

# DIMENSIONALITY AND DISAGREEMENT:

## Asymptotic belief divergence in response to common information

Isaac Loh\*

Gregory Phelan<sup>†</sup>

This version: April 18, 2017

### Abstract

When information is of lower dimension than the model generating the data, Bayesian updating of marginal probabilities need not converge to the truth. Because the information is of lower dimension than the model, agents face an identification problem, affecting the role of data in inference. We provide conditions under which learning is asymptotically inconsistent with positive probability, and sometimes almost surely. Robustly, two agents with differing priors who observe identical, unambiguous data may disagree forever, with stronger disagreement the more data is observed. Agents use common observations to differentially update marginal beliefs about different parameters.

**Keywords:** Heterogeneous beliefs, polarization, learning, Bayesian updating, bounded rationality.

**JEL classification:** D10, D80, D84.

---

\*Northwestern University. Email: IsaacLoh2015@u.northwestern.edu .

<sup>†</sup>Department of Economics, Williams College. Email: gp4@williams.edu.

We are grateful for feedback from Matthew Chao, Asim Khwaja, Rachel Kranton, Bruce Sacerdote, Larry Samuelson, and Arunava Sen.

# 1 Introduction

People disagree and sometimes in big, persistent ways. People hold very different views about which policies best achieve outcomes;<sup>1</sup> justices often come to opposing decisions about cases; professional forecasters and central bankers make different forecasts for macroeconomic variables (Crump et al., 2014; Andrade et al., 2016);<sup>2</sup> many patterns in financial markets strongly suggest investor disagreement (Hong and Stein, 2007). Since Aumann (1976) it is well-known that when beliefs are common knowledge and people have common priors, then people cannot agree to disagree, and Geanakoplos and Polemarchakis (1982) show that people cannot disagree forever so long as they can communicate their beliefs (over finite spaces). However, many instances of disagreement are publicly known and persistent (e.g., political and social beliefs), suggesting that priors are different. Furthermore, disagreements sometimes persist or even grow, even when people face the same “evidence” and especially if they see lots of evidence. In reality people often see the same information and continue to disagree.<sup>3</sup> And yet, in many cases disagreements do not grow and people appear to perceive evidence in similar ways (Gerber and Green, 1999). Why might beliefs diverge in one case and not in others?

In this paper, we present conditions for when small differences in initial beliefs can lead to growing disagreements in light of common information, which can rationally be used as “evidence” in different ways by different people. Specifically, we consider when people may use the same evidence to draw inference about different variables, and as a result interpret the same data as evidence in favor of their prior beliefs. We consider agents that use Bayes’ rule to update beliefs about marginal probabilities but assume that beliefs are independent (bounded rationality). We show that when world is of higher dimensionality than the data people receive, (i) updated posterior beliefs will converge, though not necessarily to the true values; (ii) as a result, heterogeneous priors

---

<sup>1</sup>People even disagree about whether political disagreements are about values or about facts.

<sup>2</sup>Furthermore, 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.

<sup>3</sup>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.

may diverge in light of common information; (iii) these results are likely to occur when unidentified observations occur frequently, but beliefs converge to the truth when unidentified observations are relatively infrequent. Thus, when divergence occurs, beliefs are likely to converge to values confirming initial beliefs, whether those beliefs are correct or not. Specifically, posterior beliefs converge to confirm *relative* initial beliefs: what matters is which beliefs are held more strongly, not the absolute levels of beliefs.

To illustrate how dimensionality affects updating, 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 evidence 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. What these examples illustrate is that the world is *multi-dimensional*—there are a number of factors that contribute to what we see—but the data we see is nearly always of lower dimensionality than the world. In other words, we live with identification problems. When our data are of lower dimensionality than that of the generating process, some 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.

What is most interesting is not that people can disagree after a single observation, but that with *many* observations people may continue to disagree, and disagree further, perhaps even becoming completely convinced of their beliefs. For example, economists rarely switch from hawks to doves.<sup>4</sup> In other words, disagreement can persist or increase when people observe many identical data points. Belief divergence can occur even if many observations are definitively revealing about the state of the world, so long as enough observations can rationally be used to make inference about different underlying variables.<sup>5</sup>

To show these results, we consider a simple model in which there are two state variables,  $\theta$  and  $\sigma$ , that determine whether realization variables  $t$  and  $s$  will be successes or failures. What is observed is not  $t$  and  $s$  but the sum: agents see the number of successes but not the composition. Hence, observations of “two” or “zero” successes are clearly identified, while observations of “one” are “unidentified.” Agents hold beliefs about the states  $\theta$  and  $\sigma$ , which can be good or bad, and update their beliefs about the states in light of observations. In other words, the marginal beliefs about  $\theta$  and  $\sigma$  are updated sequentially, but not the joint distribution, implying that *beliefs* about the two states are always independent. Accordingly, when interpreting “unidentified observations” (ones), the likelihood function for one state (say  $\theta$ ) depends on beliefs about the other state (say  $\sigma$ ). Thus, agents update beliefs differently when observing unidentified observations. Thus, differing priors lead to divergent likelihood functions and divergent posteriors.

When unidentified observations are relatively unlikely, beliefs converge to the truth almost surely for any priors. However, when unidentified observations are relatively likely, then robustly

---

<sup>4</sup>Malmendier et al. (2017) show that personal experiences of inflation strongly influence the hawkish or dovish leanings of central bankers. This is evidence that priors influence how central bankers interpret the same data as FOMC members.

<sup>5</sup>Our paper is about how *repeated* observations can lead to belief divergence. Thus, an example would be that different people interpret multiple cases of police shootings in ways that lead to increased polarization (for evidence see Pew Research Center, 2014).

with positive probability beliefs converge, though not necessarily to the truth. Instead, if a person has a prior that it is *more* likely that the good aggregate state is  $\theta$ , then asymptotically with positive probability the person will believe *with certainty* that  $\theta$  is the good state and that  $\sigma$  is the bad state (vice versa if the prior is that  $\sigma$  is more likely to be good). Thus, if two people disagree about which state is more likely to be good—regardless of the overall level of beliefs about each state—then robustly beliefs diverge asymptotically with positive probability. Furthermore, under certain conditions beliefs diverge asymptotically with probability one. In these cases, more data leads to greater polarization rather than greater agreement.

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 (polarization). In addition, we provide simulation results illustrating that it is almost always the case that polarization can occur with positive probability when unidentified observations are relatively likely. Thus, high-dimensionality may lead to asymptotic divergence when unidentified observations are likely, but will lead to asymptotic agreement when identified observations are relatively likely. Or put differently, when unidentified observations are likely, agents in an economy with *common* priors may have collective beliefs converge to something other than the truth.

The remainder of this section discusses our connection to the existing literature. The rest of the paper is structured as follows. Section 2 presents the baseline binary model. We present theoretical results and simulation results to characterize the limiting properties of beliefs. Section 3 presents a more general  $N$ -dimensional discrete multivariate setting. We provide generalizing results for the properties of the binary model. Section 4 discusses implications. The final section concludes.

**Related Literature** Our mechanism, that agents use their prior beliefs to make different inferences, is shared by others in the literature. While we emphasize asymptotic results, several papers consider 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).<sup>6</sup> Similarly, Baliga et al. (2013) show that 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 consider that agents receive information about only one of two states, rather than information that is a function of the two states, and must make inference on what action is preferred, in light of beliefs about the other state (disagreement persists because agents discount information filtered through the actions of others).

We show that even with infinitely many signals, beliefs may diverge even when many of those signals are in fact extreme. In fact, it is possible that disagreement grows with the number of realizations, diverging to completely opposite views of the world. Like the intermediate results of Benoît and Dubra (2014) and Baliga et al. (2013), our asymptotic result holds when the true state of the world is in an “intermediate region”, which will lead toward a higher frequency of unidentified signals. Thus, disagreement persists asymptotically only when the fundamentals are not extreme, leading to many unidentified observations.

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.<sup>7</sup>

Fryer et al. (2015) consider a model in which agents may receive “ambiguous” signals, which

---

<sup>6</sup>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).

<sup>7</sup>Nonetheless, our model can also deliver the result of asymptotic disagreement even when learning occurs. See Proposition 1 and Lemma 4.

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. Crucially, if ambiguous signals could be stored as such, learning (and agreement) would occur. Their setup shares many features with ours; however, our paper is fundamentally about the difference between the *dimensionality* of observations and the generating parameters, not about ambiguity. In our model, a signal of one is “ambiguous” only in the sense that the underlying values cannot be *identified*, and the way that agents update their posteriors depends on their beliefs—but agents *unambiguously know* that a one is a one and store the observation as a one; the observation is perfectly clear, but how it is interpreted differs because of the identification problem.

Our treatment of dimensionality in learning is applicable to the broad class of inference problems in economics in which agents are tasked with learning the distribution of a multidimensional  $x$  by observing oscillations in the one-dimensional function (or functional)  $g(x)$  (this was the case in the two examples discussed earlier in the introduction). In our model, agents are completely rational, but there is a fundamental limitation, inherent to the model but not agents, in the ability of information to identify the model.<sup>8</sup>

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

---

<sup>8</sup>In the model of Fryer et al. (2015), agents interpret observations differently because of bounded rationality and because ambiguous signals are sufficiently likely. Therefore, divergence occurs when beliefs about the states are greater than or less than one-half and ambiguous signals occur with a sufficiently high frequency. In our model, divergence can occur when beliefs about one state are *greater* than the other, regardless of how high each belief is, whenever unidentified observations occur relatively frequently. Thus, divergence occurs when *relative* beliefs differ and the identification problem is sufficiently severe. Furthermore, as we discuss below, in our model, both agents will converge to the truth when the fundamentals in the model are not polarized (states are either (1,1) or (0,0)), whereas polarization occurs when fundamentals are, in a sense, polarized (states are (1,0) or (0,1)).

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.<sup>9</sup>

## 2 The Model

We first provide a simple example to illustrate how Bayesian learning is affected by the identification problem. We then present the dynamic model with binary states and characterize asymptotic properties. Section 3 generalizes our results to an  $N$ -dimensional discrete setting.

### 2.1 A Simple Static Example

Let there be two underlying state variables  $t$  and  $s$  that jointly determine the aggregate realization  $y(t, s)$ . The two state variables can each take two values in  $\{0, 1\}$ . The aggregate realization is defined by the following function:

$$y(t, s) = t + s. \tag{1}$$

In other words, a realization of  $y = 2$  or  $y = 0$  definitively says the value of the underlying variables, but a realization of 1 yields an identification problem: there are two possible sets of parameters that would provide that realization.<sup>10</sup>

Let there be two agents denoted  $i = 1, 2$ . Agent  $i$  holds beliefs  $\Pr_i(t = 1) = P_i$  and  $\Pr_i(s = 1) = Q_i$ . Agents are rational: when they see the realization  $y$  they update their beliefs  $P_i$  and  $Q_i$  according to Bayes' Rule. Denote posterior beliefs by  $\hat{P}_i$  and  $\hat{Q}_i$ . Consider two agents faced with a moderate realization  $y = 1$  trying to learn the effectiveness of the stimulus. From Bayes' Rule the

---

<sup>9</sup>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.

<sup>10</sup>For example, one could interpret  $t$  as the strength of the economy and  $s$  as the effectiveness of a stimulus package. A strong economy and an effective stimulus lead to high output (similarly for weak variables), but a mix of strong economy and ineffective stimulus, or the reverse, will lead to a moderate realization.

posteriors, given observation  $y = 1$ , are given by

$$\begin{aligned} \Pr(t = 1, s = 0) &= \hat{P} = \frac{P(1-Q)}{P(1-Q) + (1-P)Q}, \\ \Pr(t = 0, s = 1) &= \hat{Q} = \frac{Q(1-P)}{P(1-Q) + (1-P)Q}. \end{aligned}$$

Crucially, the posterior for  $P$  depends on beliefs about  $Q$  (vice versa). In fact, the likelihood ratio for the odds ratio for  $P$  is  $\frac{1-Q}{Q}$  and for  $Q$  is  $\frac{1-P}{P}$ , which is to say the likelihood ratio is completely determined by beliefs about the other variables. How might beliefs diverge in light of information? Suppose beliefs satisfy  $P_1 > \frac{1}{2} > P_2$  and  $Q_2 > \frac{1}{2} > Q_1$ . Thus, agent 1 is optimistic about  $t$  but not  $s$ , and agent 2 is optimistic about  $s$  and not  $t$ . In this case beliefs will diverge after  $y = 1$ . However, beliefs can also diverge relative to each other without moving farther apart. We say that divergence increases *relatively* if  $\frac{\hat{P}_1}{\hat{P}_2} > \frac{P_1}{P_2}$  and  $\frac{\hat{Q}_2}{\hat{Q}_1} > \frac{Q_2}{Q_1}$ .

**Lemma 1.** *Suppose  $P_1 > \frac{1}{2} > P_2$  and  $Q_2 > \frac{1}{2} > Q_1$ . Then posteriors diverge further if  $y = 1$ :  $\hat{P}_1 > P_1 > P_2 > \hat{P}_2$  and  $\hat{Q}_2 > Q_2 > Q_1 > \hat{Q}_1$ . Relative divergence occurs if*

$$\frac{\Pr_1(s = 0)}{\Pr_2(s = 0)} > \frac{\Pr_1(y = 1)}{\Pr_2(y = 1)}, \quad \frac{\Pr_2(t = 0)}{\Pr_1(t = 0)} > \frac{\Pr_2(y = 1)}{\Pr_1(y = 1)}.$$

The results follow immediately from the equations for posteriors. When beliefs differ in this sense, the same information leads to divergence. Agent 1 becomes relatively more optimistic about  $t$  and less optimistic about  $s$ , and the reverse is true for agent 2. Thus, agents draw divergent conclusions from the same data; the observation works as evidence in favor of greater disagreement.

## 2.2 Dynamic Model

In the static setup a realization of  $y = 0$  or  $y = 2$ —evidence that cannot be interpreted differently—would cause beliefs to converge, and such a realization would cause both agents to admit that they were wrong about something. Thus, one might wonder if given enough observations beliefs must eventually converge. We now extend the model to answer precisely this question and show that, when observations of  $y = 1$  are relatively likely, more observations increase the likelihood

of polarization. We characterize the limiting behavior of posterior beliefs with theoretical and simulation results. Proofs not in the text are given in the appendix.

**Setup** Suppose now that there are two aggregate states,  $\theta$  and  $\sigma$ , which determine the frequency of  $t$  and  $s$  respectively, and the aggregate states take binary values (think “good” or “bad”). In particular,  $t$ -successes occur with higher probability when  $\theta = 1$  ( $\theta$  is “good”), and  $s$ -successes occur with higher probability when  $\sigma = 1$ . Thus,

$$\Pr(t = 1 | \theta = 1) = p_H > p_L = \Pr(t = 1 | \theta = 0), \text{ and} \quad (2)$$

$$\Pr(s = 1 | \sigma = 1) = q_H > q_L = \Pr(s = 1 | \sigma = 0), \quad (3)$$

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 aggregate states  $\theta$  and  $\sigma$  determine the frequency of observations  $y \in \{0, 1, 2\}$ . The mapping from  $\theta$  and  $\sigma$  to  $t$  and  $s$  is common knowledge: agents only disagree about the likelihood of  $\theta$  and  $\sigma$ , but not about how those states translate into realizations  $y$ .<sup>11</sup>

**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 aggregate states 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.

There are two important ways in which agents are boundedly rational in how they form beliefs. First, since agents update sequentially, agents are not strictly Bayesian in the classical sense of Doob in which agents obtain belief  $P_{n+1}$  by updating a fixed *initial prior*  $P_0$  in conjunction with all  $n + 1$  observations at once. Since the learning process is path dependent (Lemma 3), sequential updating matters for the results. Second, the state-space is truly 4-dimensional  $\{0, 1\} \times \{0, 1\}$ —the two possible values for  $\theta$  and  $\sigma$ —but agents only store the marginal probabilities for the variables

---

<sup>11</sup>Contrast this assumption with the information structure in Acemoglu et al. (2016), where a one-dimensional state  $\theta$  produces signals with probability  $p_\theta$ , and agents potentially disagree about  $p_\theta$  as well as the probability of  $\theta$ .

$\theta$  and  $\sigma$ . In other words, agents ought to have beliefs  $\pi_{ij} = Pr(\theta = i, \sigma = j)$  and update  $\pi_{ij}$  using Bayes' Rule. Since our agents only update  $P$  and  $Q$ , the correlation of beliefs are lost upon updating, and agents always assume independent *beliefs*.

Both of these deviations can be rationalized by a sparsity-based model of bounded rationality following Gabaix (2014, 2016). Sequential updating dramatically saves on information storage because agents can store only the beliefs  $P_n$  and  $Q_n$  rather than  $n$  observations together with initial beliefs.<sup>12</sup> Similarly, updating on marginal beliefs and storing  $P_n$  and  $Q_n$  instead of the full set of beliefs  $\pi_{ij,n}$  saves on storage information, and this is more true the larger is the state-space. With  $N$  underlying binary variables instead of 2, the full belief distribution would be  $N \times N$  values. Thus, storing only  $N$  marginal probabilities greatly reduces the information necessary to store.

What that in mind, we can derive the evolution of beliefs using odds ratios. Define  $\bar{p} = Pp_H + (1 - P)p_L$  and  $\bar{q} = Qq_H + (1 - Q)q_L$  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). It is convenient to work with odds ratios. Define  $O^P = \frac{P}{1-P}$ , and  $O^Q = \frac{Q}{1-Q}$ . Let  $O^P(y), O^Q(y)$  denote the updated odds ratio after observing  $y$ . Then Bayes' Theorem applied to marginal probabilities yields

$$O_{n+1}^P(2) = \frac{p_H}{p_L} O_n^P, \quad O_{n+1}^Q(2) = \frac{q_H}{q_L} O_n^Q, \quad (4)$$

$$O_{n+1}^P(0) = \frac{1 - p_H}{1 - p_L} O_n^P, \quad O_{n+1}^Q(0) = \frac{1 - q_H}{1 - q_L} O_n^Q, \quad (5)$$

$$O_{n+1}^P(1) = \frac{(p_H(1 - \bar{q}_n) + (1 - p_H)\bar{q}_n)}{(p_L(1 - \bar{q}_n) + (1 - p_L)\bar{q}_n)} O_n^P, \quad O_{n+1}^Q(1) = \frac{(q_H(1 - \bar{p}_n) + (1 - q_H)\bar{p}_n)}{(q_L(1 - \bar{p}_n) + (1 - q_L)\bar{p}_n)} O_n^Q, \quad (6)$$

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  $P, Q$ . Denote the likelihood ratios by  $L_y^P(Q)$  and  $L_y^Q(P)$ . The following result is immediate.

---

<sup>12</sup>Sequential updating may reflect bounded rationality, but it need not. Even though sequential updating is not the classically Bayesian way to update beliefs, sequential updating can be rational in the sense that it need not reflect behavioral biases or limitations (in other words, it is consistent with a rational use of information).

**Lemma 2.** *Posteriors behave as follows:*

*if  $y_{n+1} = 2$  then  $P_{n+1} > P_n$ , and  $Q_{n+1} > Q_n$*

*if  $y_{n+1} = 0$  then  $P_{n+1} < P_n$ , and  $Q_{n+1} < Q_n$*

*if  $y_{n+1} = 1$  then  $P_{n+1} > P_n$  iff  $\bar{q}_n < \frac{1}{2}$ , and  $Q_{n+1} > Q_n$  iff  $\bar{p}_n < \frac{1}{2}$ .*

In other words, 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 polarized.

**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 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. Furthermore, our agents can only possibly disagree about initial priors; divergence would be even more likely if agents also disagreed about the model parameters. The generalization to  $N$  discrete states is presented in Section 3.

Finally, 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  $\theta$  and  $\sigma$  separately. Or there may be an additional variable

whose realization will be made known in the far future, whose value could simply equal  $\theta$  or  $\sigma$ . 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  $\theta$  and  $\sigma$  separately would still matter.

## 2.3 Theoretical Results

The main results of this section are as follows: when unidentified observations are relatively unlikely, beliefs converge to the truth (Proposition 1); when unidentified observations are relatively likely and the aggregate states are symmetric, then *with probability one* beliefs converge to confirm relative initial priors (Proposition 2); robustly, when unidentified observations are relatively likely there is a *positive probability* that beliefs converge to confirm relative initial priors (Proposition 5).

First, this process is path dependent: the order of observations matters for determining the posterior.

**Lemma 3** (Path Dependence). *The stochastic process defined by the posterior is path dependent. In other words, the posterior at any time  $T > 1$  depends on the order of observations.*

The result is very intuitive. Since the evolution of posteriors depends on likelihood ratios, and since the likelihood ratios depend on priors, how an agent updates after seeing a one depends on what information has already been received. Naturally, path dependence is precisely a feature of the process. Agents interpret information differently—they have different likelihood functions—because their likelihood functions depend on their beliefs about the other variable. But these beliefs evolve as well. Thus, while our agents are completely rational and Bayesian (we have not made any behavioral assumptions) our model delivers learning that has a flavor of psychological short-sightedness: an agent need not recall the entire history of observations in order to update the posterior, but simply consults the current prior, which includes the information from the history of observations.

Our first result concerns a case where individuals converge asymptotically to the correct beliefs. Specifically, when  $\theta = \sigma = 1$  or  $\theta = \sigma = 0$ , then beliefs converge to the truth. 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 1.** *If  $\theta = \sigma = 1$  or  $\theta = \sigma = 0$ , then (respectively)  $P_n, Q_n \xrightarrow{a.s.} 1$  and  $P_n, Q_n \xrightarrow{a.s.} 0$ .*

However, when intermediate observations ( $y = 1$ ) are relatively likely, as occurs when  $\theta \neq \sigma$ , beliefs need not converge to the truth. We now focus on these results.

### 2.3.1 Incorrect Beliefs in the Symmetric Case

While path dependence makes it more difficult to characterize the limiting properties of the process, we can still prove an important result when the states  $\theta$  and  $\sigma$  are symmetric. Crucially, asymptotically beliefs converge to certainty confirming initial relative priors: extreme polarization occurs asymptotically.

**Lemma 4.** *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.<sup>13</sup>

**Proposition 2.** *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  $P_n \xrightarrow{a.s.} 1$  and  $Q_n \xrightarrow{a.s.} 0$  (vice versa if  $P_0 < Q_0$ ).*

*Proof.* By Lemma 4 the ratio of the odds-ratios,  $O^P/O^Q$  converges to infinity almost surely. This implies that  $P_n \xrightarrow{a.s.} 1$  almost surely. Since  $O^P/O^Q \rightarrow \infty$ , at least one of  $P_n \xrightarrow{a.s.} 1$  or  $Q_n \xrightarrow{a.s.} 0$ . 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

---

<sup>13</sup>It is worth comparing this result with symmetry to Fryer et al. (2015). In their paper, 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 this section, with symmetric parameters, the two polarized states are (1,0) and (0,1), and agents' beliefs will always converge to confirm their priors. In our model, polarization occurs because unidentified observations are frequent because the fundamental states are polarized. When the states are not polarized ((1,1) or (0,0)), then polarization will never happen.

Martingale Convergence Theorem, this problem converges to the truth ( $Q = 1$ ) almost surely.<sup>14</sup> By symmetry, if  $Q_n < \varepsilon$  then  $P_n \xrightarrow{a.s.} 1$ .  $\square$

**Corollary 1** (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  $P_n \xrightarrow{a.s.} 1$  and  $Q_n \xrightarrow{a.s.} 0$ , and the reverse for the other.*

An implication of Proposition 1 and Lemma 4 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  $\theta = 0 = \sigma$ .

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  $\theta$  and  $\sigma$ , disagreement would still matter. Nonetheless, this result is helpful to consider how divergence can occur even when states are asymmetric.<sup>15</sup> 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.

### 2.3.2 Incorrect Beliefs in the Asymmetric Case

Individuals can converge asymptotically to false beliefs for any parameters. For any 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

<sup>14</sup>See Diaconis and Freedman (1986).

<sup>15</sup> Notice, for all parameter values,  $L_1^P(Q)/L_1^Q(P)$  is increasing in  $P$  (this follows immediately from differentiation). Thus, there exists  $P/Q$  such that  $O^P/O^Q$  increases after a  $y = 1$ , and thus  $P/Q$  increases. When the  $p$ 's and  $q$ 's are close, then a low  $P/Q$  is required for  $L_1^P(Q)/L_1^Q(P) > 1$  (for example, we saw in the symmetric case that this is always true). Farther from symmetry requires a higher  $P/Q$ . If beliefs fall in this region, if likelihoods for  $y = 0, 2$  are not so different, if  $y = 1$  are sufficiently likely—then beliefs are likely to diverge as in the symmetric case. Thus, asymptotic disagreement is less likely, and is only likely to occur if the initial priors are very different.

result for the asymptotic behavior of  $(P_n, Q_n)$ . Proposition 3 states conditions on parameters that guarantee when priors exist 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. (The conditions are technical, so the reader may want to skip the equations and instead consider Figure 1.)

**Proposition 3.** *The following hold:*

1. *If  $\mathbb{E} [\Delta \log O_{n+1}^P(q_L)] > 0$  and  $\mathbb{E} [\Delta \log O_{n+1}^Q(p_H)] < 0$ , then there exist  $(P_0, Q_0)$  such that  $P_n \rightarrow 1$  and  $Q_n \rightarrow 0$  with positive probability tending to 1 as  $Q_0 \rightarrow 0$  and  $P_0 \rightarrow 1$*
2. *If  $\mathbb{E} [\Delta \log O_{n+1}^P(q_L)] < 0$ , then  $P \xrightarrow{a.s.} 0$ ; if additionally  $\mathbb{E} [\Delta \log O_{n+1}^Q(p_L)] < 0$  then  $Q \xrightarrow{a.s.} 0$ , whereas if  $\mathbb{E} [\Delta \log O_{n+1}^Q(p_L)] > 0$  then  $Q \xrightarrow{a.s.} 1$ .*
3. *If  $\mathbb{E} [\Delta \log O_{n+1}^Q(p_H)] > 0$ , then  $Q \xrightarrow{a.s.} 1$ ; if additionally  $\mathbb{E} [\Delta \log O_{n+1}^P(q_L)] < 0$  then  $P \xrightarrow{a.s.} 0$ , whereas if  $\mathbb{E} [\Delta \log O_{n+1}^P(q_L)] > 0$  then  $P \xrightarrow{a.s.} 1$ .*

Most importantly, if parameters satisfy condition 1 of Proposition 3 then there exist priors (correctly ordered) 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  $f(p, q) = \Pr((P_n, Q_n) \rightarrow (1, 0) | P_0 = p, Q_0 = q)$  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 Proposition 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).

**Proposition 4.** *The convergence probability  $f(p, q)$  is continuous in priors  $p, q$ .*

**Corollary 2.** *If*

$$\frac{\log \frac{q_H}{q_L}}{\log \frac{p_H}{p_L}} < \frac{\log \frac{1-q_H}{1-q_L}}{\log \frac{1-p_H}{1-p_L}}, \quad (7)$$

*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 Proposition 3 as well as equation (7), then for any (rightly ordered) priors, beliefs will diverge with positive probability. By simple counting of parameters and conditions, we have the following result.

**Proposition 5.** *There is a robust 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. Figure 1 plots the sets of  $(p_H, p_L)$  parameters that satisfy the conditions, with each frame holding fixed a different  $(q_H, q_L)$ . Yellow sections indicate that any priors will converge to (1,0) with positive probability. Green sections indicate that some priors are guaranteed to converge to (1,0) with positive probability. The blue sections are not guaranteed to never converge to (1,0), but the regions do not satisfy the conditions guaranteed that they might. The figures indicate that divergence is most likely to occur in a neighborhood of parameters around symmetry. We investigate this further in the next section using simulations.

## 2.4 Simulation Results


Simulations evidence is useful for characterizing the limiting behavior of this process to see precisely how the probability of divergence depends on underlying parameters and prior beliefs. In light of our theoretical results, we consider simulations with  $\theta = 0$  and  $\sigma = 1$  (the order does not matter), and consider two sets of parameter robustness.<sup>16</sup> We first set  $q_H = .65, q_L = .4, p_L = .4,$

---


<sup>16</sup>Remember that as with standard Bayesian learning, beliefs converge to 1 or 0 with sufficient observations; however, the belief need not converge to the *true* value. As proved, when  $\theta = \sigma = 1$  so that values of  $y = 2$  are very likely (or when  $\theta = \sigma = 0$  so that values of  $y = 0$  are very likely), then beliefs converge to the truth, regardless of initial priors. However, convergence is rarely the case when  $\theta \neq \sigma$ .




(a)  $q_H = .6, q_L = .4$ .



(b)  $q_H = .6, q_L = .2$ .



(c)  $q_H = .8, q_L = .4$ .



(d)  $q_H = .8, q_L = .2$ .

Figure 1: Parameters satisfying Proposition 3 and equation (7). Yellow: all priors will converge to  $(1,0)$  with positive probability. Green: some priors will converge to  $(1,0)$  with positive probability.

and we do parameter sensitivity varying  $p_H \in [.5, .9]$ . We set prior beliefs to  $P_0 = .8$  and consider  $Q_0 = .4, .6, .75$ . For each set of parameters we solve one million simulations, running simulations until beliefs converge to zero or one, and calculate the fraction of simulations converging to one (the wrong value). Figure 2 plots the frequency of simulations that converge to the incorrect value. The simulation evidence suggests that for a broad range of parameters  $P \rightarrow 1$  with high probability. The closer to symmetric are the states (the  $q$ 's and  $p$ 's are more similar), the more likely do beliefs converge to the wrong value. This occurs whenever  $P_0 > Q_0$ , regardless of the level. However, the more different are  $P_0$  and  $Q_0$ , the greater probability of converging to  $P = 1$  (the curve is a wider

bell).




Figure 2: Probability of converging to wrong beliefs when starting with  $P_0 > Q_0$ .

If the states are not so different and  $P_0 > Q_0$  (the probability that  $\theta = 1$  is greater than the probability that  $\sigma = 1$ ), then the simulations suggest that asymptotically  $P \rightarrow 1$  and  $Q \rightarrow 0$  with high probability. This effect induces asymptotic disagreement when one agent has prior  $P_0 > Q_0$  and the other has prior with  $P_0 < Q_0$ , and the disagreement is greater with more observations. When the parameters differ ( $p$ 's and  $q$ 's), the frequency of convergence decreases (continuously) as they differ by more. And thus the probability of permanent belief divergence decreases.<sup>17</sup>

When the states are far from symmetric, the qualitative result stands (there is a robust region with bell-shaped positive probability of converging to the wrong beliefs) but there need not be a set of parameters where asymptotic disagreement occurs with probability one. To further investigate

<sup>17</sup>Compare these results to Fryer et al. (2015). In their model, the probability of polarization is an increasing function of the probability of ambiguous signals. As our simulation results suggest, in our model the probability of polarization 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 polarization, whereas in their model the severity of ambiguity determines the probability of polarization.

this possibility, we run simulations varying  $p_H$  with  $q_H = .4, q_L = .2$  but this time with  $p_L = 1 - p_H$ , again considering  $Q_0 = .4, .6, .75$ . Figure 3 plots these results for when the states are quite asymmetric. As one expects, the probability of converging to the wrong belief is greater when the initial belief  $Q_0$  is lower.




Figure 3: Probability of converging to wrong beliefs when starting with  $P_0 > Q_0$ .  $q_H = .4, q_L = .2, p_L = 1 - p_H$ .

Figure 4 plots the convergence probability over a grid of  $(p_H, p_L)$ , holding fixed priors and parameters  $q_H = .65, q_L = .4$ . (Remember that the convergence probability is continuous in priors.) Not surprisingly given the previous figures, the convergence probability looks “multivariate Normal” as a function of the parameters.

While the values of the aggregate states are complements, posteriors need not converge to complements (i.e., they could both converge to 1 or to 0). Of course, this is not all that surprising since posteriors are not converging to the truth. The frequency of these convergences—of  $P_n, Q_n \xrightarrow{\text{a.s.}} 1, 1$  and  $P_n, Q_n \xrightarrow{\text{a.s.}} 0, 0$  are plotted in Figure 5 in the appendix.




Figure 4: Probability of converging to wrong beliefs given prior  $(P_0, Q_0) = (.8, .6)$ , with  $q_H = .65, q_L = .4$ , varying  $p_H, p_L$ .

### 3 A General Model

In this section we generalize the previous analysis to an  $N$ -dimensional discrete multivariate setting. The main results (Proposition 6 and Corollary 3) shows that wrong beliefs can be locally attractive even if agents don't necessarily combine beliefs over marginals by multiplying them (assuming independence), but as long as the aggregation rule looks enough like multiplication in that the ratio of beliefs is bounded.

#### 3.1 The Setup

Let there be a multivariate state of the world  $\theta = (\theta_1, \dots, \theta_d)$  taking values in a discrete set  $\Theta = \prod_{i=1}^d \Theta_i$ . There are also random variables  $X_1, \dots, X_d$  which are distributed as  $X_i \sim F_{\theta_i}$  for  $i = 1, \dots, d$ . In every period,  $(X_1, \dots, X_d)$  are independently drawn from their respective distributions, independent of observations in all other periods. An agent observes  $\Phi(X_1, \dots, X_d)$  for a possibly

non-injective function  $\Phi$ . For now, we consider  $\Phi(X_1, \dots, X_d) = \sum_{i=1}^d X_i$  but the results generalize. Let  $X_{in}$  denote the particular realization of  $X_i$  in period  $n$ , and  $y_n = \Phi(X_{1n}, \dots, X_{dn})$  be the period  $n$  observation.

As before 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  $\Theta$ . Specifically, let  $W_n^i(t_i)$  denote the agent's belief that  $\theta_i = t_i$ ,  $W_n^i$  the resulting measure over  $\Theta_i$ , and

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

where for every  $A \subset \Theta_i$ ,  $W_n(\theta_i \in A) = \sum_{\theta_i \in A} W_n(\theta_i = t_i)$ , since  $W_n$  is a measure. For convenience, we will use  $W_n^i(\cdot) = W_n(\theta_i = \cdot)$  interchangeably. This implies

$$\log \frac{W_{n+1}(\theta_i = t_i)}{1 - W_{n+1}(\theta_i = t_i)} = \log \frac{W_n(y_{n+1} | \theta_i = t_i)}{W_n(y_{n+1} | \theta_i \neq t_i)} + \log \frac{W_n(\theta_i = t_i)}{1 - W_n(\theta_i = t_i)}. \quad (8)$$

The agent views the events  $\theta_i = t_i$  as independent with distribution  $W_n(\theta_i = t_i)$ , in which case  $\Psi(W_n^1, \dots, W_n^d)(t_1, \dots, t_d) = \prod_{i=1}^d W_n^i(t_i)$ . Two remaining calculations remain, which are

$$W_n(y_{n+1} | \theta_i = t_i) = \frac{\sum_{\vartheta \in \Theta: \vartheta_i = t_i} W_n(\theta = \vartheta) f_\vartheta(y_{n+1})}{\sum_{\vartheta \in \Theta: \vartheta_i = t_i} W_n(\theta = \vartheta)} \quad (9)$$

$$W_n(y_{n+1} | \theta_i \neq t_i) = \frac{\sum_{\vartheta \in \Theta: \vartheta_i \neq t_i} W_n(\theta = \vartheta) f_\vartheta(y_{n+1})}{\sum_{\vartheta \in \Theta: \vartheta_i \neq t_i} W_n(\theta = \vartheta)} \quad (10)$$

We impose  $\sum_{\vartheta \in \Theta: \vartheta_i = t_i} W_n(\theta = \vartheta) = W_n(\theta_i = t_i)$  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  $n$ .

## 3.2 Results

To generalize earlier results, we are interested in situations in which a state  $t \in \Theta$  may be attracting equilibrium, in that as prior beliefs  $W_0(\theta_i = t'_i)$  approach  $\delta_{t_i}$  for every  $i$ ,  $\Pr(W_n(\theta_i = t_i) \rightarrow 1)$  approaches 1. We measure by the distance between  $W_0^i$  and  $\delta_{t_i}$  by the scaled total variation distance  $\|W_0^i - \delta_{t_i}\| = 1 - W_0(\theta_i = t_i)$ . We also make the following assumption on  $\Psi$  which is that the agent regards the  $\theta_i$  as roughly independent events conditional upon observed signals:

$$0 < \inf_{\substack{x \in (0,1]^d \\ V^i \in \Delta(\Theta_i), \\ 1 \leq i \leq d}} \frac{\Psi(V^1, \dots, V^d)(t_1, \dots, t_d)}{\prod_{i=1}^d V^i(t_i)} \leq \sup_{\substack{x \in (0,1]^d \\ V^i \in \Delta(\Theta_i), \\ 1 \leq i \leq d}} \frac{\Psi(V^1, \dots, V^d)(t_1, \dots, t_d)}{\prod_{i=1}^d V^i(t_i)} < \infty. \quad (11)$$

We require a boundedness condition for densities, which is that

$$\sup_{y, \theta, \theta'} \frac{f_\theta(y)}{f_{\theta'}(y)} < \infty. \quad (12)$$

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

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

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

Proposition 6 gives a sufficient condition for  $t$  to be attracting in the sense previously given.

**Proposition 6.** *Let  $t \in \Theta$ , and let  $\theta_0$  be the true state of the world. Suppose that for all  $i$ , (11) and (12) hold as well as*

$$\text{for all } f \in \text{co}(\{f_\theta : \theta \in G_i(t)\}), D_{KL}(f_{\theta_0} \| f) > D_{KL}(f_{\theta_0} \| f_t). \quad (13)$$

*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  $\Pr(W_n(\theta_i = t_i) \rightarrow 1) \geq 1 - \varepsilon$ .*

Essentially, the proposition shows that under some conditions  $t$  can be an attracting state if it is

a local maximizer of  $D_{\text{KL}}(f_{\theta_0} || \cdot)$ .  $KL$  divergence measures how far apart probability distributions  $p$  and  $q$  are by taking  $D_{\text{KL}}(p||q) = \int \log(p/q)p$ . The following corollary is a straightforward consequence of Gibb’s inequality:

**Corollary 3.** *If (11) and (12) hold and  $f_{\theta_0} \notin \text{co}(\{f_{\theta} : \theta \in G_i(\theta_0)\})$  for all  $i$ , then the belief  $\theta = \theta_0$  is locally attracting in the sense of Proposition 6.*

The condition of this proposition says that if  $p$  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.

## 4 Implications

Our analysis has several implications for the effectiveness of attempts to relieve identification problems, divergence and information disclosure, and model specification and dynamics.

### 4.1 Actions

First, incomplete attempts to relieve the identification problem need not completely cure it. Suppose agents can choose with rare frequency to take an action leading to some realization: action  $a$  has a greater success rate when  $\theta = 1$  and action  $b$  has a greater success rate when  $\sigma = 1$ . Since agents’ beliefs are unlikely to converge, agents are likely to disagree about the best action to take. If agents could choose actions frequently enough, then their “experimentation” could provide richer data, perhaps even relieving the identification problem. However, this is precisely the problem that we find ourselves in the real world: we could know the effect of policies on macro variables if we could experiment frequently enough, but of course we cannot. Given unidentified models, it is no surprise that economists and politicians will continue to make contrasting recommendations. In some cases it is not surprising that recommendations will become increasingly polarized with more observations.

In fact, even frequent experimentation would still likely lead to polarization. First, if agents first observe many “unidentified observations,” then beliefs will already be polarized and so identified observations will have small effects changing beliefs. Second, if the frequency of identified observations is exceeded by the frequency of unidentified observations (again, this is the scenario of concern), then polarization should still occur in the long run. An implication 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 is unidentified.<sup>18</sup>

## 4.2 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 polarization. This may explain why 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.

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.<sup>19</sup> 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

---

<sup>18</sup>We acknowledge Bruce Sacerdote for this observation.

<sup>19</sup>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.

accommodate information that was initially at odds.

### 4.3 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).

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. First, if agents instead updated beliefs in the classical sense of Doob, “suspending judgment” until all the information has arrived, or completely revisiting their beliefs in light of all information, then with probability one polarization would not occur. (When agents form  $P_{n+1}$  using  $P_0$  and  $n + 1$  observations, then  $P_{n+1}$  converges to the truth for all but a measure zero of parameters—for example, the perfectly symmetric case.) 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 of their most recent 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.

## 5 Conclusion

The world is multi-dimensional—there are a number of factors that contribute to what we see—but the data we see is nearly always less dimensional than the world. We live with identification problems. As a result, Bayesian updating of marginal probabilities need not converge to the truth, and so 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 more frequent, then asymptotically people’s initial beliefs get reinforced: whatever state they initially believed was more likely to be good will with positive probability be believed to be the good state with certainty. 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.
- 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.

- 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-VILLAYERDE, 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.
- GEANAKOPOLOS, J. D. AND H. M. POLEMARCHAKIS (1982): “We can’t disagree forever,” *Journal of Economic Theory*, 28, 192–200.
- 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.
- PEW RESEARCH CENTER (2014): “Stark Racial Divisions in Reactions to Ferguson Police Shooting,” .
- 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.
- 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 2.1.* The posterior  $\hat{P} > P$  if and only if  $Q < \frac{1}{2}$ , and the posterior  $\hat{Q} > Q$  if and only if  $P < \frac{1}{2}$ . From the posteriors,  $\hat{P} > P$  if  $(1 - Q) > P(1 - Q) + (1 - P)Q$ , which is true if  $1 - Q > Q$ . Similarly,  $\hat{Q} > Q$  if  $(1 - P) > P(1 - Q) + (1 - P)Q$ , which is true if  $1 - P > P$ .  $\square$

*Proof of Lemma 3.* Consider the likelihood ratios for  $P$  (the results are symmetric for  $Q$ ). Then the likelihood ratios are given by

$$L_2^P = \frac{p_H}{p_L}, \quad L_0^P = \frac{1 - p_H}{1 - p_L}, \quad L_1^P(Q) = \frac{p_H(1 - \bar{q}) + (1 - p_H)\bar{q}}{p_L(1 - \bar{q}) + (1 - p_L)\bar{q}}. \quad (14)$$

Thus, the likelihood from observing  $y = 0$  or  $y = 2$  is independent of  $Q$ , but the likelihood from observing  $y = 1$  depends on  $Q$ . Thus, the odds ratios after observing  $(0, 1)$  and  $(1, 0)$  are given by

$$O_2^P(0, 1) = L_0^P L_1^P(Q_1) O_0^P, \quad \text{and} \quad O_2^P(1, 0) = L_0^P L_1^P(Q_0) O_0^P, \quad (15)$$

where in the first case the likelihood ratio for the  $y = 1$  observation depends on the updated belief about  $Q$  after the  $y = 0$  observation, while in the second case the likelihood ratio for the  $y = 1$  observation depends on the initial prior belief about  $Q$ , which has not been updated.  $\square$

*Proof of Lemma 4.* First, since  $p_H = q_H$  and  $p_L = q_L$ , the likelihood ratios given observations  $y = 0$  and  $y = 2$  are equal:

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

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.  $\square$

*Proof of Proposition 1.* We prove for  $P_n$  (the proof for  $Q_n$  is similar). If  $\theta = \sigma = 1$ , then the inverse-odds ratio  $(O_n^P)^{-1} = \frac{1 - P_n}{P_n}$  is upper bounded on  $\Omega$  by the random variable  $\prod_{i=1}^n r_i$  defined such that

$$r_i(\omega) = \begin{cases} \frac{p_L}{p_H} & \text{if } y_i(\omega) = 2 \\ \frac{1 - p_L}{1 - p_H} & \text{if } y_i(\omega) = 0 \\ \frac{p_L(1 - q_H) + (1 - p_L)q_H}{p_H(1 - q_H) + (1 - p_H)q_H} & \text{if } y_i(\omega) = 1 \end{cases}$$

Note that in this case

$$\mathbb{E} \left[ \prod_{i=1}^{n+1} r_i \mid r_n, r_{n-1}, \dots \right] = \mathbb{E} [r_{n+1} \mid r_n, r_{n-1}, \dots] \prod_{i=1}^n r_i = \prod_{i=1}^n r_i. \quad (16)$$

But by the Hewitt-Savage 0-1 Law,  $\liminf_{n \rightarrow \infty} \prod_{i=1}^n r_i$  is a constant  $c \in [0, \infty)$ , and it is clear that such a constant is invariant under multiplication by  $\frac{p_L}{p_H}, \frac{1-p_L}{1-p_H}$ . Hence  $c \in \{0, \infty\}$ , but as (16) and the law of iterated expectations imply  $\mathbb{E} [\prod_{i=1}^n r_i]$  is bounded, it must be true that  $\liminf_{n \rightarrow \infty} \prod_{i=1}^n r_i = 0$ . Hence,  $\liminf_{n \rightarrow \infty} (O_n^P)^{-1} = 0$ , and thus  $\limsup_{n \rightarrow \infty} P_n = 1$  (all equalities hold a.e.). Then (20) implies that  $P_n \xrightarrow{\text{a.s.}} 1$ .  $\square$

Suppose for now that  $\theta = \sigma = 1$  and  $p_H, q_H > 0$ . Consider the following recursion relations for  $(P_n)$  as a function of  $Q_n$ :

$$\begin{aligned} P_{n+1}(2) &= \frac{P_n p_H}{P_n p_H + (1 - P_n) p_L} \\ P_{n+1}(0) &= \frac{P_n (1 - p_H)}{P_n (1 - p_H) + (1 - p) (1 - p_L)} \\ P_{n+1}(1) &= \frac{P_n (p_H (1 - \bar{q}_n) + (1 - p_H) \bar{q}_n)}{P_n (p_H (1 - \bar{q}_n) + (1 - p_H) \bar{q}_n) + (1 - P_n) (p_L (1 - \bar{q}_n) + (1 - p_L) \bar{q}_n)}, \end{aligned}$$

where  $\bar{q}_n = Q_n q_H + (1 - Q_n) q_L \in [q_L, q_H]$ . One can determine that:

$$\begin{aligned} &\bar{q}_n (P_n p_H + (1 - P_n) p_L) p_{n+1}(2) \\ &+ (1 - \bar{q}_n) (P_n (1 - p_H) + (1 - p) (1 - p_L)) p_{n+1}(0) \\ &+ \left( P_n (p_H (1 - \bar{q}_n) + (1 - p_H) \bar{q}_n) + (1 - P_n) (p_L (1 - \bar{q}_n) + (1 - p_L) \bar{q}_n) \right) p_{n+1}(1) \\ &= P_n, \end{aligned} \quad (17)$$

i.e. the agent believes that his estimate of  $P_n$  will be correct in expectation in the next period. In addition, the inequality

$$P_{n+1}(2) \geq P_{n+1}(1) \geq P_{n+1}(0) \quad (18)$$

naturally holds because  $\frac{p_H}{1-p_H} > \frac{p_L}{1-p_L}$ . Note that in comparison to (17), one has

$$\begin{aligned} q_H p_H &\geq \bar{q}_n (P_n p_H + (1 - P_n) p_L) \\ (1 - q_H)(1 - p_H) &\leq (1 - \bar{q}_n) (P_n(1 - p_H) + (1 - p)(1 - p_L)). \end{aligned}$$

Hence, the imposition of (18) implies

$$\begin{aligned} \mathbb{E}_n(P_{n+1}) &= q_H p_H P_{n+1}(2) + (1 - q_H)(1 - p_H) P_{n+1}(0) \\ &\quad + ((1 - q_H)p_H + q_H(1 - p_H)) P_{n+1}(1) \\ &\geq P_n, \end{aligned} \tag{19}$$

where  $\mathbb{E}_n(\cdot)$  is the time  $n$  expectation given  $P_n, p_{n-1}, \dots$ . So  $(P_n)$  is a supermartingale, and by Doob's Martingale Convergence Theorem it is true that

$$(P_n) \xrightarrow{\text{a.s.}} \limsup_{n \in \mathbb{N}} P_n : [0, 1] \rightarrow [0, 1]. \tag{20}$$

In the following lemma, let  $\mathcal{F}_n$  be the filtration of  $\mathcal{F}_\infty$  which contain (as cylinder sets of  $\mathcal{F}_\infty$ ) sets in the probability space  $\Omega$  on which  $(P_n, Q_n)$  are defined. For  $\omega \in \Omega$  write  $\omega = \omega_1 \omega_2 \dots$ , where  $\omega_n \in \{0, 1, 2\}$  according to the value of  $y_n(\omega)$ .

**Lemma 5.** *If  $\lim_{n \rightarrow \infty} P_n$  exists in  $[0, 1]$  almost everywhere and  $p_H, p_L \in (0, 1)$ , then*

$$Pr\left(\omega : \lim_{n \rightarrow \infty} P_n(\omega) \in (0, 1)\right) = 0. \tag{21}$$

*Proof.* We claim that any point  $\omega$  such that  $\lim_{n \rightarrow \infty} P_n(\omega) \in (0, 1)$  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  $\lim_{n \rightarrow \infty} p_n(\omega) = c \in (0, 1)$ . Then taking such  $m$  arbitrarily high,  $(P_{n_m+1}) \rightarrow \frac{c p_H}{c p_H + (1-c) p_L} > c$ , a contradiction. The proof when  $\omega_n = 0$  infinitely often is similar. (21) follows immediately by the Second Borel-Cantelli Lemma.  $\square$

### Proof of Proposition 3

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

$$\log O_{n+1}^P(2) = \log \frac{p_H}{p_L} + \log O_n^P \quad (22)$$

$$\log O_{n+1}^P(0) = \log \frac{1-p_H}{1-p_L} + \log O_n^P \quad (23)$$

$$\log O_{n+1}^P(1) = \log \frac{p_H(1-\bar{q}_n) + (1-p_H)\bar{q}_n}{p_L(1-\bar{q}_n) + (1-p_L)\bar{q}_n} + \log O_n^P. \quad (24)$$

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

$$\begin{aligned} \mathbb{E} [\Delta \log O_{n+1}^P(\bar{q})] &= p_L q_H \log \frac{p_H}{p_L} + (1-p_L)(1-q_H) \log \frac{1-p_H}{1-p_L} \\ &\quad + (p_L(1-q_H) + (1-p_L)q_H) \log \frac{p_H(1-\bar{q}) + (1-p_H)\bar{q}}{p_L(1-\bar{q}) + (1-p_L)\bar{q}}. \end{aligned}$$

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 6.** *Let  $\theta = 0$  and  $\sigma = 1$ . If  $\mathbb{E} [\Delta \log O_{n+1}^P(q_L)] > 0$  and  $\mathbb{E} [\Delta \log O_{n+1}^Q(p_H)] < 0$ , then there exist  $(P_0, Q_0)$  such that  $P_0 \rightarrow 1$  and  $Q_0 \rightarrow 0$  with positive probability tending to 1 as  $Q_0 \rightarrow 0$  and  $P_0 \rightarrow 1$ .*

*Proof.* First, recall Hoeffding's inequality: let  $X_1, \dots, X_n$  be independent and  $a_i \leq X_i \leq b_i$  for all  $i$ .

Then

$$\Pr \left( \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \geq n\varepsilon \right) \leq \exp \left( \frac{-2n^2 \varepsilon^2}{\sum_{i=1}^n (b_i - a_i)} \right). \quad (25)$$

By continuity, there is a  $\delta_1 > 0$  such that if  $\bar{q} < q_L + \delta_1$ , then  $\mathbb{E} [\Delta \log O_{n+1}^P(\bar{q})] > \varepsilon_1 > 0$ . Consider the process  $(X_n)$  which has  $\Delta(X_n) = \Delta(\log O_n^P)$  when  $\omega_n = 0, 2$ , and  $\Delta(X_n) = \log \frac{p_H(1-q_L-\delta_1) + (1-p_H)(q_L+\delta_1)}{p_L(1-q_L-\delta_1) + (1-p_L)(q_L+\delta_1)}$  when  $\omega_n = 1$ . Thus,  $\mathbb{E} [\Delta(X_n)] > \varepsilon_1$ . Also, if  $\bar{q}_n \in [q_L, q_L +$

$\delta_1]$  for all values of  $n$ , we have the relation  $\log O_n^P \geq X_n$ . Set  $Q^*$  to satisfy  $Q^*q_H + (1 - Q^*)q_L = q_L + \delta_1$ . By the analogous assumption for  $Q_n$ , there exists some  $\delta_2 > 0$  and  $\varepsilon_2 > 0$  such that  $\mathbb{E} \left[ \Delta \log O_{n+1}^Q(\bar{p}) \right] < -\varepsilon_2$  if  $\bar{p} > p_H - \delta_2$ . Construct the random walk  $(Y_n)$  as an upper bound for  $\log O_n^Q$  just as  $X_n$  was constructed as a lower bound for  $\log O_n^P$ , with  $\bar{P}$  replaced with  $p_H - \delta_2$ . Set  $P^*$  to satisfy  $P^*p_H + (1 - P^*)p_L = p_H - \delta_2$ . Also, let  $X^* = \log \frac{P^*}{1-P^*}$  and  $Y^* = \log \frac{Q^*}{1-Q^*}$ . Note that if  $X_n \geq X^*$  and  $Y_n \leq Y^*$  for all  $n$ , then  $P_n \geq P^*$  and  $Q_n \leq Q^*$  for all  $n$  and consequentially  $X_n$  is always a lower bound for  $P_n$  whereas  $Y_n$  is always an upper bound for  $Q_n$ . Fix  $X_0 = \log O_0^P$  and  $Y_0 = \log O_0^Q$ . By (25), for  $\nu > 0$  we have

$$\begin{aligned} \Pr(X_n - X_0 \leq n\mathbb{E}[X_i] - n\nu) &= \Pr\left(\sum_{i=1}^n (\Delta X_i - \mathbb{E}[\Delta X_i]) \leq -n\nu\right) \\ &\leq \exp\left(\frac{-2n^2\nu^2}{\sum_{i=1}^n \left(\log \frac{p_H}{p_L} - \log \frac{1-p_H}{1-p_L}\right)}\right). \end{aligned}$$

Pick  $X_0 > X^*$ . Of course,  $\Delta X_i > \varepsilon_1$ , so picking  $\nu = \varepsilon_1/2$ , we have

$$\Pr\left(X_n \leq X^* + n\frac{\varepsilon_1}{2}\right) \leq \Pr\left(X_n \leq X_0 + n\frac{\varepsilon_1}{2}\right) \leq \exp\left(\frac{-\varepsilon_1^2 n\nu^2}{2\left(\log \frac{p_H}{p_L} - \log \frac{1-p_H}{1-p_L}\right)}\right),$$

and if  $n \leq \left\lfloor \frac{X_0 - X^*}{-\log \frac{1-p_H}{1-p_L} + \varepsilon_1/2} \right\rfloor$ ,  $\Pr(X_n \leq X^*) = 0$ . Hence,

$$\Pr\left(X_n \leq X^* + \frac{n\varepsilon_1}{2} \text{ for some } n\right) \leq \sum_{\left\lfloor \frac{X_0 - X^*}{-\log \frac{1-p_H}{1-p_L} + \varepsilon_1/2} \right\rfloor}^{\infty} \exp\left(\frac{-\varepsilon_1^2 n\nu^2}{2\left(\log \frac{p_H}{p_L} - \log \frac{1-p_H}{1-p_L}\right)}\right).$$

Obviously,  $X_0 = \log O_0^P$  can be chosen high enough so that this quantity is strictly less than  $\frac{1}{2}$ .

Similarly, we can choose  $Y_0 = \log O_0^Q$  so low such that  $\Pr(Y_n \geq Y^* - \frac{n\epsilon_2}{2} \text{ for some } n) < \frac{1}{2}$ . Then

$$\begin{aligned} \Pr(P_n \rightarrow 1 \text{ and } Q_n \rightarrow 0) &\geq \Pr\left(X_n \geq X^* + \frac{n\epsilon_1}{2} \text{ and } Y_n \leq Y^* - \frac{n\epsilon_2}{2} \text{ for all } n\right) \\ &\geq 1 - \Pr\left(X_n \leq X^* + \frac{n\epsilon_1}{2} \text{ for some } n\right) - \Pr\left(Y_n \geq Y^* - \frac{n\epsilon_2}{2} \text{ for some } n\right) \\ &> 0. \end{aligned}$$

Indeed, it is clear that as  $Q_0 \rightarrow 0$  and  $P_0 \rightarrow 1$  for a fixed  $P^*$  and  $Q^*$ , this probability tends to 1, despite the fact that  $\theta = 0$  and  $\sigma = 1$ . □

**Lemma 7.** *If  $\mathbb{E}[\Delta \log O_{n+1}^P(q_L)] < 0$ , then  $P \xrightarrow{\text{a.s.}} 0$ , and if  $\mathbb{E}[\Delta \log O_{n+1}^P(p_H)] > 0$ , then  $Q \xrightarrow{\text{a.s.}} 1$*

*Proof.* The proof follows the method of Lemma 6 by setting  $Q^* = 1$  and  $P^* = 0$ . □

**Lemma 8.** *Let  $\mathbb{E}[\Delta \log O_{n+1}^P(q_L)] < 0$ . Then  $P \xrightarrow{\text{a.s.}} 0$ , and if  $\mathbb{E}[\Delta \log O_{n+1}^Q(p_L)] < 0$  then  $Q \xrightarrow{\text{a.s.}} 0$ , whereas if  $\mathbb{E}[\Delta \log O_{n+1}^Q(p_L)] > 0$  then  $Q \xrightarrow{\text{a.s.}} 1$ .*

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

*Proof.* The convergence  $P \xrightarrow{\text{a.s.}} 0$  follows from Lemma 7. As for convergence of  $Q$ , if  $\mathbb{E}[\Delta \log O_{n+1}^P(p_L)] < 0$ , then  $\mathbb{E}[\Delta \log O_{n+1}^P(\bar{p})]$  is upper-bounded by a strictly negative number everywhere, and a law of large numbers argument suffices. On the other hand, if  $\mathbb{E}[\Delta \log O_{n+1}^P(p_L)] > 0$ , then as  $P$  concentrates almost surely at 0,  $\mathbb{E}[\Delta \log O_{n+1}^Q]$  becomes lower bounded (for all sufficiently high  $n$ ) by a positive  $\epsilon \in (0, \mathbb{E}[\Delta \log O_{n+1}^P(p_L)])$  almost surely, whence  $Q \xrightarrow{\text{a.s.}} 1$ . An obvious symmetry establishes the second claim. □

## Proof of Proposition 4

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,

as in (Schenk-Hoppé and Schmalfuß, 2001). Our result of continuity also holds on a larger set of random dynamical systems with a Bernoulli shift as a random component.

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,

$$f(p, q) = p_L q_H f(P'((p, q), 2), Q'((p, q), 2)) + (1 - p_L)(1 - q_H) (P'((p, q), 0), Q'((p, q), 0)) \\ + \left( p_L(1 - q_H) + (1 - p_L)q_H \right) f(P'((p, q), 1), Q'((p, q), 1)).$$

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

**Lemma 9.** *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

$$P'_n(\omega) \geq P_n(\omega), \quad Q'_n(\omega) \leq Q_n(\omega). \quad (26)$$

The argument is accomplished with the observation that

1. If  $y_n(\omega) = 0, 2$ , the ordering is preserved by monotonicity of the relevant functions.
2. 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  $O_{n+1}^P$  increases by more when  $y_n = 1$  and  $Q_n$  is smaller. Conversely,  $O_{n+1}^Q$  increases by less when  $y_n = 1$  and  $P_n$  is larger. Because  $O_{n+1}^P \geq O_{n+1}^{P'}$  if and only if  $P_{n+1} \geq P'_{n+1}$ , the inequalities in (26) are indeed preserved.

□

*Proof of Corollary 2.* Note that

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

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

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  $\ell$  becomes arbitrarily large,  $\log O_{\ell(n+m)}^P$  becomes arbitrarily large and  $\log O_{\ell(n+m)}^Q$  becomes arbitrarily small from any initial prior  $(P, Q)$ . In particular, if  $f$  does not vanish for all  $(p, q)$ , by Lemma 9, 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  $O^P \rightarrow \infty$ . Because this critical threshold can be reached in a finite number of steps from any prior, any prior has a positive probability of  $O^P \rightarrow \infty$ . The second claim is proved similarly by exchanging ‘ $P$ ’ and ‘ $Q$ ’.  $\square$

**Lemma 10.** 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  $\sum_{k=0}^\infty p_k = 1$ ,

$$\left| D_{KL} \left( f_0 \left\| \sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^\infty p_k \right. \right) - D_{KL} \left( f_0 \left\| \sum_{k=0}^\infty p_k f_k \right. \right) \right| \quad (27)$$

$$\leq \max \left\{ \log \frac{\beta_1}{\beta_1 - (\beta_2 - \beta_1) \sum_{k=n_0}^\infty p_k}, -\log \frac{\beta_1}{\beta_1 + (\beta_2 - \beta_1) \sum_{k=n_0}^\infty p_k} \right\}. \quad (28)$$

*Proof.* This follows from the calculation:

$$\begin{aligned} & \left| \int \left( \log \frac{f_0}{\sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^{\infty} p_k} - \log \frac{f_0}{\sum_{k=0}^{\infty} p_k f_k} \right) f_0 \right| dy \\ & \leq \int \left| \log \frac{\sum_{k=0}^{\infty} p_k f_k}{\sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^{\infty} p_k} \right| f_0 dy. \end{aligned} \quad (29)$$

where

$$\begin{aligned} \log \frac{\sum_{k=0}^{\infty} p_k f_k}{\sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^{\infty} p_k} & \leq \log \frac{\sum_{k=0}^{\infty} p_k f_k}{\sum_{k=0}^{\infty} p_k f_k - (\beta_2 - \beta_1) \sum_{k=n_0}^{\infty} p_k f_0} \\ & \leq \log \frac{\beta_1}{\beta_1 - (\beta_2 - \beta_1) \sum_{k=n_0}^{\infty} p_k} \end{aligned} \quad (30)$$

and similarly

$$\log \frac{\sum_{k=0}^{\infty} p_k f_k}{\sum_{k=0}^{n_0-1} p_k f_k + f_1 \sum_{k=n_0}^{\infty} p_k} \geq \log \frac{\beta_1}{\beta_1 + (\beta_2 - \beta_1) \sum_{k=n_0}^{\infty} p_k} \quad (31)$$

Since  $\int f_0 dy = 1$ , (29) is bounded in magnitude by the larger of the magnitudes of (30) and (31).  $\square$

*Proof of Proposition 6.* Denote lower and upper limits in (11) by  $\alpha_1$  and  $\alpha_2$  respectively, and let the left side of (12) be  $\beta$ . For any fixed  $i$ , we claim that there is a  $\delta \in (0, 1)$  such that if

$$W_n^j(t_j) \geq \delta \text{ for all } j \quad (32)$$

then  $\mathbb{E} \left[ \log \frac{W_n(y_{n+1} | \theta_i = t_i)}{W_n(y_{n+1} | \theta_i \neq t_i)} \right] > 0$ . First, note that (13) implies

$$L \equiv \inf_{f \in \text{co}(\{f_{\theta} : \theta \in G_i(t)\})} D_{\text{KL}}(f_{\theta_0} || f) > D_{\text{KL}}(f_{\theta_0} || f_t). \quad (33)$$

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

and corresponding densities  $f^m \rightarrow f \in \text{co}(\{f_\theta : \theta \in G_i(t)\})$ . But by (11),

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

which allows use of the DCT to conclude that (13) is met with equality at  $f$ , a contradiction.

Now by (9),  $W_n(y_{n+1} | \theta_i = t_i)$  is a weighted average of  $f_\vartheta(y_{n+1})$  such that  $\vartheta_i = t_i$ ; furthermore, by (11) one has  $W_n(\theta = t) \geq \alpha_1 \prod_{j=1}^d W_n(t_j)$  and:

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

From (10),  $W_n(y_{n+1} | \theta_i \neq t_i)$  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 (11),  $W_n(\theta_i = s, \theta_{-i} = t_{-i}) \geq \alpha_1 W_n^i(s) \prod_{j \neq i} W^j(t_j)$  and

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

Thus, in calculating (9) the agent places little weight on  $\vartheta$  that are not  $t$ , and in calculating (10), the agent places little weight on  $\vartheta$  such that  $\vartheta_{-i} \neq t_{-i}$ . In particular, the weight placed on densities  $f_\vartheta$  with  $\vartheta \neq t$  in (9) 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  $f_\vartheta$  with  $\vartheta_{-i} \neq t_{-i}$  in (10) 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  $\delta$  satisfying  $\frac{\alpha_2 d \delta}{\alpha_1 (1-\delta)^d} < \eta$  such that when  $W_n^j(t_j) \geq 1 - \delta$  for all  $j$  then the agent's beliefs in these ancillary states of the world which differ from  $t$  is bounded by  $\eta$ , conditioning either on  $\theta_i = t_i$  or  $\theta_i \neq t_i$ . Suppose that (32) holds with  $\delta > 0$  meeting this condition.

Now, application of (28) in Lemma 10 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

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

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

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

where  $f \in \text{co}(\{f_\theta : \theta \in G_i(t)\})$ . By (33), if one choose  $\delta$  and hence  $\eta$  to be small enough so that  $K(\beta, \eta) < \frac{L - D_{\text{KL}}(f_{\theta_0} \| f_t)}{3} \equiv C$ , the triangle inequality implies

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

where  $C > 0$ . Thus, by (8), (32) implies,

$$\mathbb{E} \left[ \log \frac{W_n(y_{n+1} | \theta_i = t_i)}{W_n(y_{n+1} | \theta_i \neq t_i)} \right] = \int \left( \log \frac{W_n(y_{n+1} | \theta_i = t_i)}{f_{\theta_0}} - \frac{\log W_n(y_{n+1} | \theta_i \neq t_i)}{f_{\theta_0}} \right) f_{\theta_0} dy > C. \quad (34)$$

Since the number of marginals to be considered is  $d < \infty$ , one can pick  $\delta$  and  $C$  above small enough so that (34) holds for all  $i$ . Let  $\Omega = Y^{\mathbb{Z}_+}$  where  $\Omega$  is equipped with the product  $\sigma$ -algebra and product measure inherited from  $f_{\theta_0}$  (since the observations  $y_n$  are iid  $\sim f_{\theta_0}$ ). For  $1 \leq i \leq d$ ,

define the following stochastic process  $(X_n^i) : \Omega \rightarrow [0, 1]$ :

$$X_0^i = \log \frac{W_0^i(t_i)}{1 - W_0^i(t_i)}$$

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

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

$$\begin{aligned} & \mathbb{P}_{\theta_0} \left( \bigcap_{i=1}^d \left\{ \log \frac{W_1^i(t_i)}{1 - W_1^i(t_i)} \geq \log \frac{1-\delta}{\delta} + \frac{C}{2} \right\} \right) \\ & \geq 1 - \sum_{i=1}^d \mathbb{P}_{\theta_0} \left( \log \frac{W_1^i(t_i)}{1 - W_1^i(t_i)} < \log \frac{W_0^i(t_i)}{1 - W_0^i(t_i)} - \gamma + \frac{C}{2} \right) \\ & \geq 1 - \sum_{i=1}^d \mathbb{P}_{\theta_0} \left( X_1^i - X_0^i < -\gamma - \frac{C}{2} \right) \\ & \geq 1 - d \exp \left( \frac{-(\gamma + C/2)^2}{2 \log \beta} \right). \end{aligned} \tag{35}$$

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

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

Since  $P_{\theta_0}(E_n) = P_{\theta_0}(E_1) \prod_{\ell=1}^{n-1} P_{\theta_0}(E_{\ell+1}|E_\ell)$  with  $P_{\theta_0}(E_1)$  as in (35), we obtain:

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

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,  $\lim_{\gamma \rightarrow \infty} d \sum_{\ell=0}^{\infty} \exp\left(-\frac{(\gamma + (\ell+1)C/2)^2}{2(\ell+1)\log\beta}\right) = 0$ . It follows that  $\gamma$ , and thus  $c$ , can be chosen high enough so that  $P_{\theta_0}(\bigcap_{n=1}^{\infty} E_n) > 1 - \varepsilon$ , which concludes.  $\square$

## B Additional Figures

When states are not symmetric, posteriors can converge to values that are not even “symmetric” with the truth. Figure 5 shows the likelihood that both beliefs converge to zero in this case. The curve shifts up (more likely) when  $P_0$  is more different from  $Q_0$ , down the closer they are, but is not much affected by the levels. This result is not surprising: the  $Q$ -belief will get pushed to zero when there are a sequence of  $y = 0$  which updates  $Q$  down, which is likely to occur if  $P$  is high.

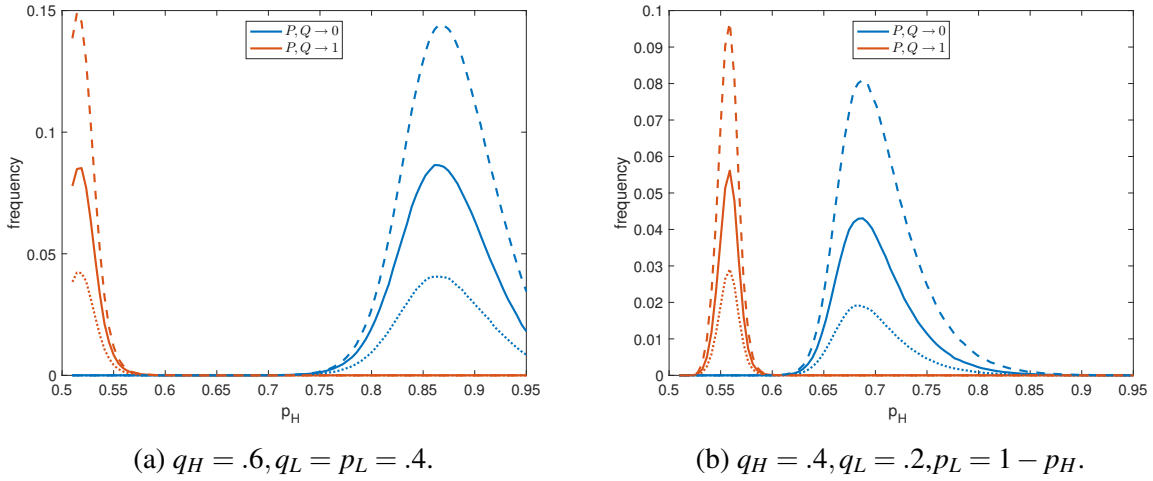


Figure 5: Probability both  $P \rightarrow 0, Q \rightarrow 0$  when starting with  $P_0 > Q_0$ .