

The Hazards and Benefits of Condescension in Social Learning

Itai Arieli* Yakov Babichenko† Stephan Müller‡
Farzad Pourbabaee§ Omer Tamuz¶

February 17, 2024

Abstract

In a misspecified social learning setting, agents are condescending if they perceive their peers as having private information that is of lower quality than it is in reality. Applying this to a standard sequential model, we show that outcomes improve when agents are mildly condescending. In contrast, too much condescension leads to worse outcomes, as does anti-condescension.

*Technion, E-mail: iarieli@technion.ac.il

†Technion, E-mail: yakovbab@technion.ac.il

‡University of Göttingen, E-mail: stephan.mueller@wiwi.uni-goettingen.de

§Caltech, E-mail: far@caltech.edu

¶Caltech, E-mail: omertamuz@gmail.com

1 Introduction

Most human decisions are made under uncertainty and in a social context. Understanding how economic agents use their private and social information to form beliefs is a prerequisite for the understanding of important phenomena such as the diffusion of ideas, the adoption of technologies, or the formation of political opinions. In particular, agents’ beliefs about their peers’ information is an important factor that can play a decisive role in the social outcome.

We study the effect of condescension on social learning outcomes: What happens when agents, through misspecification, underestimate the quality of the information that their peers have? Our main result is that condescension can lead to improved social outcomes, as long as it is mild. In contrast, anti-condescension, in which agents overestimate their peers’ quality of information, leads to bad outcomes, as does too much condescension.

We study a misspecified version of the classical sequential social learning model of [Bikhchandani et al. \(1992\)](#) and [Banerjee \(1992\)](#), with unbounded signals, as introduced by [Smith and Sørensen \(2000\)](#). Our notion of a good social learning outcome is that of efficient learning ([Rosenberg and Vieille, 2019](#)), which is said to occur when the number of agents who choose the incorrect action has finite expectation. In the well-specified setting, the number of agents who choose the incorrect action is always finite ([Smith and Sørensen, 2000](#)), but its expectation can be finite or infinite ([Rosenberg and Vieille, 2019](#)).

In our misspecified setting agents perfectly understand and interpret their own signal but misperceive the quality of their predecessors’ signals. When agents are mildly condescending, efficient learning occurs. Because agents underestimate the quality of others’ signals, they put too little weight on their predecessors’ actions. In consequence, their actions are suboptimal, but reveal more of their own private information. When this is done in moderation more is gained than lost, and in the long run, the result is quick convergence to the correct action. This occurs even with signal distributions that would have induced inefficient learning for well-specified agents. Of course, since agents are misspecified, each agent attains lower expected utility than they would if they were not, *ceteris paribus*. Nevertheless, their behavior has positive externalities on later agents, with improved asymptotic outcomes.

When agents are too condescending they put so little weight on their predecessors’ actions that no herd forms and both actions are taken infinitely often, *i.e.*, asymptotic learning is not obtained. When agents are anti-condescending they put too much weight on their predecessors’ actions. In consequence, wrong herds form with positive probability, and again asymptotic learning is not obtained. Interestingly, it follows that asymptotic learning is equivalent to efficient learning across all misspecified regimes.

Our proof techniques follow those introduced by [Hann-Caruthers et al. \(2018\)](#) and [Rosenberg and Vieille \(2019\)](#) who approximate the discrete time dynamics of the public belief using a continuous time differential equation. Due to the misspecified nature of our model, our analysis deviates from theirs in a number of places. For example, we need to circumvent the fact that the misspecified belief is not a martingale. In their model, asymptotic learning is guaranteed by this martingale property ([Smith and Sørensen, 2000](#)), whereas in our model we need to prove it by other means.

Related Literature. A closely related paper is [Bernardo and Welch \(2001\)](#). They study a cascade setting with binary signals, but where some fraction of the agents are *overconfident*: They do not put enough weight (in Bayesian terms) on the public information. Through mostly numerical analysis, the authors reach a conclusion that is similar to ours: Moderate overconfidence is beneficial for society.

The empirical literature on social learning supports the idea of overweighting the private information relative to the public information. For example, [Weizsäcker \(2010\)](#) finds in a meta-analysis of 13 social learning experiments that subjects underweight their social observations relative to the payoff-maximizing strategy. [Duffy et al. \(2021\)](#) find sizable proportions of both behavioral types, i.e. relative over- and underweighting of the private information. Condescension provides one possible mechanism which leads to individually suboptimal low weights of public information. Other mechanisms that may cause a distortion in the optimal weighting of private and public information are, for example, cursedness ([Eyster and Rabin, 2005](#)) and naïveté ([Eyster and Rabin, 2010](#)).

Our paper is related to the literature on social learning with misspecification. Most of this literature documents a detrimental effect of misspecification on asymptotic outcomes: For example, the gambler’s fallacy leads to incorrect learning almost surely ([He, 2022](#)), and misinterpreting peers’ preferences can lead to incorrect ([Frick, Iijima, and Ishii, 2020](#)) and cyclical ([Gagnon-Bartsch, 2016](#)) learning, or entrenched disagreement ([Bohren and Hauser, 2019](#)). [Frick et al. \(2023\)](#) analyze belief convergence in a general setting. They demonstrate that in the sequential social learning environment arbitrarily small amounts of misspecification can lead to extreme failures of learning. [Bohren \(2016\)](#) studies agents with misspecification regarding the correlation between others’ actions, and shows that various undesirable outcomes are possible, depending on the degree and direction of misspecification; see also [Bohren and Hauser \(2021\)](#) for a more general setting that subsumes a number of previous ones. Our results highlight potential positive welfare effects, i.e., misspecifications may increase the efficiency of learning. Similarly to some of these papers, our results show

that learning outcomes depend on the behavior of the signal distributions near the fixed points of the learning dynamics, which in our case are the extremal beliefs.

The literature has identified other channels—not requiring misspecification—by which agents may put more weight on their own signals, as compared to the signals of others. For example, idiosyncratic taste shocks (see, e.g., [Goeree et al., 2006](#); [Lobel and Sadler, 2016](#)) imply that an agent’s own signal carries more information about their own payoff relevant state than do the signals of the others. As in our model, this causes agents to reveal more of their private information through their actions, which in turn can improve information aggregation in the long term.

A paper that is slightly further away in its goal—but closer in techniques and some of the results—is [Chen \(2022\)](#), who studies ambiguity in sequential social learning. In his model, agents have ambiguity about the distributions of the other agents’ signals. The main conclusion is that information cascades are a robust outcome that occurs whenever there is sufficient ambiguity. Technically, similar observations to our Propositions [1](#) and [2](#) appear in [Chen \(2022\)](#), but our paper is focused on the speed of learning (i.e., efficient learning), which is not studied in that paper.

Our work also complements a burgeoning literature which analyzes the rationale for the persistence of misspecifications (e.g., [He and Libgober, 2020](#); [Ba, 2021](#); [Fudenberg and Lanzani, 2022](#)). That is, in a sequential learning environment misspecified agents might have an evolutionary advantage over correctly specified agents by learning the true state of the world faster. Consequently, misspecifications caused by intermediate levels of condescension might persist in the long run.

2 Model

2.1 Social Learning with Misspecification

There is a binary state of nature $\theta \in \Theta = \{\ell, h\}$, chosen at time zero, and equal to h with probability $\pi \in (0, 1)$. A countably infinite set of agents $= \{1, 2, \dots\}$ arrive sequentially. Each agent n , in turn, takes an action $a_n \in \{\ell, h\}$, with utility 1 if $a = \theta$ and 0 otherwise. Before choosing her action, agent n observes her predecessors’ actions $I_n = (a_1, \dots, a_{n-1})$.

Each agent also observes a private signal $s_n \in S$. Here S is some measurable set of possible signal realizations. Signals are independent and identically distributed conditioned on the state. We denote probabilities by \mathbb{P} , and explicitly write \mathbb{P}_π when we want to highlight varying values of the prior π . We further use the notation \mathbb{P}_h to refer to $\mathbb{P}(\cdot \mid \theta = h)$, the

probability measure \mathbb{P} conditional on the realized state being h . We define $\mathbb{P}_{\pi,h}$ analogously.

Let $q_n = \mathbb{P}(\theta = h|s_n)$ be the (random) *private posterior*: The belief induced by observing the private signal of agent n . By a standard direct revelation argument, we can assume that $s_n = q_n$, since q_n is a sufficient statistic for θ given s_n . Denote by F_ℓ and F_h the cumulative distribution functions of q_n , conditioned on $\theta = \ell$ and $\theta = h$, respectively. We define $F = \frac{1}{2}(F_\ell + F_h)$. This is the cumulative distribution function of q_n , for prior $\pi = 1/2$.

So far, this model matches the standard herding model (Bikhchandani et al., 1992; Banerjee, 1992; Smith and Sørensen, 2000). We deviate from these models by introducing a misspecification regarding others' private signals: Agents correctly observe their own type q_n , but have a (common) misspecified prior about the distribution of types. Namely, each agent believes that all the others' private posteriors have conditional distributions \tilde{F}_ℓ and \tilde{F}_h . Furthermore, it is common knowledge that these are the agents' beliefs. Note that agents still interpret their own private signals correctly, with agent n calculating q_n from s_n according to $q_n = \mathbb{P}(\theta = h|s_n)$. We denote by $\tilde{\mathbb{P}}$ the posterior probabilities calculated according to the agents' misspecified beliefs.

In equilibrium, agents choose actions a_n to maximize their subjective expected utilities:

$$a_n = \arg \max_{a \in \{\ell, h\}} \tilde{\mathbb{P}}(\theta = a | I_n, q_n).$$

We will below restrict ourselves to q_n with non-atomic distributions, i.e., we assume that F_ℓ and F_h are continuous. This will ensure that agents are never indifferent and the maximum above is unique. We will likewise assume that \tilde{F}_ℓ and \tilde{F}_h are continuous.

A pair of conditional CDFs (F_ℓ, F_h) is *symmetric* (around $q = 1/2$) if $F_\ell(q) + F_h(1-q) = 1$. This in turn implies $F(q) + F(1-q) = 1$. To simplify our exposition we will make the following assumption.

Assumption 1 (Symmetry). We assume throughout that (F_ℓ, F_h) and $(\tilde{F}_\ell, \tilde{F}_h)$ are symmetric.

When the prior is $\pi = 1/2$, this is equivalent to requiring that the model is invariant with respect to renaming the states.

2.2 Efficiency

To study efficiency in this setting, we follow Rosenberg and Vieille (2019) and introduce some additional notation. Let $W := \#\{n : a_n \neq \theta\}$ be the (random) number of agents who take the incorrect action.

The next definition includes two notions of efficiency of social learning.

Definition 1. 1. *Asymptotic learning* holds if all agents, except finitely many, choose the correct action. That is, if W is finite \mathbb{P} -almost surely.

2. *Efficient learning* holds if $\mathbb{E}[W] < \infty$.

Note that asymptotic learning is equivalent to the sequence of actions a_n converging to θ , which is again equivalent to $a_n = \theta$ for all n large enough. Note also that efficient learning implies asymptotic learning.

2.3 The Well-Specified Case

Without misspecification, the classical herding result of [Bikhchandani et al. \(1992\)](#) is that asymptotic learning does not hold for any finitely supported private signal distribution in which no signal is revealing. This is an outcome that displays extreme inefficiency: With positive probability, all but finitely many agents choose incorrectly, and in particular there is no asymptotic or efficient learning. [Smith and Sørensen \(2000\)](#) show that asymptotic learning holds if and only if signals are *unbounded*: That is, if the support of the distribution of the private posteriors q_n includes 0 and 1. Thus, when signals are sufficiently informative, the extreme inefficiency of the wrong herds of [Bikhchandani et al. \(1992\)](#) is overturned.

Nevertheless, this result left open the possibility that many agents choose incorrectly before the correct herd arrives. To quantify this intuition, [Sørensen \(1996\)](#) gave an example in which learning is not efficient: $\mathbb{E}[W]$, the expected number of agents who choose incorrectly, is infinite. He also conjectured that this is the case for every signal distribution. This conjecture was shown to be false by [Hann-Caruthers et al. \(2018\)](#) and [Rosenberg and Vieille \(2019\)](#). In particular, [Rosenberg and Vieille \(2019\)](#) give an elegant necessary and sufficient condition for efficient learning, showing that efficient learning holds if and only if $\int_0^1 \frac{1}{F(x)} dx < \infty$.

2.4 Condensation

We use our misspecified social learning framework to study how outcomes change when agents are *condensing*, or think that others' signals are less informative than they really are. To formalize and quantify this notion, we restrict ourselves to signals that are *tail-regular*: A pair of symmetric conditional CDFs (F_ℓ, F_h) is tail-regular if there exists $\alpha > 0$ such that $F(q) = (F_\ell(q) + F_h(q))/2$ behaves like q^α near $q = 0$. Formally, if

$$0 < \liminf_{q \rightarrow 0} \frac{F(q)}{q^\alpha} \leq \limsup_{q \rightarrow 0} \frac{F(q)}{q^\alpha} < \infty.$$

We use Landau notation and write

$$F(q) = \Theta(q^\alpha)$$

as a shorthand for the expression above.¹

Assumption 2 (Tail-Regularity). We assume throughout that (F_ℓ, F_h) and $(\tilde{F}_\ell, \tilde{F}_h)$ are tail-regular.

The exponent α associated with a symmetric, tail-regular signal is unique, and given by

$$\alpha = \lim_{q \rightarrow 0} \frac{\log F(q)}{\log q}.$$

It captures a notion of the thinness of the tail of the signals: For high α there is a small chance of very informative signals, as compared to low α . Thus, in an asymptotic sense, signals are less informative for higher α . Note that by a standard argument (Lemma 4) if $F(q) = \Theta(q^\alpha)$ then $F_\ell(q) = \Theta(q^\alpha)$ and $F_h(q) = \Theta(q^{\alpha+1})$.

As an example of tail-regular signals, consider the family of beta distributions, which are commonly used in the applied literature to model the distribution of posterior beliefs (see, e.g., McKelvey and Palfrey, 1992; Nyarko et al., 2006; Bosch-Domènech et al., 2010; Çelen et al., 2020). This is a family of probability distributions on the interval $[0, 1]$ parametrized by $\alpha, \beta > 0$, with probability density function given by $g_{\alpha, \beta}(q) = Cq^{\alpha-1}(1-q)^{\beta-1}$, where C is a normalization constant. The parameters β and α describe the thickness of the distribution around 0 and 1, respectively. Suppose that private signals have conditional densities $f_\ell = g_{\alpha, \alpha+1}$ and $f_h = g_{\alpha+1, \alpha}$. Then, the unconditional density is again a beta distribution with probability density function $f = g_{\alpha, \alpha}$. This is easily seen to be symmetric and tail-regular, with exponent α .

An important example of signals that are not tail-regular is that of conditionally Gaussian signals, with distributions $\mathcal{N}(-m, 1)$ and $\mathcal{N}(+m, 1)$, depending on the state. In this case $\lim_{q \rightarrow 0} \frac{\log F(q)}{\log q} = \infty$, and the tail-regularity condition is violated, as $F(q)$ decays faster than any polynomial.

In a misspecified model we denote by α and $\tilde{\alpha}$ the exponents associated with (F_ℓ, F_h) and $(\tilde{F}_\ell, \tilde{F}_h)$, respectively. When $\tilde{\alpha} > \alpha$ we say that agents are *condescending*: They believe that

¹More generally in Landau notation, given two functions $f(x)$ and $g(x)$, one writes $f(x) = \Theta(g(x))$ if

$$0 < \liminf_{x \rightarrow 0} \frac{f(x)}{g(x)} \leq \limsup_{x \rightarrow 0} \frac{f(x)}{g(x)} < \infty.$$

others' signals are less informative than they really are. Conversely, when $\tilde{\alpha} < \alpha$ agents are *anti-condescending*. Thus, $\tilde{\alpha} - \alpha$ is a measure of how condescending the agents are. Note that our definition of condescension is relatively mild, in the sense that it only depends on the tail properties and not on the bulk of the distribution.

3 Results

Our first result characterizes the efficiency of learning outcomes for condescending and anti-condescending agents. The formal proof appears in the appendix. The intuition and dynamics behind this result are presented in detail in Section 4.

Theorem 1. *Suppose $\tilde{\alpha} \neq \alpha$. Then the following are equivalent: (i) asymptotic learning; (ii) efficient learning; (iii) $\tilde{\alpha} - \alpha \in (0, 1)$.*

The regime $\tilde{\alpha} - \alpha \in (0, 1)$ describes agents who are condescending ($\tilde{\alpha} - \alpha > 0$) but not overly condescending ($\tilde{\alpha} - \alpha < 1$). In this regime the agents' mild condescension causes them to slightly discount the actions of their predecessors, resulting in an increased chance that wrong herds are overturned and a correct herd starts. Since there are infinitely many agents there will be infinitely many fresh chances to do this, and eventually one will succeed, even if each has a very low probability. Moreover, the success probabilities are bounded from below, rendering the expected number of mistakes finite.

In a model without misspecification, the results of [Rosenberg and Vieille \(2019\)](#) imply that efficient learning occurs if and only if $\alpha < 1$. In contrast, Theorem 1 shows that efficiency can be regained under potentially small misspecification, for any α , as long as agents are condescending, but not too condescending. Put differently, even for large values of α , i.e., when highly informative signals are extremely rare, learning can still be efficient under mild condescension. The technical reason for why a difference of exactly 1 between $\tilde{\alpha}$ and α is the boundary between mild condescension and over-condescension is related to the summability of certain sequences that determine whether or not a herd can start immediately with positive probability. We explain this in detail in Section 4.

Note that the exponent α captures a tail property of the private signals rather than a parameter that determines the entire distribution. Hence two distributions can be very different even if their exponents are very close or the same. When $\tilde{\alpha} \neq \alpha$, the finer properties of the distributions do not play a role, and efficient learning is solely determined by the tail exponents. When $\tilde{\alpha} = \alpha$, we conjecture that it is possible for learning to either be efficient or inefficient, depending on finer properties of the distributions; see footnote 5 for the technical

details. We leave it for future work to identify these finer properties that determine efficient learning in this regime.

The next two propositions shed light into why learning fails when agents are either anti-condescending or overly condensing. These results also appear in the supplementary material to [Chen \(2022\)](#). We provide proofs for these claims in Appendix E for completeness.

The first proposition concerns the anti-condescension regime, in which $\tilde{\alpha} < \alpha$.

Proposition 1. *Suppose that $\tilde{\alpha} < \alpha$. Then, with \mathbb{P} -positive probability, a wrong herd forms, i.e., from some point on, all agents take the wrong action.*

In the case $\tilde{\alpha} < \alpha$ agents are anti-condescending: They believe that others have signals that are more informative than they really are. In consequence, they are more easily swayed by other’s actions, and tend to more often ignore their private signals.² Thus, wrong herds can form. This is despite the fact that signals are unbounded, which, without misspecification, would rule out wrong herds.

Our next result tackles the question of why learning fails when agents are overly condensing, i.e., when $\tilde{\alpha} - \alpha \geq 1$.

Proposition 2. *Suppose that $\tilde{\alpha} \geq \alpha + 1$. Then, \mathbb{P} -almost surely, both actions are taken by infinitely many agents.*

In the case $\tilde{\alpha} \geq \alpha + 1$, agents are very condensing: They think that others have very uninformative signals. In consequence, they follow their own signals too much, and herds—wrong or right—do not form: Given enough time, an agent will come along who will overturn her predecessor’s action.

To sum, Proposition 1 shows that if $\tilde{\alpha} < \alpha$, then when public belief assigns a low probability to the realized state, incorrect cascades remain stable with positive probability. Proposition 2 shows that if $\tilde{\alpha} \geq \alpha + 1$, then correct cascades are unstable, and thus even if the public belief assigns high probability to the realized state, herds on the correct action almost certainly break down. As we shall show, these conditions are also determinants of efficient learning.

²The effect of anti-condescension is similar to that of naïveté in [Eyster and Rabin \(2010\)](#). The mechanism is, however, very different. While naïve agents fail to realize that previous movers’ also infer from still earlier actions, anti-condescending agents are fully aware of this and take it into account but believe others to have better information than they actually do.

4 Dynamics

In this section, we study how agents update their beliefs and choose their actions under misspecification. We define the public belief and derive its equations of motion. We show that two properties of the well-specified model—stationarity and the overturning principle—still hold in our misspecified environment.

4.1 Belief Updating

An important tool in social learning is the public belief (or social belief) at time n :

$$\pi_n = \mathbf{P}(\theta = h | a_1, \dots, a_{n-1}).$$

In our case, however, it is also important to consider the misspecified public belief, which is given by

$$\tilde{\pi}_n = \tilde{\mathbf{P}}(\theta = h | a_1, \dots, a_{n-1}).$$

The public belief π_n is the belief held by a well-specified observer who sees the agents' actions but not their signals. In contrast, $\tilde{\pi}_n$ is the belief held by an observer who holds the same misspecified beliefs as the agents, and again sees only actions.

Let $p_n = \mathbf{P}(\theta = h | I_n, q_n)$ be the posterior belief held by a well-specified agent who observes all the information available to agent n . The actual, misspecified, posterior of agent n is denoted $\tilde{p}_n = \tilde{\mathbf{P}}(\theta = h | I_n, q_n)$. Then by Bayes' Law

$$\begin{aligned} \frac{p_n}{1-p_n} &= \frac{\pi_n}{1-\pi_n} \times \frac{q_n}{1-q_n}, \\ \frac{\tilde{p}_n}{1-\tilde{p}_n} &= \frac{\tilde{\pi}_n}{1-\tilde{\pi}_n} \times \frac{q_n}{1-q_n}. \end{aligned}$$

It follows that the action a_n chosen by agent n is equal to h if $\tilde{\pi}_n + q_n \geq 1$, and to ℓ otherwise.³ Thus, conditioned on θ , the probability that agent n chooses the low action is $F_\theta(1 - \tilde{\pi}_n)$.

This implies that when agent n chooses the low action, the public beliefs $\{\pi_n\}$ and $\{\tilde{\pi}_n\}$ evolve as follows:

$$\frac{\pi_{n+1}}{1-\pi_{n+1}} = \frac{\pi_n}{1-\pi_n} \times \frac{F_h(1-\tilde{\pi}_n)}{F_\ell(1-\tilde{\pi}_n)}, \quad (4.1a)$$

$$\frac{\tilde{\pi}_{n+1}}{1-\tilde{\pi}_{n+1}} = \frac{\tilde{\pi}_n}{1-\tilde{\pi}_n} \times \frac{\tilde{F}_h(1-\tilde{\pi}_n)}{\tilde{F}_\ell(1-\tilde{\pi}_n)}. \quad (4.1b)$$

³As we note above, indifference occurs with probability zero because we assume that the distribution of q_n is non-atomic.

When agent n chooses the high action,

$$\frac{\pi_{n+1}}{1-\pi_{n+1}} = \frac{\pi_n}{1-\pi_n} \times \frac{1-F_h(1-\tilde{\pi}_n)}{1-F_\ell(1-\tilde{\pi}_n)}, \quad (4.2a)$$

$$\frac{\tilde{\pi}_{n+1}}{1-\tilde{\pi}_{n+1}} = \frac{\tilde{\pi}_n}{1-\tilde{\pi}_n} \times \frac{1-\tilde{F}_h(1-\tilde{\pi}_n)}{1-\tilde{F}_\ell(1-\tilde{\pi}_n)}. \quad (4.2b)$$

4.2 Stationarity and the Overturning Principle

The above equations of motion imply that as in the well-specified case, our model is stationary, with $\tilde{\pi}_n$ capturing all the relevant information about the past.

Lemma 1 (Stationarity). *For any fixed sequence b_1, \dots, b_k of actions in $\{\ell, h\}$ and any $\tilde{\pi} \in (0, 1)$,*

$$\mathbf{P}(a_{n+1} = b_1, \dots, a_{n+k} = b_k \mid \tilde{\pi}_{n+1} = \tilde{\pi}) = \mathbf{P}_{\tilde{\pi}}(a_1 = b_1, \dots, a_k = b_k).$$

That is, suppose that at time n the misspecified public belief $\tilde{\pi}_n$ was equal to some $\tilde{\pi}$. Then the probability that the subsequent actions are b_1, \dots, b_k is the same as the probability of observing this sequence of actions at time 1, when the prior is $\tilde{\pi}$.

Another important observation that generalizes to the misspecified setting is Sørensen's overturning principle.

Lemma 2 (Overturning principle). *The misspecified public belief $\tilde{\pi}_{n+1}$ in period $n+1$ is greater than or equal to $1/2$ if and only if $a_n = h$.*

Proof. Observe that by the law of total expectation

$$\begin{aligned} \tilde{\pi}_{n+1} &= \tilde{\mathbf{E}} \left[\mathbb{1}_{\{\theta=h\}} \mid a_1, \dots, a_n \right] \\ &= \tilde{\mathbf{E}} \left[\tilde{\mathbf{E}} \left[\mathbb{1}_{\{\theta=h\}} \mid a_1, \dots, a_n, q_n \right] \mid a_1, \dots, a_n \right] \\ &= \tilde{\mathbf{E}} \left[\tilde{p}_n \mid a_1, \dots, a_n \right]. \end{aligned}$$

Therefore, $a_n = h$ is equivalent to $\tilde{p}_n \geq 1/2$, and hence equivalent to $\tilde{\pi}_{n+1} \geq 1/2$. □

4.3 Asymptotic Learning and Immediate Herding

In the misspecified setting the public belief $\tilde{\pi}_n$ is not a martingale under the correct measure \mathbf{P} . This martingale property is an important tool in the proof of asymptotic learning for unbounded signals in the well-specified case ([Smith and Sørensen, 2000](#)). In our case,

asymptotic learning indeed does not always hold, and in particular, we need different tools to analyze it.

We denote by $a_n \rightarrow h$ the asymptotic event that the sequence of actions converges to h . Namely, that $a_n = h$ for all n large enough, or that a high action herd forms eventually. We denote by $\bar{a} = h$ the event $\{a_1 = h, a_2 = h, \dots\}$ that all agents took the high action; this is the event that a high action herd formed immediately.

Asymptotic learning occurs when $\mathbb{P}_\ell(a_n \rightarrow \ell) = 1$ and $\mathbb{P}_h(a_n \rightarrow h) = 1$. To study the asymptotic events $a_n \rightarrow \ell$ and $a_n \rightarrow h$, we study the immediate herding events $\bar{a} = h$ and $\bar{a} = \ell$. These are easier to analyze because conditioned on $\bar{a} = \ell$ or on $\bar{a} = h$, the sequence of misspecified public beliefs $\{\tilde{\pi}_n\}$ is deterministic and given recursively by (4.1b) or (4.2b), respectively.

To see the connection between asymptotic learning and immediate herding, condition on $\theta = h$ and consider the event $a_n \rightarrow h$ of a good herd forming eventually. In our setting, we show that this event has probability 1 if and only if two conditions are met:⁴

- (i) The event $\bar{a} = \ell$ of an immediate bad herd has probability 0.
- (ii) The event $\bar{a} = h$ of an immediate good herd has positive probability.

The first condition is clearly necessary for asymptotic learning: If bad herd can form then the probability of a good herd is less than 1. The reason that the second condition is necessary is related to the stationarity of the process; for a good herd to form eventually, it must have a positive probability to be formed at any point in time, and hence also in the beginning.

To see that these conditions are sufficient for asymptotic learning, note that again applying stationarity, the first condition implies that it is impossible for a bad herd to start at any point in time. This implies that the high action will be taken infinitely often. Hence, there will be infinitely many chances for a good herd to start, and thus, by the second condition (and again stationarity) a good herd will form eventually.

To apply stationarity, we need these two conditions to hold for any prior, and moreover uniformly so. This is done formally in Appendix C.

Having reduced the problem of asymptotic learning to that of immediate herding, we turn to calculating the probability of the events $\bar{a} = \ell$ and $\bar{a} = h$. Condition on $\bar{a} = h$. Then the public belief $\tilde{\pi}_n$ evolves deterministically according to (4.2b). It will be useful to consider the

⁴We omit some technical details in the statements of these two conditions. A complete formal treatment is presented in Appendix C, Lemmas 10 and 11.

misspecified public log-likelihood ratio $\tilde{r}_n := \log \frac{\tilde{\pi}_n}{1-\tilde{\pi}_n}$. In terms of \tilde{r}_n , the equation of motion (4.2b) becomes

$$\tilde{r}_{n+1} = \tilde{r}_n + \log \frac{1 - \tilde{F}_h\left(\frac{1}{1+e^{\tilde{r}_n}}\right)}{1 - \tilde{F}_\ell\left(\frac{1}{1+e^{\tilde{r}_n}}\right)}.$$

It starts at the initial level $\tilde{r}_1 = \log \frac{\pi}{1-\pi}$. When $(\tilde{F}_\ell, \tilde{F}_h)$ is tail-regular with exponent $\tilde{\alpha}$, we can for small q approximate $\tilde{F}_\ell(q)$ with $q^{\tilde{\alpha}}$ and $\tilde{F}_h(q)$ with $q^{\tilde{\alpha}+1}$ (neglecting constants). And since \tilde{r}_n tends to infinity when $\bar{a} = h$, this equation of motion is well approximated by

$$\tilde{r}_{n+1} \approx \tilde{r}_n + e^{-\tilde{\alpha}\tilde{r}_n}. \quad (4.3)$$

Intuitively, after each observed high action the misspecified public log-likelihood increases by an amount $e^{-\tilde{\alpha}\tilde{r}_n}$ that becomes smaller as \tilde{r}_n increases. More importantly, $e^{-\tilde{\alpha}\tilde{r}_n}$ is also smaller when $\tilde{\alpha}$ is higher, i.e., when the signals are less informative: After many high actions, agents are less surprised to see another high action when signals are less likely to be very informative.

The asymptotic behavior of this discrete time equation can in turn be approximated by the differential equation $\frac{d\tilde{r}(t)}{dt} = e^{-\tilde{\alpha}\tilde{r}(t)}$ whose solution is $\tilde{r}(t) = \tilde{\alpha}^{-1} \log(1 + \tilde{\alpha}t)$; this is shown formally in Appendix B, Lemmas 8 and 9. Thus, conditioned on the event $\bar{a} = h$, the misspecified public log-likelihood \tilde{r}_n takes the sequence of deterministic values \tilde{r}_n^h , which we can approximate by $\tilde{r}_n^h \approx \tilde{\alpha}^{-1} \log(1 + \tilde{\alpha}n)$. Transforming this back to public beliefs, we get

$$\tilde{\pi}_n^h \approx 1 - n^{-1/\tilde{\alpha}}. \quad (4.4)$$

Thus, the sequence of misspecified public beliefs converges to 1, and it does so more slowly for higher $\tilde{\alpha}$, i.e., for less informative signals.

We remind the reader that $a_n = h$ if and only if $q_n \geq 1 - \tilde{\pi}_n$: The agent takes the high action if her private posterior $q_n = \mathbb{P}(\theta = h | s_n)$ exceeds $1 - \tilde{\pi}_n$. Hence the event $\bar{a} = h$ is the event that $q_n \geq 1 - \tilde{\pi}_n^h$ for all n . Conditioned on $\theta = h$ the (actual, not misspecified) probability of this event is

$$1 - F_h(1 - \tilde{\pi}_n^h) \approx 1 - F_h(n^{-1/\tilde{\alpha}}) \approx 1 - n^{-\frac{\alpha+1}{\tilde{\alpha}}},$$

where the first approximation uses (4.4) and the second uses $F_h(q) = \Theta(q^{\alpha+1})$.

Since the random variables q_n are independent conditioned on the state, we get that the probability of $\bar{a} = h$ is

$$\mathbb{P}_h(\bar{a} = h) = \prod_{n=1}^{\infty} (1 - F_h(1 - \tilde{\pi}_n^h)) \approx \prod_{n=1}^{\infty} \left(1 - n^{-\frac{\alpha+1}{\tilde{\alpha}}}\right).$$

Crucially, we are only interested in whether this probability is positive or zero. As we show formally in Lemma 12 in Appendix D, the approximations we perform are good enough, in the sense that the first product vanishes if and only if the second one does. Thus, by an elementary argument we get that $P_h(\bar{a} = h) > 0$ if and only if $\tilde{\alpha} - \alpha < 1$.

This argument shows that immediate good herds can form if and only if $\tilde{\alpha} - \alpha < 1$, i.e., agents are not overly condescending. A similar line of reasoning shows that $P_h(\bar{a} = \ell) = 0$ if and only if $\tilde{\alpha} - \alpha \geq 0$, i.e., immediate bad herds are excluded when agents are condescending.

The assumption of tail regularity is used in the approximation $\tilde{r}_{n+1} \approx \tilde{r}_n + e^{-\tilde{\alpha}\tilde{r}_n}$ of the evolution of the public log-likelihood ratio, made in (4.3). Tail regularity is also crucial for showing that the differential equation $\frac{d\tilde{r}(t)}{dt} = e^{-\tilde{\alpha}\tilde{r}(t)}$ is a good approximation of this discrete time dynamics; see the proofs of Lemmas 8 and 9 in Appendix B.

4.4 Efficient Learning

In the previous section we explained why asymptotic learning holds only in the regime $\tilde{\alpha} - \alpha \in [0, 1)$. This immediately implies that outside this range there is also no efficient learning. In this section we explain why efficient learning does hold when $\tilde{\alpha} - \alpha \in (0, 1)$.

Suppose $\tilde{\alpha} - \alpha \in (0, 1)$. As asymptotic learning holds, we know that the agents will take the high action from some point on. Until then, there will be *runs* of wrong actions, or sequences of consecutive agents who make the wrong choice. These will be separated by runs of agents who make the correct choice.

The argument for efficient learning includes two parts. First, we show that the expected number of bad runs is finite. Second, we show that the expected length of each bad run is finite. Moreover, the expected length of a bad run is uniformly bounded, regardless of the history that came before that run. It follows that the total number of agents W who take the wrong action has a finite expectation.

The reason that the number of bad runs has finite expectation is that regardless of the history, there is a uniform lower bound δ on the probability that a good herd continues forever. This implies that the distribution of the number of bad runs is stochastically dominated by a geometric distribution, which has a finite expectation. This holds whenever $\tilde{\alpha} - \alpha < 1$, i.e., whenever agents are not overly condescending. The argument is similar to the one from the previous section, which showed that in this range the probability of $\bar{a} = h$ is positive in the high state.

To show that the expected length of each bad run is finite, we again follow the line of argument from the previous section showing that $\bar{a} = \ell$ has zero probability in the high state.

This holds whenever $\tilde{\alpha} > \alpha$, i.e., when agents are condescending. Moreover, we show that the expected length of a bad run is uniformly bounded, regardless of the history that came before it started. This is a consequence of the fact that the public belief at the onset of a run cannot be arbitrarily high or low, but is bounded away from 0 and 1. This is a consequence of tail regularity (see Proposition 7).

We note that this last step is obtained in the well-specified setting of [Rosenberg and Vieille \(2019\)](#) by appealing to the overturning principle and the fact that $\{\frac{1-\pi_n}{\pi_n}\}$ is a martingale under the correct conditional measure (that is, P_h in the high state). In our misspecified case, the overturning principle still holds, but the public likelihood $\{\frac{1-\tilde{\pi}_n}{\tilde{\pi}_n}\}$ is not a martingale under the correct measure. Thus, we have to apply a mechanical method that appeals to tail-regularity.

4.5 Expected Time of the First Correct Action

[Rosenberg and Vieille \(2019\)](#) consider another notion of the efficiency of learning, which is briefly discussed in this section. Let τ be the first time that the correct action is taken:

$$\tau = \min\{n : a_n = \theta\}.$$

This is a random time that takes values in $\mathbb{N} \cup \{\infty\}$.

[Rosenberg and Vieille \(2019\)](#) show that in the well-specified setting, the finiteness of the expectation of τ coincides with efficient learning, or the finiteness of the expectation of W .

In our model, when agents are condescending, i.e., when $\tilde{\alpha} > \alpha$, the expectation of τ is finite (see Proposition 7). This holds even when agents are over-condescending (i.e., $\tilde{\alpha} \geq \alpha + 1$), and efficient learning does not hold. In the latter regime, there is no learning because the agents' condescension causes them to put too much weight on their own signals, resulting in both actions being taken infinitely often, and also in small expected τ . When agents are anti-condescending there is a positive probability of $\tau = \infty$ (see Proposition 5), and in particular, τ has an infinite expectation.

5 Conclusion

In this paper, we study social learning with condescending agents who underestimate the quality of their peers' information. We show that mild condescension can have positive externalities that result in efficient learning. In particular, there are private signal distributions

for which learning is not efficient in the well-specified case, but is efficient with even very small levels of condescension.

We make several simplifying assumptions for expositional purposes. For example, relaxing symmetry (Assumption 1) yields the same type of results, but where the exponent α needs to be defined for each of the two states (corresponding to the left and right tails), and outcomes can be different in each state. We believe that our results also hold for private belief distributions that are not continuous, but currently, our proof techniques only apply in the continuous case.

A more substantial assumption is that all agents have the same misspecified beliefs about others. We see this as the smallest possible deviation from the well-specified case, involving misspecification only about the distribution of agents' types, and nothing else. In particular, because all agents have the same prior, higher order beliefs are trivial, which makes the model tractable. A natural avenue for future work is to relax this assumption. Indeed, higher order beliefs play an interesting and important role in the misspecified social learning literature (Bohren, 2016; Bohren and Hauser, 2021).

Our analysis of social welfare is restricted to the question of whether the expected number of incorrect actions is finite or not. A more nuanced question is to study how this expectation changes as the actual and perceived distributions of private signals vary. In particular, for private signal distributions where this expectation is finite in the well-specified case, it is interesting to understand how misspecification alters this expectation; this is possible even when $\tilde{\alpha} = \alpha$. It is furthermore natural to consider a discounted sum of the number of incorrect actions. These are interesting questions that currently seem to be beyond what is technically tractable.

References

- Cuimin Ba (2021). “Robust Model Misspecification and Paradigm Shifts,” *arXiv preprint arXiv:2106.12727*.
- Abhijit Banerjee (1992). “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107(3): 797–817.
- Antonio E. Bernardo and Ivo Welch (2001). “On the Evolution of Overconfidence and Entrepreneurs,” *Journal of Economics & Management Strategy*, 10(3): 301–330.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch (1992). “A Theory of Fads, Fashion,

- Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100(5): 992–1026.
- J. Aislinn Bohren (2016). “Informational Herding with Model Misspecification,” *Journal of Economic Theory*, 163: 222–247.
- J. Aislinn Bohren and Daniel N. Hauser (2019). “Misinterpreting Social Outcomes and Information Campaigns,” *Working Paper*.
- (2021). “Learning with Heterogeneous Misspecified Models: Characterization and Robustness,” *Econometrica*, 89(6): 3025–3077.
- Antoni Bosch-Domènech, José G. Montalvo, Rosemarie Nagel, and Albert Satorra (2010). “A Finite Mixture Analysis of Beauty-Contest Data Using Generalized Beta Distributions,” *Experimental Economics*, 13: 461–475.
- Boğaçhan Çelen, Sen Geng, and Huihui Li (2020). “Belief Error and Non-Bayesian Social Learning: Experimental Evidence,” *arXiv preprint arXiv:2011.09640*.
- Jaden Yang Chen (2022). “Sequential Learning under Informational Ambiguity,” *Available at SSRN 3480231*.
- John Duffy, Ed Hopkins, and Tatiana Kornienko (2021). “Lone Wolf or Herd Animal? Information Choice and Learning from Others,” *European Economic Review*, 134: 103690.
- Erik Eyster and Matthew Rabin (2005). “Cursed equilibrium,” *Econometrica*, 73(5): 1623–1672.
- (2010). “Naive herding in rich-information settings,” *American Economic Journal: Microeconomics*, 2(4): 221–243.
- Mira Frick, Ryota Iijima, and Yuhta Ishii (2020). “Misinterpreting Others and the Fragility of Social Learning,” *Econometrica*, 88(6): 2281–2328.
- (2023). “Belief Convergence under Misspecified Learning: A Martingale Approach,” *The Review of Economic Studies*, 90(2): 781–814.
- Drew Fudenberg and Giacomo Lanzani (2022). “Which Misperceptions Persist?” *Theoretical Economics*.

- Tristan Gagnon-Bartsch (2016). “Taste Projection in Models of Social Learning,” *Working Paper*.
- Jacob K Goeree, Thomas R Palfrey, and Brian W Rogers (2006). “Social Learning with Private and Common Values,” *Economic Theory*, 28: 245–264.
- Wade Hann-Caruthers, Vadim V. Martynov, and Omer Tamuz (2018). “The Speed of Sequential Asymptotic Learning,” *Journal of Economic Theory*, 173: 383–409.
- Kevin He (2022). “Mislearning from Censored Data: The Gambler’s Fallacy and Other Correlational Mistakes in Optimal-Stopping Problems,” *Theoretical Economics*, 17(3): 1269–1312.
- Kevin He and Jonathan Libgober (2020). “Evolutionarily Stable (Mis)specifications: Theory and Applications,” *arXiv:2012.15007*.
- Ilan Lobel and Evan Sadler (2016). “Preferences, Homophily, and Social Learning,” *Operations Research*, 64(3): 564–584.
- Richard D. McKelvey and Thomas R. Palfrey (1992). “An Experimental Study of the Centipede Game,” *Econometrica*: 803–836.
- Yaw Nyarko, Andrew Schotter, and Barry Sopher (2006). “On the Informational Content of Advice: A Theoretical and Experimental Study,” *Economic Theory*, 29(2): 433–452.
- Dinah Rosenberg and Nicolas Vieille (2019). “On the Efficiency of Social Learning,” *Econometrica*, 87(6): 2141–2168.
- Lones Smith and Peter Sørensen (2000). “Pathological Outcomes of Observational Learning,” *Econometrica*, 68(2): 371–398.
- Peter Norman Sørensen (1996). “Rational Social Learning,” Ph.D. thesis, Massachusetts Institute of Technology.
- Georg Weizsäcker (2010). “Do We Follow Others When We Should? A Simple Test of Rational Expectations,” *American Economic Review*, 100(5): 2340–60.

A Preliminaries

The following lemma is a standard result, with proofs given, for example, in Appendix A of [Hann-Caruthers et al. \(2018\)](#) or [Rosenberg and Vieille \(2019\)](#).

Lemma 3. *Let G_ℓ and G_h be two cumulative distribution functions on $[0, 1]$, with the Radon-Nikodym derivative dG_h/dG_ℓ satisfying the iterated likelihood principle $\frac{dG_h}{dG_\ell}(q) = q/(1-q)$. Then it holds that:*

$$\begin{aligned} G_h(q) &= 2 \left(qG(q) - \int_0^q G(x) dx \right), \\ G_\ell(q) &= 2 \left((1-q)G(q) + \int_0^q G(x) dx \right). \end{aligned}$$

where $G = \frac{1}{2}(G_\ell + G_h)$. These in turn imply that $G_h(q) \leq 2qG(q)$ and $|G_\ell(q) - 2G(q)| \leq 3qG(q)$. Therefore, $\lim_{q \rightarrow 0} G_h(q)/G(q) = 0$ and $\lim_{q \rightarrow 0} G_\ell(q)/G(q) = 2$.

We use this lemma to prove the following additional lemma which relates the exponent of G to the exponents of G_ℓ and G_h .

Lemma 4. *Suppose $G(q) = \Theta(q^\alpha)$. Then $G_\ell(q) = \Theta(q^\alpha)$ and $G_h(q) = \Theta(q^{\alpha+1})$.*

Proof. Lemma 3 immediately implies that $G_\ell(q) = \Theta(q^\alpha)$ whenever $G(q) = \Theta(q^\alpha)$. To see that $G_h(q) = \Theta(q^{\alpha+1})$, note that $G(q) = \Theta(q^\alpha)$ implies there are constants $C \geq c > 0$ such that for all $q \in [0, 1]$, one has $cq^\alpha \leq G(q) \leq Cq^\alpha$. The previous lemma thus implies that $G_h(q) \leq 2Cq^{\alpha+1}$. Next, let us define $m := (c/2C)^{1/\alpha}$, and observe that

$$G(mq) \leq C(mq)^\alpha = \frac{c}{2} q^\alpha \leq \frac{1}{2} G(q).$$

Since G is increasing, then $G(x) \leq G(q)/2$ for all $x \leq mq$, and therefore, $\int_0^{mq} G(x) dx \leq mqG(q)/2$. In addition, G being increasing implies that $\int_{mq}^q G(x) dx \leq (1-m)qG(q)$. Therefore, one obtains the following upper bound for the integral:

$$\begin{aligned} \int_0^q G(x) dx &= \int_0^{mq} G(x) dx + \int_{mq}^q G(x) dx \\ &\leq \frac{m}{2} qG(q) + (1-m)qG(q) = \left(1 - \frac{m}{2}\right) qG(q). \end{aligned}$$

It follows from the expression for G_h in the previous lemma that $G_h(q) \geq mqG(q)$, and hence $G_h(q) \geq mcq^{\alpha+1}$. Therefore, we have shown that $G_h(q) = \Theta(q^{\alpha+1})$. \square

B The Evolution of the Public Log-Likelihood

Define the misspecified public log-likelihood ratio by

$$\tilde{r}_n = \log \frac{\tilde{\pi}_n}{1 - \tilde{\pi}_n},$$

and the well-specified public log-likelihood ratio by

$$r_n = \log \frac{\pi_n}{1 - \pi_n}.$$

At $n = 1$, it holds that $r_1 = \tilde{r}_1 = \log \frac{\pi}{1 - \pi}$. Conditioned on the event $\bar{a} = h$, $\tilde{\pi}_n$ satisfies the recursive equation (4.2b), and thus $\tilde{\pi}_n$ is deterministic and equals to some $\tilde{\pi}_n^h$. We accordingly denote $\tilde{r}_n^h = \log \frac{\tilde{\pi}_n^h}{1 - \tilde{\pi}_n^h}$.

Lemma 5. $\lim_{n \rightarrow \infty} \tilde{\pi}_n^h = 1$.

Proof. The perceived distributions \tilde{F}_h and \tilde{F}_ℓ satisfy the iterated likelihood principle, that is,

$$\frac{d\tilde{F}_h}{d\tilde{F}_\ell}(q) = \frac{q}{1 - q}.$$

This relation implies that $\tilde{F}_h - \tilde{F}_\ell$ is strictly decreasing on $[0, 1/2]$ and strictly increasing on $[1/2, 1]$. Therefore, for every $\pi \in (0, 1/2]$ it must be that

$$\tilde{F}_h(\pi) - \tilde{F}_\ell(\pi) < \tilde{F}_h(0) - \tilde{F}_\ell(0) = 0 \Rightarrow \tilde{F}_h(\pi) < \tilde{F}_\ell(\pi),$$

and for every $\pi \in [1/2, 1)$ one has

$$\tilde{F}_h(\pi) - \tilde{F}_\ell(\pi) < \tilde{F}_h(1) - \tilde{F}_\ell(1) = 0 \Rightarrow \tilde{F}_h(\pi) > \tilde{F}_\ell(\pi).$$

Observe that due to equation (4.2b) the sequence $\tilde{\pi}_n$ is strictly increasing. Now, assume by contradiction that $\tilde{\pi}_n \rightarrow \hat{\pi} \in (0, 1)$, then it must be that

$$\frac{1 - \tilde{F}_h(1 - \hat{\pi})}{1 - \tilde{F}_\ell(1 - \hat{\pi})} = 1,$$

which is in contrast with the previous two implications about $\{\tilde{F}_\ell, \tilde{F}_h\}$. □

In this section we provide asymptotic results for the evolution of r_n and \tilde{r}_n on the high action path. As discussed above, on this path these random variables are deterministic, and equal to some constants r_n^h and \tilde{r}_n^h , respectively. These constants satisfy the following reformulation of expressions in (4.2a) and (4.2b):

$$r_{n+1}^h = r_n^h + U(\tilde{r}_n^h), \tag{B.1}$$

$$\tilde{r}_{n+1}^h = \tilde{r}_n^h + \tilde{U}(\tilde{r}_n^h), \tag{B.2}$$

where

$$U(r) := \log \frac{1 - F_h\left(\frac{1}{1+e^r}\right)}{1 - F_\ell\left(\frac{1}{1+e^r}\right)},$$

$$\tilde{U}(r) := \log \frac{1 - \tilde{F}_h\left(\frac{1}{1+e^r}\right)}{1 - \tilde{F}_\ell\left(\frac{1}{1+e^r}\right)}.$$

One can readily show that both U and \tilde{U} are decreasing functions. In addition, they always take positive values, because $F_h \geq F_\ell$ and $\tilde{F}_h \geq \tilde{F}_\ell$ in first order stochastic dominance. This means not only $\{\tilde{r}_n^h\}$, but also $\{r_n^h\}$ is an increasing sequence.

Lemma 5 implies that $\lim_n \tilde{r}_n^h = \infty$. Thus, to study the public belief at large times n , we need to understand $U(r)$ and $\tilde{U}(r)$ for large r . The next lemma provides the asymptotic behavior of these functions.

Lemma 6. *For large r , one has $U(r) = \Theta(e^{-\alpha r})$ and $\tilde{U}(r) = \Theta(e^{-\tilde{\alpha} r})$, that is*

$$0 < \liminf_{r \rightarrow \infty} \frac{U(r)}{e^{-\alpha r}} \leq \limsup_{r \rightarrow \infty} \frac{U(r)}{e^{-\alpha r}} < \infty,$$

and

$$0 < \liminf_{r \rightarrow \infty} \frac{\tilde{U}(r)}{e^{-\tilde{\alpha} r}} \leq \limsup_{r \rightarrow \infty} \frac{\tilde{U}(r)}{e^{-\tilde{\alpha} r}} < \infty.$$

Proof. Define $\mu = \frac{1}{1+e^r}$. We first propose an upper bound on U . To this end, note that

$$U(r) = \log \frac{1 - F_h(\mu)}{1 - F_\ell(\mu)} \leq -\log(1 - F_\ell(\mu)).$$

Due to Lemma 3, $F_\ell(q) \leq 2F(q)$, therefore $U(r) \leq -\log(1 - 2F(\mu))$. Since for small enough x , one has $-\log(1-x) \leq x + x^2$, then

$$U(r) \leq 2F(\mu)(1 + 2F(\mu)),$$

thereby establishing an upper bound.

Before proceeding with a lower bound, we introduce the Landau notations, $o(\cdot)$ and $O(\cdot)$: We say $f(x) = o(g(x))$ if $\lim_{x \rightarrow 0} \frac{f(x)}{g(x)} = 0$, and $f(x) = O(g(x))$ if $\limsup_{x \rightarrow 0} \frac{f(x)}{g(x)} < \infty$.

To propose a lower bound, observe that because of Lemma 3, $F_h(q) \leq 2qF(q)$ and $F_\ell(q) \geq 2(1-q)F(q)$, therefore,

$$\begin{aligned} e^{U(r)} &\geq \frac{1 - 2\mu F(\mu)}{1 - 2(1-\mu)F(\mu)} \\ &\geq (1 - 2\mu F(\mu))(1 + 2(1-\mu)F(\mu) + 2(1-\mu)^2 F(\mu)^2) \\ &= 1 + 2F(\mu) - 4\mu F(\mu) + 2(1 - O(\mu))F(\mu)^2. \end{aligned}$$

Since $\log(1+x) \geq x - x^2/2$, then

$$\begin{aligned} U(r) &\geq 2F(\mu) - 4\mu F(\mu) + (2 - O(\mu))F(\mu)^2 + o(F(\mu)^2) \\ &\geq 2F(\mu) \left(1 - 2\mu + \frac{3}{2}F(\mu) \right). \end{aligned}$$

The above upper and lower bounds imply that $\lim_{r \rightarrow \infty} \frac{U(r)}{2F(\mu)} = 1$. In addition, because of tail regularity (Assumption 2), it holds that $F(\mu) = \Theta(e^{-\alpha r})$, thereby justifying the lemma's first claim. A similar argument implies that $\tilde{U}(r) = \Theta(e^{-\tilde{\alpha} r})$. \square

So far our results sidestepped the role of the prior π , which determines the initial value for the sequence $\{\tilde{r}_n^h\}$, and only looked at the asymptotics as $n \rightarrow \infty$. In the next lemma, we establish a property of this sequence, that will prove useful for uniform convergence results. We use the notation $\tilde{r}_n^h(\pi)$ to refer to the value of \tilde{r}_n^h when the initial belief was π , that is, when $\tilde{r}_1^h = \log\left(\frac{\pi}{1-\pi}\right)$. The rest of the sequence evolves according to (B.2).

Lemma 7. *For every $\bar{r} \geq 0$, there exists n_0 such that $\tilde{r}_n^h(\pi) \geq \bar{r}$ for all $n \geq n_0$ and importantly for all initial $\pi \geq 1/2$.*

Proof. The idea is similar to the proof of Lemma 12 in Rosenberg and Vieille (2019). Let us introduce the mapping $\Psi(r) := r + \tilde{U}(r)$, and show its n -times composition by Ψ^n . Hence, one has $\tilde{r}_n^h = \Psi^{n-1}(\tilde{r}_1)$. First observe that since $\tilde{U} > 0$, if $\tilde{r}_1 \geq \bar{r}$, then $\Psi^n(\tilde{r}_1) \geq \bar{r}$. Now assume by contradiction that the conclusion of the lemma does not hold. Then, for every $n \in \mathbb{N}$, there exists an initial belief $\pi^{(n)}$ such that $\tilde{r}_n^h(\pi^{(n)}) \leq \bar{r}$. Also, one has $\Psi^{m-1}(\tilde{r}_1(\pi^{(n)})) = \tilde{r}_m^h(\pi^{(n)}) \leq \bar{r}$ for all $m \leq n$. Since the interval $[0, \bar{r}]$ is compact, there is a subsequence of initial values $\{\tilde{r}_1(\pi^{(k)})\}$, that we index by k , which is converging to $r_* \in [0, \bar{r}]$. Since the mapping Ψ^n is continuous for every fixed n , then one has

$$\Psi^n(r_*) = \lim_{k \rightarrow \infty} \Psi^n(\tilde{r}_1(\pi^{(k)})) \leq \bar{r}.$$

The above inequality holds for every n , hence it leads to a contradiction, because for every initial prior $\pi > 0$, the induced sequence $\{\tilde{r}_n^h\}$ increases to infinity (this follows from Lemma 5). \square

In the next two lemmas we calculate the asymptotic behavior of \tilde{r}_n^h and show that $e^{\tilde{r}_n^h} = \Theta(n^{1/\tilde{\alpha}})$. We first establish a lower bound for \tilde{r}_n^h . We achieve this by introducing a lower bound for the increments of \tilde{r}_n^h in (B.2). We then approximate the resulting lower envelope with the solution to a differential equation.

Lemma 8 (Lower Envelope). *The misspecified public log-likelihood satisfies*

$$\liminf_{n \rightarrow \infty} \frac{e^{\tilde{r}_n^h}}{n^{1/\tilde{\alpha}}} > 0. \quad (\text{B.3})$$

Proof. By Lemma 6, there exists $c > 0$ such that for all sufficiently large n (say $n \geq \bar{n}$), one has $\tilde{U}(\tilde{r}_n^h) \geq c e^{-\tilde{\alpha}\tilde{r}_n^h}$. Additionally, observe that the mapping $z \mapsto z + c e^{-\tilde{\alpha}z}$ is increasing for all sufficiently large z (say $z \geq \bar{z}$). Since $\tilde{r}_n^h \rightarrow \infty$, one can choose $N \geq \bar{n}$ such that $\tilde{r}_N \geq \bar{z}$. For all $n \geq N$ it holds that

$$\tilde{r}_{n+1}^h - \tilde{r}_n^h = \tilde{U}(\tilde{r}_n^h) \geq c e^{-\tilde{\alpha}\tilde{r}_n^h}. \quad (\text{B.4})$$

We show that this discrete time equation can be bounded from below by the following differential equation:

$$\frac{dz(t)}{dt} = c e^{-\tilde{\alpha}z(t)}.$$

This equation has the solution form $z(t) = \tilde{\alpha}^{-1} \log(\kappa + c\tilde{\alpha}t)$, where the initial condition parameter κ is chosen so that at $n = N$, we have $z(N) = \tilde{r}_N^h$. Next, we inductively show $\tilde{r}_n^h \geq z(n)$ for all $n \geq N$, which in turn establishes the claim in (B.3). The base step holds by definition. Suppose the claim also holds at some $n > N$, i.e., $\tilde{r}_n^h \geq z(n)$. Then, observe that because of the mean value theorem, there exists $t \in [n, n+1]$ such that

$$z(n+1) - z(n) = c e^{-\tilde{\alpha}z(t)} \leq c e^{-\tilde{\alpha}z(n)},$$

where the inequality follows because $z(t)$ is increasing. Therefore, one has

$$\begin{aligned} z(n+1) &\leq z(n) + c e^{-\tilde{\alpha}z(n)} \\ &\leq \tilde{r}_n^h + c e^{-\tilde{\alpha}\tilde{r}_n^h} \leq \tilde{r}_{n+1}^h. \end{aligned}$$

The second inequality holds because $z \mapsto z + c e^{-\tilde{\alpha}z}$ is increasing for $z \geq \bar{z}$, and $z(n) \geq \bar{z}$ for $n \geq N$. The third inequality holds because of (B.4). This justifies the claim in (B.3). \square

The next lemma posits an upper bound for the increments of \tilde{r}_n^h . Its proof strategy is similar to that of the previous lemma, with some additional technical considerations.

Lemma 9 (Upper Envelope). *The misspecified public log-likelihood satisfies*

$$\limsup_{n \rightarrow \infty} \frac{e^{\tilde{r}_n^h}}{n^{1/\tilde{\alpha}}} < \infty. \quad (\text{B.5})$$

Proof. As in the proof of the previous lemma—but changing the direction of the inequalities—there exists $C > 0$ and \bar{n} such that for all $n \geq \bar{n}$ one has $\tilde{U}(\tilde{r}_n^h) \leq C e^{-\tilde{\alpha}\tilde{r}_n^h}$. Likewise, the

mapping $z \mapsto z + Ce^{-\tilde{\alpha}z}$ is increasing for all $z \geq \bar{z}$, and we can choose $N \geq \bar{n}$ such that $\tilde{r}_N^h \geq \bar{z}$. Then for all $n \geq N$ it holds that

$$\tilde{r}_{n+1}^h - \tilde{r}_n^h = \tilde{U}(\tilde{r}_n^h) \leq Ce^{-\tilde{\alpha}\tilde{r}_n^h}. \quad (\text{B.6})$$

Take the following differential equation as an upper envelope for the above difference equation:

$$\frac{dz(t)}{dt} = 2Ce^{-\tilde{\alpha}z(t)},$$

with the solution form $z(t) = \tilde{\alpha}^{-1} \log(\kappa + 2C\tilde{\alpha}t)$. Observe that for all $\kappa > 0$ and $n \geq 1$ one has

$$2e^{-\tilde{\alpha}z(n+1)} \geq e^{-\tilde{\alpha}z(n)}. \quad (\text{B.7})$$

Therefore, one can choose κ large enough such that $z(N) \geq \tilde{r}_N^h$. Next, we inductively show $\tilde{r}_n^h \leq z(n)$ for all $n \geq N$, which in turn establishes the claim in (B.5). The base step holds by definition. Suppose the claim also holds at some $n > N$. Then, observe that because of the mean value theorem, there exists $t \in [n, n+1]$ such that

$$z(n+1) - z(n) = 2Ce^{-\tilde{\alpha}z(t)} \geq 2Ce^{-\tilde{\alpha}z(n+1)} \geq Ce^{-\tilde{\alpha}z(n)},$$

where the first inequality above holds because $z(t)$ is increasing and the second inequality follows from (B.7). Since $z \mapsto z + Ce^{-\tilde{\alpha}z}$ is increasing for all $z \geq \bar{z}$, and $n \geq N$, then $z(n) \geq \tilde{r}_n^h \geq \bar{z}$ implies that

$$z(n+1) \geq z(n) + Ce^{-\tilde{\alpha}z(n)} \geq \tilde{r}_n^h + Ce^{-\tilde{\alpha}\tilde{r}_n^h} \geq \tilde{r}_{n+1}^h,$$

where the last inequality follows from equation (B.6). This concludes the induction and thus establishes the asymptotic upper bound for $e^{\tilde{r}_n^h}$ in (B.5). \square

The previous two lemmas jointly imply that $e^{\tilde{r}_n^h} = \Theta(n^{1/\tilde{\alpha}})$. Importantly this holds regardless of the initial belief π (i.e., the initial level \tilde{r}_1). Of course, the implied constants may depend on π .

C Characterization of Asymptotic Learning

In Section 4.3 we drew a connection between asymptotic learning and immediate herding. In this section we formalize this, establishing necessary and sufficient conditions for asymptotic learning in terms of immediate herding. Note that the results of this section, Lemmas 10 and 11 do not require tail-regularity and apply more broadly.

The first lemma states that asymptotic learning in the high state implies that immediate herding on the high action happens with positive probability for some prior π' , and immediate herding on the low action cannot occur.

Lemma 10 (Necessary condition). *Assume $a_n \rightarrow h$, \mathbb{P}_h -almost surely. Then, the following two conditions hold:*

$$(i) \exists \pi' < 1 \text{ such that } \mathbb{P}_{\pi',h}(\bar{a} = h) > 0,$$

$$(ii) \mathbb{P}_h(\bar{a} = \ell) = 0.$$

Proof. Condition on $\theta = h$, and let σ be the random time of the last incorrect action, which has to be finite, because $a_n \rightarrow h$. Since $a_\sigma = \ell$, then it must be that $\tilde{\pi}_{\sigma+1} < 1/2$, by the overturning principle. Therefore,

$$\begin{aligned} 1 = \mathbb{P}_h(a_n \rightarrow h) &= \sum_{k=0}^{\infty} \mathbb{P}_h(\sigma = k) \\ &= \sum_{k=0}^{\infty} \mathbb{P}_h(a_m = h \forall m > k, \tilde{\pi}_{k+1} < 1/2). \end{aligned}$$

Applying the law of total expectations, this is equal to

$$\begin{aligned} &= \sum_{k=0}^{\infty} \mathbb{E}_h [\mathbb{P}_h(a_m = h \forall m > k, \tilde{\pi}_{k+1} < 1/2 \mid \tilde{\pi}_{k+1})] \\ &= \sum_{k=0}^{\infty} \mathbb{E}_h [\mathbb{P}_h(a_m = h \forall m > k \mid \tilde{\pi}_{k+1}) \mathbb{1}_{\{\tilde{\pi}_{k+1} < 1/2\}}] \tag{C.1} \\ &= \sum_{k=0}^{\infty} \mathbb{E}_h [\mathbb{P}_{\tilde{\pi}_{k+1},h}(\bar{a} = h) \mathbb{1}_{\{\tilde{\pi}_{k+1} < 1/2\}}], \end{aligned}$$

where the last equality is an application of stationarity. To show (i), assume by contradiction that $\mathbb{P}_{\pi',h}(\bar{a} = h) = 0$ for every $\pi' \in [0, 1)$. Then the right hand side is equal to zero, thereby resulting in a contradiction.

Since the event $a_n \rightarrow h$ is disjoint from the event $\bar{a} = \ell$, the assumption that the former happens with probability one implies that the latter has probability zero, and thus we have shown (ii). \square

The next lemma shows that asymptotic learning in the high state is implied by uniformly positive probability (over priors at least one half) for immediate herding on the high action, and zero probability (for any prior at least one half) for immediate herding on the low action.

Lemma 11 (Sufficient condition). *The following two conditions are sufficient for $a_n \rightarrow h$, \mathbb{P}_h -almost surely:*

(i) $\inf_{\pi' \geq 1/2} \mathbb{P}_{\pi',h}(\bar{a} = h) > 0$,

(ii) $\mathbb{P}_{\pi',h}(\bar{a} = \ell) = 0$ for all $\pi' > 0$.

Proof. To show asymptotic learning, we first rule out convergence to the wrong action, that is we claim $\mathbb{P}_h(a_n \rightarrow \ell) = 0$. Let σ be the last time that agents take the correct action h , hence $\tilde{\pi}_{\sigma+1} \geq 1/2$. Then, $a_n \rightarrow \ell$ iff $\sigma < \infty$. Therefore, applying the same logic of equation (C.1) leads to

$$\mathbb{P}_h(a_n \rightarrow \ell) = \sum_{k=0}^{\infty} \mathbb{P}_h(\sigma = k) = \sum_{k=0}^{\infty} \mathbb{E}_h \left[\mathbb{P}_{\tilde{\pi}_{k+1},h}(\bar{a} = \ell) \mathbb{1}_{\{\tilde{\pi}_{k+1} \geq 1/2\}} \right].$$

Since $\mathbb{P}_{\pi',h}(\bar{a} = \ell) = 0$ for all $\pi' > 0$ the above expression implies that $\mathbb{P}_h(a_n \rightarrow \ell) = 0$.

As we have shown that a_n does not converge to ℓ , it follows that the sequence of actions a_n has some number of bad runs: consecutive agents who take the wrong action, flanked by agents who take the correct action. To show asymptotic learning it suffices to show that the number of bad runs is finite. Let

$$\delta = \inf_{\pi' \geq \frac{1}{2}} \mathbb{P}_{\pi',h}(\bar{a} = h).$$

At the end of a bad run the next action is h , and so the misspecified public belief is at least $1/2$. Hence, by stationarity, there is a chance of at least δ of never having another bad run. Since signals are independent conditioned on the state, this implies that the probability of having m bad runs is at most $(1-\delta)^m$. In particular, the probability of infinitely many bad runs is zero. \square

D Proof of Theorem 1

We divide the proof of Theorem 1 into two: Proposition 3 characterizes asymptotic learning, and Proposition 6 characterizes efficient learning. Jointly, they imply the theorem.

D.1 Parametric Characterization for Asymptotic Learning

In this section we characterize the range of condensation where asymptotic learning is achieved.

Proposition 3. *The following are equivalent:*

(i) *Asymptotic learning*;

(ii) $\tilde{\alpha} - \alpha \in [0, 1)$.

To prove this statement we leverage the necessary and sufficient conditions found in Lemmas 10 and 11, as well as prove the following two propositions that relate the probability of immediate herding events to the underlying parameters α and $\tilde{\alpha}$:

Proposition 4. *The following are equivalent:*

(i) $\inf_{\pi' \geq 1/2} \mathbf{P}_{\pi', h}(\bar{a} = h) > 0$;

(ii) $\exists \pi' < 1$ such that $\mathbf{P}_{\pi', h}(\bar{a} = h) > 0$;

(iii) $\tilde{\alpha} - \alpha < 1$.

Proposition 5. *The following are equivalent:*

(i) $\mathbf{P}_{\pi', h}(\bar{a} = \ell) = 0$ for all $\pi' > 0$;

(ii) $\mathbf{P}_{\pi', h}(\bar{a} = \ell) = 0$ for some $\pi' < 1$;

(iii) $\tilde{\alpha} - \alpha \geq 0$.

Proposition 4 implies that the probability of an immediate good herd is *uniformly* positive (over all initial beliefs larger than 1/2) if and only if $\tilde{\alpha} - \alpha < 1$, i.e., when agents are not overly condescending. Proposition 5 claims that the probability of an immediate wrong herd is zero for all positive initial beliefs if and only if $\tilde{\alpha} - \alpha \geq 0$, i.e., when agents are condescending.

We use these propositions to prove Proposition 3, before proceeding with their proofs.

Proof of Proposition 3. Suppose that asymptotic learning holds, i.e., $a_n \rightarrow \theta$, \mathbf{P} -almost surely. Then $a_n \rightarrow h$, \mathbf{P}_h -almost surely. By Lemma 10, this implies that condition (ii) of both Propositions 4 and 5 hold. Hence, by these propositions, conditions (iii) in the two propositions hold, and $\tilde{\alpha} - \alpha \in [0, 1)$.

Suppose that $\tilde{\alpha} - \alpha \in [0, 1)$. Then condition (iii) of both Propositions 4 and 5 hold. Therefore, condition (i) of both propositions hold. Hence, by Lemma 11, we have asymptotic learning. \square

D.1.1 Proof of Proposition 4

The proof of the first implication, namely (i) \Rightarrow (ii), is immediate. The next lemma establishes the second implication, i.e., (ii) \Rightarrow (iii). In fact, it shows a stronger statement.

Lemma 12. *The following are equivalent:*

$$(i) \quad \tilde{\alpha} - \alpha < 1;$$

$$(ii) \quad \mathbf{P}_h(\bar{a} = h) > 0.$$

Proof. Conditioned on the event $\bar{a} = h$, the public belief $\tilde{\pi}_n$ is deterministic and equals $\tilde{\pi}_n^h$. Thus the event $\bar{a} = h$ is equal to the event $\{q_n + \tilde{\pi}_n^h \geq 1, \forall n\}$. Since the random variables q_n are independent conditioned on $\theta = h$, we have that

$$\mathbf{P}_h(\bar{a} = h) = \prod_n \mathbf{P}_h(q_n + \tilde{\pi}_n^h \geq 1) = \prod_n (1 - F_h(1 - \tilde{\pi}_n^h)). \quad (\text{D.1})$$

This implies that $\mathbf{P}_h(\bar{a} = h) > 0$ if and only if $-\sum_n \log(1 - F_h(1 - \tilde{\pi}_n^h)) < \infty$. For two sequences f_n and g_n , we say $f_n \sim g_n$ if $\frac{f_n}{g_n} \rightarrow 1$ as $n \rightarrow \infty$. Since $\tilde{\pi}_n^h \rightarrow 1$, then $-\log(1 - F_h(1 - \tilde{\pi}_n^h)) \sim F_h(1 - \tilde{\pi}_n^h)$, and the previous sum is finite if and only if

$$\sum_n F_h(1 - \tilde{\pi}_n^h) < \infty.$$

Observe that Lemma 4 implies that $F_h(q) = \Theta(q^{\alpha+1})$. Also as $n \rightarrow \infty$, we have $e^{-\tilde{r}_n^h} \sim 1 - \tilde{\pi}_n^h$, therefore, the above sum is finite if and only if

$$\sum_n e^{-(\alpha+1)\tilde{r}_n^h} < \infty. \quad (\text{D.2})$$

Because of the Lemmas 8 and 9, one has $e^{-(\alpha+1)\tilde{r}_n^h} = \Theta\left(n^{-\frac{\alpha+1}{\tilde{\alpha}}}\right)$. Thus, the sum in (D.2) is finite if and only if $\tilde{\alpha} - \alpha < 1$. \square

The following two lemmas are aimed at proving the third and final implication in Proposition 4, that is (iii) \Rightarrow (i). In the first one, we show that the sum in (D.2) can be made arbitrarily small if the initial value \tilde{r}_1 is large enough. Often in the following expressions, we use the notation $\tilde{r}_n^h(r)$ to refer to the process initiated at $\tilde{r}_1 = r$. Also, recall our former notation, where we used $\tilde{r}_n^h(\pi)$ to refer to the process initiated at $\tilde{r}_1 = \log \frac{\pi}{1-\pi}$. We use both of these notations interchangeably depending on the context.

Lemma 13. *Assume $\tilde{\alpha} - \alpha < 1$. Then for every $\varepsilon > 0$, there exists $\bar{r} \geq 0$ such that for all $r \geq \bar{r}$ one has $\sum_{n \geq 0} e^{-(\alpha+1)\tilde{r}_n^h(r)} < \varepsilon$.*

Proof. We appeal to the idea used in the proof of Lemma 8. Since $\tilde{U}(r) = \Theta(e^{-\tilde{\alpha}r})$, then there exists $c > 0$, and correspondingly a threshold \bar{r} , such that $\tilde{U}(r) \geq c e^{-\tilde{\alpha}r}$ for every $r \geq \bar{r}$ and the mapping $r \mapsto r + c e^{-\tilde{\alpha}r}$ is increasing on $[\bar{r}, \infty)$. In particular, since \tilde{r}_n^h is increasing in n , starting the process at any $\tilde{r}_1 = r \geq \bar{r}$, implies

$$\tilde{U}(\tilde{r}_n^h(r)) \geq c e^{-\tilde{\alpha}\tilde{r}_n^h(r)}.$$

Next, we recall the continuous time process $z(t)$ such that $z(0) = \bar{r}$, and

$$\frac{dz(t)}{dt} = c e^{-\tilde{\alpha}z(t)}.$$

The solution to this differential equation takes the form

$$z(t) = \frac{1}{\tilde{\alpha}} \log \left(e^{\tilde{\alpha}\bar{r}} + c\tilde{\alpha}t \right).$$

Using induction, similar to the one used in Lemma 8, we can show $\tilde{r}_n^h(r) \geq z(n)$ for every initial value $r \geq \bar{r}$. Therefore, for every $r \geq \bar{r}$, it holds that

$$\sum_{n \geq 0} e^{-(\alpha+1)\tilde{r}_n^h(r)} \leq \sum_{n \geq 0} e^{-(\alpha+1)z(n)} = \sum_{n \geq 0} \left(e^{\tilde{\alpha}\bar{r}} + c\tilde{\alpha}n \right)^{-\frac{\alpha+1}{\tilde{\alpha}}}.$$

Since $\alpha+1 > \tilde{\alpha}$, for a given $\varepsilon > 0$, we can choose \bar{r} large enough, such that the above sum is less than ε . \square

Let $\eta: [1/2, 1] \rightarrow [0, 1]$,

$$\eta(\pi) = \mathbb{P}_{\pi, h}(\bar{a} = h),$$

be the probability of immediate herding on the high action, conditioned on the high state, when the prior is π . By (D.1),

$$\eta(\pi) = \prod_n (1 - F_h(1 - \tilde{\pi}_n^h(\pi))).$$

Lemma 14. *Assume $\tilde{\alpha} - \alpha < 1$. Then η is continuous.*

Proof. Let

$$\eta_n(\pi) = \mathbb{P}_{\pi, h}(a_1 = h, \dots, a_n = h) = \prod_{k=1}^n (1 - F_h(1 - \tilde{\pi}_k^h(\pi))), \quad (\text{D.3})$$

be the probability that the first n agents take the high action, conditioned on the high state, when the prior is π . By definition, $\eta(\pi) = \lim_n \eta_n(\pi)$. Since the distribution of the private

posteriors q_n is non-atomic, each η_n is continuous. Thus, we prove that η is continuous by showing that η_n converges uniformly to η .

First, Lemma 4 implies that $F_h(e^{-r}) = \Theta(e^{-(\alpha+1)r})$, and hence there exists $C > 0$ such that $F_h(e^{-r}) \leq Ce^{-(\alpha+1)r}$. Second, because of Lemma 13, for a given $\varepsilon_1 > 0$, there exists $\bar{r} \geq 0$ such that for all $r_1 \geq \bar{r}$, one has

$$\sum_{n \geq 0} e^{-(\alpha+1)\bar{r}_n^h(r_1)} \leq \varepsilon_1.$$

Then, because of Lemma 7 there exists $n_0 \equiv n_0(\bar{r})$ such that $\tilde{r}_n^h(\pi) \geq \bar{r}$ for all initial $\pi \geq 1/2$, and $n \geq n_0$. By (D.3),

$$\eta_{n+1}(\pi) = \eta_n(\pi)(1 - F_h(1 - \tilde{\pi}_n^h(\pi))),$$

so that $|\eta_{n+1}(\pi) - \eta_n(\pi)| \leq F_h(1 - \tilde{\pi}_n^h(\pi))$. Hence, for every $k > 0$, and $\pi \geq 1/2$,

$$|\eta_{n_0+k}(\pi) - \eta_{n_0}(\pi)| \leq \sum_{n=n_0}^{\infty} F_h(e^{-\tilde{r}_n^h(\pi)}) \leq C \sum_{n \geq n_0} e^{-(\alpha+1)\tilde{r}_n^h(\pi)} \leq C\varepsilon_1. \quad (\text{D.4})$$

The third inequality above holds because $\tilde{r}_{n_0}^h(\pi) \geq \bar{r}$, and thus Lemma 13 implies the sum is smaller than ε_1 . Since ε_1 was chosen independently, the final term above can be made arbitrarily small, by taking n_0 large enough. This implies the sequence $\{\eta_n\}$ is Cauchy w.r.t. the sup-norm in $C[1/2, 1]$, and thus it converges uniformly to η . Therefore, η is continuous. \square

Using the above lemma, we can now conclude the proof of the last implication in Proposition 4, namely (iii) \Rightarrow (i). Assume by contradiction, that condition (i) does not hold, then $\inf_{\pi \geq 1/2} \eta(\pi) = 0$. By the previous lemma η is a continuous function, hence there must exist $\hat{\pi} \in [1/2, 1]$ such that $\eta(\hat{\pi}) = 0$. Since, $\eta(1) \neq 0$, then $\hat{\pi} \in [1/2, 1)$ and Lemma 12 implies that $\tilde{\alpha} - \alpha \geq 1$. This violates the initial assumption (i.e., $\tilde{\alpha} - \alpha < 1$) and hence concludes the proof of Proposition 4.

D.1.2 Proof of Proposition 5

The first implication, namely (i) \Rightarrow (ii), is immediate. For the remaining two implications, define

$$\xi(\pi) = \mathbf{P}_{\pi, h}(\bar{a} = \ell),$$

and appeal to the next lemma.

Lemma 15. *For every $\pi \in (0, 1)$, one has $\xi(\pi) = 0$ if and only if $\tilde{\alpha} - \alpha \geq 0$.*

Proof. Because of the symmetry assumption, we have $\xi(\pi) = \mathbf{P}_{1-\pi, \ell}(\bar{a} = h)$. Let $\tilde{\pi}_n^h = \tilde{\pi}_n^h(1-\pi)$ be the misspecified public belief on the high action path, initiated at $\tilde{\pi}_1 = 1-\pi$. Then, following the same argument of Lemma 12, one has

$$\xi(\pi) = \prod_n (1 - F_\ell(1 - \tilde{\pi}_n^h(1-\pi))).$$

Therefore, $\xi(\pi) > 0$ if and only if $-\sum_n \log(1 - F_\ell(1 - \tilde{\pi}_n^h)) < \infty$. Since on the high action path $\tilde{\pi}_n^h \rightarrow 1$, then $-\log(1 - F_\ell(1 - \tilde{\pi}_n^h)) \sim F_\ell(1 - \tilde{\pi}_n^h)$, and the previous sum is finite if and only if

$$\sum_n F_\ell(1 - \tilde{\pi}_n^h) < \infty.$$

Lemma 3 implies that $F_\ell(q) = \Theta(q^\alpha)$. Also, as $n \rightarrow \infty$, we have $1 - \tilde{\pi}_n^h \sim e^{-\tilde{r}_n^h}$, therefore the above sum is finite if and only if

$$\sum_n e^{-\alpha \tilde{r}_n^h} < \infty.$$

It was shown in Lemmas 8 and 9 that $e^{\tilde{r}_n} = \Theta(n^{1/\tilde{\alpha}})$, thus one can deduce that the above sum is finite if and only if $\alpha > \tilde{\alpha}$. Therefore, $\xi(\pi) = 0$ if and only if $\tilde{\alpha} - \alpha \geq 0$. \square

Observe that $\xi(1) = 0$. Therefore, the second and the third implications of Proposition 5, namely (ii) \Rightarrow (iii) \Rightarrow (i), respectively follow from the above lemma, thereby concluding the proof of Proposition 5.

D.2 Parametric Characterization for Efficient Learning

In this section we characterize the range of condensation where efficient learning is achieved.

Proposition 6. *Assume $\tilde{\alpha} \neq \alpha$. The following are equivalent:*

- (i) *Efficient learning;*
- (ii) $\tilde{\alpha} - \alpha \in (0, 1)$.

Proof. The implication (i) \Rightarrow (ii) follows immediately from Proposition 3, since efficient learning implies asymptotic learning.

Towards sufficiency, assume $\tilde{\alpha} - \alpha < 1$. Following the same logic as in the proof of Lemma 11, one obtains that conditioned on $\theta = h$, the probability of having m bad runs is at most $(1-\delta)^m$ for some $\delta > 0$, and hence the number of bad runs has a finite expectation.

By Proposition 7 below, conditioned on the high state, $\tilde{\alpha} - \alpha > 0$ implies that the expected length of the first bad run is bounded by $C_0 \frac{1-\pi}{\pi}$, for some constant $C_0 > 0$. This proposition

also implies, by stationarity, that conditioned on $\theta = h$ and on any prior history, the expected length of any future bad run is at most $C_0 B$, where $B > 0$ is another constant. It thus follows from the fact that signals are conditionally i.i.d. that the expected total number of low actions in the high state is finite. The argument is analogous to the one that appears in Appendix B.3 of [Rosenberg and Vieille \(2019\)](#).

Finally, by symmetry, the expected number of high actions in the low state is also finite, and thus we have efficient learning. □

We end this second with the following proposition, which is the main ingredient of the proof above. It shows that $\tilde{\alpha} - \alpha > 0$ implies that the expected length of a bad run is uniformly bounded.

Define $\tau_\theta := \min\{n : a_n = \theta\}$. Note that conditioned on $\theta = h$, τ_h is the length of the first bad run.

Proposition 7. *Assume $\tilde{\alpha} - \alpha > 0$, then the following statements hold:*

(i) *Let $\pi \leq 1/2$. There exists a constant $C_0 > 0$ (independent of π) such that*

$$\mathbf{E}_{\pi,h} [\tau_h] \leq C_0 \frac{1-\pi}{\pi} .$$

(ii) *Let $\tilde{\pi}_{n+1}$ be the misspecified public belief after observing a history ending with $a_{n-1} = h$ and $a_n = \ell$. Then $\frac{1-\tilde{\pi}_{n+1}}{\tilde{\pi}_{n+1}} \leq B$ for some constant $B < \infty$ that does not depend on the history.*

Proof. To see (i), observe that because of symmetry, one has $\mathbf{E}_{\pi,h} [\tau_h] = \mathbf{E}_{1-\pi,\ell} [\tau_\ell]$. Also, it holds that

$$\mathbf{E}_{1-\pi,\ell} [\tau_\ell] = 1 + \sum_{n \geq 1} \mathbf{P}_{1-\pi,\ell} (\tau_\ell > n) .$$

By Bayes Law

$$\mathbf{P}_{1-\pi,\ell} (\tau_\ell > n) = \frac{1-\pi}{\pi} \times \frac{1-\pi_{n+1}^h}{\pi_{n+1}^h} \mathbf{P}_{1-\pi,h} (a_1 = \dots = a_n = h) ,$$

where π_n^h is the correctly specified public belief on the high action path, starting at $1-\pi$ and following the dynamics in equation (4.2a). Recall that r_n^h represents the correctly specified public log-likelihood on the high action path, which follows the dynamics in equation (B.1), namely $r_{n+1}^h - r_n^h = U(\tilde{r}_n^h)$. Hence, the above expression implies that

$$\mathbf{E}_{1-\pi,\ell} [\tau_\ell] \leq 1 + \frac{1-\pi}{\pi} \sum_{n \geq 1} e^{-r_{n+1}^h} . \tag{D.5}$$

Since the initial belief is set to $1-\pi$ and it is assumed in part (i) that $\pi \leq 1/2$, then $r_n^h \geq 0$. Next, observe that on the path $\bar{a} = h$, the misspecified public log-likelihood follows the difference equation (B.2), namely $\tilde{r}_{n+1}^h - \tilde{r}_n^h = \tilde{U}(\tilde{r}_n^h)$. Additionally, because of Lemma 6, there exists $C > 0$ such that $\tilde{U}(\tilde{r}_n^h) \leq Ce^{-\tilde{\alpha}\tilde{r}_n^h}$. We continue by finding a continuous time upper envelope for \tilde{r}_n^h —analogous to Lemma 9 with a slight catch in selecting the initial condition. Choose $\bar{r} > 0$ such that the mapping $r \mapsto r + Ce^{-\tilde{\alpha}r}$ becomes increasing on $[\bar{r}, \infty)$. Let $n_0 := \min\{n : \tilde{r}_n^h \geq \bar{r}\}$ that is finite because $\tilde{r}_n \rightarrow \infty$ on the high action path. Since $\tilde{U}(\cdot)$ is a decreasing function, then $\bar{r} \leq \tilde{r}_{n_0}^h \leq \bar{r} + \tilde{U}(0)$. Let $z(t)$ be the solution to the following differential equation

$$\frac{dz(t)}{dt} = 2Ce^{-\tilde{\alpha}z(t)},$$

starting at $z(0) = \bar{r} + \tilde{U}(0)$. Therefore, $z(t) = \tilde{\alpha}^{-1} \log \left(e^{\bar{r} + \tilde{U}(0)} + 2C\tilde{\alpha}t \right)$. Following the recipe of Lemma 9, one can show by induction that $z(k) \geq \tilde{r}_{k_0+n}^h$ for all $k \geq 0$. Next, we examine Raabe's criterion⁵ for the infinite sum $\sum_{n \geq 1} e^{-r_n^h}$, that is to examine the limit of the following expression:

$$n \left(\frac{e^{-r_n^h}}{e^{-r_{n+1}^h}} - 1 \right) = n \left(e^{U(\tilde{r}_n^h)} - 1 \right) \geq nU(\tilde{r}_n^h).$$

Note that U is decreasing, that $z(k) \geq \tilde{r}_{k_0+n}^h$ for all $k \geq 0$, and that there exists $c > 0$ such that $U(z) \geq ce^{-\alpha z}$. Hence

$$\liminf_{n \rightarrow \infty} n \left(\frac{e^{-r_n^h}}{e^{-r_{n+1}^h}} - 1 \right) \geq \limsup_{k \rightarrow \infty} ck e^{-\alpha z(k)} = \limsup_{k \rightarrow \infty} \frac{ck}{\left(e^{\bar{r} + \tilde{U}(0)} + 2C\tilde{\alpha}k \right)^{\alpha/\tilde{\alpha}}}.$$

Since $\tilde{\alpha} > \alpha$, the limit superior on the right hand side above is infinite and thus the sum $\sum_{n \geq 1} e^{-r_n^h}$ is convergent. Together with (D.5), this implies that there exists a constant $C_0 > 0$ such that $\mathbf{E}_{\pi,h}[\tau_h] \leq C_0 \left(\frac{1-\pi}{\pi} \right)$. This establishes (i).

To see (ii), condition on the event $\{a_{n-1} = h, a_n = \ell\}$. Equivalently, $\tilde{\pi}_n \geq 1/2$ and $\tilde{\pi}_{n+1} \leq 1/2$, by the overturning principle. Then, Bayes law implies

$$\frac{1 - \tilde{\pi}_{n+1}}{\tilde{\pi}_{n+1}} = \frac{1 - \tilde{\pi}_n}{\tilde{\pi}_n} \times \frac{\tilde{F}_\ell(1 - \tilde{\pi}_n)}{\tilde{F}_h(1 - \tilde{\pi}_n)}.$$

Since $\tilde{F}_\ell(q) = \Theta(q^{\tilde{\alpha}})$ and $\tilde{F}_h(q) = \Theta(q^{\tilde{\alpha}+1})$, then, there are constants $C > 0$ and $c > 0$ such that $\tilde{F}_\ell(q) \leq Cq^{\tilde{\alpha}}$ and $\tilde{F}_h(q) \geq cq^{\tilde{\alpha}+1}$ for all $q \in [0, 1/2]$. Therefore,

$$\frac{1 - \tilde{\pi}_{n+1}}{\tilde{\pi}_{n+1}} \leq \frac{1 - \tilde{\pi}_n}{\tilde{\pi}_n} \frac{C(1 - \tilde{\pi}_n)^{\tilde{\alpha}}}{c(1 - \tilde{\pi}_n)^{\tilde{\alpha}+1}} \leq \frac{2C}{c}.$$

⁵ Raabe's criterion for convergence of sums states that $\sum_n \nu_n$ converges if $\liminf_n \rho_n > 1$ and diverges if $\limsup_n \rho_n < 1$, where $\rho_n = n(\nu_n/\nu_{n+1} - 1)$. It is inconclusive when $\lim_n \rho_n = 1$. This latter case corresponds to $\tilde{\alpha} = \alpha$.

This establishes (ii). □

E Proofs of Propositions 1 and 2

Proof of Proposition 1. Suppose that $\tilde{\alpha} < \alpha$. Then, by Proposition 5 one has $P_h(\bar{a} = \ell) > 0$, so that a wrong herd forms immediately with positive probability. □

Proof of Proposition 2. Suppose that $\tilde{\alpha} \geq \alpha + 1$. Condition on the high state. Then, by Proposition 4, for any prior $\pi' < 1$, the probability of an immediate herd on the high action is zero. Hence, by stationarity, the probability that $a_n \rightarrow h$ is zero. By Proposition 5, the probability of an immediate herd on the low action is also zero, and hence, again by stationarity, the probability that $a_n \rightarrow \ell$ is zero. Thus the agents take both actions infinitely many times. The same argument applies when conditioning on the low state. □