof is based on constructing the nonnegative supermartingales (9), which does not require any order structure over states.

²¹That is, for each $\omega, \omega' \in \Omega$, $\mu \in \Delta(\Omega)$ and $z \in Z$, we have that $p_\mu(z)$, $\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')}$, and $p_\mu(z) \frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')}$ are continuous in μ . Belief continuity can be dropped in Theorem 2 and Proposition 1, up to slightly strengthening the corresponding dominance requirements; see also footnote 33.

²²Call B a **neighborhood of a set** $M \subseteq \Delta(\Omega)$ if there exists $\varepsilon > 0$ such that $B_\varepsilon(\mu) \subseteq B$ for all $\mu \in M$.

²³He (2021) allows for two-dimensional states, but proves that the analysis can be reduced to the one-dimensional case in the long-run. As noted, Theorem 2 extends to any compact metric space of states.

3.3.2 Global stability via uniform local dominance

Theorem 2 requires that eliminated states are dominated at *all* beliefs in a subsimplex $\Delta(S^k(\Omega))$, which is restrictive in some applications. In such settings, an alternative approach to obtain global stability is to restrict the prediction accuracy order only locally, near point-mass beliefs, but to impose more structure on how states are ranked. The following result provides one formalization of this approach that is useful for our applications in Section 4.2:

Proposition 1. *Suppose that belief continuity holds and states $\Omega = \{\omega_1, \dots, \omega_N\}$ can be enumerated in such a way that*

- (i) *for each ω , there exists $q > 0$ and a neighborhood $B \ni \delta_\omega$ such that for all m and n with $m > n$, we have $\omega_n \succ_\mu^q \omega_m$ for all $\mu \in B \setminus \{\delta_\omega\}$;*
- (ii) *for all $n \neq N$ and mixed μ , there is $z \in \text{supp}P_\mu$ with $\hat{p}_\mu(z|\omega_n) > \hat{p}_\mu(z|\omega_m)$ for all $m > n$.*

Then δ_{ω_1} is globally stable.

Condition (i) requires that, near all point-mass beliefs δ_ω , the prediction accuracy ranking is the same: states with a lower index dominate higher states.²⁴ For binary Ω , (i) amounts to imposing the local stability condition (6) from Theorem 1 on δ_{ω_1} and the instability condition (7) on δ_{ω_2} . However, when $|\Omega| > 2$, (i) is more demanding than imposing local stability on δ_{ω_1} and instability on all other δ_{ω_n} ; we explain the role of this added strength below. Condition (ii) is relatively weak, in that it does not restrict the prediction accuracy ranking. One natural condition that implies (ii) is if perceived signal distributions satisfy the monotone likelihood ratio property, as is the case in many applications.

When Ω is binary, the logic behind Proposition 1 is analogous to [Bohren \(2016\)](#), who derived a similar result (under a strengthening of condition (i) that requires strict KL-dominance at point-mass beliefs, ruling out slow learning). By condition (i), there are neighborhoods $B_1 \ni \delta_{\omega_1}$ and $B_2 \ni \delta_{\omega_2}$ such that from any initial belief in B_1 , μ_t converges to δ_{ω_1} with positive probability, while from any initial belief in B_2 , μ_t a.s. leaves B_2 . By condition (ii), one can find some T such that with positive probability, μ_t reaches B_1 within T periods from any initial belief $\mu \notin B_1 \cup B_2$. Combining these observations, a simple recursive argument shows that μ_t converges to δ_{ω_1} a.s. from any initial belief.

Beyond binary states, say if $\Omega = \{\omega_1, \omega_2, \omega_3\}$, a complication with the above argument is the following.²⁵ Even if δ_{ω_1} is locally stable and δ_{ω_2} and δ_{ω_3} are unstable, condition (ii) is

²⁴For example, if $Z = \{0, 1\}$, then by the same logic as in Example 1, this is the case if near all δ_ω , we have $p_\mu(1) \leq \hat{p}_\mu(1|\omega_1) < \dots < \hat{p}_\mu(1|\omega_N)$; as we will see, this arises naturally in the applications in Section 4.2.

²⁵[Bohren and Hauser \(2021\)](#) address related challenges under binary states but heterogeneous models.

consistent with beliefs getting stuck in a neighborhood of the subsimplex $\Delta(\{\omega_2, \omega_3\})$ and cycling forever between δ_{ω_2} and δ_{ω_3} . However, this is ruled out by the uniform ranking over states that condition (i) imposes near point-mass beliefs. Indeed, as we show using similar martingale arguments as before, the latter ensures that whenever beliefs approach δ_{ω_2} or δ_{ω_3} , they must escape in the direction of δ_{ω_1} with positive probability.

Finally, one might also be interested in a weak form of global stability, which only requires that from all initial beliefs, process μ_t converges to δ_{ω_1} with *positive* probability (rather than with probability one, as ensured by our results). Using similar arguments as above, it can be shown that δ_{ω_1} is globally stable in this weak sense if it satisfies the local stability condition (6) and if condition (ii) in Proposition 1 is only imposed for $n = 1$. Note that under this weak notion, multiple beliefs δ_ω can be globally stable.

4 Applications

We now apply the preceding stability results to two classes of economic applications.

4.1 Active learning under one-dimensional states

First, we consider single-agent active learning under rich one-dimensional states, $\Omega \subseteq \mathbb{R}$, as in many important applications in the literature. We show how the iterated elimination criterion in Theorem 2 is straightforward to verify in this setting, providing a simple and unified method to establish global stability.

For ease of exposition, we assume that $\Omega = [\underline{\omega}, \bar{\omega}]$ is a compact interval; as noted, Appendix B.1 shows that Theorem 2 remains valid in this case.²⁶ Consider an active learning environment as in Remark 1. Recall that G_a and $(\hat{G}_a(\cdot|\omega))_{\omega \in \Omega}$ denote the true and perceived signal distributions when action a is chosen. Assume that the agent's action set $A \subseteq \mathbb{R}$ is an interval, that her action choices $a : \Delta(\Omega) \rightarrow A$ are FOSD-increasing and continuous, and that $\text{KL}(G_a, \hat{G}_a(\cdot|\omega))$ is strictly quasi-convex in ω and continuous in (a, ω) .²⁷

These assumptions ensure that for each ω , there is a unique state $m(\omega)$ that is KL-dominant at δ_ω , i.e., $m(\omega) \succ_{\delta_\omega}^{\text{KL}} \omega'$ for all $\omega' \neq m(\omega)$. Observe that ω is a fixed point of the one-dimensional map $m : \Omega \rightarrow \Omega$ if and only if δ_ω is a strict BeNE.

²⁶Similar analysis goes through whenever Ω is a finite but sufficiently dense subset of $[\underline{\omega}, \bar{\omega}]$, as in this case $S^\infty(\Omega)$ approximates $S^\infty([\underline{\omega}, \bar{\omega}])$ (see Appendix F of the previous version, Frick, Iijima, and Ishii, 2020b).

²⁷A natural setting that satisfies strict quasi-convexity is the following. Suppose $G_a = H_{\hat{\phi}(a)}$ and $\hat{G}_a(\cdot|\omega) = H_{\hat{\phi}(a, \omega)}$ for some family $(H_\theta)_{\theta \in \Theta} \in \Delta(\mathbb{R})$ of distributions that satisfy the strict monotone likelihood ratio property with respect to the parameter $\theta \in \Theta \subseteq \mathbb{R}$. If $\hat{\phi}$ is strictly increasing in ω , then $-\text{KL}(H_{\hat{\phi}(a)}, H_{\hat{\phi}(a, \omega)})$ is strictly single-peaked in ω , and hence $\text{KL}(H_{\hat{\phi}(a)}, H_{\hat{\phi}(a, \omega)})$ is strictly quasi-convex in ω .

The following result shows that iterated elimination of dominated states corresponds to iterated application of the map m . Moreover, simple conditions that only involve considering the fixed points of the maps m or m^2 yield that $S^\infty(\Omega) = \{\hat{\omega}\}$ is a singleton, in which case Theorem 2 implies that δ_ω is globally stable:²⁸

Proposition 2. *For all $k = 1, 2, \dots, \infty$, we have $S^k(\Omega) = m^k(\Omega)$. Moreover:*

1. *Suppose m is weakly increasing. Then $S^\infty(\Omega) = \{\hat{\omega}\}$ if and only if $\hat{\omega}$ is the unique fixed point of m .*
2. *Suppose m is weakly decreasing. Then $S^\infty(\Omega) = \{\hat{\omega}\}$ if and only if $\hat{\omega}$ is the unique fixed point of m^2 .*

Deriving m is straightforward in many applications in the literature, and many natural forms of misspecification that are considered induce an increasing or decreasing m . To nest these applications, assume further that $Z = \mathbb{R}$ and that action a induces the true signal distribution according to $z = g(a) + \varepsilon$, but the agent perceives signals in state ω to follow $z = \hat{g}(a, \omega) + \varepsilon$, where $g : A \rightarrow \mathbb{R}$ and $\hat{g} : A \times \Omega \rightarrow \mathbb{R}$ are continuously differentiable with $\frac{\partial \hat{g}}{\partial \omega} > 0$, and the mean-zero noise term ε is distributed according to a log-concave and strictly positive density on \mathbb{R} .²⁹ Then, letting $a(\omega) := a(\delta_\omega)$, any interior $m(\omega) \in (\underline{\omega}, \bar{\omega})$ solves

$$\hat{g}(a(\omega), m(\omega)) = g(a(\omega)). \quad (10)$$

Thus, m is weakly increasing if and only if $\frac{dg}{da}(a(\omega)) \geq \frac{\partial \hat{g}}{\partial a}(a(\omega), m(\omega))$, and decreasing if and only if $\frac{dg}{da}(a(\omega)) \leq \frac{\partial \hat{g}}{\partial a}(a(\omega), m(\omega))$ for all ω , capturing that the agent either under- or overstates the marginal effect of her actions on signals.

For example, based on this, one can establish global stability in the following applications:

- **Misspecified monopoly pricing:** Consider a monopolist who is learning about his demand function. Here, $z = g(a) + \varepsilon$ represents the true demand faced by the monopolist when he sets price $a \in A = \mathbb{R}_+$, where $g(a) = g(a, \omega^*) = \omega^* - \beta a$. The true intercept of demand $\omega^* \in [\underline{\omega}, \bar{\omega}] \subseteq \mathbb{R}_+$ is unknown to the monopolist. In updating beliefs about ω^* , he misperceives the slope of demand β to be $\hat{\beta}$, where $\beta, \hat{\beta} > 0$. Thus, $\hat{g}(a, \omega) = \omega - \hat{\beta}a$ for each ω . Each period t , the monopolist myopically maximizes

²⁸Parts 1 and 2 of Proposition 2 parallel conditions for dominance solvability in games with strategic complements (Milgrom and Roberts, 1990) and substitutes (Zimper, 2007), respectively.

²⁹The assumption that signals can take any real values is not essential. For example, the same conclusion holds under binary signals, $Z = \{0, 1\}$, where the true and perceived probabilities of signal 1 are respectively $g(a)$ and $\hat{g}(a, \omega)$ for each a and ω .

expected revenue, i.e., his price a_t as a function of his belief is $a(\mu_t)$, where

$$a(\mu) = \operatorname{argmax}_{a \in \mathbb{R}_+} a \times \left(\mathbb{E}_\mu[\omega] - \hat{\beta}a \right) = \frac{\mathbb{E}_\mu[\omega]}{2\hat{\beta}}.$$

In particular, $a(\omega) = \frac{\omega}{2\hat{\beta}}$. By the above, map m is increasing/decreasing if the monopolist over-/underestimates the slope of demand β , and $m(\omega) = \omega^* + \frac{\omega}{2\hat{\beta}}(\hat{\beta} - \beta)$ when this is interior. If $|\frac{\hat{\beta}-\beta}{2\hat{\beta}}| < 1$, then m and m^2 are contractions, and thus admit a unique fixed point $\hat{\omega}$, where $\hat{\omega} = \frac{2\hat{\beta}\omega^*}{\hat{\beta}+\beta}$ when this is interior. Hence, $\delta_{\hat{\omega}}$ is globally stable by Proposition 2 and Theorem 2. While Esponda and Pouzo (2016) and Heidhues, Kőszegi, and Strack (2021) establish analogous results using stochastic approximation arguments that rely on Gaussian signal distributions, our approach does not require this parametric assumption.

- **Effort choice under overconfidence:** In Heidhues, Kőszegi, and Strack (2018) (HKS), g and \hat{g} take the form $g(a) = Q(a, \beta, \omega^*)$ and $\hat{g}(a, \omega) = Q(a, \hat{\beta}, \omega)$ for some function Q . Here, signals z can be interpreted as output, actions a as effort choice, states ω as project quality (with true quality ω^*), and β and $\hat{\beta}$ as the agent's true and perceived ability. The agent chooses $a(\mu)$ to maximize expected output. When the agent is overconfident ($\hat{\beta} > \beta$), the natural assumptions that HKS impose on the output function Q ensure that m is increasing with a unique fixed point $\hat{\omega}$, where $\hat{\omega} < \omega^*$ (see Online Appendix C.2.1). Thus, Proposition 2 and Theorem 2 immediately imply HKS's result that the pessimistic belief $\delta_{\hat{\omega}}$ is globally stable.
- **Optimal stopping under the gambler's fallacy:** Similar reasoning yields the global stability result in He (2021), where m can again be seen to be increasing and admit a unique fixed point (see Online Appendix C.2.2).

Esponda, Pouzo, and Yamamoto (2021) (Section 7) consider a similar one-dimensional state setting and provide conditions for local/global stability and instability.³⁰ While we consider continuous actions in this section, their results assume finite actions; however, the more recent paper by Murooka and Yamamoto (2021) extends their approach to continuous actions. The approaches in Esponda, Pouzo, and Yamamoto (2021) and Murooka and Yamamoto (2021) are based on characterizing limiting action frequencies by means of a differential inclusion; different and complementary to this, we provide an approach based on iterated elimination that uses martingale arguments to establish belief convergence.

³⁰Our iterated elimination approach can also be extended to study local stability in the current setting; see Appendix B.2.

4.2 Slow learning and fragility of long-run beliefs

Next, we present two applications that illustrate how to apply our results to environments with slow learning. Our analysis highlights how slow learning can render long-run beliefs fragile against misspecification. Section 4.2.3 contrasts these findings with other recent work that has examined the robustness of learning outcomes to misspecification.

Throughout, we consider finite state spaces $\Omega = \{\omega_1, \dots, \omega_N\} \subseteq \mathbb{R}_+$, with $\omega_1 < \dots < \omega_N$.

4.2.1 Costly information acquisition

Consider a single-agent active learning setting, where the agent learns about some state (e.g., her ability) by acquiring costly information (e.g., seeking out expert feedback). The fixed and unknown true state is $\omega^* \in \Omega$. Each period t , the agent chooses a precision parameter $\gamma_t \in [0, \bar{\gamma}]$ at cost $C(\gamma_t)$. She then observes a signal z_t that is 1 (“good news”) with probability $\gamma_t \omega^* + \beta$ and 0 (“bad news”) otherwise; here, β is the state-independent base rate of the high signal, over which the agent has no control. In updating her beliefs $\mu_t \in \Delta(\Omega)$, the agent misperceives the base rate β to be $\hat{\beta}$. For example, if $\hat{\beta} < \beta$, this implies a form of “ego-biased” belief-updating: the agent overreacts to good news about her ability, but underreacts to bad news (e.g., Eil and Rao, 2011; Mobius, Niederle, Niehaus, and Rosenblat, 2014).

Note that true and perceived signal distributions are (Blackwell-)more informative the greater γ_t and are uninformative when $\gamma_t = 0$. Assume the agent has positive value to information, as captured by a utility $v : \Delta(\Omega) \rightarrow \mathbb{R}$ that is continuous and strictly convex in her current belief.³¹ Each period, she chooses γ_t as a function of her current belief μ_t to myopically maximize expected utility net of the cost (myopia is assumed for simplicity):³²

$$\gamma_t = \gamma_{\hat{\beta}}(\mu_t) \in \operatorname{argmax}_{\gamma \in [0, \bar{\gamma}]} \hat{\mathbb{E}}_{\mu_t}[v(\mu_{t+1}(\gamma))] - C(\gamma), \quad (11)$$

where $\mu_{t+1}(\gamma)$ denotes the agent’s random posterior following period- t signal realizations and the expectation $\hat{\mathbb{E}}_{\mu_t}$ is with respect to the perceived signal distribution. Assume $\bar{\gamma} \in (0, 1)$ and $\beta, \hat{\beta} \in (0, 1 - \bar{\gamma})$ are such that true and perceived signal probabilities $p_{\mu}(1|\omega^*) = \gamma_{\hat{\beta}}(\mu)\omega^* + \beta$ and $\hat{p}_{\mu}(1|\omega) = \gamma_{\hat{\beta}}(\mu)\omega + \hat{\beta}$ are always well-defined and nondegenerate. We also assume that $\gamma_{\hat{\beta}}(\mu)$ is continuous in μ .³³

³¹For example, suppose that $v(\mu) = \max_{a \in \mathbb{R}} \mathbb{E}_{\mu}[-(a - \omega)^2]$ is the indirect utility to a prediction problem that the agent must solve at the end of each period (where realized payoffs are not observed until some exogenously distributed stopping time).

³²All results generalize to forward-looking agents. In particular, note Lemma 5 (slow learning) remains valid with the same proof, as the continuation value is continuous in μ .

³³Without continuity, the main result (Proposition 3) remains valid under the following assumption: for any compact set K of mixed beliefs, $\inf_{\mu \in K} \gamma_{\hat{\beta}}(\mu) > 0$. This is slightly stronger than the current assumption

First, suppose that the agent incurs the same constant cost $C(\gamma) = c$ for any precision choice γ , so information is effectively costless. Then learning is successful when the agent is correctly specified ($\hat{\beta} = \beta$) and successful learning is robust to small amounts of misspecification. Formally, say that **learning is successful** at ω^* if, when the true state is ω^* , we have $\mathbb{P}_\mu[\mu_t \rightarrow \delta_{\omega^*}] = 1$ for all beliefs $\mu \in \Delta(\Omega)$ with $\mu(\omega^*) > 0$.

Lemma 4. *Suppose C is constant. For any β , there exists $\varepsilon > 0$ such that for any $\hat{\beta}$ with $|\hat{\beta} - \beta| \leq \varepsilon$, learning is successful in all true states ω^* .*

When information is costless, then for all $\hat{\beta}$, the agent always chooses the maximal precision $\bar{\gamma}$. This implies that when $\hat{\beta} = \beta$, the true state ω^* strictly dominates all other states ω at all beliefs μ , where, importantly, the relative prediction accuracy $\sum_z p_\mu(z) \left(\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega^*)} \right)^q < 1$ is independent of μ . Given this, the same is true whenever $\hat{\beta}$ is sufficiently close to β , based on which we conclude that learning is successful.

Next, suppose information is costly, in the sense that C is strictly increasing in γ . The key departure this introduces is the following:

Lemma 5. *Suppose C is strictly increasing. For any $\hat{\beta}$, $\lim_{\mu \rightarrow \delta_\omega} \gamma_{\hat{\beta}}(\mu) = 0$ for every ω .*

That is, if information is (even slightly) costly, then the agent stops acquiring information in the limit as she becomes confident in any particular state ω , because her value to information vanishes as she grows confident. Lemma 5 implies that costly information leads to slow learning, since the agent's perceived signal probabilities satisfy

$$\lim_{\mu \rightarrow \delta_\omega} \hat{p}_\mu(1|\omega') = \lim_{\mu \rightarrow \delta_\omega} \gamma_{\hat{\beta}}(\mu)\omega' + \hat{\beta} = \hat{\beta}, \quad \forall \omega, \omega'.$$

Based on this, we show that learning under costly information is fragile against misspecification: Suppose learning is successful whenever the agent is correctly specified. Then, in sharp contrast with Lemma 4, arbitrarily small amounts of misspecification not only break successful learning, but indeed render the agent's long-run belief *independent* of the true state ω^* : If $\beta > \hat{\beta}$ (resp. $\beta < \hat{\beta}$), then regardless of ω^* , she becomes confident in the highest (resp. lowest) possible state.

Proposition 3. *Suppose C is strictly increasing and for any $\beta, \hat{\beta}$ with $\beta = \hat{\beta}$, learning is successful at all states ω^* . Then:*

1. *For any $\beta, \hat{\beta}$ with $\beta > \hat{\beta}$, δ_{ω_N} is globally stable in all true states ω^* .*
2. *For any $\beta, \hat{\beta}$ with $\beta < \hat{\beta}$, δ_{ω_1} is globally stable in all true states ω^* .*




Figure 1: Prediction accuracy ranking of ω_1 vs. ω_2 as a function of μ when $\omega^* = \omega_1$. Left: $\hat{\beta} = \beta$. Right: $\beta > \hat{\beta}$. Here $D_\mu^{KL}(\omega_2, \omega_1) := \text{KL}(P_\mu(\cdot|\omega^*), \hat{P}_\mu(\cdot|\omega_2)) - \text{KL}(P_\mu(\cdot|\omega^*), \hat{P}_\mu(\cdot|\omega_1))$.

When feedback is costless, then, by Lemma 4, a small propensity for ego-biased interpretation of signals does not prevent the agent from learning her ability. In contrast, Proposition 3 shows that if obtaining feedback requires just a slight amount of effort, then even arbitrarily small amounts of this bias are greatly amplified over time, leading to drastic overconfidence in the long run.

To see the idea, suppose that $\Omega = \{\omega_1, \omega_2\}$ and the true state is ω_1 . For any $\hat{\beta}$, the fact that learning is successful at all states when $\beta = \hat{\beta}$ means that $\gamma_{\hat{\beta}}(\mu) > 0$ for all mixed μ ; otherwise the agent’s belief would get stuck at some initial mixed beliefs. At the same time, by Lemma 5, $\lim_{\mu \rightarrow \delta_\omega} \gamma_{\hat{\beta}}(\mu) = 0$. As a result, when $\beta = \hat{\beta}$, the true state ω_1 strictly dominates ω_2 at all mixed beliefs, but in contrast with costless learning, the gap in prediction accuracy now vanishes as beliefs approach δ_{ω_1} or δ_{ω_2} . As shown in Figure 1, this makes the prediction accuracy ranking near point-mass beliefs highly sensitive to misspecification.³⁴

Indeed, if $\beta > \hat{\beta}$, the ranking between ω_1 and ω_2 is reversed: Since γ is very small near point-mass beliefs, the true probability $\gamma\omega_1 + \beta$ of the high signal exceeds the perceived probabilities $\gamma\omega_2 + \hat{\beta}$, $\gamma\omega_1 + \hat{\beta}$ in both states, but because $\omega_2 > \omega_1$, the perceived probability in state ω_2 comes closer to the truth. By the logic in Example 1, this implies $\omega_2 \succ_\mu^q \omega_1$ for all $q \in (0, 1)$ and μ near δ_{ω_1} and δ_{ω_2} . Intuitively, if signals are precise (γ is high), the true state always explains the agent’s observations best, but if signals are sufficiently imprecise (γ is low), then overestimating the state can partly compensate for underestimating the base rate of the high signal. Finally, since ω_2 strictly dominates ω_1 near both point-mass beliefs and the probabilities of the high signal are increasing in states, Proposition 1 applies up to relabeling states in decreasing order. Thus, when $\beta > \hat{\beta}$, δ_{ω_2} is globally stable.³⁵

(“successful learning at all states when $\hat{\beta} = \beta$ ”), which is equivalent to the requirement that $\gamma_{\hat{\beta}}(\mu) > 0$ for all mixed μ (Lemma 6). The robustness of costless learning (Lemma 4) does not rely on continuity.

³⁴The figure uses KL-dominance for the sake of graphical illustration, but the proof relies on q -dominance.

³⁵Proposition 3 does not rely on the specific true and perceived signal distributions in the text: Indeed, writing signal probabilities as functions of the agent’s choice γ and assuming slow learning, part 1 (resp. 2) generalizes as long as (i) $\hat{p}_\gamma(1|\omega_N) < p_\gamma(1|\omega^*)$ (resp. $\hat{p}_\gamma(1|\omega_1) > p_\gamma(1|\omega^*)$) for all $\omega^* \in \Omega$ and small enough $\gamma > 0$, and (ii) $\hat{p}_\gamma(1|\omega_n)$ is strictly increasing in n at each $\gamma > 0$. Unlike the specification in the text, this allows for specifications where the true and perceived long-run signal distributions (i.e., at $\gamma = 0$) coincide.

Finally, to understand when Proposition 3 applies, we clarify which cost functions lead to successful learning when the agent is correctly specified. To state this, we slightly strengthen the requirement that the utility $v : \Delta(\Omega) \rightarrow \mathbb{R}$ is strictly convex, as follows:

Lemma 6. *Suppose v is twice continuously differentiable with a positive-definite Hessian. Fix any $\hat{\beta}$. For any twice continuously differentiable cost function C with $C'(0) = C''(0) = 0$,*

$$\gamma_{\hat{\beta}}(\mu) > 0 \text{ for all mixed } \mu. \quad (12)$$

Moreover, (12) is necessary and sufficient for learning to be successful at all ω^* when $\beta = \hat{\beta}$.

Lemma 6 provides “Inada” conditions on C which ensure that small amounts of information are very cheap. Thus, the agent remains willing to acquire a positive amount of information whenever she is not completely certain about the state. These conditions are satisfied, for example, by any power function $C(\gamma) = \gamma^d$ with $d > 2$.³⁶

4.2.2 Sequential social learning

Consider the sequential social learning setting from Remark 1, with the following additional assumptions. Private signals s_t at each state ω are drawn according to a positive and continuous density $\phi(\cdot|\omega)$ that satisfies the monotone likelihood ratio property. True and perceived type distributions F and \hat{F} admit positive densities over \mathbb{R} . The utility difference $v(\theta, \omega) := u(1, \theta, \omega) - u(0, \theta, \omega)$ between the two actions is strictly increasing and continuous in types and states (θ, ω) , with $\lim_{\theta \rightarrow -\infty} v(\theta, \omega) < 0$ and $\lim_{\theta \rightarrow +\infty} v(\theta, \omega) > 0$; that is, sufficiently low (risk-averse) types always prefer action 0 (not adopt) and sufficiently high (risk-tolerant) types always prefer action 1 (adopt).

Then the true and perceived probabilities of observing action 0 at public belief μ are

$$p_{\mu}(0|\omega^*) = \int F(\theta^*(\mu^s))\phi(s|\omega^*) ds, \quad \hat{p}_{\mu}(0|\omega) = \int \hat{F}(\theta^*(\mu^s))\phi(s|\omega) ds,$$

where $\mu^s \in \Delta(\Omega)$ denotes the Bayesian update of μ following private signal realization s , and $\theta^*(\nu)$ denotes the type who is indifferent between action 0 and 1 at belief ν . Note that $\theta^*(\nu)$ exists and is unique for each ν by the above assumptions. We write $\theta_{\omega}^* := \theta^*(\delta_{\omega})$ and $\theta_i^* := \theta_{\omega_i}^*$. Observe that θ_i^* is strictly decreasing in i , as $\omega_1 < \dots < \omega_N$.

We first note that when agents are correctly specified, learning is successful:

Lemma 7. *Suppose that $\hat{F} = F$. Then learning is successful in all true states ω^* .*

³⁶The restriction $C''(0) = 0$ on the second derivative is related to the Radner-Stiglitz non-concavity in the value of information (Chade and Schlee, 2002). Since the agent’s marginal value of information is zero at $\gamma = 0$, the restriction $C'(0) = 0$ on the first derivative is not enough to ensure a positive choice of γ .

An analogous result is established by [Goeree, Palfrey, and Rogers \(2006\)](#). Observe that herding is ruled out here due to rich preference heterogeneity (in particular, the existence of dominant types), despite the fact that private signals need not have unbounded precision.

However, we observe next that sequential social learning leads to slow learning:

Lemma 8. *For all \hat{F} , ω , and ω' , we have $\lim_{\mu \rightarrow \delta_\omega} \int \hat{F}(\theta^*(\mu^s)) \phi(s|\omega') ds = \hat{F}(\theta_\omega^*)$.*

Lemma 8 shows that as the public belief becomes confident in any given state ω , the perceived probability of observing action 0, $\lim_{\mu \rightarrow \delta_\omega} \hat{p}_\mu(0|\omega') = \hat{F}(\theta_\omega^*)$, is the same in all states ω' ; that is, (2) holds. This reflects the familiar slow-learning logic under sequential social learning that we discussed in Section 2.3.

Similar to costly information acquisition, this again leads successful learning to be highly fragile against misspecification. The following result classifies possible learning outcomes:

Proposition 4. *Fix any F and \hat{F} . In each true state ω^* :*

1. δ_{ω_N} is globally stable if $F(\theta_i^*) < \hat{F}(\theta_i^*)$ for all i , locally stable if $F(\theta_N^*) < \hat{F}(\theta_N^*)$, and unstable if $F(\theta_N^*) > \hat{F}(\theta_N^*)$.
2. δ_{ω_1} is globally stable if $F(\theta_i^*) > \hat{F}(\theta_i^*)$ for all i , locally stable if $F(\theta_1^*) > \hat{F}(\theta_1^*)$, and unstable if $F(\theta_1^*) < \hat{F}(\theta_1^*)$.
3. For each $n \in \{2, \dots, N-1\}$, δ_{ω_n} is unstable if $F(\theta_n^*) \neq \hat{F}(\theta_n^*)$.

Depending on the nature of misspecification, Proposition 4 highlights three general possibilities. First, beliefs might converge globally to a point-mass on the highest (resp. lowest) state. Similar to Proposition 3, this occurs if agents systematically underestimate (resp. overestimate) the type distribution (e.g., extent of risk tolerance in the population), no matter how close \hat{F} is to F and regardless of the true state ω^* . Second, the extreme beliefs δ_{ω_1} and/or δ_{ω_N} might be locally stable, if agents overestimate the share of very high types (above θ_1^*) and/or of very low types (below θ_N^*). Finally, if agents underestimate both the shares of very high types and of very low types (i.e., underestimate type heterogeneity), then generically *all* point-mass beliefs are unstable, so beliefs cycle.³⁷

To see the idea, consider any ω_i . If $F(\theta_i^*) < \hat{F}(\theta_i^*)$, then Lemma 8 implies that at all public beliefs μ close to the point-mass belief δ_{ω_i} , the perceived probability of action 0, $\hat{p}_\mu(0|\omega) \approx \hat{F}(\theta_i^*)$, is strictly higher in all states ω than the actual probability $p_\mu(0|\omega^*) \approx F(\theta_i^*)$. At the

³⁷Relatedly, [Gagnon-Bartsch \(2017\)](#) considers sequential social learning with “taste projection” and shows that a point-mass on the true state can be unstable under arbitrarily small misspecification. His environment can be seen to also feature slow learning, but due to the difference in the nature of misspecification, his setting requires large misspecification in order for a point-mass on an incorrect state to be locally/globally stable.

same time, by the assumptions on signals and utilities, $\hat{p}_\mu(0|\omega)$ is strictly decreasing in ω at all mixed μ . Thus, at all mixed μ close to δ_{ω_i} , the perceived action distribution comes closest to the actual one at the highest state ω_N . Analogously, if $F(\theta_i^*) > \hat{F}(\theta_i^*)$, then the lowest state ω_1 dominates all other states near δ_{ω_i} . Based on this, the local stability and instability results follow from Theorem 1, while Proposition 1 implies the global stability results.

4.2.3 Discussion

Our finding that slow learning can lead to fragility against misspecification complements other recent work. [Bohren and Hauser \(2021\)](#) (BH) establish a general *robustness* result for misspecified learning in their setting: If learning is successful under correct specification, then learning is also successful whenever agents’ perceptions are close enough to the true model. The key difference is that they consider environments that do not feature slow learning, because, even near point-mass beliefs, agents take actions that generate non-vanishingly informative signals. For instance, this is naturally the case under costless learning as well as the examples analyzed in Section 4.1. Intuitively, robustness in these settings results from the fact that, under correct specification, the difference in prediction accuracy between the true state ω^* and all other states is bounded away from zero; given this, the same remains true under small enough amounts of misspecification, similar to the logic in Lemma 4. By contrast, when learning is slow, as in Sections 4.2.1–4.2.2, then differences in prediction accuracy vanish near point-mass beliefs. As illustrated above, this renders the prediction accuracy ranking, and hence stable beliefs, highly sensitive to small amounts of misspecification.

Even under costly information acquisition or social learning, the usual slow-learning logic might hold only approximately if other offsetting forces are introduced: For example, agents might have access to small amounts of exogenous costless information each period (similarly, under social learning, BH introduce a small fraction of “autarkic” agents, who act solely based on their private information, ignoring others’ actions). For a fixed positive amount of such exogenous information, the results in BH imply that learning is successful whenever agents’ perceptions are within some small enough threshold $\varepsilon > 0$ of the true model. Complementary to this, our analysis implies that the smaller the amount of exogenous information, the smaller is ε (i.e., the more sensitive is learning to misspecification), and in the limit as there is no exogenous information, vanishingly small amounts of misspecification can generate extreme failures of learning. The following example illustrates this point:

Example 2. Consider the setting in Section 4.2.1. Suppose that true and perceived probabilities are $\beta + (\gamma(\mu) + \alpha)\omega^*$ and $\hat{\beta} + (\gamma(\mu) + \alpha)\omega$, where $\alpha > 0$ captures exogenous information. Then for any $\hat{\beta} > \beta$ (resp. $\hat{\beta} < \beta$), there exists $\bar{\alpha} > 0$ such that whenever $\alpha < \bar{\alpha}$, then δ_{ω_N}

(resp. δ_{ω_1}) is globally stable at all ω^* . Here, $\bar{\alpha}$ can be chosen to be decreasing in $\varepsilon = |\hat{\beta} - \beta|$, with $\bar{\alpha}(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$. ▲

Taken together, these results suggest that some policy interventions, such as releasing additional public signals or shutting down some agents’ observations of others’ actions, might be used to “robustify” learning against misspecification, but that the effectiveness of such interventions would depend on the relative strength of additional information and agents’ amount of misspecification.

The slow learning channel we highlight also complements other fragility results in the literature. Frick, Iijima, and Ishii (2020a) (FII20) study a different social learning model, with a continuum of states and continuum of agents, who each *privately* observe the action of a random other agent each period. Importantly, the fact that action observations are private means that the setting in FII20 is not nested by the current paper, nor by BH, as these papers require a public belief process. As a result, the preceding discussion on robustness/fragility without/with slow learning does not apply. Indeed, as FII20 show, their setting does not feature slow learning: Agents view their new private action observations as non-vanishingly informative, no matter how confident they themselves have become in a particular state.³⁸ Yet, despite the absence of slow learning, FII20 establish that arbitrarily small misspecification about the type distribution F can lead beliefs to converge to a state-independent point-mass, similar to the current fragility result in Proposition 4. The mechanism behind the two fragility results is quite different.³⁹ One notable manifestation of this difference is that the fragility result in FII20 relies on a continuous state space: FII20 show that, in their setting, successful learning is robust if the state space is finite, in contrast with Proposition 4.

Cho and Kasa (2017) consider single-agent learning under a Markovian fundamental. Their setting is also not nested by ours and does not feature slow learning, but they show that long-run beliefs can be discontinuous against the details of the agent’s misspecification. Their discontinuity result holds away from the correctly specified benchmark and relies on intertemporal correlation in the signal process.

³⁸This is because, under random matching in a continuum population, the history observed by an agent’s new match in each period almost surely has no overlap with her own history of observations.

³⁹Specifically, in Section 4.2.2, slow learning implies that *all* point-mass beliefs are BeNE, and the logic behind Proposition 4 is that misspecification can discontinuously change which of these beliefs are *stable*. By contrast, FII20 highlight a discontinuity at the level of the equilibrium correspondence: all point-mass beliefs are BeNE under correct specification, but misspecification can discontinuously shrink the BeNE set to a *single* state-independent point-mass.

5 Concluding remarks

This paper presents an approach to analyze belief convergence in a broad class of misspecified learning environments, including single-agent and social learning. The key ingredients underlying our approach are (i) a novel prediction accuracy order over subjective models, q -dominance, and (ii) the observation that throughout any region of beliefs where q -dominance obtains, standard martingale arguments from the correctly specified setting can be applied locally. Based on this, we obtain conditions for local/global stability or instability of long-run beliefs. One difference with existing approaches is that our results can be applied to study the impact of misspecification when learning is slow. When this is the case, as is natural under costly information acquisition or social learning, we illustrate that successful learning can be highly fragile against misspecification. We also apply our results to unify and generalize various convergence results in previously studied settings.

Fruitful directions in which to extend our results include multi-agent settings with heterogeneous beliefs (partially addressed in Appendix G of the previous version, [Frick, Iijima, and Ishii, 2020b](#)) and Markov decision problems. Another important direction that we leave open is to analyze when a mixed belief (or region of mixed beliefs) μ^* is stable: This can be seen as an extreme form of slow learning, where belief-updating ceases completely before agents have become confident in any given state, and arises in some important economic applications (see Section 2.2). We expect that stability conditions for this case might be obtained by again requiring a suitable transformation of the posterior ratio process to be a nonnegative supermartingale near μ^* .⁴⁰

Appendix

Appendix A contains all proofs for Section 3 (Lemma 1 is immediate from Theorem 1). Appendix B extends the stability analysis to infinite state spaces. The proofs for the applications in Section 4, as well as all supplemental material referenced in the text, appear in Online Appendices C–D.

⁴⁰More specifically, our q -dominance condition for local stability of δ_ω in Theorem 1 ensures that $f(\mu_t) = \left(\frac{\mu_t(\omega')}{\mu_t(\omega)}\right)^q$ is a nonnegative supermartingale at beliefs near δ_ω . To establish the local stability of a mixed belief μ^* , a similar approach would be to construct a function f that is minimized at μ^* and such that $f(\mu_t)$ strictly decreases in expectation near μ^* (similar arguments establish stability of a region of mixed beliefs μ^*). The key new step would be to identify suitable conditions on P and \hat{P} that yield such a function f .

A Proofs for Section 3

A.1 Preliminary results

Say *belief continuity holds at* $M \subseteq \Delta(\Omega)$ if for each $\omega, \omega' \in \Omega$, $\mu \in M$ and $z \in Z$, we have that $p_\mu(z)$, $\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')}$, and $p_\mu(z) \frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')}$ are continuous in μ .

Lemma 9. *Assume belief continuity holds at* $M \subseteq \Delta(\Omega)$. *Pick* $q^* > 0$ *as in Assumption 1.2. For all* ω, ω' *and* q *with* $0 < q \leq \min\{q^*, 1\}$, $\int \left(\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')}\right)^q p_\mu(z) d\nu(z)$ *is continuous in* μ *on* M .

Proof. Fix ω, ω' and $q \in (0, q^*]$. Consider $\mu \in M$ and a sequence $\mu_n \rightarrow \mu$. Observe that belief continuity implies that, for each z ,

$$\left(\frac{\hat{p}_{\mu_n}(z|\omega)}{\hat{p}_{\mu_n}(z|\omega')}\right)^q p_{\mu_n}(z) \rightarrow \left(\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')}\right)^q p_\mu(z).$$

When $p_\mu(z) > 0$, the claim is clear. When $p_\mu(z) = 0$, this holds because $\frac{\hat{p}_{\mu_n}(z|\omega)}{\hat{p}_{\mu_n}(z|\omega')} p_{\mu_n}(z) \rightarrow \frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')} p_\mu(z) = 0$ and $q \leq 1$.

Given the above observation, the desired continuity holds by the dominated convergence theorem, as $\left(\frac{\hat{p}_\mu(\cdot|\omega)}{\hat{p}_\mu(\cdot|\omega')}\right)^q p_\mu(\cdot)$ is dominated by the ν -integrable $h(\cdot)$ (by Assumption 1.2 and $q \leq q^*$). \square

The following result shows that mixed beliefs are unstable under an identification condition. The argument is similar to Theorem B.1 in Smith and Sørensen (2000):

Lemma 10. *Take any compact set* $K \subseteq \Delta(\Omega)$ *at which belief continuity holds. Suppose there exist* ω, ω' *such that for each* $\mu \in K$, *we have (i) $\mu(\omega), \mu(\omega') > 0$ and (ii) $\hat{p}_\mu(z|\omega) \neq \hat{p}_\mu(z|\omega')$ for some* $z \in \text{supp}(P_\mu)$. *Then for any initial belief* μ_0 , $\mathbb{P}_{\mu_0}[\exists \tau < \infty \text{ s.t. } \mu_t \in K \forall t \geq \tau, \text{ and } \exists \lim_t \frac{\mu_t(\omega)}{\mu_t(\omega')}] = 0$.

Proof. For each $\mu \in K$, (ii) yields some $z_\mu \in \text{supp}(P_\mu)$ such that $\left|\log \frac{\hat{p}_\mu(z_\mu|\omega)}{\hat{p}_\mu(z_\mu|\omega')}\right| > 0$. Since perceived signal densities are continuous in z , there exists a neighborhood $Z_\mu \ni z_\mu$ with

$$\inf_{z \in Z_\mu} \left| \log \frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')} \right| > 0, \quad P_\mu(Z_\mu) > 0.$$

By belief continuity at K , there exists a neighborhood $B_\mu \ni \mu$ such that

$$\inf_{z \in Z_\mu, \mu' \in B_\mu} \left| \log \frac{\hat{p}_{\mu'}(z|\omega)}{\hat{p}_{\mu'}(z|\omega')} \right| > 0, \quad \inf_{\mu' \in B_\mu} P_{\mu'}(Z_\mu) > 0.$$

By compactness of K , there is a finite subcover $(B_{\mu^i})_{i=1}^n$ of K . Thus, there is $\gamma > 0$ such that

$$\inf_{z \in Z_{\mu^i}, \mu' \in B_{\mu^i}} \left| \log \frac{\hat{p}_{\mu'}(z|\omega)}{\hat{p}_{\mu'}(z|\omega')} \right| > \gamma, \quad \inf_{\mu' \in B_{\mu^i}} P_{\mu'}(Z_{\mu^i}) > \gamma, \quad \text{for all } i = 1, \dots, n.$$

Suppose for a contradiction that $\mathbb{P}_{\mu_0}[\exists \tau < \infty \text{ s.t. } \mu_t \in K \forall t \geq \tau, \text{ and } \exists \lim_t \frac{\mu_t(\omega)}{\mu_t(\omega')}] > 0$ for some initial belief μ_0 . Since the belief process is Markov, there exists an initial belief $\mu'_0 \in K$ such that

$\mathbb{P}_{\mu'_0}[\mu_t \in K \forall t, \text{ and } \exists \lim_t \frac{\mu_t(\omega)}{\mu_t(\omega')}] > 0$. Given this initial belief μ'_0 , take ℓ from the support of the distribution of $\lim_t \frac{\mu_t(\omega)}{\mu_t(\omega')}$ conditional on the event $\{\mu_t \in K \forall t, \text{ and } \exists \lim_t \frac{\mu_t(\omega)}{\mu_t(\omega')}\}$. Then

$$\mathbb{P}_{\mu'_0} \left[\mu_t \in K \forall t \text{ and } \exists T < \infty \text{ s.t. } \left| \log \frac{\mu_t(\omega)}{\mu_t(\omega')} - \ell \right| \leq \gamma/2 \forall t \geq T \right] > 0. \quad (13)$$

But for any t , if $\mu_t \in K$ and $\left| \log \frac{\mu_t(\omega)}{\mu_t(\omega')} - \ell \right| \leq \gamma/2$, then there exists i such that $\mu_t \in B_{\mu^i}$. Hence, by construction, there is probability at least $\gamma > 0$ that $\left| \log \frac{\mu_{t+1}(\omega)}{\mu_{t+1}(\omega')} - \ell \right| > \gamma/2$. Since the process is Markov, this implies that the event in (13) occurs with zero probability, a contradiction. \square

A.2 Proof of Lemma 2

Consider the random variable $\log \frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)}$, where z is distributed according to P_μ . The corresponding moment-generating function $M(q) := \int \left(\frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)} \right)^q dP_\mu(z)$ is well-defined for $q \in [-q^*, q^*]$ by Assumption 1.2. Note that $M'(0) = \int \log \frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)} dP_\mu(z)$ and that M is convex with $M(0) = 1$.

Part 1. If $\omega \succ_\mu^q \omega'$ for some $q > 0$, then $M(q) < 1 = M(0)$. Thus, convexity of M implies for all $q' \in (0, q)$ that $M(q') \leq \frac{q'}{q}M(q) + (1 - \frac{q'}{q})M(0) < 1$, i.e., $\omega \succ_\mu^{q'} \omega'$. By convexity of M , we also have $M'(0) \leq \frac{1}{q}(M(q) - M(0)) < 0$, whence $\omega \succ_\mu^{\text{KL}} \omega'$.

Part 2. If $\omega \succ_\mu^{\text{KL}} \omega'$, then $M'(0) < 0$. Thus, for all small enough $q > 0$, $M(q) < M(0) = 1$, i.e., $\omega \succ_\mu^q \omega'$. \square

A.3 Proof of Theorem 1

First part: Suppose there exist $q > 0$ and $B \ni \delta_\omega$ such that (6) holds. We can (i) choose B small enough that belief continuity holds at B (by Assumption 1.3), and (ii) assume that $q < 1$ (by Lemma 2). For any initial belief μ_0 with induced probability measure \mathbb{P}_{μ_0} over sequences of beliefs and each $\omega' \neq \omega$, define the stochastic process $\ell_t(\omega') := \left(\frac{\mu_{\min\{t, \tau\}}(\omega')}{\mu_{\min\{t, \tau\}}(\omega)} \right)^q$, where $\tau := \inf\{s : \mu_s \notin B\}$. By (6) and Lemma 3, each $\ell_t(\omega')$ is a nonnegative supermartingale. Thus, by Doob's convergence theorem, there exists an L^∞ -random variable $\ell_\infty(\omega')$ such that $\ell_t(\omega') \rightarrow \ell_\infty(\omega')$ occurs a.s.

To prove that δ_ω is locally stable, it suffices to show the following two claims:

Claim 1: For any initial belief μ_0 , $\mathbb{P}_{\mu_0}[\mu_t \in B \forall t \text{ and } \mu_t \rightarrow \delta_\omega] = \mathbb{P}_{\mu_0}[\mu_t \in B \forall t]$.

Proof of Claim 1. Consider any initial belief μ_0 such that $\mathbb{P}_{\mu_0}[\mu_t \in B, \forall t] > 0$. We show that $\mathbb{P}_{\mu_0}[\mu_t \rightarrow \delta_\omega | \mu_t \in B \forall t] = 1$. Conditional on the event $\{\mu_t \in B, \forall t\}$, we have $\tau = \infty$, so the fact that $\ell_t(\omega') \rightarrow \ell_\infty(\omega')$ a.s. implies that each $\frac{\mu_t(\omega')}{\mu_t(\omega)}$ converges a.s. to a finite value. Suppose for a contradiction that for some $\omega' \neq \omega$, $\mathbb{P}_{\mu_0}[\lim_t \frac{\mu_t(\omega')}{\mu_t(\omega)} > 0 | \tau = \infty] > 0$. Then there exists a compact $K \subseteq B$ such that $\mu(\omega'), \mu(\omega) > 0$ for all $\mu \in K$ and $\mathbb{P}_{\mu_0}[\exists T \text{ s.t. } \mu_t \in K \forall t \geq T \text{ and } \exists \lim_t \frac{\mu_t(\omega')}{\mu_t(\omega)} | \tau = \infty] > 0$. But this contradicts Lemma 10, because for any $\mu \in B \setminus \{\delta_\omega\}$, (6) yields some $z \in \text{supp}P_\mu$ with $\hat{p}_\mu(z|\omega) \neq \hat{p}_\mu(z|\omega')$. Hence, we have $\mathbb{P}_{\mu_0}[\lim_t \frac{\mu_t(\omega')}{\mu_t(\omega)} = 0 | \tau = \infty] = 1$ for all $\omega' \neq \omega$. Thus, $\mathbb{P}_{\mu_0}[\mu_t \rightarrow \delta_\omega | \tau = \infty] = 1$, as claimed. \square

Claim 2: For any $\gamma > 0$, there exists a neighborhood $B' \subseteq B$ of δ_ω such that $\mathbb{P}_{\mu_0}[\mu_t \in B, \forall t] \geq \gamma$ for any initial belief $\mu_0 \in B'$.

Proof of Claim 2. Fix any $\gamma > 0$. Pick $\varepsilon_+ > 0$ such that $\{\mu \in \Delta(\Omega) : \sum_{\omega' \neq \omega} \left(\frac{\mu(\omega')}{\mu(\omega)}\right)^q < \varepsilon_+\} \subseteq B$. Pick $\varepsilon_- > 0$ such that $\frac{\varepsilon_-}{\varepsilon_+} \leq 1 - \gamma$. For any $\mu_0 \in B' := \{\mu \in \Delta(\Omega) : \sum_{\omega' \neq \omega} \left(\frac{\mu(\omega')}{\mu(\omega)}\right)^q < \varepsilon_-\}$, we have

$$\mathbb{P}_{\mu_0}[\exists t, \mu_t \notin B] \leq \mathbb{P}_{\mu_0}\left[\sum_{\omega' \neq \omega} \ell_\infty(\omega') \geq \varepsilon_+\right] \leq \mathbb{E}_{\mu_0}\left[\sum_{\omega' \neq \omega} \ell_\infty(\omega')\right]/\varepsilon_+ \leq \frac{\varepsilon_-}{\varepsilon_+},$$

where the second inequality uses Markov's inequality and the third follows from Fatou's lemma and the fact that each $\ell_t(\omega')$ is a nonnegative supermartingale. Thus, $\mathbb{P}_{\mu_0}[\mu_t \in B, \forall t] \geq \gamma$. \square

Second part: Suppose there exist $q > 0$ and a neighborhood $B \ni \delta_\omega$ such that (7) holds for some $\omega' \neq \omega$, where we again assume without loss that belief continuity holds at B and $q < 1$. Up to restricting to a subneighborhood of B , we can assume that there exists $\varepsilon > 0$ such that $\mu(\omega) > \varepsilon$ for all $\mu \in B$. Fix any initial belief $\mu_0 \in B \setminus \{\delta_\omega\}$. Let $\tau := \inf\{s : \mu_s \notin B\}$. To prove instability of δ_ω , it suffices to show that $\mathbb{P}_{\mu_0}[\tau < \infty] = 1$. Consider the process $\ell_t := \left(\frac{\mu_{\min\{t, \tau\}}(\omega)}{\mu_{\min\{t, \tau\}}(\omega')}\right)^q$, which is a non-negative supermartingale by (7) and Lemma 3. Hence, Doob's convergence theorem yields an L^∞ -random variable ℓ_∞ such that $\ell_t \rightarrow \ell_\infty$ a.s.

Suppose for a contradiction that with positive probability, we have $\tau = \infty$. Conditional on $\tau = \infty$, we have $\left(\frac{\mu_t(\omega)}{\mu_t(\omega')}\right)^q = \ell_t$ for all t . Thus, conditional on $\tau = \infty$, $\frac{\mu_t(\omega)}{\mu_t(\omega')}$ converges a.s. to an L^∞ random limit $\lim_t \frac{\mu_t(\omega)}{\mu_t(\omega')}$, which must be strictly positive since $\mu(\omega) > \varepsilon$ for all $\mu \in B$. Hence, there exists some compact set $K \subseteq B \setminus \{\delta_\omega\}$ such that $\mu(\omega), \mu(\omega') > 0$ for all $\mu \in K$ and $\mathbb{P}_{\mu_0}[\exists T \text{ s.t. } \mu_t \in K \forall t \geq T \text{ and } \exists \lim_t \frac{\mu_t(\omega)}{\mu_t(\omega')} \mid \tau = \infty] > 0$. But this contradicts Lemma 10, because (7) implies that for each $\mu \in K$, there exists $z \in \text{supp}P_\mu$ with $\hat{p}_\mu(z_\mu|\omega) \neq \hat{p}_\mu(z_\mu|\omega')$. \square

A.4 Proof of Theorem 2

This result is a special case of Theorem 3 in Appendix B.

A.5 Proof of Proposition 1

We call $K \subseteq \Delta(\Omega)$ an **unstable set** if there exists a neighborhood B of K such that $\mathbb{P}_{\mu_0}[\exists t, \mu_t \notin B] = 1$ for every initial belief $\mu_0 \in B \setminus K$. We call $K \subseteq \Delta(\Omega)$ **transient** if $\mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \mu_t \notin K] = 1$ for any initial belief $\mu_0 \in K$. We invoke the following lemma, which we prove in Appendix A.5.1.

Lemma 11. *Suppose that belief continuity holds. Consider $\Omega = \{\omega_1, \dots, \omega_N\}$ and suppose that*

- (i) δ_{ω_1} satisfies the condition for local stability in Theorem 1;
- (ii) $\Delta(\{\omega_2, \dots, \omega_N\})$ is unstable;
- (iii) for any mixed $\mu \in \Delta(\Omega)$, there is $z \in \text{supp}(P_\mu)$ with $\hat{p}_\mu(z|\omega_1) > \hat{p}_\mu(z|\omega_n)$ for all $n \neq 1$.

Then δ_{ω_1} is globally stable.

To prove Proposition 1, we verify the assumptions in Lemma 11. Assumptions (i) and (iii) in Lemma 11 follow from assumptions (i) and (ii) in Proposition 1 applied with $n = 1$. Thus, it remains to show that $\Delta(\{\omega_2, \dots, \omega_N\})$ is unstable. We prove inductively that $\Delta(\{\omega_{N-m}, \dots, \omega_N\})$ is unstable for all $m = 0, \dots, N-2$. For $m = 0$, this holds since δ_{ω_N} is unstable by assumption (i) in Proposition 1 and Theorem 1. For the inductive step, we prove the following lemma; this completes the proof, because assumptions (i)–(ii) in Proposition 1 imply assumptions (i)–(iii) in the lemma.

Lemma 12. *Fix any $n \in \{2, \dots, N-1\}$. Suppose that the set $\Delta(\{\omega_{n+1}, \dots, \omega_N\})$ is unstable and belief continuity holds at some neighborhood of this set. Assume that (i) there exist $q > 0$ and a neighborhood $B_n \ni \delta_{\omega_n}$ such that $\omega_n \succ_{\mu}^q \omega_k$ for all $k > n$ and $\mu \in B_n \setminus \{\delta_{\omega_n}\}$; (ii) δ_{ω_n} is unstable; and (iii) for each mixed belief $\mu \in \Delta(\{\omega_n, \dots, \omega_N\})$, there exists $z \in \text{supp}(P_{\mu})$ such that $\hat{p}_{\mu}(z|\omega_n) > \hat{p}_{\mu}(z|\omega_k)$ for all $k > n$. Then $\Delta(\{\omega_n, \dots, \omega_N\})$ is unstable.*

Proof. Note first that since $\Delta(\{\omega_{n+1}, \dots, \omega_N\})$ is unstable, there exists $\varepsilon_{n+1} > 0$ such that $\Delta_{n+1} := \{\mu \in \Delta(\Omega) : \mu(\{\omega_{n+1}, \dots, \omega_N\}) \geq 1 - \varepsilon_{n+1}\}$ is transient. Moreover, we can assume that B_n in assumption (i) takes the form $\{\mu \in \Delta(\Omega) : \mu(\omega_n) > 1 - \kappa\}$ for some $\kappa > 0$, where, by choosing κ sufficiently small, assumption (ii) ensures that B_n is transient.

We claim that we can choose $\varepsilon > 0$, $\gamma \in (0, 1)$, and $\varepsilon_n \in (0, \varepsilon_{n+1})$ such that, defining

$$\Delta_n := \{\mu \in \Delta(\Omega) : \mu(\{\omega_n, \dots, \omega_N\}) \geq 1 - \varepsilon_n\}, \quad B'_n := \{\mu \in \Delta_n : \sum_{k>n} \left(\frac{\mu(\omega_k)}{\mu(\omega_n)} \right)^q \leq \varepsilon\},$$

the following three properties are satisfied:

$$B'_n \subseteq B_n \tag{14}$$

$$\forall \mu \in \Delta_n \setminus (\Delta_{n+1} \cup B'_n), \exists Z_{\mu} \subseteq Z \text{ with } P_{\mu}(Z_{\mu}) \geq \gamma \text{ and } \inf_{z \in Z_{\mu}} \frac{\hat{p}_{\mu}(z|\omega_n)}{\hat{p}_{\mu}(z|\omega_k)} - 1 \geq \gamma \text{ for all } k > n \tag{15}$$

$$\frac{\varepsilon_{n+1}}{\varepsilon_{n+1} - \varepsilon_n} \leq 1 + \gamma. \tag{16}$$

Indeed, first pick $\varepsilon > 0$ sufficiently small that $\mu(\omega_n) \geq 1 - \kappa/2$ holds for every $\mu \in \Delta(\{\omega_n, \dots, \omega_N\})$ with $\sum_{k>n} \left(\frac{\mu(\omega_k)}{\mu(\omega_n)} \right)^q \leq \varepsilon$. Then (14) is satisfied for all sufficiently small $\varepsilon_n \in (0, \varepsilon_{n+1})$. To show (15), note that by assumption (iii) and continuity of signal densities in z , for all $\mu \in \Delta(\{\omega_n, \dots, \omega_N\}) \setminus \{\delta_{\omega_n}, \dots, \delta_{\omega_N}\}$, there exists $Z_{\mu} \subseteq Z$ with $P_{\mu}(Z_{\mu}) > 0$ and $\inf_{z \in Z_{\mu}} \frac{\hat{p}_{\mu}(z|\omega_n)}{\hat{p}_{\mu}(z|\omega_k)} - 1 > 0$ for all $k > n$. By belief continuity, for each such μ , there exists an open neighborhood $B_{\mu} \ni \mu$ such that $\inf_{\mu' \in B_{\mu}} P_{\mu'}(Z_{\mu}) > 0$ and $\inf_{z \in Z_{\mu}, \mu' \in B_{\mu}} \frac{\hat{p}_{\mu'}(z|\omega_n)}{\hat{p}_{\mu'}(z|\omega_k)} - 1 > 0$ for all $k > n$. Moreover, given $\varepsilon > 0$, but independent of the choice of ε_n , $\mu(\omega_n), \dots, \mu(\omega_N)$ are bounded away from 1 for all $\mu \in \Delta(\{\omega_n, \dots, \omega_N\}) \setminus (\Delta_{n+1} \cup B'_n)$. Thus, $\Delta(\{\omega_n, \dots, \omega_N\}) \setminus (\Delta_{n+1} \cup B'_n)$ is contained in some compact set $K \subset \Delta(\{\omega_n, \dots, \omega_N\}) \setminus \{\delta_{\omega_n}, \dots, \delta_{\omega_N}\}$. Hence, by taking a finite subcover $(B_{\mu_i})_{i=1, \dots, I}$ of K , there is $\gamma \in (0, 1)$ such that $\inf_{\mu' \in B_{\mu_i}} P_{\mu'}(Z_{\mu_i}) \geq \gamma$ and $\inf_{z \in Z_{\mu_i}, \mu' \in B_{\mu_i}} \frac{\hat{p}_{\mu'}(z|\omega_n)}{\hat{p}_{\mu'}(z|\omega_k)} - 1 \geq \gamma$ for all

$k > n$ and $i \in 1, \dots, I$. For all small enough ε_n , we can then ensure that (15) and (16) hold, where the former is guaranteed by requiring $\Delta_n \setminus (\Delta_{n+1} \cup B'_n)$ to be included in the cover $(B_{\mu_i})_{i=1, \dots, I}$.

For ε , γ , and ε_n as chosen above, we establish the following two claims:

Claim 1: There exists $T \in \mathbb{N}$ such that $\mathbb{P}_{\mu_0}[\exists t \leq T \text{ s.t. } \mu_t \in B'_n \cup \Delta_n^c] \geq \gamma^T$ for every initial belief $\mu_0 \in \Delta_n \setminus \Delta_{n+1}$.

Proof of Claim 1. Observe first that $\frac{\mu_0(\omega_{n+1})}{\mu_0(\omega_n)}, \dots, \frac{\mu_0(\omega_N)}{\mu_0(\omega_n)}$ are uniformly bounded from above for all $\mu_0 \in \Delta_n \setminus \Delta_{n+1}$, as $\mu_0(\omega_n) \geq \varepsilon_{n+1} - \varepsilon_n > 0$. Thus, there exists T with $\sum_{k>n} \left(\frac{\mu_0(\omega_k)}{\mu_0(\omega_n)} (1 + \gamma) \right)^{-T} \leq \varepsilon$.

Starting with any initial belief $\mu_0 \in \Delta_n \setminus \Delta_{n+1}$, we recursively construct sequences of signal realizations $z_0, z_1, \dots, z_{T'}$ with $T' \leq T - 1$ and corresponding updated beliefs $\mu_1, \mu_2, \dots, \mu_{T'+1}$. Suppose we have constructed z_0, \dots, z_{t-1} for some $t \in \{0, \dots, T\}$. We distinguish two cases:

(a) If $\mu_t \in B'_n \cup \Delta_n^c$, set $T' = t - 1$ and terminate the construction of the signal sequence.

(b) Suppose $\mu_t \in \Delta_n \setminus (\Delta_{n+1} \cup B'_n)$. Then by (15), we can pick any signal $z_t \in Z_{\mu_t}$, which satisfies $\frac{\hat{p}_{\mu_t}(z_t|\omega_n)}{\hat{p}_{\mu_t}(z_t|\omega_k)} - 1 \geq \gamma$ for all $k > n$. We claim that the updated belief μ_{t+1} satisfies $\mu_{t+1}(\{\omega_{n+1}, \dots, \omega_N\}) \leq \mu_t(\{\omega_{n+1}, \dots, \omega_N\})$, so $\mu_{t+1} \notin \Delta_{n+1}$. Indeed, suppose to the contrary that $\mu_{t+1}(\{\omega_{n+1}, \dots, \omega_N\}) > \mu_t(\{\omega_{n+1}, \dots, \omega_N\})$. By choice of z_t , we have $\frac{\mu_{t+1}(\omega_k)}{\mu_{t+1}(\omega_n)} \geq \frac{\mu_t(\omega_k)}{\mu_t(\omega_n)} (1 + \gamma)$ for each $k > n$. Thus, $\frac{\mu_{t+1}(\omega_n)}{\mu_t(\omega_n)} \geq \max_{k>n} \frac{\mu_{t+1}(\omega_k)}{\mu_t(\omega_k)} (1 + \gamma) \geq \frac{\mu_{t+1}(\{\omega_{n+1}, \dots, \omega_N\})}{\mu_t(\{\omega_{n+1}, \dots, \omega_N\})} (1 + \gamma) > 1 + \gamma$. At the same time,

$$\frac{\mu_{t+1}(\omega_n)}{\mu_t(\omega_n)} \leq \frac{1 - \mu_{t+1}(\{\omega_{n+1}, \dots, \omega_N\})}{1 - \mu_t(\{\omega_{n+1}, \dots, \omega_N\}) - \varepsilon_n} < \frac{1 - \mu_t(\{\omega_{n+1}, \dots, \omega_N\})}{1 - \mu_t(\{\omega_{n+1}, \dots, \omega_N\}) - \varepsilon_n} \leq \frac{\varepsilon_{n+1}}{\varepsilon_{n+1} - \varepsilon_n}$$

where the first inequality holds because $\mu_t \in \Delta_n$ and the third because $\mu_t \notin \Delta_{n+1}$. Thus, $\frac{\mu_{t+1}(\omega_n)}{\mu_t(\omega_n)} > 1 + \gamma$, which contradicts (16).

Note that the construction above ensures that case (a) must occur at the latest at $t = T$, so that $T' \leq T - 1$. Indeed, if (b) holds for all $t < T$, then $\mu_T \in B'_n$, as $\sum_{k>n} \left(\frac{\mu_T(\omega_k)}{\mu_T(\omega_n)} \right)^q \leq \sum_{k>n} \left(\frac{\mu_0(\omega_k)}{\mu_0(\omega_n)} (1 + \gamma) \right)^{-T} \leq \varepsilon$ by (b) and the choice of T . This proves Claim 1, as by construction and (15), signal realizations $(z_0, \dots, z_{T'})$ of the above form occur with probability at least $\gamma^{T'+1}$. \square

Claim 2: Let $\tau := \inf\{t : \mu_t \notin B'_n\}$. There exists $\xi \in [0, 1)$ such that $\mathbb{P}_{\mu_0}[\tau < \infty] = 1$ and $\mathbb{P}_{\mu_0}[\mu_\tau \in \Delta_n \setminus B'_n] \leq \xi$ for every initial belief $\mu_0 \in B'_n$.

Proof of Claim 2. Note that $\mathbb{P}_{\mu_0}[\tau < \infty] = 1$ is immediate from (14) and the fact that B_n is transient. To show the existence of ξ , define $\ell_t := \sum_{k>n} \left(\frac{\mu_{\min\{t, \tau\}}(\omega_k)}{\mu_{\min\{t, \tau\}}(\omega_n)} \right)^q$. By (14) and assumption (ii), ℓ_t is a nonnegative supermartingale, and in particular $\mathbb{E}_{\mu_0}[\ell_1] < \ell_0 \leq \varepsilon$ for every initial belief $\mu_0 \in B'_n$. Since $\mathbb{E}_{\mu_0}[\ell_1]$ is continuous in μ_0 by Lemma 9 and B'_n is compact, there exists $\xi \in [0, 1)$ such that $\mathbb{E}_{\mu_0}[\ell_1] \leq \xi \varepsilon$ holds for every initial belief $\mu_0 \in B'_n$. Hence,

$$\mathbb{P}_{\mu_0}[\mu_\tau \in \Delta_n \setminus B'_n] \varepsilon + \mathbb{P}_{\mu_0}[\mu_\tau \notin \Delta_n \setminus B'_n] \cdot 0 \leq \mathbb{E}_{\mu_0}[\ell_\tau] \leq \mathbb{E}_{\mu_0}[\ell_1] \leq \xi \varepsilon,$$

where the first inequality holds by definition of B'_n . Thus, $\mathbb{P}_{\mu_0}[\mu_\tau \in \Delta_n \setminus B'_n] \leq \xi$. \square

To complete the proof of Lemma 12, for each initial belief μ_0 , define $g(\mu_0) := \mathbb{P}_{\mu_0}[\mu_t \in \Delta_n \forall t]$. We verify that $\sup_{\mu_0 \in \Delta_n} g(\mu_0) = 0$. First, take any $\mu_0 \in \Delta_n \cap \Delta_{n+1}$ and set $\tau' := \inf\{t : \mu_t \notin \Delta_{n+1}\}$, which satisfies $\mathbb{P}_{\mu_0}[\tau' < \infty] = 1$ since Δ_{n+1} is transient. By the Markov property,

$$g(\mu_0) = \mathbb{P}_{\mu_0}[\mu_{\tau'} \in \Delta_n] \mathbb{E}_{\mu_0}[g(\mu_{\tau'}) | \mu_{\tau'} \in \Delta_n] + \mathbb{P}_{\mu_0}[\mu_{\tau'} \notin \Delta_n] \cdot 0 \leq \sup_{\mu \in \Delta_n \setminus \Delta_{n+1}} g(\mu).$$

This implies that

$$\sup_{\mu_0 \in \Delta_n} g(\mu_0) = \sup_{\mu_0 \in \Delta_n \setminus \Delta_{n+1}} g(\mu_0). \quad (17)$$

Next, take any $\mu_0 \in B'_n$ and define $\tau := \inf\{t : \mu_t \notin B'_n\}$ as in Claim 2. By the Markov property,

$$g(\mu_0) = \mathbb{P}_{\mu_0}[\mu_\tau \in \Delta_n] \mathbb{E}_{\mu_0}[g(\mu_\tau) | \mu_\tau \in \Delta_n] \leq \xi \sup_{\mu \in \Delta_n} g(\mu) = \xi \sup_{\mu \in \Delta_n \setminus \Delta_{n+1}} g(\mu)$$

where the inequality holds by Claim 2 and the equality by (17). Thus,

$$\sup_{\mu \in B_n} g(\mu) \leq \xi \sup_{\mu \in \Delta_n \setminus \Delta_{n+1}} g(\mu). \quad (18)$$

Last, take $\mu_0 \in \Delta_n \setminus \Delta_{n+1}$ and let $\tau'' := \inf\{\min\{t : \mu_t \in \Delta_n^c \cup B'_n\}, T+1\}$. By the Markov property,

$$\begin{aligned} g(\mu_0) &= \mathbb{P}_{\mu_0}[\tau'' \leq T] \mathbb{E}_{\mu_0}[g(\mu_{\tau''}) | \tau'' \leq T] + \mathbb{P}_{\mu_0}[\tau'' > T] \mathbb{E}_{\mu_0}[g(\mu_{\tau''}) | \tau'' > T] \\ &\leq \mathbb{P}_{\mu_0}[\tau'' \leq T] \sup_{\mu \in B_n} g(\mu) + \mathbb{P}_{\mu_0}[\tau'' > T] \sup_{\mu \in \Delta_n} g(\mu) \\ &\leq \gamma^T \sup_{\mu \in B_n} g(\mu) + (1 - \gamma^T) \sup_{\mu \in \Delta_n} g(\mu) \leq (\gamma^T \xi + 1 - \gamma^T) \sup_{\mu \in \Delta_n \setminus \Delta_{n+1}} g(\mu), \end{aligned}$$

where the second inequality follows from Claim 1 and the fact that $\sup_{\mu \in B_n} g(\mu) \leq \sup_{\mu \in \Delta_n} g(\mu)$ by (18), and the final inequality holds by (17)–(18). Thus, $\sup_{\mu \in \Delta_n \setminus \Delta_{n+1}} g(\mu) = 0$ and the desired conclusion follows from (17). \square

A.5.1 Proof of Lemma 11

Fix any $\gamma \in (0, 1)$. Given assumption (i), Claims 1 and 2 in the proof of Theorem 1 ensure that there exist neighborhoods $B_1 \supseteq B'_1 \ni \delta_{\omega_1}$ such that

$$\mathbb{P}_{\mu_0}[\mu_t \in B_1 \forall t] = \mathbb{P}_{\mu_0}[\mu_t \in B'_1 \forall t, \text{ and } \mu_t \rightarrow \delta_{\omega_1}] \geq \gamma \text{ for all initial beliefs } \mu_0 \in B'_1. \quad (19)$$

By assumption (ii), $\Delta(\{\omega_2, \dots, \omega_N\})$ admits a neighborhood Δ_2 such that $\mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \mu_t \notin \Delta_2] = 1$ for all initial beliefs $\mu_0 \in \Delta_2 \setminus \Delta(\{\omega_2, \dots, \omega_N\})$. Since initial beliefs have full support, we equivalently have that $\mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \mu_t \notin \Delta_2] = 1$ for all initial beliefs $\mu_0 \in \Delta_2$. Thus, Δ_2 is transient.

Observe that there exist $T \in \mathbb{N}$ and $\eta > 0$ such that, for every initial belief $\mu_0 \notin \Delta_2$,

$$\mathbb{P}_{\mu_0}[\exists t \leq T \text{ s.t. } \mu_t \in B'_1] \geq \eta \quad (20)$$

Indeed, pick $L > 1$ large enough that (i) $\mu \in B'_1$ for all μ with $\log \frac{\mu(\omega_1)}{\mu(\omega_n)} \geq L$ for each $n > 1$, and (ii) $\log \frac{\mu(\omega_1)}{\mu(\omega_n)} \geq 1/L$ for all $\mu \notin \Delta_2$ and $n > 1$. By continuity of $p_\mu(z)$, $\frac{\hat{p}_\mu(z|\omega_1)}{\hat{p}_\mu(z|\omega_n)}$ in (z, μ) and assumption (iii), there exists $\varepsilon > 0$ such that for all μ in the compact set $\{\mu \in \Delta(\Omega) : L \geq \min_{n>1} \log \frac{\mu(\omega_1)}{\mu(\omega_n)} \geq 1/L\}$, there is $Z_\mu \subseteq Z$ such that $P_\mu(Z_\mu) > \varepsilon$ and $\log \frac{\hat{p}_\mu(z|\omega_1)}{\hat{p}_\mu(z|\omega_n)} > \varepsilon$ for all $n \neq 1$ and $z \in Z_\mu$. Starting from any initial belief $\mu_0 \notin \Delta_2$, consider any realization of signals (z_t) and corresponding beliefs (μ_t) such that $z_t \in Z_{\mu_t}$. This ensures $\log \frac{\mu_t(\omega_1)}{\mu_t(\omega_n)} \geq 1/L + t\varepsilon$ for each $n > 1$ and t . Along this sequence, $\mu_{t'} \in B'_1$ for some $t' \leq \frac{L-1/L}{\varepsilon}$. Thus, claim (20) holds by choosing $T \geq \frac{L-1/L}{\varepsilon}$ and $\eta = \varepsilon^T$.

For each initial belief μ_0 , define $h(\mu_0) := \mathbb{P}_{\mu_0}[\mu_t \rightarrow \delta_{\omega_1}]$. To show global stability of δ_{ω_1} , we will prove that $\inf_{\mu \in \Delta^\circ(\Omega)} h(\mu) = 1$. Note first that for any initial belief μ_0 , $\tau := \inf\{t : \mu_t \notin \Delta_2\}$ satisfies $\mathbb{P}_{\mu_0}[\tau < \infty] = 1$ as Δ_2 is transient. Thus, by the Markov property of μ_t , we have $h(\mu_0) = \mathbb{E}_{\mu_0}[h(\mu_\tau)] \geq \inf_{\mu \in \Delta^\circ(\Omega) \setminus \Delta_2} h(\mu)$, whence

$$\inf_{\mu \in \Delta^\circ(\Omega)} h(\mu) = \inf_{\mu \in \Delta^\circ(\Omega) \setminus \Delta_2} h(\mu). \quad (21)$$

Next, take any initial belief $\mu_0 \in B'_1$ and $\tau' := \inf\{t : \mu_t \notin B_1\}$. By the Markov property and (19),

$$\begin{aligned} h(\mu_0) &= \mathbb{P}_{\mu_0}[\tau' = \infty] \mathbb{P}_{\mu_0}[\mu_t \rightarrow \delta_{\omega_1} | \tau' = \infty] + \mathbb{P}_{\mu_0}[\tau' < \infty] \mathbb{E}_{\mu_0}[h(\mu_{\tau'}) | \tau' < \infty] \\ &= \mathbb{P}_{\mu_0}[\tau' = \infty] + \mathbb{P}_{\mu_0}[\tau' < \infty] \mathbb{E}_{\mu_0}[h(\mu_{\tau'}) | \tau' < \infty] \geq \gamma + (1 - \gamma) \inf_{\mu \in \Delta^\circ(\Omega)} h(\mu). \end{aligned}$$

Combining this with (21) yields

$$\inf_{\mu \in B'_1} h(\mu) \geq \gamma + (1 - \gamma) \inf_{\mu \in \Delta^\circ(\Omega) \setminus \Delta_2} h(\mu). \quad (22)$$

Finally, consider any initial belief $\mu_0 \notin \Delta_2$ and let $\tau'' := \min\{\inf\{t : \mu_t \in B'_1\}, T + 1\}$. Then, by the Markov property and (20)-(22), we have

$$\begin{aligned} h(\mu_0) &= \mathbb{P}_{\mu_0}[\tau'' \leq T] \mathbb{E}_{\mu_0}[h(\mu_{\tau''}) | \tau'' \leq T] + \mathbb{P}_{\mu_0}[\tau'' > T] \mathbb{E}_{\mu_0}[h(\mu_{\tau''}) | \tau'' > T] \\ &\geq \mathbb{P}_{\mu_0}[\tau'' \leq T] \inf_{\mu \in B'_1} h(\mu) + \mathbb{P}_{\mu_0}[\tau'' > T] \inf_{\mu \in \Delta^\circ(\Omega)} h(\mu) \\ &\geq \eta \inf_{\mu \in B'_1} h(\mu) + (1 - \eta) \inf_{\mu \in \Delta^\circ(\Omega)} h(\mu) \geq \eta\gamma + (1 - \eta\gamma) \inf_{\mu \in \Delta^\circ(\Omega) \setminus \Delta_2} h(\mu). \end{aligned}$$

This holds for all $\mu_0 \notin \Delta_2$, so $\inf_{\mu \in \Delta^\circ(\Omega) \setminus \Delta_2} h(\mu) = 1$. By (21), $\inf_{\mu \in \Delta^\circ(\Omega)} h(\mu) = 1$. \square

B General states

We provide local and global stability conditions for infinite state spaces, by extending the martingale approach in the main text. Assume Ω is a compact metric space and endow $\Delta(\Omega)$ with the Prokhorov metric d . In addition to Assumption 1, we impose the following standard assumption, which is automatically satisfied if Ω is finite:

Assumption 2 (Continuity in states). For each $\mu \in \Delta(\Omega)$ and $z \in Z$, $\hat{p}_\mu(z|\omega)$ is continuous in ω .

As in Section 2, given any full-support initial belief μ_0 , the belief process μ_t is induced by (P_μ) and $(\hat{P}_\mu(\cdot|\omega))$ using Bayes' rule. In particular, after signal z_t is drawn according to p_{μ_t} , μ_t is updated to μ_{t+1} by setting $\mu_{t+1}(\Omega') = \frac{\int_{\Omega'} \hat{p}_{\mu_t}(z_t|\omega) d\mu_t(\omega)}{\int_{\Omega} \hat{p}_{\mu_t}(z_t|\omega) d\mu_t(\omega)}$ for each measurable $\Omega' \subseteq \Omega$.

B.1 Global iterated dominance

For global stability, we extend Theorem 2. For each nonempty $\Omega' \subseteq \Omega$, let

$$S(\Omega') := \{\omega \in \overline{\Omega'} : \exists \omega' \in \overline{\Omega'} \text{ s.t. } \omega' \succ_{\mu}^{\text{KL}} \omega \text{ for all } \mu \in \Delta(\overline{\Omega'})\},$$

where $\overline{\Omega'}$ denotes the closure of Ω' in Ω . Under belief continuity, $S(\Omega')$ is nonempty and compact (Lemma 14). Thus, $S^\infty(\Omega') := \bigcap_{k \in \mathbb{N}} S^k(\Omega')$ is nonempty and compact by Cantor's intersection theorem. The following result shows that Theorem 2 remains true unchanged:

Theorem 3. *Assume belief continuity holds. Then $\Delta(S^\infty(\Omega))$ is globally stable.*

We prove Theorem 3 in Appendix B.4. All proofs in Appendix B rely on Lemma 15, which extends our supermartingale construction via q -dominance to infinite state spaces.

B.2 Local iterated dominance

To obtain a condition for local stability, we also use the above iterated dominance approach. We consider a set-valued notion of local stability: $M \subseteq \Delta(\Omega)$ is a **locally stable set** if for any $\gamma < 1$, there exists a neighborhood B of M such that $\mathbb{P}_{\mu_0}[\inf_{\nu \in M} d(\mu_t, \nu) \rightarrow 0] \geq \gamma$ from each initial belief $\mu_0 \in B$. We also generalize the notion of strict BeNE to sets of beliefs: For each nonempty measurable $\Omega' \subseteq \Omega$, call $\Delta(\Omega')$ a **strict BeNE set** if for all $\omega \notin \Omega'$, there exists $\omega' \in \overline{\Omega'}$ such that

$$\omega' \succ_{\mu}^{\text{KL}} \omega \text{ for all } \mu \in \Delta(\overline{\Omega'}).$$

Note that if $\Omega' = \{\omega'\}$ is a singleton, this definition reduces to $\delta_{\omega'}$ being a strict BeNE. We prove the following result in Appendix B.5:

Theorem 4. *Suppose Ω' is open and belief continuity holds at some neighborhood of $\Delta(\Omega')$. If $\Delta(\Omega')$ is a strict BeNE set, then $\Delta(S^k(\Omega'))$ is locally stable for all $k = 0, 1, \dots, \infty$.*

Theorem 4 implies Corollary 1 when Ω is finite. However, a strict BeNE δ_ω need not be locally stable under general Ω , as $\{\omega\}$ need not be open.

Similar to the application of Theorem 3 in Section 4.1, Theorem 4 is straightforward to apply under one-dimensional states, because in this case local iterated dominance again corresponds to iterated application of the map m :

Example 3. Consider the environment in Section 4.1. Proposition 5 (Online Appendix C.1) generalizes Proposition 2 by showing that if $\Omega' \subseteq \Omega$ is an open interval such that $m(\overline{\Omega}') \subseteq \Omega'$, then $S^k(\Omega') = m^k(\overline{\Omega}')$ for all $k = 0, 1, \dots, \infty$. For any such Ω' , the fact that $S(\Omega') = m(\overline{\Omega}') \subseteq \Omega'$ implies that Ω' is a strict BeNE set. Thus, by Theorem 4, $\Delta(m^\infty(\overline{\Omega}'))$ is locally stable.

For example, consider any BeNE $\delta_{\hat{\omega}}$. Then if m is continuously differentiable near $\hat{\omega}$ with $|m'(\hat{\omega})| < 1$, this implies that $\delta_{\hat{\omega}}$ is locally stable, because for some small enough open interval $\Omega' \ni \hat{\omega}$, we have $m(\overline{\Omega}') \subseteq \Omega'$ and $m^\infty(\overline{\Omega}') = \{\hat{\omega}\}$. \blacktriangle

B.3 Preliminary results for the proofs of Theorems 3–4

Lemma 13. *Pick q^* as in Assumption 1.2. For each μ and $q \in (0, q^*]$, $\int \left(\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\tilde{\omega})}\right)^q p_\mu(z) d\nu(z)$ is continuous in ω and $\tilde{\omega}$.*

Proof. For all z such that $p_\mu(z) > 0$, $\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\tilde{\omega})}$ is continuous in $\omega, \tilde{\omega}$ by Assumptions 1.1 and 2. Thus, $\int \left(\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\tilde{\omega})}\right)^q p_\mu(z) d\nu(z)$ is continuous in ω and $\tilde{\omega}$ by the dominated convergence theorem, as $\left(\frac{\hat{p}_\mu(\cdot|\omega)}{\hat{p}_\mu(\cdot|\tilde{\omega})}\right)^q p_\mu(\cdot)$ is dominated by the ν -integrable function $h(\cdot)$ (Assumption 1.2). \square

Lemma 14. *Take any nonempty $\Omega' \subseteq \Omega$ such that belief continuity holds at $\Delta(\overline{\Omega}')$. Then $S(\Omega')$ is nonempty and compact.*

Proof. Take any $\omega \in \overline{\Omega}' \setminus S(\Omega')$. Then, by definition of $S(\Omega')$, there is $\phi(\omega) \in \overline{\Omega}'$ such that $\int \log \frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\phi(\omega))} dP_\mu(z) < 0$ for each $\mu \in \Delta(\overline{\Omega}')$. Thus, for each $\mu \in \Delta(\overline{\Omega}')$, Lemma 2 yields $q_\mu \in (0, q^*]$ such that, for all $q \in (0, q_\mu]$,

$$\int \left(\frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\phi(\omega))}\right)^q p_\mu(z) d\nu(z) < 1.$$

By belief continuity, the LHS is continuous in μ at $\Delta(\overline{\Omega}')$ (Lemma 9). Thus, $\int \left(\frac{\hat{p}_{\mu'}(z|\omega)}{\hat{p}_{\mu'}(z|\phi(\omega))}\right)^{q_\mu} p_{\mu'}(z) d\nu(z) < 1$ for all μ' in some neighborhood B_μ of μ . Since $\Delta(\overline{\Omega}')$ is compact, by taking a finite subcover of $\{B_\mu : \mu \in \Delta(\overline{\Omega}')$, we can choose $q_\mu =: q$ to be independent of μ . Thus, at $\omega' = \omega$, we have

$$\max_{\mu \in \Delta(\overline{\Omega}')} \int \left(\frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\phi(\omega))}\right)^q p_\mu(z) d\nu(z) < 1. \quad (23)$$

Since the LHS of (23) is continuous in ω' by Lemma 13 and the maximum theorem, there is a neighborhood $B_\omega \ni \omega$ such that for all $\omega' \in B_\omega \cap \overline{\Omega}'$, $\max_{\mu \in \Delta(\overline{\Omega}')} \int \left(\frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\phi(\omega))}\right)^q p_\mu(z) d\nu(z) < 1$. By

Lemma 2, this implies $\phi(\omega) \succ_{\mu}^{\text{KL}} \omega'$ for all $\mu \in \Delta(\overline{\Omega}')$ and $\omega' \in B_{\omega} \cap \overline{\Omega}'$. Thus, $\overline{\Omega}' \setminus S(\Omega')$ is open in $\overline{\Omega}'$, which implies that $S(\Omega')$ is closed in $\overline{\Omega}'$ and hence compact.

Next, suppose that $S(\Omega')$ is empty. Then the above observation shows that for each $\omega \in \overline{\Omega}'$, there exists $\phi(\omega) \in \overline{\Omega}'$ and a neighborhood B_{ω} of ω such that $\phi(\omega) \succ_{\mu}^{\text{KL}} \omega'$ for all $\mu \in \Delta(\overline{\Omega}')$ and $\omega' \in B_{\omega} \cap \overline{\Omega}'$. By compactness of $\overline{\Omega}'$, $\{B_{\omega} : \omega \in \overline{\Omega}'\}$ admits a finite subcover $\{B_{\omega_i} : i = 1, \dots, I\}$. Then for each $i \in \{1, \dots, I\}$, there exists $j \in \{1, \dots, I\}$ such that $\phi(\omega_j) \succ_{\mu}^{\text{KL}} \phi(\omega_i)$ for all $\mu \in \Delta(\overline{\Omega}')$. By transitivity of KL dominance, this yields $i \in \{1, \dots, I\}$ such that $\phi(\omega_i) \succ_{\mu}^{\text{KL}} \phi(\omega_i)$, which is impossible. Thus, $S(\Omega')$ is nonempty. \square

The following lemma extends the supermartingale construction via q -dominance to general Ω . For any $M \subseteq \Delta(\Omega)$ and $\varepsilon > 0$, let $B_{\varepsilon}(M) := \{\nu \in \Delta(\Omega) : \inf_{\mu \in M} d(\mu, \nu) < \varepsilon\}$. Note that (24) below ensures that each $\ell_t^i := \left(\frac{\mu_{\min\{t, \tau\}}(A_i)}{\mu_{\min\{t, \tau\}}(A'_i)}\right)^{q_i}$ with $\tau := \inf\{s : \mu_s \notin B_{\varepsilon}(D)\}$ is a nonnegative supermartingale. Moreover, the lemma shows that $\ell_t^i \rightarrow 0$ a.s. conditional on $\tau = \infty$.

Lemma 15. *Suppose belief continuity holds at a neighborhood of some nonempty compact set $D \subseteq \Delta(\Omega)$. Let $\Omega' \subseteq \Omega$ be a compact set such that for any $\omega' \in \Omega'$, there exists $\omega \in \Omega$ with $\omega \succ_{\mu}^{\text{KL}} \omega'$ for all $\mu \in D$. Then there exist a family of measurable sets of states $\{A_i\}_{i=1}^I$, a family of open sets of states $\{A'_i\}_{i=1}^I$, $\varepsilon > 0$, and $q_i > 0$ for each i such that $\bigcup_i A_i = \Omega'$ and*

$$\int \left(\frac{\int_{A_i} \hat{p}_{\mu}(z|\omega) d\mu(\omega) / \mu(A_i)}{\int_{A'_i} \hat{p}_{\mu}(z|\omega) d\mu(\omega) / \mu(A'_i)} \right)^{q_i} dP_{\mu}(z) \leq 1 - \varepsilon \quad (24)$$

for each i and $\mu \in B_{\varepsilon}(D)$ with $\mu(A_i), \mu(A'_i) > 0$. Moreover, for any initial belief μ_0 ,

$$\mathbb{P}_{\mu_0}[\mu_t(\Omega') \rightarrow 0, \mu_t \in B_{\varepsilon}(D) \forall t] = \mathbb{P}_{\mu_0}[\mu_t \in B_{\varepsilon}(D) \forall t]. \quad (25)$$

Proof. By assumption, for each $\omega \in \Omega'$, there exists $\phi(\omega) \in \Omega$ such that, for all $\mu \in D$,

$$\int \log \frac{\hat{p}_{\mu}(z|\omega)}{\hat{p}_{\mu}(z|\phi(\omega))} p_{\mu}(z) d\nu(z) < 0. \quad (26)$$

Claim 1: For each $\omega \in \Omega'$, there exist $q_{\omega} \in (0, q^*]$ and $\zeta_{\omega} > 0$ such that, for all $\mu \in \overline{B}_{\zeta_{\omega}}(D)$,

$$\int \left(\frac{\hat{p}_{\mu}(z|\omega)}{\hat{p}_{\mu}(z|\phi(\omega))} \right)^{q_{\omega}} p_{\mu}(z) d\nu(z) \leq 1 - \zeta_{\omega}. \quad (27)$$

Proof of Claim 1. For each $\omega \in \Omega'$ and $\mu \in D$, (26) and Lemma 2 yield $q_{\omega, \mu} \in (0, q^*]$ such that

$$\int \left(\frac{\hat{p}_{\mu}(z|\omega)}{\hat{p}_{\mu}(z|\phi(\omega))} \right)^q p_{\mu}(z) d\nu(z) < 1$$

for all $q \in (0, q_{\omega, \mu}]$. By belief continuity, the LHS is continuous in μ in a neighborhood of D (Lemma 9). Thus, $\int \left(\frac{\hat{p}_{\mu'}(z|\omega)}{\hat{p}_{\mu'}(z|\phi(\omega))} \right)^{q_{\omega, \mu}} p_{\mu'}(z) d\nu(z) < 1$ for all μ' in some neighborhood B_{μ} of μ .

Since D is compact, by taking a finite subcover of $\{B_\mu : \mu \in D\}$, we can choose $q_{\omega,\mu} =: q_\omega$ to be independent of μ . Since the subcover of D is open, (27) holds for $\zeta_\omega > 0$ sufficiently small. \square

Claim 2: For each $\omega \in \Omega'$, there exists $\varepsilon_\omega > 0$ such that, for any $\mu \in \overline{B}_{\zeta_\omega}(D)$ with $\mu(B_{\varepsilon_\omega}(\omega) \cap \Omega'), \mu(B_{\varepsilon_\omega}(\phi(\omega))) > 0$, we have

$$\int \left(\frac{\int_{B_{\varepsilon_\omega}(\omega) \cap \Omega'} \hat{p}_\mu(z|\omega') d\mu(\omega') / \mu(B_{\varepsilon_\omega}(\omega) \cap \Omega')}{\int_{B_{\varepsilon_\omega}(\phi(\omega))} \hat{p}_\mu(z|\omega') d\mu(\omega') / \mu(B_{\varepsilon_\omega}(\phi(\omega)))} \right)^{q_\omega} p_\mu(z) d\nu(z) \leq 1 - \zeta_\omega/2. \quad (28)$$

Proof of Claim 2. Fix $\omega \in \Omega'$. For each $\mu \in \overline{B}_{\zeta_\omega}(D)$, we first observe that

$$\max_{\hat{\mu} \in \Delta(\overline{B}_\varepsilon(\omega)), \hat{\mu}' \in \Delta(\overline{B}_\varepsilon(\phi(\omega)))} \int \left(\frac{\int \hat{p}_\mu(z|\omega') d\hat{\mu}(\omega')}{\int \hat{p}_\mu(z|\omega') d\hat{\mu}'(\omega')} \right)^{q_\omega} p_\mu(z) d\nu(z) \quad (29)$$

is continuous in ε by the maximum theorem: Indeed, $\left(\frac{\int \hat{p}_\mu(z|\omega') d\hat{\mu}(\omega')}{\int \hat{p}_\mu(z|\omega') d\hat{\mu}'(\omega')} \right)^{q_\omega}$ is continuous in $\hat{\mu}, \hat{\mu}'$, since for each z , $\hat{p}_\mu(z|\cdot)$ is continuous and bounded (by Assumption 2 and compactness of Ω). Thus, $\int \left(\frac{\int \hat{p}_\mu(z|\omega') d\hat{\mu}(\omega')}{\int \hat{p}_\mu(z|\omega') d\hat{\mu}'(\omega')} \right)^{q_\omega} p_\mu(z) d\nu(z)$ is continuous in $\hat{\mu}, \hat{\mu}'$ by the dominated convergence theorem, as $\left(\frac{\int \hat{p}_\mu(\cdot|\omega') d\hat{\mu}(\omega')}{\int \hat{p}_\mu(\cdot|\omega') d\hat{\mu}'(\omega')} \right)^{q_\omega} p_\mu(\cdot)$ is dominated by $h(\cdot)$ (Assumption 1.2). Therefore, by (27), there exists $\varepsilon_{\omega,\mu} > 0$ such that (29) is strictly less than $1 - \zeta_\omega/2$ for all $\varepsilon \in (0, \varepsilon_{\omega,\mu}]$.

Moreover (29) is continuous in μ by the maximum theorem, as $\int \left(\frac{\int \hat{p}_\mu(z|\omega') d\hat{\mu}(\omega')}{\int \hat{p}_\mu(z|\omega') d\hat{\mu}'(\omega')} \right)^{q_\omega} p_\mu(z) d\nu(z)$ is continuous in μ by belief continuity (using the same argument as in Lemma 9). Therefore, $\max_{\hat{\mu} \in \Delta(\overline{B}_{\varepsilon_{\omega,\mu}}(\omega)), \hat{\mu}' \in \Delta(\overline{B}_{\varepsilon_{\omega,\mu}}(\phi(\omega)))} \int \left(\frac{\int \hat{p}_{\mu'}(z|\omega') d\hat{\mu}(\omega')}{\int \hat{p}_{\mu'}(z|\omega') d\hat{\mu}'(\omega')} \right)^{q_\omega} p_{\mu'}(z) d\nu(z) < 1 - \zeta_\omega/2$ for all μ' in some neighborhood B_μ of μ . Since $\overline{B}_{\zeta_\omega}(D)$ is compact, by taking a finite subcover of $\{B_\mu : \mu \in \overline{B}_{\zeta_\omega}(D)\}$, we can choose $\varepsilon_{\omega,\mu} =: \varepsilon_\omega$ to be independent of μ . This establishes (28). \square

Since $\{B_{\varepsilon_\omega}(\omega) \cap \Omega' : \omega \in \Omega'\}$ covers the compact set Ω' , there is a finite subcover $\{B_{\varepsilon_{\omega_i}}(\omega_i) \cap \Omega' : i = 1, \dots, I\}$. Thus by setting $A_i := B_{\varepsilon_{\omega_i}}(\omega_i) \cap \Omega'$, $A'_i := B_{\varepsilon_{\omega_i}}(\phi(\omega_i))$, $q_i := q_{\omega_i}$ for each i , and $\varepsilon := \min_i \min\{\varepsilon_{\omega_i}, \zeta_{\omega_i}/2\}$, we obtain (24) for each i and any $\mu \in \overline{B}_\varepsilon(D)$ with $\mu(A_i), \mu(A'_i) > 0$.

For the “moreover” part, define $\ell_t^i := \left(\frac{\mu_{\min\{t,\tau\}}(A_i)}{\mu_{\min\{t,\tau\}}(A'_i)} \right)^{q_i}$ for each $i = 1, \dots, I$, where $\tau := \inf\{s : \mu_s \notin B_\varepsilon(D)\}$. For any initial belief μ_0 , ℓ_t^i is a nonnegative supermartingale by (24). Thus, Doob’s convergence theorem yields an L^∞ random variable ℓ_∞^i such that $\ell_t^i \rightarrow \ell_\infty^i$ a.s. Observe that, for any initial belief $\mu_0 \in B_\varepsilon(D)$, Markov’s inequality and (24) imply

$$\mathbb{P}_{\mu_0}[\ell_1^i \geq (1 - \varepsilon/2)\ell_0^i] \leq \frac{\mathbb{E}_{\mu_0}[\ell_1^i]}{(1 - \varepsilon/2)\ell_0^i} \leq \frac{1 - \varepsilon}{1 - \varepsilon/2}.$$

Thus, conditional on any $\mu_t \in B_\varepsilon(D)$, the probability that ℓ_{t+1}^i is less than $(1 - \varepsilon/2)\ell_t^i$ is at least $\frac{\varepsilon/2}{1 - \varepsilon/2}$. This implies that $\mathbb{P}_{\mu_0}[\ell_\infty^i > 0, \tau = \infty] = 0$ for any initial belief. Since, conditional on $\tau = \infty$, we have $\ell_t^i = \frac{\mu_t(A_i)}{\mu_t(A'_i)}$ for each i and t , this ensures the desired claim. \square

B.4 Proof of Theorem 3

Call $M \subseteq \Delta(\Omega)$ *Lyapunov stable* if for any neighborhood B of M and $\gamma < 1$, there exists a neighborhood B' of M such that $\mathbb{P}_{\mu_0}[\mu_t \in B \forall t] \geq \gamma$ for every initial belief $\mu_0 \in B'$. We start with a preliminary lemma:

Lemma 16. *Let $\Omega' \subseteq \Omega$ be a nonempty and measurable set such that $\Delta(\Omega')$ is Lyapunov stable and belief continuity holds at a neighborhood of $\Delta(\Omega')$. Then $\Delta(S(\Omega'))$ is Lyapunov stable.*

Proof. Write $\Omega'' := S(\Omega')$, which is nonempty and compact by Lemma 14. If $\Omega'' = \overline{\Omega'}$, the claim is immediate, so assume $\Omega'' \subsetneq \overline{\Omega'}$. Take any neighborhood B of $\Delta(\Omega'')$ and any $\gamma < 1$. Pick N large enough that $\overline{\Delta(B_{1/N}(\Omega''))} \subseteq B$. By Lemma 15, there exist a family of measurable sets of states $\{A_i\}_{i=1}^I$, a family of open sets of states $\{A'_i\}_{i=1}^I$, $\varepsilon > 0$, and $q_i > 0$ for each i such that $\bigcup_i A_i = \overline{\Omega'} \setminus B_{1/N}(\Omega'')$ and (24) holds for each i and $\mu \in B_\varepsilon(\Delta(\overline{\Omega'}))$ with $\mu(A_i), \mu(A'_i) > 0$.

Define $C := \{\mu \in B_{\varepsilon'}(\Delta(\overline{\Omega'})) : \sum_i \left(\frac{\mu(A_i)}{\mu(A'_i)}\right)^{q_i} \leq \varepsilon'\}$, where by construction of $\{A_i\}_{i=1}^I$, we can choose $\varepsilon' \in (0, \varepsilon)$ small enough that $C \subseteq B$. Set $\tau := \inf\{t : \mu_t \notin C\}$. Then from any initial belief, each $\ell_t^i := \left(\frac{\mu_{\min\{t, \tau\}}(A_i)}{\mu_{\min\{t, \tau\}}(A'_i)}\right)^{q_i}$ is a nonnegative supermartingale by (24), and thus a.s. converges to an L^∞ limit ℓ_∞^i .

For each $\eta > 0$, define $C'_\eta := \{\mu \in \Delta(\Omega) : \sum_i \left(\frac{\mu(A_i)}{\mu(A'_i)}\right)^{q_i}, \mu(\Omega \setminus \overline{\Omega'}) \leq \eta\}$, which is a neighborhood of $\Delta(\Omega'')$. For any initial belief $\mu_0 \in C'_\eta$, we have

$$\mathbb{P}_{\mu_0}[\tau < \infty] \leq \mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \mu_t \notin B_{\varepsilon'}(\Delta(\overline{\Omega'}))] + \mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \sum_i \left(\frac{\mu_t(A_i)}{\mu_t(A'_i)}\right)^{q_i} > \varepsilon', \mu_s \in B_{\varepsilon'}(\Delta(\overline{\Omega'})) \forall s \leq t].$$

By Lyapunov stability of $\Delta(\Omega')$, we can pick η sufficiently small that the first term is less than $\frac{1-\gamma}{2}$ for all $\mu_0 \in C'_\eta$. Moreover, the second term is less than $\mathbb{P}_{\mu_0}[\sum_i \ell_\infty^i > \varepsilon'] \leq \mathbb{E}_{\mu_0}[\sum_i \ell_1^i] / \varepsilon' \leq \eta / \varepsilon'$ by Markov's inequality, Fatou's lemma and the fact that $\sum_i \ell_t^i$ is a nonnegative supermartingale. Thus, by taking η sufficiently small, $\mathbb{P}_{\mu_0}[\mu_t \in B \forall t] \geq \mathbb{P}_{\mu_0}[\mu_t \in C \forall t] \geq \gamma$ for every initial belief $\mu_0 \in C'_\eta$. \square

Proof of Theorem 3. Let $\Omega^k := S^k(\Omega)$ for $k = 0, 1, \dots$, which is a nested sequence of nonempty compact sets (Lemma 14). We inductively show that $\mathbb{P}_{\mu_0}[\mu_t(\Omega^k) \rightarrow 1] = 1$ for all initial beliefs μ_0 and every $k \geq 0$. Case $k = 0$ is true by definition.

Suppose the claim is true for all $k = 0, 1, \dots, \kappa - 1$ and consider $k = \kappa$. Take any N with $\Omega \setminus B_{1/N}(\Omega^\kappa)$ nonempty. By Lemma 15 applied with $\Omega' = \Omega \setminus B_{1/N}(\Omega^\kappa)$ and $D = \Delta(\Omega^{\kappa-1})$, there exists $\varepsilon > 0$ such that (25) holds for each initial belief μ_0 .

Take any $\gamma < 1$. Then by Lyapunov stability of $\Delta(\Omega^{\kappa-1})$ (Lemma 16) there exists a neighborhood B of $\Delta(\Omega^{\kappa-1})$ such that $\mathbb{P}_{\mu_0}[\mu_t \in B_\varepsilon(\Delta(\Omega^{\kappa-1})) \forall t] \geq \gamma$ for every initial belief $\mu_0 \in B$. Thus, for any initial belief μ_0 , (25) and the inductive hypothesis that $\mathbb{P}_{\mu_0}[\mu_t(\Omega^{\kappa-1}) \rightarrow 1] = 1$ imply

$$\mathbb{P}_{\mu_0}[\mu_t(\Omega \setminus B_{1/N}(\Omega^\kappa)) \rightarrow 0] \geq \mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \mu_t \in B] \gamma = \gamma.$$

Since this holds for all $\gamma < 1$ and N large enough, we have

$$\mathbb{P}_{\mu_0}[\mu_t(\Omega^k) \rightarrow 1] = \mathbb{P}_{\mu_0}[\mu_t(B_{1/N}(\Omega^k)) \rightarrow 1 \forall N] = 1,$$

for all initial beliefs μ_0 , completing the inductive step. Finally, for all initial beliefs μ_0 ,

$$\mathbb{P}_{\mu_0}[\mu_t(S^\infty(\Omega)) \rightarrow 1] = \mathbb{P}_{\mu_0}[\mu_t(\Omega^k) \rightarrow 1 \forall k] = 1.$$

Thus, $\Delta(S^\infty(\Omega))$ is globally stable. \square

B.5 Proof of Theorem 4

Lemma 17. *Suppose Ω'' is open and belief continuity holds at some neighborhood of $\Delta(\Omega'')$. If $\Delta(\Omega'')$ is a strict BeNE set, then $\Delta(\Omega'')$ is locally stable and Lyapunov stable.*

Proof. Based on the fact that $\Delta(\Omega'')$ is a strict BeNE set, we can apply Lemma 15 with $\Omega' = \Omega \setminus \Omega''$ and $D = \Delta(\overline{\Omega''})$. This yields measurable sets of states $\{A_i\}_{i=1}^I$ with $\bigcup_i A_i = \Omega \setminus \Omega''$, open sets of states $\{A'_i\}_{i=1}^I$, $\varepsilon > 0$, and $q_i > 0$ for each i such that (24) holds for each i and $\mu_0 \in B_\varepsilon(\Delta(\overline{\Omega''}))$ with $\mu(A_i), \mu(A'_i) > 0$, and (25) holds for each initial belief μ_0 .

To show Lyapunov stability of $\Delta(\Omega'')$, take any $\gamma < 1$ and neighborhood B of $\Delta(\Omega'')$. Given any $\eta > 0$, consider the neighborhood of $\Delta(\Omega'')$ of the form

$$C_\eta := \left\{ \mu \in \Delta(\Omega) : \sum_i \left(\frac{\mu(A_i)}{\mu(A'_i)} \right)^{q_i} < \eta \right\}.$$

Pick $\eta_+, \eta_- > 0$ small enough that $C_{\eta_+} \subseteq B \cap B_\varepsilon(\Delta(\Omega''))$ and $\frac{\eta_-}{\eta_+} \leq 1 - \gamma$. For any i and any initial belief μ_0 , $\ell_t^i := \left(\frac{\mu_{\min\{t, \tau\}}(A_i)}{\mu_{\min\{t, \tau\}}(A'_i)} \right)^{q_i}$ with $\tau := \inf\{s : \mu_s \notin C_{\eta_+}\}$ is a nonnegative supermartingale by (24), so Doob's convergence theorem yields an L^∞ random variable ℓ_∞^i such that $\ell_t^i \rightarrow \ell_\infty^i$ a.s. For any initial belief $\mu_0 \in C_{\eta_-}$,

$$\mathbb{P}_{\mu_0}[\exists t, \mu_t \notin B] \leq \mathbb{P}_{\mu_0}[\sum_i \ell_\infty^i \geq \eta_+] \leq \mathbb{E}_{\mu_0}[\sum_i \ell_\infty^i] / \eta_+ \leq \frac{\eta_-}{\eta_+},$$

where the second inequality uses Markov's inequality and the third follows from Fatou's lemma and the fact that each ℓ_t^i is a nonnegative supermartingale. Thus, $\mathbb{P}_{\mu_0}[\mu_t \in B \forall t] \geq \gamma$ for all $\mu_0 \in C_{\eta_-}$, proving that $\Delta(\Omega'')$ is Lyapunov stable.

To show that $\Delta(\Omega'')$ is locally stable, take any $\gamma < 1$. Since $\Delta(\Omega'')$ is Lyapunov stable, there exists a neighborhood B of $\Delta(\Omega'')$ such that $\mathbb{P}_{\mu_0}[\mu_t \in B_\varepsilon(\Delta(\Omega'')) \forall t] \geq \gamma$ for any initial belief $\mu_0 \in B$. Thus, (25) implies that for any initial belief μ_0 in B ,

$$\mathbb{P}_{\mu_0}[\mu_t(\Omega'') \rightarrow 1] \geq \mathbb{P}_{\mu_0}[\mu_t(\Omega'') \rightarrow 1, \mu_t \in B_\varepsilon(\Delta(\Omega'')) \forall t] \geq \gamma,$$

showing that $\Delta(\Omega'')$ is locally stable. \square

Proof of Theorem 4. For Ω' as in the theorem, let $\Omega^k := S^k(\Omega')$ for each $k = 0, 1, \dots, \infty$. Suppose $\Delta(\Omega')$ is a strict BeNE set. Then $\Delta(\Omega')$ is Lyapunov stable (Lemma 17), which combined with Lemma 16 implies that $\Delta(\Omega^k)$ is Lyapunov stable for each $k \in \mathbb{N}$.

Fix any $\gamma < 1$. By Lemma 17, $\Delta(\Omega')$ is locally stable. Thus, there exists a neighborhood B_0 of $\Delta(\Omega')$ such that $\mathbb{P}_{\mu_0}[\mu_t(\Omega') \rightarrow 1] \geq \gamma$ for any initial belief $\mu_0 \in B_0$. We show inductively that for each $k \in \mathbb{N}$, $\mathbb{P}_{\mu_0}[\mu_t(\Omega^k) \rightarrow 1] \geq \gamma$ for any initial belief $\mu_0 \in B_0$.

For $k = 0$, the claim is true by choice of B_0 . Thus, suppose the claim holds for $k \leq \kappa - 1$ and consider the case $k = \kappa$. Take any $N > 0$ such that $\Omega \setminus B_{1/N}(\Omega^\kappa)$ is nonempty. By Lemma 15 applied with $D = \Delta(\Omega^{\kappa-1})$, there exists $\varepsilon > 0$ such that for all initial beliefs μ_0 ,

$$\mathbb{P}_{\mu_0}[\mu_t(\Omega \setminus B_{1/N}(\Omega^\kappa)) \rightarrow 0, \mu_t \in B_\varepsilon(D) \forall t] = \mathbb{P}_{\mu_0}[\mu_t \in B_\varepsilon(D) \forall t]. \quad (30)$$

Since $\Delta(\Omega^{\kappa-1})$ is Lyapunov stable, for any $\eta < 1$, there exists a neighborhood C of $\Delta(\Omega^{\kappa-1})$ such that, for any initial belief $\mu_0 \in C$, $\mathbb{P}_{\mu_0}[\mu_t \in B_\varepsilon(\Delta(\Omega^{\kappa-1})) \forall t] \geq \eta$. Thus, for any initial belief $\mu_0 \in B_0$,

$$\mathbb{P}_{\mu_0}[\mu_t(B_{1/N}(\Omega^\kappa)) \rightarrow 1] \geq \mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \mu_t \in C] \eta \geq \gamma \eta,$$

where the first inequality uses (30) and the second uses the inductive hypothesis that $\mathbb{P}_{\mu_0}[\mu_t(\Omega^{\kappa-1}) \rightarrow 1] \geq \gamma$. Since η can be chosen arbitrarily close to 1, $\mathbb{P}_{\mu_0}[\mu_t(B_{1/N}(\Omega^\kappa)) \rightarrow 1] \geq \gamma$. Since N can be chosen arbitrarily large, this implies $\mathbb{P}_{\mu_0}[\mu_t(\Omega^\kappa) \rightarrow 1] = \mathbb{P}_{\mu_0}[\mu_t(B_{1/N}(\Omega^\kappa)) \rightarrow 1 \forall N \in \mathbb{N}] \geq \gamma$, as claimed.

This shows that $\Delta(S^k(\Omega'))$ is locally stable for all $k \in \mathbb{N}$. Finally, to complete the proof, observe that, for any initial belief $\mu_0 \in B_0$,

$$\mathbb{P}_{\mu_0}[\mu_t(S^\infty(\Omega')) \rightarrow 1] = \mathbb{P}_{\mu_0}[\mu_t(\Omega^k) \rightarrow 1 \forall k] \geq \gamma.$$

Thus, $\Delta(S^\infty(\Omega'))$ is also locally stable. \square

References

- BANERJEE, A. (1992): “A simple model of herd behavior,” *The Quarterly Journal of Economics*, 107(3), 797–817.
- 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.
- BOHREN, J. A. (2016): “Informational herding with model misspecification,” *Journal of Economic Theory*, 163, 222–247.
- BOHREN, J. A., AND D. HAUSER (2021): “Learning with Heterogeneous Misspecified Models: Characterization and Robustness,” *Econometrica*, forthcoming.
- BOHREN, J. A., A. IMAS, AND M. ROSENBERG (2019): “The Dynamics of Discrimination: Theory and Evidence,” *American Economic Review*, 109(10), 3395–3436.

- BOLTON, P., AND C. HARRIS (1999): “Strategic experimentation,” *Econometrica*, 67(2), 349–374.
- BUSHONG, B., AND T. GAGNON-BARTSCH (2019): “Learning with Misattribution of Reference Dependence,” working paper.
- CHADE, H., AND E. SCHLEE (2002): “Another look at the Radner–Stiglitz nonconcavity in the value of information,” *Journal of Economic Theory*, 107(2), 421–452.
- CHAMLEY, C. (2004): *Rational herds: Economic models of social learning*. Cambridge University Press.
- CHO, I.-K., AND K. KASA (2017): “Gresham’s Law of Model Averaging,” *American Economic Review*, 107(11), 3589–3616.
- CONG, T. (2019): “A Bayesian Learning Model of Cyclical Actions under misspecification,” working paper.
- DASARATHA, K., AND K. HE (2020): “Network structure and naive sequential learning,” *Theoretical Economics*, 15(2), 415–444.
- EIL, D., AND J. M. RAO (2011): “The good news-bad news effect: asymmetric processing of objective information about yourself,” *American Economic Journal: Microeconomics*, 3(2), 114–38.
- 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,” *Theoretical Economics*, forthcoming.
- ESPONDA, I., D. POUZO, AND Y. YAMAMOTO (2021): “Asymptotic Behavior of Bayesian Learners with Misspecified Models,” working paper.
- EYSTER, E., AND M. RABIN (2005): “Cursed equilibrium,” *Econometrica*, 73(5), 1623–1672.
- (2010): “Naive herding in rich-information settings,” *American Economic Journal: Microeconomics*, 2(4), 221–43.
- FRICK, M., R. IJIMA, AND Y. ISHII (2020a): “Misinterpreting Others and the Fragility of Social Learning,” *Econometrica*, 88(6), 2281–2328.
- (2020b): “Stability and Robustness in Misspecified Learning Models,” Cowles Foundation Discussion Paper.
- FUDENBERG, D., G. LANZANI, AND P. STRACK (2021a): “Limit Points of Endogenous Misspecified Learning,” *Econometrica*, forthcoming.
- (2021b): “Pathwise Concentration Bounds for Misspecified Bayesian Beliefs,” Working paper.
- 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.
- GOEREE, J. K., T. R. PALFREY, AND B. W. ROGERS (2006): “Social learning with private and common values,” *Economic Theory*, 28(2), 245–264.
- HANN-CARUTHERS, W., V. V. MARTYNOV, AND O. TAMUZ (2018): “The speed of sequential asymptotic learning,” *Journal of Economic Theory*, 173, 383–409.
- HE, K. (2021): “Mislearning from Censored Data: The Gambler’s Fallacy in Optimal-Stopping Problems,” *Theoretical Economics*, forthcoming.

- HEIDHUES, P., B. KŐSZEGI, AND P. STRACK (2018): “Unrealistic expectations and misguided learning,” *Econometrica*, 86(4), 1159–1214.
- (2021): “Convergence in models of misspecified learning,” *Theoretical Economics*, 16(1), 73–99.
- JEHIEL, P. (2005): “Analogy-based expectation equilibrium,” *Journal of Economic Theory*, 123(2), 81–104.
- LEVY, G., AND R. RAZIN (2018): “Information diffusion in networks with the Bayesian Peer Influence heuristic,” *Games and Economic Behavior*, 109, 262–270.
- MCLENNAN, A. (1984): “Price dispersion and incomplete learning in the long run,” *Journal of Economic Dynamics and Control*, 7(3), 331–347.
- MILGROM, P., AND J. ROBERTS (1990): “Rationalizability, learning, and equilibrium in games with strategic complementarities,” *Econometrica*, pp. 1255–1277.
- MOBIUS, M. M., M. NIEDERLE, P. NIEHAUS, AND T. ROSENBLAT (2014): “Managing Self-Confidence: Theory and Experimental Evidence,” working paper.
- MOLAVI, P. (2019): “Macroeconomics with Learning and Misspecification: A General Theory and Applications,” working paper.
- MUROOKA, T., AND Y. YAMAMOTO (2021): “Multi-Player Bayesian Learning with Misspecified Models,” Working paper.
- 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.
- ROSENBERG, D., AND N. VIELLE (2019): “On the efficiency of social learning,” *Econometrica*, 87(6), 2141–2168.
- SCHWARTZSTEIN, J. (2014): “Selective attention and learning,” *Journal of the European Economic Association*, 12(6), 1423–1452.
- 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.
- 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.
- ZIMPER, A. (2007): “A fixed point characterization of the dominance-solvability of lattice games with strategic substitutes,” *International Journal of Game Theory*, 36(1), 107–117.

Online Appendix to “Belief Convergence under Misspecified Learning: A Martingale Approach”

Mira Frick, Ryota Iijima, and Yuhta Ishii

C Proofs for Section 4

C.1 Proof of Proposition 2

Consider the setting in Section 4.1. We prove a slight generalization of Proposition 2 that can also be combined with Theorem 4 (Appendix B) to show local stability of $\delta_{\hat{\omega}}$:

Proposition 5. *Take any compact interval $\Omega' \subseteq \Omega$ such that $m(\Omega') \subseteq \Omega'$. Then $S^k(\Omega') = m^k(\Omega')$ for all $k = 0, 1, \dots, \infty$. Moreover:*

1. *If m is increasing on Ω' , then $S^\infty(\Omega') = \{\hat{\omega}\}$ iff $\hat{\omega}$ is the unique fixed point of m in Ω' .*
2. *If m is decreasing on Ω' , then $S^\infty(\Omega') = \{\hat{\omega}\}$ iff $\hat{\omega}$ is the unique fixed point of m^2 in Ω' .*

Proof. For each ω , let $a(\omega) := a(\delta_\omega)$. Since $\text{KL}(G_a, \hat{G}_a(\cdot|\omega))$ is continuous in a and $a(\omega)$ is continuous in ω , the map m is continuous. Take any compact interval $\Omega' \subseteq \Omega$ such that $m(\Omega') \subseteq \Omega'$. We first show by induction that for all $n = 0, 1, \dots, \infty$, $S^n(\Omega') = m^n(\Omega') =: \Omega_n$ for some sequence of compact intervals Ω_n that is decreasing with respect to set inclusion. For $n = 0$, $S^0(\Omega') := \Omega' =: m^0(\Omega')$, so there is nothing to prove. Suppose the claim holds for all $n \leq k$. We show that $S(\Omega_k) = m(\Omega_k)$.

To see that $m(\Omega_k) \subseteq S(\Omega_k)$, take any $\omega \in m(\Omega_k)$. Then there is $\omega' \in \Omega_k$ with $\omega \succ_{\delta_{\omega'}}^{\text{KL}} \omega''$ for all $\omega'' \in \Omega \setminus \{\omega\}$. Thus, there does not exist $\omega'' \in \Omega$ such that $\omega'' \succ_{\mu}^{\text{KL}} \omega$ for all $\mu \in \Delta(\Omega_k) = \Delta(S^k(\Omega'))$. Moreover, $\omega \in \Omega'$, as $\Omega_k = S^k(\Omega') \subseteq \Omega'$ and $m(\Omega') \subseteq \Omega'$ by assumption. This implies $\omega \in S^{k+1}(\Omega') = S(\Omega_k)$.

To see that $S(\Omega_k) \subseteq m(\Omega_k)$, take any $\omega \in \Omega_k \setminus m(\Omega_k)$. Since m is continuous and Ω_k is a compact interval, so is $m(\Omega_k)$, say $m(\Omega_k) = [\underline{\omega}_{k+1}, \bar{\omega}_{k+1}]$. Thus, either (i) $\omega < \underline{\omega}_{k+1}$ or (ii) $\omega > \bar{\omega}_{k+1}$. Consider case (i); a symmetric argument applies to case (ii). For any $\omega'' \in \Omega_k$, we have $\omega < \underline{\omega}_{k+1} \leq m(\omega'')$, which implies $\text{KL}(G_{a(\omega'')}, \hat{G}_{a(\omega'')}(\cdot|\omega)) > \text{KL}(G_{a(\omega'')}, \hat{G}_{a(\omega'')}(\cdot|\underline{\omega}_{k+1}))$ by the strict quasi-convexity assumption on KL. Moreover, for any $\mu \in \Delta(\Omega_k)$, the intermediate value theorem yields $\omega'' \in \Omega_k$ such that $a(\mu) = a(\omega'')$, as $a(\cdot)$ is FOSD-increasing and continuous. Thus, for all $\mu \in \Delta(\Omega_k)$, $\text{KL}(G_{a(\mu)}, \hat{G}_{a(\mu)}(\cdot|\omega)) > \text{KL}(G_{a(\mu)}, \hat{G}_{a(\mu)}(\cdot|\underline{\omega}_{k+1}))$, i.e., $\underline{\omega}_{k+1} \succ_{\mu}^{\text{KL}} \omega$. Since $\underline{\omega}_{k+1} \in m(\Omega_k) \subseteq S(\Omega_k)$ by the previous paragraph, this shows $\omega \notin S(\Omega_k)$.

Thus, by induction, $S^k(\Omega') = m^k(\Omega') =: \Omega_k$ for all $k \in \mathbb{N}$, and hence also $S^\infty(\Omega') := \bigcap_k S^k(\Omega') = \bigcap_k m^k(\Omega') =: m^\infty(\Omega)$. Moreover, since the $\Omega_k = [\underline{\omega}_k, \bar{\omega}_k]$ form a decreasing sequence of compact intervals, $S^\infty(\Omega') = [\underline{\omega}_\infty, \bar{\omega}_\infty]$ is nonempty, with $\underline{\omega}_\infty = \lim_k \underline{\omega}_k$ and $\bar{\omega}_\infty = \lim_k \bar{\omega}_k$.

For the “moreover” part, suppose m is increasing. Then $S^k(\Omega') = [\underline{\omega}_k, \bar{\omega}_k] = [m(\underline{\omega}_{k-1}), m(\bar{\omega}_{k-1})]$ for all $k \geq 1$. By continuity of m , this implies that $\underline{\omega}_\infty$ and $\bar{\omega}_\infty$ are fixed points of m in Ω' . Thus, the “if” direction holds. For the “only if” direction, suppose $\underline{\omega}_\infty = \bar{\omega}_\infty =: \hat{\omega}$. Then for any fixed point $\omega \in \Omega'$ of m , we have $\omega \in m^k(\Omega') = S^k(\Omega')$ for all $k \in \mathbb{N}$, so $\underline{\omega}_\infty \leq \omega \leq \bar{\omega}_\infty$, i.e., $\omega = \hat{\omega}$.

Finally, suppose m is decreasing. Then $S^k(\Omega') = [\underline{\omega}_k, \bar{\omega}_k] = [m^2(\underline{\omega}_{k-2}), m^2(\bar{\omega}_{k-2})]$ for all $k \geq 2$. By continuity of m , this implies that $\underline{\omega}_\infty$ and $\bar{\omega}_\infty$ are fixed points of m^2 . Thus, the “if” direction holds. For the “only if” direction, suppose $\underline{\omega}_\infty = \bar{\omega}_\infty =: \hat{\omega}$. Then for any fixed point $\omega \in \Omega'$ of m^2 , we have $\omega \in m^k(\Omega') = S^k(\Omega')$ for all even $k \in \mathbb{N}$, so again $\underline{\omega}_\infty \leq \omega \leq \bar{\omega}_\infty$, i.e., $\omega = \hat{\omega}$. \square

C.2 Details for the applications in Section 4.1

C.2.1 Effort choice under overconfidence

As in [Heidhues, Kőszegi, and Strack \(2018\)](#) (HKS), assume Q is twice continuously differentiable with (i) $Q_{aa} < 0$, and $Q_a(\underline{a}, \beta, \omega) > 0 > Q_a(\bar{a}, \beta, \omega)$ for all (β, ω) ; (ii) $Q_\beta, Q_\omega > 0$; (iii) $Q_{a\omega} > 0$; (iv) $Q_{a\beta} \leq 0$; (v) $|Q_\omega| < \kappa$ for some constant $\kappa > 0$. Then standard arguments guarantee that the optimal action $a(\mu)$ is continuous and FOSD-increasing. Moreover, if $\hat{\beta} > \beta$, any state $\omega > \omega^*$ is dominated by ω^* at all beliefs, because $0 > Q(a, \beta, \omega^*) - Q(a, \hat{\beta}, \omega^*) > Q(a, \beta, \omega^*) - Q(a, \hat{\beta}, \omega)$ for all a . Thus, $m(\omega) \leq \omega^*$ for all ω . Hence, for all ω , $Q_a(a(\omega), \beta, \omega^*) - Q_a(a(\omega), \hat{\beta}, m(\omega)) \geq Q_a(a(\omega), \beta, \omega^*) - Q_a(a(\omega), \hat{\beta}, \omega^*) > 0$, which by (10) implies that m is increasing. HKS also assume that m has a unique fixed point $\hat{\omega}$; their Proposition 1 shows that this obtains under several specific functional forms Q , or if $\hat{\beta} - \beta$ is sufficiently small, Q_β is bounded and Q_ω is bounded away from 0. Given this, Proposition 2 and Theorem 2 imply that $\delta_{\hat{\omega}}$ is globally stable.

C.2.2 Optimal stopping under the gambler’s fallacy

In [He \(2021\)](#), each period consists of a two-stage decision problem. In the first stage, output x_1 follows $\mathcal{N}(m_1^*, \sigma^2)$. If the realized x_1 is lower than the agent’s stopping threshold a , then second-stage output x_2 is observed, which follows $\mathcal{N}(m_2^*, \sigma^2)$. The agent knows the first-stage mean m_1^* and the variance σ^2 in both stages, but is uncertain about the second-stage mean m_2 . Thus, the state space $\Omega = [\underline{m}_2, \bar{m}_2]$ represents values of second-stage means, with true state $\omega^* = m_2^*$.⁴¹ While in reality there is no correlation between x_1 and x_2 , the agent perceives negative correlation. That is, her perceived distribution of x_2 given m_2 and conditional on period-1 realization x_1 is $\mathcal{N}(m_2 - \gamma(x_1 - m_1^*), \sigma^2)$, where $\gamma \geq 0$ captures the extent of the agent’s bias. Given current

⁴¹[He \(2021\)](#) also considers the case in which the agent updates beliefs about both m_1 and m_2 , assuming that the state space Ω is a bounded parallelogram in \mathbb{R}^2 whose left and right edges are parallel to the y -axis and whose top and bottom edges have slope $-\gamma$. In this case, any $\omega = (m_1, m_2)$ with $m_1 \neq m_1^*$ is dominated by $\omega' := (m_1 + d, m_2 - \gamma d)$ such that $|m_1 - m_1^*| > |m_1 + d - m_1^*|$ for some d . This is because ω' yields a lower KL-divergence for the first-stage, while it provides the same second-stage prediction as ω after any realization of x_1 . Therefore, after one round of elimination, we can focus on the one-dimensional state space that corresponds to values of m_2 .

belief $\mu \in \Delta(\Omega)$, the agent chooses the threshold $a \in \mathbb{R}$ to maximize the expected value of $u : \mathbb{R} \times (\mathbb{R} \cup \{\emptyset\}) \rightarrow \mathbb{R}$, where $u(x_1, x_2)$ denotes the utility when she draws (x_1, x_2) , and $u(x_1, \emptyset)$ denotes the utility when she only draws x_1 . Under the assumptions in He (2021), $a(\cdot)$ is continuous and FOSD-increasing in μ .

This model maps to the setting in Section 4.1 by letting $g(a) := \omega^*$ and $\hat{g}(a, \omega) := \omega - \gamma(\mathbb{E}[x_1 | x_1 \leq a] - m_1^*)$. By (10), m is increasing, as $g'(a) - \frac{\partial \hat{g}}{\partial a}(a, \omega) = \gamma \frac{\partial \mathbb{E}[x_1 | x_1 \leq a]}{\partial a} \geq 0$ for all a . As He (2021) shows, there is a unique BeNE $\delta_{\hat{\omega}}$, where $\hat{\omega} < \omega^*$. Thus, Proposition 2 and Theorem 2 imply that $\delta_{\hat{\omega}}$ is globally stable.

C.3 Preliminary results for Section 4.2

The following result shows that δ_{ω} is globally stable whenever ω strictly q -dominates all other states at all mixed beliefs.

Proposition 6. *Consider any $\omega \in \Omega$. Suppose that belief continuity holds and for some $q > 0$, we have $\omega \succsim_{\mu}^q \omega'$ for all $\omega' \neq \omega$ and all μ , with strict dominance for all mixed μ . Then δ_{ω} is globally stable.*

Proof. Fix any initial belief μ_0 . By Lemma 3, $\ell_t(\omega') := \left(\frac{\mu_t(\omega')}{\mu_t(\omega)}\right)^q$ is a nonnegative supermartingale for each $\omega' \neq \omega$, since $\omega \succsim_{\mu}^q \omega'$ for all μ . Thus, by Doob's convergence theorem, there exists an L^∞ random variable $\ell_\infty(\omega')$ such that $\ell_t(\omega') \rightarrow \ell_\infty(\omega') \geq 0$ a.s. Hence, the belief process μ_t converges a.s. Let μ_∞ denote the limit. Suppose for a contradiction that $\mu_\infty \neq \delta_{\omega}$ with positive probability, which implies that for some $\omega' \neq \omega$, $\ell_\infty(\omega') > 0$ with positive probability. Then there exists a compact set $K \subseteq \Delta(\Omega)$ with $\mu(\omega), \mu(\omega') > 0$ for each $\mu \in K$ such that $\mathbb{P}_{\mu_0}[\exists \tau \text{ s.t. } \mu_t \in K \forall t \geq \tau \text{ and } \exists \lim_t \frac{\mu_t(\omega')}{\mu_t(\omega)}] > 0$. But for each $\mu \in K$, we have $\omega \succ_{\mu}^q \omega'$, which implies that $\hat{p}_{\mu}(z|\omega) > \hat{p}_{\mu}(z|\omega')$ for some $z \in \text{supp}(P_{\mu})$. This yields a contradiction with Lemma 10. \square

A corollary of Proposition 6 is that if the true signal distribution coincides with the perceived signal distribution at some state ω^* (i.e., the environment is correctly specified), then δ_{ω^*} is globally stable under an appropriate identification condition at mixed beliefs:

Corollary 2. *Suppose belief continuity holds and for some $\omega^* \in \Omega$, (i) $P_{\mu} = \hat{P}_{\mu}(\cdot|\omega^*)$ for all $\mu \in \Delta(\Omega)$, and (ii) $\hat{P}_{\mu}(\cdot|\omega^*) \neq \hat{P}_{\mu}(\cdot|\omega)$ for all $\omega \neq \omega^*$ and all mixed μ . Then δ_{ω^*} is globally stable.*

Proof. Take any $q \in (0, 1)$ and $\omega \neq \omega^*$. For each belief μ , we have

$$\int \left(\frac{\hat{p}_{\mu}(z|\omega)}{p_{\mu}(z)}\right)^q p_{\mu}(z) d\nu(z) \leq \left(\int_{\text{supp}(P_{\mu})} \hat{p}_{\mu}(z|\omega) d\nu(z)\right)^q = (\hat{P}_{\mu}(\text{supp}P_{\mu}|\omega))^q \leq 1, \quad (31)$$

where the first inequality holds by Jensen's inequality applied to the concave function x^q . Since $P_{\mu} = \hat{P}_{\mu}(\cdot|\omega^*)$ by (i), this shows $\omega^* \succsim_{\mu}^q \omega$. Consider any mixed μ . If the second inequality in (31) holds with equality, then (ii) ensures $\frac{\hat{p}_{\mu}(z|\omega)}{p_{\mu}(z)} \neq \frac{\hat{p}_{\mu}(z'|\omega)}{p_{\mu}(z')}$ for some $z, z' \in \text{supp}P_{\mu}$, in which case

the first inequality in (31) is strict. In either case, $\omega^* \succ_{\mu}^q \omega$. Thus, the conclusion follows from Proposition 6. \square

C.4 Proof of Lemma 4

Fix any $q \in (0, 1)$ and true state $\omega^* \in \Omega$. We will find $\varepsilon > 0$ such that learning is successful at ω^* for any $\hat{\beta}$ with $|\hat{\beta} - \beta| < \varepsilon$. This ensures the desired conclusion by finiteness of Ω . Consider any $\hat{\beta}$. Since C is constant and v is strictly convex, we have $\gamma_{\hat{\beta}}(\mu) = \bar{\gamma}$ for all mixed μ . Thus, for each mixed μ , the true and perceived probabilities of signal 1 satisfy $p_{\mu}(1) = \beta + \bar{\gamma}\omega^*$ and $\hat{p}_{\mu}(1|\omega) = \hat{\beta} + \bar{\gamma}\omega$. If $\hat{\beta} = \beta$, then Jensen's inequality implies that for any $\omega \neq \omega^*$ and mixed μ ,

$$\sum_z p_{\mu}(z) \left(\frac{\hat{p}_{\mu}(z|\omega)}{\hat{p}_{\mu}(z|\omega^*)} \right)^q < 1, \quad (32)$$

where the value of the left-hand side is independent of μ . Hence, there exists $\varepsilon > 0$ such that for any $\hat{\beta}$ with $|\hat{\beta} - \beta| < \varepsilon$ and any mixed μ and $\omega \neq \omega^*$, (32) continues to hold, so that $\omega^* \succ_{\mu}^q \omega$. Thus, for any $\hat{\beta}$ with $|\hat{\beta} - \beta| < \varepsilon$, Proposition 6 implies that δ_{ω^*} is globally stable in any state space $\Omega' \subseteq \Omega$ with $\omega^* \in \Omega'$, i.e., learning is successful at ω^* . \square

C.5 Proof of Lemma 5

Consider any strictly increasing cost function C . We will prove the following: Fix any $\hat{\beta}$, $\omega \in \Omega$, and $\tilde{\gamma} > 0$. Then there exists a neighborhood $B \ni \delta_{\omega}$ such that $\gamma_{\hat{\beta}}(\mu) < \tilde{\gamma}$ for all $\mu \in B$.

At any belief μ , let $V_{\mu}(\gamma)$ denote the agent's expected payoff to precision γ ; that is,

$$V_{\mu}(\gamma) = \left(\hat{\beta} + \gamma \mu \cdot \omega \right) v(\bar{\mu}(\gamma)) + \left(1 - \hat{\beta} - \gamma \mu \cdot \omega \right) v(\underline{\mu}(\gamma)), \quad (33)$$

where $\omega := (\omega_1, \dots, \omega_N)' \in \mathbb{R}^N$ and $\bar{\mu}(\gamma)$ and $\underline{\mu}(\gamma)$ denote the posteriors updated from μ under precision choice γ and perception $\hat{\beta}$ following signals 1 and 0, respectively. By (11), $\gamma_{\hat{\beta}}(\mu) \in \operatorname{argmax}_{\gamma \in [0, \bar{\gamma}]} V_{\mu}(\gamma) - C(\gamma)$ for all μ .

Since C is strictly increasing, $C(\tilde{\gamma}) > C(0)$. Thus, by continuity of v , there exists a neighborhood $B \ni \delta_{\omega}$ such that for each $\mu \in B$ and $\gamma \in \{0, \bar{\gamma}\}$,

$$|V_{\mu}(\gamma) - v(\delta_{\omega})| < \frac{C(\tilde{\gamma}) - C(0)}{2}. \quad (34)$$

Note that $V_{\mu}(\gamma)$ is increasing in γ for all μ . Thus, it follows that (34) holds for each $\mu \in B$ and $\gamma \in [0, \bar{\gamma}]$. This implies that for any $\gamma \in [0, \bar{\gamma}]$ and $\mu \in B$,

$$V_{\mu}(\gamma) - V_{\mu}(0) \leq |V_{\mu}(\gamma) - v(\delta_{\omega})| + |V_{\mu}(0) - v(\delta_{\omega})| < C(\tilde{\gamma}) - C(0). \quad (35)$$

Hence, for all $\gamma \geq \tilde{\gamma}$ and $\mu \in B$, we have

$$V_\mu(\gamma) - C(\gamma) \leq V_\mu(\tilde{\gamma}) - C(\tilde{\gamma}) < V_\mu(0) - C(0),$$

where the first inequality uses the fact that C is increasing and the second inequality uses (35). Thus, for all $\mu \in B$, we have $\gamma_{\hat{\beta}}(\mu) < \tilde{\gamma}$, as claimed. \square

C.6 Proof of Proposition 3

Fix any true state $\omega^* \in \Omega$ and consider any $\hat{\beta}$. The assumption that learning is successful at all states if $\hat{\beta} = \beta$ implies that for all mixed μ , we have $\gamma_{\hat{\beta}}(\mu) > 0$. Now suppose that $\beta < \hat{\beta}$; the argument for $\beta > \hat{\beta}$ is analogous.

Consider any $\omega \in \Omega$. By Lemma 5, there exists $B \ni \delta_\omega$ such that $\gamma_{\hat{\beta}}(\mu) < \frac{\hat{\beta} - \beta}{\omega_N - \omega_1}$ for all $\mu \in B$. Consider any $\omega', \omega'' \in \Omega$ with $\omega' < \omega''$. Then, for any $\mu \in B \setminus \{\delta_\omega\}$, we have $\beta + \gamma_{\hat{\beta}}(\mu)\omega^* < \hat{\beta} + \gamma_{\hat{\beta}}(\mu)\omega' < \hat{\beta} + \gamma_{\hat{\beta}}(\mu)\omega''$. By the same argument as in Example 1 (see footnote 19), this implies that for all $q \in (0, 1)$ and $\mu \in B \setminus \{\delta_\omega\}$, we have $\omega' \succ_\mu^q \omega''$. Note also that for each mixed μ , $\gamma_{\hat{\beta}}(\mu) > 0$ implies $\hat{p}_\mu(0|\omega') > \hat{p}_\mu(0|\omega'')$. Hence, Proposition 1 implies that δ_{ω_1} is globally stable. \square

C.7 Proof of Lemma 6

Fix any $\hat{\beta}$. We begin with some preliminary observations about the agent's expected value $V_\mu(\gamma)$ of precision γ at current belief μ , as given by (33). Note that the posteriors $\bar{\mu}(\gamma)$ and $\underline{\mu}(\gamma)$ of μ under signal realizations 1 and 0, respectively, assign probabilities

$$\bar{\mu}_n(\gamma) = \frac{\mu_n(\hat{\beta} + \gamma\omega_n)}{\hat{\beta} + \gamma\mu \cdot \omega}, \quad \underline{\mu}_n(\gamma) = \frac{\mu_n(1 - \hat{\beta} - \gamma\omega_n)}{1 - \hat{\beta} - \gamma\mu \cdot \omega},$$

to each state ω_n . The first and second derivatives with respect to γ satisfy

$$\begin{aligned} \bar{\mu}'_n(\gamma) &= \mu_n \frac{\hat{\beta}(\omega_n - \mu \cdot \omega)}{(\hat{\beta} + \gamma\mu \cdot \omega)^2}, & \underline{\mu}'_n(\gamma) &= -\mu_n \frac{(1 - \hat{\beta})(\omega_n - \mu \cdot \omega)}{(1 - \hat{\beta} - \gamma\mu \cdot \omega)^2}, \\ \bar{\mu}''_n(\gamma) &= -2\mu_n \mu \cdot \omega \frac{\hat{\beta}(\omega_n - \mu \cdot \omega)}{(\hat{\beta} + \gamma\mu \cdot \omega)^3}, & \underline{\mu}''_n(\gamma) &= -2\mu_n \mu \cdot \omega \frac{(1 - \hat{\beta})(\omega_n - \mu \cdot \omega)}{(1 - \hat{\beta} - \gamma\mu \cdot \omega)^3}. \end{aligned}$$

Thus, the marginal value of γ at μ satisfies

$$\begin{aligned} V'_\mu(\gamma) &= \mu \cdot \omega (v(\bar{\mu}(\gamma)) - v(\underline{\mu}(\gamma))) + (\hat{\beta} + \gamma\mu \cdot \omega) \sum_n \partial_n v(\bar{\mu}(\gamma)) \bar{\mu}'_n(\gamma) \\ &\quad + (1 - \hat{\beta} - \gamma\mu \cdot \omega) \sum_n \partial_n v(\underline{\mu}(\gamma)) \underline{\mu}'_n(\gamma), \end{aligned}$$

where $\partial_n v(\mu)$ denotes the partial derivative of v with respect to the n th coordinate. Since $\bar{\mu}(0) = \underline{\mu}(0) = \mu$ and $\hat{\beta}\bar{\mu}'_n(0) + (1 - \hat{\beta})\underline{\mu}'_n(0) = 0$ for each n , this yields

$$V'_\mu(0) = 0. \quad (36)$$

The second derivative satisfies

$$\begin{aligned} V''_\mu(\gamma) &= 2\mu \cdot \omega \sum_n \left(\partial_n v(\bar{\mu}(\gamma)) \bar{\mu}'_n(\gamma) - \partial_n v(\underline{\mu}(\gamma)) \underline{\mu}'_n(\gamma) \right) \\ &\quad + (\hat{\beta} + \gamma\mu \cdot \omega) \left(\sum_{n,m} \partial_{n,m}^2 v(\bar{\mu}(\gamma)) \bar{\mu}'_n(\gamma) \bar{\mu}'_m(\gamma) + \sum_n \partial_n v(\bar{\mu}(\gamma)) \bar{\mu}''_n(\gamma) \right) \\ &\quad + (1 - \hat{\beta} - \gamma\mu \cdot \omega) \left(\sum_{n,m} \partial_{n,m}^2 v(\underline{\mu}(\gamma)) \underline{\mu}'_n(\gamma) \underline{\mu}'_m(\gamma) + \sum_n \partial_n v(\underline{\mu}(\gamma)) \underline{\mu}''_n(\gamma) \right). \end{aligned}$$

Evaluating this at $\gamma = 0$ yields

$$V''_\mu(0) = \frac{1}{\hat{\beta}(1 - \hat{\beta})} \sum_{n,m} \partial_{n,m}^2 v(\mu) \mu_n (\omega_n - \mu \cdot \omega) \mu_m (\omega_m - \mu \cdot \omega). \quad (37)$$

To prove Lemma 6, consider any twice continuously differentiable C with $C'(0) = C''(0) = 0$. For any mixed μ , we have $V'_\mu(0) = 0 = C'(0)$ by (36), but $V''_\mu(0) > 0 = C''(0)$ by (37) and the fact that the Hessian of v is positive definite. Thus, by Taylor approximation,

$$V_\mu(\gamma) - C(\gamma) > V_\mu(0) - C(0)$$

for all sufficiently small $\gamma > 0$. Hence, for all mixed μ , $\gamma_{\hat{\beta}}(\mu) > 0$, as required.

For the ‘‘moreover’’ part of Lemma 6, it is clear that (12) is necessary for learning to be successful at all states ω^* when $\hat{\beta} = \beta$. To see that (12) is sufficient, fix any true state ω^* and suppose that $\hat{\beta} = \beta$. Then $P_\mu = \hat{P}_\mu(\cdot|\omega^*)$ for all μ , and by (12), $\hat{P}_\mu(\cdot|\omega^*) \neq \hat{P}_\mu(\cdot|\omega)$ for all $\omega \neq \omega^*$ and mixed μ . Thus, by Corollary 2, δ_{ω^*} is globally stable at ω^* in any state space $\Omega' \subseteq \Omega$ with $\omega^* \in \Omega'$. Hence, learning is successful at ω^* . \square

C.8 Proof of Lemma 7

Consider any true state $\omega^* \in \Omega$. Since $F = \hat{F}$, we have $P_\mu(\cdot) = \hat{P}_\mu(\cdot|\omega^*)$ for all μ . Moreover, for any mixed μ , the monotone likelihood ratio assumption on private signals ensures that $\frac{\mu^s(\omega)}{\mu^s(\omega')}$ is strictly increasing in s for any states $\omega > \omega'$ in $\text{supp}(\mu)$, which implies that $\theta^*(\mu^s)$ is strictly decreasing in s . Thus, for all mixed μ , $\hat{p}_\mu(0|\omega) = \int \hat{F}(\theta^*(\mu^s)) \phi(s|\omega) ds$ is strictly decreasing in ω , so that $\hat{P}_\mu(\cdot|\omega) \neq \hat{P}_\mu(\cdot|\omega^*)$ for all $\omega \neq \omega^*$. Hence, by Corollary 2, δ_{ω^*} is globally stable at ω^* in every state space $\Omega' \subseteq \Omega$ with $\omega^* \in \Omega'$. This shows that learning is successful at ω^* . \square

C.9 Proof of Lemma 8

Let $\Phi(\cdot|\omega)$ denote the cumulative distribution function of private signals conditional on ω . Since $\delta_\omega^s = \delta_\omega$ for each ω and s , the bounded convergence theorem implies that $\lim_{\mu \rightarrow \delta_\omega} \int \hat{F}(\theta^*(\mu^s)) d\Phi(s|\omega') = \hat{F}(\theta^*)$ for each ω, ω' , as claimed. \square

C.10 Proof of Proposition 4

We will invoke the following lemma:

Lemma 18. *Fix any true state $\omega^* \in \Omega$, $q \in (0, 1)$, and $n \in \{1, \dots, N\}$. If $F(\theta_n^*) > \hat{F}(\theta_n^*)$, then there exists a neighborhood $B_n \ni \delta_{\omega_n}$ such that $\omega_\ell \succ_\mu^q \omega_k$ for all ℓ, k with $\ell < k$ and all mixed $\mu \in B_n$. If $F(\theta_n^*) < \hat{F}(\theta_n^*)$, then there exists a neighborhood $B_n \ni \delta_{\omega_n}$ such that $\omega_k \succ_\mu^q \omega_\ell$ for all ℓ, k with $\ell < k$ and all mixed $\mu \in B_n$.*

Proof. Suppose $F(\theta_n^*) > \hat{F}(\theta_n^*)$; the argument when $F(\theta_n^*) < \hat{F}(\theta_n^*)$ is analogous. By Lemma 8, there exists a neighborhood $B_n \ni \delta_{\omega_n}$ such that for all $\mu \in B_n$ and ω' , we have $|p_\mu(0) - F(\theta_n^*)|, |\hat{p}_\mu(0|\omega') - \hat{F}(\theta_n^*)| < \frac{F(\theta_n^*) - \hat{F}(\theta_n^*)}{2}$. Hence, $p_\mu(0) > \hat{p}_\mu(0|\omega')$ for all $\mu \in B_n$ and ω' . Consider any ℓ, k with $\ell < k$. By the monotone likelihood ratio assumption on private signals, $\hat{p}_\mu(0|\omega_k) < \hat{p}_\mu(0|\omega_\ell)$ for all mixed μ . Thus, for any mixed $\mu \in B_n$, $p_\mu(0) > \hat{p}_\mu(0|\omega_\ell) > \hat{p}_\mu(0|\omega_k)$. By the same argument as in Example 1 (see footnote 19), this implies that for all $q \in (0, 1)$ and mixed $\mu \in B_n$, $\omega_\ell \succ_\mu^q \omega_k$, as claimed. \square

We now prove Proposition 4. Fix any $q \in (0, 1)$. For the first part, note that if $F(\theta_N^*) < \hat{F}(\theta_N^*)$, then Lemma 18 yields some neighborhood $B \ni \delta_{\omega_N}$ such that $\omega_N \succ_\mu^q \omega_k$ for all $k \neq N$ and mixed $\mu \in B$, while if $F(\theta_N^*) > \hat{F}(\theta_N^*)$, then Lemma 18 yields a neighborhood $B \ni \delta_{\omega_N}$ such that $\omega_1 \succ_\mu^q \omega_N$ for all mixed $\mu \in B$. Thus, by Theorem 1, δ_{ω_N} is locally stable in the former case and unstable in the latter. Finally, if $F(\theta_n^*) < \hat{F}(\theta_n^*)$ for each n , then Lemma 18 implies that for each n , there is a neighborhood $B_n \ni \delta_{\omega_n}$ such that $\omega_k \succ_\mu^q \omega_\ell$ for all $\ell > k$ and mixed $\mu \in B_n$. Moreover, $\hat{p}_\mu(1|\omega)$ is strictly increasing in ω by the monotone likelihood ratio assumption on private signals and the monotonicity of utilities. Hence, up to reversing indices of states, Proposition 1 implies δ_{ω_N} is globally stable. The arguments for the second part are analogous.

Finally, for the third part, note that if $F(\theta_n^*) \neq \hat{F}(\theta_n^*)$, then Lemma 18 implies that for some neighborhood $B_n \ni \delta_{\omega_n}$, we either have $\omega_1 \succ_\mu^q \omega_n$ for all mixed $\mu \in B_n$ or $\omega_N \succ_\mu^q \omega_n$ for all mixed $\mu \in B_n$. In either case, δ_{ω_n} is unstable by Theorem 1, as claimed. \square

D Additional results

D.1 A derivative condition for Theorem 1

Under slow learning, we provide a way to verify the conditions in Theorem 1 by only considering the derivatives of the difference in KL-prediction accuracy at the belief δ_ω . Let $\Delta(\Omega) - \Delta(\Omega) :=$

$\{\nu_1 - \nu_2 : \nu_1, \nu_2 \in \Delta(\Omega) \subseteq \mathbb{R}^{|\Omega|}\}$. Denote by $\nabla_m g(\mu)$ the directional derivative of $g : \Delta(\Omega) \rightarrow \mathbb{R}$ at μ in the direction of $m \in \Delta(\Omega) - \Delta(\Omega)$ whenever this is well-defined.

Corollary 3. *Assume slow learning holds. Suppose that, at $\mu = \delta_\omega$, $p_\mu(z)$ and $\frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)}$ admit ν -integrable directional derivatives for each $\omega' \neq \omega$ and ν -almost all z .*

1. *If at $\mu = \delta_\omega$, we have $\nabla_{\delta_{\omega''} - \delta_\omega} \int \log \frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)} dP_\mu(z) < 0$ for every $\omega', \omega'' \neq \omega$, then condition (6) in Theorem 1 holds, so δ_ω is locally stable.*
2. *If at $\mu = \delta_\omega$, we have $\nabla_{\delta_{\omega''} - \delta_\omega} \int \log \frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)} dP_\mu(z) > 0$ for some ω' and every $\omega'' \neq \omega$, then condition (7) in Theorem 1 holds, so δ_ω is unstable.*

To interpret, note that by slow learning, $\text{KL}(P_\mu, \hat{P}_\mu(\cdot|\omega)) - \text{KL}(P_\mu, \hat{P}_\mu(\cdot|\omega')) = \int \log \frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)} dP_\mu(z) = 0$ at $\mu = \delta_\omega$. The first condition ensures that for all μ close enough to δ_ω , $\text{KL}(P_\mu, \hat{P}_\mu(\cdot|\omega)) - \text{KL}(P_\mu, \hat{P}_\mu(\cdot|\omega')) < 0$, i.e., $\omega \succ_\mu^{\text{KL}} \omega'$, and that this difference has a first-order magnitude as $\mu \approx \delta_\omega$.

Proof. We only prove the first part; the second part is analogous. Fix any $\omega' \neq \omega$. Since Ω is finite, it suffices to find a neighborhood $B \ni \delta_\omega$ and $q > 0$ such that $\omega \succ_\mu^q \omega'$ for all $\mu \in B \setminus \{\delta_\omega\}$.

For each μ and $q > 0$, define

$$\gamma(\mu) := \int \log \frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)} p_\mu(z) d\nu(z), \quad \gamma^q(\mu) := \int \frac{\left(\frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)}\right)^q - 1}{q} p_\mu(z) d\nu(z).$$

We first show that $\lim_{q \rightarrow 0} \nabla_m \gamma^q(\mu) = \nabla_m \gamma(\mu)$ for all directions m . Denote $\ell(z|\mu) := \frac{\hat{p}_\mu(z|\omega')}{\hat{p}_\mu(z|\omega)}$. Then

$$\nabla_m \gamma^q(\mu) = \int \left(\nabla_m p_\mu(z) \frac{(\ell(z|\mu))^q - 1}{q} + p_\mu(z) (\ell(z|\mu))^{q-1} \nabla_m \ell(z|\mu) \right) d\nu(z).$$

As $q \rightarrow 0$, this converges to

$$\nabla_m \gamma(\mu) = \int \left(\nabla_m p_\mu(z) \log \ell(z|\mu) + p_\mu(z) (\ell(z|\mu))^{-1} \nabla_m \ell(z|\mu) \right) d\nu(z).$$

By assumption,

$$\max_{\mu \in \Delta(\Omega \setminus \{\omega\})} \nabla_{\mu - \delta_\omega} \gamma(\delta_\omega) = \max_{\omega'' \neq \omega} \nabla_{\delta_{\omega''} - \delta_\omega} \gamma(\delta_\omega) < 0.$$

Thus, by the above convergence, there exists $q > 0$ such that

$$\max_{\mu \in \Delta(\Omega \setminus \{\omega\})} \nabla_{\mu - \delta_\omega} \gamma^q(\delta_\omega) = \max_{\omega'' \neq \omega} \nabla_{\delta_{\omega''} - \delta_\omega} \gamma^q(\delta_\omega) < 0.$$

This implies that there exists a neighborhood $B \ni \delta_\omega$ such that for all $\mu \in B \setminus \{\delta_\omega\}$,

$$\gamma^q(\mu) < \gamma^q(\delta_\omega) = 0,$$

where the equality holds by slow learning. Thus, $\omega \succ_\mu^q \omega'$ for all $\mu \in B \setminus \{\delta_\omega\}$, as desired. \square

D.2 Details for Remark 2

The following example shows that one cannot replace q -dominance with KL-dominance in the local stability condition in Theorem 1. That is, condition (8) in Remark 2 does not ensure that δ_ω is locally stable:

Example 4. Let $\Omega = \{\omega, \omega'\}$ and $Z = \{\bar{z}, \underline{z}\}$. Set

$$p_\mu(\bar{z}) = \begin{cases} f(\log \frac{\mu(\omega)}{\mu(\omega')}) & \text{for all mixed } \mu \\ 1/2 & \text{otherwise,} \end{cases}$$

$$\hat{p}_\mu(\bar{z}|\omega) = \frac{e}{e+1}, \quad \hat{p}_\mu(\bar{z}|\omega') = \frac{1}{e+1} \text{ for all } \mu,$$

where $f : \mathbb{R} \rightarrow (0, 1)$ is any continuous function such that $f(x) = \frac{\sqrt{x} - \sqrt{x-1}}{\sqrt{x+1} - \sqrt{x-1}}$ for all $x \geq 1$, $f(x) > \frac{1}{2}$ for all $x < 1$, and $\lim_{x \rightarrow -\infty} f(x) = 1/2$. Note that $\lim_{x \rightarrow +\infty} f(x) = 1/2$, whence belief continuity holds. For each mixed μ , observe that

$$\sum_z p_\mu(z) \log \frac{\hat{p}_\mu(z|\omega)}{\hat{p}_\mu(z|\omega')} = f(\log \frac{\mu(\omega)}{\mu(\omega')}) \log e + \left(1 - f(\log \frac{\mu(\omega)}{\mu(\omega')})\right) \log \frac{1}{e} = 2f\left(\log \frac{\mu(\omega)}{\mu(\omega')}\right) - 1 > 0,$$

so $\omega \succ_\mu^{\text{KL}} \omega'$. Thus, condition (8) is satisfied.

However, δ_ω is unstable. To see this, fix any initial belief μ_0 and let $\ell_t := \sqrt{\log \frac{\mu_{\min\{t, \tau\}}(\omega)}{\mu_{\min\{t, \tau\}}(\omega')}}}$ where $\tau := \inf\{t : \log \frac{\mu_t(\omega)}{\mu_t(\omega')} < 1\}$. Then (ℓ_t) is a nonnegative martingale. This is because

$$\mathbb{E}[\ell_{t+1} | (\mu_s)_{s \leq t}] = \begin{cases} \ell_t & \text{if } \log \frac{\mu_{t'}(\omega)}{\mu_{t'}(\omega')} < 1 \text{ for some } t' \leq t \\ f(\log \frac{\mu_t(\omega)}{\mu_t(\omega')}) \sqrt{\log \frac{\mu_t(\omega)}{\mu_t(\omega')} + 1} + (1 - f(\log \frac{\mu_t(\omega)}{\mu_t(\omega')})) \sqrt{\log \frac{\mu_t(\omega)}{\mu_t(\omega')} - 1} = \sqrt{\log \frac{\mu_t(\omega)}{\mu_t(\omega')}} & \text{otherwise.} \end{cases}$$

Thus, by Doob's convergence theorem, there is an L^∞ random variable ℓ_∞ such that $\ell_t \rightarrow \ell_\infty$ a.s. Since, by construction, $\left| \log \frac{\mu_{t+1}(\omega)}{\mu_{t+1}(\omega')} - \log \frac{\mu_t(\omega)}{\mu_t(\omega')} \right| = 1$ for all t along all paths of signal realizations, there is probability zero that μ_t converges to a mixed belief. Thus, $\tau < \infty$ a.s. Hence, there a.s. exists some t such that $\log \frac{\mu_t(\omega)}{\mu_t(\omega')} < 1$. This implies that δ_ω is unstable.

Finally, observe that, even though $\omega \sim_{\delta_\omega}^{\text{KL}} \omega'$, this example does not feature slow learning. \blacktriangle