

Learning and Disagreement in an Uncertain World*

Daron Acemoglu, Victor Chernozhukov, and Muhamet Yildiz[†]

First version: September, 2006.

This version: July, 2007.

Abstract

Most economic analyses presume that there are limited differences in the prior beliefs of individuals, an assumption most often justified by the argument that sufficient common experiences and observations will eliminate disagreements. We investigate this claim using a simple model of Bayesian learning. Two individuals with different priors observe the same infinite sequence of signals about some underlying parameter. Existing results in the literature establish that when individuals *know* the interpretation of signals, under very mild conditions, there will be asymptotic agreement—their assessments will eventually agree. In contrast, we look at an environment in which individuals are *uncertain* about the interpretation of signals, meaning that they have non-degenerate probability distributions over the conditional distribution of signals given the underlying parameter. When priors on the parameter and the conditional distribution of signals have full support, we show the following: (1) Individuals will never agree, even after observing the same infinite sequence of signals. (2) Before observing the signals, they believe with probability 1 that their posteriors about the underlying parameter will fail to converge. (3) Observing the same (infinite) sequence of signals may lead to a divergence of opinion rather than the typically-presumed convergence. We then characterize the conditions for asymptotic agreement under “approximate certainty”—i.e., as we look at the limit where uncertainty about the interpretation of the signals disappears. When the family of probability distributions of signals given the parameter has “rapidly-varying tails” (such as the normal or the exponential distributions), approximate certainty restores asymptotic agreement. However, when the family of probability distributions has “regularly-varying tails” (such as the Pareto, the log-normal, and the t-distributions), asymptotic agreement does not obtain even in the limit as the amount of uncertainty disappears. We also discuss how lack of common priors implied by the type of learning in this paper interacts with economic behavior in various different situations, including games of common interest, coordination, asset trading and bargaining.

Keywords: asymptotic disagreement, Bayesian learning, merging of opinions.

JEL Classification: C11, C72, D83.

*We thank Eduardo Faingold, Greg Fisher, Bengt Holmstrom, Matthew Jackson, Drew Fudenberg, Alessandro Lizzeri, Giuseppe Moscarini, Marciano Sinscalchi, Robert Wilson and seminar participants at MIT, the University of British Columbia, University of Illinois at Urbana-Champaign, Stanford and Yale for useful comments and suggestions.

[†]Department of Economics, Massachusetts Institute of Technology.

1 Introduction

The common prior assumption is one of the cornerstones of modern economic analysis. Most models postulate that the players in a game have the “same model of the world,” or more precisely, that they have a common prior about the game form and payoff distributions—for example, they all agree that some payoff-relevant parameter vector θ is drawn from a known distribution G , even though each may also have additional information about some components of θ . The typical justification for the common prior assumption comes from *learning*; individuals, through their own experiences and the communication of others, will have access to a history of events informative about the vector θ , and this process will lead to “agreement” among individuals about the distribution of the vector θ . A strong version of this view is expressed in Savage (1954, p. 48) as the statement that a Bayesian individual, who does not assign zero probability to “the truth,” will learn it eventually as long as the signals are informative about the truth. An immediate implication of this result is that two individuals who observe the same sequence of signals will ultimately agree, even if they start with very different priors. A more sophisticated version of this conclusion also follows from Blackwell and Dubins’ (1962) theorem about the “merging of opinions”.¹

Despite these powerful intuitions and theorems, disagreement is the rule rather than the exception in practice. Just to mention a few instances, there is typically considerable disagreement even among economists working on a certain topic. For example, economists routinely disagree about the role of monetary policy, the impact of subsidies on investment or the magnitude of the returns to schooling. Similarly, there are deep divides about religious beliefs within populations with shared experiences, and finally, there was recently considerable disagreement among experts with access to the same data about whether Iraq had weapons of mass destruction. In none of these cases can the disagreements be traced to individuals having access to different histories of observations. Rather it is their *interpretations* that differ. In particular, it seems that an estimate showing that subsidies increase investment is interpreted very differently by two economists starting with different priors; for example, an economist believing that subsidies have no effect on investment appears more likely to judge the data or the methods leading to this estimate to be unreliable and thus to attach less importance to this evidence. Similarly, those who believed in the existence of weapons of mass destruction in Iraq

¹Blackwell and Dubins’ (1962) theorem shows that if two probability measures are absolutely continuous with respect to each other (meaning that they assign positive probability to the same events), then as the number of observations goes to infinity, their predictions about future frequencies will agree. This is also related to Doob’s (1948) consistency theorem for Bayesian posteriors, which we discuss and use below.

presumably interpreted the evidence from inspectors and journalists indicating the opposite as biased rather than informative.

In this paper, we show that this type of behavior will be the outcome of learning by Bayesian individuals with different priors when they are *uncertain* about the informativeness of signals. Even though Bayesian individuals will learn the asymptotic frequency of signals, they may not always be able to infer the payoff-relevant (state) variables because of an *identification problem*. The same long run frequency of signals may result from different combinations of payoff-relevant variables and different interpretations of the signals. Our main objectives in this paper are to determine when such an identification problem will prevent agreement among individuals and to provide a full characterization of when a small amount of uncertainty will lead to failure of identification and lack of agreement.

We consider the following simple environment: one or two individuals with given priors observe a sequence of signals, $\{s_t\}_{t=0}^n$, and form their posteriors about some underlying state variable (or parameter) θ . The only non-standard feature of the environment is that these individuals may be uncertain about the distribution of signals conditional on the underlying state. In the simplest case where the state and the signal are binary, e.g., $\theta \in \{A, B\}$, and $s_t \in \{a, b\}$, this implies that $\Pr(s_t = \theta \mid \theta) = p_\theta$ is not a known number, but individuals also have a prior over p_θ , say given by F_θ . We refer to this distribution F_θ as individuals' *subjective probability distribution* and to its density f_θ as *subjective (probability) density*. This distribution, which can differ among individuals, is a natural measure of their uncertainty about the informativeness of signals. When subjective probability distributions are non-degenerate, individuals will have some latitude in interpreting the sequence of signals they observe.

We first identify general conditions under which Bayesian updating leads to *asymptotic learning* (individuals learning, or believing that they will be learning, the true value of θ with probability 1 after observing infinitely many signals) and *asymptotic agreement* (convergence between their assessments of the value of θ). For the case that both individuals attach probability 1 to the event that $p_\theta > 1/2$ for $\theta \in \{A, B\}$, we show that there will always be asymptotic learning. Nevertheless, asymptotic learning is not sufficient for asymptotic agreement. We therefore also characterize the conditions for asymptotic agreement. A simple condition sufficient both for asymptotic learning and asymptotic agreement is for each individual $i = 1, 2$ to be certain that $p_\theta = p^i$ for some known number $p^i > 1/2$ (with possibly $p^1 \neq p^2$).

These positive results do not hold, however, when there is a positive probability that p_θ might be less than $1/2$. In particular, when F_θ has a full support for each θ , we show that:

1. There will not be asymptotic learning. Instead each individual’s posterior of θ continues to be a function of his prior.
2. There will not be asymptotic agreement; two individuals with different priors observing the *same* sequence of signals will reach different posterior beliefs even after observing infinitely many signals. Moreover, individuals attach *ex ante probability 1* that they will disagree after observing the sequence of signals.
3. Two individuals may *disagree more* after observing a common sequence of signals than they did so previously. In fact, for any model of learning under uncertainty that satisfies the full support assumption, there exists an open set of pairs of priors such that the disagreement between the two individuals will necessarily grow starting from these priors.

While it may appear plausible that the individuals should not attach zero probability to the event that $p_\theta < 1/2$, it is also reasonable to expect that the probability of such events should be relatively small. This raises the question of whether the results regarding the lack of asymptotic learning and agreement under uncertainty survive when there is a small amount of uncertainty.

Our most important results concern whether the asymptotic learning and agreement results under certainty are robust to a small amount of uncertainty. We investigate this issue by studying learning under “approximate certainty,” i.e., by considering a family of subjective density functions $\{f_m\}$ that become more and more concentrated around a single point—thus converging to full certainty. It is straightforward to see that as each individual becomes more and more certain about the interpretation of the signals, asymptotic learning obtains. Interestingly, however, even though each individual expects to learn the payoff-relevant parameters, asymptotic agreement may fail to obtain. This implies that asymptotic agreement under certainty may be a *discontinuous limit point* of a general model of learning under uncertainty. We show that whether or not this is the case depends on the tail properties of the family of subjective density functions $\{f_m\}$. When this family has *regularly-varying tails* (such as the Pareto or the log-normal distributions), even under approximate certainty there will be asymptotic disagreement. When $\{f_m\}$ has rapidly-varying tails (such as the normal distribution), there will be asymptotic agreement under approximate certainty.

Intuitively, approximate certainty is sufficient to make each individual believe that they will learn the payoff-relevant parameter, but they may still believe that the other individual will fail to learn. Whether or not they believe this depends on how an individual reacts when

a frequency of signals different from the one he expects with “almost certainty” occurs. If this event prevents the individual from learning, then there will be asymptotic disagreement under approximate certainty. This is because under approximate certainty, each individual trusts his own model of the world and thus expects the limiting frequencies to be consistent with his model. When the other individual’s model of the world differs, he expects the other individual to be surprised by the limiting frequencies of the signals. Then whether or not asymptotic agreement will obtain depends on whether this surprise is sufficient to prevent the other individual from learning, which in turn depends on the tail properties of the family of subjective density functions $\{f_m\}$.

Lack of asymptotic agreement has important implications for a range of economic situations. We illustrate some of these by considering a number of simple environments where two individuals observe the same sequence of signals before or while playing a game. In particular, we discuss the implications of learning in uncertain environments for games of coordination, games of common interest, bargaining, games of communication and asset trading. We show how, when they are learning under uncertainty, individuals will play these games differently than they would in environments with common priors—and also differently than in environments without common priors but where learning takes place under certainty. For example, we establish that contrary to standard results, individuals may wish to play games of common interests before receiving more information about payoffs.

We also show how the possibility of observing the same sequence of signals may lead individuals to trade *only after* they observe the public information. This result contrasts with both standard no-trade theorems (e.g., Milgrom and Stokey, 1982) and with existing results on asset trading without common priors, which assume learning under certainty (Harrison and Kreps, 1978, and Morris, 1996). Harris and Raviv (1993) have already shown that public signals can lead to greater disagreement among individuals in the “short run” and have used this observation to generate asset trading. Recent independent work by Dixit and Weibull (2007) also shows how this kind of short-run disagreement can arise both in general and also in political situations. The source of divergence of beliefs in our environment and thus the reason for the resulting asset trades are quite different, however. Existing results focus on learning under certainty. We know from Theorem 1 below that, in this case, individuals’ beliefs will eventually agree (and trade will necessarily stop). Instead when learning is under uncertainty, individuals may continue to interpret public signals differently even asymptotically (and may prefer to trade after observing infinitely many signals). The fact that the contrast between the

implications of learning under certainty and uncertainty is most transparent asymptotically is one of our main motivations for focusing on asymptotic results in this paper (i.e., learning, agreement and divergence of opinions after individuals observe *infinitely* many signals).

Our results cast doubt on the idea that the common prior assumption may be justified by learning. In many environments, even when there is little uncertainty so that each individual believes that he will learn the true state, learning does not necessarily imply agreement about the relevant parameters. Consequently, the strategic outcomes may be significantly different from those in the common-prior environments.² Whether this assumption is warranted therefore depends on the specific setting and what type of information individuals are trying to glean from the data.

Relating our results to the famous Blackwell-Dubins (1962) theorem may help clarify their essence. As briefly mentioned in Footnote 1, this theorem shows that when two agents agree on zero-probability events (i.e., their priors are absolutely continuous with respect to each other), asymptotically, they will make the same predictions about future frequencies of signals. Our results do not contradict this theorem, since we impose absolute continuity. Instead, as pointed out above, our results rely on the fact that agreeing about future frequencies is not the same as agreeing about the underlying payoff-relevant variables, because of the identification problem that arises in the presence of uncertainty.³ This identification problem leads to different possible interpretations of the same signal sequence by individuals with different priors. In most economic situations, what is important is not future frequencies of signals but some payoff-relevant parameter. For example, what was essential for the debate on the weapons of mass destruction was not the frequency of news about such weapons but whether or not they existed. What is relevant for economists trying to evaluate a policy is not the frequency of estimates on the effect of similar policies from other researchers, but the impact of this specific policy when (and if) implemented. Similarly, what may be relevant in trading assets is not the frequency of information about the dividend process, but the actual dividend that the asset will pay. Thus, many situations in which individuals need to learn about a parameter or state that will determine their ultimate payoff as a function of their action falls within the realm of the analysis here.

In this respect, our work differs from papers, such as Freedman (1963, 1965) and Miller

²For previous arguments on whether game-theoretic models should be formulated with all individuals having a common prior, see, for example, Aumann (1986, 1998) and Gul (1998). Gul (1998), for example, questions whether the common prior assumption makes sense when there is no ex ante stage.

³In this respect, our paper is also related to Kurz (1994, 1996), who considers a situation in which agents agree about long-run frequencies, but their beliefs fail to merge because of the non-stationarity of the world.

and Sanchirico (1999), that question the applicability of the absolute continuity assumption in the Blackwell-Dubins theorem in statistical and economic settings (see also Diaconis and Freedman, 1986, Stinchcombe, 2005). Similarly, a number of important theorems in statistics, for example, Berk (1966), show that when individuals place zero probability on the true data generating process, limiting posteriors will have their support on the set of all identifiable values (though they may fail to converge to a limiting distribution). Our results are different from those of Berk both because in our model individuals always place positive probability on the truth and also because we provide a tight characterization of the conditions for lack of asymptotic learning and agreement.⁴

Our paper is also closely related to recent independent work by Cripps, Ely, Mailath and Samuelson (2006), who study the conditions under which there will be “common learning” by two agents observing correlated private signals. Cripps, et al. focus on a model in which individuals start with *common priors* and then learn from *private signals* under *certainty* (though they note that their results could be extended to the case of non-common priors). They show that individual learning ensures “approximate common knowledge” when the signal space is finite, but not necessarily when it is infinite. In contrast, we focus on the case in which the agents start with *heterogenous priors* and learn from *public signals* under *uncertainty* or under *approximate certainty*. Since all signals are public in our model, there is no difficulty in achieving approximate common knowledge.⁵

The rest of the paper is organized as follows. Section 2 provides all our main results in the context of a two-state two-signal setup. Section 3 provides generalizations of these results to an environment with K states and $L \geq K$ signals. Section 4 considers a variety of applications of our results, and Section 5 concludes.

2 The Two-State Model

2.1 Environment

We start with a two-state model with binary signals. This model is sufficient to establish all our main results in the simplest possible setting. These results are generalized to arbitrary

⁴In dynamic games, another source of non-learning (and thus lack of convergence to common prior) is that some subgames are never visited along the equilibrium path and thus players do not observe information that contradict their beliefs about payoffs in these subgames (see, Fudenberg and Levine, 1993, Fudenberg and Kreps, 1995). Our results differ from those in this literature, since individuals fail to learn or fail to reach agreement despite the fact that they receive signals about all payoff-relevant variables.

⁵Put differently, we investigate whether, under approximate certainty, a player thinks that the other player will learn, whereas Cripps et al. ask whether a player i thinks that the other player j thinks that i thinks that j thinks that ... a player will learn.

number of states and signal values in Section 3.

There are two individuals, denoted by $i = 1$ and $i = 2$, who observe a sequence of signals $\{s_t\}_{t=0}^n$ where $s_t \in \{a, b\}$. The underlying state is $\theta \in \{A, B\}$, and agent i assigns ex ante probability $\pi^i \in (0, 1)$ to $\theta = A$. The individuals believe that, given θ , the signals are exchangeable, i.e., they are independently and identically distributed with an unknown distribution.⁶ That is, the probability of $s_t = a$ given $\theta = A$ is an unknown number p_A ; likewise, the probability of $s_t = b$ given $\theta = B$ is an unknown number p_B —as shown in the following table:

	A	B
a	p_A	$1 - p_B$
b	$1 - p_A$	p_B

Our main departure from the standard models is that we allow the individuals to be uncertain about p_A and p_B . We denote the cumulative distribution function of p_θ according to individual i —i.e., his *subjective probability distribution*—by F_θ^i . In the standard models, F_θ^i is degenerate and puts probability 1 at some \hat{p}_θ^i . In contrast, for most of the analysis, we will impose the following assumption:

Assumption 1 *For each i and θ , F_θ^i has a continuous, non-zero and finite density f_θ^i over $[0, 1]$.*

The assumption implies that F_θ^i has *full support* over $[0, 1]$. It is worth noting that while this assumption allows $F_\theta^1(p)$ and $F_\theta^2(p)$ to differ, for many of our results it is not important whether or not this is so (i.e., whether or not the two individuals have a common prior about the distribution of p_θ). Moreover, as discussed in Remark 2, Assumption 1 is stronger than necessary for our results, but simplifies the exposition.

In addition, throughout we assume that π^1 , π^2 , F_θ^1 and F_θ^2 are known to both individuals.⁷

We consider infinite sequences $s \equiv \{s_t\}_{t=1}^\infty$ of signals and write S for the set of all such sequences. The posterior belief of individual i about θ after observing the first n signals $\{s_t\}_{t=1}^n$ is

$$\phi_n^i(s) \equiv \Pr^i(\theta = A \mid \{s_t\}_{t=1}^n),$$

⁶See, for example, Billingsley (1995). If there were only one state, then our model would be identical to De Finetti’s canonical model (see, for example, Savage, 1954). In the context of this model, De Finetti’s theorem provides a Bayesian foundation for classical probability theory by showing that exchangeability (i.e., invariance under permutations of the order of signals) is equivalent to having an independent identical unknown distribution and implies that posteriors converge to long-run frequencies. De Finetti’s decomposition of probability distributions is extended by Jackson, Kalai and Smorodinsky (1999) to cover cases without exchangeability.

⁷The assumption that player 1 knows the prior and probability assessment of player 2 regarding the distribution of signals given the state is used in the “asymptotic agreement” results and in applications. Since our purpose is to understand whether learning justifies the common prior assumption, we assume that agents do not change their views because the beliefs of others differ from theirs.

where $\Pr^i(\theta = A | \{s_t\}_{t=1}^n)$ denotes the posterior probability that $\theta = A$ given a sequence of signals $\{s_t\}_{t=1}^n$ under prior π^i and subjective probability distribution F_θ^i . Since the sequence of signals, s , is generated by an exchangeable process, the order of the signals does not matter for the posterior. It only depends on

$$r_n(s) \equiv \# \{t \leq n | s_t = a\},$$

the number of times $s_t = a$ out of first n signals.⁸ By the strong law of large numbers, $r_n(s)/n$ converges to some $\rho(s) \in [0, 1]$ almost surely according to both individuals. Defining the set

$$\bar{S} \equiv \{s \in S : \lim_{n \rightarrow \infty} r_n(s)/n \text{ exists}\}, \quad (1)$$

this observation implies that $\Pr^i(s \in \bar{S}) = 1$ for $i = 1, 2$. We will often state our results for all sample paths s in \bar{S} , which equivalently implies that these statements are true almost surely or with probability 1. Now, a straightforward application of the Bayes rule gives

$$\phi_n^i(s) = \frac{1}{1 + \frac{1 - \pi^i}{\pi^i} \frac{\Pr^i(r_n | \theta = B)}{\Pr^i(r_n | \theta = A)}}, \quad (2)$$

where $\Pr^i(r_n | \theta)$ is the probability of observing the signal $s_t = a$ exactly r_n times out of n signals with respect to the distribution F_θ^i .

Throughout, without loss of generality, we suppose that in reality $\theta = A$. The two questions of interest for us are:

1. **Asymptotic learning:** whether $\Pr^i(\lim_{n \rightarrow \infty} \phi_n^i(s) = 1 | \theta = A) = 1$ for $i = 1, 2$.
2. **Asymptotic agreement:** whether $\Pr^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| = 0) = 1$ for $i = 1, 2$.

Notice that both asymptotic learning and agreement are defined in terms of the ex ante probability assessments of the two individuals. Therefore, asymptotic learning implies that an individual believes that he or she will ultimately learn the truth, while asymptotic agreement implies that both individuals believe that their assessments will eventually converge.⁹

⁸Given the definition of $r_n(s)$, the probability distribution \Pr^i on $\{A, B\} \times S$ is

$$\begin{aligned} \Pr^i(E^{A,s,n}) &\equiv \pi^i \int_0^1 p^{r_n(s)} (1-p)^{n-r_n(s)} f_A^i(p) dp, \text{ and} \\ \Pr^i(E^{B,s,n}) &\equiv (1-\pi^i) \int_0^1 (1-p)^{r_n(s)} p^{n-r_n(s)} f_B^i(p) dp \end{aligned}$$

at each event $E^{\theta,s,n} = \{(\theta, s') | s'_t = s_t \text{ for each } t \leq n\}$, where $s \equiv \{s_t\}_{t=1}^\infty$ and $s' \equiv \{s'_t\}_{t=1}^\infty$.

⁹We formulate asymptotic learning in terms of each individual's initial probability measure so as not to take a position on what the "objective" for "true" probability measure is.

In terms of asymptotic agreement, we will see that $\Pr^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| = 0) = 1$ also implies $\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| = 0$ for almost all sample paths, thus individual beliefs that there will be asymptotic agreement coincide with asymptotic agreement (and vice versa).

2.2 Asymptotic Learning and Disagreement

The following is a well-known result, which applies when Assumption 1 does *not* hold. A version of this result is stated in Savage (1954) and also follows from Blackwell and Dubins' (1962) more general theorem applied to this case.

Theorem 1 (Savage) *Assume that each F_θ^i puts probability 1 on \hat{p}_θ for some $\hat{p}_\theta > 1/2$, i.e., $F_\theta^i(\hat{p}_\theta) = 1$ and $F_\theta^i(p) = 0$ for each $p < \hat{p}_\theta$. Then, for each $i = 1, 2$,*

1. $\Pr^i(\lim_{n \rightarrow \infty} \phi_n^i(s) = 1 | \theta = A) = 1$.
2. $\Pr^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| = 0) = 1$.

This standard result states that when the individuals know the conditional distributions of the signals (and hence they agree what those distributions are), they will learn the truth with experience (almost surely as $n \rightarrow \infty$) and two individuals observing the same sequence will necessarily come to agree what the underlying state, θ , is. A simple intuition for this result is that the underlying state θ is *fully identified* from the limiting frequencies, so that both individuals can infer the underlying state from the observation of the limiting frequencies of signals.

However, there is more to this theorem than the simple intuition. Each individual is sure that they will be confronted either with a limiting frequency of a signals equal to \hat{p}_A , in which case they will conclude that $\theta = A$, or they will observe a limiting frequency of $1 - \hat{p}_B$, and they will conclude that $\theta = B$; and they attach zero probability to the events that they will observe a different asymptotic frequency. What happens if an individual observes a frequency ρ of signals different from \hat{p}_A and $1 - \hat{p}_B$ in a large sample of size n ? The answer to this question will provide the intuition for some of the results that we will present next. Observe that this event has zero probability under the individual's beliefs at the limit $n = \infty$. However, for $n < \infty$ he will assign a strictly positive (but small) probability to such a frequency of signals resulting from *sampling variation*. Moreover, it is straightforward to see that there exists a unique $\hat{\rho}(\hat{p}_A, \hat{p}_B) \in (1 - \hat{p}_B, \hat{p}_A)$ such that when $\rho > \hat{\rho}(\hat{p}_A, \hat{p}_B)$, the required sampling variation that leads to ρ under $\theta = B$ is *infinitely greater* (as $n \rightarrow \infty$) than the one under $\theta = A$. This cutoff value $\hat{\rho}(p_A, p_B)$ is clearly the solution to the equation $p_A^\rho (1 - p_A)^{1-\rho} = p_B^{1-\rho} (1 - p_B)^\rho$, given by

$$\hat{\rho}(p_A, p_B) \equiv \frac{\log(p_B / (1 - p_A))}{\log(p_B / (1 - p_A)) + \log(p_A / (1 - p_B))} \in (1 - p_B, p_A). \quad (3)$$

Consequently, when $\rho > \hat{\rho}(\hat{p}_A, \hat{p}_B)$, the individual will asymptotically assign probability 1 to the event that $\theta = A$. Conversely, when $\rho < \hat{\rho}(\hat{p}_A, \hat{p}_B)$, he will assign probability 1 to $\theta = B$.

In Theorem 1, the assumption that $\hat{p}_\theta > 1/2$ assures that $\hat{p}_A \neq 1 - \hat{p}_B$, so that θ is fully identified. Our next result determines exactly how far we can generalize asymptotic learning and agreement results of Theorem 1. Let us denote the *support* of a distribution F by $\text{supp}F$ (that is, the smallest set such that the distribution F assigns zero probability to all events outside this set). Let us define $\inf(\text{supp}F)$ to be the infimum of the set $\text{supp}F$ (i.e., the largest p such that $F(p) = 0$).

Theorem 2 (Generalized Asymptotic Learning and Agreement) *Let $\hat{\rho}(p_A, p_B)$ be as defined in (3). Assume that for each θ and i , $p_{\theta,i} = \inf(\text{supp}F_\theta^i) \in (1/2, 1)$ and $1 - p_{B,i} \neq \hat{\rho}(p_{A,j}, p_{B,j}) \neq p_{A,i}$ for all $i \neq j$.*

1. $\Pr^i(\lim_{n \rightarrow \infty} \phi_n^i(s) = 1 | \theta = A) = 1$;
2. $\Pr^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| = 0) = 1$ if and only if $1 - p_{B,i} < \hat{\rho}(p_{A,j}, p_{B,j}) < p_{A,i}$.

Proof. We will first prove that for any $s \in \bar{S}$,

$$\lim_{n \rightarrow \infty} \phi_n^i(s) = \begin{cases} 1 & \text{if } r_n(s)/n \rightarrow \rho > \hat{\rho}(p_{A,i}, p_{B,i}) \\ 0 & \text{if } r_n(s)/n \rightarrow \rho < \hat{\rho}(p_{A,i}, p_{B,i}) \end{cases}. \quad (4)$$

Both parts of the theorem follow readily from (4). First, since $p_{\theta,i} = \inf(\text{supp}F_\theta^i) \in (1/2, 1)$, (4) implies that conditional on $\theta = A$, agent i assigns probability 1 to the event that $r_n(s)/n \rightarrow \rho \geq p_{A,i} > \hat{\rho}(p_{A,i}, p_{B,i})$, where the last inequality follows from (3). This implies that $\Pr^i(\lim_{n \rightarrow \infty} \phi_n^i(s) = 1 | \theta = A) = 1$, establishing part 1.

For part 2, suppose that $\hat{\rho}(p_{A,j}, p_{B,j}) < p_{A,i}$. Then conditional on $\theta = A$, (4) implies that $\phi_n^j(s)$ also converges to 1, and therefore $|\phi_n^1(s) - \phi_n^2(s)| \rightarrow 0$. Next, when $\hat{\rho}(p_{A,j}, p_{B,j}) > 1 - p_{B,i}$, we have that $\phi_n^j(s) \rightarrow 0$ and also conditional on $\theta = B$, $\phi_n^i(s) \rightarrow 0$, thus $|\phi_n^1(s) - \phi_n^2(s)| \rightarrow 0$, proving sufficiency. To prove necessity, suppose $p_{A,i} < \hat{\rho}(p_{A,j}, p_{B,j})$. Then, i assigns strictly positive probability to the event that $r_n(s)/n \rightarrow \rho \in [p_{A,i}, \hat{\rho}(p_{A,j}, p_{B,j})]$. But then (4) implies $\phi_n^i(s) \rightarrow 1$ and $\phi_n^j(s) \rightarrow 0$, so that $|\phi_n^1(s) - \phi_n^2(s)| \rightarrow 1$, establishing that $1 - p_{B,i} < \hat{\rho}(p_{A,j}, p_{B,j}) < p_{A,i}$ is necessary for $\Pr^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| = 0) = 1$.

To complete the proof, we need to derive (4). Let us define

$$R_n^i(r_n) \equiv \frac{\Pr^i(r_n | \theta = B)}{\Pr^i(r_n | \theta = A)} = \frac{\int (1-p)^{r_n} p^{n-r_n} dF_B^i}{\int p^{r_n} (1-p)^{n-r_n} dF_A^i}.$$

Take any $\rho > \hat{\rho}(p_{A,i}, p_{B,i})$. Since $1 - p_{B,i} < p_{A,i}$, we have

$$(1 - p_{B,i})^\rho p_{B,i}^{1-\rho} < p_{A,i}^\rho (1 - p_{A,i})^{1-\rho}. \quad (5)$$

The function $p^\rho (1 - p)^{1-\rho}$ is continuous and concave in p , and reaches its maximum at $p = \rho$. Then (5) implies that there exists $\varepsilon > 0$ and $\hat{p} > p_{A,i}$ such that for all $\tilde{p} \in \text{supp}F_B^i$, $p \in [p_{A,i}, \hat{p}]$, $r_n/n \in (\rho - \varepsilon, \rho + \varepsilon)$,

$$(1 - \tilde{p})^{r_n} \tilde{p}^{n-r_n} \leq (1 - p_{B,i})^{r_n} p_{B,i}^{n-r_n} < p^{r_n} (1 - p)^{n-r_n} \leq \hat{p}^{r_n} (1 - \hat{p})^{n-r_n}.$$

Hence,

$$\int (1 - p)^{r_n} p^{n-r_n} dF_B^i \leq (1 - p_{B,i})^{r_n} p_{B,i}^{n-r_n}$$

and

$$\int p^{r_n} (1 - p)^{n-r_n} dF_A^i \geq \int_{p \leq \hat{p}} p^{r_n} (1 - p)^{n-r_n} dF_A^i \geq F_A^i(\hat{p}) \hat{p}^{r_n} (1 - \hat{p})^{n-r_n}.$$

Thus,

$$0 \leq R_n^i(r_n) \leq \frac{1}{F_A^i(\hat{p})} \left(\frac{(1 - p_{B,i})^{r_n/n} p_{B,i}^{1-r_n/n}}{\hat{p}^{r_n/n} (1 - \hat{p})^{1-r_n/n}} \right)^n.$$

When $r_n/n \in [\rho - \varepsilon/2, \rho + \varepsilon/2]$, the expression in the paranthesis is smaller than 1 and therefore the right-hand side converges to 0 as $n \rightarrow \infty$ and $r_n/n \rightarrow \rho$. This implies $R_n^i(r_n) \rightarrow 0$ and thus $\phi_n^i(s) \rightarrow 1$. The same argument (switching A and B) implies that when $\rho < \hat{\rho}(p_{A,i}, p_{B,i})$, $\phi_n^i(s) \rightarrow 0$. This establishes (4) and completes the proof. ■

The characterization in part 2 generalizes the asymptotic agreement result of Theorem 1 in several directions. The following corollary spells out some of the implications of part 2 of Theorem 2.

Corollary 1 (Sufficient Conditions for Asymptotic Agreement) *Under the assumptions of Theorem 2, there will be asymptotic agreement whenever any of the following conditions hold:*

1. certainty (with symmetry): each F_θ^i puts probability 1 on some $\hat{p}^i > 1/2$;
2. symmetric support: $\text{supp}F_A^i = \text{supp}F_B^i$ for each i ;
3. common support: $\text{supp}F_\theta^1 = \text{supp}F_\theta^2$ for each θ .

Proof. Part 1 of the corollary is a special case of part 2. Under symmetric support assumption, we have $\hat{\rho}(p_{A,i}, p_{B,i}) = 1/2$ for each i , so that part 2 of the corollary follows from part 2 of Theorem 2. Finally, part 3 of the corollary follows from the fact that under the common support assumption $\hat{\rho}(p_{A,j}, p_{B,j}) = \hat{\rho}(p_{A,i}, p_{B,i}) \in (1 - p_{B,i}, p_{A,i})$. ■

Theorem 2 provides a precise answer to one of our main questions: it shows that under the “full identification assumption” that $p_{\theta,i} > 1/2$ for each θ and i , asymptotic learning always obtains. Furthermore, asymptotic agreement depends on the lowest value $p_{\theta,i}$ of p_{θ} that individual $i = 1, 2$ assigns positive probability. This is intuitive in view of the discussion preceding Theorem 2. When individual i observes a frequency $\rho \in (1 - p_{B,i}, p_{A,i})$, he presumes that this has resulted from sampling variation, and decides whether frequency ρ is more likely under $\theta = A$ or under $\theta = B$. In particular, for each θ , the lowest sampling variation that leads to ρ is attained at $p_{\theta,i}$, and the asymptotic beliefs depend only on how large these variations are. When $\rho > \hat{\rho}(p_{A,i}, p_{B,i})$ (and as $n \rightarrow \infty$) the necessary sampling variation is infinitely smaller under $\theta = A$ than under $\theta = B$. Consequently, the individual believes with probability 1 that $\theta = A$. Conversely, when $\rho < \hat{\rho}(p_{A,i}, p_{B,i})$, he believes with probability 1 that $\theta = B$. Whether there will be asymptotic agreement then purely depends on whether and how different the cutoff values $\hat{\rho}(p_{A,1}, p_{B,1})$ and $\hat{\rho}(p_{A,2}, p_{B,2})$ are. When they are close, both individuals will interpret the limiting frequency of signals, ρ , similarly, even when this is a frequency to which they initially assigned zero probability, and will reach asymptotic agreement. In contrast, when these cutoff values are far apart, so that $\hat{\rho}(p_{A,j}, p_{B,j}) \notin (1 - p_{B,i}, p_{A,i})$, both players assign positive probability to the event that their beliefs will diverge to the extremes: $\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| = 1$.

Corollary 1 shows that various reasonable conditions ensure asymptotic agreement. Asymptotic agreement is implied, for example, by certainty, symmetric support or common support assumptions. In particular, certainty (with symmetry), which corresponds to both individuals believing that limiting frequencies have to be \hat{p}^i or $1 - \hat{p}^i$ (but with $\hat{p}^1 \neq \hat{p}^2$) is sufficient for asymptotic agreement. In this case, each individual is certain about what the limiting frequency will be and therefore believes that the frequency expected by the other individual will not be realized (creating a discrepancy between that individual’s initial belief and observation). Nevertheless, with the same reasoning as in the discussion preceding Theorem 2, each individual also believes that the other individual will ascribe this discrepancy to sampling variation and reach the same conclusion as himself. This is sufficient for asymptotic agreement.

Theorem 2 and Corollary 1 therefore show that results on asymptotic learning and agreement are substantially more general than Savage’s original theorem (Theorem 1). Nevertheless, these results rely on the feature that $F_\theta^i(1/2) = 0$ for each $i = 1, 2$ and each θ . This implies that both individuals attach zero probability to a range of possible models of the world—i.e., they are certain that p_θ cannot be less than $1/2$. There are two reasons for considering situations in which this is not the case. First, the preceding discussion illustrates why assigning zero probability to certain models of the world is important; it enables individuals to ascribe any frequency of signals that are unlikely under these models to sampling variability. This kind of inference may be viewed as somewhat unreasonable, since individuals are reaching very strong conclusions based on events that have vanishingly small probabilities (since sampling variability vanishes as $n \rightarrow \infty$). Second, our motivation of investigating learning under uncertainty suggests that individuals may attach positive (albeit small) probabilities to all possible values of p_θ . This latter feature is what the full support assumption, Assumption 1, implies. We next impose this assumption and show that under the more general circumstances where F_θ^i has full support, there will be neither asymptotic learning nor asymptotic agreement.

Theorem 3 (*Lack of Asymptotic Learning and Agreement*) *Suppose Assumption 1 holds for $i = 1, 2$. Then,*

1. $\Pr^i(\lim_{n \rightarrow \infty} \phi_n^i(s) \neq 1 | \theta = A) = 1$ for $i = 1, 2$;
2. $\Pr^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| \neq 0) = 1$ whenever $\pi^1 \neq \pi^2$ and $F_\theta^1 = F_\theta^2$ for each $\theta \in \{A, B\}$.

This theorem contrasts with Theorems 1 and 2 and implies that, with probability 1, each individual will fail to learn the true state. The second part of the theorem states that if the individuals’ prior beliefs about the state differ (but they interpret the signals in the same way), then their posteriors will eventually disagree, and moreover, they will both attach probability 1 to the event that their beliefs will eventually diverge. Put differently, this implies that there is “agreement to eventually disagree” between the two individuals, in the sense that they both believe ex ante that after observing the signals they will fail to agree. This feature will play an important role in the applications in Section 4 below.

Remark 1 The assumption that $F_\theta^1 = F_\theta^2$ in this theorem is adopted for simplicity. Even in the absence of this condition, there will typically be no asymptotic agreement. Theorem 6 in the next section generalizes this theorem to a situation with multiple states and multiple signals

and also dispenses with the assumption that $F_\theta^1 = F_\theta^2$. It establishes that the set of priors and subjective probability distributions that leads to asymptotic agreement is of “measure zero”.

Remark 2 Assumption 1 is considerably stronger than the necessary conditions for Theorem 3. It is adopted only for simplicity. It can be verified that for lack of asymptotic learning it is sufficient (but not necessary) that the measures generated by the distribution functions $F_A^i(p)$ and $F_B^i(1-p)$ be absolutely continuous with respect to each other. Similarly, for lack of asymptotic agreement, it is sufficient (but not necessary) that the measures generated by $F_A^1(p)$, $F_B^1(1-p)$, $F_A^2(p)$ and $F_B^2(1-p)$ be absolutely continuous with respect each other. For example, if both individuals believe that p_A is either 0.3 or 0.7 (with the latter receiving greater probability) and that p_B is also either 0.3 or 0.7 (with the former receiving greater probability), then there will be neither asymptotic learning nor asymptotic agreement. Throughout we use Assumption 1 both because it simplifies the notation and because it is a natural assumption when we turn to the analysis of asymptotic agreement under approximate certainty below.

We next provide a proof of Theorem 3. The next lemma provides a useful formula for $\phi_\infty^i(s) \equiv \lim_{n \rightarrow \infty} \phi_n^i(s)$ for all sample paths s in \bar{S} and also introduces the concept of the asymptotic likelihood ratio, which will be used throughout the rest of the paper.

Lemma 1 *Suppose Assumption 1 holds. Then for all $s \in \bar{S}$,*

$$\phi_\infty^i(\rho(s)) \equiv \lim_{n \rightarrow \infty} \phi_n^i(s) = \frac{1}{1 + \frac{1-\pi^i}{\pi^i} R^i(\rho(s))}, \quad (6)$$

where $\rho(s) = \lim_{n \rightarrow \infty} r_n(s)/n$, and $\forall \rho \in [0, 1]$,

$$R^i(\rho) \equiv \frac{f_B^i(1-\rho)}{f_A^i(\rho)} \quad (7)$$

is the asymptotic likelihood ratio.

Proof. Write

$$\begin{aligned} \frac{\Pr^i(r_n | \theta = B)}{\Pr^i(r_n | \theta = A)} &= \frac{\int_0^1 p^{r_n} (1-p)^{n-r_n} f_B(1-p) dp}{\int_0^1 p^{r_n} (1-p)^{n-r_n} f_A(p) dp} \\ &= \frac{\int_0^1 p^{r_n} (1-p)^{n-r_n} f_B(1-p) dp}{\int_0^1 p^{r_n} (1-p)^{n-r_n} dp} \\ &= \frac{\int_0^1 p^{r_n} (1-p)^{n-r_n} f_A(p) dp}{\int_0^1 p^{r_n} (1-p)^{n-r_n} dp} \\ &= \frac{\mathbb{E}^\lambda[f_B(1-p) | r_n]}{\mathbb{E}^\lambda[f_A(p) | r_n]}. \end{aligned}$$

Here, the first equality is obtained by dividing the numerator and the denominator by the same term. The resulting expression on the numerator is the conditional expectation of $f_B(1-p)$ given r_n under the flat (Lebesgue) prior on p and the Bernoulli distribution on $\{s_t\}_{t=0}^n$. Denoting this by $\mathbb{E}^\lambda[f_B(1-p)|r_n]$, and the denominator, which is similarly defined as the conditional expectation of $f_A(p)$, by $\mathbb{E}^\lambda[f_A(p)|r_n]$, we obtain the last equality. By Doob's consistency theorem for Bayesian posterior expectation of the parameter, as $r_n \rightarrow \rho$, we have that $\mathbb{E}^\lambda[f_B(1-p)|r_n] \rightarrow f_B(1-\rho)$ and $\mathbb{E}^\lambda[f_A(p)|r_n] \rightarrow f_A(\rho)$ (see, e.g., Doob, 1949, Ghosh and Ramamoorthi, 2003, Theorem 1.3.2). This establishes

$$\frac{\Pr^i(r_n|\theta = B)}{\Pr^i(r_n|\theta = A)} \rightarrow R^i(\rho),$$

as defined in (7). Equation (6) then follows from (2). ■

In equation (7), $R^i(\rho)$ is the *asymptotic likelihood ratio* of observing frequency ρ of a when the true state is B versus when it is A . Lemma 1 states that, asymptotically, individual i uses this likelihood ratio and Bayes rule to compute his posterior beliefs about θ .

An immediate implication of Lemma 1 is that given any $s \in \bar{S}$,

$$\phi_\infty^1(\rho(s)) = \phi_\infty^2(\rho(s)) \text{ if and only if } \frac{1-\pi^1}{\pi^1}R^1(\rho(s)) = \frac{1-\pi^2}{\pi^2}R^2(\rho(s)). \quad (8)$$

The proof of Theorem 3 now follows from Lemma 1 and equation (8).

Proof of Theorem 3. Since $f_B^i(1-\rho(s)) > 0$ and $f_A(\rho(s))$ is finite, $R^i(\rho(s)) > 0$. Hence, by Lemma 1, $\phi_\infty^i(\rho(s)) \neq 1$ for each s , establishing the first part. The second part follows from equation (8), since $\pi^1 \neq \pi^2$ and $F_\theta^1 = F_\theta^2$ implies that for each $s \in \bar{S}$, $\phi_\infty^1(s) \neq \phi_\infty^2(s)$, and thus $\Pr^i(|\phi_\infty^1(s) - \phi_\infty^2(s)| \neq 0) = 1$ for $i = 1, 2$. ■

Intuitively, when Assumption 1 (in particular, the full support feature) holds, an individual is never sure about the exact interpretation of the sequence of signals he observes and will update his views about p_θ (the informativeness of the signals) as well as his views about the underlying state. For example, even when signal a is more likely in state A than in state B , a very high frequency of a will not necessarily convince him that the true state is A , because he may infer that the signals are not as reliable as he initially believed, and they may instead be biased towards a . Therefore, the individual never becomes certain about the state, which is captured by the fact that $R^i(\rho)$ defined in (7) never takes the value zero or infinity. Consequently, as shown in (6), his posterior beliefs will be determined by his prior beliefs about the state and also by R^i , which tells us how the individual updates his beliefs about the

informativeness of the signals as he observes the signals. When two individuals interpret the informativeness of the signals in the same way (i.e., $R^1 = R^2$), the differences in their priors will always be reflected in their posteriors.

In contrast, if an individual were certain about the informativeness of the signals (i.e., if i were sure that $p_\theta = p_\theta^i$ for some $p_\theta^i > 1/2$) as in Theorem 1 and Corollary 1, then he would never question the informativeness of the signals, even when the limiting frequency of a converges to a value different from p_A^i or $1 - p_B^i$, and would interpret such discrepancies as resulting from sampling variation. This would be sufficient for asymptotic agreement. The full support assumption in Assumption 1 prevents this type of reasoning and ensures asymptotic disagreement.

As noted above, an important implication of Theorem 3 is that there will typically be “agreement to eventually disagree” between the individuals. In other words, given their priors, both individuals will agree that after seeing the same infinite sequence of signals they will still disagree (with probability 1). This implication is interesting in part because the common prior assumption, typically justified by learning, leads to the celebrated “no agreement to disagree” result (Aumann, 1976, 1998), which states that if the individuals’ posterior beliefs are common knowledge, then they must be equal.¹⁰ In contrast, in the limit of the learning process here, individuals’ beliefs are common knowledge (as there is no private information), but they are different with probability 1. This is because in the presence of uncertainty and full support as in Assumption 1, both individuals understand that their priors will have an effect on their beliefs even asymptotically; thus they expect to disagree. Many of the applications we discuss in Section 4 exploit this feature.

2.3 Divergence of Opinions

Theorem 3 established that the differences in priors are reflected in the posteriors even in the limit as $n \rightarrow \infty$. It does not, however, quantify the possible disagreement between the two individuals. The rest of this section investigates different aspects of this question. We first show that two individuals that observe the same sequence of signals may have diverging posteriors, so that common information can increase disagreement.

Theorem 4 (*Divergence of Opinions*) *Suppose that subjective probability distributions are given by F_θ^1 and F_θ^2 that satisfy Assumption 1 and that there exists $\epsilon > 0$ such that*

¹⁰Note, however, that the “no agreement to disagree” result derives from individuals’ updating their beliefs because those of others differ from their own (Geanakoplos and Polemarchakis, 1982), whereas here individuals only update their beliefs by learning.

$|R^1(\rho) - R^2(\rho)| > \epsilon$ for each $\rho \in [0, 1]$. Then, there exists an open set of priors π^1 and π^2 , such that for all $s \in \bar{S}$,

$$\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| > |\pi^1 - \pi^2|;$$

in particular,

$$\Pr^i \left(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| > |\pi^1 - \pi^2| \right) = 1.$$

Proof. Fix F_θ^1 and F_θ^2 and take $\pi^1 = \pi^2 = 1/2$. By Lemma 1 and the hypothesis that $|R^1(\rho) - R^2(\rho)| > \epsilon$ for each $\rho \in [0, 1]$, $\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| > \epsilon'$ for some $\epsilon' > 0$, while $|\pi^1 - \pi^2| = 0$. Since both expressions are continuous in π^1 and π^2 , there is an open neighborhood of $1/2$ such that the above inequality uniformly holds for each ρ whenever π^1 and π^2 are in this neighborhood. The last statement follows from the fact that $\Pr^i(s \in \bar{S}) = 1$.

■

Remark 3 The assumption that R^1 and R^2 are different for each ρ is not very restrictive but implies that both R^1 and R^2 cannot be continuous everywhere. It is straightforward to extend Theorem 4 such that for $\epsilon > 0$ the hypothesis that the (Lebesgue) measure of the set over which $|R^1(\rho) - R^2(\rho)| > \epsilon$ is greater than $1 - \epsilon$ implies $\Pr^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| > |\pi^1 - \pi^2|) \geq 1 - \epsilon$. We do not state this generalization of Theorem 4 to economize on space and notation.

Intuitively, even a small difference in priors ensures that individuals will interpret signals differently, and if the original disagreement is relatively small, after almost all sequences of signals, the disagreement between the two individuals grows. Consequently, the observation of a common sequence of signals causes an initial difference of opinion between individuals to widen (instead of the standard merging of opinions under certainty). Theorem 4 also shows that both individuals are certain ex ante that their posteriors will diverge after observing the same sequence of signals, because they understand that they will interpret the signals differently. This strengthens our results further and shows that for some priors individuals will “agree to eventually disagree even more”.

An interesting implication of Theorem 4 is also worth noting. As demonstrated by Theorems 1 and 2, when there is learning under certainty individuals initially disagree, but each individual also believes that they will eventually agree (and in fact, that they will converge to his beliefs). This implies that each individual expects the other to “learn more”. More specifically, let $\mathbf{I}_{\theta=A}$ be the indicator function for $\theta = A$ and $\Lambda^i = (\pi^i - \mathbf{I}_{\theta=A})^2 - (\phi_\infty^i - \mathbf{I}_{\theta=A})^2$ be a measure of learning for individual i , and let \mathbb{E}^i be the expectation of individual i (under

the probability measure \Pr^i). Under certainty, Theorem 1 implies that $\phi_\infty^i = \phi_\infty^j = \mathbf{I}_{\theta=A}$, so that $\mathbb{E}^i[\Lambda^i - \Lambda^j] = -(\pi^i - \pi^j)^2 < 0$ and thus $\mathbb{E}^i[\Lambda^i] < \mathbb{E}^i[\Lambda^j]$. Under uncertainty, this is not necessarily true. In particular, Theorem 4 implies that, under the assumptions of the theorem, there exists an open subset of the interval $[0, 1]$ such that whenever π^1 and π^2 are in this subset, we have $\mathbb{E}^i[\Lambda^i] > \mathbb{E}^i[\Lambda^j]$, so that individual i would expect to learn more than individual j . The reason is that individual i is not only confident about his initial guess π^i , but also expects to *learn more* from the sequence of signals than individual j , because he believes that individual j has the “wrong model of the world.” The fact that an individual may expect to learn more than others will play an important role in some of the applications in Section 4.

2.4 Non-monotonicity of the Likelihood Ratio and Discontinuity of Asymptotic Agreement

Theorems 3 and 4 show how Bayesian individuals may fail to agree even after observing an infinite sequence of common signals. One may be concerned that these results rely on a significant degree of uncertainty on the part of the individuals (i.e., on significantly dispersed subjective probability distributions). In this subsection, we show by means of an example that even a very small amount of uncertainty can lead to significant asymptotic disagreement. This topic is then investigated in greater detail in the next subsection.

Our example in this subsection will serve two purposes. First, it will show the possibility of discontinuity of asymptotic agreement. Second, it will illustrate that the important role that the non-monotonicity of the asymptotic likelihood ratio, $R^i(\rho)$, plays in the behavior of individual beliefs and the possibility of asymptotic agreement under a small amount of uncertainty. When the asymptotic likelihood ratio is non-monotone, a high frequency of signals taking the value a may imply that the signals could be biased towards a and induce the individual to put lower probability on state A than he would have done with a lower frequency of a among the signals.

We start with a simple result. Inspection of expression (6) establishes the following:

Lemma 2 *For any $s \in \bar{S}$, $\phi_\infty^i(s)$ is decreasing at $\rho(s)$ if and only if R^i is increasing at $\rho(s)$.*

Proof. This follows immediately from equation (6) above. ■

When R^i is non-monotone, even a small amount of uncertainty about the informativeness of the signals may lead to significant differences in limit posteriors. The next example illustrates this point, while the second example shows that there can be “reversals” in individuals’

assessments, meaning that after observing a sequence “favorable” to state A , the individual may have a lower posterior about this state than his prior.

Example 1 (*Non-monotonicity*) Each individual i thinks that with probability $1 - \epsilon$, p_A and p_B are in a δ -neighborhood of some $\hat{p}^i > (1 + \delta)/2$, but with probability $\epsilon > 0$, the signals are not informative. More precisely, for $\hat{p}^i > (1 + \delta)/2$, $\epsilon > 0$ and $\delta < |\hat{p}^1 - \hat{p}^2|$, we have

$$f_{\theta}^i(p) = \begin{cases} \epsilon + (1 - \epsilon)/\delta & \text{if } p \in (\hat{p}^i - \delta/2, \hat{p}^i + \delta/2) \\ \epsilon & \text{otherwise} \end{cases} \quad (9)$$

for each θ and i . Now, by (7), the asymptotic likelihood ratio is

$$R^i(\rho(s)) = \begin{cases} \frac{\epsilon\delta}{1-\epsilon(1-\delta)} & \text{if } \rho(s) \in (\hat{p}^i - \delta/2, \hat{p}^i + \delta/2) \\ \frac{1-\epsilon(1-\delta)}{\epsilon\delta} & \text{if } \rho(s) \in (1 - \hat{p}^i - \delta/2, 1 - \hat{p}^i + \delta/2) \\ 1 & \text{otherwise.} \end{cases}$$

This and other relevant functions are plotted in Figure 1 for $\epsilon \rightarrow 0$ and $\delta \rightarrow 0$. The likelihood ratio $R^i(\rho(s))$ is 1 when $\rho(s)$ is small, takes a very high value at $1 - \hat{p}^i$, goes down to 1 afterwards, becomes nearly zero around \hat{p}^i , and then jumps back to 1. By Lemmas 1 and 2, $\phi_{\infty}^i(s)$ will also be non-monotone: when $\rho(s)$ is small, the signals are not informative, thus $\phi_{\infty}^i(s)$ is the same as the prior, π^i . In contrast, around $1 - \hat{p}^i$, the signals become very informative suggesting that the state is B , thus $\phi_{\infty}^i(s) \cong 0$. After this point, the signals become uninformative again and $\phi_{\infty}^i(s)$ goes back to π^i . Around \hat{p}^i , the signals are again informative, but this time favoring state A , so $\phi_{\infty}^i(s) \cong 1$. Finally, signals again become uninformative and $\phi_{\infty}^i(s)$ falls back to π^i . Intuitively, when $\rho(s)$ is around $1 - \hat{p}^i$ or \hat{p}^i , the individual assigns very high probability to the true state, but outside of this region, he sticks to his prior, concluding that the signals are not informative.

The first important observation is that even though ϕ_{∞}^i is equal to the prior for a large range of limiting frequencies, as $\epsilon \rightarrow 0$ and $\delta \rightarrow 0$ each individual attaches probability 1 to the event that he will learn θ . This is because as illustrated by the discussion after Theorem 1, as $\epsilon \rightarrow 0$ and $\delta \rightarrow 0$, each individual becomes convinced that the limiting frequencies will be either $1 - \hat{p}^i$ or \hat{p}^i .

However, asymptotic learning is considerably weaker than asymptotic agreement. Each individual also understands that since $\delta < |\hat{p}^1 - \hat{p}^2|$, when the long-run frequency is in a region where he learns that $\theta = A$, the other individual will conclude that the signals are uninformative and adhere to his prior belief. Consequently, he expects the posterior beliefs of the other individual to be always far from his. Put differently, as $\epsilon \rightarrow 0$ and $\delta \rightarrow 0$, each




Figure 1: The three panels show, respectively, the approximate values of $R^i(\rho)$, ϕ_∞^i , and $|\phi_\infty^1 - \phi_\infty^2|$ as $\epsilon \rightarrow 0$.

individual believes that he will learn the value of θ himself but that the other individual will fail to learn, thus attaches probability 1 to the event that they disagree. This can be seen from the third panel of Figure 1; at each sample path in \bar{S} , at least one of the individuals will fail to learn, and the difference between their limiting posteriors will be uniformly higher than the following “objective” bound

$$\min \{ \pi^1, \pi^2, 1 - \pi^1, 1 - \pi^2, |\pi^1 - \pi^2| \}.$$

When $\pi^1 = 1/3$ and $\pi^2 = 2/3$, this bound is equal to $1/3$. In fact, the belief of each individual regarding potential disagreement can be greater than this; each individual believes that he will learn but the other individual will fail to do so. Consequently, for each i , $\Pr^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| \geq Z) \geq 1 - \epsilon$, where as $\epsilon \rightarrow 0$, $Z \rightarrow \min \{ \pi^1, \pi^2, 1 - \pi^1, 1 - \pi^2 \}$. This “subjective” bound can be as high as $1/2$.

Example 1 illustrates how asymptotic agreement may be a “discontinuous” limit point of models with small amount of uncertainty. The important feature leading to this result is the non-monotonicity of the asymptotic likelihood ratio $R^i(\rho)$. However, in this example, $R^i(\rho)$ is not only non-monotonic but also discontinuous. In the next subsection, we will see that, under some additional conditions, similar results obtain when $R^i(\rho)$ is a continuous function of ρ . Before presenting these results, we also observe that an even more extreme phenomenon, whereby a high frequency of $s = a$ among the signals may reduce the individual’s posterior that $\theta = A$ below his prior, is also possible.

Example 2 (Reversal) Now suppose that individuals' subjective probability densities are given by

$$f_{\theta}^i(p) = \begin{cases} (1 - \epsilon - \epsilon^2) / \delta & \text{if } \hat{p}^i - \delta/2 \leq p \leq \hat{p}^i + \delta/2 \\ \epsilon & \text{if } p < 1/2 \\ \epsilon^2 & \text{otherwise} \end{cases}$$

for each θ and $i = 1, 2$, where $\epsilon > 0$, $\hat{p}^i > 1/2$, and $0 < \delta < \hat{p}^1 - \hat{p}^2$. Clearly, as $\epsilon \rightarrow 0$, (7) gives:

$$R^i(\rho(s)) \cong \begin{cases} 0 & \begin{array}{l} \text{if } \rho(s) < 1 - \hat{p}^i - \delta/2, \\ \text{or } 1 - \hat{p}^i + \delta/2 < \rho(s) < 1/2, \\ \text{or } \hat{p}^i - \delta/2 \leq \rho(s) \leq \hat{p}^i + \delta/2 \end{array} \\ \infty & \text{otherwise.} \end{cases}$$

Hence, the asymptotic posterior probability that $\theta = A$ is

$$\phi_{\infty}^i(\rho(s)) \cong \begin{cases} 1 & \begin{array}{l} \text{if } \rho(s) < 1 - \hat{p}^i - \delta/2, \\ \text{or } 1 - \hat{p}^i + \delta/2 < \rho(s) < 1/2, \\ \text{or } \hat{p}^i - \delta/2 \leq \rho(s) \leq \hat{p}^i + \delta/2 \end{array} \\ 0 & \text{otherwise.} \end{cases}$$

Consequently, in this case observing a sufficiently high frequency of $s = a$ may reduce the posterior that $\theta = A$ below the prior. Moreover, the individuals assign probability $1 - \epsilon$ that there will be extreme asymptotic disagreement in the sense that $|\phi_{\infty}^1(\rho(s)) - \phi_{\infty}^2(\rho(s))| \cong 1$.

In both examples, it is crucial that the likelihood ratio R^i is not monotone. If R^i were monotone, at least one of the individuals would expect that their beliefs will asymptotically agree. To see this, take $\hat{p}^i \geq \hat{p}^j$. Given the form of $R^i(\rho)$, individual i is almost certain that, when the state is A , $\rho(s)$ will be close to \hat{p}^i . He also understands that j would assign a very high probability to the event that $\theta = A$ when $\rho(s) = \hat{p}^j \geq \hat{p}^i$. If R^j were monotone, individual j would assign even higher probability to A at $\rho(s) = \hat{p}^i$ and thus his probability assessment on A would also converge to 1 as $\epsilon \rightarrow 0$. Therefore, in this case i will be almost certain that j will learn the true state and that their beliefs will agree asymptotically.

Theorem 1 and Corollary 1 established that there will be asymptotic agreement under certainty when $\hat{p}^1, \hat{p}^2 > 1/2$. One might have thought that as $\epsilon \rightarrow 0$ and uncertainty disappears, the same conclusion would apply. In contrast, the above examples show that even as each F_{θ}^i converges to a Dirac distribution (that assigns a unit mass to a point) at $\hat{p}^i > 1/2$, there may be significant asymptotic disagreement between the two individuals. Notably this is true not only when there is negligible uncertainty, i.e., $\epsilon \rightarrow 0$ and $\delta \rightarrow 0$, but also when the individuals' subjective distributions are nearly identical, i.e., as $\hat{p}^1 - \hat{p}^2 \rightarrow 0$ as in Theorem 1. This shows

that the result of asymptotic agreement in Theorem 1 is not necessarily a continuous limit point of the more general model of learning under uncertainty (which allows for nondegenerate subjective probability distributions F_θ^i 's).¹¹ We will see in the next subsection that when R^i 's are continuous in ρ , whether or not there is asymptotic agreement under approximate certainty (i.e., as F_θ^i becomes more and more concentrated around a point) is determined by the tail properties of the family of distributions F_θ^i .

2.5 Agreement and Disagreement with Approximate Certainty

In this subsection, we characterize the conditions under which “approximate certainty” ensures asymptotic agreement. More specifically, we will study the behavior of asymptotic beliefs as the subjective probability distribution F_θ^i converges to a Dirac distribution and the uncertainty about the interpretation of the signals disappears. As already illustrated by Example 1, as F_θ^i converges to a Dirac distribution, each individual will become increasingly convinced that he will learn the true state. However, because asymptotic agreement is considerably more demanding than asymptotic learning, this does not guarantee that the individuals will believe that they will also agree on θ . We will demonstrate that whether or not there is asymptotic agreement in the limit depends on the family of distributions converging to certainty—in particular, on their *tail properties*. For many natural distributions, a small amount of uncertainty about the informativeness of the signals is sufficient to lead to significant differences in posteriors.

To state and prove our main result in this case, consider a *family* of subjective probability density functions $f_{\theta,m}^i$ for $i = 1, 2$, $\theta \in \{A, B\}$ and $m \in \mathbb{Z}_+$, such that as $m \rightarrow \infty$, we have that $F_{\theta,m}^i \rightarrow F_{\theta,\infty}^i$ where $F_{\theta,\infty}^i$ assigns probability 1 to $p = \hat{p}^i$ for some $\hat{p}^i \in (1/2, 1)$. Naturally, there are many different ways in which a family of subjective probability distributions may converge to such a limiting distribution. Both for tractability and to make the analysis more concrete, we focus on families of subjective probability distributions $\{f_{\theta,m}^i\}$ parameterized by a *determining* density function f . We impose the following conditions on f :

- (i) f is symmetric around zero;
- (ii) there exists $\bar{x} < \infty$ such that $f(x)$ is decreasing for all $x \geq \bar{x}$;

¹¹Nevertheless, it is also *not* the case that asymptotic agreement under approximate certainty requires the support of the distribution of each F_θ^i to converge to a set as in Theorem 2 (that does not assign positive probability to $p_\theta^i < 1/2$). See Theorem 5 below.

(iii)

$$\tilde{R}(x, y) \equiv \lim_{m \rightarrow \infty} \frac{f(mx)}{f(my)} \quad (10)$$

exists in $[0, \infty]$ at all $(x, y) \in \mathbb{R}_+^2$.¹²

Conditions (i) and (ii) are natural and serve to simplify the notation. Condition (iii) introduces the function $\tilde{R}(x, y)$, which will arise naturally in the study of asymptotic agreement and has a natural meaning in asymptotic statistics (see Definitions 1 and 2 below).

In order to vary the amount of uncertainty, we consider mappings of the form $x \mapsto (x - y)/m$, which scale down the real line around y by the factor $1/m$. The family of subjective densities for individuals' beliefs about p_A and p_B , $\{f_{\theta, m}^i\}$, will be determined by f and the transformation $x \mapsto (x - \hat{p}^i)/m$.¹³ In particular, we consider the following family of densities

$$f_{\theta, m}^i(p) = c^i(m) f(m(p - \hat{p}^i)) \quad (11)$$

for each θ and i where $c^i(m) \equiv 1/\int_0^1 f(m(p - \hat{p}^i)) dp$ is a correction factor to ensure that $f_{\theta, m}^i$ is a proper probability density function on $[0, 1]$ for each m . We also define $\phi_{\infty, m}^i \equiv \lim_{n \rightarrow \infty} \phi_{n, m}^i(s)$ as the limiting posterior distribution of individual i when he believes that the probability density of signals is $f_{\theta, m}^i$. In this family of subjective densities, the uncertainty about p_A is scaled down by $1/m$, and $f_{\theta, m}^i$ converges to unit mass at \hat{p}^i as $m \rightarrow \infty$, so that individual i becomes sure about the informativeness of the signals in the limit. In other words, as $m \rightarrow \infty$, this family of subjective probability distributions leads to approximate certainty (and ensures asymptotic learning; see the proof of Part 1 of Theorem 5).

The next theorem characterizes the class of determining functions f for which the resulting family of the subjective densities $\{f_{\theta, m}^i\}$ leads to asymptotic agreement under approximate certainty.

Theorem 5 (*Asymptotic Agreement and Disagreement Under Approximate Certainty*) *Suppose that Assumption 1 holds. For each $i = 1, 2$, consider the family of subjective densities $\{f_{\theta, m}^i\}$ defined in (11) for some $\hat{p}^i > 1/2$, with f satisfying conditions (i)-(iii) above. Suppose that $f(mx)/f(my)$ uniformly converges to $\tilde{R}(x, y)$ over a neighborhood of $(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|)$. Then,*

¹²Convergence will be uniform in most cases in view of the results discussed following Definition 1 below (and of Egorov's Theorem, which links pointwise convergence of a family of functions to a limiting function to uniform convergence, see, for example, Billingsley, 1995, Section 13).

¹³This formulation assumes that \hat{p}_A^i and \hat{p}_B^i are equal. We can easily assume these to be different, but do not introduce this generality here to simplify the exposition. Theorem 8 allows for such differences in the context of the more general model with multiple states and multiple signals.

1. $\lim_{m \rightarrow \infty} (\phi_{\infty, m}^i(\hat{p}^i) - \phi_{\infty, m}^j(\hat{p}^i)) = 0$ if and only if $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) = 0$.
2. Suppose that $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) = 0$. Then for every $\epsilon > 0$ and $\delta > 0$, there exists $\bar{m} \in \mathbb{Z}_+$ such that

$$\Pr^i \left(\lim_{n \rightarrow \infty} |\phi_{n, m}^1(s) - \phi_{n, m}^2(s)| > \epsilon \right) < \delta \quad (\forall m > \bar{m}, i = 1, 2).$$

3. Suppose that $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) \neq 0$. Then there exists $\epsilon > 0$ such that for each $\delta > 0$, there exists $\bar{m} \in \mathbb{Z}_+$ such that:

$$\Pr^i \left(\lim_{n \rightarrow \infty} |\phi_{n, m}^1(s) - \phi_{n, m}^2(s)| > \epsilon \right) > 1 - \delta \quad (\forall m > \bar{m}, i = 1, 2).$$

Proof. (Proof of Part 1) Let $R_m^i(\rho)$ be the asymptotic likelihood ratio as defined in (7) associated with subjective density $f_{\theta, m}^i$. One can easily check that $\lim_{m \rightarrow \infty} R_m^i(\hat{p}^i) = 0$. Hence, by (8), $\lim_{m \rightarrow \infty} (\phi_{\infty, m}^i(\hat{p}^i) - \phi_{\infty, m}^j(\hat{p}^i)) = 0$ if and only if $\lim_{m \rightarrow \infty} R_m^j(\hat{p}^i) = 0$. By definition, we have:

$$\begin{aligned} \lim_{m \rightarrow \infty} R_m^j(\hat{p}^i) &= \lim_{m \rightarrow \infty} \frac{f(m(1 - \hat{p}^1 - \hat{p}^2))}{f(m(\hat{p}^1 - \hat{p}^2))} \\ &= \tilde{R}(1 - \hat{p}^1 - \hat{p}^2, \hat{p}^1 - \hat{p}^2) \\ &= \tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|), \end{aligned}$$

where the last equality follows by condition (i), the symmetry of the function f . This establishes that $\lim_{m \rightarrow \infty} R_m^i(\hat{p}^i) = 0$ (and thus $\lim_{m \rightarrow \infty} (\phi_{\infty, m}^i(\hat{p}^i) - \phi_{\infty, m}^j(\hat{p}^i)) = 0$) if and only if $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) = 0$.

(Proof of Part 2) Take any $\epsilon > 0$ and $\delta > 0$, and assume that $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) = 0$. By Lemma 1, there exists $\epsilon' > 0$ such that $\phi_{\infty, m}^i(\rho(s)) > 1 - \epsilon$ whenever $R^i(\rho(s)) < \epsilon'$. There also exists x_0 such that

$$\Pr^i(\rho(s) \in (\hat{p}^i - x_0/m, \hat{p}^i + x_0/m) | \theta = A) = \int_{-x_0}^{x_0} f(x) dx > 1 - \delta. \quad (12)$$

Let $\kappa = \min_{x \in [-x_0, x_0]} f(x) > 0$. Since f monotonically decreases to zero in the tails (see (ii) above), there exists x_1 such that $f(x) < \epsilon' \kappa$ whenever $|x| > |x_1|$. Let $m_1 = (x_0 + x_1) / (2\hat{p}^i - 1) > 0$. Then, for any $m > m_1$ and $\rho(s) \in (\hat{p}^i - x_0/m, \hat{p}^i + x_0/m)$, we have $|\rho(s) - 1 + \hat{p}^i| > x_1/m$, and hence

$$R_m^i(\rho(s)) = \frac{f(m(\rho(s) + \hat{p}^i - 1))}{f(m(\rho(s) - \hat{p}^i))} < \frac{\epsilon' \kappa}{\kappa} = \epsilon'.$$

Therefore, for all $m > m_1$ and $\rho(s) \in (\hat{p}^i - x_0/m, \hat{p}^i + x_0/m)$, we have that

$$\phi_{\infty, m}^i(\rho(s)) > 1 - \epsilon. \quad (13)$$

Again, by Lemma 1, there exists $\epsilon'' > 0$ such that $\phi_{\infty,m}^j(\rho(s)) > 1 - \epsilon$ whenever $R_m^j(\rho(s)) < \epsilon''$. Now, for each $\rho(s)$,

$$\lim_{m \rightarrow \infty} R_m^j(\rho(s)) = \tilde{R}(\rho(s) + \hat{p}^j - 1, |\rho(s) - \hat{p}^j|). \quad (14)$$

Moreover, by the uniform convergence assumption, there exists $\eta > 0$ such that $R_m^j(\rho(s))$ uniformly converges to $\tilde{R}(\rho(s) + \hat{p}^j - 1, |\rho(s) - \hat{p}^j|)$ on $(\hat{p}^i - \eta, \hat{p}^i + \eta)$ and

$$\tilde{R}(\rho(s) + \hat{p}^j - 1, |\rho(s) - \hat{p}^j|) < \epsilon''/2$$

for each $\rho(s)$ in $(\hat{p}^i - \eta, \hat{p}^i + \eta)$. Moreover, uniform convergence also implies that \tilde{R} is continuous at $(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|)$ (and in this part of the proof, by hypothesis, it takes the value 0). Hence, there exists $m_2 < \infty$ such that for all $m > m_2$ and $\rho(s) \in (\hat{p}^i - \eta, \hat{p}^i + \eta)$,

$$R_m^j(\rho(s)) < \tilde{R}(\rho(s) + \hat{p}^j - 1, |\rho(s) - \hat{p}^j|) + \epsilon''/2 < \epsilon''.$$

Therefore, for all $m > m_2$ and $\rho(s) \in (\hat{p}^i - \eta, \hat{p}^i + \eta)$, we have

$$\phi_{\infty,m}^j(\rho(s)) > 1 - \epsilon. \quad (15)$$

Set $\bar{m} \equiv \max\{m_1, m_2, \eta/x_0\}$. Then, by (13) and (15), for any $m > \bar{m}$ and $\rho(s) \in (\hat{p}^i - x_0/m, \hat{p}^i + x_0/m)$, we have $|\phi_{\infty,m}^i(\rho(s)) - \phi_{\infty,m}^j(\rho(s))| < \epsilon$. Then, (12) implies that $\Pr^i(|\phi_{\infty,m}^i(\rho(s)) - \phi_{\infty,m}^j(\rho(s))| < \epsilon | \theta = A) > 1 - \delta$. By the symmetry of A and B , this establishes that $\Pr^i(|\phi_{\infty,m}^i(\rho(s)) - \phi_{\infty,m}^j(\rho(s))| < \epsilon) > 1 - \delta$ for $m > \bar{m}$.

(Proof of Part 3) Since $\lim_{m \rightarrow \infty} R_m^j(\hat{p}^i) = \tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|)$ is assumed to be strictly positive, $\lim_{m \rightarrow \infty} \phi_{\infty,m}^j(\hat{p}^i) < 1$. We set $\epsilon = (1 - \lim_{m \rightarrow \infty} \phi_{\infty,m}^j(\hat{p}^i))/2$ and use similar arguments to those in the proof of Part 2 to obtain the desired conclusion. ■

Theorem 5 provides a complete characterization of the conditions under which approximate certainty will lead to asymptotic agreement. In particular, it shows that, while approximate certainty ensures asymptotic learning, it may not be sufficient to guarantee asymptotic agreement. This contrasts with the result in Theorems 1 that there will always be asymptotic agreement under full certainty. Theorem 5, instead, shows that even a small amount of uncertainty may be sufficient to cause disagreement between the individuals.

The first part of the theorem provides a simple condition on the tail of the distribution f that determines whether the asymptotic difference between the posteriors is small under approximate uncertainty. This condition can be expressed as:

$$\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) \equiv \lim_{m \rightarrow \infty} \frac{f(m(\hat{p}^1 + \hat{p}^2 - 1))}{f(m(\hat{p}^1 - \hat{p}^2))} = 0. \quad (16)$$

The theorem shows that if this condition is satisfied, then as uncertainty about the informativeness of the signals disappears the difference between the posteriors of the two individuals will become negligible. Notice that condition (16) is symmetric and does not depend on i .

Intuitively, condition (16) is related to the beliefs of one individual on whether the other individual will learn. Under approximate certainty, we always have that $\lim_{m \rightarrow \infty} R_m^i(\hat{p}^i) = 0$, so that each agent believes that he will learn the value of θ with probability 1. Asymptotic agreement (or lack thereof) depends on whether he also believes the other individual will learn the value of θ . When $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) = 0$, an individual who expects a limiting frequency of \hat{p}^2 in the asymptotic distribution will still learn the true state when the limiting frequency is \hat{p}^1 . Therefore, individual 1, who is almost certain that the limiting frequency will be \hat{p}^1 , still believes that individual 2 will reach the same inference as himself. In contrast, when $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) \neq 0$, individual 1 is still certain that limiting frequency of signals will be \hat{p}^1 and thus expects to learn himself. However, he understands that, when $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) \neq 0$, an individual who expects a limiting frequency of \hat{p}^2 will fail to learn the true state when limiting frequency happens to be \hat{p}^1 . Since he is almost certain that the limiting frequency will be \hat{p}^1 (or $1 - \hat{p}^1$), he expects the other agent not to learn the truth and thus he expects the disagreement between them to persist asymptotically.

Parts 2 and 3 of the theorem then exploit this result and the continuity of \tilde{R} to show that the individuals will attach probability arbitrarily close to 1 to the event that the asymptotic difference between their beliefs will disappear when (16) holds, and they will attach probability 1 to asymptotic disagreement when (16) fails to hold. Thus the behavior of asymptotic beliefs under approximate certainty are completely determined by condition (16).

Theorem 5 establishes that whether or not there will be asymptotic agreement depends on whether $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|)$ is equal to 0. We next investigate what this condition means for determining distributions f . Clearly, this will depend on the tail behavior of f , which, in turn, determines the behavior of the family of subjective densities $\{f_{\theta, m}^i\}$. Suppose $x \equiv \hat{p}^1 + \hat{p}^2 - 1 > \hat{p}^1 - \hat{p}^2 \equiv y > 0$. Then, condition (16) can be expressed as

$$\lim_{m \rightarrow \infty} \frac{f(mx)}{f(my)} = 0.$$

This condition holds for distributions with exponential tails, such as the exponential or the normal distributions. On the other hand, it fails for distributions with polynomial tails. For example, consider the Pareto distribution, where $f(x)$ is proportional to $|x|^{-\alpha}$ for some $\alpha > 1$.




Figure 2: $\lim_{n \rightarrow \infty} \phi_n^i(s)$ for Pareto distribution as a function of $\rho(s)$ [$\alpha = 2$, $\hat{p}^i = 3/4$.]

Then, for each m ,

$$\frac{f(mx)}{f(my)} = \left(\frac{x}{y}\right)^{-\alpha} > 0.$$

This implies that for the Pareto distribution, individuals' beliefs will fail to converge even when there is a negligible amount of uncertainty. In fact, for this distribution, the asymptotic beliefs will be independent of m (since R_m^i does not depend on m). If we take $\pi^1 = \pi^2 = 1/2$, then the asymptotic posterior probability of $\theta = A$ according to i is

$$\phi_{\infty,m}^i(\rho(s)) = \frac{(\rho(s) - \hat{p}^i)^{-\alpha}}{(\rho(s) - \hat{p}^i)^{-\alpha} + (\rho(s) + \hat{p}^i - 1)^{-\alpha}}$$

for any m .

As illustrated in Figure 2, in this case $\phi_{\infty,m}^i$ is not monotone (in fact, the discussion in the previous subsection explained why it had to be non-monotone for asymptotic agreement to breakdown). To see the magnitude of asymptotic disagreement, consider $\rho(s) \cong \hat{p}^i$. In that case, $\phi_{\infty,m}^i(\rho(s))$ is approximately 1, and $\phi_{\infty,m}^j(\rho(s))$ is approximately $y^{-\alpha} / (x^{-\alpha} + y^{-\alpha})$. Hence, both individuals believe that the difference between their asymptotic posteriors will be

$$|\phi_{\infty,m}^1 - \phi_{\infty,m}^2| \cong \frac{x^{-\alpha}}{x^{-\alpha} + y^{-\alpha}}.$$

This asymptotic difference is increasing with the difference $y \equiv \hat{p}^1 - \hat{p}^2$, which corresponds to the difference in the individuals' views on which frequencies of signals are most likely. It is also clear from this expression that this asymptotic difference will converge to zero as $y \rightarrow 0$ (i.e., as $\hat{p}^1 \rightarrow \hat{p}^2$). This last statement is indeed generally true when \tilde{R} is continuous and implies that differences in the subjective probability distributions between the two individuals put limits on their potential asymptotic disagreement:

Proposition 1 (Limits to Asymptotic Disagreement) *In Theorem 5, in addition, assume that \tilde{R} is continuous on the set $D = \{(x, y) \mid -1 \leq x \leq 1, |y| \leq \bar{y}\}$ for some $\bar{y} > 0$. Then for every $\epsilon > 0$ and $\delta > 0$, there exist $\lambda > 0$ and $\bar{m} \in (0, \infty)$ such that whenever $|\hat{p}^1 - \hat{p}^2| < \lambda$,*

$$\Pr^i \left(\lim_{n \rightarrow \infty} |\phi_{n,m}^1 - \phi_{n,m}^2| > \epsilon \right) < \delta \quad (\forall m > \bar{m}, i = 1, 2).$$

Proof. To prove this proposition, we modify the proof of Part 2 of Theorem 5 and use the notation in that proof. Since \tilde{R} is continuous on the compact set D and $\tilde{R}(x, 0) = 0$ for each x , there exists $\lambda > 0$ such that $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) < \epsilon''/4$ whenever $|\hat{p}^1 - \hat{p}^2| < \lambda$. Fix any such \hat{p}^1 and \hat{p}^2 . Then, by the uniform convergence assumption, there exists $\eta > 0$ such that $R_m^j(\rho(s))$ uniformly converges to $\tilde{R}(\rho(s) + \hat{p}^j - 1, |\rho(s) - \hat{p}^j|)$ on $(\hat{p}^i - \eta, \hat{p}^i + \eta)$ and

$$\tilde{R}(\rho(s) + \hat{p}^j - 1, |\rho(s) - \hat{p}^j|) < \epsilon''/2$$

for each $\rho(s)$ in $(\hat{p}^i - \eta, \hat{p}^i + \eta)$. The rest of the proof is identical to the proof of Part 2 in Theorem 5. ■

This proposition implies that if the individuals are almost certain about the informativeness of signals, then any significant difference in their asymptotic beliefs must be due to a significant difference in their subjective densities regarding the signal distribution (i.e., it must be the case that $|\hat{p}^1 - \hat{p}^2|$ is not small). In particular, the continuity of \tilde{R} in Proposition 1 implies that when $\hat{p}^1 = \hat{p}^2$, we must have $\tilde{R}(\hat{p}^1 + \hat{p}^2 - 1, |\hat{p}^1 - \hat{p}^2|) = 0$, and thus, from Theorem 5, there will be no significant differences in asymptotic beliefs. Notably, however, the requirement that $\hat{p}^1 = \hat{p}^2$ is rather strong. For example, Corollary 1 established that under certainty there will be asymptotic agreement for all $\hat{p}^1, \hat{p}^2 > 1/2$.

It is also worth noting that the assumption that \tilde{R} or $\lim_{m \rightarrow 0} R_m^i(\rho)$ is continuous in the relevant range is important for the results in Proposition 1. In particular, recall that Example 1 illustrated a situation in which this assumption failed and the asymptotic differences remained bounded away from zero, irrespective of the gap between \hat{p}^1 and \hat{p}^2 .

We next focus on the case where $\hat{p}^1 \neq \hat{p}^2$ and provide a further characterization of which classes of determining functions lead to asymptotic agreement under approximate certainty. We first define:

Definition 1 *A density function f has regularly-varying tails if it has unbounded support and satisfies*

$$\lim_{m \rightarrow \infty} \frac{f(mx)}{f(m)} = H(x) \in \mathbb{R}$$

for any $x > 0$.

The condition in Definition 1 that $H(x) \in \mathbb{R}$ is relatively weak, but nevertheless has important implications. In particular, it implies that $H(x) \equiv x^{-\alpha}$ for $\alpha \in (0, \infty)$. This follows from the fact that in the limit, the function $H(\cdot)$ must be a solution to the functional equation $H(x)H(y) = H(xy)$, which is only possible if $H(x) \equiv x^{-\alpha}$ for $\alpha \in (0, \infty)$.¹⁴ Moreover, Seneta (1976) shows that the convergence in Definition 1 holds locally uniformly, i.e., uniformly for x in any compact subset of $(0, \infty)$. This implies that if a density f has regularly-varying tails, then the assumptions imposed in Theorem 5 (in particular, the uniform convergence assumption) are satisfied. In fact, we have that, in this case, \tilde{R} defined in (10) is given by the same expression as for the Pareto distribution,

$$\tilde{R}(x, y) = \left(\frac{x}{y}\right)^{-\alpha},$$

and is everywhere continuous. As this expression suggests, densities with regularly-varying tails behave approximately like power functions in the tails; indeed a density $f(x)$ with regularly-varying tails can be written as $f(x) = \mathcal{L}(x)x^{-\alpha}$ for some *slowly-varying* function \mathcal{L} (with $\lim_{m \rightarrow \infty} \mathcal{L}(mx)/\mathcal{L}(m) = 1$). Many common distributions, including the Pareto, log-normal, and t-distributions, have regularly-varying densities. We also define:

Definition 2 *A density function f has rapidly-varying tails if it satisfies*

$$\lim_{m \rightarrow \infty} \frac{f(mx)}{f(m)} = x^{-\infty} \equiv \begin{cases} 0 & \text{if } x > 1 \\ 1 & \text{if } x = 1 \\ \infty & \text{if } x < 1 \end{cases}$$

for any $x > 0$.

As in Definition 1, the above convergence holds locally uniformly (uniformly in x over any compact subset that excludes 1). Examples of densities with rapidly-varying tails include the exponential and the normal densities.

From these definitions, the following corollary to Theorem 5 is immediate and links asymptotic agreement under approximate certainty to the tail behavior of the determining density function.

Corollary 2 (Tail Properties and Asymptotic Disagreement Under Approximate Certainty) *Suppose that Assumption 1 holds and $\hat{p}^1 \neq \hat{p}^2$.*

¹⁴To see this, note that since $\lim_{m \rightarrow \infty} (f(mx)/f(m)) = H(x) \in \mathbb{R}$, we have

$$H(xy) = \lim_{m \rightarrow \infty} \left(\frac{f(mxy)}{f(m)}\right) = \lim_{m \rightarrow \infty} \left(\frac{f(mxy)}{f(my)} \frac{f(my)}{f(m)}\right) = H(x)H(y).$$

See de Haan (1970) or Feller (1971).

1. Suppose that in Theorem 5 f has regularly-varying tails. Then there exists $\epsilon > 0$ such that for each $\delta > 0$, there exists $\bar{m} \in \mathbb{Z}_+$ such that

$$\Pr^i \left(\lim_{n \rightarrow \infty} |\phi_{n,m}^1(s) - \phi_{n,m}^2(s)| > \epsilon \right) > 1 - \delta \quad (\forall m > \bar{m}, i = 1, 2).$$

2. Suppose that in Theorem 5 f has rapidly-varying tails. Then for every $\epsilon > 0$ and $\delta > 0$, there exists $\bar{m} \in \mathbb{Z}_+$ such that

$$\Pr^i \left(\lim_{n \rightarrow \infty} |\phi_{n,m}^1(s) - \phi_{n,m}^2(s)| > \epsilon \right) < \delta \quad (\forall m > \bar{m}, i = 1, 2).$$

This corollary therefore implies that whether there will be asymptotic agreement depends on whether the family of subjective densities converging to “certainty” has regularly or rapidly-varying tails (provided that $\hat{p}^1 \neq \hat{p}^2$).

Returning to the intuition above, Corollary 2 and the previous definitions make it clear that the failure of asymptotic agreement is related to disagreement between the two individuals about limiting frequencies, i.e., $\hat{p}^1 \neq \hat{p}^2$, together with sufficiently thick tails of the subjective probability distribution so that an individual who expects \hat{p}^2 should have sufficient uncertainty when confronted with a limiting frequency of \hat{p}^1 . Along the lines of the intuition given there, this is sufficient for both individuals to believe that they will learn the true value of θ themselves, but that the other individual will fail to do so. Rapidly-varying tails imply that individuals become relatively certain of their model of the world and thus when individual i observes a limiting frequency ρ close to but different from \hat{p}^i , he will interpret this as being driven by sampling variation and attach a high probability to $\theta = A$. This will guarantee asymptotic agreement between the two individuals. In contrast, with regularly-varying tails, even under approximate certainty, limiting frequencies different from \hat{p}^i will be interpreted not as a sampling variation, but as potential evidence for $\theta = B$, preventing asymptotic agreement.

3 Generalizations

The previous section provided our main results in an environment with two states and two signals. In this section, we show that our main results generalize to an environment with $K \geq 2$ states and $L \geq K$ signals. The main results parallel those of Section 2 and all the proofs for this section are contained in the Appendix. To economize on space, we do not provide analogs of Theorems 1 and 2, since these require additional notation.

To generalize our results to this environment, let $\theta \in \Theta$, where $\Theta \equiv \{A^1, \dots, A^K\}$ is a set containing $K \geq 2$ distinct elements. We refer to a generic element of the set by A^k . Similarly,

let $s_t \in \{a^1, \dots, a^L\}$, with $L \geq K$ signal values. As before, define $s \equiv \{s_t\}_{t=1}^\infty$, and for each $l = 1, \dots, L$, let

$$r_n^l(s) \equiv \#\{t \leq n | s_t = a^l\}$$

be the number of times the signal $s_t = a^l$ out of first n signals. Once again, the strong law of large numbers implies that, according to both individuals, for each $l = 1, \dots, L$, $r_n^l(s)/n$ almost surely converges to some $\rho^l(s) \in [0, 1]$ with $\sum_{l=1}^L \rho^l(s) = 1$. Define $\rho(s) \in \Delta(L)$ as the vector $\rho(s) \equiv (\rho^1(s), \dots, \rho^L(s))$, where $\Delta(L) \equiv \{p = (p^1, \dots, p^L) \in [0, 1]^L : \sum_{l=1}^L p^l = 1\}$, and let the set \bar{S} be

$$\bar{S} \equiv \left\{ s \in S : \lim_{n \rightarrow \infty} r_n^l(s)/n \text{ exists for each } l = 1, \dots, L \right\}. \quad (17)$$

With analogy to the two-state-two-signal model in Section 2, let $\pi_k^i > 0$ be the prior probability individual i assigns to $\theta = A^k$, $\pi^i \equiv (\pi_1^i, \dots, \pi_K^i)$, and p_θ^l be the frequency of observing signal $s = a^l$ when the true state is θ . When players are certain about p_θ^l 's as in usual models, immediate generalizations of Theorems 1 and 2 apply. With analogy to before, we define F_θ^i as the *joint subjective probability distribution* of conditional frequencies $p \equiv (p_\theta^1, \dots, p_\theta^L)$ according to individual i . Since our focus is learning under uncertainty, we impose an assumption similar to Assumption 1.

Assumption 2 *For each i and θ , the distribution F_θ^i over $\Delta(L)$ has a continuous, non-zero and finite density f_θ^i over $\Delta(L)$.*

This assumption can be weakened along the lines discussed in Remark 2 above.

We also define $\phi_{k,n}^i(s) \equiv \Pr^i(\theta = A^k | \{s_t\}_{t=0}^n)$ for each $k = 1, \dots, K$ as the posterior probability that $\theta = A^k$ after observing the sequence of signals $\{s_t\}_{t=0}^n$, and

$$\phi_{k,\infty}^i(\rho(s)) \equiv \lim_{n \rightarrow \infty} \phi_{k,n}^i(s).$$

Given this structure, it is straightforward to generalize the results in Section 2. Let us now define the transformation $T_k : \mathbb{R}_+^K \rightarrow \mathbb{R}_+^{K-1}$, such that

$$T_k(x) = \left(\frac{x_{k'}}{x_k}; k' \in \{1, \dots, K\} \setminus k \right).$$

Here $T_k(x)$ is taken as a column vector. This transformation will play a useful role in the theorems and the proofs. In particular, this transformation will be applied to the vector π^i of priors to determine the ratio of priors assigned the different states by individual i . Let us also define the norm $\|x\| = \max_l |x|^l$ for $x = (x^1, \dots, x^L) \in \mathbb{R}^L$.

The next lemma generalizes Lemma 1:

Lemma 3 *Suppose Assumption 2 holds. Then for all $s \in \bar{S}$,*

$$\phi_{k,\infty}^i(\rho(s)) = \frac{1}{1 + \frac{\sum_{k' \neq k} \pi_{k'}^i f_{A^{k'}}^i(\rho(s))}{\pi_k^i f_{A^k}^i(\rho(s))}}. \quad (18)$$

Our first theorem in this section parallels Theorem 3 and shows that under Assumption 2 there will be lack of asymptotic learning, and under a relatively weak additional condition, there will also asymptotic disagreement.

Theorem 6 (Generalized Lack of Asymptotic Learning and Agreement) *Suppose Assumption 2 holds for $i = 1, 2$, then for each $k = 1, \dots, K$, and for each $i = 1, 2$,*

1. $\Pr^i(\phi_{k,\infty}^i(\rho(s)) \neq 1 | \theta = A^k) = 1$, and
2. $\Pr^i(|\phi_{k,\infty}^1(\rho(s)) - \phi_{k,\infty}^2(\rho(s))| \neq 0) = 1$ whenever $\Pr^i((T_k(\pi^1) - T_k(\pi^2))' T_k(f^i(\rho(s)) = 0) = 0$ and $F_\theta^1 = F_\theta^2$ for each $\theta \in \Theta$.

The additional condition in part 2 of Theorem 6, that $\Pr^i((T_k(\pi^1) - T_k(\pi^2))' T_k(f^i(\rho(s)) = 0) = 0$, plays the role of differences in priors in Theorem 3 (here “ $'$ ” denotes the transpose of the vector in question). In particular, if this condition did not hold, then at some $\rho(s)$, the relative asymptotic likelihood of some states could be the same according to two individuals with different priors and they would interpret at least some sequences of signals in a similar manner and achieve asymptotic agreement. It is important to note that the condition that $\Pr^i((T_k(\pi^1) - T_k(\pi^2))' T_k(f^i(\rho(s)) = 0) = 0$ is relatively weak and holds generically—i.e., if it did not hold, a small perturbation of π^1 or π^2 would restore it.¹⁵ The Part 2 of Theorem 6 therefore implies that asymptotic disagreement occurs *generically*.

The next theorem shows that small differences in priors can again widen after observing the same sequence of signals.

Theorem 7 (Generalized Divergence of Opinions) *Suppose Assumption 2 holds and also assume that $\mathbf{1}'(T_k((f_\theta^1(\rho))_{\theta \in \Theta}) - T_k((f_\theta^2(\rho))_{\theta \in \Theta})) \neq 0$ for each $\rho \in [0, 1]$, each $k =$*

¹⁵More formally, the set of solutions $\mathcal{S} \equiv \{(\pi^1, \pi^2, \rho) \in \Delta(L)^2 : (T_k(\pi^1) - T_k(\pi^2))' T_k(f^i(\rho)) = 0\}$ has Lebesgue measure 0. This is a consequence of the Preimage Theorem and Sard’s Theorem in differential topology (see, for example, Guillemin and Pollack, 1974, pp. 21 and 39). The Preimage Theorem implies that if y is a regular value of a map $f : X \rightarrow Y$, then $f^{-1}(y)$ is a submanifold of X with dimension equal to $\dim X - \dim Y$. In our context, this implies that if 0 is a regular value of the map $(T_k(\pi^1) - T_k(\pi^2))' T_k(f^i(\rho))$, then the set \mathcal{S} is a two dimensional submanifold of $\Delta(L)^3$ and thus has Lebesgue measure 0. Sard’s theorem implies that 0 is generically a regular value.

$1, \dots, K$ (where $\mathbf{1} \equiv (1, \dots, 1)'$). Then, there exists an open set of prior vectors π^1 and π^2 , such that

$$|\phi_{k,\infty}^1(\rho(s)) - \phi_{k,\infty}^2(\rho(s))| > |\pi_k^1 - \pi_k^2| \text{ for each } k = 1, \dots, K \text{ and } s \in \bar{S}$$

and

$$\Pr^i (|\phi_{k,\infty}^1(\rho(s)) - \phi_{k,\infty}^2(\rho(s))| > |\pi_k^1 - \pi_k^2|) = 1 \text{ for each } k = 1, \dots, K.$$

The condition $\mathbf{1}' \left(T_k \left((f_\theta^1(\rho))_{\theta \in \Theta} \right) - T_k \left((f_\theta^2(\rho))_{\theta \in \Theta} \right) \right) \neq 0$ is similar to the additional condition in part 2 of Theorem 6, and as with that condition, it is relatively weak and holds generically. Finally, the following theorem generalizes Theorem 5. The appropriate construction of the families of probability densities is also provided in the theorem.

Theorem 8 (Generalized Asymptotic Agreement and Disagreement Under Approximate Certainty) Suppose that Assumption 2 holds. For each $\theta \in \Theta$ and $m \in \mathbb{Z}_+$, define the subjective density $f_{\theta,m}^i$ by

$$f_{\theta,m}^i(p) = c(i, \theta, m) f(m(p - \hat{p}(i, \theta))) \quad (19)$$

where $c(i, \theta, m) \equiv 1 / \int_{p \in \Delta(L)} f(m(p - \hat{p}(i, \theta))) dp$, $\hat{p}(i, \theta) \in \Delta(L)$ with $\hat{p}(i, \theta) \neq \hat{p}(i, \theta')$ whenever $\theta \neq \theta'$, and $f: \mathbb{R}^L \rightarrow \mathbb{R}$ is a positive, continuous probability density function that satisfies the following conditions:

(i) $\lim_{h \rightarrow \infty} \max_{\{x: \|x\| \geq h\}} f(x) = 0,$

(ii)

$$\tilde{R}(x, y) \equiv \lim_{m \rightarrow \infty} \frac{f(mx)}{f(my)} \quad (20)$$

exists at all x, y , and

(iii) convergence in (20) holds uniformly over a neighborhood of each

$$(\hat{p}(i, \theta) - \hat{p}(j, \theta'), \hat{p}(i, \theta) - \hat{p}(j, \theta)).$$

Also let $\phi_{k,\infty,m}^i(\rho(s)) \equiv \lim_{n \rightarrow \infty} \phi_{k,n,m}^i(s)$ be the asymptotic posterior of individual i with subjective density $f_{\theta,m}^i$. Then,

1. $\lim_{m \rightarrow \infty} \left(\phi_{k,\infty,m}^i(\hat{p}(i, A^k)) - \phi_{k,\infty,m}^j(\hat{p}(i, A^k)) \right) = 0$ if and only if $\tilde{R}(\hat{p}(i, A^k) - \hat{p}(j, A^{k'}), \hat{p}(i, A^k) - \hat{p}(j, A^k)) = 0$ for each $k' \neq k$.

2. Suppose that $\tilde{R}(\hat{p}(i, \theta) - \hat{p}(j, \theta'), \hat{p}(i, \theta) - \hat{p}(j, \theta)) = 0$ for each distinct θ and θ' . Then for every $\epsilon > 0$ and $\delta > 0$, there exists $\bar{m} \in \mathbb{Z}_+$ such that

$$\Pr^i (\|\phi_{\infty, m}^1(s) - \phi_{\infty, m}^2(s)\| > \epsilon) < \delta \quad (\forall m > \bar{m}, i = 1, 2).$$

3. Suppose that $\tilde{R}(\hat{p}(i, \theta) - \hat{p}(j, \theta'), \hat{p}(i, \theta) - \hat{p}(j, \theta)) \neq 0$ for each distinct θ and θ' . Then there exists $\epsilon > 0$ such that for each $\delta > 0$, there exists $\bar{m} \in \mathbb{Z}_+$ such that

$$\Pr^i (\|\phi_{\infty, m}^1(s) - \phi_{\infty, m}^2(s)\| > \epsilon) > 1 - \delta \quad (\forall m > \bar{m}, i = 1, 2).$$

These theorems therefore show that the results about lack of asymptotic learning and asymptotic agreement derived in the previous section do not depend on the assumption that there are only two states and binary signals. It is also straightforward to generalize Proposition 1 and Corollary 2 to the case with multiple states and signals; we omit this to avoid repetition.

The results in this section are stated for the case in which both the number of signal values and states are finite. They can also be generalized to the case of a continuum of signal values and states, but this introduces a range of technical issues that are not central to our focus here.

4 Applications

In this section we discuss a number of applications of the results derived so far. The applications are chosen to show various different economic consequences from learning and disagreement under uncertainty. Throughout, we strive to choose the simplest examples. The first example illustrates how learning under uncertainty can overturn some simple insights from basic game theory. The second example shows how such learning can act as an equilibrium selection device as in Carlsson and van Damme (1993). The third example is the most substantial application and shows how learning under uncertainty affects speculative asset trading. The fourth example illustrates how learning under uncertainty can affect the timing of agreement in bargaining. Finally, the last example shows how a special case of our model of learning under uncertainty can arise when there is information transmission by a potentially biased media outlet.¹⁶

¹⁶In this section, except for the example on equilibrium selection and the last example of the game of belief manipulation, we study complete-information games with possibly non-common priors. Formally, information and belief structure in these games can be described as follows. Fix the state space $\Omega = \Theta \times \bar{S}$, and for each $n < \infty$, consider the information partition $I^n = \{I^n(s) = \{(\theta, s') \mid s'_t = s_t \forall t \leq n\} \mid s \in \bar{S}\}$ that is common for both players. For $n = \infty$, we introduce the common information partition $I^\infty = \{I^\infty(s) = \Theta \times \{s\} \mid s \in \bar{S}\}$.

4.1 Value of Information in Common-Interest Games

Consider a common-interest game in which the players have identical payoff functions. Typically in common interest games information is valuable in the sense that with more information about underlying parameters, the value of the game in the best equilibrium will be higher. We would therefore expect players to collect or at least wait for the arrival of additional information before playing such games. We now show that when there is learning under uncertainty, additional information can be harmful in common-interest games, and thus the agents may prefer to play the game *before* additional information arrives.

To illustrate these issues, consider the payoff matrix

	α	β
α	$2\theta, 2\theta$	$1/2, 1/2$
β	$1/2, 1/2$	$1 - \theta, 1 - \theta$

where $\theta \in \{0, 1\}$, and the agents have a common prior on θ according to which probability of $\theta = 1$ is $\pi \in (1/2, 1)$. When there is no information, α strictly dominates β (since the expected value of the payoff from (α, α) is strictly greater than $1/2$ and the expected value of the payoff from (β, β) is strictly less than $1/2$). In the dominant-strategy equilibrium, (α, α) , each player receives 2θ with probability π , thus achieve an expected payoff of $2\pi > 1$.

First, consider the implications of learning under certainty. Suppose that the agents are allowed to observe an infinite sequence of signals $s = \{s_t\}_{t=1}^{\infty}$, where each agent believes that $\Pr^i(s_t = \theta | \theta) = p^i > 1/2$. Theorem 1 then implies that after observing the sequence of signals, the agents will learn θ . If the frequency $\rho(s)$ of signal with $s_t = 1$ is greater than $1/2$, they will learn that $\theta = 1$; otherwise they will learn that $\theta = 0$. If $\rho(s) < 1/2$, β strictly dominates α , and hence (β, β) is the dominant strategy equilibrium. If $\rho(s) > 1/2$, α strictly dominates β and (α, α) is the dominant strategy equilibrium. Consequently, when they learn under certainty before playing the game, the expected payoff to each player is $2\pi + (1 - \pi) > 2\pi$. This implies that, if they have the option, the players would prefer to wait for the arrival of public information before playing the game.

Let us next turn to learning under uncertainty. In particular, suppose that the agents do not know the signal distribution and their subjective densities are similar to those in Example

At each $I^n(s)$, player $i = 1, 2$ assigns probability $\phi_n^i(s)$ to the state $\theta = A$ and probability $1 - \phi_n^i(s)$ to the state $\theta = B$. Since the players have a common partition at each s and n , their beliefs are common knowledge. Notice that, under certainty, $\phi_\infty^1(s) = \phi_\infty^2(s) \in \{0, 1\}$, so that after observing s , both players assign probability 1 to the same θ . In that case, there will be *common certainty* of θ , or loosely speaking, θ becomes “common knowledge.” This is not necessarily the case under uncertainty.

2:

$$f_{\theta}^i(p) = \begin{cases} (1 - \epsilon - \epsilon^2) / \delta & \text{if } \hat{p}^i - \delta/2 \leq p \leq \hat{p}^i + \delta/2 \\ \epsilon & \text{if } p < 1/2 \\ \epsilon^2 & \text{otherwise} \end{cases} \quad (21)$$

for each θ , where $0 < \delta < \hat{p}^1 - \hat{p}^2$ and ϵ and δ are taken to be arbitrarily small (i.e., we consider the limit where $\epsilon \rightarrow 0$ and $\delta \rightarrow 0$, or loosely speaking, where $\epsilon \cong 0$ and $\delta \cong 0$). Recall from Example 2 that when $\epsilon \cong 0$ and $\delta \cong 0$, the asymptotic posterior probability of $\theta = 1$ is

$$\phi_{\infty}^i(\rho(s)) \cong \begin{cases} 1 & \begin{aligned} &\text{if } \rho(s) < 1 - \hat{p}^i - \delta/2, \\ &\text{or } 1 - \hat{p}^i + \delta/2 < \rho(s) < 1/2, \\ &\text{or } \hat{p}^i - \delta/2 \leq \rho(s) \leq \hat{p}^i + \delta/2, \end{aligned} \\ 0 & \text{otherwise.} \end{cases}$$

As discussed above, when $\epsilon \cong 0$ and $\delta \cong 0$, each agent believes that he will learn the true value of θ , while the other agent will reach the opposite conclusion. This implies that both agents expect that one of them will have $\phi_{\infty}^i(\rho(s)) \cong 1$ while the other has $\phi_{\infty}^i(\rho(s)) \cong 0$. Consequently, the unique equilibrium will be (α, β) , giving both agents an ex ante expected payoff of $1/2$, which is strictly less than the expected payoff to playing the game before the arrival of information (which is 2π). Therefore, when there is learning under uncertainty, both agents may prefer to play the game before the arrival of public information.

4.2 Selection in Coordination Games

The initial difference in players' beliefs about the signal distribution need not be due to lack of common prior; it may be due to private information. Building on an example by Carlsson and van Damme (1993), we now illustrate that when the players are uncertain about the signal distribution, small differences in beliefs, combined with learning, may have a significant effect on the outcome of the game and may select one of the multiple equilibria of the game.

Consider a game with the payoff matrix

	I	N
I	θ, θ	$\theta - 1, 0$
N	$0, \theta - 1$	$0, 0$

where $\theta \sim \mathcal{N}(0, 1)$. The players observe an infinite sequence of public signals $s \equiv \{s_t\}_{t=0}^{\infty}$, where $s_t \in \{0, 1\}$ and

$$\Pr(s_t = 1 | \theta) = 1 / (1 + \exp(-(\theta + \eta))), \quad (22)$$

with $\eta \sim \mathcal{N}(0, 1)$. In addition, each player observes a *private* signal

$$x_i = \eta + u_i$$

where u_i is uniformly distributed on $[-\epsilon/2, \epsilon/2]$ for some small $\epsilon > 0$.

Let us define $\kappa \equiv \log(\rho(s)) - \log(1 - \rho(s))$. Equation (22) implies that after observing s , the players infer that $\theta + \eta = \kappa$. For small ϵ , conditional on x_i , η is distributed approximately uniformly on $[x_i - \epsilon/2, x_i + \epsilon/2]$ (see Carlsson and van Damme, 1993). This implies that conditional on x_i and s , θ is approximately uniformly distributed on $[\kappa - x_i - \epsilon/2, \kappa - x_i + \epsilon/2]$. Now note that with the reverse order on x_i , the game is supermodular. Therefore, there exist extremal rationalizable strategy profiles, which also constitute monotone, symmetric Bayesian Nash Equilibria. In each equilibrium, there is a cutoff value, x^* , such that the equilibrium action is I if $x_i < x^*$ and N if $x_i > x^*$. This cutoff, x^* , is defined such that player i is indifferent between the two actions, i.e.,

$$\kappa - x^* = \Pr(x_j > x^* | x_i = x^*) = 1/2 + O(\epsilon),$$

where $O(\epsilon)$ is such that $\lim_{\epsilon \rightarrow 0} O(\epsilon) = 0$. This establishes that

$$x^* = \kappa - 1/2 - O(\epsilon).$$

Therefore, when ϵ is small, the game is dominance solvable, and each player i plays I if $x_i < \kappa - 1/2$ and N if $x_i > \kappa + 1/2$.

In this game, learning under certainty has very different implications from those above. Suppose instead that the players knew the conditional signal distribution (i.e., they knew η), so that we are in a world of learning under certainty. Then after s is observed, θ would become common knowledge, and there would be multiple equilibria whenever $\theta \in (0, 1)$. This example therefore illustrates how learning under uncertainty can lead to the selection of one of the equilibria in a coordination game.

4.3 A Simple Model of Asset Trade

One of the most interesting applications of the ideas developed here is to models of asset trading. Models of assets trading with different priors have been studied by, among others, Harrison and Kreps (1978) and Morris (1996). These works assume different priors about the dividend process and allow for learning under certainty. They establish the possibility of “speculative asset trading”.¹⁷ We now investigate the implications of learning under uncertainty for the pattern of speculative asset trading.

¹⁷Fudenberg and Levine (2005) show that learning by traders may be sufficient to prevent asset trading even when it does not lead to common knowledge about the gains from trade.

Consider an asset that pays 1 if the state is A and 0 if the state is B . Assume that Agent 2 owns the asset, but Agent 1 may wish to buy it. We have two dates, $\tau = 0$ and $\tau = 1$, and the agents observe a sequence of signals between these dates. For simplicity, we again take this to be an infinite sequence $s \equiv \{s_t\}_{t=1}^{\infty}$. We also simplify this example by assuming that Agent 1 has all the bargaining power: at either date, if he wants to buy the asset, Agent 1 makes a take-it-or-leave-it price offer P_{τ} , and trade occurs at price P_{τ} if Agent 2 accepts the offer. Assume also that $\pi^1 > \pi^2$, so that Agent 1 is more optimistic. This assumption ensures that Agent 1 would like to purchase the asset. We are interested in subgame-perfect equilibrium of this game.¹⁸

Let us start with the case in which there is learning under certainty. Suppose that each agent is certain that $p_A = p_B = p^i$ for some number $p^i > 1/2$. In that case, from Theorem 1, both agents recognize at $\tau = 0$ that at $\tau = 1$, for each $\rho(s)$, the value of the asset will be the same for both of them: it will be worth 1 if $\rho(s) > 1/2$ and 0 if $\rho(s) < 1/2$. Hence, at $\tau = 1$ the agents will be indifferent between trading the asset (at price $P_1 = \phi_{\infty}^1(\rho(s)) = \phi_{\infty}^2(\rho(s))$) at each history $\rho(s)$. Therefore, if trade does not occur at $\tau = 0$, the continuation value of Agent 1 is 0, and the continuation value of Agent 2 is π^2 . If they trade at price P_0 , then the continuation value of agents 1 and 2 will be $\pi^1 - P_0$ and P_0 , respectively. This implies that at date 0, Agent 2 accepts an offer if and only if $P_0 \geq \pi^2$. Since $\pi^1 > \pi^2$, Agent 1 is happy to offer the price $P_0 = \pi^2$ at date $\tau = 0$ and trade takes place. Therefore, with learning under certainty, there will be immediate trade at $\tau = 0$.

We next turn to the case of learning under uncertainty and suppose that the agents do not know p_A and p_B . In contrast to the case of learning under certainty, the agents now have an incentive to delay trading. To illustrate this, we first consider a simple example where subjective densities are as in Example 1, with $\epsilon \rightarrow 0$. Now, at date 1, if $\hat{p}^1 - \delta/2 < \rho(s) < \hat{p}^1 + \delta/2$, then the value of the asset for Agent 2 is $\phi_{\infty}^2(\rho(s)) = \pi^2$, and the value of the asset for Agent 1 is approximately 1. Hence, at such $\rho(s)$, Agent 1 buys the asset from Agent 2 at price $P_1(\rho(s)) = \pi^2$, enjoying gains from trade equal to $1 - \pi^2$. Since the equilibrium payoff of Agent 1 must be non-negative in all other contingencies, this shows that when they do not trade at date 0, his continuation value is at least

$$\pi^1 (1 - \pi^2)$$

(when $\epsilon \rightarrow 0$). The continuation value of Agent 2 must be at least π^2 , as he has the option

¹⁸We focus on a game between two agents for convenience. All the results presented in this subsection easily generalize to the competitive equilibrium of a model with many agents on each side.

of never selling his asset. Therefore, they can trade at date 0 only if the total payoff from trading, which is π^1 , exceeds the sum of these continuation values, $\pi^1 (1 - \pi^2) + \pi^2$. Since this is impossible, there will be no trade at $\tau = 0$. Instead, Agent 1 will wait for the information to buy the asset at date 1 (provided that $\rho(s)$ turns out to be in a range where he concludes that the asset pays 1).

This example exploits the general intuition discussed after Theorem 4: if the agents are uncertain about the informativeness of the signals, each agent may expect to *learn more* from the signals than the other agent. In fact, this example has the extreme feature whereby each agent believes that he will definitely learn the true state, but the other agent will fail to do so. This induces the agents to wait for the arrival of additional information before trading. This contrasts with the intuition that observation of common information should take agents towards common beliefs and make trades less likely. This intuition is correct in models of learning under certainty and is the reason why previous models have generated speculative trade at the beginning (Harrison and Kreps, 1978, and Morris, 1996). Instead, here there is delayed speculative trading.

The next result characterizes the conditions for delayed asset trading more generally:

Proposition 2 *In any subgame-perfect equilibrium, trade is delayed to $\tau = 1$ if and only if*

$$\mathbb{E}^2 [\phi_\infty^2] = \pi^2 > \mathbb{E}^1 [\min \{\phi_\infty^1, \phi_\infty^2\}].$$

That is, when $\pi^2 > \mathbb{E}^1 [\min \{\phi_\infty^1, \phi_\infty^2\}]$, Agent 1 does not buy at $\tau = 0$ and buys at $\tau = 1$ if $\phi_\infty^1(\rho(s)) > \phi_\infty^2(\rho(s))$; when $\pi^2 < \mathbb{E}^1 [\min \{\phi_\infty^1, \phi_\infty^2\}]$, Agent 1 buys at $\tau = 0$.

Proof. In any subgame-perfect equilibrium, Agent 2 is indifferent between trading and not, and hence his valuation of the asset is $\Pr^2(\theta = A | \text{Information})$. Therefore, trade at $\tau = 0$ can take place at the price $P_0 = \pi^2$, while trade at $\tau = 1$ will be at the price $P_1(\rho(s)) = \phi_\infty^2(\rho(s))$. At date 1, Agent 1 buys the asset if and only if $\phi_\infty^1(\rho(s)) \geq \phi_\infty^2(\rho(s))$, yielding the payoff of $\max\{\phi_\infty^1(\rho(s)) - \phi_\infty^2(\rho(s)), 0\}$. This implies that Agent 1 is willing to buy at $\tau = 0$ if and only if

$$\begin{aligned} \pi^1 - \pi^2 &\geq \mathbb{E}^1 [\max \{\phi_\infty^1(\rho(s)) - \phi_\infty^2(\rho(s)), 0\}] \\ &= \mathbb{E}^1 [\phi_\infty^1(\rho(s)) - \min \{\phi_\infty^1(\rho(s)), \phi_\infty^2(\rho(s))\}] \\ &= \pi^1 - \mathbb{E}^1 [\min \{\phi_\infty^1(\rho(s)), \phi_\infty^2(\rho(s))\}], \end{aligned}$$

as claimed. ■

Since $\pi^1 = \mathbb{E}^1 [\phi_\infty^1] \geq \mathbb{E}^1 [\min \{\phi_\infty^1, \phi_\infty^2\}]$, this result provides a cutoff value for the initial difference in beliefs, $\pi^1 - \pi^2$, in terms of the differences in the agents' interpretation of the signals. The cutoff value is $\mathbb{E}^1 [\max \{\phi_\infty^1(\rho(s)) - \phi_\infty^2(\rho(s)), 0\}]$. If the initial difference is lower than this value, then agents will wait until $\tau = 1$ to trade; otherwise, they will trade immediately. Consistent with the above example, delay in trading becomes more likely when the agents interpret the signals more differently, which is evident from the expression for the cutoff value. This reasoning also suggests that if $F_\theta^1 = F_\theta^2$ for each θ (so that the agents interpret the signals in a similar fashion),¹⁹ then trade should occur immediately. The next lemma shows that each agent believes that additional information will bring the other agent's expectations closer to his own and will be used to prove that $F_\theta^1 = F_\theta^2$ indeed implies immediate trading.

Lemma 4 *If $\pi^1 > \pi^2$ and $F_\theta^1 = F_\theta^2$ for each θ , then*

$$\mathbb{E}^1 [\phi_\infty^2] \geq \pi^2.$$

Proof. Recall that ex ante expectation of individual i regarding ϕ_∞^j can be written as

$$\begin{aligned} \mathbb{E}^i [\phi_\infty^j] &= \int_0^1 [\pi^i f_A^i(\rho) \phi_\infty^j(\rho) + (1 - \pi^i) f_B^i(1 - \rho) \phi_\infty^j(\rho)] d\rho \\ &= \int_0^1 \frac{\pi^i f_A(\rho) + (1 - \pi^i) f_B(1 - \rho)}{\pi^j f_A(\rho) + (1 - \pi^j) f_B(1 - \rho)} f_A(\rho) d\rho, \end{aligned} \quad (23)$$

where the first line uses the definition of ex ante expectation under the probability measure Pr^i , while the second line exploits equations (6) and (7) and the fact that since $F_\theta^1 = F_\theta^2$, $f_\theta^1(\rho) = f_\theta^2(\rho) = f_\theta(\rho)$ for all ρ . Now define

$$I(\pi) \equiv \int_0^1 \frac{\pi f_A(\rho) + (1 - \pi) f_B(1 - \rho)}{\pi^2 f_A(\rho) + (1 - \pi^2) f_B(1 - \rho)} f_A(\rho) d\rho.$$

From (23), $\mathbb{E}^1 [\phi_\infty^2] = I(\pi^1)$ and $\pi^2 = \mathbb{E}^2 [\phi_\infty^2] = I(\pi^2)$. Hence, it suffices to show that I is increasing in π . Now,

$$I'(\pi) = \int_0^1 \frac{f_A(\rho)}{\pi^2 f_A(\rho) + (1 - \pi^2) f_B(1 - \rho)} (f_A(\rho) - f_B(1 - \rho)) d\rho.$$

Moreover, $f_A(\rho) / [\pi^2 f_A(\rho) + (1 - \pi^2) f_B(1 - \rho)] \geq 1$ if and only if $f_A(\rho) \geq f_B(1 - \rho)$.

¹⁹Recall from Theorem 3 that even when $F_\theta^1 = F_\theta^2$, agents interpret signals differently because $\pi^1 \neq \pi^2$.

Hence,

$$\begin{aligned}
I'(\pi) &= \int_{f_A \geq f_B} \frac{f_A(\rho)}{\pi^2 f_A(\rho) + (1 - \pi^2) f_B(1 - \rho)} (f_A(\rho) - f_B(1 - \rho)) d\rho \\
&\quad - \int_{f_A < f_B} \frac{f_A(\rho)}{\pi^2 f_A(\rho) + (1 - \pi^2) f_B(1 - \rho)} (f_B(1 - \rho) - f_A(\rho)) d\rho \\
&\geq \int_{f_A \geq f_B} (f_A(\rho) - f_B(1 - \rho)) d\rho - \int_{f_A < f_B} (f_B(1 - \rho) - f_A(\rho)) d\rho \\
&= \int_0^1 (f_A(\rho) - f_B(1 - \rho)) d\rho = 0.
\end{aligned}$$

■

Together with the previous proposition, this lemma yields the following result establishing that delay in asset trading can only occur when subjective probability distributions differ across individuals.

Proposition 3 *If $F_\theta^1 = F_\theta^2$ for each θ , then in any subgame-perfect equilibrium, trade occurs at $\tau = 0$.*

Proof. Since $\pi^1 > \pi^2$ and $R^1 = R^2$, Lemma 1 implies that $\phi_\infty^1(\rho(s)) \geq \phi_\infty^2(\rho(s))$ for each $\rho(s)$. Then, $\mathbb{E}^1[\min\{\phi_\infty^1, \phi_\infty^2\}] = \mathbb{E}^1[\phi_\infty^2] \geq \pi^2$, where the last inequality is by Lemma 4. Therefore, by Proposition 2, Agent 1 buys at $\tau = 0$. ■

This proposition establishes that when the two agents have the same subjective probability distributions, there will be no delay in trading. However, as the example above illustrates, when $F_\theta^1 \neq F_\theta^2$, delayed speculative trading is possible. The intuition is given by Lemma 4: when agents have the same subjective probability distribution but different priors, each will believe that additional information will bring the other agent's beliefs closer to his own. This leads to early trading. However, when the agents differ in terms of their subjective probability distributions, they expect to learn more from new information (because, as discussed after Theorem 4 above, they believe that they have the “correct model of the world”). Consequently, they delay trading.

Learning under uncertainty does not necessarily lead to additional delay in economic transactions, however. Whether it does so or not depends on the effect of the extent of disagreement on the timing of economic transactions. We will next see that, in the context of bargaining, the presence of learning under uncertainty may be a force towards immediate agreement rather than delay.

4.4 Bargaining With Outside Options

Consider two agents bargaining over the division of a dollar. There are two dates, $\tau \in \{0, 1\}$, and Agent 2 has an outside option $\theta \in \{\theta_L, \theta_H\}$ that expires at the end of date 1, where $\theta_L < \theta_H < 1$ and the value of θ is initially unknown. Between the two dates, the agents observe an infinite sequence of public signals $s \equiv \{s_t\}_{t=1}^\infty$ with $s_t \in \{a_L, a_H\}$, where the signal a_L can be thought to be more likely under θ_L .

Bargaining follows a simple protocol: at each date τ , Agent 1 offers a share w_τ to Agent 2. If Agent 2 accepts the offer, the game ends, Agent 2 receives the proposal, w_τ , and Agent 1 receives the remaining $1 - w_\tau$. If Agent 2 rejects the offer, she decides whether to take her outside option, terminating the game, or wait for the next stage of the game. We assume that delay is costly, so that if negotiations continue until date $\tau = 1$, Agent 1 incurs a cost $c > 0$.

Finally, as in Yildiz (2003), the agents are assumed to be “optimistic,” in the sense that

$$y \equiv \mathbb{E}^2[\theta] - \mathbb{E}^1[\theta] > 0.$$

In other words, they differ in their expectations of θ on the outside option of Agent 2—with Agent 2 believing that her outside option is higher than Agent 1’s assessment of this outside option—and y parameterizes the extent of optimism in this game.

We assume that the game form and beliefs are common knowledge and look for the subgame-perfect equilibrium of this simple bargaining game.

By backward induction, at date $\tau = 1$, for any $\rho(s)$, the value of outside option for Agent 1 is $\mathbb{E}^2[\theta|\rho(s)] < 1$, and hence she accepts an offer w_1 if and only if $w_1 \geq \mathbb{E}^2[\theta|\rho(s)]$. Agent 2 therefore offers $w_1 = \mathbb{E}^2[\theta|\rho(s)]$. If there is no agreement at date 0, the continuation values of the two agents are:

$$V^1 = 1 - c - \mathbb{E}^1[\mathbb{E}^2[\theta|\rho(s)]] \quad \text{and} \quad V^2 = \mathbb{E}^2[\mathbb{E}^2[\theta|\rho(s)]] = \mathbb{E}^2[\theta],$$

which uses the fact that there is no cost of delay for Agent 2. Since they have 1 dollar in total, the agents will delay the agreement to date $\tau = 1$ if and only if

$$\mathbb{E}^2[\theta] - \mathbb{E}^1[\mathbb{E}^2[\theta|\rho(s)]] > c.$$

Here, $\mathbb{E}^1[\mathbb{E}^2[\theta|\rho(s)]]$ is Agent 1’s expectation about how Agent 2 will update her beliefs after observing the signals s . If Agent 1 expects that the information will reduce Agent 2’s expectation of her outside option more than the cost of waiting, then Agent 1 is willing to wait. This description makes it clear that whether there will be agreement at date $\tau = 0$ depends on Agent 1’s assessment of how Agent 2 will interpret the (public) signals.

When each agent is certain about the informativeness of the signals, they agree ex ante that they will interpret the information correctly. Consequently, as in Lemma 4 in the previous subsection, Agent 1’s Bayesian updating will indicate that the public information will reveal him to be right. Yildiz (2004) has shown that this reasoning gives Agent 1 an incentive to “wait to persuade” Agent 2 that her outside option is relatively low. More specifically, assume that each agent i is certain that $\Pr^i(s_t = \theta|\theta) = \hat{p}^i > 1/2$ for some \hat{p}^1 and \hat{p}^2 , where \hat{p}^1 and \hat{p}^2 may differ. Then, from Theorem 1, the agents agree that Agent 2 will learn her outside option, i.e., $\Pr^i(\mathbb{E}^2[\theta|\rho(s)] = \theta) = 1$ for each i . Hence, $\mathbb{E}^1[\mathbb{E}^2[\theta|\rho(s)]] = \mathbb{E}^1[\theta]$. Therefore, Agent 1 delays the agreement to date $\tau = 1$ if and only if

$$y > c,$$

i.e., if and only if the level of optimism is higher than the cost of waiting. This discussion therefore indicates that the arrival of public information can create a reason for delay in bargaining games.

We now show that when agents are uncertain about the informativeness of the signals, this motive for delay is reduced and there can be immediate agreement. Intuitively, each agent understands that the same signals will be interpreted differently by the other agent and thus expects that they are less likely to persuade the other agent. This decreases the incentives to delay agreement.

This result is illustrated starkly here, with an example where a small amount of uncertainty about the informativeness of signals removes all incentives to delay agreement. Suppose that the agents’ beliefs are again as in Example 1 with ϵ small. Now Agent 1 assigns probability more than $1 - \epsilon$ to the event that that $\rho(s)$ will be either in $[\hat{p} - \delta/2, \hat{p}^1 + \delta/2]$ or in $[1 - \hat{p} - \delta/2, 1 - \hat{p}^1 + \delta/2]$, inducing Agent 2 to stick to her prior. Hence, Agent 1 expects that Agent 2 will not update her prior by much. In particular, we have

$$\mathbb{E}^1[\mathbb{E}^2[\theta|\rho(s)]] = \mathbb{E}^2[\theta] + O(\epsilon).$$

Thus

$$\mathbb{E}^2[\theta] - \mathbb{E}^1[\mathbb{E}^2[\theta|\rho(s)]] = -O(\epsilon) < c.$$

This implies that agents will agree at the beginning of the game. Therefore, the same forces that led to delayed asset trading in the previous subsection can also induce immediate agreement in bargaining when agents are “optimistic”.

4.5 Manipulation and Uncertainty

Our final example is intended to show how the pattern of uncertainty used in the body of the paper can result from game theoretic interactions between an agent and an informed party, for example as in cheap talk games (Crawford and Sobel, 1982). Since our purpose is to illustrate this possibility, we choose the simplest environment to communicate these ideas and limit the discussion to the single agent setting. The generalization to the case with two or more agents, which would enable us to reiterate results related to asymptotic agreement, is straightforward and is omitted to economize on space.

The environment is as follows. The state of the world is $\theta \in \{0, 1\}$, and the agent starts with a prior belief $\pi \in (0, 1)$ that $\theta = 1$ at $t = 0$. At time $t = 1$, this agent has to make a decision $x \in [0, 1]$, and his payoff is $-(x - \theta)^2$. Thus the agent would like to form as accurate an expectation about θ as possible.

The other player is a media outlet, M , which observes a large (infinite) number of signals $s' \equiv \{s'_t\}_{t=1}^\infty$ with $s'_t \in \{0, 1\}$, and makes a sequence of reports to the agent $s \equiv \{s_t\}_{t=1}^\infty$ with $s_t \in \{0, 1\}$. The reports s can be thought of as contents of newspaper articles, while s' correspond to the information that the newspaper collects before writing the articles. Since s' is an exchangeable sequence, we can represent it, as before, with the fraction of signals that are 1's, denoted by $\rho' \in [0, 1]$, and similarly s is represented by $\rho \in [0, 1]$. This is convenient as it allows us to model the mixed strategy of the media as a mapping

$$\sigma_M : [0, 1] \rightarrow \Delta([0, 1]),$$

where $\Delta([0, 1])$ is the set of probability distributions on $[0, 1]$. Let \mathbf{i} be the strategy that puts probability 1 on the identity mapping, thus corresponding to M reporting truthfully. Otherwise, i.e., if $\sigma_M \neq \mathbf{i}$, there is manipulation (or misreporting) on the part of the media outlet M .²⁰

We also assume for simplicity that ρ' has a continuous distribution with density g_1 when $\theta = 1$ and g_0 when $\theta = 0$, such that $g_1(\rho) = 0$ for all $\rho \leq \bar{\rho}$ and $g_1(\rho) > 0$ for all $\rho > \bar{\rho}$, while $g_0(\rho) > 0$ for all $\rho \leq \bar{\rho}$ and $g_0(\rho) = 0$ for all $\rho > \bar{\rho}$. This assumption implies that if M reports truthfully, i.e., $\sigma_M = \mathbf{i}$, then Theorem 2 applies and there will be asymptotic learning (and also asymptotic agreement when there are more than one agent).

Now suppose instead that there are three different types of player M (unobservable to the agent). With probability $\lambda_H \in (0, 1)$, the media is honest and can only play $\sigma_M^H = \mathbf{i}$ (where

²⁰See Ottaviani and Sorenson (2006a,b) for more general analyses of potentially biased professional advice and Baron (2004) and Gentzkow and Shapiro (2006) for related models of media bias.

the superscript is for type H —honest). With probability $\lambda_\alpha \in (0, 1 - \lambda_H)$, the media outlet is of type α and is biased towards 1. Type α media outlet receives utility equal to x irrespective of ρ' , and hence would like to manipulate the agent to choose high values of x . With the complementary probability $\lambda_\beta = 1 - \lambda_\alpha - \lambda_H$, the media outlet is of type β and is biased towards 0, and receives utility equal to $1 - x$.

Let us now look for the perfect Bayesian equilibrium of the game between the media outlet and the agent. The perfect Bayesian equilibrium can be represented by two reporting functions $\sigma_M^\alpha : [0, 1] \rightarrow \Delta([0, 1])$ and $\sigma_M^\beta : [0, 1] \rightarrow \Delta([0, 1])$ for the two biased types of M , and updating function $\phi : [0, 1] \rightarrow [0, 1]$, which determines the belief of the agent that $\theta = 1$ when the sequence of reports is ρ , and an action function $x : [0, 1] \rightarrow [0, 1]$, which determines the choice of the agent as a function of ρ (there is no loss of generality here in restricting to pure strategies).

In equilibrium, x must be optimal for the agent given ϕ ; ϕ must be derived from Bayes rule given σ_M^α , σ_M^β and the prior π ; and σ_M^α and σ_M^β must be optimal for the two biased media outlets given x .

Note first that since the payoff to the biased media outlet does not depend on the true ρ' , without loss of generality, we can restrict σ_M^α and σ_M^β not to depend on ρ' . Then, with a slight abuse of notation, let $\sigma_M^\alpha(\rho)$ and $\sigma_M^\beta(\rho)$ be the respective densities with which these two types report ρ .

Second, the optimal choice of the agent after observing a sequence of signals with fraction ρ being equal to 1 is

$$x(\rho) = \phi(\rho),$$

for all $\rho \in [0, 1]$, i.e., the agent will choose an action equal to his belief $\phi(\rho)$.

Third, an application of Bayes' rule implies the following belief for the agent:

$$\phi(\rho) = \begin{cases} \frac{(\lambda_\alpha \sigma_M^\alpha(\rho) + \lambda_\beta \sigma_M^\beta(\rho))\pi}{(1-\pi)\lambda_H g_0(\rho) + \lambda_\alpha \sigma_M^\alpha(\rho) + \lambda_\beta \sigma_M^\beta(\rho)} & \text{if } \rho \leq \bar{\rho} \\ \frac{(\lambda_H g_1(\rho) + \lambda_\alpha \sigma_M^\alpha(\rho) + \lambda_\beta \sigma_M^\beta(\rho))\pi}{\pi \lambda_H g_1(\rho) + \lambda_\alpha \sigma_M^\alpha(\rho) + \lambda_\beta \sigma_M^\beta(\rho)} & \text{if } \rho > \bar{\rho}. \end{cases} \quad (24)$$

The following lemma shows that any (perfect Bayesian) equilibrium has a very simple form:

Lemma 5 *In any equilibrium, there exist $\phi_A > \pi$ and $\phi_B < \pi$ such that $\phi(\rho) = \phi_B$ for all $\rho < \bar{\rho}$ and $\phi(\rho) = \phi_A$ for all $\rho > \bar{\rho}$.*

Proof. From (24), $\phi(\rho) < \pi$ when $\rho < \bar{\rho}$, and $\phi(\rho) > \pi$ when $\rho > \bar{\rho}$. Since the media type α maximizes $x(\rho) = \phi(\rho)$, we have $\sigma_M^\alpha(\rho) = 0$ for $\rho < \bar{\rho}$. Now suppose that the lemma is false and there exists $\rho_1, \rho_2 \leq \bar{\rho}$ such that $\phi(\rho_1) > \phi(\rho_2)$. Then we also have $\sigma_M^\beta(\rho_1) = 0$ —since media type β minimizes $x(\rho) = \phi(\rho)$. But in that case, equation (24) implies that $\phi(\rho_1) = 0$, contradicting the hypothesis. Therefore, $\phi(\rho)$ is constant over $\rho \in [0, \bar{\rho}]$. The proof for $\phi(\rho)$ being constant over $\rho \in (\bar{\rho}, 1]$ is analogous. ■

It follows immediately from this lemma that equilibrium beliefs will take the form given in the next proposition:

Proposition 4 *Suppose that $\rho \neq \bar{\rho}$, then the unique equilibrium actions and beliefs are:*

$$\sigma_M^\alpha(\rho) = g_1(\rho) \tag{25}$$

$$\sigma_M^\beta(\rho) = g_0(\rho) \tag{26}$$

$$x(\rho) = \phi(\rho) = \begin{cases} \frac{\lambda_\beta \pi}{(1-\pi)\lambda_H + \lambda_\beta} & \text{if } \rho < \bar{\rho} \\ \frac{\pi(\lambda_H + \lambda_\alpha)}{\pi\lambda_H + \lambda_\alpha} & \text{if } \rho > \bar{\rho}. \end{cases} \tag{27}$$

Proof. Consider the case $\rho < \bar{\rho}$. As in the proof of Lemma 5, $\sigma_M^\alpha(\rho) = 0$. Since $\phi(\rho)$ is constant over $\rho \in [0, \bar{\rho}]$ (by Lemma 5), equation (24) implies that σ_M^β is proportional to g_0 on this range. Since this range is the common support of the densities σ_M^β and g_0 , it must be that $\sigma_M^\beta = g_0$. Similarly, $\sigma_M^\alpha = g_1$. Substituting these equalities in (24), we obtain (27). ■

This proposition implies that the unique equilibrium of the game between the media outlet and the agent leads to a special case of our model of learning under uncertainty. In particular, the beliefs in (27) can be obtained by the appropriate choice of the functions $f_A(\cdot)$ and $f_B(\cdot)$ from equation (6) in Section 2. This illustrates that the type of learning under uncertainty analyzed in this paper is likely to emerge in game-theoretic situations where one of the players is trying to manipulate the beliefs of others.

5 Concluding Remarks

A key assumption of most theoretical analyses is that individuals have a “common prior,” meaning that they have beliefs consistent with each other regarding the game forms, institutions, and possible distributions of payoff-relevant parameters. This presumption is often justified by the argument that sufficient common experiences and observations, either through individual observations or transmission of information from others, will eliminate disagreements, taking

agents towards common priors. This presumption receives support from a number of well-known theorems in statistics and economics, for example, Savage (1954) and Blackwell and Dubins (1962).

Nevertheless, existing theorems apply to environments in which learning occurs under *certainty*, that is, individuals are certain about the meaning of different signals. Certainty is sufficient to ensure that payoff-relevant variables can be *identified* from limiting frequencies of signals. In many situations, individuals are not only learning about a payoff-relevant parameter but also about the interpretation of different signals. This takes us to the realm of environments where learning takes place under *uncertainty*. For example, many signals favoring a particular interpretation might make individuals suspicious that the signals come from a biased source. We show that learning in environments with uncertainty may lead to a situation in which there is lack of *full identification* (in the standard sense of the term in econometrics and statistics). In such situations, information will be useful to individuals but may not lead to full learning.

This paper investigates the conditions under which learning under uncertainty will take individuals towards common priors and asymptotic agreement. We consider an environment in which two individuals with different priors observe the same infinite sequence of signals informative about some underlying parameter. Learning is under *uncertainty*, however, because each individual has a non-degenerate subjective probability distribution over the likelihood of different signals given the values of the parameter. We show that when subjective probability distributions of both individuals have full support, they will never agree, even after observing the same infinite sequence of signals. We also show that this corresponds to a result of “agreement to eventually disagree”; individuals will agree, before observing the sequence of signals, that their posteriors about the underlying parameter will not converge. This common understanding that more information may not lead to similar beliefs for agents has important implications for a variety of games and economic models. On the other hand, when there is no full support in subjective probability distributions, asymptotic learning and agreement may obtain.

An important implication of this analysis is that after observing the same (infinite) sequence of signals, two Bayesian individuals may end up disagreeing more than they originally did. This result contrasts with the common presumption that shared information and experiences will ensure *eventual* agreement.

We also systematically investigate whether asymptotic agreement obtain as the amount of

uncertainty in the environment diminishes (i.e., as we look at families of subjective probability distributions converging to degenerate limit distributions with all their mass at one point). We provide a complete characterization of the conditions under which this will be the case. Asymptotic disagreement may prevail even under “approximate certainty,” as long as the family of subjective probability distributions converging to a degenerate distribution (and thus to an environment with certainty) has regularly-varying tails (such as for the Pareto, the log-normal or the t-distributions). In contrast, with rapidly-varying tails (such as the normal and the exponential distributions), convergence to certainty leads to asymptotic agreement.

Lack of common beliefs and common priors has important implications for economic behavior in a range of circumstances. We illustrate how the type of learning outlined in this paper interacts with economic behavior in various different situations, including games of coordination, games of common interest, bargaining, asset trading and games of communication. For example, we show that contrary to standard results, individuals may wish to play common-interest games *before* rather than after receiving more information about payoffs. Similarly, we show how the possibility of observing the same sequence of signals may lead to “speculative delay” in asset trading among individuals that start with similar beliefs. We also provide a simple example illustrating why individuals may be uncertain about informativeness of signals—the strategic behavior of other agents trying to manipulate their beliefs.

6 Appendix: Omitted Proofs

Proof of Lemma 3. The proof is identical to that of Lemma 1. ■

Proof of Theorem 6.

(Part1) This part immediately follows from Lemma 3, as each $\pi_{k'}^i f_{A^{k'}}(\rho(s))$ is positive, and $\pi_k^i f_{A^k}(\rho(s))$ is finite.

(Part 2) Assume $F_\theta^1 = F_\theta^2$ for each $\theta \in \Theta$. Then, by Lemma 3, $\phi_{k,\infty}^1(\rho) - \phi_{k,\infty}^2(\rho) = 0$ if and only if $(T_k(\pi^1) - T_k(\pi^2))' T_k \left((f_\theta^1(\rho))_{\theta \in \Theta} \right) = 0$. The latter inequality has probability 0 under both probability measures Pr^1 and Pr^2 by hypothesis. ■

Proof of Theorem 7. Define $\bar{\pi} = (1/K, \dots, 1/K)$. First, take $\pi^1 = \pi^2 = \bar{\pi}$. Then,

$$\frac{\sum_{k' \neq k} \pi_{k'}^1 f_{A^{k'}}^1(\rho(s))}{\pi_k^1 f_{A^k}^1(\rho(s))} - \frac{\sum_{k' \neq k} \pi_{k'}^2 f_{A^{k'}}^1(\rho(s))}{\pi_k^2 f_{A^k}^1(\rho(s))} = \mathbf{1}' \left(T_k \left((f_\theta^1(\rho(s)))_{\theta \in \Theta} \right) - T_k \left((f_\theta^2(\rho(s)))_{\theta \in \Theta} \right) \right) \neq 0,$$

where $\mathbf{1} \equiv (1, \dots, 1)'$, and the inequality follows by the hypothesis of the theorem. Hence, by Lemma 3, $|\phi_{k,\infty}^1(\rho(s)) - \phi_{k,\infty}^2(\rho(s))| > 0$ for each $\rho(s) \in [0, 1]$. Since $[0, 1]$ is compact and $|\phi_{k,\infty}^1(\rho(s)) - \phi_{k,\infty}^2(\rho(s))|$ is continuous in $\rho(s)$, there exists $\epsilon > 0$ such that $|\phi_{k,\infty}^1(\rho(s)) - \phi_{k,\infty}^2(\rho(s))| > \epsilon$ for each $\rho(s) \in [0, 1]$. Now, since $|\phi_{k,\infty}^1(\rho(s)) - \phi_{k,\infty}^2(\rho(s))|$ is continuous in π^1 and π^2 , there exists a neighborhood $N(\bar{\pi})$ of $\bar{\pi}$ such that

$$|\phi_{k,\infty}^1(\rho(s)) - \phi_{k,\infty}^2(\rho(s))| > |\pi_k^1 - \pi_k^2| \text{ for each } k = 1, \dots, K \text{ and } s \in \bar{S}$$

for all $\pi^1, \pi^2 \in N(\bar{\pi})$. Since $\text{Pr}^i(\bar{S}) = 1$, the last statement in the theorem follows. ■

Proof of Theorem 8. Our proof utilizes the following two lemmas.

Lemma A.

$$\lim_{m \rightarrow \infty} \phi_{k,\infty,m}^i(p) = \frac{1}{1 + \sum_{k' \neq k} \frac{\pi_{k'}^i}{\pi_k^i} \tilde{R}(p - \hat{p}(i, A^{k'}), p - \hat{p}(i, A^k))}.$$

Proof. By condition (i), $\lim_{m \rightarrow \infty} c(i, A^k, m) = 1$ for each i and k . Hence, for every distinct k and k' ,

$$\lim_{m \rightarrow \infty} \frac{f_{A^{k'}}^i(p)}{f_{A^k}^i(p)} = \lim_{m \rightarrow \infty} \frac{c(i, A^{k'}, m)}{c(i, A^k, m)} \lim_{m \rightarrow \infty} \frac{f(m(p - \hat{p}(i, A^{k'})))}{f(m(p - \hat{p}(i, A^k)))} = \tilde{R}(p - \hat{p}(i, A^{k'}), p - \hat{p}(i, A^k)).$$

Then, Lemma A follows from Lemma 3. ■

Lemma B. For any $\tilde{\epsilon} > 0$ and $h > 0$, there exists \tilde{m} such that for each $m > \tilde{m}$, $k \leq K$, and each $\rho(s)$ with $\|\rho(s) - \hat{p}(i, A^k)\| < h/m$,

$$\left| \phi_{k,\infty,m}^i(\rho(s)) - \lim_{m \rightarrow \infty} \phi_{k,\infty,m}^i(\hat{p}(i, A^k)) \right| < \tilde{\epsilon}. \quad (28)$$

Proof. Since, by hypothesis, \tilde{R} is continuous at each $(\hat{p}(i, \theta) - \hat{p}(j, \theta'), \hat{p}(i, \theta) - \hat{p}(j, \theta))$, by Lemma A, there exists $h' > 0$, such that

$$\left| \lim_{m \rightarrow \infty} \phi_{k,\infty,m}^i(\rho(s)) - \lim_{m \rightarrow \infty} \phi_{k,\infty,m}^i(\hat{p}(i, A^k)) \right| < \tilde{\epsilon}/2 \quad (29)$$

and by condition (iii), there exists $\tilde{m} > h/h'$ such that

$$\left| \phi_{k,\infty,m}^i(\rho(s)) - \lim_{m \rightarrow \infty} \phi_{k,\infty,m}^i(\rho(s)) \right| < \tilde{\epsilon}/2. \quad (30)$$

holds uniformly in $\|\rho(s) - \hat{p}(i, A^k)\| < h'$. The inequalities in (29) and (30) then imply (28). \blacksquare

(Proof of Part 1) Since $\tilde{R}(\hat{p}(i, A^k) - \hat{p}(i, A^{k'}), 0) = 0$ for each $k' \neq k$ (by condition (i)), Lemma A implies that $\lim_{m \rightarrow \infty} \phi_{k,\infty,m}^i(\hat{p}(i, A^k)) = 1$. Hence, $\lim_{m \rightarrow \infty} (\phi_{k,\infty,m}^i(\hat{p}(i, A^k)) - \phi_{k,\infty,m}^j(\hat{p}(i, A^k))) = 0$ if and only if $\lim_{m \rightarrow \infty} \phi_{k,\infty,m}^j(\hat{p}(i, A^k)) = 1$. Since each ratio $\pi_{k'}^j/\pi_k^j$ is positive, by Lemma A, the latter holds if and only if $\tilde{R}(\hat{p}(i, A^k) - \hat{p}(j, A^{k'}), \hat{p}(i, A^k) - \hat{p}(j, A^k)) = 0$ for each $k' \neq k$, establishing Part 1.

(Proof of Part 2) Fix $\epsilon > 0$ and $\delta > 0$. Fix also any i and k . Since each $\pi_{k'}^j/\pi_k^j$ is finite, by Lemma 3, there exists $\epsilon' > 0$, such that $\phi_{k,\infty,m}^i(\rho(s)) > 1 - \epsilon$ whenever $f_{A^{k'}}^i(\rho(s))/f_{A^k}^i(\rho(s)) < \epsilon'$ holds for every $k' \neq k$. Now, by (i), there exists $h_{0,k} > 0$, such that

$$\Pr^i(\|\rho(s) - \hat{p}(i, A^k)\| \leq h_{0,k}/m | \theta = A^k) = \int_{\|x\| \leq h_{0,k}} f(x) dx > (1 - \delta).$$

Let

$$Q_{k,m} = \{p \in \Delta(L) : \|p - \hat{p}(i, A^k)\| \leq h_{0,k}/m\}$$

and $\kappa \equiv \min_{\|x\| \leq h_{0,k}} f(x) > 0$. By (i), there exists $h_{1,k} > 0$ such that, whenever $\|x\| > h_{1,k}$, $f(x) < \epsilon'\kappa/2$. There exists a sufficiently large constant $m_{1,k}$ such that for any $m > m_{1,k}$, $\rho(s) \in Q_{k,m}$, and any $k' \neq k$, we have $\|\rho(s) - \hat{p}(i, A^{k'})\| > h_{1,k}/m$, and

$$\frac{f(m(\rho(s) - \hat{p}(i, A^{k'})))}{f(m(\rho(s) - \hat{p}(i, A^k)))} < \frac{\epsilon'\kappa}{2} \frac{1}{\kappa} = \frac{\epsilon'}{2}.$$

Moreover, since $\lim_{m \rightarrow \infty} c(i, \theta, m) = 1$ for each i and θ , there exists $m_{2,k} > m_{1,k}$ such that $c(i, A^{k'}, m)/c(i, A^k, m) < 2$ for every $k' \neq k$ and $m > m_{2,k}$. This implies

$$f_{A^{k'}}^i(\rho(s))/f_{A^k}^i(\rho(s)) < \epsilon',$$

establishing that

$$\phi_{k,\infty,m}^i(\rho(s)) > 1 - \epsilon. \quad (31)$$

Now, for $j \neq i$, assume that $\tilde{R}(\hat{p}(i, \theta) - \hat{p}(j, \theta'), \hat{p}(i, \theta) - \hat{p}(j, \theta)) = 0$ for each distinct θ and θ' . Then, by Lemma A, $\lim_{m \rightarrow \infty} \phi_{k,\infty,m}^j(\hat{p}(i, A^k)) = 1$, and hence by Lemma B, there exists $m_{3,k} > m_{2,k}$ such that for each $m > m_{3,k}$, $\rho(s) \in Q_{k,m}$,

$$\phi_{k,\infty,m}^j(\rho(s)) > 1 - \epsilon. \quad (32)$$

Notice that when (31) and (32) hold, we have $\|\phi_{\infty,m}^1(s) - \phi_{\infty,m}^2(s)\| < \epsilon$. Then, setting $\bar{m} = \max_k m_{4,k}$, we obtain the desired inequality for each $m > \bar{m}$:

$$\begin{aligned} \Pr^i(\|\phi_{\infty,m}^1(s) - \phi_{\infty,m}^2(s)\| < \epsilon) &= \sum_{k \leq K} \Pr^i(\|\phi_{\infty,m}^1(s) - \phi_{\infty,m}^2(s)\| < \epsilon | \theta = A^k) \Pr^i(\theta = A^k) \\ &\geq \sum_{k \leq K} \Pr^i(\rho(s) \in Q_{k,m} | \theta = A^k) \Pr^i(\theta = A^k) \\ &\geq \sum_{k \leq K} (1 - \delta) \pi_k^i \\ &= 1 - \delta. \end{aligned}$$

(Proof of Part 3) Assume that $\tilde{R}(\hat{p}(i, \theta) - \hat{p}(j, \theta'), \hat{p}(i, \theta) - \hat{p}(j, \theta)) \neq 0$ for each distinct θ and θ' . Then, since each $\pi_{k'}^j / \pi_k^j$ is positive, Lemma A implies that $\lim_{m \rightarrow \infty} \phi_{k, \infty, m}^j(\hat{p}(i, A^k)) < 1$ for each k . Let

$$\epsilon = \min_k \left\{ 1 - \lim_{m \rightarrow \infty} \phi_{k, \infty, m}^j(\hat{p}(i, A^k)) \right\} / 3 > 0.$$

Then, by part 2, for each k , there exists $m_{2, k}$ such that for every $m > m_{2, k}$ and $\rho(s) \in Q_{k, m}$, we have $\phi_{k, \infty}^i(\rho(s)) > 1 - \epsilon$. By Lemma B, there also exists $m_{5, k} > m_{2, k}$ such that for every $m > m_{5, k}$ and $\rho(s) \in Q_{k, m}$,

$$\phi_{k, \infty, m}^j(\rho(s)) < \lim_{m \rightarrow \infty} \phi_{k, \infty, m}^j(\hat{p}(i, A^k)) + \epsilon \leq 1 - 2\epsilon < \phi_{k, \infty}^i(\rho(s)) - \epsilon.$$

This implies that $\|\phi_{\infty, m}^1(\rho(s)) - \phi_{\infty, m}^2(\rho(s))\| > \epsilon$. Setting $\bar{m} = \max_k m_{5, k}$ and changing $\|\phi_{\infty, m}^1(s) - \phi_{\infty, m}^2(s)\| < \epsilon$ at the end of the proof of Part 2 to $\|\phi_{\infty, m}^1(s) - \phi_{\infty, m}^2(s)\| > \epsilon$, we obtain the desired inequality. ■

7 References

- Aumann, Robert (1986): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1-18.
- Aumann, Robert (1998): "Common priors: A Reply to Gul," *Econometrica*, 66-4, 929-938.
- Baron, David (2004): "Persistent Media Bias," Stanford Graduate School of Business, Mimeo.
- Berk, Robert (1966): "Limiting Behavior of Posterior Distributions When the Model Is Incorrect," *Annals of Mathematical Statistics*, 37, 51-58.
- Billingsley, Patrick (1995): *Probability and Measure*, Third Edition, Wiley and Sons, New York.
- Blackwell, David and Lester Dubins (1962): "Merging of Opinions with Increasing Information," *Annals of Mathematical Statistics*, 33, 882-887.
- Carlsson, Hans and Eric van Damme (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61, 981-1018.
- Crawford, Vincent and Joel Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50:1431-52.
- Cripps, Martin, Jeffrey Ely, George Mailath and Larry Samuelson (2006): "Common Learning" Cowles Foundation Discussion Paper No. 1575, August 2006.
- Doob, J. L. (1949): "Application of the Theory of Martingales," in *Le Calcul des Probabilités et ses Applications*, 23-27, Paris.
- Diaconis Persi and David Freedman (1986): "On Inconsistent Bayesian Estimates of Location" *Annals of Statistics*, 14, 68-87.
- Dixit, Avinash and Jorgen Weibull (2007): "Political Polarization," Mimeo, Princeton University and Stockholm School of Economics.
- Feller, William (1971): *An Introduction to Probability Theory and Its Applications*. Volume II, 2nd ed. Wiley, New York.
- Freedman, David (1963): "On the Asymptotic Behavior of Bayes Estimates in the Discrete Case I" *Annals of Mathematical Statistics*, 34, 1386-1403.
- Freedman, David (1965): "On the Asymptotic Behavior of Bayes Estimates in the Discrete Case II" *Annals of Mathematical Statistics*, 36, 454-456.
- Fudenberg Drew and David Kreps (1995): "Learning in Extensive Games, I: Self-Confirming Equilibria," *Games and Economic Behavior*, 8, 20-55.
- Fudenberg, Drew and David Levine (1993): "Self-Confirming Equilibria" *Econometrica*, 61, 523- 546.

- Fudenberg, Drew and David Levine (2005): "Learning and Belief Based Trade," *Latin American Journal of Economics*, 42, 199-27.
- Geanakoplos, John and Heraklis Polemarchakis (1982): "We Can't Disagree Forever," *Journal of Economic Theory*, 28, 192-200.
- Gentzkow, Matthew and Jesse M. Schapiro (2006): "Media Bias and Reputation," *Journal of Political Economy*, 114, 280-317.
- Ghosh J. K. and R. V. Ramamoorthi (2003): *Bayesian Nonparametrics*, Springer-Verlag, New York.
- Gul, Faruk (1998): "A Comment on Aumann's Bayesian view", *Econometrica*, 66-4, 923-927.
- Guillemin Victor and Alan Pollack (1974): *Differential Topology*, Prentice Hall, New Jersey.
- Haan, L. de (1970): *On Regular Variation and Its Application to the Weak Convergence of Sample Extremes*, Mathematical Centre Tract 32, Mathematics Centre, Amsterdam, Holland.
- Harris, Milton and Artur Raviv (1993): "Differences of Opinion Make a Horse Race," *Review of Financial Studies*, 6, 473-506.
- Harrison, Michael and David Kreps (1978): "Speculative Investor Behavior in a Stock Market with Heterogenous Expectations," *Quarterly Journal of Economics*, 92, 323-36.
- Jackson, Matthew O., Ehud Kalai and Rann Smorodinsky (1999): "Bayesian Representation of Stochastic Processes under Learning: De Finetti Revisited," *Econometrica*, 67, 875-893.
- Ottaviani, Marco and Peter Sorensen (2006a): "Professional Advice," *Journal of Economic Theory*, 126, 120-141.
- Ottaviani, Marco and Peter Sorensen (2006b): "The Strategy of Professional Forecasting," *Journal of Financial Economics*, 81, 441-466.
- Kurz, Mordecai (1994): "On the Structure and Diversity of Rational Beliefs," *Economic Theory*, 4, 877-900.
- Kurz, Mordecai (1996): "Rational Beliefs and Endogenous Uncertainty: An Introduction," *Economic Theory*, 8, 383-397.
- Milgrom, Paul and Nancy Stokey (1982): "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26, 177-227.
- Miller, Ronald I. and Chris W. Sanchirico (1999): "The Role of Option with Continuity in Merging of Opinions and Rational Learning," *Games and Economic Behavior*, 29, 170-190.
- Morris, Stephen (1996): "Speculative Investor Behavior and Learning," *Quarterly Journal of Economics*, Vol. 111, No. 4, 1111-33.
- Savage, Leonard J. (1954): *The Foundations of Statistics*, Dover reprint, New York, 1972.

Seneta, Eugène (1976): *Regularly Varying Functions*, Lecture Notes in Mathematics, 508, Springer, Berlin.

Stinchcombe, Maxwell (2005): “The Unbearable Flightness of Bayesians: Generically Erratic Updating,” University of Texas, Austin, mimeo.

Yildiz, Muhamet (2003): “Bargaining without a Common Prior – An Immediate Agreement Theorem,” *Econometrica*, 71 (3), 793-811.

Yildiz, Muhamet (2004): “Waiting to Persuade,” *Quarterly Journal of Economics*, 119 (1), 223-249.