The Hazards and Benefits of Condescension in Social Learning

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

January 26, 2023

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 bad outcomes, as does anti-condescension.

JEL Classification numbers: .

Keywords: .

*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 follows 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 overestimate the informativeness of their own signals. Through mostly numerical analysis, the authors reach a conclusion that is similar to ours: Moderate overconfidence is beneficial for society.

A large empirical literature has documented overconfidence. 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. Overconfidence is closely related to condescension, and may be hard to distinguish empirically, and so some of these findings could be interpreted as evidence for condescension.

Our paper is related to two additional, recent strands of the literature. First, we contribute to the social learning literature 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). 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. Our result highlights potential positive welfare effects, i.e., misspecifications may increase the efficiency of learning.

Second, we complement a burgeoning literature which analyzes the rationale for the presence of persistent misspecifications from an evolutionary perspective (e.g., He and Libgober, 2020; 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 $= \{1, 2, \ldots\}$ 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 $I_n = (a_1, \ldots, a_{n-1})$.

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 P, and explicitly write P_π when we want to highlight varying values of the prior π. We further use the notation P_h to refer to $\mathsf{P}(\cdot \mid \theta = h)$, the probability measure P conditional on the realized state being h. We define $\mathsf{P}_{\pi,h}$ analogously.

Let $q_n = \mathsf{P}(\theta = h | s_n)$ 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. Namely, each agent believes that the others' private posteriors have conditional distributions $\widetilde{F}_\ell$ and $\widetilde{F}_h$. 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 $q_n = \mathsf{P}(\theta = h|s_n)$; since they observe q_n, it is immaterial what their beliefs are regarding its prior distribution. We denote by $\widetilde{\mathsf{P}}$ posterior probabilities calculated according to the agents' misspecified beliefs.

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

$$a_n = \underset{a \in \{\ell, h\}}{\arg\max}\, \widetilde{\mathsf{P}}(\theta = a | I_n, q_n).$$

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 $\widetilde{F}_\ell$ and $\widetilde{F}_h$ 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 2.1 (Symmetry)**.** We assume throughout that (F_ℓ, F_h) and $(\widetilde{F}_\ell, \widetilde{F}_h)$ 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 2.2.

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

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

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: $\mathsf{E}[W]$, 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 $\int_0^1 \frac{1}{F(x)} \mathrm{d}x < \infty$.

2.4 Condescension

We use our misspecified social learning framework to study how outcomes change when agents are *condescending*, 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

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

We use Landau notation and write

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

as a shorthand for the expression above.[1]

Assumption 2.3 (Tail-Regularity)**.** We assume throughout that (F_ℓ, F_h) and $(\widetilde{F}_\ell, \tilde{F}_h)$ are tail-regular.

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

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

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 A.2) if $F(q) = \Theta(q^\alpha)$ then $F_\ell(q) = \Theta(q^\alpha)$ and $F_h(q) = \Theta(q^{\alpha+1})$.

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

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

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

3 Results

Our first result characterizes the efficiency of learning outcomes for condescending and anti-condescending agents.

Theorem 3.1. *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 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 3.1 shows that efficiency can be regained under potentially small misspecification, for any α, as long as agents are condescending, but not too condescending.

Our next result tackles the question of why learning fails when agents are anti-condescending, i.e., when $\tilde{\alpha} < \alpha$.

Theorem 3.2. *Suppose that $\tilde{\alpha} < \alpha$. Then, with 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. Thus, wrong herds can form. This is despite the fact that signals are unbounded, which, without misspecification, would rule out wrong herds.

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

Theorem 3.3. *Suppose that $\tilde{\alpha} \geqslant \alpha + 1$. Then, P-almost surely, both actions are taken by infinitely many agents.*

In the case $\tilde{\alpha} \geqslant \alpha + 1$, agents are very condescending: 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.

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:

$$\pi_n = \mathsf{P}\left(\theta = h \middle| a_1, \ldots, a_{n-1}\right).$$

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

$$\tilde{\pi}_n = \widetilde{\mathsf{P}}\left(\theta = h \middle| a_1, \ldots, a_{n-1}\right).$$

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 $p_n = \mathsf{P}(\theta = h|I_n, q_n)$ 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 $\tilde{p}_n = \widetilde{\mathsf{P}}(\theta = h|I_n, q_n)$. Then by Bayes' Law

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

It follows that the action a_n chosen by agent n is equal to h if $\tilde{\pi}_n + q_n \geqslant 1$, and to ℓ otherwise.[2] 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:

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

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

[2] As 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,

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

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

4.2 Stationarity and the Overturning Principle

These 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 4.1 (Stationarity)**.** *For any fixed sequence $b_1, \ldots, b_k$ of actions in $\{\ell, h\}$ and any $\tilde{\pi} \in (0,1)$,*

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

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 $b_1, \ldots, b_k$ 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 4.2 (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

$$\begin{aligned}
\tilde{\pi}_{n+1} &= \widetilde{\mathsf{E}}\left[\mathbb{1}_{\{\theta=h\}} \middle| a_1, \ldots, a_n\right] \\
&= \widetilde{\mathsf{E}}\left[\widetilde{\mathsf{E}}\left[\mathbb{1}_{\{\theta=h\}} \middle| a_1, \ldots, a_n, q_n\right] \middle| a_1, \ldots, a_n\right] \\
&= \widetilde{\mathsf{E}}\left[\tilde{p}_n \middle| a_1, \ldots, a_n\right].
\end{aligned}$$

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

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 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 $a_n \to h$ 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 $\{a_1 = h, a_2 = h, \ldots\}$ that all agents took the high action; this is the event that a high action herd formed immediately.

Asymptotic learning occurs when $\mathsf{P}_\ell(a_n \to \ell) = 1$ and $\mathsf{P}_h(a_n \to h) = 1$. To study the asymptotic events $a_n \to \ell$ and $a_n \to h$, 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 and immediate learning, condition on $\theta = h$ and consider the event $a_n \to h$ 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:[3]

(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 positive probability to form 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 the appendix.

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

[3]We omit some technical details in the statements of these two conditions. A complete formal treatment is presented in the appendix.

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

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

It starts at the initial level $\tilde{r}_1 = \log \frac{\pi}{1-\pi}$. When $(\widetilde{F}_\ell, \widetilde{F}_h)$ is tail-regular with exponent $\tilde{\alpha}$, we can for small q approximate $\widetilde{F}_\ell(q)$ with $q^{\tilde{\alpha}}$ and $\widetilde{F}_h(q)$ 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

$$\tilde{r}_{n+1} \approx \tilde{r}_n + \mathrm{e}^{-\tilde{\alpha}\tilde{r}_n}\,.$$

Intuitively, after each observed high action the misspecified public log-likelihood increases by an amount $\mathrm{e}^{-\tilde{\alpha}\tilde{r}_n}$ that become smaller as $\tilde{r}_n$ increases. More importantly, $\mathrm{e}^{-\tilde{\alpha}\tilde{r}_n}$ 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 $\frac{\mathrm{d}\tilde{r}(t)}{\mathrm{d}t} = \mathrm{e}^{-\tilde{\alpha}\tilde{r}(t)}$ whose solution is $\tilde{r}(t) = \tilde{\alpha}^{-1}\log(1+\tilde{\alpha}t)$. 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

$$\tilde{\pi}_n^h \approx 1-n^{-1/\tilde{\alpha}}. \tag{4.3}$$

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 $q_n \geqslant 1-\tilde{\pi}_n$: The agent takes the high action if her private posterior $q_n = \mathsf{P}(\theta = h|s_n)$ exceeds $1-\tilde{\pi}_n$. Hence the event $\bar{a} = h$ is the event that $q_n \geqslant 1-\tilde{\pi}_n^h$ 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 (4.3) 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

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

Crucially, we are only interested in whether this probability is positive or zero. As we show formally in the appendix, 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 $\mathsf{P}_h(\bar{a} = h) > 0$ 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 $\mathsf{P}_h(\bar{a} = \ell) = 0$ if and only if $\tilde{\alpha} - \alpha \geqslant 0$, i.e., immediate bad herds are excluded when agents are condescending.

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 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 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.

We note that this last step is obtained in the well-specified setting of Rosenberg and Vieille (2019) by appealing to 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 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 D.9). This holds even when agents are over-condescending (i.e., $\tilde{\alpha} \geqslant \alpha+1$), 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 positive probability of $\tau = \infty$ (see Proposition D.3), and in particular τ has 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.

References

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*.

Mira Frick, Ryota Iijima, and Yuhta Ishii (2020). "Misinterpreting Others and the Fragility of Social Learning," *Econometrica*, 88(6): 2281–2328.

Drew Fudenberg and Giacomo Lanzani (2022). "Which Misperceptions Persist?" *Theoretical Economics*.

Tristan Gagnon-Bartsch (2016). "Taste Projection in Models of Social Learning," *Working Paper*.

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*.

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

$$G_h(q) = 2\left(qG(q) - \int_0^q G(x)\,\mathrm{d}x\right),$$
$$G_\ell(q) = 2\left((1-q)G(q) + \int_0^q G(x)\,\mathrm{d}x\right).$$

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

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 A.2. *Suppose $G(q) = \Theta(q^\alpha)$. Then $G_\ell(q) = \Theta(q^\alpha)$ and $G_h(q) = \Theta(q^{\alpha+1})$.*

Proof. Lemma A.1 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 $C \geqslant c > 0$ such that for all $q \in [0,1]$, one has: $cq^\alpha \leqslant G(q) \leqslant Cq^\alpha$. The previous lemma thus implies that $G_h(q) \leqslant 2Cq^{\alpha+1}$. Next, let us define $m := \left(c/2C\right)^{1/\alpha}$, and observe that

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

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

$$\int_0^q G(x)\mathrm{d}x = \int_0^{mq} G(x)\mathrm{d}x + \int_{mq}^{q} G(x)\mathrm{d}x$$
$$\leqslant \frac{m}{q} qG(q) + (1-m)qG(q) = \left(1-\frac{m}{2}\right) qG(q)\,.$$

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

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 π_n^h. We accordingly denote $\tilde{r}_n^h = \log \frac{\pi_n^h}{1-\pi_n^h}$.

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

Proof. The perceived distributions $\widetilde{F}_h$ and $\widetilde{F}_\ell$ satisfy the iterated likelihood principle, that is,

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

This relation implies that $\widetilde{F}_h - \widetilde{F}_\ell$ 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

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

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

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

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

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

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

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):

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

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

where

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

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

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

Lemma B.1 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 $\widetilde{U}(r)$ for large r. The next lemma provides the asymptotic behavior of these functions.

Lemma B.2. *For large r, one has $U(r) = \Theta(\mathrm{e}^{-\alpha r})$ and $\widetilde{U}(r) = \Theta(\mathrm{e}^{-\tilde{\alpha} r})$, that is*

$$0 < \liminf_{r\to\infty} \frac{U(r)}{\mathrm{e}^{-\alpha r}} \leqslant \limsup_{r\to\infty} \frac{U(r)}{\mathrm{e}^{-\alpha r}} < \infty\,,$$

and

$$0 < \liminf_{r\to\infty} \frac{\widetilde{U}(r)}{\mathrm{e}^{-\tilde{\alpha} r}} \leqslant \limsup_{r\to\infty} \frac{\widetilde{U}(r)}{\mathrm{e}^{-\tilde{\alpha} r}} < \infty\,.$$

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

$$U(r) = \log \frac{1-F_h(\mu)}{1-F_\ell(\mu)} \leqslant -\log\left(1-F_\ell(\mu)\right)\,.$$

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

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

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 $\lim_{x\to 0} \frac{f(x)}{g(x)} = 0$, and $f(x) = O(g(x))$ if $\limsup_{x\to 0} \frac{f(x)}{g(x)} < \infty$.

To propose a lower bound, observe that because of Lemma A.1, $F_h(q) \leqslant 2qF(q)$ and $F_\ell(q) \geqslant 2(1-q)F(q)$, therefore,

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

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

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

The above upper and lower bounds imply that $\lim_{r\to\infty} \frac{U(r)}{2F(\mu)} = 1$. In addition, because of tail regularity (Assumption 2.3), it holds that $F(\mu) = \Theta(e^{-\alpha r})$, thereby justifying the lemma's first claim. A similar argument implies that $\widetilde{U}(r) = \Theta(e^{-\tilde{\alpha} r})$. □

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 $n \to \infty$. 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 B.3. *For every $\bar{r} \geqslant 0$, there exists n_0 such that $\tilde{r}_n^h(\pi) \geqslant \bar{r}$ for all $n \geqslant n_0$ and importantly for all initial $\pi \geqslant 1/2$.*

Proof. The idea is similar to the proof of Lemma 12 in Rosenberg and Vieille (2019). Let us introduce the mapping $\Psi(r) := r+\widetilde{U}(r)$, 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 $\widetilde{U} > 0$, if $\tilde{r}_1 \geqslant \bar{r}$, then $\Psi^n(\tilde{r}_1) \geqslant \bar{r}$. 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 $\tilde{r}_n^h(\pi^{(n)}) \leqslant \bar{r}$. Also, one has $\Psi^{m-1}\big(\tilde{r}_1(\pi^{(n)})\big) = \tilde{r}_m^h(\pi^{(n)}) \leqslant \bar{r}$ for all $m \leqslant n$. Since the interval $[0, \bar{r}]$ is compact, there is a subsequence of initial values $\{\tilde{r}_1(\pi^{(n)})\}$, 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

$$\Psi^n(r^*) = \lim_{k\to\infty} \Psi^n\big(\tilde{r}_1(\pi^{(k)})\big) \leqslant \bar{r}\,.$$

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 B.1). □

In the next two lemmas we calculate the asymptotic behavior of $\tilde{r}_n^h$ and show that $\mathrm{e}^{\tilde{r}_n^h} = \Theta(n^{1/\tilde{\alpha}})$. 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 B.4 (Lower Envelope)**.** *The misspecified public log-likelihood satisfies*

$$\liminf_{n\to\infty} \frac{\mathrm{e}^{\tilde{r}_n^h}}{n^{1/\tilde{\alpha}}} > 0\,. \tag{B.3}$$

Proof. By Lemma B.2, there exists $c > 0$ such that for all sufficiently large n (say $n \geqslant \bar{n}$), one has $\widetilde{U}(\tilde{r}_n^h) \geqslant c\,\mathrm{e}^{-\tilde{\alpha}\tilde{r}_n^h}$. Additionally, observe that the mapping $z \mapsto z + c\,\mathrm{e}^{-\tilde{\alpha}z}$ is increasing for all sufficiently large z (say $z \geqslant \bar{z}$). Since $\tilde{r}_n^h \to \infty$, one can choose $N \geqslant \bar{n}$ such that $\tilde{r}_N \geqslant \bar{z}$. For all $n \geqslant N$ it holds that

$$\tilde{r}_{n+1}^h - \tilde{r}_n^h = \widetilde{U}(\tilde{r}_n^h) \geqslant c\,\mathrm{e}^{-\tilde{\alpha}\tilde{r}_n^h}\,. \tag{B.4}$$

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

$$\frac{\mathrm{d}z(t)}{\mathrm{d}t} = c\,\mathrm{e}^{-\tilde{\alpha}z(t)}\,.$$

This equation has the solution form $z(t) = \tilde{\alpha}^{-1}\log\big(\kappa + c\tilde{\alpha}t\big)$, 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 $\tilde{r}_n^h \geqslant z(n)$ for all $n \geqslant N$, 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., $\tilde{r}_n^h \geqslant z(n)$. Then, observe that because of the mean value theorem, there exists $t \in [n, n+1]$ such that

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

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

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

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

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 B.5 (Upper Envelope)**.** *The misspecified public log-likelihood satisfies*

$$\limsup_{n\to\infty} \frac{\mathrm{e}^{\tilde{r}_n}}{n^{1/\tilde{\alpha}}} < \infty\,. \tag{B.5}$$

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 $n \geqslant \bar{n}$ one has $\widetilde{U}(\tilde{r}_n^h) \leqslant C\mathrm{e}^{-\tilde{\alpha}\tilde{r}_n^h}$. Likewise, the mapping $z \mapsto z+C\mathrm{e}^{-\tilde{\alpha}z}$ is increasing for all $z \geqslant \bar{z}$, and we can choose $N \geqslant \bar{n}$ such that $\tilde{r}_N^h \geqslant \bar{z}$. Then for all $n \geqslant N$ it holds that

$$\tilde{r}_{n+1}^h-\tilde{r}_n^h = \widetilde{U}(\tilde{r}_n^h) \leqslant C\mathrm{e}^{-\tilde{\alpha}\tilde{r}_n^h}\,. \tag{B.6}$$

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

$$\frac{\mathrm{d}z(t)}{\mathrm{d}t} = 2C\mathrm{e}^{-\tilde{\alpha}z(t)}\,,$$

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

$$2\mathrm{e}^{-\tilde{\alpha}z(n+1)} \geqslant \mathrm{e}^{-\tilde{\alpha}z(n)}. \tag{B.7}$$

Therefore, one can choose κ large enough such that $z(N) \geqslant \tilde{r}_N^h$. Next, we inductively show $\tilde{r}_n^h \leqslant z(n)$ for all $n \geqslant N$, 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

$$z(n+1)-z(n) = 2C\mathrm{e}^{-\tilde{\alpha}z(t)} \geqslant 2C\mathrm{e}^{-\tilde{\alpha}z(n+1)} \geqslant C\mathrm{e}^{-\tilde{\alpha}z(n)}\,,$$

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

$$z(n+1) \geqslant z(n)+C\mathrm{e}^{-\tilde{\alpha}z(n)} \geqslant \tilde{r}_n^h+C\mathrm{e}^{-\tilde{\alpha}\tilde{r}_n^h} \geqslant \tilde{r}_{n+1}^h\,,$$

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

The previous two lemmas jointly imply that $\mathrm{e}^{\tilde{r}_n^h} = \Theta(n^{1/\tilde{\alpha}})$. 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.

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 C.1 (Necessary condition). *Assume $a_n \to h$, P_h-almost surely. Then, the following two conditions hold:*

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

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

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

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

Applying the law of total expectations, this is equal to

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

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

Since the event $a_n \to h$ 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). □

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 C.2 (Sufficient condition)**.** *The following two conditions are sufficient for $a_n \to h$, P_h-almost surely:*

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

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

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

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

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

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

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

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. □

D Proof of Theorem 3.1

We divide the proof of Theorem 3.1 into two: Proposition D.1 characterizes asymptotic learning, and Proposition D.8 characterizes efficient learning. Jointly, they imply the theorem.

D.1 Parametric Characterization for Asymptotic Learning

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

Proposition D.1. *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 C.1 and C.2, 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 D.2. *The following are equivalent:*

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

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

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

Proposition D.3. *The following are equivalent:*

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

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

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

Proposition D.2 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 D.3 claims that the probability of an immediate wrong herd is zero for all positive initial beliefs if and only if $\tilde{\alpha} - \alpha \geqslant 0$, i.e., when agents are condescending.

We use these propositions to prove Proposition D.1, before proceeding with their proofs.

Proof of Proposition D.1. Suppose that asymptotic learning holds, i.e., $a_n \to \theta$, P-almost surely. Then $a_n \to h$, P_h-almost surely. By Lemma C.1, this implies that condition (ii) of both Propositions D.2 and D.3 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 D.2 and D.3 hold. Therefore, condition (i) of both propositions hold. Hence, by Lemma C.2, we have asymptotic learning. $\square$

D.1.1 Proof of Proposition D.2

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

Lemma D.4. *The following are equivalent:*

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

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

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 $\{q_n + \tilde{\pi}_n^h \geqslant 1, \forall n\}$. Since the random variables q_n are independent conditioned on $\theta = h$, we have that

$$\mathsf{P}_h(\bar{a} = h) = \prod_n \mathsf{P}_h(q_n + \tilde{\pi}_n^h \geqslant 1) = \prod_n \left(1 - F_h(1-\tilde{\pi}_n^h)\right) . \tag{D.1}$$

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

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

Observe that Lemma A.2 implies that $F_h(q) = \Theta(q^{\alpha+1})$. Also as $n \to \infty$, we have $\mathrm{e}^{-\tilde{r}_n^h} \sim 1-\tilde{\pi}_n^h$, therefore, the above sum is finite if and only if

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

Because of the Lemmas B.4 and B.5, one has $\mathrm{e}^{-(\alpha+1)\tilde{r}_n^h} = \Theta\left(n^{-\frac{\alpha+1}{\tilde{\alpha}}}\right)$. 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 D.2, that is (iii) ⇒ (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 $\tilde{r}^h_n(r)$ to refer to the process initiated at $\tilde{r}_1 = r$. Also, recall our former notation, where we used $\tilde{r}^h_n(\pi)$ 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 D.5. *Assume $\tilde{\alpha}-\alpha < 1$. Then for every $\varepsilon > 0$, there exists $\bar{r} \geqslant 0$ such that for all $r \geqslant \bar{r}$ one has $\sum_{n\geqslant 0} \mathrm{e}^{-(\alpha+1)\tilde{r}^h_n(r)} < \varepsilon$.*

Proof. We appeal to the idea used in the proof of Lemma B.4. Since $\widetilde{U}(r) = \Theta(\mathrm{e}^{-\tilde{\alpha}r})$, then there exists $c > 0$, and correspondingly a threshold $\bar{r}$, such that $\widetilde{U}(r) \geqslant c\,\mathrm{e}^{-\tilde{\alpha}r}$ for every $r \geqslant \bar{r}$ and the mapping $r \mapsto r + c\,\mathrm{e}^{-\tilde{\alpha}r}$ is increasing on $[\bar{r}, \infty)$. In particular, since $\tilde{r}^h_n$ is increasing in n, starting the process at any $\tilde{r}_1 = r \geqslant \bar{r}$, implies

$$\widetilde{U}(\tilde{r}^h_n(r)) \geqslant c\,\mathrm{e}^{-\tilde{\alpha}\tilde{r}^h_n(r)}.$$

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

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

The solution to this differential equation takes the form

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

Using induction, similar to the one used in Lemma B.4, we can show $\tilde{r}^h_n(r) \geqslant z(n)$ for every initial value $r \geqslant \bar{r}$. Therefore, for every $r \geqslant \bar{r}$, it holds that

$$\sum_{n\geqslant 0} \mathrm{e}^{-(\alpha+1)\tilde{r}^h_n(r)} \leqslant \sum_{n\geqslant 0} \mathrm{e}^{-(\alpha+1)z(n)} = \sum_{n\geqslant 0} \left(\mathrm{e}^{\tilde{\alpha}\bar{r}} + c\tilde{\alpha}n\right)^{-\frac{\alpha+1}{\tilde{\alpha}}}.$$

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. □

Let $\eta\colon [1/2, 1] \to [0, 1]$,

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

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

$$\eta(\pi) = \prod_n \left(1 - F_h(1 - \tilde{\pi}^h_n(\pi))\right).$$

Lemma D.6. *Assume $\tilde{\alpha}-\alpha < 1$. Then η is continuous.*

Proof. Let

$$\eta_n(\pi) = \mathsf{P}_{\pi,h}(a_1 = h, \ldots, a_n = h) = \prod_{k=1}^{n} \left(1-F_h(1-\tilde{\pi}_k^h(\pi))\right), \tag{D.3}$$

be the probability that the first n agents take the high action, conditioned on the high state, when the prior is π. By definition, $\eta(\pi) = \lim_n \eta(\pi)$. 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 A.2 implies that $F_h(\mathrm{e}^{-r}) = \Theta(\mathrm{e}^{-(\alpha+1)r})$, and hence there exists $C > 0$ such that $F_h(\mathrm{e}^{-r}) \leqslant C\mathrm{e}^{-(\alpha+1)r}$. Second, because of Lemma D.5, for a given $\varepsilon_1 > 0$, there exists $\bar{r} \geqslant 0$ such that for all $r_1 \geqslant \bar{r}$, one has

$$\sum_{n\geqslant 0} \mathrm{e}^{-(\alpha+1)\tilde{r}_n^h(r_1)} \leqslant \varepsilon_1 .$$

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

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

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

$$|\eta_{n_0+k}(\pi)-\eta_{n_0}(\pi)| \leqslant \sum_{n=n_0}^{\infty} F_h\left(\mathrm{e}^{-\tilde{r}_n^h(\pi)}\right) \leqslant C \sum_{n\geqslant n_0} \mathrm{e}^{-(\alpha+1)\tilde{r}_n^h(\pi)} \leqslant C\varepsilon_1 . \tag{D.4}$$

The third inequality above holds because $\tilde{r}_{n_0}^h(\pi) \geqslant \bar{r}$, and thus Lemma D.5 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. □

Using the above lemma, we can now conclude the proof of the last implication in Proposition D.2, namely (iii) ⇒ (i). Assume by contradiction, that condition (i) does not hold, then $\inf_{\pi\geqslant 1/2} \eta(\pi) = 0$. 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 D.4 implies that $\tilde{\alpha}-\alpha \geqslant 1$. This violates the initial assumption (i.e., $\tilde{\alpha}-\alpha < 1$) and hence concludes the proof of Proposition D.2.

D.1.2 Proof of Proposition D.3

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

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

and appeal to the next lemma.

Lemma D.7. *For every $\pi \in (0,1)$, one has $\xi(\pi) = 0$ if and only if $\tilde{\alpha} - \alpha \geqslant 0$.*

Proof. Because of the symmetry assumption, we have $\xi(\pi) = \mathsf{P}_{1-\pi,\ell}(\bar{a} = h)$. 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 D.4, one has

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

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

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

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

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

It was shown in Lemmas B.4 and B.5 that $\mathrm{e}^{\tilde{r}_n} = \Theta(n^{1/\tilde{\alpha}})$, 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 $\tilde{\alpha} - \alpha \geqslant 0$. □

Observe that $\xi(1) = 0$. Therefore, the second and the third implications of Proposition D.3, namely (ii) $\Rightarrow$ (iii) $\Rightarrow$ (i), respectively follow from the above lemma, thereby concluding the proof of Proposition D.3.

D.2 Parametric Characterization for Efficient Learning

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

Proposition D.8. *Assume $\tilde{\alpha} \neq \alpha$. The following are equivalent:*

(i) Efficient learning;

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

Proof. The implication (i) ⇒ (ii) follows immediately from Proposition D.1, since efficient learning implies asymptotic learning.

Towards sufficiency, assume $\tilde{\alpha} - \alpha < 1$. Following the same logic as in the proof of Lemma C.2, 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 D.9 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 D.9. *Assume $\tilde{\alpha} - \alpha > 0$, then the following statements hold:*

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

$$\mathsf{E}_{\pi,h}\left[\tau_h\right] \leqslant C_0 \, \frac{1-\pi}{\pi} \, .$$

(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 $\frac{1-\tilde{\pi}_{n+1}}{\tilde{\pi}_{n+1}} \leqslant B$ for some constant $B < \infty$ that does not depend on the history.

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

$$\mathsf{E}_{1-\pi,\ell}\left[\tau_\ell\right] = 1 + \sum_{n \geqslant 1} \mathsf{P}_{1-\pi,\ell}\left(\tau_\ell > n\right) \, .$$

By Bayes Law

$$\mathsf{P}_{1-\pi,\ell}\left(\tau_\ell > n\right) = \frac{1-\pi}{\pi} \times \frac{1-\pi^h_{n+1}}{\pi^h_{n+1}}\, \mathsf{P}_{1-\pi,h}\left(a_1 = \cdots = a_n = h\right)\,,$$

where π^h_n 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^h_n represents the correctly specified public log-likelihood on the high action path, which follows the dynamics in equation (B.1), namely $r^h_{n+1} - r^h_n = U(\tilde{r}^h_n)$. Hence, the above expression implies that

$$\mathsf{E}_{1-\pi,\ell}\left[\tau_\ell\right] \leqslant 1 + \frac{1-\pi}{\pi} \sum_{n\geqslant 1} \mathrm{e}^{-r^h_{n+1}}\,. \tag{D.5}$$

Since the initial belief is set to $1-\pi$ and it is assumed in part (i) that $\pi \leqslant 1/2$, then $r^h_n \geqslant 0$. Next, observe that on the path $\bar{a} = h$, the misspecified public log-likelihood follows the difference equation (B.2), namely $\tilde{r}^h_{n+1} - \tilde{r}^h_n = \widetilde{U}(\tilde{r}^h_n)$. Additionally, because of Lemma B.2, there exists $C > 0$ such that $\widetilde{U}(\tilde{r}^h_n) \leqslant C\mathrm{e}^{-\tilde{\alpha}\tilde{r}^h_n}$. We continue by finding a continuous time upper envelope for $\tilde{r}^h_n$—analogous to Lemma B.5 with a slight catch in selecting the initial condition. Choose $\bar{r} > 0$ such that the mapping $r \mapsto r + C\mathrm{e}^{-\tilde{\alpha} r}$ becomes increasing on $[\bar{r}, \infty)$. Let $n_0 := \min\{n : \tilde{r}^h_n \geqslant \bar{r}\}$ that is finite because $\tilde{r}_n \to \infty$ on the high action path. Since $\widetilde{U}(\cdot)$ is a decreasing function, then $\bar{r} \leqslant \tilde{r}^h_{n_0} \leqslant \bar{r} + \widetilde{U}(0)$. Let $z(t)$ be the solution to the following differential equation

$$\frac{\mathrm{d}z(t)}{\mathrm{d}t} = 2C\mathrm{e}^{-\tilde{\alpha} z(t)}\,,$$

starting at $z(0) = \bar{r} + \widetilde{U}(0)$. Therefore, $z(t) = \tilde{\alpha}^{-1}\log\left(\mathrm{e}^{\bar{r}+\widetilde{U}(0)} + 2C\tilde{\alpha} t\right)$. Following the recipe of Lemma B.5, one can show by induction that $z(k) \geqslant \tilde{r}^h_{k_0+n}$ for all $k \geqslant 0$. Next, we examine Raabe's criterion[4] for the infinite sum $\sum_{n\geqslant 1} \mathrm{e}^{-r^h_n}$, that is to examine the limit of the following expression:

$$n\left(\frac{\mathrm{e}^{-r^h_n}}{\mathrm{e}^{-r^h_{n+1}}} - 1\right) = n\left(\mathrm{e}^{U(\tilde{r}^h_n)} - 1\right) \geqslant nU(\tilde{r}^h_n)\,.$$

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

$$\liminf_{n\to\infty} n\left(\frac{\mathrm{e}^{-r^h_n}}{\mathrm{e}^{-r^h_{n+1}}} - 1\right) \geqslant \limsup_{k\to\infty} ck\,\mathrm{e}^{-\alpha z(k)} = \limsup_{k\to\infty} \frac{c\,k}{\left(\mathrm{e}^{\bar{r}+\widetilde{U}(0)} + 2C\tilde{\alpha}\,k\right)^{\alpha/\tilde{\alpha}}}\,.$$

[4] 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 $\rho_n = n\big(\nu_n/\nu_{n+1} - 1\big)$. It is inconclusive when $\lim_n \rho_n = 1$. This latter case corresponds to $\tilde{\alpha} = \alpha$.

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

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

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

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

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

This establishes (ii). □

E Proofs of Theorems 3.2 and 3.3

Proof of Theorem 3.2. Suppose that $\tilde{\alpha} < \alpha$. Then, by Proposition D.3 one has $\mathsf{P}_h(\bar{a} = \ell) > 0$, so that a wrong herd forms immediately with positive probability. □

Proof of Theorem 3.3. Suppose that $\tilde{\alpha} \geqslant \alpha+1$. Condition on the high state. Then, by Proposition D.2, for any prior $\pi' < 1$, the probability of an immediate herd on the high action is zero. Hence, by stationarity, the probability that $a_n \to h$ is zero. By Proposition D.3, the probability of an immediate herd on the low action is also zero, and hence, again by stationarity, the probability that $a_n \to \ell$ is zero. Thus the agents take both actions infinitely many times. The same argument applies when conditioning on the low state. □