DIMENSIONALITY AND DISAGREEMENT:

Asymptotic belief divergence in response to common information

Isaac Loh* Gregory Phelan†

This version: July 11, 2017

Abstract

We provide a model of boundedly rational, multidimensional learning in which divergence occurs when hard to identify observations are frequent, but convergence occurs when such observations are rare. Because the information is of lower dimension than the model, agents face an identification problem affecting the role of data in inference. Because the model is high-dimensional, boundedly rational agents maintain marginal beliefs rather than beliefs over the whole state space. With this dimension-reducing modification, Bayesian learning ceases to be almost surely consistent and incorrect beliefs might be infinitely persistent. We show that consistency occurs when easily identifiable observations are frequent. However, we also provide conditions under which this learning process is dependent on initial conditions and asymptotically inconsistent with positive probability. Robustly, two agents with differing priors who observe identical, unambiguous data may disagree forever, with stronger disagreement the more data observed.

Keywords: Heterogeneous beliefs, divergence, learning, Bayesian updating, bounded rationality, sparsity.

JEL classification: D10, D80, D84.

*Northwestern University. Email: IsaacLoh2015@u.northwestern.edu .
†Department of Economics, Williams College. Email: gp4@williams.edu.
We are grateful for feedback from Matthew Chao, Eddie Dekel, Yingni Guo, Asim Khwaja, Rachel Kranton, Bruce Sacerdote, Larry Samuelson, and Arunava Sen.

1 Introduction

People disagree and sometimes in big, persistent ways: people disagree about which policies best achieve outcomes; professional forecasters and central bankers make different forecasts for macroeconomic variables (Crump et al., 2014; Andrade et al., 2016);[1] and many patterns in financial markets strongly suggest investor disagreement (Hong and Stein, 2007). Furthermore, disagreements sometimes grow as people see the same information and continue to disagree.[2] Many instances of disagreement are publicly known and persistent (e.g., political and social beliefs), suggesting that priors are different.[3] And yet, in many cases disagreements do *not* grow, and people appear to perceive evidence in similar ways (Gerber and Green, 1999). In this paper, we provide a model that can explain why disagreement grows in some cases but not others. As a result of difficulties arising in multidimensional learning, common observations can lead to permanent divergence in arbitrarily small initial disagreements.

Consider a multi-armed bandit learning problem with d arms, each of which is randomly chosen to be good (high probability of payoff) or bad (low probability) and then fixed in its selected state. An agent wishes to learn the state of each arm, but suppose that in every period she is constrained to pull all d arms in unison and observe their *collective payoff* rather than the payoff to each arm. Importantly, the integer signal observed by the agent is one-dimensional but the underlying model is multidimensional. This introduces an identification problem for doing inference. Suppose 2 arms are pulled. Observations of "two" or "zero" successes are clearly "identified" (both success or both failure), while observations of "one" are "unidentified." When interpreting

[1] Crump et al. (2014) find that FOMC members disagree about projections for the Fed Funds Rate and other variables even more than forecasters in the Survey of Primary Dealers. Furthermore, economists rarely switch from hawks to doves (Malmendier et al., 2017).

[2] There is ample empirical evidence that people sometimes interpret evidence differently: religion (Batson, 1975); death penalty (Lord et al., 1979); nuclear power (Plous, 1991); caffeine (Chaiken et al., 1992); sexuality (Munro and Ditto, 1997); affirmative action and gun violence (Taber and Lodge, 2006). Darley and Gross (1983) provide evidence of people interpreting the same evidence differently in light of earlier evidence, and Hirshleifer and Teoh (2003) document how the presentation of accounting information affects its interpretation. Malmendier et al. (2017) show that personal experiences of inflation strongly influence the hawkish or dovish leanings of central bankers, which is evidence that priors influence how FOMC members interpret the same data.

[3] Since Aumann (1976) it is well-known that people cannot "agree to disagree" when beliefs are common knowledge and people have common priors, and Geanakoplos and Polemarchakis (1982) show that people with common priors cannot disagree forever so long as they can communicate their beliefs (the result holds over finite spaces).

these "unidentified observations" (ones), the likelihood function for one arm depends on belief about the other arm. Hence, some observations will be rationally used as "evidence" in different ways depending on an agent's beliefs.

A typical assumption in multidimensional learning problems is that agents maintain their priors over the entire joint distribution of variables giving the state of the world. A canonical Bayesian agent operating within this setup maintains a prior over the joint distribution of possible arm configurations: this agent must at each instance recall 2^d numbers describing the believed probability that arm 1 is good/arm 2 is bad/.../arm d is good, arm 1 is good/arm 2 is good/.../arm d is good, and so on. Of course this rapidly becomes unfeasible as d grows large. In contrast, the signals we observe every day are shaped by innumerable variables constituting the state of the world, and storing a joint prior distribution across all possible states is in many cases impractical. One appealing deviation that boundedly rational agents may take is to maintain priors only over marginal distributions. The agent that we consider only recalls the updated probability that arm $i \in \{1,\dots,d\}$ is good. As an example, people very commonly hold beliefs on whether the Democratic and Republic parties individually are "good" (or "bad"). But one might suppose it likely that people are more hard-pressed to give their stored beliefs on a state like Democrats good/Republicans bad. In addition, it may be plausible that most would simply combine their marginal beliefs in some way to obtain the joint distribution if it was ever required. (Indeed, we show that maintaining beliefs in this way can be approximately optimal.)

In this reduced-dimension setting, we show that (i) updated posterior beliefs may converge but not necessarily to the true values (divergence), and (ii) divergence is likely to occur when unidentified observations occur frequently, but beliefs converge to the truth when unidentified observations are limited. When divergence occurs, beliefs are likely to converge to values confirming initial beliefs whether those beliefs are correct or not. Whether disagreements can persist depends on the extent to which agents can use the same observations to draw inference about different variables. Since agents update beliefs differently when observing unidentified observations, the likelihood that disagreements persist depends on the severity of the identification problem. Indeed, we show that when observations of this type are frequent, with *many* observations people may continue to disagree, and disagree further, perhaps even becoming completely convinced of their beliefs. In other words, disagreement can persist or increase even when people observe identical evidence. Furthermore, agents with common priors may have collective beliefs converge to something other than the truth, and heterogeneous priors may diverge in light of common information

To show these results, we present a general version of this simple model described and show that the given intuitions hold in a general setting. We theoretically characterize the limiting properties of beliefs in our model and show that there are robust conditions under which beliefs converge to the wrong values, leading to asymptotic divergence. Furthermore, under certain conditions beliefs diverge asymptotically with probability one. In these cases, more data leads to greater divergence rather than greater agreement. We show that high-dimensionality may lead to asymptotic divergence when unidentified observations are likely, but will lead to asymptotic agreement when identified observations are more abundant.

The rest of the paper is structured as follows. The remainder of this section discusses our connection to the existing literature. Section 2 presents the general multi-dimensional discrete multivariate setting, provides a simple microfoundation for updating marginal probabilities (independence copula), and characterizes the limiting properties of beliefs. We also include more detailed results in a simplified 2x2 setting. Section 3 discusses the setup and our results, including implications of our results for persuasion, information release and design, and model dynamics. The final section concludes.

Related Literature

While we emphasize asymptotic results on belief divergence, several papers have considered how observing a small number of signals can increase belief disagreement. Benoît and Dubra (2014) show that as a result of identification problems of the type we consider, disagreement increases for intermediate values of information, but not for extreme values of information (which are more informative of the underlying identification structure).[4] Similarly, Baliga et al. (2013) show that

[4]Similar observations are made by Dixit and Weibull (2007) and Jern et al. (2014). Kondor (2012) shows that public information can increase disagreement when traders have private information and different trading horizons (they update higher-order expectations about beliefs of others).

polarization can occur as an optimal response to ambiguity aversion, a possibility only at signal with an intermediate likelihood ratio. Andreoni and Mylovanov (2012) also consider polarization about optimal actions when agents receive two-dimensional information to form a one-dimensional opinion. They hypothesize that agents receive information about only one of two states, rather than information that is a function of the two states. Disagreement persists because agents discount information filtered through the actions of others).

Our asymptotic result holds when the true state of the world is in an "intermediate region" which will lead toward a higher frequency of hard to identify signals. Thus, disagreement is more likely to persist asymptotically when information is not extremal, balanced by many unidentified observations. This accords with the analogous intermediate results of Benoît and Dubra (2014) and Baliga et al. (2013)

Our asymptotic result differs from Acemoglu et al. (2016), who show that *agreement* does not necessarily follow even when learning does. They consider a model in which agents have different priors about the state, and they also have different beliefs about how the state maps to signals; in other words, they disagree about how much noise there is. As a result, agents also face an identification problem regarding how to interpret signals. They show that even though agents learn the true state asymptotically, this does not mean that agents will *agree* asymptotically because the likelihood ratios of their beliefs need not converge. In contrast, in our model agents need not learn the truth asymptotically, and disagreement can be exacerbated by the same information. The results differ because the nature of the identification problems we consider are different. Nonetheless, our model can also deliver the result of asymptotic disagreement even when learning occurs. (See Proposition 3 and Lemma 2.)

Fryer et al. (2015) consider a model in which agents may receive "ambiguous" signals, which are interpreted in light of current priors (as in our model) and "stored" as an unambiguous signal. Their convergence result is driven by bounded memory, where agents do not retain whether signals were ambiguous or not, but remember the interpreted value. In the model, if ambiguous signals could be stored as such, learning (and agreement) would occur. In contrast, in our model signals can be considered ambiguous only insofar as they may be interpreted as evidence for several states of the world depending on prevailing prior beliefs. This flexibility of a signal evolves over time according to an agent's past observations. However, given a history of observations and joint distribution over states (which we suppose is reconstructed from stored marginal beliefs), the prescribed interpretation of a signal is clear and fixed; updating occurs by Bayes's rule. Proposition 3 of Fryer et al. (2015) requires agents 1 and 2 to start out with beliefs in certain regions in order to disagree asymptotically, whereas we find that for two agents the initial conditions presaging a possibility of disagreement are quite broad.

A critical ingredient in our model is that agents only store marginal beliefs, not the joint distribution of beliefs over the full state space (agents are boundedly rational, see for example Gabaix (2014)). Many papers provide theories with a behavioral assumption in which beliefs converge to incorrect beliefs. Rabin and Schrag (1999) consider confirmatory bias; Eyster and Rabin (2010) show how "social confirmation bias" leads to herding with positive probability to incorrect actions; Ortoleva and Snowberg (2015) consider the role of overconfidence in political behavior to explain ideological extremeness; Schwartzstein (2014) shows that selective attention to information can lead to persistently incorrect beliefs (see also Sundaresan and Turban, 2014). Heidhues et al. (2015) consider a model in which agents are (overly) optimistic about their ability, choose an effort level, and observe the outcome. Beliefs about fundamentals diverge from reality precisely because agents change their effort level in a self-defeating, misguided way. Thus, incorrect learning follows because agents endogenously observe different signals based on their effort. In our model agents observe exactly the same signals and yet learn differently.[5]

Finally, models allowing for asymptotic disagreement with positive probability where Bayes' rule is used clasically must have carefully selected prior beliefs (disjoint support). Cromwell's rule may be violated to break asymptotic consistency á la Doob: if agents put belief zero on the truth, then they will never learn the truth. In contrast to this strategy, our setup has several advantages: (i) in our model, divergence is driven by the relative frequency of easy to identify observations,

[5]See Eyster et al. (2014) for a model in which learning follows by observing the actions of others, which are endogenously determined as a function of expected payoffs. Similarly, Sethi and Yildiz (2016) provide a model in which disagreement can persist asymptotically because agents decide which expert opinions to target, and the set of experts can be disjoint over time. However, when the precision of initial beliefs about the perspectives of others is above a certain threshold, path-dependency disappears and long-run efficiency occurs, which is similar in spirit to the result in our paper that convergence occurs for some states but not others.

which is a property of the fundamental state, not the coincidence of initial priors; (ii) our model allows us to say when agents are guaranteed to learn the truth, regardless of their priors; (iii) our results are robust in the sense that they accommodate an open set with positive measure of priors.

2 Main Model

This section provides a multivariate model with d arms (underlying states) that can take k values (probabilities of success). After the setup, we provide a bounded rationality microfoundation for updating marginals (Lemma 1). We then provide two main results. Proposition 1 provides conditions for when learning is asymptotically consistent, which occurs when easy to identify observations are most frequent. Proposition 2 and Corollary 1 show that wrong beliefs can be locally attractive even if agents don't necessarily combine beliefs over marginals by multiplying them (i.e., by assuming independence), but as long as the aggregation rule looks enough like multiplication in that the ratio of beliefs is bounded.

2.1 Setup

Let there be a multivariate state of the world $\theta = (\theta_1, \ldots, \theta_d)$ taking values in a discrete set $\Theta = \prod_{i=1}^{d} \Theta_i$. There are also random variables $X_1, \ldots, X_d$ which are distributed as $X_i \sim G_{\theta_i}$ for $i = 1, \ldots, d$. In every period, $(X_1, \ldots, X_d)$ are independently drawn from their respective distributions, independent of observations in all other periods. An agent observes $\Phi(X_1, \ldots, X_d)$ for a possibly non-injective function Φ. Let X_{in} denote the particular realization of X_i in period n, and $y_n = \Phi(X_{1n}, \ldots, X_{dn})$ be the period n observation. Let f_θ denote the density of $y = \Phi(X_1, \ldots, X_d)$ denote the probability density of y when the state of the world is θ.

We suppose that in every period the agent updates only his marginal probabilities and retains only those marginals to calculate the joint distributions in the next period. In place of standard probability notation, we will use W_n to denote the agent's beliefs over marginals as a measure and induced measure on the probability space Θ. Specifically, let $W_n^i(t_i)$ denote the agent's belief that

$\theta_i = t_i$, W_n^i the resulting measure over Θ_i, and

$$\begin{aligned}
W_n\left(\theta=(t_1,\ldots,t_d)\right) =& \Psi\left(W_n^1,\ldots,W_n^d\right)(t_1,\ldots,t_d) \\
W_{n+1}\left(\theta_i=t_i\right) \equiv& W_{n+1}\left(\theta_i=t_i \middle| y_{n+1}\right) \\
=& \frac{W_n\left(\theta_i=t_i\right) w_n\left(y_{n+1}\middle|\theta_i=t_i\right)}{w_n\left(y_{n+1}\right)} \\
=& \frac{W_n\left(\theta_i=t_i\right) w_n\left(y_{n+1}\middle|\theta_i=t_i\right)}{W_n\left(\theta_i=t_i\right) w_n\left(y_{n+1}\middle|\theta_i=t_i\right)+W_n\left(\theta_i\neq t_i\right) w_n\left(y_{n+1}\middle|\theta_i\neq t_i\right)},
\end{aligned}$$

where for every $A \subset \Theta_i$, $W_n\left(\theta_i \in A\right) = \sum_{t_i \in A} W_n\left(\theta_i = t_i\right)$, since W_n is a measure. The notation is used in lower-case with respect to y to evoke that y is continuously distributed with respect to a density, whereas θ is discretely valued with a probability mass function. For convenience, we will use $W_n^i(\cdot) = W_n(\theta_i = \cdot)$ interchangeably. This implies

$$\log \frac{W_{n+1}\left(\theta_i=t_i\right)}{1-W_{n+1}\left(\theta_i=t_i\right)} = \log \frac{w_n\left(y_{n+1}\middle|\theta_i=t_i\right)}{w_n\left(y_{n+1}\middle|\theta_i\neq t_i\right)} + \log \frac{W_n\left(\theta_i=t_i\right)}{1-W_n\left(\theta_i=t_i\right)}. \tag{1}$$

The agent views the events $\theta_i = t_i$ as independent with distribution $W_n\left(\theta_i = t_i\right)$, in which case $\Psi\left(W_n^1,\ldots,W_n^d\right)(t_1,\ldots,t_d) = \prod_{i=1}^d W_n^i(t_i)$. Two remaining calculations remain, which are analogous to the normal Bayesian version:

$$w_n\left(y_{n+1}|\theta_i=t_i\right) = \frac{\sum_{\vartheta\in\Theta:\vartheta_i=t_i} W_n\left(\theta=\vartheta\right) f_\vartheta\left(y_{n+1}\right)}{\sum_{\vartheta\in\Theta:\vartheta_i=t_i} W_n\left(\theta=\vartheta\right)}, \tag{2}$$

$$w_n\left(y_{n+1}|\theta_i\neq t_i\right) = \frac{\sum_{\vartheta\in\Theta:\vartheta_i\neq t_i} W_n\left(\theta=\vartheta\right) f_\vartheta\left(y_{n+1}\right)}{\sum_{\vartheta\in\Theta:\vartheta_i\neq t_i} W_n\left(\theta=\vartheta\right)}. \tag{3}$$

We impose $\sum_{\vartheta\in\Theta:\theta_i=t_i} W_n(\theta=\vartheta) = W_n(\theta_i=t_i)$ for every i. Then, when $\sum_{t_i\in\Theta_i} W_0(\theta_i=t_i)=1$, one can verify that $\sum_{t_i\in\Theta_i} W_n(\theta_i=t_i)=1$ for every n.

2.2 Approximate Optimality of the Independence Copula

The belief process of agents is boundedly rational. For Doob's well-known Bayesian consistency result to apply, agents must sequentially update their priors over the entire state space by calculating the joint posterior distribution after every observation. Our deviation from this benchmark can be rationalized by a sparsity-based model of bounded rationality following Gabaix (2014, 2016). Consider d state variables that can take on k values. Then regular Bayesian updating (Doob) requires storing k^d real numbers in the joint distribution. In practice agents might only maintain and update the marginal distributions over each dimension, thereby reducing the memory burden to kd real numbers. To update the marginal distributions from their observations and knowledge of the data generating process, the agents need to attempt to reconstruct the joint distribution of state variables. This saves the agent from remembering only kd real numbers but performing k^d calculations every time marginal beliefs are updated. In this section we assume that agents apply the independence copula to construct the joint distribution, but in Proposition 2 this assumption is relaxed.[6]

Suppose that the agent is concerned with estimating a signal $y_n \in \mathbb{R}^d$ as accurately as possible in the sense that he wishes to minimize a discounted square loss function

$$\min_{\{\hat{y}_n\}_{n=1}^{\infty}} \mathbb{E}\left[\sum_{n=1}^{\infty} \delta^{n-1} |y_n - \hat{y}_n|^2\right].$$

The major behavioral assumption of this paper is that agents retain only their marginal beliefs $W_n^i, i = 1, \ldots, d$ from period n. In the next period, in order to obtain W_{n+1}^i, the agents must reconstruct the joint distribution $W_n = W_n(\theta = (t_1, \ldots, t_d))$ using some choice of copula. It is straightforward to see that the independence copula, given by

$$W_n(\theta = (t_1, \ldots, t_d)) = \prod_{i=1}^{d} W_n^i(t_i), \tag{4}$$

is approximately optimal in the following sense, which resembles Proposition 1 of Fryer et al. (2015):

Lemma 1. *Suppose that $y_n \in K \subset \mathbb{R}^d$ a.s., where K is a compact set. Let $\hat{y}_n^* = \mathbb{E}[y|W_n^*]$, where W_n^* is generated by application of the independence copula* (4) *in every period. Then for every $\varepsilon > 0$,*

[6]This means that agents in our model get the *correlation of beliefs* wrong. While the underlying states (bandit arms) are independent, beliefs about the states should be correlated. See also Ellis and Piccione (2017), which provides a model that allows the misperception of correlated risks and illustrates implications for decision making.

there is a $\overline{\delta} \in (0,1)$ such that for $\delta \leq \overline{\delta}$,

$$\mathbb{E}\left[\sum_{n=1}^{\infty} \delta^{n-1} |y_n - \hat{y}_n^*|^2\right] \leq \inf_{\{\hat{y}_n\}_{n=1}^{\infty} \subset K} \mathbb{E}\left[\sum_{n=1}^{\infty} \delta^{n-1} |y_n - \hat{y}_n|^2\right] + \varepsilon.$$

Hence, for a sufficiently impatient agent, the strategy of retaining only marginals and reconstructing joint distributions via multiplication is arbitrarily close to optimal. The strategy also allows the agent to retain less information on the distribution of the θ's from period to period.

2.3 Results on Consistency

Consistent Learning

It is sometimes the case that agents will almost surely learn the truth asymptotically when they use the independence copula to reconstruct their joint beliefs. The most general setting that we will consider is when period n signals y take values in some subset of $\mathbb{R}$ on which densities $f_\theta(y)$ are nonzero for every θ. For learning to be asymptotically consistent requires monotonicity conditions such that observations y are sufficiently informative of the true state of the world, which occurs when "easier to identify" observations are sufficiently likely. Intuitively, consistency requires that the true state of the world θ^* be easy to learn in the sense that higher (or lower) observations of y provide a stronger signal that the true state is θ^*, and θ^* generates a relatively high frequency of high (or low) observations of y. In this setting, a state of the world with many "easier to identify" observations corresponds to the following 3 conditions holding:

A1: There is some state $\theta^* = (\theta_1^*, \ldots, \theta_d^*)$ such that the likelihood ratio:

$$\frac{f_{\theta^*}(y)}{f_\theta(y)}$$

is non-decreasing for all $\theta \in \Theta$.

A2: For any $\theta = (\theta_1, \ldots, \theta_d)$, let $T_i(\theta) = (\theta_1, \ldots, \theta_i^*, \ldots, \theta_d)$ be θ with the i^{th} coordinate replaced by θ_i^*. Then $F_{T_i(\theta)}(y)$ first-order stochastically dominates $F_\theta(y)$ for all θ and i.

The first condition is a stronger version of first order stochastic dominance, which is already implied by the second for θ^*. The second assumption is that switching a state θ_i to θ_i^* tends to shift the resulting distribution of signals y to the right. Note that these assumptions do not preclude the case that θ has no bearing on the distribution of y, i.e. $f_\theta(y) = f(y)$ for all θ. To strengthen our result, we may also make another assumption on the existence of strong signals that perturb agents' beliefs with any prior.

- **A3**: With respect to the push-forward F_{θ^*} measure, there is a set $A \subset \mathbb{R}$ and $\varepsilon > 0$ such that for all $y \in A$, $i \in \{1, \ldots, d\}$, and $\theta \in \Theta$,

$$f_{T_i(\theta)}(y) \geq f_\theta(y) \quad \text{and } f_{\theta^*}(y) \geq f_\vartheta(y) + \varepsilon$$

whenever $\vartheta \in T_i^{-1}(\theta^*)$. In addition, $\sup_{\theta \in \Theta} \left\| f_\theta(y) \mathbf{1}_{y \in A} \right\|_\infty < \infty$.

Given these conditions, agents will learn the true state θ^* asymptotically. States satisfying these conditions are sufficiently observationally different from other states that updating of marginal probabilities is nonetheless sufficient to distinguish these states from other states. As will be evident in the proof, assumption A3 is only necessary insofar as it provides perturbations to marginal beliefs in the interior of the d dimensional unit cube: in this respect, it may be generalized a number of ways. Lemma 6 is an example of such a perturbation argument.

Proposition 1 (Consistency)**.** *Suppose that $f_\theta(y)$ is continuously differentiable for every θ, agents update marginals W_n according to the independence cupola, and the true state is θ^*. If easy to identify observations are frequent in the sense that A1–A3 hold, then for every i, $W_n(\theta_i = \theta_i^*) \overset{a.s.}{\to} 1$.*

Inconsistent Learning

Having established when learning is consistent, we now turn to asymptotic inconsistency. We are interested in situations in which a state $t \in \Theta$ may be an attracting equilibrium, in that as prior beliefs $W_0(\theta_i = t_i')$ approach δ_{t_i} for every i, $\Pr(W_n(\theta_i = t_i) \to 1)$ approaches 1. Attractiveness might hold even if t is not the true state of the world. All that is necessary is for a state to be *locally* the best explanation for the world (the explicit condition is (7)). In other words, a state $(\theta_1, \ldots, \theta_d)$

is attractive if perturbing one of the θ_i's into θ_i' while keeping the rest of the θ_{-i} static results in a very poor explanation of the world. This evokes the notion of "local thinking" (Gennaioli and Shleifer, 2010), although the microfoundations and updating dynamics are quite different.

To illustrate the preceding setting, suppose that $d = 2$ for binary states with θ_1="the president is great or terrible," and θ_2= "the media is truthful or lies." Consider an agent with initial prior heavily weighing θ_1=the president is great and θ_2= the media lies. Hearing frequent negative news about the president, it turns out that an agent in our model essentially weighs the prior belief about the state of the world against one-dimensional perturbations (president good, media truthful) and (president bad, media lies), both of which are inconsistent with the narrative of "bad Mr. President." The agent, because of his priors, would not consider the state "president bad and media truthful," which is nonetheless consistent with the world. As a result, the agent is drawn continually deeper into a possibly false perception of the world.

We measure the distance between W_0^i and δ_{t_i} by the scaled total variation distance $\left\|W_0^i - \delta_{t_i}\right\| = 1 - W_0\left(\theta_i = t_i\right)$. We also make the following assumption on Ψ which is that the agent regards the θ_i as roughly independent events conditional upon observed signals:

$$0 < \inf_{\substack{x\in(0,1]^d \\ V^i\in\Delta(\Theta_i), \\ 1\leq i\leq d}} \frac{\Psi\big(V^1,\ldots,V^d\big)(t_1,\ldots,t_d)}{\prod_{i=1}^d V^i(t_i)} \leq \sup_{\substack{x\in(0,1]^d \\ V^i\in\Delta(\Theta_i), \\ 1\leq i\leq d}} \frac{\Psi\big(V^1,\ldots,V^d\big)(t_1,\ldots,t_d)}{\prod_{i=1}^d V^i(t_i)} < \infty. \tag{5}$$

We will require a boundedness condition for densities

$$\sup_{y,\theta,\theta'} \frac{f_\theta(y)}{f_{\theta'}(y)} < \infty, \tag{6}$$

which requires that no realization of y rule out any state θ.

In this context, learning is not asymptotically consistent (divergence occurs) when some state t other than the truth can be locally attracting in the sense that beliefs may converge to t. This can occur when easy to identify observations are not sufficiently likely (the conditions of the Proposition 1 do not hold). To give sufficient conditions for a state t to be attracting, we confine our attention to neighborhoods of t that may be obtained by changing one of its elements. Namely, for every $i = 1, \dots, d$, we set

$$G_i(t) = \{\theta \in \Theta : \theta_{-i} = t_{-i}, \theta_i \neq \theta_{-i}\},$$

i.e. $\theta_j = t_j$ for every $j \neq i$.

Proposition 2 gives a sufficient condition for t to be locally attracting. Given probability densities f, g on a set X with dominating measure μ, we let the *Kullback-Leibler* divergence of f and g be given by

$$D_{\mathrm{KL}}(f \,||\, g) = \int_X f(x) \log \frac{f(x)}{g(x)} d\mu.$$

The basic premise of the proposition is that a state t only needs to be a better explanation for the real world than small perturbations of t in order to be locally attracting. Specifically, suppose that θ_0 is the true state of the world and the divergence of f_t from f_{θ_0} is smaller than the divergence of f from f_{θ_0} whenever f is in the convex hull of $f_{t'}$ for t' which agree with t at all but one coordinate,

$$\text{for all } f \in \mathrm{co}\left(\{f_\theta : \theta \in G_i(t)\}\right), D_{\mathrm{KL}}\left(f_{\theta_0} \,||\, f\right) > D_{\mathrm{KL}}\left(f_{\theta_0} \,||\, f_t\right). \tag{7}$$

Then t has the local attracting property in the sense that an agent who updates along marginal distributions can be made to have arbitrarily high probability of asymptotically believing t is true by shifting the agent's initial beliefs very close to t. Notice (7) can only hold because observations are not clearly identifiable. Intuitively, this condition states that information is less informative of the actual true state of the world and as a result individuals are more readily swayed into incorrect beliefs.[7]

Proposition 2. *Let θ_0 be the true state of the world and assume* (5) *and* (6)*. Consider $t \in \Theta$ such that for all i* (7) *holds. Then for every $\varepsilon > 0$, there exists a $c > 0$ such that if $W_0(\theta_i = t_i) \geq 1 - c$ for all i, then $Pr(W_n(\theta_i = t_i) \to 1) \geq 1 - \varepsilon$.*

[7] This condition can be understood intuitively for the 2x2 case presented below. Symmetry (i.e. $p_H \approx q_H$, $p_L \approx q_L$) tends to predict inconsistent beliefs. In the phrasing of (7), the true state of the world θ_0 is $(\theta = 0, \sigma = 1)$, the false state of the world is $t = (\theta = 1, \sigma = 0)$, and taking $(p_H, p_L) \to (q_H, q_L)$ tends to push $D_{\mathrm{KL}}\left(f_{\theta_0} \,||\, f_t\right)$ towards its lowest possible value of 0.

The following corollary then is a straightforward consequence of Gibb's inequality:

Corollary 1. *If* (5) *and* (6) *hold and* $f_{\theta_0} \notin co(\{f_\theta : \theta \in G_i(\theta_0)\})$ *for all i, then the belief* $\theta = \theta_0$ *is locally attracting in the sense of Proposition 2.*

The condition of this proposition says that if distribution p is "closer" to the true distribution p_0 of the signal y then all small perturbations of p, along with the other conditions, then p can be attracting. When agents see a distribution of y they may only look for potential explanations of the distribution that are small perturbations of the present way of thinking (dismissing larger changes), and if none of them explain y better, then they will converge to a potentially incorrect belief.

2.4 A 2x2 Model

We have shown that learning is asymptotically consistent when observations are highly informative of the underlying state, but inconsistent otherwise. In light of these general results, we now consider the simple case with two arms that can take two values (good or bad) to provide greater intuition for how the previous theoretical results connect to real examples. Given the simpler structure, we adopt slightly different notation to ease exposition.

2.4.1 Setup

There are two underlying state variables, θ and σ, which determine the frequency of outcomes t and s respectively, and the state variables take binary values (think "good" or "bad"). In particular, t-successes occur with higher probability when $\theta = 1$ (θ is "good"), and s-successes occur with higher probability when $\sigma = 1$. Thus,

$$\Pr(t = 1|\theta = 1) = p_H > p_L = \Pr(t = 1|\theta = 0), \text{ and} \tag{8}$$

$$\Pr(s = 1|\sigma = 1) = q_H > q_L = \Pr(s = 1|\sigma = 0), \tag{9}$$

so that $t = 1$ is more likely when $\theta = 1$ and $s = 1$ is more likely when $\sigma = 1$. Let all probabilities p, q lie strictly between zero and one. The realization function is $y(t,s) = t + s$, as before. Thus, the states variables θ and σ determine the frequency of observations $y \in \{0,1,2\}$. The mapping from θ and σ to t and s is common knowledge: agents only disagree about the likelihood of θ and σ, but not about how those states translate into realizations y.[8]

2.4.2 Belief Updating

Let an agent hold beliefs $P_n = Pr(\theta = 1)$ and $Q_n = Pr(\sigma = 1)$ (the probability the states are good), where n denotes the number of signals observed (also the period). Agents observe signal y_{n+1} in period $n+1$ and use Bayes' Rule to update their marginal beliefs about the two states variables to P_{n+1} and Q_{n+1}. Thus, agents update beliefs sequentially, using only the current prior and the current information, together with Bayes' Rule, to form their posteriors.

We derive the evolution of beliefs using odds ratios. Define $\overline{p} = Pp_H + (1-P)p_L$ and $\overline{q} = Qq_H + (1-Q)q_L$ to be the ex-ante expected realizations of t and s given beliefs P, Q (notice we will suppress the sub-n subscripts when all variables share the same n value). Define $O^P = \frac{P}{1-P}$, and $O^Q = \frac{Q}{1-Q}$, and let $O^P(y), O^Q(y)$ denote the updated odds ratio after observing y. Then Bayes' Theorem applied to marginal probabilities yields

$$O^P_{n+1}(2) = \frac{p_H}{p_L} O^P_n, \quad O^Q_{n+1}(2) = \frac{q_H}{q_L} O^Q_n, \tag{10}$$

$$O^P_{n+1}(0) = \frac{1-p_H}{1-p_L} O^P_n, \quad O^Q_{n+1}(0) = \frac{1-q_H}{1-q_L} O^Q_n, \tag{11}$$

$$O^P_{n+1}(1) = \frac{(p_H(1-\overline{q}_n) + (1-p_H)\overline{q}_n)}{(p_L(1-\overline{q}_n) + (1-p_L)\overline{q}_n)} O^P_n, \quad O^Q_{n+1}(1) = \frac{(q_H(1-\overline{p}_n) + (1-q_H)\overline{p}_n)}{(q_L(1-\overline{p}_n) + (1-q_L)\overline{p}_n)} O^Q_n, \tag{12}$$

where we have written the posterior odds ratio as a product of the prior odds ratio and the likelihood ratio, which depends on y, and, importantly for $y = 1$, on P, Q. It is immediate that when the realization gives identified information about t and s, both agents update posteriors in the same way, as standard models suggest, though posteriors will still differ because priors differ. However, unidentified realizations ($y_{n+1} = 1$) can lead to divergent posteriors when beliefs are sufficiently divergent.

[8] Contrast this assumption with the information structure in Acemoglu et al. (2016), where a one-dimensional state θ produces signals with probability p_θ, and agents potentially disagree about p_θ as well as the probability of θ.

2.4.3 Results in the 2×2 Example

Individuals converge asymptotically to the correct beliefs when $\theta = \sigma = 1$ or $\theta = \sigma = 0$. This is because identified observations ($y = 2$ or $y = 0$) are relatively most likely in these cases. Thus, for the most frequent observations, the likelihood ratios do not depend on both beliefs P and Q, and so agents tend to interpret observations the same most of the time.

Proposition 3. *If $\theta = \sigma = 1$ or $\theta = \sigma = 0$, then (respectively) $P_n, Q_n \overset{a.s.}{\to} 1$ and $P_n, Q_n \overset{a.s.}{\to} 0$.*

In contrast, when intermediate observations ($y = 1$) are relatively likely, as occurs when $\theta \neq \sigma$, beliefs needs not converge to the truth. It is instructive to first consider when the states θ and σ are symmetric. In this case, beliefs converge to certainty confirming initial relative priors: extreme polarization occurs asymptotically.

Lemma 2. *Let $p_H = q_H$ and $p_L = q_L$. Then almost surely the ratio O^P/O^Q diverges to infinity if $P > Q$ and converges to zero if $P < Q$.*

The result follows because likelihood ratios get updated only after observing a one, going up or down depending only on the initial prior. The convergence of posteriors to reinforce initial priors immediately follows from this lemma. An implication of Proposition 3 and Lemma 2 is that when $\theta = 1 = \sigma$, both P and Q converge to 1. However, agents starting with different priors will continue to hold different *relative* beliefs about P and Q asymptotically—and those relative beliefs will diverge—even as beliefs converge to the truth (similar to Acemoglu et al., 2016). The symmetric result holds when $\theta = 0 = \sigma$.

Proposition 4. *Let $p_H = q_H$ and $p_L = q_L$. Suppose the truth is $\theta = 0$ and $\sigma = 1$. If $P_0 > Q_0$, then $P_n \overset{a.s.}{\to} 1$ and $Q_n \overset{a.s.}{\to} 0$ (vice versa if $P_0 < Q_0$).*

Corollary 2 (Asymptotic belief divergence)**.** *Let $p_H = q_H$ and $p_L = q_L$. Suppose one agent holds prior beliefs with $P_0 > Q_0$ and the other has priors with $P_0 < Q_0$. Then if $\theta \neq \sigma$, with probability 1 agents' beliefs will asymptotically diverge to complete polarization, the first with $P_n \overset{a.s.}{\to} 1$ and $Q_n \overset{a.s.}{\to} 0$, and the reverse for the other.*

The symmetric case is perhaps least interesting because the states $\theta = 1, \sigma = 0$ and $\sigma = 1, \theta = 0$ produce observationally equivalent outcomes. However, if agents could take actions whose payoffs depended on the values of θ and σ, disagreement would still matter. Furthermore, this result is helpful to consider how divergence can occur even when states are asymmetric. In particular, divergence may occur when the states $\theta = 1, \sigma = 0$ and $\sigma = 1, \theta = 0$ produce outcomes that are observationally similar, even if not identical.

Individuals can converge asymptotically to false beliefs even when states are asymmetric. For many parameters, there exist priors such that almost surely beliefs converge to the wrong values. Crucially, there is a robust set of parameters such that beliefs diverge. Throughout suppose $\theta = 0$ and $\sigma = 1$. (By symmetry all results also hold for $\theta = 1$ and $\sigma = 0$.)

First, we can characterize conditions on beliefs P, Q such that beliefs converge incorrectly almost surely. Ignoring a (measurably) small set of parameter values, we obtain a classification result for the asymptotic behavior of (P_n, Q_n). Lemma 3 states conditions on parameters that guarantee when priors exists such that beliefs will asymptotically converge to the endpoints (zero or one) with positive probability. Specifically, the proposition states that the asymptotic behavior of (P_n, Q_n) behavior can be stated in terms of the expectation of the transition function (i.e., the log odds ratio, defined in the Appendix) in neighborhoods of the extremal points of $[p_L, p_H] \times [q_L, q_H]$, because asymptotically individuals will accumulate in these neighborhoods with positive probability.

Lemma 3. *The following hold:*

1. *If* $\mathbb{E}\left[\Delta \log O^P_{n+1}(q_L)\right] > 0$ *and* $\mathbb{E}\left[\Delta \log O^Q_{n+1}(p_H)\right] < 0$*, then there exist* (P_0, Q_0) *such that* $P_n \to 1$ *and* $Q_n \to 0$ *with positive probability tending to* 1 *as* $Q_0 \to 0$ *and* $P_0 \to 1$

2. *If* $\mathbb{E}\left[\Delta \log O^P_{n+1}(q_L)\right] < 0$*, then* $P \overset{a.s.}{\to} 0$*; if additionally* $\mathbb{E}\left[\Delta \log O^Q_{n+1}(p_L)\right] < 0$ *then* $Q \overset{a.s.}{\to} 0$*, whereas if* $\mathbb{E}\left[\Delta \log O^Q_{n+1}(p_L)\right] > 0$ *then* $Q \overset{a.s.}{\to} 1$.

3. *If* $\mathbb{E}\left[\Delta \log O^Q_{n+1}(p_H)\right] > 0$*, then* $Q \overset{a.s.}{\to} 1$*; if additionally* $\mathbb{E}\left[\Delta \log O^P_{n+1}(q_L)\right] < 0$ *then* $P \overset{a.s.}{\to} 0$*, whereas if* $\mathbb{E}\left[\Delta \log O^P_{n+1}(q_L)\right] > 0$ *then* $P \overset{a.s.}{\to} 1$.

Most importantly, if parameters satisfy condition 1 of Lemma 3 then there exist priors such that beliefs will diverge with positive probability. Numerically evaluating over all combinations

$(p_H, p_L, q_H, q_L) \in (0,1)^4$, together with ordering restrictions, approximately 27.26% of parameters satisfy condition 1, meaning that for at least this many parameters there is a positive probability of converging to the wrong values for some priors.

Next, given parameters p_H, p_L, q_H, q_L, let $f(p,q) = \Pr((P_n, Q_n) \to (1,0) | P_0 = p, Q_0 = q)$ be the probability of that beliefs converge to $(1,0)$, which are the wrong beliefs. We have the following important results, including a finding on continuity in Lemma 4 that extends to a broader class of random dynamical systems with a Bernoulli Shift as a random component (see Proposition 1.6 in the Online Appendix).

Lemma 4. *If* $\frac{\log \frac{q_H}{q_L}}{\log \frac{p_H}{p_L}} \neq \frac{\log \frac{1-q_H}{1-q_L}}{\log \frac{1-p_H}{1-p_L}}$*, then the convergence probability* $f(p,q)$ *is continuous in priors* p,q.

The following Corollary is a related result which follows from Lemma 10, which is stated and proved in the appendix.

Corollary 3. *If*

$$\frac{\log \frac{q_H}{q_L}}{\log \frac{p_H}{p_L}} < \frac{\log \frac{1-q_H}{1-q_L}}{\log \frac{1-p_H}{1-p_L}}, \tag{13}$$

then either $f(p,q) > 0$ *for all* $(p,q) \in (0,1) \times (0,1)$*, or it vanishes for all* (p,q).

Thus, if parameters satisfy condition 1 of Lemma 3 as well as equation (13), then for any (rightly ordered) priors, beliefs will diverge with positive probability. The results of this section—namely, guaranteed divergence for symmetry together with continuity of the divergence probability—suggest that divergence is more likely to occur when the states $(1,0)$ and $(0,1)$ are more observationally similar (closer to symmetry). The intuition is that the more observationally similar are the states, the less the data can distinguish between competing beliefs when agents update only marginals. With exact observational equivalence, learning is guaranteed to converge to reinforce initial priors (this is also true with perfect Bayesian learning).[9]

[9]While we do not provide the results, we have investigated this issue using simulations. Indeed, we find that divergence is most likely to occur close to symmetry and decreases continuously. Compare these results to Fryer

By evaluating the set of parameters which meet the bounds presented in Corollary 3 as well as those in Lemma 3, we find the following:

Proposition 5. *There is a positive measure set of parameters such that for any priors* $0 \leq Q_0 < P_0 \leq 1$ *beliefs diverge with positive probability.*

Approximately 8.33% of all parameters satisfy both sets of conditions. The sets of parameters for which (i) any priors will converge to (1,0) with positive probability, or for which (ii) some priors are guaranteed to converge to (1,0) with positive probability, indicate that divergence is most likely to occur in a neighborhood of parameters around symmetry.

3 Discussion and Implications

Our analysis has several implications for the effectiveness of attempts to relieve identification problems, divergence and information disclosure, and model specification and dynamics. We first discuss the setup of our model and the interpretation of the results.

3.1 Discussion

The reader may wonder about the simplicity of the model and signals received and whether our results are robust to more general setups. Our results are completely driven by the following three assumptions: (i) a partially identified model, with (ii) relatively frequent observations that do not completely identify the model, and (iii) rational, sequential updating of marginal beliefs but not the

et al. (2015). In their model, there are two symmetric states (a and b) and agents' beliefs converge to one of these states to confirm their prior when ambiguous signals are sufficiently frequent. In their model, the probability of polarization is an increasing function of the probability of ambiguous signals. Our simulation results suggest that in our model the probability of divergence is a function of how different are parameters for each state. The more parameters differ, the easier it is to "statistically identify" observations of 1, which are otherwise unidentifiable. In their model ambiguous signals are completely unidentified. Thus, in our model, fundamentals determine the severity of the identification problem and the probability of divergence, whereas in their model the severity of ambiguity determines the probability of divergence. Furthermore, when states are not symmetric, posteriors can converge to values that are not even "symmetric" with the truth (i.e., to $(1,1)$ or $(0,0)$. This result is important because agents would not only disagree about the underlying values of θ, σ, but they would have quite different predictions for the distribution of y_n. One could argue that polarization truly refers to agents disagreeing about the value of $\theta + \sigma$, which is what would occur in this case.

full joint distribution of beliefs. Our model is silent about where prior beliefs come from; perhaps there are behavioral or generational explanations for priors (see Bénabou and Tirole (2016) for an overview of belief production).

We have deliberately chosen the simplest model to illustrate that partial identification can lead to beliefs diverging from the truth—in fact, if anything the simplicity of our model and the set of signals makes learning the truth more likely. There are many reasons to believe that the identification problem in a higher dimensional model with a richer signal-space would be more severe, making divergent learning even more likely, since fully-identifying observations would be even less common. Also, our agents can only possibly disagree about initial priors; divergence would be even more likely if agents also disagreed about the model parameters.

Furthermore, we suppose that agents desire to know the underlying state variables for reasons beyond (just) being able to predict future observations y_n. For example, agents may want to take actions whose payoffs depend on the values of θ and σ separately. Or there may be an additional variable whose realization will be made known in the far future, whose value could simply equal θ or σ. Thus, even if knowing $\theta+\sigma$ is sufficient to predict the distribution of y_n (as is exactly true when the states are symmetric), knowledge of θ and σ separately would still matter.

Finally, we have argued that "unidentified" observations lead to greater disagreement because inference depends on initial beliefs. How might this mechanism look in reality? Consider two examples:

(i) Two economists, a Keynesian and a Neoclassical, walk into a bar. There was a recent stimulus package, and the new GDP results are sluggish. The Neoclassical says, "Goes to show that stimulus doesn't work." The Keynesian replies, "Oh no, goes to show that the economy is much worse than we thought" (perhaps later adding that the stimulus was poorly designed). Meanwhile, in another bar in another country far away, a stimulus bill has passed with apparently great results. "I guess the economy was already out of the recession," says the Neoclassical.

(ii) A liberal and a conservative are watching the news together. After a string of media coverage providing evidence that the country has become more liberal, the liberal says, "The country is coming my way." The conservative, responds, "No, more evidence of the media's liberal bias."

In each of these cases, the exact same datum is interpreted in completely different ways, even confirming the prior that each person had before. But the observers are not rejecting the data, nor are the data ambiguous and unclear: they are simply using the evidence in different, rational ways. There are a number of underlying factors that contribute to what we see, but in these cases the data we see are of lower dimensionality than the world. The observers in these examples face identification problems so that the particular observations do not identify the underlying parameters. For the economists at the bar, GDP is a function of the state of the economy, the effectiveness of stimulus, and potentially how well designed that stimulus was. For the liberal and the conservative, the politics of news reflect the underlying beliefs of the population and how accurately those are reported. Each observer uses the data—they do not ignore the observation or assume it is just "noise"—but they rationally use the data to make inference about different underlying variables. The Keynesian infers that weak GDP data means something about the economy, not the effectiveness of stimulus, while the Neoclassical infers weak GDP means the fiscal multiplier is low.

3.2 Persuasion and Information

Our results have implications for persuasion and information release. First and most immediately, our model predicts that in some cases divergence may be unavoidable when initial beliefs are relatively different. Initial heterogeneity can lead to extreme polarization (in the sense of divergent beliefs), and more information worsens this outcome. This may explain why belief polarization in many places has seemingly increased in recent decades as information becomes more plentiful. If it is possible that priors can get "reset" (perhaps with a new generation), beliefs may be determined and unchanging until such a generational change occurs.[10]

[10] Relatedly, our model can produce a sense of "fact-free learning" (Aragones et al., 2005) if agents simply consider trying on a different initial prior and then analyzing the existing data. Moving a prior from (for example) $P_0 > Q_0$ to one with $Q_0 > P_0$, the exact same observations can be used to draw a completely different conclusion. This exercise could lead to an "Of course!" moment when states are asymmetric and agents recognize that (for example) $(1,0)$ provides a better model of the distribution of y's than $(0,1)$ does, even if Bayesian updating of marginals leads to one over the other depending on the initial prior.

Second, our results demonstrate that initial priors are critical for how people infer data. Our results suggest that the stakes from being able to control or influence initial priors (if that can be done) can be very high. Accordingly, our results suggest that agents attempting to persuade others (see Kamenica and Gentzkow, 2011) may choose to purposely disclose "unidentified information" when priors are in their favor, knowing that people will infer the unidentified information in ways that are favorable to their outcome.[11] Additionally, in light of information undermining agents' initial beliefs, agents may have an incentive to "modify" an existing model by adding additional dimensions. This creates an identification problem, so that an initial hypothesis or theory can accommodate information that was initially at odds. For example, a politician may argue that a news story suggesting unethical behavior is a reflection of media bias rather than indicating unethical dealings.

3.3 Model Specification and Dynamics

Finally, our results suggest that model specification has important implications for equilibrium dynamics, whether an economy has a representative agent or heterogeneous agents. When learning processes are unidentified, then a representative agent need not learn the truth, and beliefs among heterogeneous agents may diverge rather than converging or being stable. Identification problems and learning can have important implications for macroeconomic dynamics (Collin-Dufresne et al., 2016; Milani, 2007), asset pricing (Adam et al., 2016), stability of non-rational expectations (Woodford, 2013) and strategies (Kalai and Lehrer, 1993), convergence of dynamical systems and dynamic models (Fernández-Villaverde et al., 2006; Schenk-Hoppé and Schmalfuß, 2001), and for properties of Bayesian estimation of macro models (Schorfheide, 2011). Moreover, an implication of inconsistency is that if economists produce a handful of well-designed natural experiments or instrumental variables to relieve an identification problem, the profession may nonetheless continue to hold divergent beliefs so long as there are plenty of cases in which the model leaves room for multiple interpretations. Opinions predicated on information which is unidentified, in the sense

[11] To the extent that priors differ because of initial private information, polarization can be prevented only when valuable communication can occur. A large literature since Crawford and Sobel (1982) have studied strategic information transmission.

of Proposition 4, are susceptible to prior dependence[12]

How can divergence be avoided? Our results show that, when agents face an identification problem and when the underlying fundamentals of the world are "polarized," agents' beliefs can diverge to certain but differing beliefs about fundamentals. However, a critical driver of this result is that agents update their marginal beliefs whenever they get new information. Marginal updating is less obfuscated if information is more informative. Therefore, "patient" agents who update their marginal priors only when $m \geq 2$ signals are observed in a row and together interpreted as one composite-signal are less susceptible to asymptotic inconsistency. In general, shaping signals to be less difficult to identify should favor asymptotic convergence to the truth. Second, divergence would generally not occur if agents stored the full joint distribution of beliefs over the full state space rather than just marginal beliefs. In our model, agents update their beliefs frequently and rationally (though boundedly so), and they rationally incorporate new information in light albeit with faulty conditional joint prior. Further research should determine how agents form beliefs in light of new information and under what conditions those learning processes can be manipulated or modified.

4 Conclusion

The world is multi-dimensional—there are a number of factors that contribute to what we see—but the data we see are often lower dimensional than the world. Therefore we live with identification problems. It is also likely that we retain and update only marginal beliefs on the thousands of variables which shape the signals we receive, rather than maintain a huge and costly joint distribution. However, Bayesian updating of marginal probabilities need not converge to the truth. In particular, agents with differing priors may have posteriors diverge forever, with greater divergence the more information received. When unidentified observations are sufficiently likely, beliefs are likely to converge to certainty confirming relative initial beliefs.

We have characterized the limiting properties of beliefs for a simple model in which some observations are not identified. Our main result is that when unidentifiable observations are relatively

[12]We acknowledge Bruce Sacerdote for this observation.

more frequent, then asymptotically initial beliefs are likely to become reinforced. Thus, beliefs diverge in light of common information. However, if clearly identified observations are relatively likely, beliefs converge to the truth with probability one, and divergence will not occur.

References

ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2016): "Fragility of asymptotic agreement under Bayesian learning," *Theoretical Economics*, 11, 187–225.

ADAM, K., A. MARCET, AND J. P. NICOLINI (2016): "Stock market volatility and learning," *The Journal of Finance*, 71, 33–82.

ANDRADE, P., R. K. CRUMP, S. EUSEPI, AND E. MOENCH (2016): "Fundamental disagreement," *Journal of Monetary Economics*, 83, 106–128.

ANDREONI, J. AND T. MYLOVANOV (2012): "Diverging opinions," *American Economic Journal: Microeconomics*, 209–232.

ARAGONES, E., I. GILBOA, A. POSTLEWAITE, AND D. SCHMEIDLER (2005): "Fact-Free Learning," *American Economic Review*, 95, 1355–1368.

AUMANN, R. J. (1976): "Agreeing to disagree," *The annals of statistics*, 1236–1239.

BALIGA, S., E. HANANY, AND P. KLIBANOFF (2013): "Polarization and ambiguity," *The American Economic Review*, 103, 3071–3083.

BATSON, C. D. (1975): "Rational processing or rationalization? The effect of disconfirming information on a stated religious belief." *Journal of Personality and Social Psychology*, 32, 176.

BÉNABOU, R. AND J. TIROLE (2016): "Mindful economics: The production, consumption, and value of beliefs," *The Journal of Economic Perspectives*, 30, 141–164.

BENOÎT, J.-P. AND J. DUBRA (2014): "A Theory of Rational Attitude Polarization," *Available at SSRN 2529494*.

CHAIKEN, S. ET AL. (1992): "Defensive processing of personally relevant health messages," *Personality and Social Psychology Bulletin*, 18, 669–679.

COLLIN-DUFRESNE, P., M. JOHANNES, AND L. A. LOCHSTOER (2016): "Parameter learning in general equilibrium: The asset pricing implications," *The American Economic Review*, 106, 664–698.

CRAWFORD, V. P. AND J. SOBEL (1982): "Strategic information transmission," *Econometrica: Journal of the Econometric Society*, 1431–1451.

CRUMP, R., T. DAVIG, S. EUSEPI, AND E. MOENCH (2014): "Connecting "the Dots": Disagreement in the Federal Open Market Committee," Liberty Street Economics blog.

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

DIACONIS, P. AND D. FREEDMAN (1986): "On the Consistency of Bayes Estimates," *Ann. Statist.*, 14, 1–26.

DIXIT, A. K. AND J. W. WEIBULL (2007): "Political polarization," *Proceedings of the National Academy of Sciences*, 104, 7351–7356.

ELLIS, A. AND M. PICCIONE (2017): "Correlation misperception in choice," *The American Economic Review*, 107, 1264–1292.

EYSTER, E., A. GALEOTTI, N. KARTIK, AND M. RABIN (2014): "Congested observational learning," *Games and Economic Behavior*, 87, 519–538.

EYSTER, E. AND M. RABIN (2010): "Naive herding in rich-information settings," *American economic journal: microeconomics*, 2, 221–243.

FERNÁNDEZ-VILLAVERDE, J., J. F. RUBIO-RAMÍREZ, AND M. S. SANTOS (2006): "Convergence properties of the likelihood of computed dynamic models," *Econometrica*, 74, 93–119.

FRYER, R. G., P. HARMS, AND M. O. JACKSON (2015): "Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization," *Available at SSRN 2263504*.

GABAIX, X. (2014): "A sparsity-based model of bounded rationality," *The Quarterly Journal of Economics*, 129, 1661–1710.

——— (2016): "Behavioral Macroeconomics Via Sparse Dynamic Programming," Tech. rep., National Bureau of Economic Research.

GEANAKOPLOS, J. D. AND H. M. POLEMARCHAKIS (1982): "We can't disagree forever," *Journal of Economic Theory*, 28, 192–200.

GENNAIOLI, N. AND A. SHLEIFER (2010): "What comes to mind," *The Quarterly journal of economics*, 125, 1399–1433.

GERBER, A. AND D. GREEN (1999): "Misperceptions about perceptual bias," *Annual review of political science*, 2, 189–210.

HEIDHUES, P., B. KOSZEGI, AND P. STRACK (2015): "Unrealistic Expectations and Misguided Learning," *Available at SSRN*.

HIRSHLEIFER, D. AND S. H. TEOH (2003): "Limited attention, information disclosure, and financial reporting," *Journal of accounting and economics*, 36, 337–386.

HONG, H. AND J. C. STEIN (2007): "Disagreement and the stock market," *The Journal of Economic Perspectives*, 21, 109–128.

JERN, A., K.-M. K. CHANG, AND C. KEMP (2014): "Belief polarization is not always irrational." *Psychological review*, 121, 206.

KALAI, E. AND E. LEHRER (1993): "Rational learning leads to Nash equilibrium," *Econometrica: Journal of the Econometric Society*, 1019–1045.

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian persuasion," *The American Economic Review*, 101, 2590–2615.

KONDOR, P. (2012): "The more we know about the fundamental, the less we agree on the price," *The Review of Economic Studies*, 79, 1175–1207.

LORD, C. G., L. ROSS, AND M. R. LEPPER (1979): "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence." *Journal of personality and social psychology*, 37, 2098.

MALMENDIER, U., S. NAGEL, AND Z. YAN (2017): "The Making of Hawks and Doves: Inflation Experiences on the FOMC," Tech. rep., National Bureau of Economic Research.

MILANI, F. (2007): "Expectations, learning and macroeconomic persistence," *Journal of monetary Economics*, 54, 2065–2082.

MUNRO, G. D. AND P. H. DITTO (1997): "Biased assimilation, attitude polarization, and affect in reactions to stereotype-relevant scientific information," *Personality and Social Psychology Bulletin*, 23, 636–653.

ORTOLEVA, P. AND E. SNOWBERG (2015): "Overconfidence in political behavior," *The American Economic Review*, 105, 504–535.

PLOUS, S. (1991): "Biases in the assimilation of technological breakdowns: Do accidents make us safer?" *Journal of Applied Social Psychology*, 21, 1058–1082.

RABIN, M. AND J. L. SCHRAG (1999): "First Impressions Matter: A Model of Confirmatory Bias," *The Quarterly Journal of Economics*, 114, 37–82.

SCHENK-HOPPÉ, K. R. AND B. SCHMALFUSS (2001): "Random fixed points in a stochastic Solow growth model," *Journal of Mathematical Economics*, 36, 19–30.

SCHORFHEIDE, F. (2011): "Estimation and evaluation of DSGE models: progress and challenges," Tech. rep., National Bureau of Economic Research.

SCHWARTZSTEIN, J. (2014): "Selective Attention and Learning," *Journal of the European Economic Association*, 12, 1423–1452.

SETHI, R. AND M. YILDIZ (2016): "Communication with Unknown Perspectives," *Econometrica*, 84, 2029–2069.

SUNDARESAN, S. AND S. TURBAN (2014): "Inattentive valuation and belief polarization," Tech. rep., Working paper.

TABER, C. S. AND M. LODGE (2006): "Motivated skepticism in the evaluation of political beliefs," *American Journal of Political Science*, 50, 755–769.

WOODFORD, M. (2013): "Macroeconomic analysis without the rational expectations hypothesis," Tech. rep., National Bureau of Economic Research.

Appendices

A Proofs

Proof of Lemma 1. This follows from the observation that in period $n = 1$, $\hat{y}_n^* = \mathbb{E}[y]$, where the expectation is over the prior probability distribution $Q_1 \times \cdots \times Q_d$ on Θ and on the contingent distribution of y induced by θ. Hence,

$$\mathbb{E}\left[|y_n - \hat{y}_n^*|^2\right] = \mathbb{E}\left[|y - \mathbb{E}[y]|^2\right] \leq \mathbb{E}\left[|y_n - \hat{y}_n|^2\right]$$

for all r.v. $\hat{y}_n$. The desired result follows from noting that

$$\mathbb{E}\left[\sum_{n=1}^{\infty} \delta^{n-1} |y_n - \hat{y}_n^*|^2\right] \leq \mathbb{E}\left[|y_n - \hat{y}_n^*|^2\right] + \frac{\delta C}{1-\delta}$$

for some constant $C < \infty$ depending on the set K. □

Proof of Proposition 1. Recall that $y_n \overset{\text{iid}}{\sim} F_{\theta^*}$. When the independence copula is applied, the denominator of (2) is

$$\sum_{\vartheta\in\Theta:\vartheta_i=\theta_i^*} W_n(\theta=\vartheta) = \sum_{\vartheta\in\Theta:\vartheta_i=\theta_i^*} W_n(\theta_i=\theta_i^*)\prod_{j\neq i} W_n(\theta_j=\vartheta_j)$$
$$= W_n(\theta_i=\theta_i^*)\sum_{\vartheta_{-i}\in\Theta_{-i}}\prod_{j\neq i} W_n(\theta_j=\vartheta_j) = W_n(\theta_i^*).$$

Expanding the denominator of (3) in a similar fashion allows us to derive the following conditional expectation:

$$\mathbb{E}_n\left(W_{n+1}(\theta_i^*)\right) = W_n(\theta_i=\theta_i^*)\int_{\mathbb{R}} \frac{w_n(y|\theta_i=\theta_i^*)}{w_n(y)} f_{\theta^*}(y)\,dy$$
$$= W_n(\theta_i=\theta_i^*)\int_{\mathbb{R}} \frac{\sum_{\theta:\theta_i=\theta_i^*} f_\theta(y)\prod_{j\neq i} W_n(\theta_j)}{\sum_{\theta\in\Theta} f_\theta(y)\prod_{j=1}^d W_n(\theta_j)} f_{\theta^*}(y)\,dy.$$

Integration by parts implies

$$\int_{\mathbb{R}} \frac{\sum_{\theta:\theta_i=\theta_i^*} f_\theta(y)\prod_{j\neq i} W_n(\theta_j)}{\sum_{\theta\in\Theta} f_\theta(y)\prod_{j=1}^d W_n(\theta_j)} f_{\theta^*}(y)\,dy - 1$$
$$= \int_{\mathbb{R}} \frac{f_{\theta^*}(y)}{\sum_{\theta\in\Theta} f_\theta(y)\prod_{j=1}^d W_n(\theta_j)}\left(\sum_{\theta:\theta_i=\theta_i^*} f_\theta(y)\prod_{j\neq i} W_n(\theta_j) - \sum_{\theta\in\Theta} f_\theta(y)\prod_{j=1}^d W_n(\theta_j)\right) dy$$
$$= \int_{\mathbb{R}}\left[\frac{d}{dy}\left(\frac{f_{\theta^*}(y)}{\sum_{\theta\in\Theta} f_\theta(y)\prod_{j=1}^d W_n(\theta_j)}\right)\right.$$
$$\left.\int_y^\infty \left(\sum_{\theta:\theta_i=\theta_i^*} f_\theta(z)\prod_{j\neq i} W_n(\theta_j) - \sum_{\theta\in\Theta} f_\theta(z)\prod_{j=1}^d W_n(\theta_j)\right) dz\right] dy,$$

where in the third line we have used that the term in parenthesis on the preceeding line is a difference of two probability densities on $\mathbb{R}$. By A1

$$\frac{f'_{\theta^*}(y)}{f_{\theta^*}(y)} \geq \max_{\theta\in\Theta}\frac{f'_\theta(y)}{f_\theta(y)} \geq \frac{\sum_{\theta\in\Theta} f'_\theta(y)\prod_{j=1}^d W_n(\theta_j)}{\sum_{\theta\in\Theta} f_\theta(y)\prod_{j=1}^d W_n(\theta_j)},$$

so the quotient rule implies $\frac{d}{dy}\left(\frac{f_{\theta^*}(y)}{\sum_{\theta\in\Theta} f_\theta(y)\prod_{j=1}^d W_n(\theta_j)}\right) \geq 0$. On the other hand, one can write

$$
\begin{aligned}
&\sum_{\theta:\theta_i=\theta_i^*} f_\theta(z)\prod_{j\neq i} W_n(\theta_j) - \sum_{\theta\in\Theta} f_\theta(z)\prod_{j=1}^d W_n(\theta_j)\\
&\quad= \sum_{\theta:\theta_i=\theta_i^*} f_\theta(z)\prod_{j\neq i} W_n(\theta_j) - W_n(\theta_i^*)\sum_{\theta:\theta_i=\theta_i^*} f_\theta(z)\prod_{j\neq i} W_n(\theta_i) - \sum_{\theta:\theta_i\neq\theta_i^*} f_\theta(z)\prod_{j=1}^d W_n(\theta_j)\\
&\quad=(1-W_n(\theta_i^*))\sum_{\theta:\theta_i=\theta_i^*} f_\theta(z)\prod_{j\neq i} W_n(\theta_j) - \sum_{\theta_{-i}\in\Theta_{-i}}\prod_{j\neq i} W_n(\theta_j)\sum_{\theta_i\in\Theta_i\setminus\{\theta_i^*\}} f_\theta(z)W_n(\theta_i)\\
&\quad=(1-W_n(\theta_i^*))\sum_{\substack{\theta_{-i}\in\Theta_{-i}\\ \theta_i=\theta_i^*}} f_\theta(z)\prod_{j\neq i} W_n(\theta_j)\\
&\qquad-(1-W_n(\theta_i^*))\sum_{\theta_{-i}\in\Theta_{-i}}\prod_{j\neq i} W_n(\theta_j)\sum_{\theta_i\in\Theta_i\setminus\{\theta_i^*\}} f_\theta(z)\frac{W_n(\theta_i)}{(1-W_n(\theta_i^*))}.
\end{aligned}
$$

Fix a $\theta_{-i}\in\Theta_{-i}$ and let θ^s denote the element $\theta\in\Theta$ with $\theta^s_{-i}=\theta_{-i}$ and $\theta_i^s=\theta_i^*$. Let $S\subset\Theta$ denote the set of ϑ with $\vartheta_{-i}=\theta_{-i}$ and $\vartheta_i\neq\theta_i^*$. A2 implies that for any $|S|$-dimensional probability vector $\mathbf{q}$ over S and any $y\in\mathbb{R}$,

$$
\int_y^\infty \left(f_{\theta^s}(z) - \left\langle (f_\vartheta(z))_{\vartheta\in S}, \mathbf{q}\right\rangle\right) dz \geq 0.
$$

Because $\left(\frac{W_n(\theta_i)}{1-W_n(\theta_i^*)}\right)_{\theta_i\neq\theta_i^*}$ is one such probability vector, it follows that

$$
\int_y^\infty \sum_{\theta:\theta_i=\theta_i^*} f_\theta(z)\prod_{j\neq i} W_n(\theta_j) - \sum_{\theta\in\Theta} f_\theta(z)\prod_{j=1}^d W_n(\theta_j)\, dz \geq 0.
$$

Hence, $W_n(\theta_i^*)$ is indeed a submartingale for all i. As it is bounded between 0 and 1, Doob's convergence theorem for submartingales establishes the first claim. It can similarly be shown that the inverse odds ratio $\frac{1-W_n(\theta_i^*)}{W_n(\theta_i^*)}$ is a nonnegative supermartingale. In particular, $\mathbb{E}_0\left(\frac{1-W_n(\theta_i^*)}{W_n(\theta_i^*)}\right)$ is at least weakly decreasing, so almost surely $\lim_{n\to\infty} W_n(\theta_i^*) > 0$.

Now, suppose that A1–A3 hold. Suppose for the sake of contradiction that for $i=1,\ldots,d$,

$\lim_{n\to\infty} W_n(\theta_i^*) = L_i$, where for some i it is the case that $L_i < 1$. Fixing this i, note that for any n

$$\frac{w_n(y|\theta_i = \theta_i^*)}{w_n(y|\theta_i \neq \theta_i^*)}$$

$$= \frac{\sum_{\substack{\theta_{-i}\in\Theta_{-i}\setminus\{\theta_{-i}^*\} \\ \theta_i=\theta_i^*}} f_\theta(y)\prod_{j\neq i} W_n(\theta_j) + f_{\theta^*}(y)\prod_{j\neq i} W_n(\theta_j^*)}{\sum_{\theta_{-i}\in\Theta_{-i}\setminus\{\theta_{-i}^*\}}\prod_{j\neq i} W_n(\theta_j)\sum_{\theta_i\in\Theta_i\setminus\{\theta_i^*\}} f_\theta(y)\frac{W_n(\theta_i)}{(1-W_n(\theta_i^*))} + \sum_{\theta_i\in\Theta_i\setminus\{\theta_i^*\}} f_{(\theta_i,\theta_{-i}^*)}(y)\frac{W_n(\theta_i)\prod_{j\neq i} W_n(\theta_j^*)}{(1-W_n(\theta_i^*))}}.$$

Using A3 and the arguments used to prove convergence, $y \in A$ implies

$$\begin{aligned}\sum_{\theta_{-i}\in\Theta_{-i}\setminus\{\theta_{-i}^*\}}\prod_{j\neq i} W_n(\theta_j)\sum_{\theta_i\in\Theta_i\setminus\{\theta_i^*\}} f_\theta(y)\frac{W_n(\theta_i)}{(1-W_n(\theta_i^*))} &\leq \sum_{\substack{\theta_{-i}\in\Theta_{-i}\setminus\{\theta_{-i}^*\} \\ \theta_i=\theta_i^*}} f_\theta(y)\prod_{j\neq i} W_n(\theta_j) \\ &\leq M\left(1-\prod_{j\neq i} W_n(\theta_j^*)\right) < \infty\end{aligned}$$

where $M < \infty$ is the upper bound in A3. In addition,

$$\begin{aligned}\lim_{n\to\infty} f_{\theta^*}(y)\prod_{j\neq i} W_n(\theta_j^*) = f_{\theta^*}(y)\prod_{j\neq i} L_j &\geq \sum_{\theta_i\in\Theta_i\setminus\{\theta_i^*\}} f_{(\theta_i,\theta_{-i}^*)}(y)\frac{W_n(\theta_i)\prod_{j\neq i} L_j}{(1-W_n(\theta_i^*))} + \varepsilon\prod_{j\neq i} L_j \\ &= \lim_{n\to\infty}\sum_{\theta_i\in\Theta_i\setminus\{\theta_i^*\}} f_{(\theta_i,\theta_{-i}^*)}(y)\frac{W_n(\theta_i)\prod_{j\neq i} W_n(\theta_j^*)}{(1-W_n(\theta_i^*))} + \varepsilon\prod_{j\neq i} W_n(\theta_j^*)\end{aligned}$$

By the second Borel-Cantelli Lemma, $\Pr\left(\limsup_{n\to\infty}\{y_n \in A\}\right) = 1$. Notice that the first term in the previous line is bounded by $M\prod_{j\neq i} W_n(\theta_j^*)$. Hence, almost surely

$$\limsup_{n\to\infty}\frac{w_n(y|\theta_i = \theta_i^*)}{w_n(y|\theta_i \neq \theta_i^*)} \geq \frac{M + \varepsilon\prod_{j\neq i} L_j}{M} > 1,$$

since $L_j > 0$ for all j. Recall that

$$\frac{1 - W_{n+1}(\theta_i^*)}{W_{n+1}(\theta_i^*)} = \frac{w_n(y_{n+1}|\theta_i = \theta_i^*)}{w_n(y_{n+1}|\theta_i \neq \theta_i^*)}\frac{1 - W_n(\theta_i^*)}{W_n(\theta_i^*)},$$

so it is impossible that the inverse odds ratio converges. Hence, $L_i = 1$ for all i almost surely. $\square$

The following is a simple subcase of Proposition 1 which arises when the distribution of y is a convolution of the X_i distributions:

Corollary 4. *Let $X_i \sim G_{\theta_i}$ for every i where G_{θ_i} has continuously differentiable density g_{θ_i}. Suppose that $y = \sum_{i=1}^d X_i$ and for all i, $\theta_i \in \Theta_i$, the ratio $\frac{g_{\theta_i^*}(x_i)}{g_{\theta_i}(x_i)}$ is weakly increasing. Then $W_n(\theta_i^*)$ is a submartingale for every i which converges almost surely in $(0,1]$.*

Proof. It suffices to verify that assumptions A1 and A2 hold. We do this by verifying that for every θ and every i, $\frac{f_{T_i(\theta)}(y)}{f_\theta(y)}$ is weakly increasing. Suppose without loss of generality that $i = d$. By the convolution formula and our hypothesis,

$$\begin{aligned}
f'_{T_d(\theta)}(y) f_\theta(y) =& \frac{\partial}{\partial y} \int \cdots \int g_{\theta_1}(x_1) \cdots g_{\theta_d^*}(y - x_1 - \cdots x_{d-1})\, dx_1 \cdots dx_{d-1} \\
& \cdot \int \cdots \int g_{\theta_1}(x'_1) \cdots g_{\theta_d}(y - x'_1 - \cdots x'_{d-1})\, dx'_1 \cdots dx'_{d-1} \\
=& \int \cdots \int g_{\theta_1}(x_1) g_{\theta_1}(x'_1) \cdots g'_{\theta_d^*}(y - x_1 - \cdots x_{d-1}) g_{\theta_d}(y - x'_1 - \cdots x'_{d-1}) dx_1 \cdots dx'_{d-1} \\
\geq& \int \cdots \int g_{\theta_1}(x_1) g_{\theta_1}(x'_1) \cdots g_{\theta_d^*}(y - x_1 - \cdots x_{d-1}) g'_{\theta_d}(y - x'_1 - \cdots x'_{d-1}) dx_1 \cdots dx'_{d-1} \\
=& f_{T_d(\theta)}(y) f'_\theta(y).
\end{aligned}$$

□

Lemma 5. *Let $\{f_n\}_{n=0}^\infty$ be a collection of probability densities on Y satisfying*

$$0 < \beta_1 \equiv \inf_{n,y} \frac{f_n(y)}{f_0(y)} \leq \sup_{n,y} \frac{f_n(y)}{f_0(y)} \equiv \beta_2 < \infty.$$

Then for any vector of nonnegative weights $(p_n)_{n=0}^\infty$ and any $n_0 \in \mathbb{N}$ satisfying $\sum_{k=0}^\infty p_k = 1$,

$$\left| D_{KL}\left(f_0 \middle\| \sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^\infty p_k \right) - D_{KL}\left(f_0 \middle\| \sum_{k=0}^\infty p_k f_k \right) \right| \tag{14}$$

$$\leq \max\left\{ \log \frac{\beta_1}{\beta_1 - (\beta_2 - \beta_1)\sum_{k=n_0}^\infty p_k}, -\log \frac{\beta_1}{\beta_1 + (\beta_2 - \beta_1)\sum_{k=n_0}^\infty p_k} \right\}. \tag{15}$$

Proof. This follows from the calculation:

$$\left| \int \left(\log \frac{f_0}{\sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^{\infty} p_k} - \log \frac{f_0}{\sum_{k=0}^{\infty} p_k f_k} \right) f_0 \right| dy$$
$$\leq \int \left| \log \frac{\sum_{k=0}^{\infty} p_k f_k}{\sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^{\infty} p_k} \right| f_0 \, dy. \tag{16}$$

where

$$\log \frac{\sum_{k=0}^{\infty} p_k f_k}{\sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^{\infty} p_k} \leq \log \frac{\sum_{k=0}^{\infty} p_k f_k}{\sum_{k=0}^{\infty} p_k f_k - (\beta_2 - \beta_1) \sum_{k=n_0}^{\infty} p_k f_0}$$
$$\leq \log \frac{\beta_1}{\beta_1 - (\beta_2 - \beta_1) \sum_{k=n_0}^{\infty} p_k} \tag{17}$$

and similarly

$$\log \frac{\sum_{k=0}^{\infty} p_k f_k}{\sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^{\infty} p_k} \geq \log \frac{\beta_1}{\beta_1 + (\beta_2 - \beta_1) \sum_{k=n_0}^{\infty} p_k} \tag{18}$$

Since $\int f_0 \, dy = 1$, (16) is bounded in magnitude by the larger of the magnitudes of (17) and (18).

□

Proof of Proposition 2. Denote lower and and upper limits in (5) by α_1 and α_2 respectively, and let the left side of (6) be β. For any fixed i, we claim that there is a $\delta \in (0,1)$ such that if

$$W_n^j(t_j) \geq \delta \text{ for all } j \tag{19}$$

then $\mathbb{E}\left[\log \frac{w_n\left(y_{n+1} \middle| \theta_i = t_i\right)}{w_n\left(y_{n+1} \middle| \theta_i \neq t_i\right)}\right] > 0$. First, note that (7) implies

$$L \equiv \inf_{f \in \mathrm{co}(\{f_\theta : \theta \in G_i(t)\})} D_{\mathrm{KL}}\left(f_{\theta_0} \,||f\,\right) > D_{\mathrm{KL}}\left(f_{\theta_0} \,||f_t\,\right). \tag{20}$$

Otherwise, there is a convergent subsequence of weights $(p_m^1, \ldots, p_m^{d-1}) \to (p^1, \ldots, p^{d-1}) \in \Delta(G_i(t))$

and corresponding densities $f^m \to f \in \mathrm{co}(\{f_\theta : \theta \in G_i(t)\})$. But by (5),

$$\int \left|\log \frac{f_{\theta_0}}{f}\right| f_{\theta_0}\, dy \le \int \log \beta f_{\theta_0}\, dy = \log \beta,$$

which allows use of the DCT to conclude that (7) is met with equality at f, a contradiction.

Now by (2), $w_n(y_{n+1}|\theta_i = t_i)$ is a weighted average of $f_\vartheta(y_{n+1})$ such that $\vartheta_i = t_i$; furthermore, by (5) one has $W_n(\theta = t) \ge \alpha_1 \prod_{j=1}^{d} W_n(t_j)$ and:

$$\sum_{\substack{\vartheta\in\Theta:\vartheta_i=t_i\\ \vartheta_{-i}\neq t_{-i}}} W_n(\theta=\vartheta) \le \sum_{j\neq i}\sum_{\substack{\vartheta\in\Theta:\vartheta_i=t_i\\ \vartheta_j\neq t_j}} W_n(\theta=\vartheta) \le \sum_{j\neq i}(1-W_n^j(t_j)).$$

From (3), $w_n(y_{n+1}|\theta_i \neq t_i)$ can similarly be expressed as a weighted average with weights satisfying

$$\sum_{\vartheta\in\Theta:\vartheta_i\neq t_i} W_n(\theta=\vartheta) = \sum_{s\in\Theta_i:s\neq t_i}\left(W_n(\theta_i=s,\theta_{-i}=t_{-i}) + \sum_{\substack{\vartheta\in\Theta:\theta_i=s\\ \theta_{-i}\neq t_{-i}}} W_n(\theta=\vartheta)\right).$$

By (5), $W_n(\theta_i = s, \theta_{-i} = t_{-i}) \ge \alpha_1 W_n^i(s)\prod_{j\neq i} W^j(t_j)$ and

$$\begin{aligned}\sum_{\substack{\vartheta\in\Theta:\theta_i=s\\ \theta_{-i}\neq t_{-i}}} W_n(\theta=\vartheta) &\le \alpha_2 W_n^i(s)\sum_{j\neq i}\sum_{\vartheta\in\Theta_{-i}:\vartheta_j\neq t_j}\prod_{k\neq i} W_n^k(\vartheta_k)\\ &= \alpha_2 W_n^i(s)\sum_{j\neq i}(1-W_n^j(t_j)).\end{aligned}$$

Thus, in calculating (2) the agent places little weight on ϑ that are not t, and in calculating (3), the agent places little weight on ϑ such that $\vartheta_{-i} \neq t_{-i}$. In particular, the weight placed on densities f_ϑ with $\vartheta \neq t$ in (2) is bounded by $\frac{\sum_{j\neq i}(1-W_n^j(t_j))}{\alpha_1\prod_{j=1}^{d} W_n(t_j)}$, and the weight placed on densities f_ϑ with $\vartheta_{-i} \neq t_{-i}$ in (3) is bounded above by:

$$\frac{\sum_{s\in\Theta_i:s\neq t_i}\alpha_2 W_n^i(s)\sum_{j\neq i}(1-W_n^j(t_j))}{\sum_{s\in\Theta_i:s\neq t_i}\alpha_1 W_n^i(s)\prod_{j\neq i} W^j(t_j)} = \frac{\alpha_2\sum_{j\neq i}(1-W_n^j(t_j))}{\alpha_1\prod_{j\neq i} W^j(t_j)}.$$

For every $\eta > 0$, it follows that there is a δ satisfying $\frac{\alpha_2 d\delta}{\alpha_1(1-\delta)^d} < \eta$ such that when $W_n^j(t_j) \ge 1-\delta$

for all j then the agent's beliefs in these ancillary states of the world which differ from t is bounded by η, conditioning either on $\theta_i = t_i$ or $\theta_i \neq t_i$. Suppose that (19) holds with $\delta > 0$ meeting this condition.

Now, application of (15) in Lemma 5 with $f_0 = f_{\theta_0}$, $f_1 = f_t$, $n_0 = 2$, $p_0 = 0$, and $p_1 = W_n^i(t_i) \geq 1-\delta$ along with $\beta_1 = \beta^{-1}, \beta_2 = \beta$ implies that

$$\left| D_{\mathrm{KL}}\left(f_0 \,||W_n(\cdot|\theta_i = t_i)\right) - D_{\mathrm{KL}}\left(f_0 \,||f_t\,\right)\right| \leq K(\beta,\eta),$$

where $\lim_{\eta\to 0} K(\beta,\eta) = 0$. Similarly, by letting $f_1,\ldots,f_{d-1}$ be the elements of $G_i(t)$ and $n_0 = d$, one finds

$$\left| D_{\mathrm{KL}}\left(f_0 \,||W_n(\cdot|\theta_i \neq t_i)\right) - D_{\mathrm{KL}}\left(f_0 \,||f\,\right)\right| \leq K(\beta,\eta),$$

where $f \in \mathrm{co}\left(\{f_\theta : \theta \in G_i(t)\}\right)$. By (20), if one choose δ and hence η to be small enough so that $K(\beta,\eta) < \frac{L - D_{\mathrm{KL}}\left(f_{\theta_0}||f_t\,\right)}{3} \equiv C$, the triangle inequality implies

$$D_{\mathrm{KL}}\left(f_0 \,||W_n(\cdot|\theta_i \neq t_i)\right) > C + D_{\mathrm{KL}}\left(f_0 \,||W_n(\cdot|\theta_i = t_i)\right).$$

where $C > 0$. Thus, by (1), (19) implies,

$$\mathbb{E}\left[\log \frac{w_n\left(y_{n+1}\middle|\theta_i = t_i\right)}{w_n\left(y_{n+1}\middle|\theta_i \neq t_i\right)}\right] = \int \left(\log \frac{w_n\left(y_{n+1}\middle|\theta_i = t_i\right)}{f_{\theta_0}} - \frac{\log w_n\left(y_{n+1}\middle|\theta_i \neq t_i\right)}{f_{\theta_0}}\right) f_{\theta_0}\, dy > C. \quad (21)$$

Since the number of marginals to be considered is $d < \infty$, one can pick δ and C above small enough so that (21) holds for all i. Let $\Omega = Y^{\mathbb{Z}_+}$ where Ω is equipped with the product σ-algebra and product measure inherited from f_{θ_0} (since the observations y_n are iid$\sim f_{\theta_0}$). For $1 \leq i \leq d$, define the following stochastic process $(X_n^i):\Omega\to[0,1]$:

$$X_0^i = \log\frac{W_0^i(t_i)}{1-W_0^i(t_i)}$$
$$X_{n+1}^i = \begin{cases} X_n^i & \text{if } \min_{\substack{1\le j\le d,\\ 0\le \ell\le n}} W_\ell^j(t_j) < 1-\delta \\ X_n^i + \log\frac{w_n(y_{n+1}|\theta_i=t_i)}{w_n(y_{n+1}|\theta_i\neq t_i)} - C & \text{otherwise}\end{cases}$$

Then X_n^i is a submartingale with respect to the filtration $\mathscr{F}_n = \sigma\left(\left(W_\ell^j(t_j)\right)_{\substack{1\le j\le d,\\ 0\le \ell\le n}}\right)$. Suppose that we pick $c\le\delta$ and $W_0^i(t_i)\ge 1-c$ for every i so that $X_0^i\ge\log\frac{1-c}{c}$. Then, let $\gamma=\log\frac{1-c}{c}-\log\frac{1-\delta}{\delta}$ and note that Azuma's inequality for submartingales implies that:

$$\begin{aligned}
&\mathrm{P}_{\theta_0}\left(\bigcap_{i=1}^d\left\{\log\frac{W_1^i(t_i)}{1-W_1^i(t_i)}\ge\log\frac{1-\delta}{\delta}+\frac{C}{2}\right\}\right)\\
&\quad\ge 1-\sum_{i=1}^d \mathrm{P}_{\theta_0}\left(\log\frac{W_1^i(t_i)}{1-W_1^i(t_i)}<\log\frac{W_0^i(t_i)}{1-W_0^i(t_i)}-\gamma+\frac{C}{2}\right)\\
&\quad\ge 1-\sum_{i=1}^d \mathrm{P}_{\theta_0}\left(X_1^i-X_0^i<-\gamma-\frac{C}{2}\right)\\
&\quad\ge 1-d\exp\left(\frac{-(\gamma+C/2)^2}{2\log\beta}\right).
\end{aligned}\tag{22}$$

Let E_n denote the event $\bigcap_{\substack{1\le j\le d,\\ 0\le \ell\le n}}\left\{\log\frac{W_\ell^i(t_i)}{1-W_\ell^i(t_i)}\ge\log\frac{1-\delta}{\delta}+\frac{\ell C}{2}\right\}$. Then similarly,

$$\begin{aligned}
\mathrm{P}_{\theta_0}(E_{n+1}|E_n) &= \mathrm{P}_{\theta_0}\left(\bigcap_{i=1}^d\left\{\log\frac{W_{n+1}^i(t_i)}{1-W_{n+1}^i(t_i)}\ge\log\frac{1-\delta}{\delta}+\frac{(n+1)C}{2}\right\}\middle| E_n\right)\\
&\ge 1-\sum_{i=1}^d \mathrm{P}_{\theta_0}\left(\log\frac{W_{n+1}^i(t_i)}{1-W_{n+1}^i(t_i)}<\log\frac{W_0^i(t_i)}{1-W_0^i(t_i)}-\gamma+\frac{(n+1)C}{2}\middle| E_n\right)\\
&\ge 1-\sum_{i=1}^d \mathrm{P}_{\theta_0}\left(X_{n+1}^i-X_0^i<-\gamma-\frac{(n+1)C}{2}\middle| E_n\right)\\
&\ge 1-d\exp\left(-\frac{(\gamma+(n+1)C/2)^2}{2(n+1)\log\beta}\right)\mathrm{P}_{\theta_0}(E_n)^{-1}.
\end{aligned}$$

Since $\mathrm{P}_{\theta_0}(E_n) = \mathrm{P}_{\theta_0}(E_1) \prod_{\ell=1}^{n-1} \mathrm{P}_{\theta_0}\left(E_{\ell+1} \middle| E_\ell\right)$ with $\mathrm{P}_{\theta_0}(E_1)$ as in (22), we obtain:

$$\mathrm{P}_{\theta_0}(E_n) \geq 1 - d \sum_{\ell=0}^{n-1} \exp\left(-\frac{(\gamma + (\ell+1)C/2)^2}{2(\ell+1)\log\beta}\right).$$

The sum on the right is majorized by $\sum_{\ell=0}^{\infty} \exp\left(-\frac{(\ell+1)C^2}{8\log\beta}\right)$, which converges, so by application of the DCT with respect to counting measure, $\lim_{\gamma\to\infty} d \sum_{\ell=0}^{\infty} \exp\left(-\frac{(\gamma+(\ell+1)C/2)^2}{2(\ell+1)\log\beta}\right) = 0$. It follows that γ, and thus c, can be chosen high enough so that $\mathrm{P}_{\theta_0}\left(\bigcap_{n=1}^{\infty} E_n\right) > 1 - \varepsilon$, which concludes. □

Proof of Lemma 2. First, since $p_H = q_H$ and $p_L = q_L$, the likelihood ratios given observations $y = 0$ and $y = 2$ are equal:

$$L_2^P = \frac{p_H}{p_L} = L_2^Q, \qquad L_0^P = \frac{1-p_H}{1-p_L} = L_0^Q.$$

Thus, O^P/O^Q is unchanged after these observations.

Second, O^P/O^Q increases after $y = 1$ whenever $P > Q$. Differentiating $L_1^P(Q)$ with respect to Q shows that it is decreasing in Q. By symmetry, if $P > Q$, then $L_1^P(Q)/L_1^Q(P) > 1$. Similarly, differentiating and collecting terms, $L_1^P(Q)/L_1^Q(P)$ is increasing in P. Thus, as P/Q grows, the ratio of likelihoods grows and thus O^P/O^Q increases by more. Thus, $L_1^P(Q_n)/L_1^Q(P_n) \geq L_1^P(Q_0)/L_1^Q(P_0) > 1$, and so $L_1^P(Q_n)/L_1^Q(P_n)$ is bounded below by a number strictly greater than 1. Since all probabilities are strictly positive, by the Strong Law of Large Numbers asymptotically there will be an infinite number of $y = 1$ observations, and thus O^P/O^Q increases without bound. By similar argument, if $P < Q$ then O^P/O^Q decreases to zero. □

In the following lemma, for a sequence $\omega \in \{0,1,2\}^{\mathbb{N}}$, let ω_n denote the n^{th} element of ω:

Lemma 6. *If $\lim_{n\to\infty} P_n$ exists in $[0,1]$ almost everywhere and $p_H, p_L \in (0,1)$, then*

$$Pr\left(\lim_{n\to\infty} P_n \in (0,1)\right) = 0. \tag{23}$$

Proof. Let $\Omega = \{0,1,2,\}^{\mathbb{N}}$ denote the set of all possible sequences of y signals. Then $P_n \equiv P_n(\omega)$ can be written as a function of points $\omega \in \Omega$. Ω is endowed with the natural probability measure over signal observations and the σ-algebra generated by $P_n(\omega), Q_n(\omega)$. We claim that any point

ω such that $\lim_{n\to\infty} P_n(\omega) \in (0,1)$ has $\omega_n \in \{0,2\}$ for only finitely many indices n. This is easy to see for $\omega_n = 2$: assume to the contrary that infinitely many $(n_m) \subset \mathbb{N}$ satisfy $\omega_{n_m} = 2$, and let $\lim_{n\to\infty} p_n(\omega) = c \in (0,1)$. Then taking such m arbitrarily high, $(P_{n_m+1}) \to \frac{cp_H}{cp_H+(1-c)p_L} > c$, a contradiction. The proof when $\omega_n = 0$ infinitely often is similar. (23) follows immediately by the Second Borel-Cantelli Lemma. □

Proof of Proposition 3. We prove for P_n (the proof for Q_n is similar). Suppose that $\theta = \sigma = 1$. The discrete version of Corollary 4, using summation by parts instead of integration b parts in Proposition 1, establishes that P_n converges almost surely in $(0,1]$. Lemma 6 then implies that $P_n \overset{\text{a.s.}}{\to} 1$. The proof when $\theta = \sigma = 0$ is similar. □

Proof of Proposition 4. By Lemma 2 the ratio of the odds-ratios, O^P/O^Q converges to infinity almost surely. This implies that $P_n \overset{\text{a.s.}}{\to} 1$ almost surely. Since $O^P/O^Q \to \infty$, at least one of $P_n \overset{\text{a.s.}}{\to} 1$ or $Q_n \overset{\text{a.s.}}{\to} 0$. However, if $Q = 0$ then the problem of Bayesian learning is isomorphic to learning whether $\sigma = 1$ given $P = 0$. That is, there exists $\varepsilon > 0$ such that if $P_n < \varepsilon$, then Q_n is a submartingale. By Doob's Martingale Convergence Theorem, this problem converges to the truth ($Q = 1$) almost surely (see Diaconis and Freedman (1986)). By symmetry, if $Q_n < \varepsilon$ then $P_n \overset{\text{a.s.}}{\to} 1$. □

Proof of Lemma 3

Consider the log-odds ratios, for which we have the following recursions:

$$\log O^P_{n+1}(2) = \log \frac{p_H}{p_L} + \log O^P_n \tag{24}$$

$$\log O^P_{n+1}(0) = \log \frac{1-p_H}{1-p_L} + \log O^P_n \tag{25}$$

$$\log O^P_{n+1}(1) = \log \frac{p_H(1-\overline{q}_n)+(1-p_H)\overline{q}_n}{p_L(1-\overline{q}_n)+(1-p_L)\overline{q}_n} + \log O^P_n. \tag{26}$$

Consider the random variable $\Delta\log O^P_{n+1}(\overline{q})$ (and the analogous expression for Q), whose expected value can be evaluated explicitly as:

$$\begin{aligned}\mathbb{E}\left[\Delta\log O^P_{n+1}(\overline{q})\right] =& p_L q_H \log\frac{p_H}{p_L} + (1-p_L)(1-q_H)\log\frac{1-p_H}{1-p_L} \\ &+ (p_L(1-q_H)+(1-p_L)q_H)\log\frac{p_H(1-\overline{q})+(1-p_H)\overline{q}}{p_L(1-\overline{q})+(1-p_L)\overline{q}}.\end{aligned}$$

The restriction of the transition function to the interior of the set $[p_L, p_H]\times[q_L, q_H]$ often is less relevant than its restriction to these endpoints.

We proceed in three lemmas, which deal with different cases in the parameters.

Lemma 7. *Let $\theta = 0$ and $\sigma = 1$. If $\mathbb{E}\left[\Delta\log O^P_{n+1}(q_L)\right] > 0$ and $\mathbb{E}\left[\Delta\log O^Q_{n+1}(p_H)\right] < 0$, then there exist (P_0, Q_0) such that $P_0 \to 1$ and $Q_0 \to 0$ with positive probability tending to 1 as $Q_0 \to 0$ and $P_0 \to 1$.*

Proof. The proof of this lemma will follow from Proposition 2, which is proved below. Specifically, the hypothesis of the lemma implies (7) holds with the true state of the world corresponding to $(\{\theta=0\},\{\sigma=1\})$ and the false state t corresponding to $(\{\theta=1\},\{\sigma=0\})$. □

Lemma 8. *If $\mathbb{E}\left[\Delta\log O^P_{n+1}(q_L)\right] < 0$, then $P \xrightarrow{a.s.} 0$, and if $\mathbb{E}\left[\Delta\log O^P_{n+1}(p_H)\right] > 0$, then $Q \xrightarrow{a.s.} 1$*

Proof. The proof follows the method of Lemma 7 by setting $Q^* = 1$ and $P^* = 0$. □

Lemma 9. *Let $\mathbb{E}\left[\Delta\log O^P_{n+1}(q_L)\right] < 0$. Then $P \xrightarrow{a.s.} 0$, and if $\mathbb{E}\left[\Delta\log O^Q_{n+1}(p_L)\right] < 0$ then $Q \xrightarrow{a.s.} 0$, whereas if $\mathbb{E}\left[\Delta\log O^Q_{n+1}(p_L)\right] > 0$ then $Q \xrightarrow{a.s.} 1$.*

Alternately, if $\mathbb{E}\left[\Delta\log O^Q_{n+1}(p_H)\right] > 0$, then $Q \xrightarrow{a.s.} 1$ and if $\mathbb{E}\left[\Delta\log O^P_{n+1}(q_L)\right] < 0$ then $P \xrightarrow{a.s.} 0$, whereas if $\mathbb{E}\left[\Delta\log O^P_{n+1}(q_L)\right] > 0$ then $P \xrightarrow{a.s.} 1$.

Proof. The convergence $P \xrightarrow{\text{a.s.}} 0$ follows from Lemma 8. As for convergence of Q, if $\mathbb{E}\left[\Delta\log O^P_{n+1}(p_L)\right] < 0$, then $\mathbb{E}\left[\Delta\log O^P_{n+1}(\overline{p})\right]$ is upper-bounded by a strictly negative number everywhere, and a law of large numbers argument suffices. On the other hand, if $\mathbb{E}\left[\Delta\log O^P_{n+1}(p_L)\right] > 0$, then as P concentrates almost surely at 0, $\mathbb{E}\left[\Delta\log O^Q_{n+1}\right]$ becomes lower bounded (for all sufficiently high n) by a positive $\varepsilon \in (0, \mathbb{E}\left[\Delta\log O^P_{n+1}(p_L)\right])$ almost surely, whence $Q \xrightarrow{\text{a.s.}} 1$. An obvious symmetry establishes the second claim. □

Proofs of Continuity Results

Here we provide the essential results necessary for the main claims of our paper. In the Internet Appendix we provide a proof of more general features of continuity which builds on these results. In the proof, we discuss the relevance of the random dynamical theory of cocycles to our model. Lemma 4 is detailed in the internet appendix and shown to hold on a larger set of random dynamical systems with a Bernoulli shift as a random component. (Schenk-Hoppé and Schmalfuß, 2001) expands on our discussion of random dynamical systems via cocycles.

Consider the function $f(p,q)$ giving the probability of convergence of (P_n, Q_n) to $(1,0)$ if $P_0 = p$, $Q_0 = q$ given p_H, p_L, q_H, q_L. This function has some nice properties. Suppose that $\theta = 0, \sigma = 1$. First,

$$\begin{aligned} f(p,q) = & p_L q_H f\left(P'((p,q),2), Q'((p,q),2)\right) + (1-p_L)(1-q_H)\left(P'((p,q),0), Q'((p,q),0)\right) \\ & + \Big(p_L(1-q_H) + (1-p_L)q_H\Big) f\left(P'((p,q),1), Q'((p,q),1)\right). \end{aligned}$$

Furthermore, observe that for a *fixed* q, $f(p,q)$ is monotonically increasing in p:

Lemma 10. *For $p' \geq p$ and $q' \leq q$, $f(p',q') \geq f(p,q)$. Therefore, $f(\cdot,q)$ is continuous almost everywhere at any fixed q.*

Proof. Fix a history $\omega \in \Omega$. Let $P_n(\omega), Q_n(\omega)$ correspond to initial condition (p,q) and $P'_n(\omega), Q'_n(\omega)$ be defined to correspond to (p',q'). The claim can be verified by on the hypothesis that in each n, we have

$$P'_n(\omega) \geq P_n(\omega), \quad Q'_n(\omega) \leq Q_n(\omega). \tag{27}$$

The argument is accomplished with the observation that (i) If $y_n(\omega) = 0, 2$, the ordering is preserved by monotonicity of the relevant functions. (ii) If $y_n(\omega) = 1$, then a smaller value of Q_n corresponds with a larger increase in P_n. Conversely, a larger value of P_n corresponds with a smaller increase (larger decrease) in Q_n. This can be easily verified by noting that the odds ratio O^P_{n+1} increases by more when $y_n = 1$ and Q_n is smaller. Conversely, O^Q_{n+1} increases by less when

$y_n = 1$ and P_n is larger. Because $O^P_{n+1} \geq O^{P'}_{n+1}$ if and only if $P_{n+1} \geq P'_{n+1}$, the inequalities in (27) are indeed preserved. $\square$

Proof of Corollary 3. Note that

$$n \log \frac{p_H}{p_L} + \left\lceil -n \frac{\log \frac{q_H}{q_L}}{\log \frac{1-q_H}{1-q_L}} + 1 \right\rceil \log \frac{1-p_H}{1-p_L} = n \left(\log \frac{p_H}{p_L} + \left(-\frac{\log \frac{q_H}{q_L}}{\log \frac{1-q_H}{1-q_L}} + o(1) \right) \log \frac{1-p_H}{1-p_L} \right)$$

$$n \log \frac{q_H}{q_L} + \left\lceil -n \frac{\log \frac{q_H}{q_L}}{\log \frac{1-q_H}{1-q_L}} + 1 \right\rceil \log \frac{1-q_H}{1-q_L} < 0.$$

Hence, there must exist positive integers n, $m \approx -n \frac{\log \frac{q_H}{q_L}}{\log \frac{1-q_H}{1-q_L}}$ such that $n \log \frac{p_H}{p_L} + m \log \frac{1-p_H}{1-p_L} > 0$, and $n \log \frac{q_H}{q_L} + m \log \frac{1-q_H}{1-q_L} > 0$. It follows that if $y = 2$ for $n\ell$ times and $y = 0$ for $m\ell$ times, as ℓ becomes arbitrarily large, $\log O^P_{\ell(n+m)}$ becomes arbitrarily large and $\log O^Q_{\ell(n+m)}$ becomes arbitrarily small from any initial prior (P, Q). In particular, if f does not vanish for all (p, q), by Lemma 10, there is a critical O^{P*} and O^{Q*} such that for all pairs (O^P, O^Q) with $O^P \geq O^{P*}$ and $O^Q \leq O^{Q*}$, there is a positive probability that $O^P \to \infty$. Because this critical threshold can be reached in a finite number of steps from any prior, any prior has a positive probability of $O^P \to \infty$. The second claim is proved similarly by exchanging 'P' and 'Q'. $\square$