LEARNING ABOUT INFORMATIVENESS

WANYING HUANG†

Abstract. We study whether individuals can learn the informativeness of their information technology through social learning. As in the classic sequential social learning model, rational agents arrive in order and make decisions based on the past actions of others and their private signals. There is uncertainty regarding the informativeness of the common signal-generating process. We show that in this setting asymptotic learning about informativeness is not guaranteed and depends crucially on the relative tail distributions of private beliefs induced by uninformative and informative signals. We identify the phenomenon of perpetual disagreement as the cause of learning and characterize learning in the canonical Gaussian environment.

1. Introduction

Social learning plays a vital role in the dissemination and aggregation of information. The behavior of others reflects their private knowledge about an unknown state of the world, and so by observing others, individuals can acquire additional information, enabling them to make better-informed decisions. A key assumption in most existing social learning models is the presence of an informative source that provides a useful private signal to each individual. In this paper, we explore how the possibility that the source is uninformative interferes with learning, and study the conditions under which individuals can eventually distinguish an uninformative source from an informative one. This question is particularly relevant today due to the proliferation of novel information technologies, raising concerns about the accuracy and credibility of the information they provide.

Formally, we introduce uncertainty regarding the informativeness of the source into the classic sequential social learning model (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992). As usual, a sequence of short-lived agents arrives in order, each acting once by choosing an action to match an unknown payoff-relevant state that can be either good

Date: August 28, 2025.
Acknowledgements. This work is based on the first chapter of my PhD dissertation at Caltech, and was awarded "best student paper" at EC'24. Part of the work was done while the author was supported by a PIMCO Graduate Fellowship in Data Science at Caltech.
I am indebted to my PhD advisor Omer Tamuz for his continued support and encouragement. I am grateful to (in alphabetical order) Boğaçhan Çelen, Krishna Dasaratha, Laura Doval, Federico Echenique, Amanda Friedenberg, Wade Hann-Caruthers, SangMok Lee, Kirby Nielson, Luciano Pomatto, Fan Wu, and seminar and conference participants for their helpful comments and suggestions.
†Department of Economics, Monash Business School. Email: `kate.huang@monash.edu`.

or bad. Before making their decisions, each agent observes the past actions of her predecessors and receives a private signal from a common source of information. However, unlike in the usual setting, there is uncertainty surrounding this common information source. In particular, we assume that this source can be either informative, generating private signals that are independent and identically distributed (i.i.d.) conditioned on the payoff-relevant state, or uninformative, producing private signals that are i.i.d. but independent of the payoff-relevant state. Both the payoff-relevant state and the informativeness of the source are realized independently at the outset and are assumed to be fixed throughout.

If an outside observer, who aims to evaluate the informativeness of the source, were to have access to the private signals received by the agents, he would gradually accumulate empirical evidence about the source and eventually learn its informativeness. However, when only the history of past actions is observable, his inference problem becomes more challenging—not only because there is less information available, but also because these past actions are correlated with each other. This correlation arises from social learning behavior, where agents' decisions are influenced by the inferences they draw from observing others' actions. We say that *asymptotic learning of informativeness* holds if the outside observer's belief about the source's informativeness converges to the truth, i.e., it converges almost surely to one when the source is informative and to zero when it is uninformative. The questions we aim to address are: can learning about informativeness be achieved asymptotically, and if so, under what conditions? Furthermore, what are the behavioral implications of such learning?

We consider *unbounded signals* (Smith and Sørensen, 2000) under which the agent's private belief induced by an informative signal can be arbitrarily strong. We focus on this setting since otherwise, learning about the source's informativeness can easily be precluded by agents' lack of response to their private signals. As shown by Smith and Sørensen (2000), with bounded signals, an information cascade—where agents stop responding to their private signals—would be triggered, blocking further information aggregation. In contrast, with unbounded signals, cascades do not form, thus allowing information to continue accumulating through agents' actions.

Our main result (Theorem 1) shows that even with unbounded signals, achieving asymptotic learning of informativeness is far from guaranteed. The key factor in determining such learning lies in the tail distributions of agents' private beliefs. Specifically, it depends on whether the belief distribution induced by uninformative signals has *fatter* or *thinner* tails compared to that induced by informative signals. We show that asymptotic learning of informativeness holds when uninformative signals have fatter tails than informative signals, but fails when uninformative signals have thinner tails.

For example, consider an informative source that generates Gaussian signals with unit variance and a mean of $+1$ if the payoff-relevant state is good and a mean of -1 if the state is bad. Meanwhile, the uninformative source generates Gaussian signals with mean 0, independent of the state. If the uninformative source generates signals with a variance strictly greater than one, then the uninformative signals have fatter tails, and thus asymptotic learning of informativeness holds. In contrast, when the uninformative signals have a variance strictly less than one, they exhibit thinner tails, and so asymptotic learning of informativeness fails.

As another illustration of the main result, consider the case where the informative signals follow the same distributions as before, but the uninformative signals are chosen uniformly from the bounded interval $[-\varepsilon, \varepsilon]$ for some small $\varepsilon > 0$. In this case, the distribution of private beliefs induced by these uninformative signals also has bounded support. Consequently, it can be viewed as having *extremely* thinner tails compared to the informative Gaussian signals. Hence, Theorem 1 implies that the informativeness of the source cannot be fully revealed. Nevertheless, under such an informative source, almost all agents individually learn its informativeness: once they receive a signal outside the support $[-\varepsilon, \varepsilon]$, which is highly probable for small ε, they infer that it can only come from the normal distribution, indicating that the source must be informative. However, an outside observer who only observes the agents' actions is unable to determine the informativeness of the source.

The mechanism behind the main result is as follows. First, in our model, despite information uncertainty, agents always act as if the signals are informative. Therefore, when the source is indeed informative and generates unbounded signals, agents will eventually take the correct action. Now, suppose the source is uninformative and generates signals with thinner tails. In this case, it is unlikely that agents will receive signals extreme enough to break the consensus on actions, so they usually mimic their predecessors. As a result, an outside observer who only observes agents' actions cannot discern that the source is uninformative, as an action consensus will be reached under both uninformative and informative sources. In contrast, suppose the source is uninformative but generates signals with fatter tails. In this scenario, extreme signals are more likely, allowing agents to break the consensus; in fact, both actions will be taken infinitely often, so agents will never settle on an action consensus. Hence, an outside observer who observes an infinite number of action switches learns that the source is uninformative.

For some private belief distributions, their relative tail thickness is neither thinner nor fatter. For these, we show that the same holds: asymptotic learning of informativeness is achieved if and only if conditioned on the source being uninformative, agents never settle on an action consensus (Proposition 2). In the Gaussian setting where informative signals are symmetric and the relative tail thickness is incomparable, we complement our main result by showing that learning holds if the uninformative signals have mean zero (Proposition 1). In terms of behavior, as mentioned before, when the source is informative, agents eventually choose the correct action regardless of information uncertainty;

nevertheless, we show that agents can only be certain that they are taking the correct action if and only if the source's informativeness is fully revealed (Proposition 4). In contrast, when the source is uninformative, agents are clearly not guaranteed to settle on the correct action; in fact, their actions may or may not converge at all. As demonstrated by Proposition 2, an outside observer eventually learns the informativeness of the source if and only if the agents' actions do not converge when the source is uninformative.

Related Literature. Our paper contributes to a rich literature on sequential social learning. Assuming that the common source of information is always informative, the primary focus of this literature has been on determining whether agents can eventually learn to choose the correct action. Various factors, such as the information structure (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992; Smith and Sørensen, 2000) and the observational networks (Çelen and Kariv, 2004; Acemoglu, Dahleh, Lobel, and Ozdaglar, 2011; Lobel and Sadler, 2015), have been extensively studied to analyze their impact on information aggregation, including its efficiency (Rosenberg and Vieille, 2019) and the speed of learning (e.g., Vives, 1993; Hann-Caruthers, Martynov, and Tamuz, 2018). However, the question of learning about the informativeness of the source—which is the focus of this paper—remains largely unexplored.[1]

A few papers explore the idea of agents having access to multiple sources of information in the context of social learning. For example, Liang and Mu (2020) consider a model in which agents endogenously choose from a set of correlated information sources, and the acquired information is then made public and learned by other agents. They focus on the externality in agents' information acquisition decisions and show that information complementarity can result in either efficient information aggregation or "learning traps," in which society gets stuck in choosing suboptimal information structures. In a different setting, Chen (2022) examines a sequential social learning model in which ambiguity-averse agents have access to different sources of information. Consequently, information uncertainty arises in his model because agents are unsure about the signal precision of their predecessors. He shows that under sufficient ambiguity aversion, there can be information cascades even with unbounded signals. Our paper differs from these prior works as we focus on rational agents with access to a common source of information of unknown informativeness.

Another way of viewing our model is by considering a social learning model with four states: the source is either informative with the good or bad state, or uninformative with either the good or bad state. In such multi-state settings, recent work by Arieli and Mueller-Frank (2021) demonstrates that pairwise unbounded signals are necessary and sufficient for learning, when the decision problem that agents face includes a distinct action that is uniquely optimal for each state. This is not the case in our model, because

[1]For comprehensive surveys on recent developments in the social learning literature, see e.g., Golub and Sadler (2017); Bikhchandani, Hirshleifer, Tamuz, and Welch (2021).

the same action is optimal in different states, e.g., when the source is uninformative, and so even when agents observe a very strong signal indicating that the state is uninformative, they do not reveal it in their behavior.

More recently, Kartik, Lee, Liu, and Rappoport (2022) consider a setting with multiple states and actions on general sequential observational networks. They identify a sufficient condition for learning —"excludability" —that jointly depends on the information structure and agents' preferences. Roughly speaking, this condition ensures that agents can always displace the wrong action, which is their driving force for learning. In our model, when the source is uninformative, agents cannot displace the wrong action as all signals are pure noise.[2] Conceptually, our approach differs from theirs as we are interested in identifying the uninformative state from the informative one, instead of identifying the payoff-relevant state.

Our paper is also related to the growing literature on social learning with misspecified models. Bohren (2016) investigates a model where agents fail to account for the correlation between actions, demonstrating that different degrees of misperception can lead to distinct learning outcomes. In a broader framework, Bohren and Hauser (2021) show that learning remains robust to minor misspecifications. In contrast, Frick, Iijima, and Ishii (2020) find that an incorrect understanding of other agents' preferences or types can result in a severe breakdown of information aggregation, even with a small amount of misperception. Later, Frick, Iijima, and Ishii (2023) propose a unified approach to establish convergence results in misspecified learning environments where the standard martingale approach fails to hold. On a more positive note, Arieli, Babichenko, Müller, Pourbabaee, and Tamuz (2023) illustrate that by being mildly condescending—misperceiving others as having slightly lower-quality of information—agents may perform better in the sense that on average, only finitely many of them take incorrect actions.

2. Model

There is an unknown binary state of the world $\theta \in \{\mathfrak{g}, \mathfrak{b}\}$, chosen at time 0 with equal probability. We refer to $\mathfrak{g}$ as the good state and $\mathfrak{b}$ as the bad state. Concurrently, nature chooses an additional binary state $\omega \in \{0, 1\}$, independent of θ, where $\omega = 1$ with probability $\gamma \in (0, 1)$.[3] A countably infinite set of agents indexed by time $t \in \mathbb{N} = \{1, 2, \dots\}$ arrive in order, each acting once. The action of agent t is $a_t \in A = \{\mathfrak{g}, \mathfrak{b}\}$, with a payoff of one if her action matches the state θ and zero otherwise. Before agent t chooses an action, she observes the history of her predecessors' actions $H_t = (a_1, \dots, a_{t-1})$ and receives a private signal s_t, taking values in a measurable space (S, Σ).

The informativeness of private signals depends on ω, which determines the signal-generating process for all agents. If $\omega = 1$, the source is informative, and conditional on

[2]This observation can also be seen from Theorem 2 in Kartik, Lee, Liu, and Rappoport (2022).

[3]Our results do not rely on the independence between θ and ω. They hold true as long as conditioned on $\omega = 0$, both realizations of θ are equally likely.

θ, signals are drawn i.i.d. across agents from a distribution μ_θ. If $\omega = 0$, the source is uninformative, and signals are drawn i.i.d. from a distribution μ_0 that is independent of θ. We denote by $\mathbb{P}_0[\cdot] := \mathbb{P}[\cdot|\ \omega = 0]$ and $\mathbb{P}_1[\cdot] := \mathbb{P}[\cdot|\ \omega = 1]$ the conditional probability distributions given $\omega = 0$ and $\omega = 1$, respectively. Similarly, we use the notation $\mathbb{P}_{1,\mathfrak{g}}[\cdot] := \mathbb{P}[\cdot|\ \omega = 1, \theta = \mathfrak{g}]$ to denote the conditional probability distribution given $(\omega, \theta) = (1, \mathfrak{g})$. We use an analogous notation for $(\omega, \theta) = (1, \mathfrak{b})$.

Agents' Behavior. A pure strategy of agent t is a measurable function $\sigma_t : A^{t-1} \times S \to A$ that selects an action for each possible pair of observed history and private signal. A pure strategy profile $\sigma = (\sigma_t)_{t\in\mathbb{N}}$ is a collection of pure strategies of all agents. A strategy profile is a Bayesian Nash equilibrium—referred to as equilibrium hereafter—if no agent can unilaterally deviate from this profile and obtain a strictly higher expected payoff conditioned on their information. Given that each agent acts only once, the existence of an equilibrium is guaranteed by a simple inductive argument. In equilibrium, each agent t chooses the action a_t that maximizes her expected payoff given the available information:

$$a_t \in \arg\max_{a\in A} \mathbb{E}[\mathbb{1}(\theta = a)|H_t, s_t]. \tag{1}$$

Below, we make a continuity assumption which implies that agents are never indifferent, and so there is a unique equilibrium.

We first observe that, despite the uncertainty regarding the informativeness of the source, in equilibrium, each agent chooses the action that is most likely to match the state, *conditional* on the source being informative.

Lemma 1. *The equilibrium action for each agent t is*

$$a_t \in \arg\max_{a\in A} \mathbb{P}_1[\theta = a|H_t, s_t]. \tag{2}$$

That is, agents always act as if signals are informative, irrespective of the underlying signal-generating process. Intuitively, treating signals as informative—even when they are pure noise—does not adversely affect agents' payoffs, since in the absence of any useful information, each agent with a uniform prior is indifferent between the available actions. Thus, agents always respond to their private signals. Later, in Section 6.2, we show that this responsiveness eventually breaks down under a non-uniform prior, and consequently the source's informativeness can never be fully revealed.

Information Structure. We assume that the distributions $\mu_\mathfrak{g}, \mu_\mathfrak{b}$, and μ_0 are distinct[4] and mutually absolutely continuous, so no signal fully reveals either state θ or ω. As a

[4]Formally, two distributions are distinct if there exists a positive measure set to which they assign a different measure. From an economic perspective, ensuring that μ_0 is distinct from both $\mu_\mathfrak{g}$ and $\mu_\mathfrak{b}$ prevents a situation in which there is a 'fake' information technology that perfectly mimics the informative source at one of the realizations of θ but is actually uninformative. Similarly, $\mu_\mathfrak{g}$ being distinct from $\mu_\mathfrak{b}$ means that in the informative state, the signals are actually informative.

consequence, conditioned on $\omega = 1$, the log-likelihood ratio of any signal s

$$\ell = \log \frac{d\mu_{\mathfrak{g}}}{d\mu_{\mathfrak{b}}}(s),$$

is well-defined. We refer to ℓ as the *agent's private log-likelihood ratio (LLR)*. Since regardless of the realization of ω, agents always act as if the signals are informative, it is sufficient to consider ℓ to capture their behavior.[5] Let $F_{\mathfrak{g}}$ and $F_{\mathfrak{b}}$ be the cumulative distribution functions (CDFs) of ℓ conditioned on $(\omega, \theta) = (1, \mathfrak{g})$ and $(\omega, \theta) = (1, \mathfrak{b})$, respectively. Let F_0 denote the CDF of ℓ conditioned on $\omega = 0$. Note that these conditional CDFs are mutually absolutely continuous, as $\mu_{\mathfrak{g}}, \mu_{\mathfrak{b}}$ and μ_0 are. Let $f_{\mathfrak{g}}, f_{\mathfrak{b}}$ and f_0 denote the corresponding density functions of $F_{\mathfrak{g}}$, $F_{\mathfrak{b}}$ and F_0 whenever they are differentiable.

We focus on *unbounded signals*, where the agent's private LLR can take on arbitrarily large or small values: for any $M \in \mathbb{R}$, there exists a positive probability that $\ell > M$ and a positive probability that $\ell < -M$. We informally refer to a signal s as *extreme* when the corresponding ℓ it induces has a large absolute value. A common example of unbounded private signals is the case of Gaussian signals, where s follows a normal distribution $\mathcal{N}(m_{(\omega,\theta)}, \sigma^2)$ with variance σ^2 and mean $m_{(\omega,\theta)}$ that depends on the pair of states (ω, θ). An extreme Gaussian signal is a signal that is, for example, $5 - \sigma$ away from the mean $m_{(\omega,\theta)}$.

We make two assumptions for expository simplicity. First, we assume that the pair $(F_{\mathfrak{g}}, F_{\mathfrak{b}})$ of informative conditional CDFs is symmetric around zero, i.e., $F_{\mathfrak{g}}(x) + F_{\mathfrak{b}}(-x) = 1$. Given that the prior on θ is uniform, this assumption is equivalent to requiring our model to be invariant with respect to a relabeling of the payoff-relevant state. Second, we assume that all conditional CDFs—$F_{\mathfrak{g}}$, $F_{\mathfrak{b}}$, and F_0—are continuous, so agents are never indifferent between actions.

In addition, we assume that $F_{\mathfrak{b}}$ has a differentiable left tail, i.e., it is differentiable for all x negative enough, and its probability density function $f_{\mathfrak{b}}$ satisfies the condition that $f_{\mathfrak{b}}(-x) < 1$ for all x large enough. By symmetry, this implies that $F_{\mathfrak{g}}$ also has a differentiable right tail and its density function $f_{\mathfrak{g}}$ satisfies the condition that $f_{\mathfrak{g}}(x) < 1$ for all x large enough. This is a mild technical assumption that holds for every non-atomic distribution commonly used in the literature, including the Gaussian distribution. It holds, for instance, whenever the density tends to zero at infinity.

Asymptotic Learning of Informativeness. We now define our notion of asymptotic learning. Let $q_t := \mathbb{P}[\omega = 1 | H_t]$ be the belief that an outside observer assigns to the source being informative after observing the history of agents' actions from time 1 to $t-1$. As this observer collects more information over time, his belief q_t converges almost surely since it is a bounded martingale. Formally, we say that *asymptotic learning of*

[5]Formally, the sequence of actions $a_1, \ldots, a_t$ is determined by $\ell_1, \ldots, \ell_t$.

informativeness holds if for all $\omega \in \{0, 1\}$,

$$\lim_{t\to\infty} q_t = \omega \quad \mathbb{P}_\omega\text{-almost surely.}$$

That is, the outside observer's belief eventually converges to the truth. As we explain below in Section 4.1, when all signals are publicly observable, asymptotic learning of informativeness holds regardless of the underlying information structure. In Section 6.1, we show how this notion connects to other existing concepts of learning.

3. Relative Tail Thickness

To study the conditions for achieving asymptotic learning of informativeness, it is important to understand the concept of relative tail thickness, which captures the relative likelihood of generating extreme signals from different sources. For any pair of CDFs (F_0, F_θ) where $\theta \in \{\mathfrak{g}, \mathfrak{b}\}$ and some $x \in \mathbb{R}_+$, we denote their corresponding ratios by

$$L_\theta(x) := \frac{F_0(-x)}{F_\theta(-x)} \quad \text{and} \quad R_\theta(x) := \frac{1 - F_0(x)}{1 - F_\theta(x)}.$$

For large x, these represent the left and right tail ratios of F_0 over F_θ.

Formally, we say the uninformative signals have *fatter tails* than the informative signals if there exists an $\varepsilon > 0$ such that

$$\begin{aligned} & L_\mathfrak{b}(x) \geq \varepsilon \quad \text{for all } x \text{ large enough,} \\ \text{and} \quad & R_\mathfrak{g}(x) \geq \varepsilon \quad \text{for all } x \text{ large enough.} \end{aligned}$$

We say the uninformative signals have *thinner tails* than the informative signals if there exists an $\varepsilon > 0$ such that

$$\begin{aligned} \text{either} \quad & L_\mathfrak{g}(x) \leq 1/\varepsilon \quad \text{for all } x \text{ large enough,} \\ \text{or} \quad & R_\mathfrak{b}(x) \leq 1/\varepsilon \quad \text{for all } x \text{ large enough.} \end{aligned}$$

That is, for the uninformative signals to have fatter tails, both their corresponding left *and* right tail ratios must eventually be bounded from below. Conversely, for the uninformative signals to have thinner tails, at least one of the tail ratios—either left *or* right—must eventually be bounded from above. The "and" condition is important for defining fatter tails, as it ensures that extreme signals of both types (positive and negative) occur frequently enough to prevent agents from reaching an action consensus. In contrast, the "or" condition for thinner tails requires only the infrequency of at least one type of extreme signal, thus allowing agents to reach a consensus with positive probability. As we will discuss in Section 4.2, determining whether it is possible for agents to reach a consensus under an uninformative source is key to understanding whether asymptotic learning of informativeness holds.

Note that the first condition $L_\mathfrak{b}(x) \geq \varepsilon$ implies that $L_\mathfrak{g}(x) \geq \varepsilon$. This follows from the well-known fact that $F_\mathfrak{g}$ exhibits first-order stochastic dominance over $F_\mathfrak{b}$, i.e., $F_\mathfrak{g}(x) \leq$

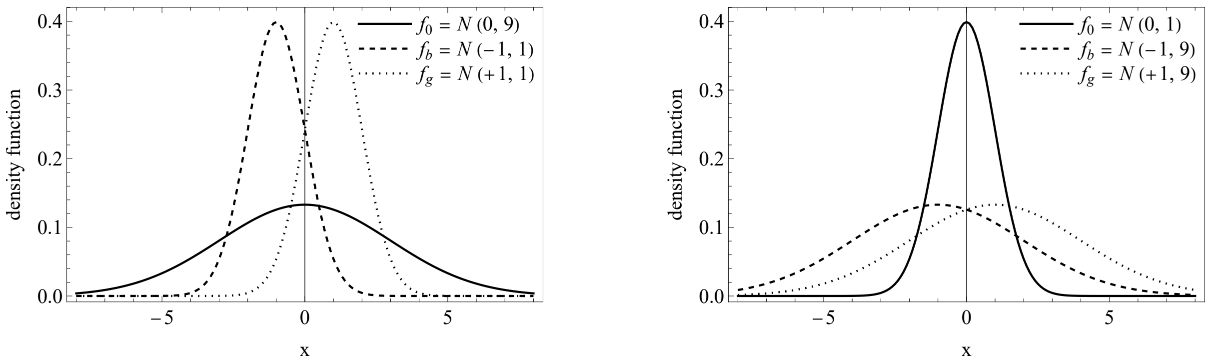

FIGURE 3.1. Probability density functions of informative Gaussian signals and their corresponding uninformative Gaussian signals with fatter tails (on the left) and with thinner tails (on the right).

$F_{\mathfrak{b}}(x)$ for all $x \in \mathbb{R}$ (see, e.g., Smith and Sørensen, 2000; Chamley, 2004; Rosenberg and Vieille, 2019). Similar statements apply to the remaining three conditions. Furthermore, notice that the uninformative signals cannot have both fatter and thinner tails simultaneously, as $F_{\mathfrak{g}}$ and $F_{\mathfrak{b}}$ represent distributions of the agent's private log-likelihood ratio.[6] However, there are distributions under which the uninformative signals have neither fatter nor thinner tails. For these cases, we characterize the conditions for asymptotic learning of informativeness in the canonical Gaussian environment (see Section 4.3).

Intuitively, compared to informative signals, uninformative signals with fatter tails are more likely to exhibit extreme values. Thus, by Bayes' Theorem, observing an extreme signal suggests that the source is uninformative. In contrast, uninformative signals with thinner tails tend to exhibit moderate values, so observing an extreme signal in this case suggests that the source is informative. Next, we provide three examples of uninformative signals with either fatter or thinner tails.

Example 1 (Gaussian Signals)**.** Consider the case where $F_{\mathfrak{g}}$ is normal with mean $+1$ and unit standard deviation and $F_{\mathfrak{b}}$ is also normal with mean -1 and unit standard deviation.

Suppose that F_0 has zero mean. If it has a standard deviation of 3, then the uninformative signals have fatter tails. In this scenario, observing an extreme signal, such as anything around 11, indicates that it is much more likely that the source is uninformative—where the signal is less than $4-\sigma$ away—than for an informative $10-\sigma$ signal to be generated under $F_{\mathfrak{g}}$. Conversely, if the standard deviation of F_0 is $1/3$, the uninformative signals have thinner tails, and thus an extreme signal suggests that the source is informative. A graphical example of uninformative Gaussian signals with fatter and thinner tails is shown in Figure 3.1.

Example 2 (First-Order Stochastically Dominated (FOSD) Signals)**.** Suppose that F_0 first-order stochastically dominates $F_{\mathfrak{g}}$—i.e., $F_0(x) \leq F_{\mathfrak{g}}(x)$ for all $x \in \mathbb{R}$. In other words, the uninformative signals are more likely to exhibit high values than the informative

[6]In particular, $F_{\mathfrak{g}}$ and $F_{\mathfrak{b}}$ satisfy the following inequality: $\mathrm{e}^{x} F_{\mathfrak{g}}(-x) \leq F_{\mathfrak{b}}(-x)$, for all $x \in \mathbb{R}$.

signals associated with the good state. This immediately implies that $L_{\mathfrak{g}}(x) \leq 1$ for all $x \in \mathbb{R}$, so the uninformative signals have thinner tails than the informative ones.

Now suppose an extreme positive signal is observed. It is highly unlikely that the source is informative and associated with the bad state. Instead, it suggests that the source is more likely to be either uninformative or informative but associated with the good state. However, one cannot exclude either possibility based on such an observation, so it does not provide conclusive evidence about the source's informativeness.

Example 3 (Mixture Signals)**.** For any pair of distributions $(F_{\mathfrak{g}}, F_{\mathfrak{b}})$ and $\alpha \in (0,1)$, let $F_0 = \alpha F_{\mathfrak{g}} + (1-\alpha) F_{\mathfrak{b}}$. Observe that the uninformative signals represented by F_0 have fatter tails than the informative signals.[7] In particular, when $\alpha = 1/2$, the corresponding mixture distribution $F_0 = (F_{\mathfrak{g}} + F_{\mathfrak{b}})/2$ coincides with the unconditional distribution of ℓ under an informative source. Thus, one can think of it as representing an uninformative source that is a priori indistinguishable from the informative one.

Alternatively, the mixture distribution F_0 can also be viewed as generating conditionally i.i.d. signals, but conditioned on a different state η rather than on θ. Specifically, in each period a new state η is randomly drawn from the same set $\{\mathfrak{g}, \mathfrak{b}\}$, independent of θ. With probability α, we have $\eta = \mathfrak{g}$ and a signal is drawn from $F_{\mathfrak{g}}$; with the complementary probability, $\eta = \mathfrak{b}$ and a signal is drawn from $F_{\mathfrak{b}}$. Thus, the mixture distribution provides a way to construct an uninformative source directly from the informative ones.

4. Main Results

4.1. **A Benchmark.** As a benchmark, we briefly discuss the case where all signals are observed by the outside observer.[8] Depending on the realizations of θ and ω, these signals are distributed according to either $\mu_{\mathfrak{g}}$, $\mu_{\mathfrak{b}}$ or μ_0. Given that these measures are distinct, as the sample size grows, this observer eventually learns which distribution is being sampled. Formally, at time t, the empirical distribution of the signals $\hat{\mu}_t$ assigns to a measurable set $B \subseteq S$ the probability

$$\hat{\mu}_t(B) := \frac{1}{t} \sum_{\tau=1}^{t} \mathbb{1}(s_\tau \in B).$$

Conditional on both states ω and θ, this is the empirical mean of i.i.d. Bernoulli random variables. Hence, by the strong law of large numbers, for every measurable set $B \subseteq S$,

$$\lim_{t \to \infty} \hat{\mu}_t(B) = \mu_{(\omega,\theta)}(B) \quad \text{almost surely,}$$

where $\mu_{(1,\mathfrak{g})} = \mu_{\mathfrak{g}}$, $\mu_{(1,\mathfrak{b})} = \mu_{\mathfrak{b}}$ and $\mu_{(\cdot,0)} = \mu_0$.

[7] To see this, fix any constant $\alpha \in (0,1)$ and let $F_0 = \alpha F_{\mathfrak{g}} + (1-\alpha) F_{\mathfrak{b}}$. Since CDFs always take nonnegative values, for any $x \in \mathbb{R}$, $F_0(x) \geq (1-\alpha) F_{\mathfrak{b}}(x)$. Similarly, $1 - F_0(x) = \alpha(1 - F_{\mathfrak{g}}(x)) + (1-\alpha)(1 - F_{\mathfrak{b}}(x)) \geq \alpha(1 - F_{\mathfrak{g}}(x))$. Let $\varepsilon = \min\{\alpha, 1-\alpha\}$. By definition, F_0 has fatter tails.

[8] Equivalently, one can let the outside observer observe all agents' actions in addition to their signals. Since actions contain no additional payoff relevant information, it suffices to only consider the signals.

Thus, regardless of the underlying signal-generating process, any uncertainty concerning the informativeness of the source is eventually resolved when all signals are publicly observable.

4.2. **Public Actions.** We now turn to our main setting, where signals remain private and only actions are observed. In contrast to the public signal benchmark, our main result (Theorem 1) shows that achieving asymptotic learning of informativeness is no longer guaranteed. Instead, the key determinant of learning is the relative tail thickness between the uninformative and informative signals, introduced in Section 3.

Theorem 1. *When the uninformative signals have fatter (thinner) tails than the informative signals, asymptotic learning of informativeness holds (fails).*

This result demonstrates that an observer learning from agents' actions can eventually distinguish between informative and uninformative sources, but only if the uninformative source generates signals with greater dispersion. In contrast, when the uninformative signals are relatively concentrated, this differentiation becomes impossible.

For example, consider informative signals that follow a normal distribution with unit variance and mean $+1$ and -1, respectively. When the uninformative signals follow a normal distribution with a variance of 2, they have fatter tails, and thus Theorem 1 implies that asymptotic learning of informativeness holds. In contrast, when the uninformative signals follow a normal distribution with a variance of $1/2$, they have thinner tails, and Theorem 1 implies that such learning fails.

The idea behind our proof of Theorem 1 is as follows. First consider the case where the source is informative. In this case, the likelihood that a single extreme signal overturns a long consensus decays rapidly. Consequently, agents will eventually choose the correct action since they always treat signals as informative—and here the signals are indeed informative. Now suppose the source is uninformative and, instead of reaching a consensus, agents continue to disagree indefinitely, leading to both actions being taken infinitely often. If this were the case, an outside observer would eventually be able to distinguish between informative and uninformative sources, as these sources induce distinct behavioral patterns among agents. We formally establish this mechanism of perpetual disagreement under an uninformative source—which drives asymptotic learning—in Section 5.1 (see Proposition 2). Whether these disagreements persist or not depends on whether the tails of the uninformative signals are thick enough to generate these overturning extreme signals.

In summary, when the tails of uninformative signals are sufficiently thick, overturning extreme signals occur frequently enough so that disagreements persist. Conversely, when the tails are relatively thin, these signals are less likely to occur, and disagreements

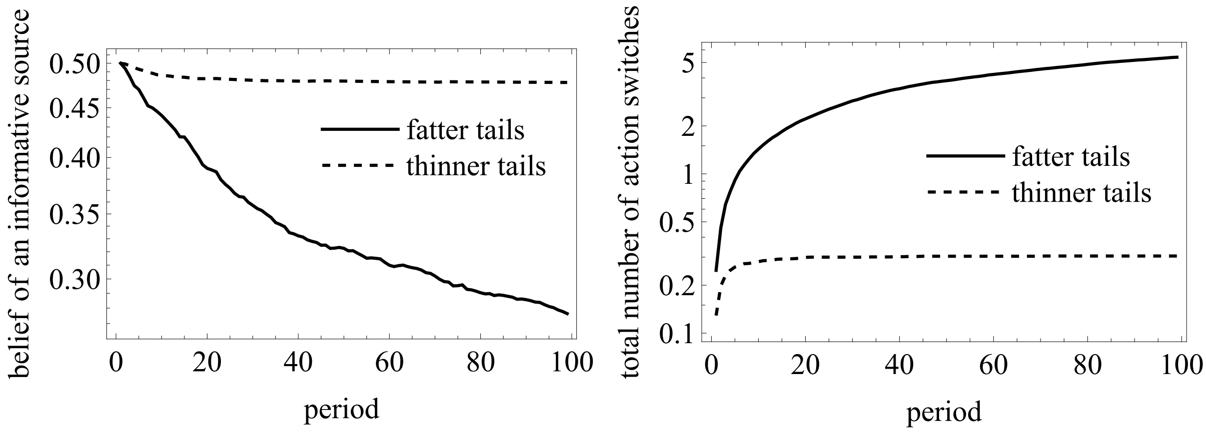

FIGURE 4.1. The belief that the source is informative (on the left) and the total number of action switches (on the right) under uninformative signals with fatter and thinner tails.

eventually cease. Hence, we conclude that the relative tail thickness between uninformative and informative signals plays an important role in determining the achievement of asymptotic learning of informativeness.

Numerical Simulation. To further illustrate our main result, we use Monte Carlo simulations to numerically examine the belief process of an outside observer and the corresponding action switches among agents. We fix the pair of informative Gaussian signal distributions to be $(\mu_{\mathfrak{g}}, \mu_{\mathfrak{b}})$ where $\mu_{\mathfrak{g}} = \mathcal{N}(+1, 2)$ and $\mu_{\mathfrak{b}} = \mathcal{N}(-1, 2)$. These processes are simulated under two scenarios: one with fatter-tailed uninformative Gaussian signals with distribution $\mu_0 = \mathcal{N}(0, 3)$, and the other with thinner-tailed uninformative Gaussian signals with distribution $\mu_0 = \mathcal{N}(0, 1)$.[9] For each uninformative source, we conduct these simulations 1,000 times, redrawing the payoff-relevant state each time, and calculate the averages for each period. This yields approximations for the expected belief and the expected total number of action switches in the presence of uninformative signals. Figure 4.1 displays the results of these simulations.

What immediately stands out is that under fatter-tailed uninformative signals, the belief of the outside observer that the source is informative decreases much faster compared to thinner-tailed uninformative signals. Under fatter-tailed uninformative signals, by period 60, this belief is approximately 0.3, which is less than two-thirds of the belief observed under thinner-tailed uninformative signals. These findings align with the predictions of Theorem 1, suggesting that in the former case the observer will eventually learn that the source is uninformative. In contrast, with thinner-tailed uninformative signals, the decline in the belief about the source's informativeness plateaus around 0.48 after period 20, suggesting that the observer will not be able to learn that the source is uninformative.

[9]Note that in this case, the agent's private log-likelihood ratio $\ell = s$, so $F_0, F_{\mathfrak{g}}$ and $F_{\mathfrak{b}}$ have the same distribution as μ_0, $\mu_{\mathfrak{g}}$ and $\mu_{\mathfrak{b}}$, respectively.

As discussed earlier, the intuition behind Theorem 1 is that the uninformative signals with fatter tails have a higher probability of being extreme, which in turn prevents agents from converging to a consensus. This phenomenon is evident in the right plot of Figure 4.1, where the average total number of action switches under fatter-tailed uninformative signals increases over time. In contrast, under the thinner-tailed uninformative signals, the total number of switches plateaus in a short amount of time, suggesting that agents eventually reach an action consensus.

4.3. **Gaussian Private Signals.** We now focus on the canonical environment in which all signals are normally distributed. In this setting, our main result applies directly whenever the variance of the uninformative signals, τ^2, differs from that of the informative signals, σ^2. This is because the relative tail thickness of normal distributions depends solely on their variances,[10] and so it immediately follows from Theorem 1 that asymptotic learning of informativeness holds if $\tau > \sigma$ and fails if $\tau < \sigma$.

When all normal signals share the same variance, however, the uninformative signals have neither fatter nor thinner tails, rendering Theorem 1 inapplicable.[11] To complement our main result, the next proposition shows that in this case, asymptotic learning of informativeness is achieved if and only if the uninformative signals are symmetric around zero.

Proposition 1. *Suppose all private signals are Gaussian with the same variance, where informative signals have means -1 and $+1$, and uninformative signals have mean $m_0 \in (-1, 1)$. Then, asymptotic learning of informativeness holds if and only if $m_0 = 0$.*

The intuition is similar to that behind our main result. When all signals share the same variance, any deviation of the mean of F_0 from zero shifts it closer to either $F_{\mathfrak{g}}$ or $F_{\mathfrak{b}}$, making the uninformative signals resemble the corresponding informative signals. As a result, agents also reach an action consensus under F_0, and this prevents an outside observer from fully distinguishing between informative and uninformative sources.

5. Analysis

In this section we first examine how agents update their beliefs. We present two standard yet useful properties of their belief updating process, namely, the *overturning principle* and *stationarity*. We then use two important phenomena—immediate agreement and perpetual disagreement—to characterize asymptotic learning of informativeness. This leads to a proof sketch of Theorem 1 at the end of this section.

[10]For more details, see Lemma 8 in the appendix.

[11]To see this, suppose F_0, $F_{\mathfrak{g}}$, and $F_{\mathfrak{b}}$ are normal distributions with the same variance and means 0, 1, and -1, respectively. As $x \to \infty$, both $F_0(-x)$ and $F_{\mathfrak{b}}(-x)$ approach zero, but the former does so much faster than the latter, leading $L_{\mathfrak{b}}(x) \to 0$. Hence, F_0 does not have fatter tails. Similarly, both $L_{\mathfrak{g}}(x)$ and $R_{\mathfrak{b}}(x)$ tend to infinity as $x \to \infty$, and so F_0 does not have thinner tails either.

5.1. **Agents' Beliefs Dynamics.** Since agents always act as if signals are informative (Lemma 1), we focus on how agents update their beliefs conditioned on an informative source. Let $\pi_t := \mathbb{P}_1[\theta = \mathfrak{g}|H_t]$ denote agent t's public belief conditional on the action history H_t and $\omega = 1$. Its log-likelihood ratio (LLR) is

$$r_t := \log \frac{\pi_t}{1-\pi_t} = \log \frac{\mathbb{P}_1[\theta = \mathfrak{g}|H_t]}{\mathbb{P}_1[\theta = \mathfrak{b}|H_t]}.$$

Similarly, let

$$L_t := \log \frac{\mathbb{P}_1[\theta = \mathfrak{g}|H_t, s_t]}{\mathbb{P}_1[\theta = \mathfrak{b}|H_t, s_t]}$$

be the LLR of agent t's posterior belief, conditional on H_t, her private signal s_t and $\omega = 1$. Recall that ℓ_t is agent t's private LLR conditional on $\omega = 1$. By Bayes' rule, we have

$$L_t = r_t + \ell_t.$$

From (2) we know that in equilibrium, agent t chooses action $\mathfrak{g}$ if $\ell_t \geq -r_t$ and action $\mathfrak{b}$ if $\ell_t < -r_t$. Hence, conditioned on θ and the event $\omega = 1$, the probability that $a_t = \mathfrak{g}$ is $1 - F_\theta(-r_t)$ and the probability that $a_t = \mathfrak{b}$ is $F_\theta(-r_t)$. Consequently, the agents' public LLRs evolve as follows:

$$r_{t+1} = r_t + D_\mathfrak{g}(r_t) \quad \text{if } a_t = \mathfrak{g}, \tag{3}$$

$$r_{t+1} = r_t + D_\mathfrak{b}(r_t) \quad \text{if } a_t = \mathfrak{b}, \tag{4}$$

where

$$D_\mathfrak{g}(r) := \log \frac{1 - F_\mathfrak{g}(-r)}{1 - F_\mathfrak{b}(-r)} \quad \text{and} \quad D_\mathfrak{b}(r) := \log \frac{F_\mathfrak{g}(-r)}{F_\mathfrak{b}(-r)}.$$

Since the agent's public belief follows a standard Bayesian updating process, it satisfies two well-known properties: the *overturning principle* (Sørensen, 1996; Smith and Sørensen, 2000) and *stationarity*. The overturning principle asserts that a single action is sufficient to change the verdict of π_t. Stationarity means that the value of π_t captures all past information about the payoff-relevant state, independent of time. For completeness, we state them below.

Lemma 2 (Overturning Principle)**.** *For each agent t, if $a_t = \mathfrak{g}$, then $\pi_{t+1} \geq 1/2$. Similarly, if $a_t = \mathfrak{b}$, then $\pi_{t+1} \leq 1/2$.*

We further write $\mathbb{P}_{\tilde{\omega},\tilde{\theta},\pi}$ to denote the conditional probability distribution given the pair of state realizations $(\tilde{\omega}, \tilde{\theta})$ while highlighting the different values of the prior π.

Lemma 3 (Stationarity)**.** *For any fixed sequence $(b_\tau)_{\tau=1}^k$ of k actions in $\{\mathfrak{g}, \mathfrak{b}\}$, any prior $\pi \in (0,1)$ and any pair $(\tilde{\omega}, \tilde{\theta}) \in \{0,1\} \times \{\mathfrak{g}, \mathfrak{b}\}$*

$$\mathbb{P}_{\tilde{\omega},\tilde{\theta}}[a_{t+1} = b_1, \ldots, a_{t+k} = b_k | \pi_t = \pi] = \mathbb{P}_{\tilde{\omega},\tilde{\theta},\pi}[a_1 = b_1, \ldots, a_k = b_k].$$

That is, regardless of the source's informativeness, if agent t's public belief is equal to π, then the probability of observing a sequence $(b_1, \ldots, b_k)$ of actions of length k is the same as observing this sequence starting from time 1, given that the agents' prior on the payoff-relevant state is π. Stationarity holds in our model because, regardless of the informativeness of the source, agents always update their public log-likelihood ratios according to either (3) or (4).

5.2. **Perpetual Disagreement and Immediate Agreement.** We now introduce two events that are important for characterizing asymptotic learning of informativeness. The first is *perpetual disagreement*. This is the event in which agents' actions never converge, which we denote by $\{S = \infty\}$, where $S = \sum_{t=1}^{\infty} \mathbb{1}(a_t \neq a_{t+1})$ is the total number of action switches.

The following proposition establishes the key mechanism underlying Theorem 1: it shows that learning about informativeness is equivalent to the almost sure occurrence of perpetual disagreement when the source is uninformative.

Proposition 2. *Asymptotic learning of informativeness holds if and only if conditioned on $\omega = 0$, the event $\{S = \infty\}$ occurs almost surely.*

Intuitively, if agents never reach a consensus under an uninformative source, then the outside observer infers that the source must be uninformative since agents would have reached a consensus under an informative source. Conversely, suppose agents could also reach a consensus when the source is uninformative. This implies that action convergence is plausible under both informative and uninformative sources. Consequently, the observer can no longer be sure of the source's informativeness.

The second important event is *immediate agreement*, where agents reach a consensus from the outset, i.e., $\{a_1 = a_2 = \ldots = a\}$. We write $\{\bar{a} = a\}$ for this event, where a is the consensus action. The next lemma shows that when the source is informative, immediate agreement on the wrong action is impossible, whereas immediate agreement on the correct action is possible. For brevity, we state this result only for the case where $\theta = \mathfrak{g}$, as the analogous statements hold for $\theta = \mathfrak{b}$ by symmetry.

Lemma 4. *Conditioned on $\omega = 1$ and $\theta = \mathfrak{g}$, the following two conditions hold:*

(i) The event $\{\bar{a} = \mathfrak{b}\}$ occurs with probability zero.

(ii) The event $\{\bar{a} = \mathfrak{g}\}$ occurs with positive probability for some prior $\pi \in (0, 1)$.

The first part of Lemma 4 holds since, in our model, conditioned on an informative source, all but finitely many agents take the correct action. This immediately implies that agents cannot reach a consensus on the wrong action from the outset. Likewise, the second part of Lemma 4 holds because if agents eventually reach a consensus on the correct action, by stationarity, they can also do so immediately at least for some prior.[12]

[12]In fact, part (ii) of Lemma 4 holds not only for some prior but also for any prior. This can be seen by applying a similar argument used in the proof of Lemma 7 in the appendix. We omit the stronger statement here, as it is not required to prove our main result.

Next, we focus on the immediate agreement event conditioned on an uninformative source. Recall that in Proposition 2, we establish the relationship between learning about informativeness and perpetual disagreement. Building on this, we now characterize asymptotic learning of informativeness in terms of immediate agreement, which is a crucial step in proving our main result.

Proposition 3. *Asymptotic learning of informativeness holds if and only if conditioned on $\omega = 0$, both events $\{\bar{a} = \mathfrak{g}\}$ and $\{\bar{a} = \mathfrak{b}\}$ occur with zero probability.*

This proposition states that achieving asymptotic learning of informativeness is equivalent to the absence of immediate agreement starting from a uniform prior given an uninformative source. The proof of Proposition 3 utilizes the idea that the agent's belief updating process is eventually monotonic—a technical property that we establish in Lemma 6 in the appendix. This property ensures that if immediate agreement is impossible for some prior, e.g., the uniform prior, it becomes impossible for *all* prior. By stationarity, this then implies that agents cannot reach a consensus on any actions for any prior. Consequently, perpetual disagreement occurs with probability one, and it follows from Proposition 2 that this is equivalent to achieving asymptotic learning of informativeness.

We have now reduced the problem of learning about informativeness to determining the possibility of immediate agreement conditioned on an uninformative source, which is much easier to analyze. Specifically, conditioned on the event $\{\bar{a} = \mathfrak{g}\}$, let $r_t^{\mathfrak{g}}$ denote the deterministic process of r_t based on (3). Recall that agent t chooses action $\mathfrak{g}$ if $\ell_t \geq -r_t$ and action $\mathfrak{b}$ otherwise. Consequently, the probability of $\{\bar{a} = \mathfrak{g}\}$ is equal to the probability that $\ell_t \geq -r_t^{\mathfrak{g}}$ for all $t \geq 1$. Moreover, conditioned on an uninformative source, since private signals are i.i.d., the corresponding private log-likelihood ratios are also i.i.d. Thus, conditioned on $\omega = 0$, the probability of immediate agreement on action $\mathfrak{g}$ is

$$\mathbb{P}_0[\bar{a} = \mathfrak{g}] = \prod_{t=1}^{\infty}(1 - F_0(-r_t^{\mathfrak{g}})).$$

To determine whether the above probability is positive or zero, by a standard approximation argument, it is equivalent to examining whether the sum of the probabilities of the following events is finite or infinite:

$$\mathbb{P}_0[\bar{a} = \mathfrak{g}] > 0 \ (= 0) \ \Leftrightarrow \ \sum_{t=1}^{\infty} F_0(-r_t^{\mathfrak{g}}) < \infty \ (= \infty). \tag{5}$$

By symmetry of the model, we also have

$$\mathbb{P}_0[\bar{a} = \mathfrak{b}] > 0 \ (= 0) \ \Leftrightarrow \ \sum_{t=1}^{\infty} \big(1 - F_0(r_t^{\mathfrak{g}})\big) < \infty \ (= \infty). \tag{6}$$

In sum, asymptotic learning of informativeness holds if and only if conditioned on the source being uninformative, the probability of generating extreme signals decreases slowly enough so that both sums in (5) and (6) are infinite. As we discuss below, for uninformative signals with fatter tails, these sums diverge, and for signals with thinner tails, at least one of these sums converges.

5.3. **Proof Sketch of Theorem 1.** We conclude this section by providing a sketch of the proof of Theorem 1. On the one hand, by part (i) of Lemma 4, the probability of generating extreme signals that match the payoff-relevant state decreases relatively slowly under informative signals. Hence, if the source is uninformative and generates signals with fatter tails, this probability declines even more slowly. Consequently, both the sums in (5) and (6) are infinite, which means that no immediate agreement is possible. Thus, by Proposition 3, asymptotic learning of informativeness holds. On the other hand, by part (ii) of Lemma 4, the probability of generating extreme signals that mismatch the state decreases relatively fast under informative signals. Hence, if the source is uninformative and generates signals with thinner tails, this probability decreases even more rapidly, at least for some type of extreme signals. Consequently, either the sum in (5) or the sum in (6) (or both) is finite, which means that immediate agreement on some action is possible. Hence, by Proposition 3, asymptotic learning of informativeness fails.

6. Discussion

6.1. **Connections to Other Notions of Learning.** We now discuss the relationship between our notion of asymptotic learning (of informativeness) and two existing notions of learning in the literature. One concerning the convergence of actions to the truth, is known as *correct herding*; the other, concerning the convergence of beliefs to the truth, is called *complete learning*. Both notions focus on the payoff-relevant state.

Formally, we say that *correct herding* holds if $\lim_{t\to\infty} a_t = \theta$ almost surely. That is, all but finitely many agents almost surely take the correct action. Recall that $H_t = (a_1, \ldots, a_{t-1})$ is the action history from time 1 to $t-1$. Define the *social belief* at time t as $p_t = \mathbb{P}[\theta = \mathfrak{g} | H_t]$. We say that *complete learning* holds if the social belief converges to the truth, i.e., p_t almost surely converges to one when $\theta = \mathfrak{g}$ and to zero when $\theta = \mathfrak{b}$.

As is well known, when the source is always informative and generates unbounded signals, both correct herding and complete learning hold. In our model with information uncertainty, however, this need not be the case. When the source is informative, correct herding still occurs, as agents' behavior mirrors that in a standard social learning model. Yet, as we show below, this no longer implies complete learning: even when agents converge to the correct action, they remain uncertain about its correctness if asymptotic learning of informativeness fails. By contrast, when the source is uninformative, achieving asymptotic learning of informativeness is equivalent to the social belief converging to the uniform prior.

Taken together, these results highlight the connections between learning about informativeness and learning about the payoff-relevant state, as formalized in the following proposition.

Proposition 4. *Asymptotic learning of informativeness holds if and only if*

(i) conditioned on $\omega = 1$, complete learning holds; and

(ii) conditioned on $\omega = 0$, $\lim_{t\to\infty} p_t = 1/2$ almost surely.

The first part of this proposition uses the idea of an outsider who only observes agents' actions, so his belief coincides with the social belief. As long as this outside observer remains uncertain about the source's informativeness, he cannot fully trust the social information derived from agents' actions. Conversely, once the source's informativeness is confirmed, this information becomes highly accurate, allowing the observer's belief to converge to the truth.

The second part follows a similar approach and applies the result of Proposition 2. One direction is straightforward: if the observer learns that the source is uninformative, then the social belief must converge to the prior, since the action history contains no information about the payoff-relevant state. For the other direction, suppose, by contraposition, that the observer does not learn that the source is uninformative. Then Proposition 2 implies that agents may eventually reach a consensus on some action. Since this is also possible under an informative source, the observer cannot fully rule out the possibility that agents' actions are informative. As a result, his belief about the payoff-relevant state is not guaranteed to remain at the prior.

6.2. **Non-Uniform Prior.** So far, we have assumed that all agents hold a uniform prior over the payoff-relevant state. This implies that agents behave as if all signals are informative and, importantly, that they always respond to their private signals. The next proposition shows that such responsiveness is essential for an outside observer seeking to infer the source's informativeness from actions.

Proposition 5. *Asymptotic learning of informativeness fails whenever the prior on the payoff-relevant state θ is non-uniform.*

Thus, when agents have a non-uniform prior—even if it is weak and close to indifference—no signal structure can guarantee asymptotic learning of informativeness. Again, the proof of this proposition uses ideas similar to those in Proposition 2. We refer to the action favored by the prior as the default action. Suppose to the contrary that asymptotic learning of informativeness holds. In that case, when the source is uninformative, agents will eventually learn this and choose the default action. However, when the source is informative, agents will also take the default action with positive probability, since it is the correct action in some states. But this contradicts the hypothesis that the outside observer's belief about the source's informativeness almost surely converges to the truth, and thus we conclude that asymptotic learning of informativeness fails.

7. Conclusion

In this paper, we study the sequential social learning problem in the presence of a potentially uninformative source. We show that achieving asymptotic learning of informativeness—where an outside observer eventually discerns whether the source is informative or uninformative—is not guaranteed, and it crucially depends on the relationship between the conditional distributions of the private signals. In particular, it hinges on the relative tail distribution of signals: when uninformative signals have fatter tails than their informative counterparts, the observer eventually learns the source's informativeness; when they have thinner tails, such learning fails. We also characterize the conditions for asymptotic learning of informativeness in the canonical case of Gaussian private signals, where the relative tail thickness is incomparable.

More generally, our analysis suggests that irregular behavior, such as an action switch or a disagreement following a prolonged sequence of identical actions, is the driving force behind learning. Unlike in the public-signal benchmark where asymptotic learning of informativeness is always achieved, when agents' signals remain private an outside observer can only infer that the source is uninformative by observing action switches. We show that, conditioned on an uninformative source, the observer eventually learns its uninformativeness if and only if action switches accumulate indefinitely (or disagreements persist). We view this characterization as the key mechanism behind our main result.[13]

We make the symmetry assumption on the informative distributions $F_{\mathfrak{g}}$ and $F_{\mathfrak{b}}$ for expositional purposes. Since our notion of relative tail thickness does not require this symmetry, our main result can be easily extended to the non-symmetric case. A more substantial assumption is that, under an uninformative source, agents are indifferent between the two actions. As shown in Section 6.2, relaxing this assumption causes asymptotic learning of informativeness to fail, as a non-uniform prior induces a default action which hinders perpetual disagreement. We see this failure as a result of agents eventually ceasing to respond to their private signals, a mechanism that is reminiscent of the one underlying information cascades.

A limitation of our results is that they apply only asymptotically. Our numerical simulations suggest that in the case of uninformative Gaussian signals with fatter tails, the asymptotics can already kick in relatively early in the process. Thus, it would also be interesting to understand the speed at which an outside observer learns about the informativeness of the source. Another promising direction for future research is to explore varying degrees of informativeness, beyond the extreme cases considered in this paper. For example, instead of assuming the source is either informative or completely uninformative, one could consider weakly or strongly informative sources and ask whether

[13]In the context of scientific paradigms, this mechanism is reminiscent of Kuhn (1962)'s idea that the accumulation of anomalies may trigger scientific revolutions and paradigm shifts. See Ba (2022) for a study on the rationale behind the persistence of a misspecified model, e.g., a wrong scientific paradigm.

learning about the source's informativeness can still be achieved. We conjecture that our current notion of asymptotic learning would fail in this case, as the "perpetual disagreement" argument would not hold, given that a weakly informative source with unbounded signals would still lead to action convergence. However, it seems plausible that if the frequency of action switches differs significantly across different informative sources, the observer could almost learn to distinguish between them.[14] We leave these interesting questions for future research.

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011): "Bayesian learning in social networks," *The Review of Economic Studies*, 78, 1201–1236.

Arieli, I., Y. Babichenko, S. Müller, F. Pourbabaee, and O. Tamuz (2023): "The Hazards and Benefits of Condescension in Social Learning," *arXiv preprint arXiv:2301.11237*.

Arieli, I. and M. Mueller-Frank (2021): "A general analysis of sequential social learning," *Mathematics of Operations Research*, 46, 1235–1249.

Ba, C. (2022): "Robust Misspecified Models and Paradigm Shift," Tech. rep., working paper.

Banerjee, A. V. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107, 797–817.

Bikhchandani, S., D. Hirshleifer, O. Tamuz, and I. Welch (2021): "Information cascades and social learning," Tech. rep., National Bureau of Economic Research.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 992–1026.

Bohren, J. A. (2016): "Informational herding with model misspecification," *Journal of Economic Theory*, 163, 222–247.

Bohren, J. A. and D. N. Hauser (2021): "Learning with heterogeneous misspecified models: Characterization and robustness," *Econometrica*, 89, 3025–3077.

Çelen, B. and S. Kariv (2004): "Observational learning under imperfect information," *Games and Economic behavior*, 47, 72–86.

Chamley, C. (2004): *Rational herds: Economic models of social learning*, Cambridge University Press.

Chen, J. Y. (2022): "Sequential learning under informational ambiguity," *Available at SSRN 3480231*.

[14]For example, one way to generate different frequencies of action switches is to allow multiple draws of the payoff-relevant state. We conjecture that this would provide more information to the outside observer, potentially speeding up learning. This is because, compared to an uninformative source, agents' actions are more likely to change under an informative source, as they are more likely to receive opposing signals when the payoff-relevant state changes.

Frick, M., R. Iijima, and Y. Ishii (2020): "Misinterpreting others and the fragility of social learning," *Econometrica*, 88, 2281–2328.

——— (2023): "Belief convergence under misspecified learning: A martingale approach," *The Review of Economic Studies*, 90, 781–814.

Golub, B. and E. Sadler (2017): "Learning in social networks," *Available at SSRN 2919146*.

Hann-Caruthers, W., V. V. Martynov, and O. Tamuz (2018): "The speed of sequential asymptotic learning," *Journal of Economic Theory*, 173, 383–409.

Herrera, H. and J. Hörner (2012): "A necessary and sufficient condition for information cascades," *Columbia and Yale mimeo.*

Kartik, N., S. Lee, T. Liu, and D. Rappoport (2022): "Beyond Unbounded Beliefs: How Preferences and Information Interplay in Social Learning," Tech. rep.

Kuhn, T. S. (1962): *The structure of scientific revolutions*, University of Chicago press.

Liang, A. and X. Mu (2020): "Complementary information and learning traps," *The Quarterly Journal of Economics*, 135, 389–448.

Lobel, I. and E. Sadler (2015): "Information diffusion in networks through social learning," *Theoretical Economics*, 10, 807–851.

Rosenberg, D. and N. Vieille (2019): "On the efficiency of social learning," *Econometrica*, 87, 2141–2168.

Smith, L. and P. Sørensen (2000): "Pathological outcomes of observational learning," *Econometrica*, 68, 371–398.

Smith, L., P. N. Sørensen, and J. Tian (2021): "Informational herding, optimal experimentation, and contrarianism," *The Review of Economic Studies*, 88, 2527–2554.

Sørensen, P. N. (1996): "Rational social learning," Ph.D. thesis, Massachusetts Institute of Technology.

Vives, X. (1993): "How fast do rational agents learn?" *The Review of Economic Studies*, 60, 329–347.

Appendix A. Proofs

Proof of Lemma 1. Recall that in equilibrium, each agent t chooses the action that is most likely to match the state θ conditioned on the available information (H_t, s_t). By Bayes' rule, the relative likelihood between the good state and the bad state for agent t is

$$\frac{\mathbb{P}[\theta = \mathfrak{g}|H_t, s_t]}{\mathbb{P}[\theta = \mathfrak{b}|H_t, s_t]} = \frac{\sum_{\tilde{\omega}\in\{0,1\}} \mathbb{P}_{\tilde{\omega}}[\theta = \mathfrak{g}|H_t, s_t] \cdot \mathbb{P}[\omega = \tilde{\omega}|H_t, s_t]}{\sum_{\tilde{\omega}\in\{0,1\}} \mathbb{P}_{\tilde{\omega}}[\theta = \mathfrak{b}|H_t, s_t] \cdot \mathbb{P}[\omega = \tilde{\omega}|H_t, s_t]}.$$

Note that $\mathbb{P}_0[\theta = \mathfrak{g}|H_t, s_t] = \mathbb{P}_0[\theta = \mathfrak{g}]$ and $\mathbb{P}_0[\theta = \mathfrak{b}|H_t, s_t] = \mathbb{P}_0[\theta = \mathfrak{b}]$ as conditioned on $\omega = 0$, neither the history H_t nor the signal s_t contains any information about the payoff-relevant state θ. Since the states ω and θ are independent of each other and the prior on θ is uniform, $\mathbb{P}_0[\theta = \mathfrak{g}|H_t, s_t] = \mathbb{P}_0[\theta = \mathfrak{b}|H_t, s_t] = 1/2$. Thus, it follows from the above equation that

$$\frac{\mathbb{P}[\theta = \mathfrak{g}|H_t, s_t]}{\mathbb{P}[\theta = \mathfrak{b}|H_t, s_t]} \geq 1 \Leftrightarrow \frac{\mathbb{P}_1[\theta = \mathfrak{g}|H_t, s_t]}{\mathbb{P}_1[\theta = \mathfrak{b}|H_t, s_t]} \geq 1.$$

That is, in equilibrium, each agent chooses the most likely action conditioned on the available information and the source being informative. □

Proof of Lemma 2. For any $t \geq 1$, one has that

$$\pi_{t+1} = \mathbb{E}_1[\mathbb{1}(\theta = \mathfrak{g})|H_{t+1}] = \mathbb{E}_1[\mathbb{E}_1[\mathbb{1}(\theta = \mathfrak{g})|H_t, s_t]|H_{t+1}]],$$

where the second equality follows from the law of the iterated expectation. Thus, if $a_t = \mathfrak{g}$, by Lemma 1, $\mathbb{P}_1[\theta = \mathfrak{g}|H_t, s_t] \geq \mathbb{P}_1[\theta = \mathfrak{b}|H_t, s_t]$. It follows from the above equation that $\pi_{t+1} \geq 1 - \pi_{t+1}$, which implies that $\pi_{t+1} \geq 1/2$. The case where $a_t = \mathfrak{b}$ implies that $\pi_{t+1} \leq 1/2$ follows from a symmetric argument. □

The following simple claim will be useful in proving Proposition 2. It employs an idea similar to the *no introspection principle* in Sørensen (1996). Recall that $q_t = \mathbb{P}[\omega = 1|H_t]$ is the belief that an outside observer assigns to the source being informative based on the history of actions from time 1 to $t-1$ and $\gamma = \mathbb{P}[\omega = 1] \in (0, 1)$ is the prior that the source is informative.

Claim 1. *For any $a \in (0, 1/2)$ and any $b \in (1/2, 1)$,*

$$\mathbb{P}_0[q_t = \tilde{q}] \leq \frac{1-a}{a}\frac{\gamma}{1-\gamma} \cdot \mathbb{P}_1[q_t = \tilde{q}], \text{ for all } \tilde{q} \in [a, 1/2];$$
$$\mathbb{P}_0[q_t = \tilde{q}] \geq \frac{1-b}{b}\frac{\gamma}{1-\gamma} \cdot \mathbb{P}_1[q_t = \tilde{q}], \text{ for all } \tilde{q} \in [1/2, b].$$

Proof. Fix a prior $\gamma \in (0, 1)$. For any $\tilde{q} \in (0, 1)$, let $\tilde{H}_t$ be a history of actions such that the associated belief q_t is equal to $\tilde{q}$. By the law of total expectation,

$$\mathbb{P}[\omega = 1|q_t = \tilde{q}] = \mathbb{E}[\mathbb{E}[\mathbb{1}(\omega = 1)|\tilde{H}_t]|q_t = \tilde{q}] = \mathbb{E}[\tilde{q}|q_t = \tilde{q}] = \tilde{q}.$$

It follows from Bayes' rule that

$$\frac{\mathbb{P}_0[q_t = \tilde{q}]}{\mathbb{P}_1[q_t = \tilde{q}]} = \frac{\mathbb{P}[\omega = 0|q_t = \tilde{q}]}{\mathbb{P}[\omega = 1|q_t = \tilde{q}]} \cdot \frac{\mathbb{P}[\omega = 1]}{\mathbb{P}[\omega = 0]} = \frac{1-\tilde{q}}{\tilde{q}}\frac{\gamma}{1-\gamma}.$$

Since for any $a \in (0, 1/2)$ and any $\tilde{q} \in [a, 1/2]$, $1 \leq \frac{1-\tilde{q}}{\tilde{q}} \leq \frac{1-a}{a}$, it thus follows from the above equation that

$$\mathbb{P}_0[q_t = \tilde{q}] = \frac{1-\tilde{q}}{\tilde{q}}\frac{\gamma}{1-\gamma} \cdot \mathbb{P}_1[q_t = \tilde{q}] \leq \frac{1-a}{a}\frac{\gamma}{1-\gamma} \cdot \mathbb{P}_1[q_t = \tilde{q}].$$

The second inequality follows from an identical argument. □

In the following proofs for Proposition 2 and 4, we use extensively the idea of an outside observer, say observer x, who observes everyone's actions. The information available to him at time t is thus $H_t = (a_1, \ldots, a_{t-1})$ and at time infinity is $H_\infty = \cup_t H_t$. He gets a utility of one if his guess about ω is correct and zero otherwise. Furthermore, we denote by $q_\infty := \mathbb{P}[\omega = 1|H_\infty]$ the belief that he has at time infinity about the event $\omega = 1$. Similarly, we denote by $p_\infty := \mathbb{P}[\theta = \mathfrak{g}|H_\infty]$ the belief that he has at time infinity about the event $\theta = \mathfrak{g}$.

Proof of Proposition 2. *(If direction)* Suppose $\mathbb{P}_0[S = \infty] = 1$. Let a^x_t be the guess that the outside observer x would make to maximize his probability of guessing ω correctly at time t. Fix a large positive integer $k \in \mathbb{N}$. Let $A_t(k)$ denote the event that there have been at least k action switches before time t and denote its complementary event by $A^c_t(k)$. Consider the following strategy $\tilde{a}^x_\infty(k)$ for x at time infinity: $\tilde{a}^x_\infty(k) = 0$ if $A_\infty(k)$ occurs and $\tilde{a}^x_\infty(k) = 1$ otherwise. The probability of x guessing it correctly under this strategy is

$$\mathbb{P}[\tilde{a}^x_\infty(k) = \omega] = \mathbb{P}_0[A_\infty(k)] \cdot \mathbb{P}[\omega = 0] + \mathbb{P}_1[A^c_\infty(k)] \cdot \mathbb{P}[\omega = 1]. \tag{A.1}$$

Since $\mathbb{P}_{1,\theta}[a_t \to \theta] = 1$, it follows that for all k large enough,

$$\mathbb{P}_1[A^c_\infty(k)] = 1. \tag{A.2}$$

By assumption, $\mathbb{P}_0[S = \infty] = 1$, and thus $\mathbb{P}_0[A_\infty(k)] = 1$ for all k. Together, (A.1) and (A.2) imply that for all k large enough, $\mathbb{P}[\tilde{a}^x_\infty(k) = \omega] = 1$.

Meanwhile, the optimal strategy for x at time infinity is the following: $a^x_\infty = 1$ if $q_\infty \geq 1/2$ and $a^x_\infty = 0$ otherwise. Since $\mathbb{P}[\tilde{a}^x_\infty(k) = \omega] = 1$ for all k large enough, the probability of guessing it correctly under a^x_∞ must also be one:

$$1 = \mathbb{P}[a^x_\infty = \omega] = \mathbb{P}_1[q_\infty \geq 1/2] \cdot \mathbb{P}[\omega = 1] + \mathbb{P}_0[q_\infty < 1/2] \cdot \mathbb{P}[\omega = 0]. \tag{A.3}$$

From (A.3), $\mathbb{P}_0[q_\infty < 1/2] = \mathbb{P}_1[q_\infty \geq 1/2] = 1$. To establish asymptotic learning of informativeness, by definition, we need to show $\mathbb{P}_0[q_\infty = 0] = 1$ and $\mathbb{P}_1[q_\infty = 1] = 1$. To this end, note that $\mathbb{P}_0[q_\infty < 1/2] = 1$ implies that $\mathbb{P}_0[q_\infty \geq 1/2] = 0$. Thus, by Claim 1, it implies that for any $b \in (1/2, 1)$ and all $\tilde{q} \in [1/2, b]$, $\mathbb{P}_1[q_\infty = \tilde{q}] = 0$. Consequently, it follows from $\mathbb{P}_1[q_\infty \geq 1/2] = 1$ that $\mathbb{P}_1[q_\infty = 1] = 1$. The case that $\mathbb{P}_0[q_\infty = 0] = 1$ follows from an identical argument.

(Only-if direction) Suppose by contraposition that $\mathbb{P}_0[S < \infty] > 0$. Again, since $\mathbb{P}_{1,\theta}[a_t \to \theta] = 1$, we have $\mathbb{P}_1[S < \infty] = 1$. Thus, there exists a history of actions at time infinity $\tilde{H}_\infty$ that is possible under both probability measures $\mathbb{P}_0$ and $\mathbb{P}_1$: $\mathbb{P}_0[\tilde{H}_\infty] > 0$ and $\mathbb{P}_1[\tilde{H}_\infty] > 0$. By Bayes' rule, $\mathbb{P}[\omega = 1|\tilde{H}_\infty] < 1$ and $\mathbb{P}[\omega = 0|\tilde{H}_\infty] < 1$. Assume without loss of generality that under $\tilde{H}_\infty$, the corresponding belief $\tilde{q}_\infty \geq 1/2$, and so the observer x would guess $a^x_\infty = 1$. Hence, the probability of x guessing correctly about ω is strictly less than one:

$$\begin{aligned}\mathbb{P}[a^x_\infty = \omega] &= \mathbb{P}[a^x_\infty = \omega, \tilde{H}_\infty] + \mathbb{P}[a^x_\infty = \omega, \tilde{H}^c_\infty] \\ &= \mathbb{P}[\omega = 1|\tilde{H}_\infty] \cdot \mathbb{P}[\tilde{H}_\infty] + \mathbb{P}[a^x_\infty = \omega, \tilde{H}^c_\infty] < \mathbb{P}[\tilde{H}_\infty] + \mathbb{P}[\tilde{H}^c_\infty] = 1.\end{aligned}$$

This contradicts $\mathbb{P}[a^x_\infty = \omega] = 1$, which is required for asymptotic learning of informativeness to hold. □

The following equation will be useful in proving Proposition 4 and later in Proposition 5. Recall that q_∞ and p_∞ are the beliefs of the observer x assigned to the events $\omega = 1$ and $\theta = \mathfrak{g}$ at time infinity, respectively. By the law of total probability,

$$p_\infty = \mathbb{P}_1[\theta = \mathfrak{g}|H_\infty] \cdot q_\infty + \mathbb{P}_0[\theta = \mathfrak{g}|H_\infty] \cdot (1 - q_\infty) \tag{A.4}$$

$$= \pi_\infty \cdot q_\infty + \frac{1}{2} \cdot (1 - q_\infty). \tag{A.5}$$

where $\pi_\infty = \mathbb{P}_1[\theta = \mathfrak{g}|H_\infty]$, and the second equality follows from the fact that, conditioned on $\omega = 0$, no action contains any information about θ and the prior on θ is uniform.

Proof of Proposition 4. We first note that because agents always act as if signals are informative, and signals are unbounded, a standard martingale convergence argument implies that the agent's public belief $\pi_t = \mathbb{P}_1[\theta = \mathfrak{g}|H_t]$ converges to the truth when $\omega = 1$. That is, (i) conditioned on $\omega = 1$ and $\theta = \mathfrak{g}$, $\pi_\infty = 1$ almost surely and (ii) conditioned on $\omega = 1$ and $\theta = \mathfrak{b}$, $\pi_\infty = 0$ almost surely.

(Only-if direction) Suppose asymptotic learning of informativeness holds. Consider the case where $\omega = 1$. By assumption, $q_\infty = 1$ almost surely. Substituting into (A.5), we get $p_\infty = \pi_\infty$. Since π_t converges to the truth, p_t also converges to the truth, i.e., complete learning holds, satisfying condition (i). Now consider the case where $\omega = 0$. By assumption, $q_\infty = 0$ almost surely. Substituting into (A.5), we get $p_\infty = 1/2$. This satisfies condition (ii).

(If direction) Suppose $\omega = 1$ and by condition (i), complete learning holds, i.e., conditioned on $\omega = 1$ and $\theta = \mathfrak{g}$, $p_\infty = 1$ almost surely, and conditioned on $\omega = 1$ and $\theta = \mathfrak{b}$, $p_\infty = 0$ almost surely. Since these also hold for π_∞, substituting into (A.5), we get that conditioned on $\omega = 1$, $1 = q_\infty + \frac{1}{2}(1 - q_\infty)$ if $\theta = \mathfrak{g}$ and $0 = \frac{1}{2}(1 - q_\infty)$ if $\theta = \mathfrak{b}$, which implies that $q_\infty = 1$ almost surely.

Now suppose $\omega = 0$. Assume towards a contradiction that $q_\infty > 0$ with some positive probability. By Proposition 2, this implies that there is a positive probability of agents reaching an action consensus. Let $\tilde{H}_\infty$ be one such consensus history. On this path, the observer belief must be $\tilde{q}_\infty \in (0, 1)$,[15] and the agent's public belief $\tilde{\pi}_\infty$ takes values in $\{0, 1\}$. If $\tilde{\pi}_\infty = 1$, then by (A.5), $\tilde{p}_\infty = \frac{1}{2}(1 + \tilde{q}_\infty) > 1/2$. If $\tilde{\pi}_\infty = 0$, then by (A.5), $\tilde{p}_\infty = \frac{1}{2}(1 - \tilde{q}_\infty) < 1/2$. In either scenario, $\tilde{p}_\infty \neq 1/2$. This means that the event $\{p_\infty \neq 1/2\}$ occurs with positive probability, which contradicts condition (ii). □

Proof of Lemma 4. The proof idea is similar to the proof of Lemma 10 in Arieli, Babichenko, Müller, Pourbabaee, and Tamuz (2023). Recall that we use $\mathbb{P}_{1,\mathfrak{g}}$ to denote the conditional probability distribution given $\omega = 1$ and $\theta = \mathfrak{g}$ and we use $\mathbb{P}_{1,\mathfrak{g},\pi}$ to denote the same conditional probability distribution while emphasizing the prior value.

[15] Note that conditioned on $\omega = 0$, q_∞ has support $\subseteq [0, 1)$ as $\mathbb{P}[\omega = 0|q_\infty = 1] = 0$.

Since conditioned on $\omega = 1$, correct herding holds, and this means that $\mathbb{P}_{1,\mathfrak{g}}[a_t \to \mathfrak{g}] = 1$. As a consequence, part (i) follows directly from the fact that the events $\{\bar{a} = \mathfrak{b}\}$ and $\{a_t \to \mathfrak{g}\}$ are disjoint, and thus $\mathbb{P}_{1,\mathfrak{g}}[\bar{a} = \mathfrak{b}] = 0$.

For part (ii), let $\tau < \infty$ denote the last random time at which the agent chooses the wrong action $\mathfrak{b}$. It is well-defined as correct herding holds. Hence, $1 = \mathbb{P}_{1,\mathfrak{g}}[a_t \to \mathfrak{g}] = \sum_{k=1}^{\infty} \mathbb{P}_{1,\mathfrak{g}}[\tau = k]$. By the overturning principle (Lemma 2), $a_\tau = \mathfrak{b}$ implies that $\pi_{\tau+1} \leq 1/2$. As a consequence,

$$
\begin{aligned}
1 = \sum_{k=1}^{\infty} \mathbb{P}_{1,\mathfrak{g}}[\tau = k] &= \sum_{k=1}^{\infty} \mathbb{P}_{1,\mathfrak{g}}[a_{k+1} = a_{k+2} = \ldots = \mathfrak{g}, \pi_{k+1} \leq 1/2] \\
&= \sum_{k=1}^{\infty} \mathbb{E}_{1,\mathfrak{g}}\left[\mathbb{P}_{1,\mathfrak{g}}[a_{k+1} = a_{k+2} = \ldots = \mathfrak{g}, \pi_{k+1} \leq 1/2 \mid \pi_{k+1}]\right] \\
&= \sum_{k=1}^{\infty} \mathbb{E}_{1,\mathfrak{g}}\left[\mathbb{1}(\pi_{k+1} \leq 1/2) \cdot \mathbb{P}_{1,\mathfrak{g}}[a_{k+1} = a_{k+2} = \ldots = \mathfrak{g} \mid \pi_{k+1}]\right] \\
&= \sum_{k=1}^{\infty} \mathbb{E}_{1,\mathfrak{g}}\left[\mathbb{1}(\pi_{k+1} \leq 1/2) \cdot \mathbb{P}_{1,\mathfrak{g},\pi_{k+1}}[\bar{a} = \mathfrak{g}]\right],
\end{aligned}
$$

where the second equality follows from the law of total expectation, and the last equality follows from the stationarity property (Lemma 3). Suppose that for all prior $\pi \in (0,1)$, $\mathbb{P}_{1,\mathfrak{g},\pi}[\bar{a} = \mathfrak{g}] = 0$. This implies that the above equation equals zero, a contradiction. $\square$

A.1. **Proof of Theorem 1.** In this section, we prove Proposition 3 and Theorem 1. To prove Proposition 3, we will first prove the following proposition (Proposition 6). Together with Proposition 2, they jointly imply Proposition 3. The proof of Theorem 1 is presented at the end of this section. We write $\mathbb{P}_{0,\pi}$ to denote the conditional probability distribution given $\omega = 0$ while highlighting the value of the prior π on θ.[16]

Proposition 6. *The following are equivalent.*

(i) For any action $a \in \{\mathfrak{g}, \mathfrak{b}\}$, $\mathbb{P}_{0,\pi}[\bar{a} = a] = 0$ for all prior $\pi \in (0,1)$.
(ii) $\mathbb{P}_0[S = \infty] = 1$.
(iii) For any action $a \in \{\mathfrak{g}, \mathfrak{b}\}$, $\mathbb{P}_{0,\pi}[\bar{a} = a] = 0$ for some prior $\pi \in (0,1)$.

To prove this proposition, we first establish some preliminary results on the process of the agents' public log-likelihood ratios conditioned on the event of immediate agreement. These results lead to Lemma 7, a crucial part in establishing the equivalence between no immediate agreement and perpetual disagreement conditioned on an uninformative source. We present the proof of Proposition 6 towards the end of this section.

Preliminaries. Recall that conditioned on $\{\bar{a} = \mathfrak{g}\}$, the process of the agent's public log-likelihood ratio r_t evolves deterministically according to (3), which we denote by $r_t^{\mathfrak{g}}$. Let

[16]We continue to omit the prior π when it is uniform.

the corresponding updating function be

$$\phi(x) := x + D_{\mathfrak{g}}(x).$$

That is, $r^{\mathfrak{g}}_{t+1} = \phi(r^{\mathfrak{g}}_t)$ for all $t \geq 1$. Since the entire sequence $(r^{\mathfrak{g}}_t)$ is determined once its initial value $r^{\mathfrak{g}}_1$ is specified, we denote the value of $r^{\mathfrak{g}}_t$ with an initial value $r^{\mathfrak{g}}_1 = r$ by $r^{\mathfrak{g}}_t(r)$. We can thus write $r^{\mathfrak{g}}_t(r) = \phi^{t-1}(r)$ for all $t \geq 1$, where ϕ^t is its t-th composition and $\phi^0(r) = r$.

We remind the readers of two standard properties of the sequence $(r^{\mathfrak{g}}_t)$, as summarized in the following lemma. The first part of this lemma states that $(r^{\mathfrak{g}}_t)$ tends to infinity as t tends to infinity, and the second part shows that it takes only some bounded time for the sequence $(r^{\mathfrak{g}}_t)$ to reach any positive value.

Lemma 5 (The Long-Run and Short-Run Behaviors of $r^{\mathfrak{g}}_t$)**.**
(i) $\lim_{t\to\infty} r^{\mathfrak{g}}_t = \infty$.
(ii) For any $\bar{r} \geq 0$*, there exists* t_0 *such that* $r^{\mathfrak{g}}_{t_0}(r) \geq \bar{r}$ *for all* $r \geq 0$.

Proof. See Lemma 6 and Lemma 12 in Rosenberg and Vieille (2019). □

Note that although the sequence $(r^{\mathfrak{g}}_t)$ eventually approaches infinity, it may not do so monotonically without additional assumptions on the distributions $F_{\mathfrak{g}}$ and $F_{\mathfrak{b}}$.[17] The next lemma shows that, under some mild technical assumptions on the left tail of $F_{\mathfrak{b}}$, the function $\phi(x)$ eventually increases monotonically.

Lemma 6 (Eventual Monotonicity)**.** *Suppose that* $F_{\mathfrak{b}}$ *has a differentiable left tail and its probability density function* $f_{\mathfrak{b}}$ *satisfies the condition that, for all* x *large enough,* $f_{\mathfrak{b}}(-x) < 1$*. Then,* $\phi(x) := x + D_{\mathfrak{g}}(x)$ *increases monotonically for all* x *large enough.*

Proof. By assumption, we can find a constant $\rho < 1$ such that for all x large enough, $f_{\mathfrak{b}}(-x) \leq \rho$. By definition, $D_g(x) = \log \frac{1-F_{\mathfrak{g}}(-x)}{1-F_{\mathfrak{b}}(-x)}$. Taking the derivative of $D_{\mathfrak{g}}$,

$$D'_{\mathfrak{g}}(x) = \frac{f_{\mathfrak{g}}(-x)}{1 - F_{\mathfrak{g}}(-x)} - \frac{f_{\mathfrak{b}}(-x)}{1 - F_{\mathfrak{b}}(-x)}.$$

Observe that the log-likelihood ratio of the agent's private log-likelihood ratio ℓ_t is the log-likelihood ratio itself (see, e.g., Chamley (2004)):

$$\log \frac{dF_{\mathfrak{g}}}{dF_{\mathfrak{b}}}(x) = x.$$

It follows that

$$-D'_{\mathfrak{g}}(x) = f_{\mathfrak{b}}(-x)\Big(\frac{1}{1 - F_{\mathfrak{b}}(-x)} - \frac{e^{-x}}{1 - F_{\mathfrak{g}}(-x)}\Big) \leq \frac{f_{\mathfrak{b}}(-x)}{1 - F_{\mathfrak{b}}(-x)}.$$

[17] In the case of binary states and actions, Herrera and Hörner (2012) show that the property of increasing hazard ratio is equivalent to the monotonicity of this updating function. See Smith, Sørensen, and Tian (2021) for a general treatment.

Fix some $\varepsilon > 0$ small enough so that $(1+\varepsilon)\rho \le 1$. It follows from the above inequality that there exists some x large enough such that $-D'_{\mathfrak{g}}(x) \le (1+\varepsilon) f_{\mathfrak{b}}(-x)$. Furthermore, for all $x' \ge x$,

$$\begin{aligned} D_{\mathfrak{g}}(x) &= D_{\mathfrak{g}}(x') - \int_x^{x'} D'_{\mathfrak{g}}(y)dy \\ &\le D_{\mathfrak{g}}(x') + (1+\varepsilon)\int_x^{x'} f_{\mathfrak{b}}(-x)dx \\ &= D_{\mathfrak{g}}(x') - (1+\varepsilon)(F_{\mathfrak{b}}(-x') - F_{\mathfrak{b}}(-x)). \end{aligned}$$

Rearranging the above equation,

$$\begin{aligned} D_{\mathfrak{g}}(x) - D_{\mathfrak{g}}(x') &\le (1+\varepsilon)(F_{\mathfrak{b}}(-x) - F_{\mathfrak{b}}(-x')) \\ &\le (1+\varepsilon)\rho(x'-x), \end{aligned}$$

where the second last inequality follows from the fact that $f_{\mathfrak{b}}(-x) \le \rho < 1$. Since $(1+\varepsilon)\rho \le 1$, the above inequality implies that there exists some x large enough such that $D_{\mathfrak{g}}(x) + x \le D_{\mathfrak{g}}(x') + x'$ for all $x' \ge x$. That is, $\phi(x)$ eventually increases monotonically.

□

Given these lemmas, we are ready to prove the following result. It shows that conditioned on an uninformative source, the possibility of immediate agreement is independent of the prior belief.

Lemma 7. *For any action $a \in \{\mathfrak{g}, \mathfrak{b}\}$, the following statements are equivalent:*
(i) $\mathbb{P}_{0,\pi}[\bar{a} = a] > 0$, for some prior $\pi \in (0,1)$;
(ii) $\mathbb{P}_{0,\pi}[\bar{a} = a] > 0$, for all prior $\pi \in (0,1)$.

Proof. The second implication, namely, $(ii) \Rightarrow (i)$ is immediate. We will show the first implication, $(i) \Rightarrow (ii)$. Fix some prior $\tilde{\pi} \in (0,1)$ such that $\mathbb{P}_{0,\tilde{\pi}}[\bar{a} = \mathfrak{g}] > 0$ and let $\tilde{r} = \log \frac{\tilde{\pi}}{1-\tilde{\pi}}$. Since $r_t^{\mathfrak{g}}(\tilde{r})$ is a deterministic process, the event $\{\bar{a} = \mathfrak{g}\}$ initiated at the prior $\pi_1 = \tilde{\pi}$ is equivalent to the event $\{\ell_t \ge -r_t^{\mathfrak{g}}(\tilde{r}), \forall t \ge 1\}$. Conditioned on $\omega = 0$, since signals are i.i.d., so are the agents' private log-likelihood ratios. Thus, we have

$$\mathbb{P}_{0,\tilde{\pi}}[\bar{a} = \mathfrak{g}] = \prod_{t=1}^{\infty}(1 - F_0(-r_t^{\mathfrak{g}}(\tilde{r})). \tag{A.6}$$

As a consequence, $\mathbb{P}_{0,\tilde{\pi}}[\bar{a} = \mathfrak{g}] > 0$ if and only if there exists $M < \infty$ such that

$$-\sum_{t=1}^{\infty} \log\left(1 - F_0(-r_t^{\mathfrak{g}}(\tilde{r}))\right) < M.$$

For two sequences (a_t) and (b_t), we write $a_t \approx b_t$ if $\lim_{t\to\infty}(a_t/b_t) = 1$. Since $r_t^{\mathfrak{g}}(\tilde{r}) \to \infty$ (this follows from part (i) of Lemma 5), $\log\left(1 - F_0(-r_t^{\mathfrak{g}}(\tilde{r}))\right) \approx -F_0(-r_t^{\mathfrak{g}}(\tilde{r}))$. Thus, the above sum is finite if and only if

$$\sum_{t=1}^{\infty} F_0(-r_t^{\mathfrak{g}}(\tilde{r})) < M. \tag{A.7}$$

By the overturning principle (Lemma 2), it suffices to show that (A.7) implies that

$$\sum_{t=1}^{\infty} F_0(-r_t^{\mathfrak{g}}(r)) < M, \quad \text{for any } r \geq 0.$$

By the eventual monotonicity of ϕ (Lemma 6) and the fact that $r_t^{\mathfrak{g}}(\tilde{r}) \to \infty$, we can find a large enough $\bar{t}$ such that $r_{\bar{t}}^{\mathfrak{g}}(\tilde{r}) := \bar{r} \geq 0$ and $\phi(r) \geq \phi(\bar{r})$ for all $r \geq \bar{r}$. By part (ii) of Lemma 5, there exists $t_0 \in \mathbb{N}$ such that $r_{t_0}^{\mathfrak{g}}(r) \geq \bar{r}$ for all $r \geq 0$. Since above $\bar{r}$, ϕ is monotonically increasing, one has $\phi(r_{t_0}^{\mathfrak{g}}(r)) \geq \phi(\bar{r})$ for any $r \geq 0$. Consequently, for all $\tau \geq 1$, $r_{\tau+t_0}^{\mathfrak{g}}(r) = \phi^{\tau}(r_{t_0}^{\mathfrak{g}}(r)) \geq \phi^{\tau}(\bar{r}) = r_{\tau+1}^{\mathfrak{g}}(\bar{r})$. Since $r_{\tau+1}^{\mathfrak{g}}(\bar{r}) = r_{\tau+1}^{\mathfrak{g}}(r_{\bar{t}}^{\mathfrak{g}}(\tilde{r})) = r_{\tau+\bar{t}}^{\mathfrak{g}}(\tilde{r})$, it follows that

$$F_0(-r_{\tau+t_0}^{\mathfrak{g}}(r)) \leq F_0(-r_{\tau+\bar{t}}^{\mathfrak{g}}(\tilde{r})).$$

Thus, it follows from (A.7) that for any $r \geq 0$, $\sum_{t=1}^{\infty} F_0(-r_t^{\mathfrak{g}}(r)) < \infty$, as required. The case for action $\mathfrak{b}$ follows from a symmetric argument. □

Now, we are ready to prove Proposition 6.

Proof of Proposition 6. We show that *(i)* ⇒ *(ii)*, *(ii)* ⇒ *(iii)*, and *(iii)* ⇒ *(i)*. To show the first implication, we prove the contrapositive statement. Suppose that $\mathbb{P}_0[S < \infty] > 0$. This implies that there exists a sequence of action realizations $(b_1, b_2, \ldots, b_{k-1}, b_k = \ldots = a)$ for some action $a \in \{\mathfrak{b}, \mathfrak{g}\}$ such that

$$\mathbb{P}_0[a_t = b_t, \forall t \geq 1] > 0.$$

By stationarity, there exists some $\pi' \in (0, 1)$ such that

$$\mathbb{P}_{0,\pi'}[\bar{a} = a] > 0,$$

which contradicts (i).

To show the second implication, suppose towards a contradiction that there exists some action $a \in \{\mathfrak{g}, \mathfrak{b}\}$ such that $\mathbb{P}_{0,\pi}[\bar{a} = a] > 0$ for all prior $\pi \in (0, 1)$. In particular, it holds for the uniform prior. Since the event $\{\bar{a} = a\}$ is contained in the event $\{S < \infty\}$,

$$0 < \mathbb{P}_0[\bar{a} = a] \leq \mathbb{P}_0[S < \infty].$$

This implies that $\mathbb{P}_0[S = \infty] < 1$, a contradiction to (ii).

Finally, we show the last implication by contraposition. Suppose that there exists some $a \in \{\mathfrak{g}, \mathfrak{b}\}$ such that $\mathbb{P}_{0,\pi}[\bar{a} = a] > 0$ for some prior $\pi \in (0, 1)$. By Lemma 7, it also holds for all prior $\pi \in (0, 1)$, which is a contradiction to (iii). This concludes the proof of Proposition 6. □

Proof of Proposition 3. By the equivalence between (ii) and (iii) in Proposition 6 and Proposition 2, we have shown Proposition 3. □

Given Proposition 3, we are now ready to prove our main result.

Proof of Theorem 1. By part (i) of Lemma 4, $\mathbb{P}_{1,\mathfrak{b}}[\bar{a} = \mathfrak{g}] = 0$ and $\mathbb{P}_{1,\mathfrak{g}}[\bar{a} = \mathfrak{b}] = 0$. Following a similar argument that led to (A.6), one has

$$0 = \mathbb{P}_{1,\mathfrak{b}}[\bar{a} = \mathfrak{g}] = \prod_{t=1}^{\infty}(1 - F_{\mathfrak{b}}(-r_t^{\mathfrak{g}})).$$

Taking the logarithm on both sides, the above equation is equivalent to $-\sum_{t=1}^{\infty}\log(1 - F_{\mathfrak{b}}(-r_t^{\mathfrak{g}})) = \infty$. Since $r_t^{\mathfrak{g}} \to \infty$, $\log(1 - F_{\mathfrak{b}}(-r_t^{\mathfrak{g}})) \approx -F_{\mathfrak{b}}(-r_t^{\mathfrak{g}})$ and the previous sum is infinite if and only if

$$\sum_{t=1}^{\infty} F_{\mathfrak{b}}(-r_t^{\mathfrak{g}}) = \infty. \tag{A.8}$$

Similarly, we have that $\mathbb{P}_{1,\mathfrak{g}}[\bar{a} = \mathfrak{b}] = 0$ if and only if $\sum_{t=1}^{\infty}(1 - F_{\mathfrak{g}}(-r_t^{\mathfrak{b}})) = \infty$, where $r_t^{\mathfrak{b}}$ denotes the deterministic process of r_t conditioned on the event $\{\bar{a} = \mathfrak{b}\}$. By symmetry, $r_t^{\mathfrak{b}} = -r_t^{\mathfrak{g}}$ for all $t \geq 1$. Hence, $\mathbb{P}_{1,\mathfrak{g}}[\bar{a} = \mathfrak{b}] = 0$ if and only if

$$\sum_{t=1}^{\infty}(1 - F_{\mathfrak{g}}(r_t^{\mathfrak{g}})) = \infty, \tag{A.9}$$

Suppose that the uninformative signals have fatter tails than the informative signals. By definition, there exists $\varepsilon > 0$ such that for all x large enough, $F_0(-x) \geq \varepsilon \cdot F_{\mathfrak{b}}(-x)$ and $1 - F_0(x) \geq \varepsilon \cdot (1 - F_{\mathfrak{g}}(x))$. It then follows from (A.8) and (A.9) that

$$\sum_{t=1}^{\infty} F_0(-r_t^{\mathfrak{g}}) = \infty \quad \text{and} \quad \sum_{t=1}^{\infty}(1 - F_0(r_t^{\mathfrak{g}})) = \infty.$$

Using the same logic we used to deduce (A.8) and (A.9), having these two divergent sums is equivalent to $\mathbb{P}_0[\bar{a} = \mathfrak{g}] = 0$ and $\mathbb{P}_0[\bar{a} = \mathfrak{b}] = 0$. Thus, by Proposition 3, asymptotic learning of informativeness holds.

By part (ii) of Lemma 4, there exist $\pi, \pi' \in (0,1)$ such that $\mathbb{P}_{1,\mathfrak{g},\pi}[\bar{a} = \mathfrak{g}] > 0$ and $\mathbb{P}_{1,\mathfrak{b},\pi'}[\bar{a} = \mathfrak{b}] > 0$. Let $r = \log\frac{\pi}{1-\pi}$ and $r' = \log\frac{\pi'}{1-\pi'}$. Following a similar argument that led to (A.7), these are equivalent to

$$\sum_{t=1}^{\infty} F_{\mathfrak{g}}(-r_t^{\mathfrak{g}}(r)) < \infty \quad \text{and} \quad \sum_{t=1}^{\infty}(1 - F_{\mathfrak{b}}(r_t^{\mathfrak{g}}(r'))) < \infty.$$

Now, suppose that the uninformative signals have thinner tails than the informative signals. By definition, there exists $\varepsilon > 0$ such that either (i) $F_0(-x) \leq (1/\varepsilon) \cdot F_{\mathfrak{g}}(-x)$ for all x large enough, or (ii) $1 - F_0(x) \leq (1/\varepsilon) \cdot (1 - F_{\mathfrak{b}}(x))$ for all x large enough. It then follows from the above inequalities that either (i) $\sum_{t=1}^{\infty} F_0(-r_t^{\mathfrak{g}}(r)) < \infty$ or (ii) $\sum_{t=1}^{\infty}(1 - F_0(r_t^{\mathfrak{g}}(r')))$. Equivalently, we have either (i) $\mathbb{P}_{0,\pi}[\bar{a} = \mathfrak{g}] > 0$ for some $\pi \in (0,1)$ or (ii) $\mathbb{P}_{0,\pi'}[\bar{a} = \mathfrak{b}] > 0$ for some $\pi' \in (0,1)$. By Lemma 7, these also hold for all prior

$\pi, \pi' \in (0,1)$, including the uniform prior. So we have either (i) $\mathbb{P}_0[\bar{a} = \mathfrak{g}] > 0$ or (ii) $\mathbb{P}_0[\bar{a} = \mathfrak{b}] > 0$. Thus, by Proposition 3, asymptotic learning of informativeness fails. $\square$

A.2. **Gaussian Private Signals.** In this section, we consider the canonical environment in which all signal distributions $\mu_\mathfrak{g}$, $\mu_\mathfrak{b}$, and μ_0 are normal. Specifically, $\mu_\mathfrak{g}$ and $\mu_\mathfrak{b}$ share the same variance σ^2 and have means $+1$ and -1, respectively, while μ_0 has mean $m_0 \in (-1,1)$ and variance τ^2. The agent's private log-likelihood ratio induced by a signal s is therefore

$$\ell_t = \log \frac{d\mu_\mathfrak{g}}{d\mu_\mathfrak{b}}(s) = \frac{2}{\sigma^2} s. \tag{A.10}$$

Since ℓ is proportional to s, the distributions F_θ and F_0 are also normal, with a variance of $4/\sigma^2$ and $4\tau^2/\sigma^4$, respectively.

We first consider the case where $\tau \neq \sigma$. Here, the relative tail thickness between the uninformative and informative signals is determined solely by their variances, as shown by the following lemma.

Lemma 8. *Suppose F and G are two Gaussian cumulative distribution functions with means μ_1 and μ_2 and variances σ_1^2 and σ_2^2, respectively. If $\sigma_1 > \sigma_2$, then F has fatter tails than G. Meanwhile, G has thinner tails than F.*

Proof. Let f and g denote the probability density functions of F and G. Their ratio evaluated at $x \in \mathbb{R}$ is

$$\frac{f(x)}{g(x)} = \frac{\sigma_2}{\sigma_1} \exp\Big((\frac{1}{\sigma_2^2} - \frac{1}{\sigma_1^2}) \frac{x^2}{2} + (\frac{\mu_1}{\sigma_1^2} - \frac{\mu_2}{\sigma_2^2}) x + \frac{1}{2}(\frac{\mu_2^2}{\sigma_2^2} - \frac{\mu_1^2}{\sigma_1^2}) \Big).$$

Suppose $\sigma_1 > \sigma_2$. It follows from the above equation that $\lim_{x\to\infty} \frac{f(-x)}{g(-x)} = \infty$ and $\lim_{x\to\infty} \frac{f(x)}{g(x)} = \infty$. This clearly implies that $F(-x) \geq G(-x)$ and $1 - F(x) \geq 1 - G(x)$ for all x large enough. Thus, by definition, F has fatter tails than G, and conversely, G has thinner tails than F. $\square$

We henceforth focus on the case where $\tau = \sigma$. As discussed in the main text, this is a special case where the uninformative signals have neither fatter nor thinner tails. When all private signals are Gaussian, Hann-Caruthers, Martynov, and Tamuz (2018) show that the sequence $r_t^\mathfrak{g}$ can be approximated by $(2\sqrt{2}/\sigma) \cdot \sqrt{\log t}$ for all t large enough (see their Theorem 4):

$$\lim_{t\to\infty} \frac{r_t^\mathfrak{g}}{(2\sqrt{2}/\sigma) \cdot \sqrt{\log t}} = 1. \tag{A.11}$$

Given this approximation and Proposition 3, we are ready to prove Proposition 1.

Proof of Proposition 1. In this proof we use the Landau notation, so that $O(g(t))$ stands for some function $f : \mathbb{N} \to \mathbb{R}$ such that there exists a positive $M \in \mathbb{R}$ and $t_0 \in \mathbb{N}$ such that $|f(t)| \leq M \cdot g(t)$ for all $t \geq t_0$.

Note that by (A.10), we can write

$$F_0(-r_t^{\mathfrak{g}}) = \mathbb{P}_0[\ell_t \le -r_t^{\mathfrak{g}}] = \mathbb{P}_0[s_t \le -(\sigma^2/2)\cdot r_t^{\mathfrak{g}}].$$

By (A.11), we have that for all t large enough,

$$F_0(-r_t^{\mathfrak{g}}) = \mathbb{P}_0[s_t \le -\sigma\sqrt{2\log t}] =: \mu_0(-\sigma\sqrt{2\log t}),$$

where μ_0 is the CDF of s_t conditioned on $\omega = 0$. Since μ_0 is the normal distribution with mean $m_0 \in (-1, 1)$ and variance σ^2, observe that $\mu_0(x) = \frac{1}{2}\operatorname{erfc}(-\frac{x-m_0}{\sigma\sqrt{2}})$, where $\operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}}\int_x^\infty \mathrm{e}^{-t^2}dt$ is the complementary error function.

Applying a standard asymptotic expansion of the complementary error function, i.e., $\operatorname{erfc}(x) = \frac{e^{-x^2}}{x\sqrt{\pi}} + O(\mathrm{e}^{-x^2}/x^3)$, we obtain that for all t large enough,

$$\mu_0(-\sigma\sqrt{2\log t}) = \frac{\mathrm{e}^{-\left(\frac{m_0}{\sigma}\sqrt{2\log t}+\frac{m_0^2}{2\sigma^2}\right)}}{t(\sqrt{\pi}\log t + \delta\cdot m_0)} + O\Big(\frac{\mathrm{e}^{-m_0\sqrt{2\log t}}}{t(\sigma\sqrt{2\log t}+m_0)^3}\Big), \tag{A.12}$$

where $\delta > 0$ is a constant.

Case (i): suppose $m_0 = 0$. Then (A.12) becomes $\frac{1}{t\cdot\sqrt{\pi\log t}} + O(\frac{1}{t\cdot(\log t)^{3/2}})$. Since the series $\frac{1}{t\log t}$ is divergent and $\frac{1}{t\log t} \le \frac{1}{t\sqrt{\log t}}$ for all $t \ge 2$, the sum of the first term also diverges. Hence, the sum of (A.12) diverges, which implies $\sum_{t=1}^\infty F_0(-r_t^{\mathfrak{g}}) = \infty$. By the same approximation argument used in the proof of Theorem 1, this is equivalent to $\mathbb{P}_0[\bar{a} = \mathfrak{g}] = 0$. Using an analogous argument, we obtain $\mathbb{P}_0[\bar{a} = \mathfrak{b}] = 0$. Together, by Proposition 3, we conclude that asymptotic learning of informativeness holds.

Case (ii): suppose $m_0 \ne 0$ and $m_0 \in (-1, 1)$. Let $c = \frac{m_0\sqrt{2}}{\sigma}$. By the change of variable $x = \sqrt{\log t}$,

$$\int_2^\infty \frac{\mathrm{e}^{-c\sqrt{\log t}}}{t\sqrt{\log t}}dt = 2\int_{\sqrt{\log 2}}^\infty \mathrm{e}^{-cx}dx.$$

If $m_0 > 0$, then $c = \frac{m_0\sqrt{2}}{\sigma} > 0$. By the integral test, the sum in (A.12) converges, and thus $\sum_{t=1}^\infty F_0(-r_t^{\mathfrak{g}}) < \infty$. Again, by the same logic that we use to deduce (A.7), this is equivalent to $\mathbb{P}_0[\bar{a} = \mathfrak{g}] > 0$. If $m_0 < 0$, it follows from a symmetric argument that $\sum_{t=1}^\infty (1 - F_0(r_t^{\mathfrak{g}})) < \infty$, which is equivalent to $\mathbb{P}_0[\bar{a} = \mathfrak{b}] > 0$. In either case, it follows from Proposition 3 that asymptotic learning of informativeness fails. □

A.3. **Non-Uniform Prior.** In this section, we prove Proposition 5. We relax the uniform prior assumption on θ while maintaining the independence assumption between θ and ω.

Recall that $p_t = \mathbb{P}[\theta = \mathfrak{g}|H_t]$ and $q_t = \mathbb{P}[\omega = 1|H_t]$ is the outside observer's belief that the source is informative at time t. Given that the prior on θ is not uniform, each agent t chooses an action according to (1). It immediately follows that p_t also satisfies the overturning principle as in Lemma 2. The following lemma is useful in proving Proposition 5. Let a^d denote the default action that is favored by the non-uniform prior.

Lemma 9. *Given a non-uniform prior on θ, the following statements hold:*

(i) Conditioned on $\lim_{t\to\infty} q_t = 0$ almost surely, actions almost surely converge to a^d.

(ii) Conditioned on $\lim_{t\to\infty} q_t = 1$ almost surely, actions almost surely converge to θ.

Proof. Fix a prior $\pi = \mathbb{P}[\theta = \mathfrak{g}] > 1/2$, and so $a^d = \mathfrak{g}$. For part (i), suppose towards a contradiction that actions do not converge to $\mathfrak{g}$ almost surely. There are two possible cases: either $\lim_{t\to\infty} a_t = \mathfrak{b}$, or the actions never converge, i.e., $S = \infty$. By the overturning principle, these scenarios imply that either $\lim_{t\to\infty} p_t = p_\infty \leq 1/2$, or $p_\infty = 1/2$. Together they imply that the event $\{p_\infty \leq 1/2\}$ occurs with positive probability. Since $\lim_{t\to\infty} q_t = q_\infty = 0$ almost surely, it follows from (A.4) that

$$p_\infty = \mathbb{P}_0[\theta = \mathfrak{g} | H_\infty] = \pi > 1/2 \quad \text{almost surely},$$

a contradiction to the event $\{p_\infty \leq 1/2\}$ occurring with positive probability. Hence, we conclude that $\lim_{t\to\infty} a_t = \mathfrak{g}$ almost surely.

For part (ii), consider the case where $\theta = \mathfrak{g}$. Assume by contradiction, that either $\lim_{t\to\infty} a_t = \mathfrak{b}$ or $S = \infty$ occurs with positive probability. Again, by the overturning principle, these imply that the event $\{p_\infty \leq 1/2\}$ occurs with positive probability. Since $\lim_{t\to\infty} q_t = 1$ almost surely, it follows from (A.4) that $\pi_\infty = p_\infty \leq 1/2$ with positive probability. By a standard martingale convergence argument with unbounded signals (Smith and Sørensen, 2000), we know that conditioned on $\theta = \mathfrak{g}$, $\pi_\infty = 1$ almost surely, which contradicts to the event $\{\pi_\infty \leq 1/2\}$ occurring with positive probability. An analogous argument applies to the case where $\theta = \mathfrak{b}$. □

Proof of Proposition 5. Suppose by contradiction that asymptotic learning of informativeness holds. By definition and Lemma 9, we have that conditioned on $\omega = 0$, $\lim_t a_t = a^d$ almost surely; similarly, conditioned on $\omega = 1$, $\lim_t a_t = \theta$ almost surely. However, since $\mathbb{P}[\theta = a^d | \omega = 1] > 0$, there exists a history of actions at time infinity H^d_∞ in which $\lim_t a_t = a^d$ and such history is possible under both probability measures $\mathbb{P}_0$ and $\mathbb{P}_1$. Following a similar argument to that used in the proof of Proposition 2 (for the only-if direction), this implies that asymptotic learning of informativeness fails, a contradiction. □