

Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Cowles Foundation Discussion Papers

Cowles Foundation

5-1-2020

Stability and Robustness in Misspecified Learning Models

Mira Frick

Ryota Iijima

Yuhta Ishii

Follow this and additional works at: <https://elischolar.library.yale.edu/cowles-discussion-paper-series>



Part of the [Economics Commons](#)

Recommended Citation

Frick, Mira; Iijima, Ryota; and Ishii, Yuhta, "Stability and Robustness in Misspecified Learning Models" (2020). *Cowles Foundation Discussion Papers*. 12.

<https://elischolar.library.yale.edu/cowles-discussion-paper-series/12>

This Discussion Paper is brought to you for free and open access by the Cowles Foundation at EliScholar – A Digital Platform for Scholarly Publishing at Yale. It has been accepted for inclusion in Cowles Foundation Discussion Papers by an authorized administrator of EliScholar – A Digital Platform for Scholarly Publishing at Yale. For more information, please contact elischolar@yale.edu.

STABILITY AND ROBUSTNESS IN MISSPECIFIED LEARNING MODELS

By

Mira Frick, Ryota Iijima, and Yuhta Ishii

May 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2235



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

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

Stability and Robustness in Misspecified Learning Models*

Mira Frick Ryota Iijima Yuhta Ishii

First version: October 1, 2019

This version: May 13, 2020

Abstract

We present an approach to analyze learning outcomes in a broad class of misspecified environments, spanning both single-agent and social learning. Our main results provide general criteria to determine—without the need to explicitly analyze learning dynamics—when beliefs in a given environment converge to some long-run belief either locally or globally (i.e., from some or all initial beliefs). The key ingredient underlying these criteria is a novel “prediction accuracy” ordering over subjective models that refines existing comparisons based on Kullback-Leibler divergence. We show that these criteria can be applied, first, to unify and generalize various convergence results in previously studied settings. Second, they enable us to identify and analyze a natural class of environments, including costly information acquisition and sequential social learning, where unlike most settings the literature has focused on so far, long-run beliefs can fail to be robust to the details of the true data generating process or agents’ perception thereof. In particular, even if agents learn the truth when they are correctly specified, vanishingly small amounts of misspecification can lead to extreme failures of learning.

*Frick: Yale University (mira.frick@yale.edu); Iijima: Yale University (ryota.iijima@yale.edu); Ishii: Pennsylvania State University (yishii1986@gmail.com). We thank Nageeb Ali, Aislinn Bohren, Christophe Chamley, In-Koo Cho, Ignacio Esponda, Erik Eyster, Drew Fudenberg, Kevin He, Marina Halac, Daniel Hauser, Paul Heidhues, Michihiro Kandori, Yves Le Yaouanq, Giacomo Lanzani, Weicheng Min, Pietro Ortoleva, Larry Samuelson, Philipp Strack, Yuichi Yamamoto, Leeat Yariv, as well as seminar audiences at Caltech, Harvard-MIT, Hitotsubashi, LMU Munich, Oxford, Princeton, QMUL, RINSE, 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 some aspects of the data generating process. Many papers have demonstrated how various specific 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, as standard properties of the correctly specified setting (e.g., the martingale property of beliefs) no longer apply when agents are misspecified. 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. Our main results provide general criteria to determine—without the need to explicitly analyze learning dynamics—when beliefs in a given environment converge to some long-run belief either locally or globally (i.e., from some or all initial beliefs). The key ingredient underlying these criteria is a novel “prediction accuracy” ordering over subjective models that refines existing comparisons based on Kullback-Leibler divergence. We show that these criteria can be applied, first, to unify and generalize various convergence results in previously studied settings. Second, they enable us to identify and analyze a natural class of environments, where unlike most settings the literature has focused on so far, long-run beliefs can fail to be robust to the details of the true data generating process or agents’ perception thereof. In particular, even if agents learn the truth when they are correctly specified, vanishingly small amounts of misspecification can lead to extreme failures of learning.

To nest a wide range of applications, Section 3 sets up an abstract framework, where agents, actions, and preferences are not explicitly modeled. Instead, we consider a belief

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](#)).

process μ_t over a finite 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 from a finite set according to a true signal distribution $P_{\mu_t}(\cdot)$ 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. Section 2 provides three simple economic examples, to which we apply our results in Section 6.

Sections 4 and 5 present criteria to determine which point-mass beliefs δ_ω are (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 .

Our criteria are based on an order over states that compares how well they predict the true signal distribution at any given belief. This order depends only on the relationship between the static primitives P_μ and \hat{P}_μ , and as such can be readily determined in a given misspecified learning environment without the need to compute belief dynamics. Specifically, we say that state ω *p-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 $P_\mu(\cdot)$ than does the perceived distribution $\hat{P}_\mu(\cdot|\omega')$ in state ω' . Importantly, instead of measuring closeness using Kullback-Leibler divergence, which features prominently in existing analyses of misspecified learning (e.g., [Berk, 1966](#); [Esponda and Pouzo, 2016](#); [Bohren and Hauser, 2018](#)), we employ a refinement based on the moment-generating function of the perceived log-likelihood ratio of states. As we will see, this refinement plays an essential role in our stability analysis, by ensuring that throughout any range of beliefs where *p*-dominance obtains, the (*p*th power of the) posterior ratio process becomes a nonnegative supermartingale, thus locally restoring standard methods from the correctly specified setting.

Based on this observation, Theorem 1 shows that δ_ω is locally stable if state ω strictly *p*-dominates all other states ω' at all beliefs μ in a neighborhood of δ_ω , *except* possibly at the belief $\mu = \delta_\omega$, where ω might be tied with some other states. Theorem 2 provides an analogous criterion for instability. As we discuss, the possibility of ties at δ_ω has the key implication that learning outcomes need not be robust to the details of the underlying environment, since slight perturbations to the true or perceived signal distributions can lead to discontinuous changes in the sets of locally stable and unstable beliefs.

While most of the literature so far has focused on settings that do not feature this possibility (we illustrate this for the monopoly pricing example in Section 2), being able to analyze environments where such failures of robustness might arise is important, as in

these environments the correctly specified benchmark may provide a poor approximation of learning outcomes under even vanishingly small amounts of misspecification. We exhibit a natural class of environments for which this possibility is relevant: These environments feature *non-identification at point-mass beliefs (NIP)*, where as beliefs grow confident in any given state, perceived signal distributions become uninformative. Examples 2 and 3 in Section 2 highlight two economic applications where agents’ behavior leads to NIP—costly information acquisition and sequential social learning. We show that vanishingly small amounts of misspecification can lead to extreme failures of learning in these settings, despite the fact that agents learn the true state when they are correctly specified.

Turning to global stability, we provide two complementary criteria. Theorem 3 shows that δ_ω is globally stable if state ω uniquely survives the iterated elimination of strictly dominated states, which we define analogously to the iterated elimination of dominated strategies in games. We use Theorem 3 to obtain unified proofs of global convergence in (discretizations of) several leading single-agent learning problems in the literature, including the aforementioned monopoly pricing example, effort choice by an overconfident agent (Heidhues, Koszegi, and Strack, 2018), and data censoring under the gambler’s fallacy (He, 2018). Theorem 4 presents an alternative criterion that is based on iteratively combining beliefs other than δ_ω into an unstable set, from which μ_t eventually escapes with probability 1. We show that this criterion can be applied to the costly information acquisition and sequential social learning examples above.

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). Section 3.3 adapts Berk-Nash equilibrium to our setting.² While it is known that any locally stable belief is a Berk-Nash equilibrium (Lemma 2 establishes this in our setting), the converse is typically not the case. Thus, our main contribution is to provide stability criteria that determine which Berk-Nash equilibria learning dynamics in a given environment converge to locally or globally, and to identify natural settings where the set of stable equilibria is not robust to the details of the environment. As we will see, our criteria refine Berk-Nash equilibrium by comparing the prediction accuracy of states at general beliefs μ , rather than only at equilibrium beliefs δ_ω , and by mea-

²Esponda and Pouzo (2016) consider a setting where multiple agents learn jointly about a payoff-relevant parameter and other agents’ behavior. Their setting nests single-agent learning, but does not nest sequential social learning.

asuring prediction accuracy using our notion of p -dominance, which is more demanding than the measure based on Kullback-Leibler divergence that underlies Berk-Nash equilibrium.

Several other papers also move beyond steady-state analysis and examine the convergence of misspecified learning dynamics in a variety of single-agent and social learning settings:

Single-agent learning. While many papers focus on specific environments and forms of misspecification (e.g., [Nyarko, 1991](#); [Heidhues, Koszegi, and Strack, 2018](#); [Bushong and Gagnon-Bartsch, 2019](#); [He, 2018](#); [Cong, 2019](#)), some recent work analyzes belief convergence in more general settings: [Fudenberg, Romanyuk, and Strack \(2017\)](#) consider a general continuous-time model with binary states and Gaussian signals. Their approach uses the fact that the belief process follows a one-dimensional stochastic differential equation. [Heidhues, Koszegi, and Strack \(2019\)](#) consider a continuous-state model with Gaussian prior and signals. Their approach is based on stochastic approximation arguments and relies on the fact that the posterior belief each period remains Gaussian. Closer to our paper, [Esponda, Pouzo, and Yamamoto \(2019\)](#) consider a general finite-action model, which like ours does not rely on any parametric assumptions. They show that, for large t , the time-average action frequency evolves according to a particular differential inclusion; in the special case where states are one-dimensional, they solve the differential inclusion under a unique identification assumption that rules out NIP. Complementary to their paper, we focus on beliefs rather than action frequencies, and we derive sufficient conditions for belief convergence that can be checked from static primitives without considering dynamics (such as their differential inclusion). In more recent work, [Fudenberg, Lanzani, and Strack \(2020\)](#) also study a general finite-action model and focus on action convergence instead of beliefs. They propose criteria for the stability of Berk-Nash equilibrium actions, partly building on the p -dominance method we introduce in this paper. In contrast to all aforementioned papers, we highlight that learning outcomes can fail to be robust in environments featuring non-identification, and we present results that are also applicable to social learning models.

Social learning. Several papers (e.g., [Eyster and Rabin, 2010](#); [Bohren, 2016](#); [Gagnon-Bartsch, 2017](#); [Bohren, Imas, and Rosenberg, 2019](#)) incorporate specific forms of misspecification into sequential social learning models à la [Bikhchandani, Hirshleifer, and Welch \(1992\)](#) and [Banerjee \(1992\)](#). Closer to our paper, [Bohren and Hauser \(2018\)](#) propose a general framework that unifies and extends many of these models, as well as certain single-agent settings.³ Relative to our paper, they focus more on the case of heterogeneous misspecifications, but only consider binary state environments. They provide criteria for local and

³Their framework can nest binary-state single-agent learning. Moving beyond binary states allows us to unify additional settings, for example those where convergence requires a sufficiently rich state space, as in [Example 1](#), [Heidhues, Koszegi, and Strack \(2018\)](#), [He \(2018\)](#).

global stability of beliefs that, like our criteria, can be checked without computing learning dynamics. Besides their focus on binary states, a key difference is that their criteria are based on Kullback-Leibler divergence and that their results and proof technique do not apply to environments featuring non-identification (see the discussion following Theorem 1). Under an assumption that rules out NIP, they obtain that successful learning is robust to small amounts of misspecification, in contrast with our Example 3 (see Remark 4). Our second global stability criterion (Theorem 4) extends ideas underlying their global stability result to a multi-state setting, while our first global stability criterion (Theorem 3) has no counterpart in their paper.

While our abstract model in Section 3, including the extension to profiles of beliefs in Appendix G, nests the single-agent and social learning settings above (subject to technical details such as finite vs. continuous states or signals), some environments in the literature are not nested, notably models with intertemporally correlated signals and social learning environments with private action observations.⁴ Within the former class of environments, [Cho and Kasa \(2017\)](#) consider single-agent learning under a Markovian fundamental and 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 is based on intertemporal correlation in the signal process.⁵ Our fragility results (Propositions 3-4) are complementary, by highlighting the discontinuity of long-run beliefs against misspecification *at* the correctly specified benchmark, based on a non-identification property (NIP) that is not present in their setting. The latter class of social learning environments includes our previous paper, [Frick, Iijima, and Ishii \(2019\)](#), which, similar to Example 3, highlights the fragility of successful learning against small amounts of misspecification about others’ preferences. However, as we discuss in Remark 4, both the logic and nature of this fragility result differs from the current paper, since the environment in [Frick, Iijima, and Ishii \(2019\)](#) does not display NIP.

2 Illustrative examples

To preview some of our insights, we present three simple economic examples, to which we will apply our results in Section 6. Remark 1 (Section 3) explains how our framework nests

⁴See, e.g., [Rabin \(2002\)](#); [Ortoleva and Snowberg \(2015\)](#); [Esponda and Pouzo \(2019\)](#); [Molavi \(2019\)](#) for the former, and [Dasaratha and He \(2019\)](#); [Levy and Razin \(2018\)](#) for the latter.

⁵In particular, the agent entertains two misspecified models, a constant-parameter (CP) and time-varying parameter (TVP) model, that both ignore the effect of her action on the evolution of the fundamental, but where the agent’s belief under the CP model becomes self-confirming in the long run if her prior is concentrated on that model. They show that even when the agent’s prior on the TVP model is arbitrarily small, her long-run belief can become confident in the TVP model with probability arbitrarily close to one.

these examples.

Example 1 (Monopoly pricing). In Section 6.1, we consider a monopolist who is learning about his demand function. Each period $t = 0, 1, \dots$, the monopolist first sets a price a_t , and then faces demand 1 with some probability $\omega^* - \beta a_t \in (0, 1)$ and demand 0 with complementary probability. The intercept of demand (“state”) $\omega^* \in \Omega = \{\omega_1, \dots, \omega_N\}$ is unknown to the monopolist, who has a full-support prior $\mu_0 \in \Delta(\Omega)$. Upon observing period- t demand, the monopolist updates his belief to μ_{t+1} . However, in so doing, the monopolist misperceives the slope of demand β to be $\hat{\beta}$, where $\beta, \hat{\beta} > 0$. The monopolist myopically maximizes expected revenue each period, i.e., his price as a function of his current belief is

$$a(\mu) = \operatorname{argmax}_{a \in \mathbb{R}_+} a \times \left(\mathbb{E}_\mu[\omega] - \hat{\beta}a \right) = \frac{\mathbb{E}_\mu[\omega]}{2\hat{\beta}}. \quad (1)$$

By applying our results, we will show that when Ω is sufficiently rich, the monopolist’s long-run belief under any prior is arbitrarily close to a point-mass on $\hat{\omega} = \frac{2\hat{\beta}\omega^*}{\hat{\beta} + \beta}$. State $\hat{\omega}$ has the property that at belief $\delta_{\hat{\omega}}$, the monopolist’s perceived probability of high demand, $\hat{\omega} - \hat{\beta}a(\delta_{\hat{\omega}})$, equals the actual probability, $\omega^* - \beta a(\delta_{\hat{\omega}})$. In contrast with [Esponda and Pouzo \(2016\)](#) and [Heidhues, Koszegi, and Strack \(2019\)](#), who establish analogous results using stochastic approximation arguments that rely on the assumption that beliefs are Gaussian, our approach does not require any distributional assumptions. We also note that successful learning is robust to small amounts of misspecification, since as $\hat{\beta}$ approximates β , $\hat{\omega}$ approximates the true state ω^* . As we will see, the reason behind this robustness is that the informativeness of observed demand about the true state ω^* is always bounded away from zero, regardless of the monopolist’s choice of price, and hence NIP does not hold. \blacktriangle

Our next two examples offer a sharp contrast to this robustness result:

Example 2 (Costly information acquisition). In Section 6.2, we consider an agent who learns about some fixed and unknown state (e.g., her ability) by acquiring costly information (e.g., seeking out expert feedback). The state $\omega^* \in \Omega = \{\omega_1, \dots, \omega_N\}$, where $0 < \omega_1 < \dots < \omega_N < 1$, and the agent has a full-support prior $\mu_0 \in \Delta(\Omega)$. Each period $t = 0, 1, \dots$, the agent observes the realization of a signal z_t that is 1 (“good news”) with some probability $q + \gamma_t \omega^* \in (0, 1)$ and 0 (“bad news”) with complementary probability. Here q is the state-independent base rate of the high signal over which the agent has no control, and $\gamma_t \in [0, \bar{\gamma}]$ is a precision parameter that the agent chooses at cost $C(\gamma_t)$. Upon observing the realized signal z_t , the agent updates her belief to μ_{t+1} . However, in so doing, she misperceives the base rate q to be \hat{q} . For example, if $\hat{q} < q$, this implies a form of “ego-biased” belief-updating,

where 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 $\gamma_t = 0$. We 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.⁶ Each period, she chooses γ_t as a function of her current belief μ_t to myopically maximize expected utility net of the cost. That is,

$$\gamma_t = \gamma(\mu_t) \in \operatorname{argmax}_{\gamma \in [0, \bar{\gamma}]} \hat{\mathbb{E}}_{\mu_t}[v(\mu_{t+1}(\gamma))] - C(\gamma), \quad (2)$$

where $\mu_{t+1}(\gamma)$ denotes the agent's random posterior following period- t signal realizations and the agent's expectation $\hat{\mathbb{E}}_{\mu_t}$ is with respect to her perceived signal distribution.

In Section 6.2, we first note that if information is costless (C is constant), then the agent's belief converges almost surely to a point-mass on the true state whenever \hat{q} is sufficiently close to q . By contrast, if information is costly, we show that successful learning is highly fragile against misspecification: Consider any strictly increasing C such that the agent learns the true state whenever she is correctly specified ($\hat{q} = q$).⁷ If $\hat{q} < q$ (resp. $\hat{q} > q$), we show that the agent's belief converges almost surely to a point-mass on the highest state ω_N (resp. lowest state ω_1) from all initial beliefs, *regardless* of the true state ω^* . Thus, in the presence of costless feedback, a small propensity for ego-biased interpretation of signals does not prevent the agent from learning her ability. But if obtaining feedback requires just a slight amount of effort, then even arbitrarily small amounts of this bias may be greatly amplified over time and lead to drastic overconfidence in the long run. We show that the key difference is that costly information acquisition leads to NIP: As the agent becomes increasingly confident in any given state, she chooses to acquire less and less precise signals, because her value to information vanishes. ▲

Example 3 (Sequential social learning). In Section 6.3, we consider social learning by a sequence of heterogeneous agents. There is a fixed and unknown state (e.g., the safety of a new product), $\omega^* \in \Omega = \{\omega_1, \dots, \omega_N\}$ with $\omega_1 < \dots < \omega_N$. Each period $t = 0, 1, \dots$, agent t chooses a one-shot action $z_t \in \{0, 1\}$ (e.g., whether or not to adopt the product) after observing a private signal $s_t \in \mathbb{R}$ about ω^* and the public sequence (z_0, \dots, z_{t-1}) of predecessors' actions. 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 .

⁶For 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).

⁷Lemma 7 clarifies under which conditions on C this is the case.

Starting with some full-support common prior $\mu_0 \in \Delta(\Omega)$, agent t chooses z_t to maximize her expected utility

$$\mathbb{E}_{\mu_t}[u(z_t, \theta_t, \omega) | \theta_t, s_t],$$

where μ_t denotes the Bayesian update of μ_0 based solely on the public action sequence (z_0, \dots, z_{t-1}) . However, in updating beliefs to μ_t , we assume that all agents misperceive the type distribution F in the population to be some cdf \hat{F} .⁸

Under standard monotonicity and richness assumptions, we first note that agents learn the true state when they are correctly specified ($\hat{F} = F$). However, by applying our results, we classify learning outcomes when $\hat{F} \neq F$, and show that successful learning is again highly non-robust to misspecification. For example, when agents even slightly underestimate (resp. overestimate) the extent of risk tolerance in the population, their beliefs converge almost surely to a point-mass on the highest (resp. lowest) safety level ω_N (resp. ω_1), no matter the true safety level ω^* ; and when agents underestimate the heterogeneity of risk attitudes, their beliefs may fail to converge and cycle between different safety levels. As we will see, the key feature behind this non-robustness is that a well-known feature of social learning again implies NIP: As previous action sequences become increasingly indicative of any particular state, agents put less and less weight on their private signals, so that new action observations become increasingly uninformative. ▲

3 Model

3.1 Setup

We conduct our general analysis in the following abstract, “reduced-form” 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 simplifies exposition by reducing notation.

Let Ω denote a finite set of **states**. At the beginning of each period $t = 0, 1, \dots$, there is a **belief** $\mu_t \in \Delta(\Omega)$ over states, where $\Delta(\Omega) := \{\mu \in \mathbb{R}_+^{|\Omega|} : \sum_{\omega} \mu(\omega) = 1\}$. The initial belief μ_0 is exogenous and has full support.⁹ The evolution of beliefs is determined as follows:

At the end of each period t , a signal z_t from a finite set of signals Z is drawn according to $P_{\mu_t}(\cdot)$, where $P_{\mu}(\cdot) \in \Delta(Z)$ denotes the **true signal distribution** at current belief μ . Upon observing signal z_t , belief μ_t is updated to belief μ_{t+1} via Bayes’ rule according to

⁸The extension of our framework in Section 7 allows different agents to hold different perceptions \hat{F} .

⁹The 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 .

a collection of conditional *perceived signal distributions*: At each current belief μ , the perceived signal distribution conditional on state ω is $\hat{P}_\mu(\cdot|\omega) \in \Delta(Z)$, and for all $\omega \in \Omega$, the updated belief satisfies

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

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 state ω^* , $P_\mu(\cdot) = \hat{P}_\mu(\cdot|\omega^*)$ for all μ , as the *correctly specified* benchmark.¹⁰

Given any initial belief $\mu_0 = \mu$, the true and perceived signal distributions jointly generate a Markov process over beliefs (μ_t) . Let \mathbb{P}_μ denote the induced probability measure over sequences of beliefs (μ_t) with $\mu_0 = \mu$. We impose the following regularity assumption:

Assumption 1.

1. For each ω and μ , $\text{supp} P_\mu(\cdot) \subseteq \text{supp} \hat{P}_\mu(\cdot|\omega)$.
2. There exists $M < \infty$ such that $\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega')} \leq M$ for all $\omega, \omega', \mu \in \Delta(\Omega)$ and $z \in \text{supp} P_\mu$.
3. For each ω , $P_\mu(\cdot)$ and $\hat{P}_\mu(\cdot|\omega)$ are continuous in μ .

The first condition is standard in the literature and rules out the possibility of belief-updating after a signal that is perceived to realize with zero probability. The second is a technical condition that rules out unbounded perceived likelihood ratios.¹¹ The third condition is not essential for our analysis, but simplifies the statements of our results.

Remark 1 (Examples). To illustrate how this abstract model maps to concrete economic applications, we first show how to nest the three simple examples from Section 2. In all three examples, $Z = \{0, 1\}$ and $\Omega \subseteq \mathbb{R}$, and the true signal distribution $P_\mu(\cdot) = P_\mu(\cdot|\omega^*)$ depends on some fixed and unknown true state $\omega^* \in \Omega$. Under monopoly pricing (Example 1), signals correspond to high or low demand realizations, and the true and perceived probabilities of high demand satisfy $P_\mu(1|\omega^*) = \omega^* - \beta a(\mu)$ and $\hat{P}_\mu(1|\omega) = \omega - \hat{\beta} a(\mu)$ for all μ and ω , where

¹⁰Here correct specification is defined only with respect to signals, as this abstract setting does not feature payoffs.

¹¹Given condition 3, this condition is automatically satisfied if, for example, P_μ and $\hat{P}_\mu(\cdot|\omega)$ have full support for all μ and ω , as is the case for the applications in Section 6.

the price $a(\mu)$ is given by (1). Likewise, under costly information acquisition (Example 2), the true and perceived probabilities of high signal realizations satisfy $P_\mu(1|\omega^*) = q + \gamma(\mu)\omega^*$ and $\hat{P}_\mu(1|\omega) = \hat{q} + \gamma(\mu)\omega$, where the precision $\gamma(\mu)$ is given by (2). Under sequential social learning (Example 3), signal z_t corresponds to agent t 's action and μ_t represents the *public* belief that is based only on the history (z_0, \dots, z_{t-1}) of past actions. Given μ_t and state ω , z_t is stochastic due to the random realization of agent t 's type θ_t and private signal s_t . Specifically, the true and perceived probabilities of action 0 satisfy

$$P_{\mu_t}(0|\omega^*) = \int F(\theta^*(\mu_t^s))\phi(s|\omega^*) ds, \quad \hat{P}_{\mu_t}(0|\omega) = \int \hat{F}(\theta^*(\mu_t^s))\phi(s|\omega) ds,$$

where $\phi(\cdot|\omega)$ is the density of private signals in state ω , $\mu_t^s \in \Delta(\Omega)$ denotes the Bayesian update of μ_t following private signal realization s , and for each $\nu \in \Delta(\Omega)$, $\theta^*(\nu)$ denotes the type who is indifferent between actions 0 and 1 at belief ν .¹²

Moving beyond Examples 1 and 2, the model nests any other single-agent active or passive learning setting where each period, the agent chooses an action as a function of her current belief (the agent's policy need not be myopic), observes a signal whose realization may depend on this action, and updates her belief based on a possibly misspecified model of the signal distribution.¹³ Action sets may be continuous or discrete; in the latter case, Assumption 1.3 can be met by assuming that the agent follows a stochastic choice rule that is continuous in her belief (e.g., by introducing payoff perturbations as in Fudenberg and Kreps, 1993; Esponda and Pouzo, 2016).¹⁴ Moving beyond Example 3, the model can incorporate other forms of homogeneous misspecification into any 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). Whereas Examples 1–3 feature misspecification in the form of a dogmatic belief in a particular parameter (e.g., $\hat{\beta}$, \hat{q} , \hat{F}), misspecification in the general model can take the form of a prior belief over a (finite) set of parameters that does not include the true parameter. Finally, Appendix G extends the model to allow for heterogeneous misspecifications across different agents, accommodating further applications (see Section 7). ▲

¹²The assumptions we impose in Section 6.3 ensure that $\theta^*(\nu)$ exists and is unique for each ν , and that agent t takes action 0 following private signal s if and only if $\theta_t < \theta^*(\mu_t^s)$.

¹³In addition to active learning, the dependence of perceived signal distributions on μ can also capture certain departures from Bayesian updating, where the agent's interpretation of signals depends on her current belief (e.g., Bushong and Gagnon-Bartsch, 2019).

¹⁴Given such a stochastic choice rule, we treat Z as the product space of realized signals and actions.

3.2 Stability notions

Given any true and perceived signal distributions, we seek to analyze which long-run beliefs μ^* can arise, in the sense that process (μ_t) converges to μ^* either locally or globally as a function of initial beliefs. Formally, we consider the following stability notions:

Definition 1. Belief $\mu^* \in \Delta(\Omega)$ is:

1. **locally stable** if for any $\gamma < 1$, there exists a neighborhood $B \ni \mu^*$ such that $\mathbb{P}_\mu[\mu_t \rightarrow \mu^*] \geq \gamma$ for each initial belief $\mu \in B$;¹⁵
2. **globally stable** if $\mathbb{P}_\mu[\mu_t \rightarrow \mu^*] = 1$ for each initial belief μ ;
3. **unstable** if there exists a neighborhood $B \ni \mu^*$ such that $\mathbb{P}_\mu[\exists t, \mu_t \notin B] = 1$ for each initial belief $\mu \in B \setminus \{\mu^*\}$.

Local stability requires that beliefs converge with positive probability to μ^* from any initial belief in some open set B around μ^* , where the probability of converging to μ^* can be made arbitrarily close to 1 as long as B is small enough. More strongly, global stability requires that beliefs converge to μ^* with probability 1 from *any* initial belief (recall that initial beliefs are assumed full-support). By contrast, μ^* is unstable if starting from any initial belief $\mu \neq \mu^*$ in some small enough neighborhood B of μ^* , beliefs eventually escape B with probability 1. Clearly, if μ^* is unstable, it is not locally stable.

We call belief $\mu \in \Delta(\Omega)$ a **point-mass** if $\mu = \delta_\omega$ assigns probability 1 to some state ω , and **mixed** otherwise. Our analysis in this paper focuses on the stability of point-mass beliefs. This is without loss of generality in environments that satisfy the following identification condition at mixed beliefs, as is the case for all applications we consider in this paper.¹⁶

Lemma 1. Consider any mixed $\mu \in \Delta(\Omega)$ such that $\hat{P}_\mu(z|\omega) \neq \hat{P}_\mu(z|\omega')$ for some $\omega, \omega' \in \text{supp}(\mu)$ and some $z \in \text{supp}P_\mu(\cdot)$. Then μ is unstable.

The condition requires that at the mixed belief μ , there is a possible signal realization that leads beliefs at μ (and hence near μ , by continuity of P and \hat{P}) to update in favor of one state in the support of μ rather than some other state. This implies that the belief process cannot settle down near μ .

¹⁵A **neighborhood** B of μ^* is a relatively open subset of $\Delta(\Omega) \subseteq \mathbb{R}^{|\Omega|}$ that contains μ^* .

¹⁶The condition can be violated in some active learning settings where the agent stops observing informative signals at some mixed belief (e.g., in McLennan, 1984, bandit problems, or under specifications of costly information acquisition that violate the condition in Lemma 7), as well as in social learning settings that feature herding or confounded learning.

3.3 Berk-Nash equilibrium

Before presenting our criteria for local and global stability of a belief δ_ω , we note a necessary condition for stability due to [Esponda and Pouzo \(2016\)](#). For any probability distributions $P, \hat{P} \in \Delta(Z)$, define the **Kullback-Leibler (KL-)divergence** of \hat{P} relative to P by $\text{KL}(P, \hat{P}) := \sum_{z \in Z} P(z) \log \frac{P(z)}{\hat{P}(z)}$.¹⁷ 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} .

Adapting [Esponda and Pouzo \(2016\)](#) to our setting, given any true and perceived signal distributions, we call belief δ_ω a **Berk-Nash equilibrium** if

$$\omega \in \underset{\omega' \in \Omega}{\operatorname{argmin}} \text{KL} \left(P_{\delta_\omega}(\cdot), \hat{P}_{\delta_\omega}(\cdot | \omega') \right). \quad (3)$$

Condition (3) is a fixed-point requirement, which says that at belief δ_ω , the perceived signal distribution that comes closest to the true signal distribution $P_{\delta_\omega}(\cdot)$ 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. Analogous to [Esponda and Pouzo \(2016\)](#), we show that this is a necessary condition for δ_ω to be locally stable:¹⁸

Lemma 2. *If δ_ω is not a Berk-Nash equilibrium, then δ_ω is unstable.*

However, while condition (3) is necessary for local stability, it is in general not sufficient. For instance, in the costly information acquisition setting of [Example 2](#), we will see that *all* point-mass beliefs δ_ω are Berk-Nash equilibria at each true state ω^* , despite the fact that at each ω^* there is a single globally stable belief. Thus, our stability criteria will take the form of refinements of Berk-Nash equilibrium.

4 Local stability, instability, and robustness

In this section, we present our criteria for local stability and instability of a belief δ_ω and point out that locally stable beliefs need not be robust to the details of the underlying environment.

¹⁷We use the convention that $\frac{0}{0} = 0$, $\frac{1}{0} = \infty$, $0 \log 0 = 0$, and $\log \infty = \infty$.

¹⁸[Esponda and Pouzo \(2016\)](#) allow for mixed Berk-Nash equilibria and show in their setting (cf. footnote 2) that if beliefs converge to μ^* with positive probability, then μ^* must be a Berk-Nash equilibrium belief (see their Lemma 2 and Theorem 2).

4.1 Prediction accuracy orders

Moving beyond the fixed-point condition (3) that underlies Berk-Nash equilibrium, our stability criteria are based on comparing the prediction accuracy of ω against other states ω' at *general* beliefs μ rather than only at belief δ_ω . However, like condition (3), these prediction accuracy comparisons are based only on the relationship between the “static” primitives P_μ and \hat{P}_μ , and do not require considering the dynamics of the belief process μ_t .

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

$$\text{KL} \left(P_\mu(\cdot), \hat{P}_\mu(\cdot|\omega) \right) - \text{KL} \left(P_\mu(\cdot), \hat{P}_\mu(\cdot|\omega') \right) := \sum_z P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega')}{\hat{P}_\mu(z|\omega)} \right) \leq 0. \quad (4)$$

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' . We write $\omega \succ_\mu^{\text{KL}} \omega'$ if inequality (4) is strict. Note that δ_ω is a Berk-Nash equilibrium if and only if $\omega \succ_{\delta_\omega}^{\text{KL}} \omega'$ for all $\omega' \neq \omega$.

Much of our stability analysis employs the following refinement of \succsim_μ^{KL} that plays an essential role in our proofs; see the discussion in Section 4.2. Given any $p > 0$, we say that ω **p -dominates** ω' **at** μ , denoted $\omega \succsim_\mu^p \omega'$, if

$$\sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega')}{\hat{P}_\mu(z|\omega)} \right)^p \leq 1, \quad (5)$$

and we write $\omega \succ_\mu^p \omega'$ if inequality (5) is strict.¹⁹ To see the connection between p -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 $P_\mu(\cdot)$. Then the left-hand side of (4) is the expectation of X , while the left-hand side of (5) is the moment-generating function $M_X(p) = \mathbb{E}[e^{pX}]$ of X evaluated at p .²⁰

Whereas \succsim_μ^{KL} is complete (by the representation on the LHS of (4)), \succsim_μ^p is in general incomplete. However, as the following lemma shows, the p -dominance orders are nested and approximate KL-dominance as $p \rightarrow 0$, in the sense that strict p -dominance implies strict KL-dominance and the converse holds for small p . This follows from the fact that $\mathbb{E}[X] = M'_X(0)$, $M_X(0) = 1$, and $M_X(p)$ is convex in p . We also note that both orders are continuous in μ .

¹⁹Note that $\omega \succ_\mu^p \omega'$ implies $\omega' \not\succsim_\mu^p \omega$. Indeed, by Lemma 3 below, $\omega \succ_\mu^p \omega'$ implies $\omega \succ_\mu^{\text{KL}} \omega'$; moreover, by the same argument, $\omega' \succsim_\mu^p \omega$ implies $\omega' \succsim_\mu^{\text{KL}} \omega$. Thus, the claim holds, as $\omega \succ_\mu^{\text{KL}} \omega'$ implies $\omega' \not\succsim_\mu^{\text{KL}} \omega$.

²⁰In the correctly specified setting, the moment-generating function of the log-likelihood ratio is also known as the Hellinger transform; see, e.g., Moscarini and Smith (2002); Deb and Ishii (2016) for applications.

Lemma 3. *Fix any belief μ and states ω, ω' .*

1. *If $\omega \succ_\mu^p \omega'$ for some $p > 0$, then $\omega \succ_\mu^{KL} \omega'$ and $\omega \succ_\mu^q \omega'$ for all $q \in (0, p)$.*
2. *If $\omega \succ_\mu^{KL} \omega'$, then there exists $p > 0$ such that $\omega \succ_\mu^p \omega'$.*
3. *The left-hand sides of (4) and (5) are continuous in μ .*

In the correctly specified benchmark, where for some ω^* , $P_\mu(\cdot) = \hat{P}_\mu(\cdot|\omega^*)$ for all μ , the state ω^* p -dominates all other states for any $p \in (0, 1]$. Indeed, Jensen's inequality implies that $\omega^* \succsim_\mu^p \omega$ for all μ and $\omega \neq \omega^*$; moreover, $\omega^* \succ_\mu^p \omega$ if $P_\mu(\cdot) \neq \hat{P}_\mu(\cdot|\omega)$ and $p \in (0, 1)$. Finally, we note that p -dominance bears some formal resemblance to a generalization of KL-divergence known as Rényi divergence. However, whereas KL-dominance amounts to comparing the KL-divergences $\text{KL}(P_\mu(\cdot), \hat{P}_\mu(\cdot|\omega))$ and $\text{KL}(P_\mu(\cdot), \hat{P}_\mu(\cdot|\omega'))$, p -dominance is not equivalent to comparing the corresponding Rényi divergences.²¹

4.2 Local stability and instability

Our first main result provides a sufficient condition for belief δ_ω to be locally stable:

Theorem 1. *Consider any $\omega \in \Omega$. Then belief δ_ω is locally stable if there exists $p > 0$ and a neighborhood $B \ni \delta_\omega$ such that*

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

That is, δ_ω is locally stable if for some $p > 0$, state ω strictly p -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.²² Thus, condition (6) strengthens Berk-Nash equilibrium, which requires that ω weakly KL-dominates all other states at the belief δ_ω itself, by comparing the prediction accuracy of ω against other states at beliefs in a neighborhood B of δ_ω and by imposing strict p -dominance rather than weak KL-dominance throughout $B \setminus \{\delta_\omega\}$.

Two features of Theorem 1 are important to note. First, the fact that (6) does not impose strict dominance at the belief δ_ω will play a key role in this paper: As we discuss in Section 4.3, this allows us to analyze settings in which locally stable beliefs can be highly sensitive to the details of the environment. This is ruled out by [Bohren \(2016\)](#) and [Bohren and Hauser's \(2018\)](#) more demanding (binary-state) criterion, whereby δ_ω is locally stable if ω strictly

²¹Formally, $\hat{P}_\mu(\cdot|\omega)$ displays lower p -Rényi divergence relative to $P_\mu(\cdot)$ than does $\hat{P}_\mu(\cdot|\omega')$ if $\sum_z P_\mu(z) \left(\frac{P_\mu(z)}{\hat{P}_\mu(z|\omega)} \right)^p \leq \sum_z P_\mu(z) \left(\frac{P_\mu(z)}{\hat{P}_\mu(z|\omega')} \right)^p$. This complete order implies (4) in the limit as $p \rightarrow 0$, but is not equivalent to the incomplete order (5) for a given $p > 0$.

²²The weak dominance $\omega \succsim_{\delta_\omega}^p \omega'$ follows from (6) and the continuity of \succsim_μ^p in μ (Lemma 3.3).

KL-dominates all other states at δ_ω : As Corollary 1 shows, Theorem 1 implies that beliefs satisfying the latter criterion are robustly locally stable, i.e., their local stability is preserved under small perturbations of the environment. The proof approach in the aforementioned papers is different from ours and does not extend to settings without strict dominance.²³

Second, the use of p -dominance, rather than KL-dominance, is essential in Theorem 1, as well as in several subsequent results. To see the idea, suppose that $\Omega = \{\omega, \omega'\}$ is binary. Starting at any $\mu_0 \in B$, the key is to consider the stopped process corresponding to the p th power of the posterior ratio of ω' vs. ω until the first time that beliefs exit B , i.e.,

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

and to note that (6) implies that ℓ_t is a nonnegative supermartingale with respect to \mathbb{P}_{μ_0} and the filtration generated by (μ_t) .²⁴ Thus, by Doob's convergence theorem, ℓ_t converges almost surely to a nonnegative random limit ℓ_∞ . Based on this, we show, first, 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 almost surely: Otherwise, the random limit belief $\mu_\infty \in B$ would be mixed with positive probability; this is impossible, as (6) together with Lemma 1 implies that all mixed beliefs in B are unstable. 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.

The role of p -dominance in the above argument is to ensure that the stopped process ℓ_t is a *nonnegative* supermartingale. Thus, locally, we are able to apply standard tools (e.g., Doob's convergence theorem), generalizing arguments in the correctly specified setting, where the unstopped process $\frac{\mu_t(\omega')}{\mu_t(\omega^*)}$ at the true state ω^* is a nonnegative martingale. Importantly, analogous arguments do not apply if p -dominance is replaced with KL-dominance: If (6) is weakened to the assumption that there exists some neighborhood $B \ni \delta_\omega$ such that

$$\omega \succ_\mu^{\text{KL}} \omega' \text{ for all } \omega' \neq \omega \text{ and } \mu \in B \setminus \{\delta_\omega\}, \quad (7)$$

²³Specifically, building on Smith and Sørensen (2000), they (locally) approximate 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. This approach 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 $\omega' \neq \omega$.

²⁴Indeed, $\mathbb{E}_{\mathbb{P}_{\mu_0}}[\ell_{t+1} | (\mu_{t'})_{t' \leq t}] = \begin{cases} \ell_t & \text{if } \mu_{t'} \notin B \text{ for some } t' \leq t \\ \sum_z P_{\mu_t}(z) \left(\frac{\hat{P}_{\mu_t}(z|\omega')}{\hat{P}_{\mu_t}(z|\omega)} \right)^p \ell_t \leq \ell_t & \text{otherwise} \end{cases}$, where the second line holds by (6).

this implies that the stopped process $\log \left(\frac{\mu_{\min\{t,\tau\}}(\omega')}{\mu_{\min\{t,\tau\}}(\omega)} \right)$ is a supermartingale.²⁵ But, crucially, this process may be *unbounded* below, since $\frac{\mu_{\min\{t,\tau\}}(\omega')}{\mu_{\min\{t,\tau\}}(\omega)}$ can be arbitrarily close to 0 as μ_t approaches δ_ω . Indeed, Example 5 in Appendix E shows that (7) does *not* imply that δ_ω is locally stable.

Our next result provides a sufficient condition for instability of δ_ω :

Theorem 2. *Consider any $\omega \in \Omega$. Then belief δ_ω is unstable if there exists a neighborhood $B \ni \delta_\omega$ such that*

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

Similar to the local stability condition (6), our instability criterion (8) requires that ω is strictly dominated by some other state ω' at all beliefs in a neighborhood of δ_ω , except possibly at belief δ_ω itself. However, in contrast with (6), condition (8) can be stated in terms of KL-dominance rather than p -dominance. This is because, unlike in the case of local stability, we can ensure that the process $\log \left(\frac{\mu_{\min\{t,\tau\}}(\omega)}{\mu_{\min\{t,\tau\}}(\omega')} \right)$, where τ is the exit time from B , is a supermartingale that is *bounded below*, by choosing the neighborhood B to be bounded away from $\delta_{\omega'}$. The proof of Theorem 2 then applies Doob's convergence theorem to this stopped process and shows that μ_t leaves B almost surely.

In Section 6.3, we will illustrate Theorems 1 and 2 by applying them to the social learning environment from Example 3.

4.3 Stability vs. robustness

The fact that Theorems 1 and 2 do not impose strict dominance at the point-mass belief δ_ω opens the door to analyzing settings where long-run beliefs can be highly sensitive to the details of an environment. To see this, we contrast our local stability criterion (6) with the stronger requirement that δ_ω is a *strict* Berk-Nash equilibrium, in the sense that $\omega \succ_{\delta_\omega}^{\text{KL}} \omega'$ for all $\omega' \neq \omega$. For such beliefs δ_ω , Theorem 1 implies that δ_ω is locally stable and that the local stability of δ_ω is robust to slightly perturbing the true or perceived signal distributions:

Corollary 1. *Suppose δ_ω is a strict Berk-Nash equilibrium at environment (P, \hat{P}) . Then there exists $\varepsilon > 0$ such that δ_ω is locally stable in every environment (Q, \hat{Q}) that is an ε -perturbation of (P, \hat{P}) , in the sense that $\text{KL}(Q_\mu, P_\mu) < \varepsilon$ and $\text{KL}(\hat{Q}_\mu(\cdot|\omega'), \hat{P}_\mu(\cdot|\omega')) < \varepsilon$ for all $\omega' \in \Omega$, $\mu \in \Delta(\Omega)$.*

²⁵By Lemma 3, (7) is equivalent to the requirement that for each $\mu \in B \setminus \{\delta_\omega\}$, there exists $p > 0$ such that $\omega \succ_\mu^p \omega'$ for all $\omega' \neq \omega$. This is weaker than (6), as p need not be uniform across $\mu \in B \setminus \{\delta_\omega\}$.

Corollary 1 holds because if $\omega \succ_{\delta_\omega}^{\text{KL}} \omega'$, then Lemma 3 yields some p such that $\omega \succ_\mu^p \omega'$ at all beliefs μ in a neighborhood of δ_ω , *including* at belief $\mu = \delta_\omega$, and these strict dominance relations are preserved under small enough perturbations of the environment. By contrast, this logic fails if (6) obtains without strict dominance at δ_ω , as in that case, even slight perturbations of the environment might reverse the prediction accuracy ranking between ω and ω' at beliefs near δ_ω , rendering δ_ω unstable. As a result, whereas much of the existing literature has focused on environments where locally stable beliefs are strict Berk-Nash equilibria and hence are robust,²⁶ Theorem 1 suggests the possibility of locally stable beliefs that are not robust.

In Sections 6.2 and 6.3, we illustrate this possibility in the context of two natural economic applications—costly information acquisition and social learning—by deriving the stark failures of robustness we previewed in Section 2. The common feature of both environments is that they display the following identification failure at point-mass beliefs:

Definition 2. Environment (P, \hat{P}) features non-identification at belief δ_ω if

$$\hat{P}_{\delta_\omega}(\cdot|\omega') = \hat{P}_{\delta_\omega}(\cdot|\omega''), \quad \forall \omega', \omega''.$$

More strongly, (P, \hat{P}) features *non-identification at point-mass beliefs (NIP)* if it features non-identification at δ_ω for all $\omega \in \Omega$.

By continuity of \hat{P} , non-identification at δ_ω is equivalent to the requirement that

$$\lim_{\mu \rightarrow \delta_\omega} \hat{P}_\mu(\cdot|\omega') = \lim_{\mu \rightarrow \delta_\omega} \hat{P}_\mu(\cdot|\omega''), \quad \forall \omega', \omega''.$$

That is, as beliefs grow confident in state ω , perceived signal distributions become increasingly uninformative, in the sense that they differ less and less across states, and they are fully uninformative at δ_ω . Clearly, this implies that all states have the *same* prediction accuracy at δ_ω , as the comparisons (4) and (5) between any two states hold with equality at δ_ω . In particular, under NIP, all point-mass beliefs δ_ω are Berk-Nash equilibria, but none are strict Berk-Nash equilibria, and thus Corollary 1 does not apply.

It is worth emphasizing that while NIP may appear “knife-edge” in the abstract class of *all* P and \hat{P} , in any concrete application, the relevant true and perceived signal distributions are determined by agents’ incentives. As we illustrate for costly information acquisition and social learning, there are natural economic environments in which NIP arises for a broad

²⁶See, e.g., Bohren (2016), Bohren and Hauser (2018), Fudenberg, Romanyuk, and Strack (2017), and Example 1, as well as (up to discretization) Heidhues, Koszegi, and Strack (2018) and He (2018). Bohren and Hauser (2018) explicitly establish the robustness of successful learning to small amounts of misspecification in their setting (see Remark 4).

range of underlying parameters, because agents' incentives to generate informative signals vanish near point-mass beliefs.

5 Global stability

Global stability is a significantly more demanding notion than local stability. For instance, even if δ_ω is the unique locally stable belief, it need not be globally stable (see Example 4 below). In this section, we present two criteria for global stability that strengthen the local stability criterion in Theorem 1 in complementary ways.

Our first approach employs a generalization of global stability to sets of beliefs: Call $K \subseteq \Delta(\Omega)$ a **globally stable set** if $\mathbb{P}_\mu[\inf_{\nu \in K} \|\mu_t - \nu\| \rightarrow 0] = 1$ 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. Formally, for each subset $\Omega' \subseteq \Omega$, let

$$\begin{aligned} S(\Omega') &:= \{\omega \in \Omega' : \nexists \omega' \in \Omega' \text{ s.t. } \omega' \succ_\mu^{\text{KL}} \omega \text{ for all } \mu \in \Delta(\Omega')\} \\ &= \{\omega \in \Omega' : \nexists \omega' \in \Omega' \text{ and } p > 0 \text{ s.t. } \omega' \succ_\mu^p \omega \text{ for all } \mu \in \Delta(\Omega')\}, \end{aligned}$$

where the equality holds by analogous arguments as in Lemma 3 and by compactness of $\Delta(\Omega')$.²⁷ Then recursively define $S^0(\Omega) := \Omega$, $S^{k+1}(\Omega) := S(S^k(\Omega))$ for all $k = 0, 1, \dots$, and $S^\infty(\Omega) := \bigcap_{k \in \mathbb{N}} S^k(\Omega)$.

Theorem 3. *The set $\Delta(S^\infty(\Omega))$ 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 3, 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 $\Delta(S(\Omega'))$. To see the idea, suppose that $S(\Omega') = \Omega' \setminus \{\omega'\}$. Then for some p and $\omega'' \in \Omega'$, we have $\omega'' \succ_\mu^p \omega'$ for all $\mu \in \Delta(\Omega')$, and hence also $\omega'' \succ_\mu^p \omega'$ for all μ in any small enough neighborhood $B \supseteq \Delta(\Omega')$.²⁸ Similar to Theorem 1, this implies that $\left(\frac{\mu_{\min\{t, \tau\}}(\omega')}{\mu_{\min\{t, \tau\}}(\omega'')}\right)^p$ with $\tau = \inf\{t' : \mu_{t'} \notin B\}$ is a nonnegative supermartingale; that from any initial $\mu \in B$, μ_t remains forever in B with positive probability; and that $\mu_t(\omega')$ converges to 0 almost surely 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 implies that $\Delta(\Omega' \setminus \{\omega'\})$ is globally stable. Note that even though the definition of iterated elimination does not feature any ties

²⁷See the proof of Theorem 3 for details.

²⁸Call B a **neighborhood of a set** $K \subseteq \Delta(\Omega)$ if there exists $\varepsilon > 0$ such that $\{\mu \in \Delta(\Omega) : \|\mu - \nu\| < \varepsilon\} \subseteq B$ for all $\nu \in K$.

in prediction accuracy, p -dominance again plays an important role in the proof of Theorem 3, by allowing us to employ martingale convergence arguments similar to Theorem 1.

In Section 6.1, we will apply Theorem 3 to the monopoly pricing example from Section 2. More generally, Appendix F highlights a class of continuous-state environments where iterated elimination of dominated states yields a unique outcome. The prediction accuracy orders in these environments display a form of complementarity or substitutability, paralleling conditions for dominance solvability in games with strategic complements or substitutes. In addition to the monopoly pricing example, these environments nest several leading single-agent settings in the literature, including Heidhues, Koszegi, and Strack (2018) and He (2018). Thus, we show that (up to discretization of the state space) Theorem 3 provides a simple and unified method to establish global stability in these settings.

The approach in Theorem 3 was to show that δ_ω is globally stable by employing a set-valued notion of global stability and iteratively considering a decreasing sequence of globally stable sets that contain δ_ω . Complementary to this, the idea behind our second global stability criterion is to employ a set-valued notion of instability and to ensure that the set $\Delta(\Omega \setminus \{\omega\})$ of beliefs that put zero probability on ω is unstable, by iteratively considering an increasing sequence of unstable sets contained in $\Delta(\Omega \setminus \{\omega\})$.

We call $K \subseteq \Delta(\Omega)$ an *unstable set* if there exists a neighborhood B of K such that $\mathbb{P}_\mu[\exists t, \mu_t \notin B] = 1$ for every initial belief $\mu \in B \setminus K$. As a preliminary result, we note that global stability holds under the following strengthening of our local stability criterion:

Lemma 4. *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 $z \in \text{supp} P_\mu(\cdot)$ 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.

Clearly, condition (ii) is necessary for global stability of δ_{ω_1} . Except when $N = 2$, condition (iii) is a stronger identification requirement than the condition for instability of mixed beliefs in Lemma 1, as for each mixed belief it postulates a possible signal realization that favors state ω_1 over *all* other states. However, the condition is still quite weak, in that it imposes no restrictions on the prediction accuracy order over states.

To see the idea behind Lemma 4, let neighborhoods $B_1 \ni \delta_{\omega_1}$ and $B_2 \supseteq \Delta(\{\omega_2, \dots, \omega_N\})$ be as given by conditions (i) and (ii). Then, first, from any initial belief in B_2 , μ_t escapes B_2 almost surely. Second, for any belief $\mu \notin B_1 \cup B_2$, condition (iii) yields a possible signal




Figure 1: Illustration of Example 4. By instability of δ_{ω_2} and δ_{ω_3} , there are neighborhoods $B_2 \ni \delta_{\omega_2}$ and $B_3 \ni \delta_{\omega_3}$ from which beliefs eventually escape a.s. By (i'), we can choose B_2 small enough that $\left(\frac{\mu_{\min\{t,\tau\}(\omega_3)}{\mu_{\min\{t,\tau\}(\omega_2)}}\right)^p$ with $\tau = \inf\{t' : \mu_{t'} \notin B_2\}$ is a supermartingale. For any small enough neighborhood $B_{23} \supseteq \Delta(\{\omega_2, \omega_3\})$, we then have the following: First, μ_t a.s. escapes $B_3 \cap B_{23}$. Second, as in Lemma 4, (iii') yields some finite T such that with positive probability, from any initial belief $\mu \in B_{23} \setminus (B_2 \cup B_3)$, μ_t either exits B_{23} or reaches B_2 within T periods. Finally, from any belief in $B_2 \cap B_{23}$, μ_t exits B_{23} with positive probability, as shown in the figure; otherwise μ_t would exit a.s. into $B_{23} \setminus B_2$, which contradicts the fact that $\left(\frac{\mu_{\min\{t,\tau\}(\omega_3)}{\mu_{\min\{t,\tau\}(\omega_2)}}\right)^p$ is a supermartingale. Combining these observations, we can show that μ_t exits B_{23} a.s.

realization z_μ that favors ω_1 over all other states, where $\inf_{\mu \notin B_1 \cup B_2} \frac{\hat{P}_\mu(z_\mu|\omega_1)}{\hat{P}_\mu(z_\mu|\omega_n)} > 1$ for all $n \neq 1$. Since $\inf_{\mu \notin B_1 \cup B_2} \mu(\omega_1) > 0$, by considering sufficiently long sequences of such signals, we can find a finite T such that with positive probability, μ_t reaches B_1 within T periods from any initial belief $\mu \notin B_1 \cup B_2$. Finally, whenever the belief process reaches B_1 , the same logic as in Theorem 1 implies that μ_t remains in B_1 with positive probability and conditional on remaining in B_1 converges to δ_{ω_1} almost surely. Combining these observations, we show that μ_t converges to δ_{ω_1} almost surely from any initial belief.

When $N = 2$, the logic above parallels global stability arguments in [Bohren \(2016\)](#), who establishes an analog of Lemma 4 (under her stronger local stability condition that requires strict KL-dominance at δ_{ω_1}). When $N > 2$, the challenge in applying Lemma 4 is to find a tractable way to verify condition (ii). The following example illustrates that it is not enough to verify that each of $\delta_{\omega_2}, \dots, \delta_{\omega_N}$ is individually unstable, and instead suggests additional conditions that ensure the instability of $\Delta(\{\omega_2, \dots, \omega_N\})$:²⁹

Example 4 (Three states). Suppose $\Omega = \{\omega_1, \omega_2, \omega_3\}$ satisfies conditions (i) and (iii) in Lemma 4 and suppose that δ_{ω_2} and δ_{ω_3} are unstable. These conditions alone do not guarantee that $\Delta(\{\omega_2, \omega_3\})$ is unstable. For example, they are consistent with the possibility that ω_1 has strictly lower prediction accuracy than both ω_2 and ω_3 at all beliefs μ that put sufficiently

²⁹[Bohren and Hauser \(2018\)](#) face related challenges in extending the global stability arguments in [Bohren \(2016\)](#) in a different direction, continuing to assume $N = 2$ but allowing for profiles of heterogeneous beliefs. In our extension to heterogeneous beliefs (Appendix G), we do not pursue a generalization of Theorem 4.

small weight on ω_1 , in which case for any small enough neighborhood B_{23} of $\Delta(\{\omega_2, \omega_3\})$, there is positive probability that μ_t remains forever in B_{23} .³⁰

However, suppose in addition that the logic in Lemma 4 applies to the comparison between ω_2 vs. ω_3 ; that is,

- (i') there exists $p > 0$ and a neighborhood $B \ni \delta_{\omega_2}$ with $\omega_2 \succ_{\mu}^p \omega_3$ for all $\mu \in B \setminus \{\delta_{\omega_2}\}$;
- (iii') for each mixed $\mu \in \Delta(\{\omega_2, \omega_3\})$, there is $z \in \text{supp} P_{\mu}$ with $\hat{P}_{\mu}(z|\omega_2) > \hat{P}_{\mu}(z|\omega_3)$.

Then Theorem 4 below shows that $\Delta(\{\omega_2, \omega_3\})$ is unstable. Figure 1 illustrates the basic idea. ▲

Theorem 4 generalizes the logic in Example 4 to arbitrary N . Suppose conditions (i) and (iii) from Lemma 4 hold and each $\delta_{\omega_2}, \dots, \delta_{\omega_N}$ is unstable. Moreover, impose analogs of conditions (i') and (iii') above to the comparison between ω_{N-1} vs. ω_N , between ω_{N-2} vs. $\{\omega_{N-1}, \omega_N\}$, ..., and between ω_2 vs. $\{\omega_3, \dots, \omega_N\}$. Then iteratively applying the three-state logic above ensures that $\Delta(\{\omega_n, \dots, \omega_N\})$ is unstable for all $n \neq 1$. Thus, δ_{ω_1} is globally stable by Lemma 4. The result reduces to Lemma 4 when Ω is binary.

Theorem 4. *Consider $\Omega = \{\omega_1, \dots, \omega_N\}$ and suppose that*

- (i) *for all $n \neq N$, there exists $p > 0$ and a neighborhood $B \ni \delta_{\omega_n}$ with*

$$\omega_n \succ_{\mu}^p \omega_k \text{ for all } k > n \text{ and } \mu \in B \setminus \{\delta_{\omega_n}\};$$

- (ii) *for all $n \neq 1$, δ_{ω_n} is unstable;*

- (iii) *for all $n \neq N$ and mixed $\mu \in \Delta(\{\omega_n, \dots, \omega_N\})$, there is $z \in \text{supp} P_{\mu}(\cdot)$ with $\hat{P}_{\mu}(z|\omega_n) > \hat{P}_{\mu}(z|\omega_k)$ for all $k > n$.*

Then δ_{ω_1} is globally stable.

Theorem 4 applies to any environment with the following order structure that is natural in several economic applications; in particular, as we will see in Sections 6.2 and 6.3, this is satisfied by the costly information acquisition and sequential social learning environments from Section 2:

³⁰If $\omega_2, \omega_3 \succ_{\mu}^p \omega_1$ whenever $\mu(\omega_1) < \varepsilon$, then for any small enough neighborhood B_{23} of $\Delta(\{\omega_2, \omega_3\})$, $\left(\frac{\mu_{\min\{t, \tau\}}(\omega_1)}{\mu_{\min\{t, \tau\}}(\omega_2)}\right)^p + \left(\frac{\mu_{\min\{t, \tau\}}(\omega_1)}{\mu_{\min\{t, \tau\}}(\omega_3)}\right)^p$ with $\tau = \inf\{t' : \mu_{t'} \notin B_{23}\}$ is a supermartingale. Similar to Theorem 1 this implies that from any initial $\mu \in B_{23}$, μ_t remains in B_{23} forever with positive probability. It is worth noting that even though, as in Lemma 4, condition (iii) ensures that from each mixed belief μ , there is a time T_{μ} by which μ_t reaches B_1 with positive probability, T_{μ} is unbounded across B_{23} , as $\inf_{\mu \in B_{23}} \mu(\omega_1) = 0$.

Remark 2. Suppose that both states and signals can be ordered in such a way that (a) perceived signal distributions $\hat{P}_\mu(\cdot|\omega)$ are strictly FOSD-increasing in ω ; and (b) near point-mass beliefs, lower states provide better prediction accuracy than higher states, i.e., for each ω , there exists a neighborhood $B \ni \delta_\omega$ and p such that $\omega' < \omega''$ implies $\omega' \succ_\mu^p \omega''$ for all $\mu \in B \setminus \{\delta_\omega\}$. Enumerating states in increasing order, property (a) implies that the lowest signal satisfies condition (iii) in Theorem 4, while (b) implies conditions (i) and (ii). \blacktriangle

Note that in binary state spaces, Theorem 3 is a strictly more demanding criterion than Theorem 4. However, in larger state spaces the two results are complementary: On the one hand, Theorem 4 only restricts the prediction accuracy order near point-mass beliefs, while Theorem 3 also restricts the order at interior beliefs (e.g., for Theorem 3 to have bite, at least one state must be dominated at all beliefs). But on the other hand, the specific form of restrictions imposed by Theorem 4 near point-mass beliefs (i.e., that ω_n dominates ω_k for all $k > n$) is not needed in order to apply Theorem 3.

Finally, while our definition of global stability requires convergence to δ_ω with probability *one* from all initial beliefs, one might also be interested in the notion of *weak* global stability, which only requires that from all initial beliefs, process μ_t converge to δ_ω with *positive* probability. Modifying the arguments above, one can show that conditions (i) and (iii) in Lemma 4 are sufficient to ensure that δ_{ω_1} is weakly globally stable.³¹ Unlike our criteria for strong global stability, the latter conditions allow for the possibility of multiple weakly globally stable beliefs.

6 Applications

To illustrate the preceding results, we apply them to the three illustrative examples from Section 2. In these examples, the true signal distribution $P_\mu(\cdot) = P_\mu(\cdot|\omega^*)$ depends on some fixed and unknown true state $\omega^* \in \Omega$, and accordingly, we index our stability results by ω^* .

6.1 Monopoly pricing

We first revisit the monopoly pricing problem from Example 1. Throughout, we fix some $\underline{\omega}, \bar{\omega} \in (0, 1)$ with $\underline{\omega} < \bar{\omega} < 2\underline{\omega}$ and consider state spaces $\Omega = \{\omega_1, \dots, \omega_N\}$ with $\underline{\omega} \leq \omega_1 < \dots < \omega_N \leq \bar{\omega}$. We also assume that $\frac{2\hat{\beta}}{\beta} > \frac{\bar{\omega}}{\underline{\omega}}$. This ensures that the true and perceived probabilities of high demand, $\omega^* - \beta a(\mu)$ and $\omega - \hat{\beta} a(\mu)$, are in $(0, 1)$ for all states and beliefs.

As noted, Esponda and Pouzo (2016) and Heidhues, Koszegi, and Strack (2019) establish global stability for variants of this problem in continuous state space settings with Gaussian

³¹This again parallels arguments in Bohren (2016) when $N = 2$.

beliefs and demand. In our finite state space setting, a simple application of Theorem 3 yields the following discrete approximation of global stability, without relying on distributional assumptions: Let $\hat{\omega} = \frac{2\hat{\beta}\omega^*}{\hat{\beta}+\beta}$, which is the unique solution to the requirement that at $\delta_{\hat{\omega}}$, the true and perceived probabilities of high demand coincide. Then if $\underline{\omega} < \hat{\omega} < \bar{\omega}$ and states in Ω are sufficiently dense in $[\underline{\omega}, \bar{\omega}]$, an arbitrarily small neighborhood around $\delta_{\hat{\omega}}$ is globally stable.³² To state the result formally, we say that Ω is δ -**dense** in $[\underline{\omega}, \bar{\omega}]$ if $\Omega \cap (\omega - \delta, \omega + \delta) \neq \emptyset$ for all $\omega \in [\underline{\omega}, \bar{\omega}]$.

Proposition 1. *Suppose $\underline{\omega} < \hat{\omega} := \frac{2\hat{\beta}\omega^*}{\hat{\beta}+\beta} < \bar{\omega}$. For any $\eta > 0$, there exists $\delta \in (0, \eta)$ such that if Ω is δ -dense in $[\underline{\omega}, \bar{\omega}]$, then $\Delta(\Omega \cap (\hat{\omega} - \eta, \hat{\omega} + \eta))$ is globally stable at ω^* .*

The proof of Proposition 1 first shows that iterated elimination of dominated states in the continuous state space $\bar{\Omega} = [\underline{\omega}, \bar{\omega}]$ yields $S^\infty(\bar{\Omega}) = \{\hat{\omega}\}$, and then establishes that $S^\infty(\Omega)$ approximates $S^\infty(\bar{\Omega})$ arbitrarily closely whenever Ω is sufficiently dense in $\bar{\Omega}$. To see why $S^\infty(\bar{\Omega}) = \{\hat{\omega}\}$, let $M(\omega)$ denote the $\succ_{\delta_\omega}^{\text{KL}}$ -maximal state in $\bar{\Omega}$, which is given by $M(\omega) := \min\{\max\{\underline{\omega}, \omega^* + \frac{(\hat{\beta}-\beta)\omega}{2\hat{\beta}}\}, \bar{\omega}\}$ for each ω . Suppose $\hat{\beta} \geq \beta$, so that $M(\cdot)$ is weakly increasing (the case $\hat{\beta} < \beta$ is analogous). Then one can verify that iterating $S(\cdot)$ over $\bar{\Omega}$ corresponds to iterated application of $M(\cdot)$ to the end-points of $\bar{\Omega}$, i.e.,

$$S^k(\bar{\Omega}) = [M^k(\underline{\omega}), M^k(\bar{\omega})] \text{ for all } k.$$

Given this, $S^\infty(\Omega) = \{\hat{\omega}\}$ follows by observing that $M(\cdot)$ is a contraction with fixed point $\hat{\omega}$.

Observe that state $\hat{\omega} = \frac{2\hat{\beta}\omega^*}{\hat{\beta}+\beta}$ converges to the true state ω^* as $\hat{\beta}$ approaches β . Thus, Proposition 1 suggests that the monopolist comes arbitrarily close to learning the true state whenever his amount of misperception is sufficiently small and states in Ω are sufficiently finely spaced. Indeed, the following result shows that in any fixed finite state space, the monopolist exactly learns the true state whenever his amount of misperception is sufficiently small. To state this, we say that **learning is successful** at ω^* if in state ω^* , we have $\mathbb{P}_\mu[\mu_t \rightarrow \delta_{\omega^*}] = 1$ for all beliefs $\mu \in \Delta(\Omega)$ with $\mu(\omega^*) > 0$.³³

Proposition 2. *Fix any $\Omega = \{\omega_1, \dots, \omega_N\}$ and β . Then there exists $\varepsilon > 0$ such that for any $\hat{\beta}$ with $|\hat{\beta} - \beta| < \varepsilon$, learning is successful at all states ω^* .*

Note that δ_{ω^*} is a strict Berk-Nash equilibrium when $\hat{\beta} = \beta$. Thus, by Corollary 1, δ_{ω^*} is locally stable under sufficiently small amounts of misspecification. More strongly, we show

³²Nyarko (1991) and Fudenberg, Romanyuk, and Strack (2017) study versions of this problem in binary state spaces and show that all beliefs are unstable when the amount of misspecification is large. In our setting, this can be established using the instability criterion in Theorem 2.

³³This is slightly stronger than the requirement that δ_{ω^*} is globally stable at ω^* , which only considers initial full-support beliefs.

that if $|\hat{\beta} - \beta|$ is small enough, then ω^* strictly KL-dominates all other states at *all* beliefs $\mu \in \Delta(\Omega)$ and hence, by Theorem 3, δ_{ω^*} is globally stable in any state space $\Omega' \subseteq \Omega$ that contains ω^* . Thus, successful learning is robust to small amounts of misspecification.

6.2 Costly information acquisition

Next, we consider the costly information acquisition problem from Example 2 and derive a stark failure of robustness when information is costly. Given any cost function C , we index the agent's precision choice $\gamma_{\hat{q}}(\mu)$ given by (2) by the perceived base rate \hat{q} . While (2) assumes for simplicity that γ is chosen myopically, our results in this section generalize to the case of a forward-looking agent who maximizes expected discounted payoffs.³⁴ Assume $\bar{\gamma} \in (0, 1)$ and $q, \hat{q} \in (0, 1 - \bar{\gamma})$, so that the true and perceived signal probabilities $q + \gamma_{\hat{q}}(\mu)\omega$ and $\hat{q} + \gamma_{\hat{q}}(\mu)\omega$ are well-defined and nondegenerate for all μ and ω . To satisfy Assumption 1.3, we also assume that $\gamma_{\hat{q}}(\mu)$ is continuous in μ .³⁵

As a benchmark, suppose first that the agent incurs the same constant cost $C(\gamma) = c$ regardless of her choice of precision, so that information is effectively costless. Then, analogous to Proposition 2, learning is successful when the agent is correctly specified ($\hat{q} = q$) and successful learning is robust to small amounts of misspecification:

Lemma 5. *Suppose C is constant. For any q , there exists $\varepsilon > 0$ such that for any \hat{q} with $|\hat{q} - q| \leq \varepsilon$, learning is successful at all states ω^* .*

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

Next, we contrast Lemma 5 with the case where information is costly, in the sense that C is strictly increasing in γ . The key departure this introduces is the following:

Lemma 6. *Suppose C is strictly increasing. For any \hat{q} , $\lim_{\mu \rightarrow \delta_{\omega}} \gamma_{\hat{q}}(\mu) = 0$ 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

³⁴To see this, note that NIP/Lemma 6 remains valid with the same proof, as the continuation value function is continuous in μ .

³⁵Without continuity, the main result (Proposition 3) remains valid up to replacing the current assumption ("successful learning at all states when $\hat{q} = q$ ") with the following assumption: for any compact set $K \subseteq \Delta(\Omega)$ of mixed beliefs, $\inf_{\mu \in K} \gamma_{\hat{q}}(\mu) > 0$. This is slightly stronger than the current assumption, which is equivalent to the requirement that $\gamma_{\hat{q}}(\mu) > 0$ for all mixed μ (Lemma 7). The robustness of costless learning (Lemma 5) does not rely on continuity.




Figure 2: Prediction accuracy ranking of ω_1 vs. ω_2 as a function of μ when $\omega^* = \omega_1$. Left-hand side: $\hat{q} = q$. Right-hand side: $q > \hat{q}$. Here $D_\mu^{\text{KL}}(\omega_2, \omega_1) := \text{KL}(P_\mu(\cdot|\omega^*), \hat{P}_\mu(\cdot|\omega_2)) - \text{KL}(P_\mu(\cdot|\omega^*), \hat{P}_\mu(\cdot|\omega_1))$.

information vanishes as she becomes confident. In the language of Section 4, Lemma 6 shows that costly information leads to NIP, since the agent's perceived signal distribution satisfies

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

As a result of NIP, the following proposition establishes that learning outcomes under costly information are highly sensitive to small amounts of misspecification: Suppose learning is successful whenever the agent is correctly specified ($\hat{q} = q$). Then, in sharp contrast with Lemma 5, introducing arbitrarily small amounts of misspecification $\hat{q} \neq q$ not only breaks successful learning, but indeed renders the agent's long-run belief *independent* of the true state ω^* : If $\hat{q} < q$, then regardless of ω^* , she becomes confident in the highest possible state ω_N ; if $\hat{q} > q$, then regardless of ω^* , she becomes confident in the lowest possible state ω_1 .

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

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

To see the idea, suppose that $\Omega = \{\omega_1, \omega_2\}$ and the true state is ω_1 . For any \hat{q} , the fact that learning is successful at all states when $q = \hat{q}$ means that $\gamma_{\hat{q}}(\mu) > 0$ for all mixed μ , as otherwise the agent's belief would get stuck at some initial mixed beliefs. At the same time, by Lemma 6, $\gamma_{\hat{q}}(\delta_\omega) = 0$ at all point-mass beliefs. As a result, the prediction accuracy ranking of ω_1 vs. ω_2 is sensitive to the relationship between q and \hat{q} .

Specifically, as shown in Figure 2, when $q = \hat{q}$, the true state ω_1 strictly dominates ω_2 at all mixed beliefs, even though the gap in prediction accuracy vanishes as beliefs approach δ_{ω_1} or δ_{ω_2} .³⁶ By contrast, if $q > \hat{q}$, then the ranking is reversed near point-mass beliefs: Indeed,

³⁶We use KL-dominance in Figure 2 for the sake of graphical illustration, but the proof of Proposition 3 relies on establishing p -dominance (for any $p \in (0, 1)$).

since γ is very small near point-mass beliefs, the true probability $\gamma\omega_1 + q$ of the high signal exceeds the perceived probabilities $\gamma\omega_2 + \hat{q}$, $\gamma\omega_1 + \hat{q}$ in both states, but because $\omega_2 > \omega_1$, the perceived probability in state ω_2 comes closer to the truth. Intuitively, when signals are very precise (γ is high), the true state always better explains the agent's observations, but when signals are sufficiently imprecise (γ is low), then overestimating the state can partly compensate for the fact that the agent underestimates the base rate of the high signal. Finally, by Remark 2, the fact that ω_2 strictly dominates ω_1 near both point-mass beliefs and that signals are FOSD-increasing in states means that Theorem 4 applies up to relabeling states in decreasing order. Thus, when $q > \hat{q}$, δ_{ω_2} is globally stable.

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 by imposing additional regularity assumptions:

Lemma 7. *Suppose v is twice continuously differentiable and admits a positive-definite Hessian. Fix any \hat{q} . For any twice continuously differentiable cost function C with $C'(0) = C''(0) = 0$, we have*

$$\gamma_{\hat{q}}(\mu) > 0 \text{ for all mixed } \mu. \quad (9)$$

Moreover, (9) is necessary and sufficient for learning to be successful at all ω^* when $q = \hat{q}$.

Lemma 7 provides “Inada” conditions on C which ensure that small amounts of information are very cheap, so that 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 $d > 2$.³⁷

Remark 3. We comment on two features of the present setting:

Persistence of misspecification. As in several misspecified learning environments in the literature, the true long-run signal distribution disagrees with the agent's perceived long-run distribution whenever $\hat{q} \neq q$, suggesting that (asymptotically) the agent might come to realize that she is misspecified.³⁸ However, this feature is not essential to the present fragility

³⁷The 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 . As is well-known, the Radner-Stiglitz non-concavity does not arise in Gaussian environments. In such settings, we can show (details available upon request) that $C'(0) = 0$ is sufficient to guarantee that learning is successful under correct specification, while the analog of Proposition 3 continues to hold. In the context of voting with costly information acquisition (but unrelated to the Radner-Stiglitz non-concavity), Martinelli (2006) also imposes the condition $C'(0) = C''(0) = 0$ to show that information aggregation obtains.

³⁸However, see Gagnon-Bartsch, Rabin, and Schwartzstein (2018) for an inattention-based rationale for the persistence of such misspecification.

result: For example, an analog of Proposition 3 can be obtained if the agent is correct about the base rate q but (even slightly) misperceives the sensitivity of the signal distribution to the agent's choice γ ; in this case the true and perceived long-run signal distributions exactly coincide (by NIP), suggesting that such misspecification might persist.

Approximate NIP. Various forces might break NIP in the present setting, for example, if the agent additionally receives some exogenous costless information each period. In that case, how sensitive learning outcomes are to misspecification depends on the amount of exogenous information. Specifically, for any $\hat{q} \neq q$, the predictions of Proposition 3 remain valid as long as the amount of exogenous information is sufficiently small.³⁹ Thus, for a given amount of misspecification, learning outcomes under NIP offer a good approximation as long as the predominant source of information is costly rather than costless. \blacktriangle

6.3 Sequential social learning

Finally, we analyze the sequential social learning environment from Example 3. We impose the following assumptions: Private signals s_t are drawn i.i.d. across agents conditional on each state ω , according to a positive and continuous density $\phi(\cdot|\omega)$ that satisfies the monotone likelihood ratio property; i.e., $\frac{\phi(s|\omega)}{\phi(s|\omega')}$ is strictly increasing in s for any $\omega > \omega'$. True and perceived type distributions F and \hat{F} admit positive densities, and the utility difference $v(\theta, \omega) := u(1, \theta, \omega) - u(0, \theta, \omega)$ between actions is strictly increasing and continuous in types and states (θ, ω) , with $\lim_{\theta \rightarrow -\infty} v(\theta, \omega) < 0$ and $\lim_{\theta \rightarrow +\infty} v(\theta, \omega) > 0$; thus, sufficiently low (risk-averse) types always prefer action 0 (not adopt) and sufficiently high (risk-tolerant) types always prefer action 1 (adopt).

Recall from Remark 1 that at each public belief μ , the true and perceived probabilities that the current-period action is 0 satisfy

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

where $\mu^s \in \Delta(\Omega)$ denotes the Bayesian update of μ following private signal realization s and for each $\nu \in \Delta(\Omega)$, $\theta^*(\nu)$ denotes the type who is indifferent between action 0 and 1 at belief ν . Note that $\theta^*(\nu)$ exists and is unique 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 .

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

Lemma 8. *Suppose that $\hat{F} = F$. Then learning is successful at all states ω^* .*

³⁹For example, suppose true and perceived probabilities are $(\gamma(\mu) + \alpha)\omega^* + q$, $(\gamma(\mu) + \alpha)\omega + \hat{q}$ for some exogenous $\alpha > 0$. Then for any $\hat{q} < q$, there exists $\varepsilon > 0$ such that δ_{ω_N} is globally stable at all ω^* whenever $\alpha < \varepsilon$; and likewise for $\hat{q} > q$.

An analogous result is established by [Goeree, Palfrey, and Rogers \(2006\)](#). Observe that herding is ruled out due to the rich preference heterogeneity of the present setting (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 NIP:

Lemma 9. *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 9 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, NIP holds. The logic is based on the following well-known feature of social learning that is sometimes referred to as the “self-correcting property” ([Vives, 2010](#)): The more indicative an agent perceives previous actions to be of a particular state ω , the less heavily she weights her own private signal s in choosing her action, and in the limit as previous actions become perfectly indicative of ω , she disregards s and acts solely based on the public belief δ_ω .

Similar to costly information acquisition, NIP again leads successful learning to be highly non-robust to the introduction of misspecification. The following result classifies possible learning outcomes depending on the nature of misspecification:

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

1. δ_{ω_N} is globally stable if $F >_{FOSD} \hat{F}$,⁴⁰ 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 $\hat{F} >_{FOSD} F$, 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 $F(\theta_n^*) \neq \hat{F}(\theta_n^*)$.

Proposition 4 highlights three general possibilities. First, beliefs might converge globally to a point-mass on the highest possible state ω_N (resp. lowest possible state ω_1). 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 under any standard notion of distance 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

⁴⁰We write $F >_{FOSD} \hat{F}$ if $F(\theta) < \hat{F}(\theta)$ for all θ .

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 (except when $F(\theta_i^*) = \hat{F}(\theta_i^*)$ for some i) *all* point-mass beliefs are unstable, so that beliefs cycle.⁴¹

To see the idea, consider any state ω_i . If $F(\theta_i^*) < \hat{F}(\theta_i^*)$, then NIP (Lemma 9) implies that at all states ω and 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 than the actual probability $P_\mu(0|\omega^*) \approx F(\theta_i^*)$. At the 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 Theorems 1 and 2, while Theorem 4 implies the global stability results.

Analogous to Remark 3, we note that when $\hat{F} \neq F$, the true long-run action distribution disagrees with agents' perceived distribution, but this feature is again not essential to our result: For instance, an analog of Proposition 4 can be obtained when agents are correct about F but misperceive the private signal distribution ϕ , and in this case true and perceived long-run action distributions coincide (by NIP).

Remark 4. We briefly contrast Proposition 4 with other recent findings in the literature:

In closely related sequential social learning environments, [Bohren \(2016\)](#) and [Bohren and Hauser \(2018\)](#) show that successful information aggregation is robust to small amounts of misspecification. A key difference relative to our setting is their assumption that either (i) agents observe an informative public signal every period, or (ii) a positive fraction of agents (so-called “autarkic types”) choose actions based solely on their own private signal. This assumption rules out NIP, because it leads the perceived action distribution $\hat{P}_\mu(\cdot|\omega)$ to be informative about the state (i.e., to depend on ω) even when μ is a point-mass. However, for a given misspecification, we note (analogous to Remark 3) that the predictions we obtained under NIP remain valid as long as the public signal in (i) is sufficiently uninformative or the fraction of autarkic types in (ii) is sufficiently small.

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

⁴¹As in the correctly specified case, mixed beliefs are unstable, since the richness of types implies the identification condition in Lemma 1.

we show that this can happen even under arbitrarily small amounts of misspecification.

Finally, [Frick, Iijima, and Ishii \(2019\)](#) (FII2019) study a different social learning model, where a continuum of agents each privately observe the action of a random other agent each period. Their setting is not nested by the present framework, as there is no public belief. They also highlight that vanishingly small amounts of misspecification about the type distribution F can lead beliefs to converge to a state-independent point-mass, but both the logic and nature of this fragility result are somewhat different. Specifically, the current setting and FII2019 have in common that for any F and \hat{F} , the set of equilibrium beliefs is “decoupled,” in the sense that this set does not depend on the true state ω^* . However, in the current setting, NIP implies that for any F and \hat{F} , this set consists of *all* point-mass beliefs, and the logic behind Proposition 4 is that misspecification can discontinuously change which of these beliefs are *stable*. By contrast, Theorem 1 in FII2019 highlights a discontinuity at the level of the equilibrium correspondence, by showing that misspecification can discontinuously shrink the set of equilibria to a *single* state-independent point-mass. The key difference is that the environment in FII2019 does *not* feature NIP, because in their private observation setting, agents view new action observations as informative, no matter how convinced they themselves have become in a particular state.⁴² Importantly, in contrast with Proposition 4, the discontinuity of the equilibrium correspondence in FII2019 relies on a continuous state space, and they show that in finite state spaces, successful learning is robust. ▲

7 Concluding remarks

This paper presents an approach to analyze learning outcomes in a broad class of misspecified environments, including single-agent and social learning. Our main results provide local and global stability criteria for long-run beliefs, using a prediction accuracy order over subjective models— p -dominance—that refines comparisons based on Kullback-Leibler divergence. Our approach makes it possible to analyze a natural class of environments, where arbitrarily small amounts of misspecification can significantly alter learning outcomes, because agents’ behavior near point-mass beliefs generates increasingly uninformative new signals. We also apply our criteria to unify and generalize various convergence results in previously studied settings.

In Appendix G, we extend the main model and results to a setting with profiles of beliefs, which makes it possible to accommodate additional classes of environments. This

⁴²To see this, note that agents believe that information aggregation is successful in the long run. Thus, no matter her own current belief, an agent expects that in each state ω' , the long-run probability of observing action 0 approximates $\hat{F}(\theta_{\omega'}^*)$ and hence is informative of the state, in contrast with Lemma 9.

includes learning in games (e.g., [Fudenberg and Kreps, 1993](#); [Esponda and Pouzo, 2016](#)), where players repeatedly choose actions, observe signals about an underlying fundamental and others' actions, and update their beliefs under the assumption that others' behavior is time-stationary and other possible misspecifications about the environment (e.g., [Eyster and Rabin, 2005](#); [Jehiel, 2005](#)). This setting can also accommodate social learning environments with heterogeneous misspecifications (e.g., [Bohren and Hauser, 2018](#); [Gagnon-Bartsch, 2017](#); [Levy, Razin, and Young, 2019](#)).

In ongoing work, we pursue extensions of our results to infinite state space settings as well as Markov decision problems. Finally, while we have maintained the assumption that agents are Bayesian (as is standard in the misspecified learning literature), our approach can also be applied to more general models of learning in which the belief μ_t follows some Markov process. For instance, for local stability, this entails suitably modifying condition (6) in Theorem 1 to ensure that the process $\left(\frac{\mu_t(\omega')}{\mu_t(\omega)}\right)^p$ is again a nonnegative supermartingale up to some stopping time.

Appendix: Main Proofs

A Preliminaries

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 $K \subseteq \Delta(\Omega)$. 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 $\left|\log \frac{\hat{P}_\mu(z_\mu|\omega)}{\hat{P}_\mu(z_\mu|\omega')}\right| > 0$. By (i) and continuity of P and \hat{P} with respect to μ (Assumption 1.3), there is a neighborhood B_μ of μ with

$$\inf_{\mu' \in B_\mu} \min \left\{ P_{\mu'}(z_\mu), \left| \log \frac{\hat{P}_{\mu'}(z_\mu|\omega)}{\hat{P}_{\mu'}(z_\mu|\omega')} \right|, \mu'(\omega), \mu'(\omega') \right\} > 0.$$

By compactness of K , there is a finite subcover $(B_{\mu^i})_{i=1}^n$ of K . Let

$$\gamma := \inf_{\substack{i=1, \dots, n, \\ \mu' \in B_{\mu^i}}} \min \left\{ P_{\mu'}(z_{\mu^i}), \left| \log \frac{\hat{P}_{\mu'}(z_{\mu^i}|\omega)}{\hat{P}_{\mu'}(z_{\mu^i}|\omega')} \right| \right\}.$$

By construction, $\gamma > 0$.

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

$$\mathbb{P}_{\mu'_0} \left[\mu_t \in K \forall t \text{ and } \exists T < \infty \text{ s.t. } \left| \log \frac{\mu_t(\omega)}{\mu_t(\omega')} - \ell \right| \leq \gamma/2 \forall t \geq T \right] > 0. \quad (10)$$

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 some i such that $\mu_t \in B_{\mu^i}$, $P_{\mu^i}(z_{\mu^i}) > \gamma$ and $\left| \log \frac{\hat{P}_{\mu^i}(z_{\mu^i}|\omega)}{\hat{P}_{\mu^i}(z_{\mu^i}|\omega')} \right| > \gamma$, and hence 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 considered in (10) occurs with zero probability, a contradiction. \square

The following result shows that δ_ω is globally stable whenever ω strictly p -dominates all other states at all mixed beliefs, but may be tied at point-mass beliefs:

Proposition 5. *Consider any $\omega \in \Omega$. Suppose that for some $p > 0$, we have $\omega \succsim_\mu^p \omega'$ 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 and consider the induced probability measure \mathbb{P}_{μ_0} over sequences of beliefs. For each $\omega' \neq \omega$, $\ell_t(\omega') := \left(\frac{\mu_t(\omega')}{\mu_t(\omega)} \right)^p$ is a non-negative supermartingale, since $\omega \succsim_\mu^p \omega'$ for all μ . Thus, by Doob's convergence theorem, there exists an L^∞ random variable $\ell_\infty(\omega')$ such that $\ell_t(\omega') \rightarrow \ell_\infty(\omega') \geq 0$ almost surely. Hence, the belief process μ_t converges almost surely. 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 $\omega \succ_\mu^p \omega'$, 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 5 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 there exists $\omega^* \in \Omega$ such that (i) $P_\mu(\cdot) = \hat{P}_\mu(\cdot|\omega^*)$ 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 $p \in (0, 1)$ and $\omega \neq \omega^*$. For each belief μ , we have

$$\sum_{z \in \text{supp} P_\mu} P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega)}{P_\mu(z)} \right)^p \leq \left(\sum_{z \in \text{supp} P_\mu} \hat{P}_\mu(z|\omega) \right)^p \leq 1, \quad (11)$$

where the first inequality holds by Jensen's inequality applied to the concave function $f(x) = x^p$. Since $P_\mu(\cdot) = \hat{P}_\mu(\cdot|\omega^*)$ by (i), this shows that $\omega^* \succsim_\mu^p \omega$. Consider any mixed μ . By Assumption 1,

$\text{supp}P_\mu \subseteq \text{supp}\hat{P}_\mu(\cdot|\omega)$. If $\text{supp}P_\mu \subsetneq \text{supp}\hat{P}_\mu(\cdot|\omega)$, then the second inequality in (11) is strict. If $\text{supp}P_\mu = \text{supp}\hat{P}_\mu(\cdot|\omega)$, then by (ii), $\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z)} \neq \frac{\hat{P}_\mu(z'|\omega)}{\hat{P}_\mu(z')}$ for some $z, z' \in \text{supp}P_\mu$, in which case the first inequality in (11) is strict. In either case, $\omega^* \succ_\mu^p \omega$. Thus, the conclusion follows from Proposition 5. \square

B Proofs for Sections 3 and 4

B.1 Proof of Lemma 1

By assumption, there exist $\omega, \omega' \in \text{supp}\mu$ and $z^* \in \text{supp}P_\mu(\cdot)$ such that $\frac{\hat{P}_\mu(z^*|\omega)}{\hat{P}_\mu(z^*|\omega')} \neq 1$. By continuity of P and \hat{P} with respect to μ , there exists $\gamma > 0$ such that $\left| \log \frac{\hat{P}_{\mu'}(z^*|\omega)}{\hat{P}_{\mu'}(z^*|\omega')} \right|, P_{\mu'}(z^*) > \gamma$ for all μ' with $\|\mu' - \mu\| < \gamma$. Take $\varepsilon > 0$ sufficiently small that $\left| \log \frac{\mu'(\omega)}{\mu'(\omega')} - \log \frac{\mu(\omega)}{\mu(\omega')} \right| \leq \gamma$ for all $\mu' \in B$, where B denotes the ε -ball around μ . Then if $\mu_t \in B$, a realization of $z_t = z^*$, which occurs with probability at least $\gamma > 0$, implies $\mu_{t+1} \notin B$. Since the belief process is Markov, we have $\mathbb{P}_{\mu_0}[\mu_t \in B \forall t] = 0$ from any initial belief $\mu_0 \in B$. Hence, μ is unstable.

B.2 Proof of Lemma 2

Suppose δ_ω is not a Berk-Nash equilibrium. Then there exists $\omega' \neq \omega$ such that $\omega' \succ_{\delta_\omega}^{\text{KL}} \omega$, whence by part 3 of Lemma 3, $\omega' \succ_\mu^{\text{KL}} \omega$ holds for all μ in a sufficiently small neighborhood $B \ni \delta_\omega$. Hence, δ_ω is unstable by Theorem 2.

B.3 Proof of Lemma 3

Consider the random variable $\log \frac{\hat{P}_\mu(z|\omega')}{\hat{P}_\mu(z|\omega)}$, where z is distributed according to P_μ . Let $M(p) := \sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega')}{\hat{P}_\mu(z|\omega)} \right)^p$ denote its moment-generating function evaluated at $p \in \mathbb{R}$. Note that M is convex with $M(0) = 1$ and $M'(0) = \sum_z P_\mu(z) \log \frac{\hat{P}_\mu(z|\omega')}{\hat{P}_\mu(z|\omega)}$.

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

Part 2. If $\omega \succ_\mu^{\text{KL}} \omega'$, then $M'(0) < 0$. Thus, for all sufficiently small $p > 0$, continuity of M' implies $M(p) < M(0) = 1$, i.e., $\omega \succ_\mu^p \omega'$.

Part 3. Let $K := \sup_{\mu \in \Delta(\Omega), z \in \text{supp} P_\mu} \frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega')}$, which is finite by Assumption 1.2. Take any $\omega, \omega' \in \Omega$, $\mu \in \Delta(\Omega)$, and any convergent sequence of beliefs $\mu^n \rightarrow \mu$. By the triangle inequality,

$$\begin{aligned} & \left| \sum_z P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega')} \right) - \sum_z P_{\mu^n}(z) \log \left(\frac{\hat{P}_{\mu^n}(z|\omega)}{\hat{P}_{\mu^n}(z|\omega')} \right) \right| \\ & \leq \left| \sum_{z \in \text{supp}(P_\mu)} \left(P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega')} \right) - P_{\mu^n}(z) \log \left(\frac{\hat{P}_{\mu^n}(z|\omega)}{\hat{P}_{\mu^n}(z|\omega')} \right) \right) \right| \\ & \quad + \left| \sum_{z \notin \text{supp}(P_\mu)} P_{\mu^n}(z) \log \left(\frac{\hat{P}_{\mu^n}(z|\omega)}{\hat{P}_{\mu^n}(z|\omega')} \right) \right|. \end{aligned}$$

On the right-hand side of the inequality, as $n \rightarrow 0$, the first term goes to zero since for each $z \in \text{supp}(P_\mu)$, we have $P_{\mu^n}(z) \rightarrow P_\mu(z) > 0$, $\hat{P}_{\mu^n}(z|\omega) \rightarrow \hat{P}_\mu(z|\omega) > 0$, and $\hat{P}_{\mu^n}(z|\omega') \rightarrow \hat{P}_\mu(z|\omega') > 0$ by Assumptions 1.1 and 1.3. Likewise, the second term goes to zero since $|\log \left(\frac{\hat{P}_{\mu^n}(z|\omega)}{\hat{P}_{\mu^n}(z|\omega')} \right)| \leq K$ for each n by choice of K and $\sum_{z \notin \text{supp}(P_\mu)} P_{\mu^n}(z) \rightarrow 0$ by Assumption 1.3. Thus, $\sum_z P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega')} \right)$ is continuous at μ . The argument for continuity of $\sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega')} \right)^p$ is analogous.

B.4 Proof of Theorem 1

Suppose there exist $p > 0$ and $B \ni \delta_\omega$ such that (6) holds. 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 $\ell_t(\omega') := \left(\frac{\mu_{\min\{t, \tau\}}(\omega')}{\mu_{\min\{t, \tau\}}(\omega)} \right)^p$, where $\tau := \inf\{t' : \mu_{t'} \notin B\}$. By (6), 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 almost surely.

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 $\ell_t(\omega') \rightarrow \ell_\infty(\omega')$ almost surely implies that each $\frac{\mu_t(\omega')}{\mu_t(\omega)}$ converges almost surely to a finite value. Suppose for a contradiction that for some $\omega' \neq \omega$, $\mathbb{P}_{\mu_0}[\lim_t \frac{\mu_t(\omega')}{\mu_t(\omega)} > 0 \mid \tau = \infty] > 0$. Then there exists a compact $K \subseteq B$ such that $\mu(\omega'), \mu(\omega) > 0$ for all $\mu \in K$ and $\mathbb{P}_{\mu_0}[\exists T \text{ s.t. } \mu_t \in K \forall t \geq T \text{ and } \exists \lim_t \frac{\mu_t(\omega')}{\mu_t(\omega)} \mid \tau = \infty] > 0$. 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 $\mathbb{P}_{\mu_0}[\lim_t \frac{\mu_t(\omega')}{\mu_t(\omega)} = 0 \mid \tau = \infty] = 1$ for all $\omega' \neq \omega$. Thus, $\mathbb{P}_{\mu_0}[\mu_t \rightarrow \delta_\omega \mid \tau = \infty] = 1$, 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 $\mu_0 \in B'$.

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

For any initial belief $\mu_0 \in B' := \{\mu \in \Delta(\Omega) : \left(\frac{\mu(\omega')}{\mu(\omega)}\right)^p < \varepsilon_-, \forall \omega' \neq \omega\} \subseteq B$, we have

$$\begin{aligned} \mathbb{P}_{\mu_0}[\exists t, \mu_t \notin B] &\leq \mathbb{P}_{\mu_0}[\exists \omega', \ell_\infty(\omega') \geq \varepsilon_+] \leq \sum_{\omega' \neq \omega} \mathbb{P}_{\mu_0}[\ell_\infty(\omega') \geq \varepsilon_+] \\ &\leq \sum_{\omega' \neq \omega} \mathbb{E}_{\mu_0}[\ell_\infty(\omega')]/\varepsilon_+ \leq (|\Omega| - 1) \frac{\varepsilon_-}{\varepsilon_+}, \end{aligned}$$

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

B.5 Proof of Theorem 2

Suppose there exists a neighborhood $B \ni \delta_\omega$ such that (8) holds for some $\omega' \neq \omega$. 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 $\mu_0 \in B \setminus \{\delta_\omega\}$ and induced probability measure \mathbb{P}_{μ_0} over sequences of beliefs. Let $\tau := \inf\{t' : \mu_{t'} \notin B\}$. To prove instability of δ_ω , it suffices to show that $\mathbb{P}_{\mu_0}[\tau < \infty] = 1$. Consider the process $\ell_t := \log\left(\frac{\mu_{\min\{t, \tau\}}(\omega)}{\mu_{\min\{t, \tau\}}(\omega')}\right)$. This is a supermartingale by (8) and is bounded below, as $\ell_t \geq \log \frac{\varepsilon}{1-\varepsilon} - \log M$ for all t , where M is taken from Assumption 1.2. Hence, Doob's convergence theorem yields an L^∞ -random variable ℓ_∞ such that $\ell_t \rightarrow \ell_\infty$ almost surely.

Suppose for a contradiction that with positive probability, we have $\tau = \infty$. Conditional on $\tau = \infty$, we have $\log\left(\frac{\mu_t(\omega)}{\mu_t(\omega')}\right) = \ell_t$ for all t . Thus, conditional on $\tau = \infty$, $\frac{\mu_t(\omega)}{\mu_t(\omega')}$ converges almost surely 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 (8) implies that for each $\mu \in K$, there exists $z \in \text{supp } P_\mu$ with $\hat{P}_\mu(z_\mu|\omega) \neq \hat{P}_\mu(z_\mu|\omega')$.

B.6 Proof of Corollary 1

Note first that if δ_ω is a strict Berk-Nash equilibrium in some environment (P, \hat{P}) , then δ_ω is locally stable at (P, \hat{P}) . Indeed, since $\omega \succ_{\delta_\omega}^{\text{KL}} \omega'$ for all $\omega' \neq \omega$, then Lemma 3.2 yields some $p > 0$ such that $\omega \succ_{\delta_\omega}^p \omega'$ for all $\omega' \neq \omega$. Hence, by Lemma 3.3, there exists a neighborhood $B \ni \delta_\omega$ such that $\omega \succ_\mu^p \omega'$ holds for all $\mu \in B$. Thus, δ_ω is locally stable by Theorem 1.

Given the previous paragraph, to prove Corollary 1, it suffices to show that if δ_ω is a strict Berk-Nash equilibrium at (P, \hat{P}) , i.e.,

$$\sum_z P_{\delta_\omega}(z) \log \frac{\hat{P}_{\delta_\omega}(z|\omega')}{\hat{P}_{\delta_\omega}(z|\omega)} < 0 \text{ for all } \omega' \neq \omega,$$

then there exists $\varepsilon > 0$ such that δ_ω remains a strict Berk-Nash equilibrium at any ε -perturbation

(Q, \hat{Q}) of (P, \hat{P}) , i.e.,

$$\sum_z Q_{\delta_\omega}(z) \log \frac{\hat{Q}_{\delta_\omega}(z|\omega')}{\hat{Q}_{\delta_\omega}(z|\omega)} < 0 \text{ for all } \omega' \neq \omega.$$

But this is immediate from the observation that (i) if $\text{KL}(Q_{\delta_\omega}(\cdot), P_{\delta_\omega}(\cdot)) < \infty$, then $\text{supp} Q_{\delta_\omega} \subseteq \text{supp} P_{\delta_\omega}$; and (ii) if $\text{KL}(Q_{\delta_\omega}(\cdot), P_{\delta_\omega}(\cdot)), \sup_{\omega''} \text{KL}(\hat{Q}_{\delta_\omega}(\cdot|\omega''), \hat{P}_{\delta_\omega}(\cdot|\omega'')) < \varepsilon$, then $\sup_z |P_{\delta_\omega}(z) - Q_{\delta_\omega}(z)|, \sup_{z, \omega''} |\hat{P}_{\delta_\omega}(z|\omega'') - \hat{Q}_{\delta_\omega}(z|\omega')| < \sqrt{\varepsilon/2}$ by Pinsker's inequality.

C Proofs for Section 5

C.1 Proof of Theorem 3

Since Ω is finite, $S^\infty(\Omega) = S^n(\Omega)$ for some n . Let $S^k := S^k(\Omega)$ for each k . For each $k = 1, \dots, n$ and $\omega_k \in S^{k-1} \setminus S^k$, there exists a state $\bar{\omega}_k \in S^{k-1}$ such that $\bar{\omega}_k \succ_\mu^{\text{KL}} \omega_k$ for all $\mu \in \Delta(S^{k-1})$. Thus, for each $\omega_k \in S^{k-1} \setminus S^k$ and $\mu \in \Delta(S^{k-1})$, Lemma 3 yields some $p_{\omega_k, \mu} > 0$ and a closed ball $D(\mu; \omega_k)$ around μ such that

$$\max_{\mu \in D(\mu; \omega_k)} \sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega_k)}{\hat{P}_\mu(z|\bar{\omega}_k)} \right)^{p_{\omega_k, \mu}} < 1.$$

For each $\omega_k \in S^{k-1} \setminus S^k$, we can find a finite subcollection $\{D(\mu_i; \omega_k) | i = 1, \dots, m_{\omega_k}\} \subseteq \{D(\mu; \omega_k) | \mu \in \Delta(S^{k-1})\}$ that covers the compact set $\Delta(S^{k-1})$. Hence, by choosing $p > 0$ small enough, we can ensure that for each $k = 1, \dots, n$ and $\omega_k \in S^{k-1} \setminus S^k$,

$$\max_{\mu \in \Delta(S^{k-1})} \sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega_k)}{\hat{P}_\mu(z|\bar{\omega}_k)} \right)^p < 1. \quad (12)$$

To show that $\Delta(S^n)$ is globally stable, it suffices to show that

$$\mathbb{P}_{\mu_0} \left[\frac{\mu_t(\omega_k)}{\mu_t(\bar{\omega}_k)} \rightarrow 0 \ \forall \omega_k \in S^{k-1} \setminus S^k \right] = 1 \quad (13)$$

for each $k = 1, \dots, n$ and initial belief μ_0 . We prove this by induction on k .

First consider the case $k = 1$. Take any $\omega_1 \in \Omega \setminus S^1(\Omega)$ and initial belief μ_0 . By (12), the process $\ell_t(\omega_1) := \left(\frac{\mu_t(\omega_1)}{\mu_t(\bar{\omega}_1)} \right)^p$ is a non-negative supermartingale. Thus, there exists an L^∞ -random variable $\ell_\infty(\omega_1)$ such that $\ell_t(\omega_1) \rightarrow \ell_\infty(\omega_1)$ almost surely. Suppose for a contradiction that $\mathbb{P}_{\mu_0}[\lim_{t \rightarrow \infty} \frac{\mu_t(\omega_1)}{\mu_t(\bar{\omega}_1)} > 0] > 0$. Then there exists a compact set $K \subseteq \Delta(\Omega)$ such that $\mu(\omega_1), \mu(\bar{\omega}_1) > 0$ for all $\mu \in K$ and $\mathbb{P}_{\mu_0}[\exists T \text{ s.t. } \mu_t \in K \ \forall t \geq T \text{ and } \exists \lim_{t \rightarrow \infty} \frac{\mu_t(\omega_1)}{\mu_t(\bar{\omega}_1)}] > 0$. This contradicts Lemma 10, because by (12), for any mixed μ , there exists $z \in \text{supp} P_\mu$ with $\hat{P}_\mu(z|\omega_1) \neq \hat{P}_\mu(z|\bar{\omega}_1)$. Hence, $\mathbb{P}_{\mu_0}[\frac{\mu_t(\omega_1)}{\mu_t(\bar{\omega}_1)} \rightarrow 0] = 1$, as claimed.

Next, suppose that (13) holds for all $k = 1, \dots, \kappa - 1$ and consider the case $k = \kappa$. For some

$\varepsilon > 0$, define

$$B_\kappa := \{\mu \in \Delta^\circ(\Omega) : \sum_{j=1}^{\kappa-1} \sum_{\omega_j \in S^{j-1} \setminus S^j} \left(\frac{\mu(\omega_j)}{\mu(\bar{\omega}_j)} \right)^p \leq \varepsilon\}.$$

Note that B_κ converges to $\Delta(S^{\kappa-1})$ in the Hausdorff distance as $\varepsilon \rightarrow 0$. Thus, by (12), Lemma 3, and the fact that $\Delta(S^{\kappa-1}) \subseteq \Delta(S^{j-1})$ for all $j \leq \kappa$, we can choose ε sufficiently small that

$$\max_{\mu \in \text{cl}(B_\kappa)} \sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega_j)}{\hat{P}_\mu(z|\bar{\omega}_j)} \right)^p < 1 \quad (14)$$

for all $j = 1, \dots, \kappa$ and $\omega_j \in S^{j-1} \setminus S^j$.

Let $\tau := \min\{t \geq 0 : \mu_t \notin B_\kappa\}$. The inductive hypothesis implies that

$$\mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \mu_t \in B_\kappa] = 1 \quad (15)$$

for any initial belief μ_0 . Given this, it suffices to show⁴³ that there exists $\gamma > 0$ such that for any initial belief $\mu_0 \in B_\kappa$, we have

$$\mathbb{P}_{\mu_0}[\tau = \infty] \geq \gamma \quad (16)$$

$$\mathbb{P}_{\mu_0} \left[\frac{\mu_t(\omega_\kappa)}{\mu_t(\bar{\omega}_\kappa)} \rightarrow 0 \ \forall \omega_\kappa \in S^{\kappa-1} \setminus S^\kappa \mid \tau = \infty \right] = 1. \quad (17)$$

To verify (16), consider

$$\ell_t^\kappa := \sum_{j=1}^{\kappa} \sum_{\omega_j \in S^{j-1} \setminus S^j} \left(\frac{\mu_{\min\{\tau, t\}}(\omega_j)}{\mu_{\min\{\tau, t\}}(\bar{\omega}_j)} \right)^p,$$

which is a nonnegative supermartingale by (14). Thus, there is an L^∞ -random variable ℓ_∞^κ such that $\ell_t^\kappa \rightarrow \ell_\infty^\kappa$ almost surely. Moreover, there exists $\gamma \in (0, 1)$ such that for any initial belief $\mu_0 \in B_\kappa$, we have

$$\begin{aligned} \mathbb{E}_{\mu_0}[\ell_1^\kappa / \ell_0^\kappa] &\leq \max_{\substack{j=1, \dots, \kappa \\ \omega_j \in S^{j-1} \setminus S^j}} \mathbb{E}_{\mu_0} \left[\left(\frac{\mu_1(\omega_j)}{\mu_1(\bar{\omega}_j)} \right)^p / \left(\frac{\mu_0(\omega_j)}{\mu_0(\bar{\omega}_j)} \right)^p \right] = \\ &\max_{\substack{j=1, \dots, \kappa \\ \omega_j \in S^{j-1} \setminus S^j}} \sum_z P_{\mu_0}(z) \left(\frac{\hat{P}_{\mu_0}(z|\omega_j)}{\hat{P}_{\mu_0}(z|\bar{\omega}_j)} \right)^p \leq 1 - \gamma, \end{aligned} \quad (18)$$

⁴³To see this, suppose (16) and (17) hold and let $\phi(\mu) := \mathbb{P}_\mu \left[\frac{\mu_t(\omega_\kappa)}{\mu_t(\bar{\omega}_\kappa)} \rightarrow 0 \ \forall \omega_\kappa \in S^{\kappa-1} \setminus S^\kappa \right]$ for each initial belief $\mu \in \Delta^\circ(\Omega)$. Then for any initial belief $\mu \notin B_\kappa$, $\phi(\mu) \geq \inf_{\mu' \in B_\kappa} \phi(\mu')$ by (15) and the Markov property. Moreover, for any initial belief $\mu \in B_\kappa$,

$$\phi(\mu) \geq \mathbb{P}_\mu[\tau = \infty] + \mathbb{P}_\mu[\tau < \infty] \inf_{\mu' \in B_\kappa} \phi(\mu') \geq \gamma + (1 - \gamma) \inf_{\mu' \in B_\kappa} \phi(\mu')$$

where the first equality uses (17) and the second uses (16). Thus, $\inf_{\mu \notin B_\kappa} \phi(\mu) \geq \inf_{\mu' \in B_\kappa} \phi(\mu') = 1$, as required.

where the last inequality holds by (14). Hence, for any initial belief $\mu_0 \in B_\kappa$

$$\mathbb{P}_{\mu_0}[\tau < \infty] \leq \mathbb{P}_{\mu_0}[\ell_\infty^\kappa \geq \varepsilon] \leq \mathbb{E}_{\mu_0}[\ell_\infty^\kappa]/\varepsilon \leq \mathbb{E}_{\mu_0}[\ell_1^\kappa]/\varepsilon \leq (1 - \gamma) \frac{\ell_0^\kappa}{\varepsilon} \leq 1 - \gamma,$$

where the second inequality uses Markov's inequality, the third inequality uses Fatou's lemma and the fact that ℓ_t^κ is a non-negative supermartingale, and the fourth inequality uses (18). This verifies (16).

Finally, to verify (17), take any $\omega_\kappa \in S^{\kappa-1} \setminus S^\kappa$ and initial belief $\mu_0 \in B_\kappa$. By (14), $\ell_t(\omega_\kappa) := \left(\frac{\mu_{\min\{t, \tau\}}(\omega_\kappa)}{\mu_{\min\{t, \tau\}}(\bar{\omega}_\kappa)} \right)^p$ is a nonnegative supermartingale. Thus, there is an L^∞ -random variable $\ell_\infty(\omega_\kappa)$ such that $\ell_t(\omega_\kappa) \rightarrow \ell_\infty(\omega_\kappa)$ almost surely. Hence, conditional on the event $\tau = \infty$, $\frac{\mu_t(\omega_\kappa)}{\mu_t(\bar{\omega}_\kappa)}$ converges to a finite value almost surely. Suppose for a contradiction that $\mathbb{P}_{\mu_0}[\lim_{t \rightarrow \infty} \frac{\mu_t(\omega_\kappa)}{\mu_t(\bar{\omega}_\kappa)} > 0 \mid \tau = \infty] > 0$. Then there exists a compact set $K \subseteq \Delta(\Omega)$ such that $\mu(\omega_k), \mu(\bar{\omega}_k) > 0$ for all $\mu \in K$ and $\mathbb{P}_{\mu_0}[\exists T \text{ s.t. } \mu_t \in K \forall t \geq T \text{ and } \exists \lim_{t \rightarrow \infty} \frac{\mu_t(\omega_1)}{\mu_t(\bar{\omega}_1)} \mid \tau = \infty] > 0$. This contradicts Lemma 10, because by (14), for any mixed $\mu \in B_\kappa$, there exists $z \in \text{supp} P_\mu$ with $\hat{P}_\mu(z|\omega_\kappa) \neq \hat{P}_\mu(z|\bar{\omega}_\kappa)$. Hence, $\mathbb{P}_{\mu_0}[\frac{\mu_t(\omega_\kappa)}{\mu_t(\bar{\omega}_\kappa)} \rightarrow 0 \mid \tau = \infty] = 1$, proving (17).

C.2 Proof of Lemma 4

We say that a set of beliefs $K \subseteq \Delta(\Omega)$ is **transient** if $\mathbb{P}_{\mu_0}[\exists t \text{ s.t. } \mu_t \notin K] = 1$ for any initial belief $\mu_0 \in K$. 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, by definition, 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.

There exist $T \in \mathbb{N}$ and $\eta > 0$ such that

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

for every initial belief $\mu_0 \notin \Delta_2$. To see this, take $L > 1$ sufficiently large so that (i) $\mu \in B'_1$ for all μ such that $\log \frac{\mu(\omega_1)}{\mu(\omega_n)} \geq L$ for each $n > 1$, (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, \hat{P} 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_μ such that $P_\mu(z_\mu) > \varepsilon$ and $\log \frac{\hat{P}_\mu(z_\mu|\omega_1)}{\hat{P}_\mu(z_\mu|\omega_n)} > \varepsilon$ for all $n \neq 1$. Starting from any initial belief $\mu_0 \notin \Delta_2$, consider the realization of the sequence (μ_t) of the form $z_t = z_{\mu_t}$, which 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 the desired claim (20) holds by choosing $T \geq \frac{L-1/L}{\varepsilon}$ and $\gamma = \varepsilon^T$.

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

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

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

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

Since this holds for all $\mu_0 \notin \Delta_2$, this implies $\inf_{\mu \in \Delta^\circ(\Omega) \setminus \Delta_2} h(\mu) = 1$. Thus, by (21), $\inf_{\mu \in \Delta^\circ(\Omega)} h(\mu) = 1$, as required.

C.3 Proof of Theorem 4

To prove that δ_{ω_1} is globally stable, it suffices to show that the assumptions in Lemma 4 are satisfied. Note that assumptions (i) and (iii) in Lemma 4 follow from assumptions (i) and (iii) in Theorem 4 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 (ii) in Theorem 4. The following lemma completes the proof, by showing that the inductive step follows from assumptions (i)–(iii) in Theorem 4.

Lemma 11. *Fix any $n \in \{2, \dots, N - 1\}$ and suppose that $\Delta(\{\omega_{n+1}, \dots, \omega_N\})$ is unstable. Assume that (i) there exist $p > 0$ and a neighborhood $B_n \ni \delta_{\omega_n}$ such that $\omega_n \succ_p^\mu \omega_k$ 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

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

the following three properties are satisfied:

$$B'_n \subseteq B_n \tag{23}$$

$$\forall \mu \in \Delta_n \setminus (\Delta_{n+1} \cup B'_n), \exists z \in Z \text{ with } P_\mu(z), \frac{\hat{P}_\mu(z|\omega_n)}{\hat{P}_\mu(z|\omega_k)} - 1 \geq \gamma \text{ for all } k > n \tag{24}$$

$$\frac{\varepsilon_{n+1}}{\varepsilon_{n+1} - \varepsilon_n} \leq 1 + \gamma. \tag{25}$$

To see this, we 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 $\sum_{k>n} \left(\frac{\mu(\omega_k)}{\mu(\omega_n)} \right)^p \leq \varepsilon$. Then (23) is satisfied for all sufficiently small $\varepsilon_n \in (0, \varepsilon_{n+1})$. To show (24), note that by assumption (iii), for all $\mu \in \Delta(\{\omega_n, \dots, \omega_N\}) \setminus \{\delta_{\omega_n}, \dots, \delta_{\omega_N}\}$, there exists $z \in Z$ with $P_\mu(z), \frac{\hat{P}_\mu(z|\omega_n)}{\hat{P}_\mu(z|\omega_k)} - 1 > 0$ 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, we can find $\gamma \in (0, 1)$ such that⁴⁴

$$\forall \mu \in \Delta(\{\omega_n, \dots, \omega_N\}) \setminus (\Delta_{n+1} \cup B'_n), \exists z \in Z \text{ with } P_\mu(z), \frac{\hat{P}_\mu(z|\omega_n)}{\hat{P}_\mu(z|\omega_k)} - 1 \geq 2\gamma \text{ for all } k > n,$$

where γ can be chosen independently of ε_n . For all sufficiently small ε_n , we can then ensure that (24) and (25) hold.⁴⁵

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

Claim 1: There exist $T \in \mathbb{N}$ 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 $\mu_0 \in \Delta_n \setminus \Delta_{n+1}$.

⁴⁴By Assumption 1, the expression $\max_{z \in Z} \min\{P_\mu(z), \frac{\hat{P}_\mu(z|\omega_n)}{\hat{P}_\mu(z|\omega_k)} - 1\}$ is finite and lower-semicontinuous in μ . Hence, since the expression is strictly positive for all μ in the compact set K , so is its minimum over K .

⁴⁵The fact that (24) holds for small enough ε_n again follows from the lower-semicontinuity of expression $\max_{z \in Z} \min\{P_\mu(z), \frac{\hat{P}_\mu(z|\omega_n)}{\hat{P}_\mu(z|\omega_k)} - 1\}$ in μ . Since the expression is greater than 2γ on $\Delta(\{\omega_n, \dots, \omega_N\}) \setminus (\Delta_{n+1} \cup B'_n)$, it must be greater than γ in some sufficiently small neighborhood of $\Delta(\{\omega_n, \dots, \omega_N\}) \setminus (\Delta_{n+1} \cup B'_n)$.

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 $\mu_0 \in \Delta_n \setminus \Delta_{n+1}$, since $\mu_0(\omega_n) \geq \varepsilon_{n+1} - \varepsilon_n > 0$. Thus, there exists T such that $\sum_{k>n} \left(\frac{\mu_0(\omega_k)}{\mu_0(\omega_n)} (1 + \gamma)^{-T} \right)^p \leq \varepsilon$.

Starting with any initial belief $\mu_0 \in \Delta_n \setminus \Delta_{n+1}$, we recursively define the following sequence of signal realizations $z_0, z_1, \dots, z_{T'}$ with $T' \leq T-1$ and corresponding updated beliefs $\mu_1, \mu_2, \dots, \mu_{T'+1}$ (that is, μ_{t+1} is the update of μ_t following signal realization z_t). Suppose we have constructed z_0, \dots, z_{t-1} for some $t \in \{0, \dots, T\}$. We distinguish two cases:

(a) Suppose $\mu_t \in B'_n \cup \Delta_n^c$. Then we 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 (24), we can pick a signal z_t such that $P_{\mu_t}(z_t), \frac{\hat{P}_{\mu_t}(z_t|\omega_n)}{\hat{P}_{\mu_t}(z_t|\omega_k)} - 1 \geq \gamma$ 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 that $\mu_{t+1} \notin \Delta_{n+1}$. To see this, 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 $\frac{\mu_{t+1}(\omega_n)}{\mu_{t+1}(\omega_k)} \geq \frac{\mu_t(\omega_n)}{\mu_t(\omega_k)} (1 + \gamma)$ 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 (25).

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 $\sum_{k>n} \left(\frac{\mu_T(\omega_k)}{\mu_T(\omega_n)} \right)^p \leq \sum_{k>n} \left(\frac{\mu_0(\omega_k)}{\mu_0(\omega_n)} (1 + \gamma)^{-T} \right)^p \leq \varepsilon$ by (b) and the choice of T . This proves Claim 1, as by construction the sequence of signal realizations $(z_0, \dots, z_{T'})$ occurs 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 (23) and the fact that B_n is transient. To show the existence of ξ , define $\ell_t := \sum_{k>n} \left(\frac{\mu_{\min\{t, \tau\}}(\omega_k)}{\mu_{\min\{t, \tau\}}(\omega_n)} \right)^p$. By (23) and assumption (ii), ℓ_t is a non-negative 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 3 and B'_n is compact, it follows that 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 11, 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

$$\sup_{\mu_0 \in \Delta_n} g(\mu_0) = \sup_{\mu_0 \in \Delta_n \setminus \Delta_{n+1}} g(\mu_0). \quad (26)$$

Next, take any $\mu_0 \in B'_n$ and define $\tau := \inf\{t : \mu_t \notin B'_n\}$ as in Claim 2. Then, 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 (26). Thus

$$\sup_{\mu \in B_n} g(\mu) \leq \xi \sup_{\mu \in \Delta_n \setminus \Delta_{n+1}} g(\mu). \quad (27)$$

Finally, take any $\mu_0 \in \Delta_n \setminus \Delta_{n+1}$ and let $\tau'' := \inf\{\min\{t : \mu_t \in \Delta_n^c \cup B'_n\}, T + 1\}$. Then, 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 (27), and the final inequality holds by (26)–(27). Thus, $\sup_{\mu \in \Delta_n \setminus \Delta_{n+1}} g(\mu) = 0$ and the desired conclusion follows from (26). \square

D Proofs for Section 6

D.1 Proof of Proposition 1

We expand the finite state space $\Omega = \{\omega_1, \dots, \omega_N\}$ to the continuous interval $\bar{\Omega} = [\underline{\omega}, \bar{\omega}]$ and show that $S^\infty(\bar{\Omega}) = \{\hat{\omega}\}$. The result then follows from Theorem 3 and Proposition 7 in Appendix F, which shows that $S^\infty(\Omega)$ approximates $S^\infty(\bar{\Omega})$ arbitrarily closely whenever Ω is sufficiently dense in $\bar{\Omega}$. Note that at belief δ_ω , the true probability $P_{\delta_\omega}(1|\omega^*) = \omega^* - \beta a(\delta_\omega)$ and perceived probability $\hat{P}_{\delta_\omega}(1|\omega') = \omega' - \hat{\beta} a(\delta_\omega)$ of high demand are equalized in state $\omega' = \omega^* + \frac{(\hat{\beta} - \beta)\omega}{2\hat{\beta}}$. Thus, for each ω ,

the $\succ_{\delta_\omega}^{\text{KL}}$ -maximal state in $\bar{\Omega}$ is given by

$$M(\omega) := \min\{\max\{\underline{\omega}, \omega^* + \frac{(\hat{\beta} - \beta)\omega}{2\hat{\beta}}\}, \bar{\omega}\}.$$

Suppose $\hat{\beta} \geq \beta$, so that the mapping $M(\cdot)$ is weakly increasing; the case $\hat{\beta} < \beta$ is analogous. Then one can verify (see Lemma 14 in Appendix F) that for each k , $S^k(\bar{\Omega})$ is given by $[M^k(\underline{\omega}), M^k(\bar{\omega})]$. Hence, at each step of the iteration, $S^k(\bar{\Omega})$ contracts by a factor of at least $\frac{\hat{\beta} - \beta}{2\hat{\beta}}$, as $M^k(\bar{\omega}) - M^k(\underline{\omega}) \leq \frac{\hat{\beta} - \beta}{2\hat{\beta}} (M^{k-1}(\bar{\omega}) - M^{k-1}(\underline{\omega}))$. Moreover, since $\hat{\omega} \in \bar{\Omega}$ and $\hat{\omega} = M(\hat{\omega})$, we must have $\hat{\omega} \in S^k(\bar{\Omega})$ for all k . This implies that $S^\infty(\bar{\Omega}) = \{\hat{\omega}\}$, as claimed.

D.2 Proof of Proposition 2

Fix any $\Omega = \{\omega_1, \dots, \omega_N\}$ and β . By finiteness of Ω , it suffices to find an appropriate ε separately for each true state $\omega^* \in \Omega$. Fix any $\omega^* \in \Omega$. By (1), the true and perceived probabilities of high demand given $\hat{\beta}$ satisfy $P_\mu(1) = \omega^* - \frac{\beta \mathbb{E}_\mu[\omega]}{2\hat{\beta}}$ and $\hat{P}_\mu(1|\omega) = \omega - \frac{\mathbb{E}_\mu[\omega]}{2}$. Thus, if $\hat{\beta} = \beta$, then $\omega^* \succ_\mu^{\text{KL}} \omega$ for all $\mu \in \Delta(\Omega)$ and $\omega \neq \omega^*$. Since $\sum_z P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega^*)} \right)$ is continuous in μ (by Lemma 3) and $\Delta(\Omega)$ is compact, this implies that for all $\omega \neq \omega^*$, $\max_{\mu \in \Delta(\Omega)} \sum_z P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega^*)} \right) < 0$ at $\hat{\beta} = \beta$. Moreover, since $\sum_z P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega^*)} \right)$ is continuous in $\hat{\beta}$, so is $\max_{\mu \in \Delta(\Omega)} \sum_z P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega^*)} \right)$. Thus, there exists $\varepsilon > 0$ such that for all $\hat{\beta}$ with $|\hat{\beta} - \beta| < \varepsilon$ and $\omega \neq \omega^*$, $\max_{\mu \in \Delta(\Omega)} \sum_z P_\mu(z) \log \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega^*)} \right) < 0$, that is, $\omega^* \succ_\mu^{\text{KL}} \omega$ for all μ . Then, for any $\hat{\beta}$ with $|\hat{\beta} - \beta| < \varepsilon$, Theorem 3 implies that δ_{ω^*} is globally stable in any state space $\Omega' \subseteq \Omega$ with $\omega^* \in \Omega'$; equivalently, learning is successful at ω^* .

D.3 Proof of Lemma 5

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

$$\sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega)}{\hat{P}_\mu(z|\omega^*)} \right)^p < 1, \quad (28)$$

where the value of the left-hand side is independent of μ . Hence, there exists $\varepsilon > 0$ such that for any \hat{q} with $|\hat{q} - q| < \varepsilon$ and any mixed μ and $\omega \neq \omega^*$, (28) continues to hold, so that $\omega^* \succ_\mu^p \omega$. Thus, for any \hat{q} with $|\hat{q} - q| < \varepsilon$, Proposition 5 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^* .

D.4 Proof of Lemma 6

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

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

$$V_\mu(\gamma) = (\hat{q} + \gamma\mu \cdot \underline{\omega}) v(\bar{\mu}(\gamma)) + (1 - \hat{q} - \gamma\mu \cdot \underline{\omega}) v(\underline{\mu}(\gamma)), \quad (29)$$

where $\underline{\omega} := (\omega_1, \dots, \omega_N)' \in \mathbb{R}^N$ and $\bar{\mu}(\gamma)$ and $\underline{\mu}(\gamma)$ denote the posteriors updated from μ under precision choice γ and perception \hat{q} following signals 1 and 0, respectively. By (2), $\gamma_{\hat{q}}(\mu) \in \arg\max_{\gamma \in [0, \bar{\gamma}]} V_\mu(\gamma) - C(\gamma)$ 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 $\gamma \in [0, \bar{\gamma}]$,

$$|V_\mu(\gamma) - v(\delta_\omega)| < \frac{C(\tilde{\gamma}) - C(0)}{2}. \quad (30)$$

Note that $V_\mu(\gamma)$ is increasing in γ for all μ . Thus, it follows that (30) holds for each $\mu \in B$ and $\gamma \in [0, \bar{\gamma}]$. This implies that for any $\gamma \in [0, \bar{\gamma}]$ and $\mu \in B$,

$$V_\mu(\gamma) - V_\mu(0) \leq |V_\mu(\gamma) - v(\delta_\omega)| + |V_\mu(0) - v(\delta_\omega)| < C(\tilde{\gamma}) - C(0). \quad (31)$$

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

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

where the first inequality uses the fact that C is increasing and the second inequality uses (31). Thus, for all $\mu \in B$, we have $\gamma_{\hat{q}}(\mu) < \tilde{\gamma}$, as claimed.

D.5 Proof of Proposition 3

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

Consider any $\omega \in \Omega$ and $p \in (0, 1)$. By Lemma 6, there exists $B \ni \delta_\omega$ such that $\gamma_{\hat{q}}(\mu) < \frac{\hat{q}-q}{\omega_N-\omega_1}$ for all $\mu \in B$. Thus, for any $\omega', \omega'' \in \Omega$ with $\omega' < \omega''$ and any $\mu \in B \setminus \{\delta_\omega\}$, we have

$$\begin{aligned} \sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega'')}{\hat{P}_\mu(z|\omega')} \right)^p &= (q + \gamma_{\hat{q}}(\mu)\omega^*) \left(\frac{\hat{q} + \gamma_{\hat{q}}(\mu)\omega''}{\hat{q} + \gamma_{\hat{q}}(\mu)\omega'} \right)^p + (1 - q - \gamma_{\hat{q}}(\mu)\omega^*) \left(\frac{1 - \hat{q} - \gamma_{\hat{q}}(\mu)\omega''}{1 - \hat{q} - \gamma_{\hat{q}}(\mu)\omega'} \right)^p \\ &< (\hat{q} + \gamma_{\hat{q}}(\mu)\omega') \left(\frac{\hat{q} + \gamma_{\hat{q}}(\mu)\omega''}{\hat{q} + \gamma_{\hat{q}}(\mu)\omega'} \right)^p + (1 - \hat{q} - \gamma_{\hat{q}}(\mu)\omega') \left(\frac{1 - \hat{q} - \gamma_{\hat{q}}(\mu)\omega''}{1 - \hat{q} - \gamma_{\hat{q}}(\mu)\omega'} \right)^p < 1, \end{aligned}$$

so that $\omega' \succ_\mu^p \omega''$. Here the first inequality uses the fact that $q + \gamma_{\hat{q}}(\mu)\omega^* < \hat{q} + \gamma_{\hat{q}}(\mu)\omega' < \hat{q} + \gamma_{\hat{q}}(\mu)\omega''$ (since $0 < \gamma_{\hat{q}}(\mu) < \frac{\hat{q}-q}{\omega_N-\omega_1}$ and $\omega' < \omega''$), and the second inequality holds by Jensen's inequality.

Note also that for each mixed μ , $\gamma_{\hat{q}}(\mu) > 0$ implies $\hat{P}_\mu(1|\omega') < \hat{P}_\mu(1|\omega'')$. Thus, conditions (a) and (b) in Remark 2 are met. Hence, Theorem 4 implies that δ_{ω_1} is globally stable.

D.6 Proof of Lemma 7

Fix any \hat{q} . 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 (29). Note that the posteriors $\bar{\mu}(\gamma)$ and $\underline{\mu}(\gamma)$ of μ under signal realizations 1 and 0, respectively, assign probabilities

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

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

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

Thus, the marginal value of γ at μ satisfies

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

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 $\hat{q}\bar{\mu}'_n(0) + (1 - \hat{q})\underline{\mu}'_n(0) = 0$ for each n , this yields

$$V'_\mu(0) = 0. \tag{32}$$

The second derivative satisfies

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

Evaluating this at $\gamma = 0$ yields

$$V''_\mu(0) = \frac{1}{\hat{q}(1 - \hat{q})} \sum_{n,m} \partial_{n,m}^2 v(\mu) \mu_n(\omega_n - \mu \cdot \omega) \mu_m(\omega_m - \mu \cdot \omega). \tag{33}$$

To prove Lemma 7, 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 (32), but $V''_\mu(0) > 0 = C''(0)$ by (33) 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 , $\gamma_{\hat{q}}(\mu) > 0$, as required.

For the “moreover” part of Lemma 7, it is clear that (9) is necessary for learning to be successful at all states ω^* when $\hat{q} = q$. To see that (9) is sufficient, fix any true state ω^* and suppose that $\hat{q} = q$. Then $P_\mu(\cdot) = \hat{P}_\mu(\cdot|\omega^*)$ for all μ , and by (9), $\hat{P}_\mu(\cdot|\omega^*) \neq \hat{P}_\mu(\cdot|\omega)$ 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^* .

D.7 Proof of Lemma 8

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^* .

D.8 Proof of Lemma 9

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

D.9 Proof of Proposition 4

We will invoke the following lemma:

Lemma 12. *Fix any true state $\omega^* \in \Omega$, $p \in (0, 1)$, and $n = 1, \dots, N$. If $F(\theta_n^*) > \hat{F}(\theta_n^*)$, then there exists a neighborhood $B_n \ni \delta_{\omega_n}$ such that $\omega_\ell \succ_\mu^p \omega_k$ 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 $\omega_k \succ_\mu^p \omega_\ell$ for all ℓ, k with $\ell < k$ and all mixed $\mu \in B_n$.*

Proof. We consider the case $F(\theta_n^*) > \hat{F}(\theta_n^*)$. The argument when $F(\theta_n^*) < \hat{F}(\theta_n^*)$ is analogous. By Lemma 9, 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' . Moreover, by the monotone likelihood ratio assumption on private signals, we have $\frac{\hat{P}_\mu(0|\omega_k)}{\hat{P}_\mu(0|\omega_\ell)} < \frac{\hat{P}_\mu(1|\omega_k)}{\hat{P}_\mu(1|\omega_\ell)}$ for any ℓ, k with $\ell < k$ and mixed μ .

Thus, for all $\mu \in B_n$ and ℓ, k with $\ell < k$, we have

$$\sum_z P_\mu(z) \left(\frac{\hat{P}_\mu(z|\omega_k)}{\hat{P}_\mu(z|\omega_\ell)} \right)^p < \sum_z \hat{P}_\mu(z|\omega_\ell) \left(\frac{\hat{P}_\mu(z|\omega_k)}{\hat{P}_\mu(z|\omega_\ell)} \right)^p < 1,$$

where the first inequality holds by the observations in the previous paragraph and the second inequality holds by Jensen's inequality. Thus, $\omega_\ell \succ_\mu^p \omega_k$, as claimed. \square

We now prove Proposition 4. For the first part, note that if $F(\theta_N^*) < \hat{F}(\theta_N^*)$, then Lemma 12 yields some neighborhood $B \ni \delta_{\omega_N}$ such that $\omega_N \succ_\mu^p \omega_k$ for all $k \neq N$ and mixed $\mu \in B$, while if $F(\theta_N^*) > \hat{F}(\theta_N^*)$, then Lemma 12 yields a neighborhood $B \ni \delta_{\omega_N}$ such that $\omega_1 \succ_\mu^p \omega_N$ for all mixed $\mu \in B$. Thus, by Theorems 1-2, δ_{ω_N} is locally stable in the former case and unstable in the latter. Finally, if $F >_{FOSD} \hat{F}$, then Lemma 12 implies that for each n , there is a neighborhood $B_n \ni \delta_{\omega_n}$ such that $\omega_k \succ_\mu^p \omega_\ell$ for all $\ell > k$ and mixed $\mu \in B_n$, so that condition (b) in Remark 2 is met when states are ranked in decreasing order. Moreover, condition (a) in Remark 2 is satisfied by the monotone likelihood ratio assumption on private signals and the monotonicity of utilities. Hence, by Theorem 4, δ_{ω_N} is globally stable.

The arguments for part 2 of Proposition 4 are analogous. Finally, for the third part, note that if $F(\theta_n^*) \neq \hat{F}(\theta_n^*)$, then Lemma 12 implies that for some neighborhood $B_n \ni \delta_{\omega_n}$, we either have $\omega_1 \succ_\mu^p \omega_n$ for all mixed $\mu \in B_n$ or $\omega_N \succ_\mu^p \omega_n$ for all mixed $\mu \in B_n$. In either case, δ_{ω_n} is unstable by Theorem 2, as claimed.

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 (2018): "Social Learning with Model Misspecification: A Framework and a Robustness Result," working paper.
- BOHREN, J. A., A. IMAS, AND M. ROSENBERG (2019): "The Dynamics of Discrimination: Theory and Evidence," *American Economic Review*, forthcoming.
- 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.
- 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 (2019): “Network Structure and Naive Sequential Learning,” .
- DEB, J., AND Y. ISHII (2016): “Reputation building under uncertain monitoring,” Cowles Foundation Discussion Paper.
- 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,” working paper.
- ESPONDA, I., D. POUZO, AND Y. YAMAMOTO (2019): “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 (2019): “Misinterpreting Others and the Fragility of Social Learning,” Cowles Foundation Discussion Paper.
- FUDENBERG, D., AND D. M. KREPS (1993): “Learning mixed equilibria,” *Games and Economic Behavior*, 5(3), 320–367.
- FUDENBERG, D., G. LANZANI, AND P. STRACK (2020): “Limits Points of Endogenous Misspecified Learning,” 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.
- GAGNON-BARTSCH, T., M. RABIN, AND J. SCHWARTZSTEIN (2018): “Channeled Attention and Stable Errors,” 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.
- HE, K. (2018): “Mislearning from Censored Data: The Gambler’s Fallacy in Optimal-Stopping Problems,” working paper.
- HEIDHUES, P., B. KOSZEGI, AND P. STRACK (2018): “Unrealistic expectations and misguided learning,” *Econometrica*, 86(4), 1159–1214.
- (2019): “Convergence in Misspecified Learning Models with Endogenous Actions,” working paper.
- 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.
- LEVY, G., R. RAZIN, AND A. YOUNG (2019): “Misspecified Politics and the Recurrence of Populism,” Discussion paper, Working Paper.

- MARTINELLI, C. (2006): “Would rational voters acquire costly information?,” *Journal of Economic Theory*, 129(1), 225–251.
- 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.
- MOSCARINI, G., AND L. SMITH (2002): “The law of large demand for information,” *Econometrica*, 70(6), 2351–2366.
- 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.
- 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.

Supplementary Appendix to “Stability and Robustness in Misspecified Learning Models”

Mira Frick, Ryota Iijima, and Yuhta Ishii

E Local stability and KL-dominance

Recall that Theorem 1 provides a condition for local stability in terms of p -dominance. The following example shows that one cannot replace p -dominance with KL-dominance, i.e., condition (7) from Section 4.2 does not ensure that δ_ω is locally stable:

Example 5. 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 P is continuous in μ . 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 that $\omega \succ_\mu^{\text{KL}} \omega'$. Thus, condition (7) 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

$$\mathbb{E}[\ell_{t+1} | (\mu_s)_{s=0, \dots, t}] = \begin{cases} \ell_t & \text{if } \log \frac{\mu_{t'}(\omega)}{\mu_{t'}(\omega')} < 1 \text{ for some } t' \leq t \\ f(\log \frac{\mu_t(\omega)}{\mu_t(\omega')}) \sqrt{\log \frac{\mu_t(\omega)}{\mu_t(\omega')} + 1} + (1 - f(\log \frac{\mu_t(\omega)}{\mu_t(\omega')})) \sqrt{\log \frac{\mu_t(\omega)}{\mu_t(\omega')} - 1} = \sqrt{\log \frac{\mu_t(\omega)}{\mu_t(\omega')}} & \text{otherwise.} \end{cases}$$

Thus, by Doob's convergence theorem, there is an L^∞ random variable ℓ_∞ such that $\ell_t \rightarrow \ell_\infty$ almost surely. 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, almost surely

$\tau < \infty$. Hence, almost surely there exists some t such that $\log \frac{\mu_t(\omega)}{\mu_t(\omega')} < 1$. This implies that δ_ω is unstable. \blacktriangle

F Iterated dominance under one-dimensional states

In this section, we focus on one-dimensional states and provide conditions under which iterated dominance yields a unique outcome.⁴⁶ Based on this, we illustrate how Theorem 3 unifies convergence results in existing models. We consider a state space of the form $\Omega = [\underline{\omega}, \bar{\omega}] \subseteq \mathbb{R}$ and allow the set of signals Z to be any measurable space. We assume that $\text{KL}(P_\mu(\cdot), \hat{P}_\mu(\cdot|\omega))$ is well-defined and continuous in μ and ω .⁴⁷

We say that **complements** holds if for all $\mu \geq_{\text{FOSD}} \mu'$ and $\omega > \omega'$, we have

$$\omega \succ_{\mu'}^{\text{KL}} \omega' \quad (\text{resp., } \omega \succ_{\mu'}^{\text{KL}} \omega') \implies \omega \succ_{\mu}^{\text{KL}} \omega' \quad (\text{resp., } \omega \succ_{\mu}^{\text{KL}} \omega'). \quad (34)$$

Likewise, **substitutes** holds if (34) is satisfied for all $\mu \geq_{\text{FOSD}} \mu'$ and $\omega' > \omega$. For each non-empty closed subset $\Omega' \subseteq \Omega$, let $m(\mu; \Omega')$ denote the set of \succ_{μ}^{KL} -maximal states in Ω' . Standard monotone comparative statics arguments imply that $m(\mu; \Omega')$ is FOSD-increasing in μ under complements, and FOSD-decreasing under substitutes.⁴⁸

Define a correspondence on Ω by $M(\omega) := m(\delta_\omega; \Omega)$ for each ω . Observe that ω is a fixed point of M if and only if δ_ω is a Berk-Nash equilibrium. The following result provides conditions under which iterated elimination of dominated states yields a unique outcome:

Proposition 6. *Suppose that the correspondence $\omega \mapsto m(\delta_\omega; \Omega')$ is single-valued for each compact interval Ω' , and one of the following holds:*

1. *complements holds and M has a unique fixed point $\hat{\omega}$;*
2. *substitutes holds and M^2 has a unique fixed point $\hat{\omega}$.*

Then $S^\infty(\Omega) = \{\hat{\omega}\}$.

Observe that the unique fixed point condition on M under complements is weaker than the analogous condition on M^2 under substitutes. These conditions are analogous to conditions for dominance solvability in games with strategic complements or substitutes, paralleling Milgrom and Roberts (1990) and Zimper (2007), respectively.

To illustrate how to verify the complements or substitutes condition in specific applications, we consider a model of single-agent active learning that nests several leading examples:

⁴⁶We thank Michihiro Kandori for suggesting this direction.

⁴⁷Here $\text{KL}(P_\mu(\cdot), \hat{P}_\mu(\cdot|\omega)) := \int_Z \log \frac{dP_\mu(\cdot)}{d\hat{P}_\mu(\cdot|\omega)} dP_\mu < \infty$ for each μ and ω , where $\frac{dP_\mu(\cdot)}{d\hat{P}_\mu(\cdot|\omega)}$ is the Radon-Nikodym derivative. The space of beliefs $\Delta(\Omega)$ is endowed with the weak convergence topology.

⁴⁸That is, $\max m(\mu; \Omega')$ and $\min m(\mu; \Omega')$ are FOSD-increasing in μ under complements, and decreasing under substitutes. Note that $m(\mu; \Omega')$ is non-empty and compact by the continuity of $\text{KL}(P_\mu, \hat{P}_\mu(\cdot|\omega))$ in ω and the compactness of Ω' .

Remark 5 (One-dimensional active learning model). Consider a model in which a single agent chooses an action a_t from an interval $A \subseteq \mathbb{R}$ every period. Assume that the true signal distribution is given by one of the following technologies: (i) $Z = \{0, 1\}$ and the probability that $z_t = 1$ is $f(a_t)$; or (ii) $Z = \mathbb{R}$ and the signal takes the form $z_t = f(a_t) + \varepsilon_t$, where f is a strictly increasing and continuously differentiable function, and ε_t is i.i.d and follows a full-support distribution that admits a log-concave and positive density. Assume that the agent's action is given by a FOSD-increasing and continuous function of the belief μ_t . To model the agent's misspecification, assume that in case (i), she perceives the probability of $z_t = 1$ to be $\hat{f}(a_t, \omega)$ at each ω , and in case (ii) she perceives the signal to take the form $z_t = \hat{f}(a_t, \omega) + \varepsilon_t$ at each ω and is correct about the distribution of ε_t , where \hat{f} is continuously differentiable and strictly increasing in (a, ω) .

Given belief μ , $\text{KL}(P_\mu(\cdot), \hat{P}_\mu(\cdot|\omega))$ is strictly quasi-convex in ω and achieves its minimum at ω if $f(a(\mu)) = \hat{f}(a(\mu), \omega)$ (by Gibbs' inequality). In particular, the correspondence $\omega \mapsto m(\delta_\omega; \Omega')$ is single-valued for each compact interval Ω' . One can verify that if $f'(a) \geq \frac{\partial \hat{f}}{\partial a}(a, \omega)$ for each (a, ω) then complements holds, while if $f'(a) \leq \frac{\partial \hat{f}}{\partial a}(a, \omega)$ for each (a, ω) then substitutes holds. \blacktriangle

The model in Remark 5 nests the monopoly pricing example (Example 1), where $f(a) = \omega^* - \beta a$ and $\hat{f}(a, \omega) = \omega - \hat{\beta} a$. Thus, complements holds when $\hat{\beta} \geq \beta$, and substitutes holds when $\hat{\beta} \leq \beta$. Given this, Proposition 1 follows from Proposition 6 combined with the finite approximation result below (Proposition 7). The model also accommodates other examples that are considered in the recent literature:

Example 6 (Effort choice by an overconfident agent (Heidhues, Koszegi, and Strack, 2018)). The state space $\Omega = [\underline{\omega}, \bar{\omega}]$ parametrizes the level of some economic fundamental, where ω^* denotes the true state. The output in period t is given by $Q_t = Q(a_t, \beta, \omega) + \varepsilon_t$ where a_t is the agent's effort in period t , β is the agent's ability, and ε_t is mean-zero noise that follows a full-support and log-concave density. While the true ability is given by β , the agent perceives it to be $\hat{\beta} > \beta$. In each period t , the agent chooses effort a_t from an interval $(\underline{a}, \bar{a}) \subseteq \mathbb{R}$ to maximize expected output.

As in Heidhues, Koszegi, and Strack (2018), assume Q is twice-continuously differentiable such that (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 $\kappa > 0$. These assumptions guarantee that the optimal action $a(\mu_t)$ is continuous and FOSD-increasing in belief $\mu_t \in \Delta(\Omega)$. Moreover, any state $\omega > \omega^*$ is dominated by ω^* , because $0 > Q(a, \beta, \omega^*) - Q(a, \hat{\beta}, \omega^*) > Q(a, \beta, \omega^*) - Q(a, \hat{\beta}, \omega)$ for all a . Therefore, up to one round of elimination, we can focus on the state space $[\underline{\omega}, \bar{\omega}]$ with $\bar{\omega} \leq \omega^*$. Using the notation in Remark 5, $f(a) = Q(a, \beta, \omega^*)$ and $\hat{f}(a, \omega) = Q(a, \hat{\beta}, \omega)$. Complements holds, because $f'(a) - \frac{\partial \hat{f}}{\partial a}(a, \omega) = Q_a(a, \beta, \omega^*) - Q_a(a, \hat{\beta}, \omega) \geq Q_a(a, \beta, \omega^*) - Q_a(a, \hat{\beta}, \omega^*) \geq 0$ as $\omega \leq \omega^*$. As in Heidhues, Koszegi, and Strack (2018), assume that there is a unique Berk-Nash equilibrium $\delta_{\hat{\omega}}$. Then the first part of Proposition 6 ensures that $S^\infty(\Omega) = \{\hat{\omega}\}$. \blacktriangle

Example 7 (Data censoring under the gambler's fallacy (He, 2018)). Each period consists of a two-stage decision problem. In the first stage, the output x_1 follows $N(m_1^*, \sigma^2)$. If the output is lower

than the agent's threshold a , then the second stage output x_2 is observed, which follows $N(m_2^*, \sigma^2)$. The agent knows the first-stage mean m_1^* and variance σ^2 in both stages, but is uncertain about the second-stage mean. Thus, the state space Ω is a closed interval $[m_2, \bar{m}_2]$ of values of second-stage means, with true state $\omega^* = m_2^*$. While in reality there is no correlation between x_1 and x_2 , the agent perceives negative correlation. In particular, the agent perceives that under $m_2 \in \Omega$, the second-stage draw conditional on the realization of x_1 follows $N(m_2 - \gamma(x_1 - m_1^*), \sigma^2)$, where $\gamma \geq 0$ captures the extent of the gambler's fallacy.

The agent chooses the cutoff $a \in \mathbb{R}$ to maximize the expected value of the period payoff $u : \mathbb{R} \times (\mathbb{R} \cup \{\emptyset\}) \rightarrow \mathbb{R}$ 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 (2018), the threshold a is FOSD-increasing in the belief $\mu \in \Delta(\Omega)$. Using the notation of Remark 5, $f(a) = \omega^*$ and $\hat{f}(a, \omega) = \omega - \gamma(\mathbb{E}[x_1 | x_1 \leq a] - m_1^*)$. Thus, complements holds because $f'(a) - \frac{\partial \hat{f}}{\partial a}(a, \omega) = \gamma \frac{\partial \mathbb{E}[x_1 | x_1 \leq a]}{\partial a} \geq 0$. As He (2018) shows, there is a unique Berk-Nash equilibrium. Hence, the first part of Proposition 6 implies that iterated elimination yields a unique outcome.⁴⁹ \blacktriangle

While many economic examples (including those above) feature a continuous state space $\Omega = [\underline{\omega}, \bar{\omega}]$, our stability results are formulated for finite state spaces. However, the following result shows that iterated elimination under a continuous state space can be approximated by that under finite state spaces. Thus, Theorem 3 applies to such economic examples up to discretizing the state space.

Proposition 7. *Suppose the assumptions in Proposition 6 are satisfied, and thus $S^\infty(\Omega) = \{\hat{\omega}\}$. Then for any $\eta > 0$, there exists $\delta > 0$ such that $S^\infty(\Omega') \subseteq [\hat{\omega} - \eta, \hat{\omega} + \eta]$ for any finite subset $\Omega' \subseteq \Omega$ that is δ -dense in Ω .*

F.1 Proofs

F.1.1 Proof of Proposition 6

Since the mapping $\omega \mapsto m(\delta_\omega; \Omega')$ is single-valued, it is also continuous by the continuity of $\text{KL}(P_\mu, \hat{P}_\mu(\cdot | \omega))$ in ω and the maximum theorem. Moreover, the mapping is increasing (resp. decreasing) under complements (resp. substitutes). We prove two preliminary lemmas:

Lemma 13. *For any compact interval $\Omega' \subseteq \Omega$, we have $S(\Omega') = \{m(\delta_\omega; \Omega') : \omega \in \Omega'\}$.*

Proof. Let $\Omega' := [\underline{\omega}', \bar{\omega}']$. Note that since $m(\cdot; \Omega')$ is continuous and monotonic, we have $\{m(\delta_\omega; \Omega') : \omega \in \Omega'\} = \text{co}\{m(\delta_{\underline{\omega}'}; \Omega'), m(\delta_{\bar{\omega}'}; \Omega')\}$, where co denotes the closed convex hull. Thus, by definition of S , it suffices to verify that $S(\Omega') \subseteq \text{co}\{m(\delta_{\underline{\omega}'}; \Omega'), m(\delta_{\bar{\omega}'}; \Omega')\}$.

⁴⁹He (2018) also considers the case in which the agent updates beliefs about the first-stage mean m_1 , 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 , and the same argument as above applies.

Take any $\hat{\omega} \notin \text{co}\{m(\delta_{\underline{\omega}'}; \Omega'), m(\delta_{\bar{\omega}'}; \Omega')\}$. If $m(\delta_{\underline{\omega}'}; \Omega') = m(\delta_{\bar{\omega}'}; \Omega') =: \hat{\omega}'$ then $m(\mu; \Omega') = \hat{\omega}'$ for all $\mu \in \Delta(\Omega')$ by the monotonicity of $m(\cdot; \Omega')$. Therefore, $\hat{\omega}' \succ_{\mu}^{\text{KL}} \hat{\omega}$ for all $\mu \in \Delta(\Omega')$, so $\hat{\omega} \notin S(\Omega')$. Now consider the case $m(\delta_{\underline{\omega}'}; \Omega') \neq m(\delta_{\bar{\omega}'}; \Omega')$. We focus on the case $\hat{\omega} > \max\{m(\delta_{\underline{\omega}'}; \Omega'), m(\delta_{\bar{\omega}'}; \Omega')\} =: \hat{\omega}'$; the argument for the remaining case $\hat{\omega} < \min\{m(\delta_{\underline{\omega}'}; \Omega'), m(\delta_{\bar{\omega}'}; \Omega')\}$ is symmetric. If $\hat{\omega}' = m(\delta_{\bar{\omega}'}; \Omega')$, then $\hat{\omega}' \succ_{\delta_{\bar{\omega}'}}^{\text{KL}} \hat{\omega}$. In this case complements holds, so $\hat{\omega}' \succ_{\mu}^{\text{KL}} \hat{\omega}$ for any $\mu \in \Delta(\Omega')$. If $\hat{\omega}' = m(\delta_{\underline{\omega}'}; \Omega')$, then $\hat{\omega}' \succ_{\delta_{\underline{\omega}'}}^{\text{KL}} \hat{\omega}$. In this case substitutes holds, so $\hat{\omega}' \succ_{\mu}^{\text{KL}} \hat{\omega}$ for any $\mu \in \Delta(\Omega')$. In either case, $\hat{\omega} \notin S(\Omega')$. \square

Lemma 14. *We have $S^k(\Omega) = \{M(\omega) : \omega \in S^{k-1}(\Omega)\}$ for each k .*

Proof. We proceed inductively. For $k = 1$, the claim is immediate from Lemma 13. Suppose the claim holds for k . For each $\omega \in S^k(\Omega)$, this implies $M(\omega) \in S^k(\Omega)$ and hence $M(\omega) = m(\delta_{\omega}; S^k(\Omega))$. Thus, $\{M(\omega) : \omega \in S^k(\Omega)\} = \{m(\delta_{\omega}; S^k(\Omega)) : \omega \in S^k(\Omega)\} = S^{k+1}(\Omega)$ by Lemma 13. \square

Proof of Proposition 6. Since M is a continuous function, Lemma 14 implies that $(S^k(\Omega))_{k \in \mathbb{N}}$ is a sequence of closed intervals that are decreasing with respect to set-inclusion, which we denote by $([\underline{\omega}_k, \bar{\omega}_k])_{k \in \mathbb{N}}$. The limits $\lim_k \underline{\omega}_k$ and $\lim_k \bar{\omega}_k$ are well-defined. To prove the proposition, it suffices to prove that $\lim_k \underline{\omega}_k = \lim_k \bar{\omega}_k$. Suppose for a contradiction that $\lim_k \underline{\omega}_k < \lim_k \bar{\omega}_k$.

Consider the case of complements. Then $S^k(\Omega) = [\underline{\omega}_k, \bar{\omega}_k] = [M(\underline{\omega}_{k-1}), M(\bar{\omega}_{k-1})]$ for every $k \geq 1$. By continuity of M , we have $\lim_k \underline{\omega}_k = M(\lim_k \underline{\omega}_k)$ and $\lim_k \bar{\omega}_k = M(\lim_k \bar{\omega}_k)$. If $\lim_k \underline{\omega}_k < \lim_k \bar{\omega}_k$ then both $\lim_k \underline{\omega}_k$ and $\lim_k \bar{\omega}_k$ are fixed points of M , a contradiction.

Consider the case of substitutes. Then $S^k(\Omega) = [\underline{\omega}_k, \bar{\omega}_k] = [M^2(\underline{\omega}_{k-2}), M^2(\bar{\omega}_{k-2})]$ for every $k \geq 2$. By continuity of M , we have $\lim_k \underline{\omega}_k = M^2(\lim_k \underline{\omega}_k)$ and $\lim_k \bar{\omega}_k = M^2(\lim_k \bar{\omega}_k)$. If $\lim_k \underline{\omega}_k < \lim_k \bar{\omega}_k$ then both $\lim_k \underline{\omega}_k$ and $\lim_k \bar{\omega}_k$ are fixed points of M^2 , a contradiction. \square

F.1.2 Proof of Proposition 7

We begin with a preliminary lemma:

Lemma 15. *Take any interval $\Omega' := [\underline{\omega}', \bar{\omega}'] \subseteq \Omega$ and $\eta > 0$. Then there exists $\delta > 0$ such that for any finite set Ω'' that is δ -dense in Ω' and included in $[\underline{\omega}' - \delta, \bar{\omega}' + \delta]$, $S(\Omega'')$ is η -dense in $S(\Omega') =: [\underline{s}, \bar{s}]$ and is included in $[\underline{s} - \eta, \bar{s} + \eta]$.*

Proof. We first show that for any small enough δ , $S(\Omega'') \subseteq [\underline{s} - \eta, \bar{s} + \eta]$. Recall from the proof of Lemma 13 that (i) $\underline{s} \succ_{\mu}^{\text{KL}} \omega$ for all $\omega \in [\underline{\omega}', \underline{s} - \eta]$ and $\mu \in \Delta(\Omega')$, and (ii) $\bar{s} \succ_{\mu}^{\text{KL}} \omega$ for all $\omega \in [\bar{s} + \eta, \bar{\omega}']$ and $\mu \in \Delta(\Omega')$. By continuity of $\text{KL}(P_{\mu}, \hat{P}_{\mu}(\cdot|\omega))$ in (ω, μ) , there exists $\delta \in (0, \eta)$ such that (i) $\omega' \succ_{\mu}^{\text{KL}} \omega$ for all $\omega \in [\underline{\omega}' - \delta, \underline{s} - \eta]$, $\omega' \in [\underline{s} - \delta, \underline{s} + \delta]$, and $\mu \in \Delta([\underline{\omega}' - \delta, \bar{\omega}' + \delta])$, and (ii) $\omega' \succ_{\mu}^{\text{KL}} \omega$ for all $\omega \in [\bar{s} + \eta, \bar{\omega}' + \delta]$, $\omega' \in [\bar{s} - \delta, \bar{s} + \delta]$, and $\mu \in \Delta([\underline{\omega}' - \delta, \bar{\omega}' + \delta])$. Thus, for any $\Omega'' \subseteq [\underline{\omega}' - \delta, \bar{\omega}' + \delta]$ that is δ -dense in Ω' , (i) there exists $\omega' \in \Omega'' \cap [\underline{s} - \delta, \underline{s} + \delta]$, which ensures $\omega \notin S(\Omega'')$ for any $\omega \leq \underline{s} - \eta$, and (ii) there exists $\omega' \in \Omega'' \cap [\bar{s} - \delta, \bar{s} + \delta]$, which ensures $\omega \notin S(\Omega'')$ for any $\omega \geq \bar{s} + \eta$. Hence, $S(\Omega'') \subseteq [\underline{s} - \eta, \bar{s} + \eta]$.

Next, we show that $S(\Omega'')$ is η -dense in $S(\Omega')$ if we take δ small enough. For this, fix any subinterval $[\underline{s}', \bar{s}']$ of $S(\Omega')$ that has length η . It suffices to show that $S(\Omega'') \cap [\underline{s}', \bar{s}'] \neq \emptyset$ for all small enough δ . Fix any $s \in (\underline{s}', \bar{s}')$. By Lemma 13 and the monotonicity and continuity of m , there exists $\omega' \in \Omega'$ such that $s = m(\delta_{\omega'}; \Omega')$. Thus, $s \succ_{\delta_{\omega'}}^{\text{KL}} \omega$ for all $\omega \in \Omega'$. By continuity of KL, for all small enough δ , we have $s \succ_{\delta_{\omega'}}^{\text{KL}} \omega$ for all $\omega \in [\underline{\omega}' - \delta, \bar{\omega}' + \delta] \setminus (\underline{s}', \bar{s}')$. By setting $\delta \in (0, \eta)$ sufficiently small, the continuity of KL ensures that $\tilde{s} \succ_{\delta_{\tilde{\omega}'}}^{\text{KL}} \omega$ for all $\tilde{s} \in [s - \delta, s + \delta]$, $\tilde{\omega}' \in [\omega' - \delta, \omega' + \delta]$, and $\omega \in [\underline{\omega}' - \delta, \bar{\omega}' + \delta] \setminus (\underline{s}', \bar{s}')$. Thus, for any $\Omega'' \subseteq [\underline{\omega}' - \delta, \bar{\omega}' + \delta]$ that is δ -dense in Ω' , since there exists $\tilde{\omega}' \in \Omega'' \cap [\omega' - \delta, \omega' + \delta]$, we have $m(\delta_{\tilde{\omega}'}; \Omega'') \subseteq \Omega'' \cap [s - \delta, s + \delta] \neq \emptyset$. This ensures that $S(\Omega'') \cap [s - \delta, s + \delta] \neq \emptyset$, which for $\delta \in (0, \eta)$ small enough implies $S(\Omega'') \cap [\underline{s}', \bar{s}'] \neq \emptyset$. \square

Proof of Proposition 7. As observed in the proof of Proposition 6, each $S^k(\Omega)$ is given by a compact interval $[\underline{\omega}_k, \bar{\omega}_k]$. Since $S^\infty(\Omega) = \{\hat{\omega}\}$, there exists $K > 0$ such that $S^K(\Omega) \subseteq [\hat{\omega} - \eta/2, \hat{\omega} + \eta/2]$. By Lemma 15, there exists $\delta_K > 0$ such that $S^K(\Omega') \subseteq [\hat{\omega} - \eta, \hat{\omega} + \eta]$ holds if $S^{K-1}(\Omega')$ is δ_K -dense in $S^{K-1}(\Omega)$ and included in $[\underline{\omega}_{K-1} - \delta_K, \bar{\omega}_{K-1} + \delta_K]$. Inductively, given $\delta_k > 0$ for $k \in \{2, \dots, K\}$, Lemma 15 ensures the existence of $\delta_{k-1} > 0$ such that $S^{k-1}(\Omega')$ is δ_k -dense in $S^{k-1}(\Omega)$ and included in $[\underline{\omega}_{k-1} - \delta_k, \bar{\omega}_{k-1} + \delta_k]$, whenever $S^{k-2}(\Omega')$ is δ_{k-1} -dense in $S^{k-2}(\Omega)$ and included in $[\underline{\omega}_{k-2} - \delta_{k-1}, \bar{\omega}_{k-2} + \delta_{k-1}]$. Hence, the desired conclusion $S^\infty(\Omega') \subseteq S^K(\Omega') \subseteq [\hat{\omega} - \eta, \hat{\omega} + \eta]$ holds by taking $\delta = \delta_1$. \square

G Extension to heterogeneous beliefs

We briefly describe an extension of our model that allows for profiles of beliefs. We omit all proofs as they follow analogous arguments as the original results. There is a finite set I and for each $i \in I$ a finite state space Ω^i . At each $t = 0, 1, 2, \dots$, there is a belief profile $\bar{\mu}_t = (\mu_t^i)_{i \in I} \in \prod_{i \in I} \Delta(\Omega^i)$. Given any initial belief profile $\bar{\mu}_0 = (\mu_0^i)_{i \in I}$, where each μ_0^i has full support on Ω^i , profile $\bar{\mu}_t$ evolves according to the following Markov process over $\prod_{i \in I} \Delta(\Omega^i)$.

At the end of each period t , a signal z_t from a finite set Z is drawn according to the true distribution $P_{\bar{\mu}_t}(\cdot) \in \Delta(Z)$, which can depend on the current belief profile. Following the realization of z_t , each belief μ_t^i is updated via Bayes' rule to the belief μ_{t+1}^i according to i 's perceived conditional signal distributions: That is, $\mu_{t+1}^i(\omega^i) = \frac{\mu_t^i(\omega^i) \hat{P}_{\bar{\mu}_t}^i(z_t | \omega^i)}{\sum_{\tilde{\omega}^i \in \Omega^i} \mu_t^i(\tilde{\omega}^i) \hat{P}_{\bar{\mu}_t}^i(z_t | \tilde{\omega}^i)}$ for each $\omega^i \in \Omega^i$, where $\hat{P}_{\bar{\mu}}^i(\cdot | \omega^i) \in \Delta(Z)$ is i 's perceived distribution conditional on state $\omega^i \in \Omega^i$ at belief profile $\bar{\mu} \in \prod_{i \in I} \Delta(\Omega^i)$.

We impose the following condition extending Assumption 1:

Assumption 2.

1. For each $i \in I$, $\omega^i \in \Omega^i$, and $\bar{\mu}$, $\text{supp} P_{\bar{\mu}}(\cdot) \subseteq \text{supp} \hat{P}_{\bar{\mu}}^i(\cdot | \omega^i)$.
2. For each $i \in I$, $\omega^i, \tilde{\omega}^i \in \Omega^i$, $\sup_{\bar{\mu}, z \in \text{supp} P_{\bar{\mu}}} \frac{\hat{P}_{\bar{\mu}}^i(z | \omega^i)}{\hat{P}_{\bar{\mu}}^i(z | \tilde{\omega}^i)} < \infty$.
3. For each $i \in I$ and $\omega^i \in \Omega^i$, $P_{\bar{\mu}}(\cdot)$ and $\hat{P}_{\bar{\mu}}^i(\cdot | \omega^i)$ are continuous in $\bar{\mu}$.

Let $\mathbb{P}_{\bar{\mu}}$ denote the probability measure over sequences $(\bar{\mu}_t)$ under the initial belief profile $\bar{\mu}$. Stability notions can be extended in the following manner:

Definition 3. Belief profile $\bar{\mu}^*$ is:

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

The following result extends Lemma 1 (instability of mixed beliefs):

Lemma 16. Consider any $j \in I$ and mixed belief $\mu^j \in \Delta(\Omega^j)$. Assume that there exist $\omega^j, \tilde{\omega}^j \in \text{supp}(\mu^j)$ such that $\hat{P}_{\bar{\mu}}^j(z|\omega^j) \neq \hat{P}_{\bar{\mu}}^j(z|\tilde{\omega}^j)$ for some $z \in \text{supp} P_{\bar{\mu}}$. Then $\bar{\mu} = (\mu^j, \mu^{-j})$ is unstable for all $\mu^{-j} \in \prod_{i \neq j} \Delta(\Omega^i)$.

For each $\bar{\mu} \in \prod_{i \in I} \Delta(\Omega^i)$, $i \in I$, and $\omega^i, \tilde{\omega}^i \in \Omega^i$, we define the **KL-dominance** order by

$$\omega^i \succsim_{\bar{\mu}}^{i, \text{KL}} \tilde{\omega}^i \iff \sum_z P_{\bar{\mu}}(z) \log \left(\frac{\hat{P}_{\bar{\mu}}^i(z|\tilde{\omega}^i)}{\hat{P}_{\bar{\mu}}^i(z|\omega^i)} \right) \leq 0,$$

and write $\omega^i \succ_{\bar{\mu}}^{i, \text{KL}} \tilde{\omega}^i$ if the inequality is strict. Likewise, given any $p > 0$, we define the **p-dominance** order by

$$\omega^i \succsim_{\bar{\mu}}^{i, p} \tilde{\omega}^i \iff \sum_z P_{\bar{\mu}}(z) \left(\frac{\hat{P}_{\bar{\mu}}^i(z|\tilde{\omega}^i)}{\hat{P}_{\bar{\mu}}^i(z|\omega^i)} \right)^p \leq 1,$$

and again write $\omega^i \succ_{\bar{\mu}}^{i, p} \tilde{\omega}^i$ if the inequality is strict.

The following result extends Theorems 1–2 (local stability and instability).

Theorem 5. Consider any $(\omega^i)_{i \in I} \in \prod_{i \in I} \Omega^i$.

1. Belief profile $(\delta_{\omega^i})_{i \in I}$ is locally stable if there exists $p > 0$ and a neighborhood $B \ni (\delta_{\omega^i})_{i \in I}$ such that

$$\omega^j \succ_{\bar{\mu}}^{j, p} \tilde{\omega}^j \text{ for all } j \in I, \tilde{\omega}^j \neq \omega^j, \text{ and } \bar{\mu} \in B \setminus \{(\delta_{\omega^i})_{i \in I}\}.$$

2. Belief profile $(\delta_{\omega^i})_{i \in I}$ is unstable if there exists a neighborhood $B \ni (\delta_{\omega^i})_{i \in I}$ such that

$$\text{for some } j \in I \text{ and } \tilde{\omega}^j \neq \omega^j, \text{ we have } \tilde{\omega}^j \succ_{\bar{\mu}}^{j, \text{KL}} \omega^j \text{ for all } \bar{\mu} \in B \setminus \{(\delta_{\omega^i})_{i \in I}\}.$$

To extend Theorem 3, we say that $K \subseteq \prod_{i \in I} \Delta(\Omega^i)$ is a **globally stable set** of belief profiles if $\mathbb{P}_{\bar{\mu}}[\inf_{\bar{\nu} \in K} \|\bar{\mu}_t - \bar{\nu}\| \rightarrow 0] = 1$ for every initial belief profile $\bar{\mu}$. For each product of subsets $\prod_{j \in I} \tilde{\Omega}^j \subseteq \prod_{j \in I} \Omega^j$ and $i \in I$, let

$$\begin{aligned} S^i \left(\prod_{j \in I} \tilde{\Omega}^j \right) &:= \left\{ \omega^i \in \tilde{\Omega}^i : \exists \tilde{\omega}^i \in \tilde{\Omega}^i \text{ s.t. } \tilde{\omega}^i \succ_{\bar{\mu}}^{i, \text{KL}} \omega^i \text{ for all } \bar{\mu} \in \prod_{j \in I} \Delta(\tilde{\Omega}^j) \right\} \\ &= \left\{ \omega^i \in \tilde{\Omega}^i : \exists \tilde{\omega}^i \in \tilde{\Omega}^i \text{ and } p > 0 \text{ s.t. } \tilde{\omega}^i \succ_{\bar{\mu}}^{i, p} \omega^i \text{ for all } \bar{\mu} \in \prod_{j \in I} \Delta(\tilde{\Omega}^j) \right\}. \end{aligned}$$

Then for $\bar{\Omega} := \prod_{j \in I} \Omega^j$, recursively define $S^{i,0}(\bar{\Omega}) := \bar{\Omega}$, $S^{i,k+1}(\bar{\Omega}) := S^i(\prod_{j \in I} S^{j,k}(\bar{\Omega}))$ for all $k = 0, 1, \dots$, and $S^{i,\infty}(\bar{\Omega}) := \bigcap_{k \in \mathbb{N}} S^{i,k}(\bar{\Omega})$.

Theorem 6. *The set $\prod_{i \in I} \Delta(S^{i,\infty}(\bar{\Omega}))$ is globally stable.*