THE SPEED OF SEQUENTIAL ASYMPTOTIC LEARNING

WADE HANN-CARUTHERS, VADIM V. MARTYNOV AND OMER TAMUZ

ABSTRACT. In the classical herding literature, agents receive a private signal regarding a binary state of nature, and sequentially choose an action, after observing the actions of their predecessors. When the informativeness of private signals is unbounded, it is known that agents converge to the correct action and correct belief. We study how quickly agents converge to the correct belief, and show that this happens more slowly than it does when agents observe signals. However, we also show that the speed of learning from actions can be arbitrarily close to the speed of learning from signals. In the canonical case of Gaussian private signals we calculate the speed of convergence precisely, and show explicitly that, in this case, learning from actions is significantly slower than learning from signals.

INTRODUCTION

When making decisions, we often rely on the decisions that others before us have made. Sequential learning models have been used to understand different phenomena that occur when many individuals rely on the observed actions of others to help them make a decision: these include herd behavior (cf. [2]), where many agents make the same choice, as well as informational cascades (e.g. [3]), where the actions of the first few agents provide such compelling evidence that later agents no longer have incentive to consider their own private information.

Such results on how information aggregation can fail are complemented by results which demonstrate that when private signals are arbitrarily strong, learning is robust to this kind of collapse [5]. In particular, in a process called asymptotic learning (see, e.g., [1]), agents will eventually choose the correct action and their beliefs will converge to the truth. A question that has not been answered in the literature

The authors would like to thank Christophe Chamley, Philipp Strack, Ye Wang and Ivo Welch for helpful comments and discussions.

is: how quickly does this happen? And how does the speed of learning compare to a setting in which agents observe signals rather than actions?[1]

We consider the classical setting of a binary state of nature and binary actions, where each of the two actions is optimal at one of the states. The agents receive private signals that are independent conditioned on the state. These signals are unbounded, in the sense that an agent's posterior belief regarding the state can be arbitrarily close to both 0 and 1. The agents are exogenously ordered, and, at each time period, a single agent takes an action, after observing the actions of her predecessors.

We measure the rate of learning by studying how the public belief evolves as more and more agents act. Consider an outside observer who observes the actions of the sequence of agents. The public belief is the posterior belief that such an outside observer assigns to the correct state of nature. It provides a measure of how well the population has learned the state. Since signals are unbounded, the public belief tends to 1 as the time t tends to infinity [5]; equivalently, the corresponding log-likelihood ratio tends to infinity.

When agents observe the *signals* (rather than actions) of all of their predecessors, this log-likelihood ratio is asymptotically linear. Thus, it cannot grow faster than linearly when the agents observe actions. Our first main finding is that when observing actions, the log-likelihood ratio always grows sub-linearly. Equivalently, the public belief converges sub-exponentially to 1. Our second main finding is that, depending on the choice of private signal distributions, the log-likelihood ratio can grow at a rate that is arbitrarily close to linear.

Finally, we analyze the specific canonical case of Gaussian private signals. Here we calculate precisely the asymptotic behavior of the log-likelihood ratio. We show that learning from actions is significantly slower than learning from signals: the log-likelihood ratio behaves asymptotically as $\sqrt{\ln t}$.

In the Gaussian case, we also consider a second measure of the speed of learning. Namely, we consider directly how the probability of choosing the incorrect action varies as an agent sees more and more of the other agents' decisions before making her own. This probability tends to 0 as time tends to infinity [5]. We find that this probability is asymptotically no less than $1/t^{1-\varepsilon}$ for any $\varepsilon > 0$. In contrast, when agents

[1]Chamley [4] gives an estimate for a particular class of private signal distributions with fat tails, and Sørensen [6] has published a related claim with an unfinished proof.

can observe the private signals of their predecessors, the probability of mistake decays exponentially, and so also in this sense learning from signals is much faster than learning from actions.

1. Model

Let $\theta \in \{-1, +1\}$ be the true state of the world, with each state a priori equally likely[2]. Each rational agent $t \in \{1, 2, \dots\}$ receives a private signal s_t. The signals are i.i.d. conditioned on θ: if $\theta = +1$ they have CDF F_+ and if $\theta = -1$ they have CDF F_-.[3] We assume that F_+ and F_- are absolutely continuous with respect to each other, so that private signals never completely reveal the state.

Let

$$L_t = \ln \frac{\mathbb{P}(\theta = +1 | s_t)}{\mathbb{P}(\theta = -1 | s_t)}$$

be the log-likelihood ratio of the belief induced by the agent's private signal. We assume that private signals are unbounded, in the sense that L_t is unbounded: for every $M \in \mathbb{R}$ the probability that $L_t > M$ is positive, as is the probability that $L_t < -M$. We denote by G_+ and G_- the conditional CDFs of L_t.

The agents act sequentially, with agent t acting after observing the actions of agents $\{1, \dots, t-1\}$. The utility of the action $a_t \in \{-1, +1\}$ is 1 if $a_t = \theta$ and 0 otherwise.

Denote the public belief by

$$\mu_t = \mathbb{P}(\theta = 1 | a_1, \dots, a_{t-1}).$$

This is the posterior held by an outside observer after recording the actions of the first $t-1$ agents. We denote by ℓ_t the log-likelihood ratio of the public belief:

$$\ell_t = \ln \frac{\mu_t}{1 - \mu_t}.$$

In equilibrium, agent t chooses $a_t = 1$ iff[4]

$$\ln \frac{\mathbb{P}(\theta = +1 | a_1, \dots, a_{t-1}, s_t)}{\mathbb{P}(\theta = -1 | a_1, \dots, a_{t-1}, s_t)} > 0.$$

[2]We make this simplification of a (1/2,1/2) prior to reduce the complexity of the presentation, but all results hold for general priors.

[3]One could consider signals that take values in a general measurable space (rather than $\mathbb{R}$), but the choice of $\mathbb{R}$ is in fact without loss of generality, since all standard measurable spaces are isomorphic.

[4]For simplicity, we assume that agents choose action 0 when indifferent. This will have no impact on our results.

A simple calculation shows that this occurs iff

$$\ell_t + L_t > 0.$$

Now, another straightforward calculation shows that when $a_t = +1$,

$$\ell_{t+1} = \ell_t + D_+(\ell_t),$$

where

$$D_+(x) = \ln \frac{1 - G_+(-x)}{1 - G_-(-x)}.$$

Likewise, when $a_t = -1$,

$$\ell_{t+1} = \ell_t + D_-(\ell_t),$$

where

$$D_-(x) = \ln \frac{G_+(-x)}{G_-(-x)}.$$

We can interpret $D_+(\ell_t)$ and $D_-(\ell_t)$ as the contributions of agent t's action to the public belief.

2. The speed of learning

Consider a baseline model, in which each agent observes the private signals of all of her predecessors. In this case the public log-likelihood ratio $\ell_t^\star$ would equal the sum

$$\ell_t^\star = \sum_{\tau=1}^{t} L_\tau.$$

Conditioned on the state this is the sum of i.i.d. random variables, and so by the law of large numbers we have that the limit $\lim_t \ell_t^\star/t$ would - conditioned on (say) $\theta = +1$ - equal the conditional expectation of L_t, which is positive.[5]

Our first main result shows that when agents observe actions rather than signals, the public log-likelihood ratio grows sub-linearly, and so learning from actions is always slower than learning from signals.

Theorem 1. *It holds with probability 1 that* $\lim_t \ell_t/t = 0$.

Our second main result shows that, depending on the choice of private signal distributions, ℓ_t can grow at a rate that is arbitrarily close to linear: given any sub-linear function r_t, it is possible to find private signal distributions so that ℓ_t grows as fast as r_t.

[5]In fact, $\mathbb{E}(L_t|\theta = +1)$ is equal to the Kullback-Leibler divergence between F_+ and F_-, which is positive as long as the two distributions are different.

Theorem 2. *For any* $r\colon \mathbb{N} \to \mathbb{R}$ *such that* $\lim_t r_t/t = 0$ *there exists a choice of CDFs* F_- *and* F_+ *such that*

$$\liminf_{t\to\infty} \frac{|\ell_t|}{r_t} > 0$$

with probability 1.

For example, for some choice of private signal distributions, ℓ_t grows asymptotically at least as fast as $t/\ln t$, which is sub-linear but (perhaps) close to linear.

3. Gaussian Private Signals

In this section we consider the case that F_+ is Normal with mean $+1$ and variance σ^2, and F_- is Normal with mean -1 and the same variance.

3.1. The Evolution of the Public Belief.

In this Gaussian case we calculate how the public log-likelihood ratio ℓ_t behaves as t tends to infinity. Since signals are unbounded, agents learn the state, so that ℓ_t tends to $+\infty$ if $\theta = +1$, and to $-\infty$ if $\theta = -1$. In particular ℓ_t stops changing sign from some t on, with probability one; all later agents choose the correct action.

By the symmetry of the model we can consider without loss of generality the case $\ell_1, \ell_2, \ldots, \ell_t$ are all positive from some t on. We first prove the following lemma.

Lemma 3. *If* ℓ_t *is positive and* $a_t = +1$ *then*

$$\ell_{t+1} = \ell_t + \frac{e^{-\sigma^2\ell_t^2/8+O(\ell_t)}}{\ell_t}. \tag{1}$$

Here $O(\ell_t)$ denotes some lower order linear term, which we do not calculate precisely. Lemma 7 in the Appendix is a rephrasing of this statement.

To get some informal intuition for how fast ℓ_t converges to $+\infty$ when $\theta = +1$, we neglect the lower order term and consider the corresponding differential equation:

$$\frac{df}{dt}(t) = \frac{e^{-f^2(t)/4}}{f(t)}$$

The family of solutions for this equation is

$$f(t) = 2\sqrt{\ln(\frac{1}{2}(\text{const} + t))}.$$

We show that the solution of (1) has the same asymptotic behavior. We note that the analysis of the discrete time dynamics poses some significant technical challenges.

Recall, that when private signals are observed, the public log-likelihood ratio is asymptotically *linear*. Thus, learning from actions is far slower than learning from signals in the Gaussian case.

Formally, as the following theorem shows, ℓ_t indeed behaves like $\sqrt{\ln t}$ as $t \to \infty$.

Theorem 4. *There exist $k_1, k_2 > 0$ such that, conditioned on $\theta = +1$, it almost surely holds that*

$$k_1\sqrt{\ln t} \geq \ell_t \geq k_2\sqrt{\ln t}$$

for all t large enough.

3.2. **Probability of taking the wrong action.** Still considering the Gaussian case, we would like to approximate $p_t = \mathbb{P}(a_t \neq \theta)$, the probability that agent t takes the wrong action, and in particular to show that it is large, as compared to a setting in agents observe signals.

To bound this probability from below we notice that the probability that agent t takes the wrong action is at least as great as the probability that the first $t-1$ agents take the right action and agent t takes the wrong action, since the latter event is contained in the former. This yields the following bound.

Theorem 5. *For any $\varepsilon > 0$, there exists $k > 0$ such that $\mathbb{P}(a_n = -1|\theta = 1) \geq ke^{-(1+\varepsilon)\ln t}$ for all sufficiently large t.*

So we get that $\mathbb{P}(a_t \neq \theta)$ goes to 0 no faster than $t^{-(1+\varepsilon)}$. For comparison, when each agent has access to all of her predecessors' private signals, the probability of making a mistake decays exponentially.

4. Conclusion

In this paper we consider a classical setting of sequential asymptotic learning from actions of others. We show that learning from actions is slow, as compared to the speed of learning when observing others' private signals. However, it is possible to approach the linear rate of learning from signals and achieve any sub-linear rate.

We calculate the speed of learning precisely in the case of Normal private signals, which already pose a significant technical challenge to analyze. Many of the techniques we use are tailored to this case, and do not obviously generalize to other signal structures. An interesting direction of research would be to understand more generally how the speed of learning depends on the private signal distribution.

Appendix A. Sub-linear learning: Proofs

Before proving our main theorems we make the observation (which has appeared before, e.g., [4]) that *the log-likelihood ratio of the log-likelihood ratio is the the log-likelihood ratio.* Formally, if ν_+ and ν_- are the conditional distributions of the private log-likelihood ratio L_t (i.e., have CDFs G_+ and G_-), then

$$\ln \frac{d\nu_+}{d\nu_-}(x) = x.$$

It follows that

$$G_+(x) = \int_{-\infty}^{x} d\nu_+(x) = \int_{-\infty}^{x} e^x \, d\nu_-(x). \tag{2}$$

Our first lemma shows that asymptotically, D_+ behaves like the left tail of G_-, and D_- behaves like the right tail of G_+.

Lemma 6. *For every $\epsilon > 0$ there exists an $x_0 > 0$ such that, for all $x > x_0$,*

$$(1-\epsilon)\cdot G_-(-x) < D_+(x) < (1+\epsilon)\cdot G_-(-x),$$

and likewise for all $x < -x_0$

$$(1-\epsilon)\cdot(1-G_+(-x)) < D_-(x) < (1+\epsilon)\cdot(1-G_+(-x)).$$

Proof. By definition,

$$D_+(x) = \ln \frac{1-G_-(-x)}{1-G_+(-x)}.$$

Since $\ln(1-z) = -z + O(z^2)$, it holds for all x large enough that

$$D_+(x) > G_-(-x) - 2\cdot G_+(-x).$$

By (2)

$$D_+(x) > \int_{-\infty}^{-x} (1-2e^z)\, d\nu_-(z),$$

and so for all x large enough,

$$D_+(x) > (1-\epsilon)\cdot \int_{-\infty}^{-x} d\nu_-(z) = (1-\epsilon)G_-(-x).$$

Using the same approximation of the logarithm, we have that

$$D_+(-x) < (1+\epsilon)G_-(-x) - G_+(-x) < (1+\epsilon)G_-(-x).$$

The corresponding bounds on D_- follow by identical arguments. $\square$

Proof of Theorem 1. Condition on $\theta = +1$. Then ℓ_t is with probability 1 positive from some point on, and all agents take action $+1$ from this point on. Hence, for all t large enough,

$$\ell_{t+1} = \ell_t + D_+(\ell_t).$$

By Lemma 6, we know that $\lim_x D_+(x) = 0$. Hence for every $\epsilon > 0$ and all t large enough, $|\ell_{t+1} - \ell_t| < \epsilon$. It follows that the limit $\lim_t \ell_t/t = 0$. The analysis of the case $\theta = -1$ is identical. □

Proof of Theorem 2. Given r_t, we will construct private signal distributions such that $\liminf_t |\ell_t|/r_t > 0$ with probability one. These distributions will furthermore have the property that $D_+(x) = -D_-(-x)$. As a consequence we have that regardless of the action chosen by the agent, as long as the sign of the action is equal to that of ℓ_t (which happens from some point on w.p. 1),

$$|\ell_{t+1}| = |\ell_t| + D_+(|\ell_t|).$$

Intuitively, if we can choose private signal distributions that make $D_+(x)$ decay very slowly, then ℓ_t will be very close to being linear.

Formally, and by elementary considerations, the theorem will follow if, for every $Q\colon \mathbb{R} \to \mathbb{R}^+$ with $\lim_{x\to\infty} Q(x) = 0$, we can find CDFs such that $D_+(x) = -D_-(-x)$ and $\liminf_{x\to\infty} D_+(x)/Q(x) > 0$.

Fix any Q such that $\lim_{x\to\infty} Q(x) = 0$, but assume without loss of generality that $Q(x)$ is monotone decreasing.[6] Define a finite measure ν on the integers by

$$\nu(n) = \frac{Q(n-1) - Q(n)}{e^n}$$

and

$$\nu(-n) = Q(n-1) - Q(n)$$

for all $n \geq 0$. Note that ν is indeed finite since

$$C := \sum_{n=-\infty}^{\infty} \nu(n) \leq 2Q(-1).$$

Note also that

$$\sum_{n=-\infty}^{\infty} \nu(n) \cdot e^n$$

is likewise equal to C.

[6]If Q is not monotone decreasing then consider instead $Q'(x) = \sup_{y \geq x} Q(y)$.

Let the private signal distributions be given by

$$\mathbb{P}(s_t = n|\theta = +1) = C^{-1}\nu(n)e^n$$

and

$$\mathbb{P}(s_t = n|\theta = -1) = C^{-1}\nu(n).$$

Then

$$L_t = \ln\frac{\mathbb{P}(s_t|\theta = +1)}{\mathbb{P}(s_t|\theta = -1)} = s_t,$$

the distribution of L_t is identical to that of s_t, and so $G_+ = F_+$ and $G_- = F_-$. By our definition of F_-, we have that for $x > 0$

$$G_-(-x) = C^{-1} \cdot Q(\lceil x\rceil - 1). \tag{3}$$

Now, by Lemma 6, we know that

$$(1-\epsilon)\cdot G_-(-x) < D_+(x) < (1+\epsilon)\cdot G_-(-x),$$

for any $\epsilon > 0$ and all x large enough. It follows that

$$\liminf_{x\to\infty}\frac{D_+(x)}{Q(x)} = \liminf_{x\to\infty}\frac{G_-(-x)}{Q(x)},$$

which, by (3) equals

$$\liminf_{x\to\infty}\frac{C^{-1}Q(\lceil x\rceil - 1)}{Q(x)} \geq C^{-1}.$$

□

Appendix B. Gaussian Private Signals: Proofs

B.1. **Preliminaries.** We start with a few observations. First, recall that L_t is the log-likelihood ratio of the belief induced by agent t's private signal, and notice that

$$L_t = \ln\frac{e^{-(s_t-1)^2/2\sigma^2}}{e^{-(s_t-(-1))^2/2\sigma^2}} = 2s_t/\sigma^2.$$

Hence L_t is normally distributed, conditioned on the state θ. Also

$$F_+(x) := \mathbb{P}(s_t \leq x|\theta = +1) = 1 - \frac{1}{2}\mathrm{erfc}\left(\frac{x-1}{\sigma\sqrt{2}}\right) \tag{4}$$

and

$$F_-(x) := \mathbb{P}(s_t \leq x|\theta = -1) = 1 - \frac{1}{2}\mathrm{erfc}\left(\frac{x+1}{\sigma\sqrt{2}}\right). \tag{5}$$

B.2. **The Evolution of the Public Belief.** Recall that when agent t takes action $+1$, the public log-likelihood ratio ℓ_t will satisfy the recurrence equation $\ell_{t+1} = \ell_t + D_+(\ell_t)$. Conditioned on $\theta = +1$, all agents eventually choose the action $+1$, and so this relation holds for all t large enough. Thus to understand the long term behavior of ℓ_t we need to solve this recurrence relation.

We will use the following lemmas to show that, conditioned on $\theta = +1$, ℓ_t behaves asymptotically as $\sqrt{\ln t}$. By symmetry, $-\ell_t$ will have the same behavior conditioned on $\theta = -1$.

In the first lemma, we note that $D_+(x)$ is asymptotically well approximated by

$$b(x) := x^{-1}e^{-\sigma^2 x^2/8} \tag{6}$$

Lemma 7. *For any $\delta > 0$,*

$$\lim_{x\to\infty} \frac{b(x)}{D_+(x)} = 0.$$

and

$$\lim_{x\to\infty} \frac{D_+(x)}{b((1-\delta)x)} = 0$$

Proof. Both follow from routine applications of L'Hôpital's rule. □

We now prove a general, technical lemma that will allow us to asymptotically compare sequences defined by recurrence relations of the type satisfied by ℓ_t. This will be useful since it will allow us to use the solutions of simpler relations to bound the solution for ℓ_t.

Lemma 8. *Let $f(x)$, $g(x)$ be positive monotonically decreasing functions such that $f(x)$, $g(x)$ and $\frac{f(x)}{g(x)}$ tend to 0 as x tends to ∞. Let s_0, $r_0 > 0$, and define the sequences (s_t), (r_t) by $s_{t+1} = s_t + g(s_t)$ and $r_{t+1} = r_t + f(r_t)$. Finally, assume that $r_t \to \infty$. Then for any $\varepsilon > 0$ and for all sufficiently large t, $s_t > (1-\varepsilon)r_t$.*

Proof. The proof proceeds in two steps. First, we show that $s_t > r_t$ for infinitely many t. Then, we show that for all sufficiently large t, $s_t > (1-\varepsilon)r_t$.

Step 1: Suppose towards contradiction that for all sufficiently large t, $s_t \le r_t$. Define the sequence (v_t) by $v_t = s_t - r_t$. By hypothesis, there exists a T such that for $t \ge T$, $s_t \le r_t$ implies $g(s_t) \ge g(r_t) \ge 2f(r_t)$.

Let $T_0 \geq T$ such that $s_t \leq r_t$ for $t \geq T_0$. Then for $t \geq T_0$,

$$\begin{aligned} v_{t+1} - v_t &= (s_{t+1} - r_{t+1}) - (s_t - r_t) \\ &= (s_{t+1} - s_t) - (r_{t+1} - r_t) \\ &= g(s_t) - f(r_t) \\ &\geq f(r_t) = r_{t+1} - r_t. \end{aligned}$$

Thus, for $t > T_0$,

$$\begin{aligned} v_t &= (v_t - v_{t-1}) + \ldots + (v_{T_0+1} - v_{T_0}) + v_{T_0} \\ &\geq (r_t - r_{t-1}) + \ldots + (r_{T_0+1} - r_{T_0}) + v_{T_0} \\ &= r_t - r_{T_0} + v_{T_0}. \end{aligned}$$

But since $r_t \to \infty$, eventually $v_t > 0$, so eventually $s_t > r_t$, contradiction. Thus, for every t such that $s_t \leq r_t$, there is a $t' > t$ such that $s_{t'} > r_{t'}$, so $s_t > r_t$ for infinitely many t.

Step 2: Now, let T be such that $\forall t \geq T$, $f(r_t) < \varepsilon r_t, g(r_t) > (1-\varepsilon)f(r_t)$ and $T' > T$ be such that $s_{T'} > r_{T'}$. First inequality: on left hand sides we have monotone decreasing function and on RHS we have monotone increasing function that goes to ∞. For the second one we use the assumption of the lemma that $\lim_{x\to\infty} \frac{f(x)}{g(x)} = 0$. Hence such T exists. We are going to prove by induction that for any $t > T'$ $s_t > (1-\varepsilon)r_t$. Base of induction: $r_{T+1} = r_T + f(r_T) < (1+\varepsilon)r_T < (1+\varepsilon)s_T < (1+\varepsilon)s_{T+1}$. And notice that $\frac{1}{1+\varepsilon} > 1-\varepsilon$. Step of induction: suppose up to some $k > T$ our assumption holds then we either have $s_k > r_k$ and we can apply the same argument as above or $r_k > s_k > (1-\varepsilon)r_k$. Then $s_{k+1} = s_k + g(s_k) > (1-\varepsilon)r_k + g(r_k) > (1-\varepsilon)r_k + (1-\varepsilon)f(r_k) = (1-\varepsilon)r_{k+1}$, where we used the fact that g is monotone decreasing in the first inequality. Thus, by induction $s_t > (1-\varepsilon)r_t \ \forall t \geq T'$.

□

We will now use the above lemmas to construct bounding sequences for ℓ_t. Recalling (6), where we define $b(x) := x^{-1}e^{-\sigma^2 x^2/8}$, let $\bar{\ell}_0 = 1$, and let

$$\bar{\ell}_{t+1} = \bar{\ell}_t + b((1-\delta)\bar{\ell}_t)$$

(for some $\delta > 0$), and let $\underline{\ell}_0 = 1$ and

$$\underline{\ell}_{t+1} = \underline{\ell}_t + b(\underline{\ell}_t).$$

Lemma 9. *For any $\varepsilon > 0$, $(1-\varepsilon)\underline{\ell}_t < \ell_t < (1+\varepsilon)\bar{\ell}_t$ for all sufficiently large t.*

Proof. Notice that if we take $D_+(x)$ and $b((1-\delta)x)$ (for some $\delta > 0$) to be the functions $f(x)$, $g(x)$, then by lemma A.1, they satisfy the conditions of lemma A.2. Applying lemma A.2 establishes the upper bound for ℓ_t. Likewise, if we take $b(x)$ and $D_+(x)$ as $f(x), g(x)$ then we get the lower bound for ℓ_t. □

It now remains to solve the recurrence relations for $\bar{\ell}_t$ and $\underline{\ell}_t$.

Lemma 10. *Let $k > 0$, $r_0 > 0$, $r_{t+1} = r_t + r_t^{-1}e^{-kr_t^2}$. Then for any $\varepsilon > 0$,*

$$\frac{1-\varepsilon}{\sqrt{k}}\sqrt{\ln t} < r_t < \frac{1+\varepsilon}{\sqrt{k}}\sqrt{\ln t}$$

for all sufficiently large t.

Proof. Let $f(x) = \alpha\sqrt{\ln\left(e^{\left(\frac{x}{\alpha}\right)^2}+1\right)} - \alpha\sqrt{\ln e^{\left(\frac{x}{\alpha}\right)^2}}$. Note that $\alpha\sqrt{\ln(t+1)} = \alpha\sqrt{\ln t} + f(\alpha\sqrt{\ln t})$.

Multiplying $f(x)$ by

$$\frac{\sqrt{\ln\left(e^{\left(\frac{x}{\alpha}\right)^2}+1\right)} + \sqrt{\ln e^{\left(\frac{x}{\alpha}\right)^2}}}{\sqrt{\ln\left(e^{\left(\frac{x}{\alpha}\right)^2}+1\right)} + \sqrt{\ln e^{\left(\frac{x}{\alpha}\right)^2}}}$$

we may rewrite $f(x)$ as:

$$f(x) = \frac{\alpha\ln\left(1+e^{-\left(\frac{x}{\alpha}\right)^2}\right)}{\sqrt{\ln\left(e^{\left(\frac{x}{\alpha}\right)^2}+1\right)} + \sqrt{\ln e^{\left(\frac{x}{\alpha}\right)^2}}}$$

For all sufficiently small $\varepsilon > 0$, we have $\frac{1}{2}\varepsilon < \ln(1+\varepsilon) < \varepsilon$. So for all sufficiently large x, we have

$$\frac{1}{2}\alpha e^{-\left(\frac{x}{\alpha}\right)^2} < \alpha\ln\left(1+e^{-\left(\frac{x}{\alpha}\right)^2}\right) < \alpha e^{-\left(\frac{x}{\alpha}\right)^2}$$

Further, for all sufficiently large x, we have

$$\frac{x}{\alpha} < \sqrt{\ln\left(e^{\left(\frac{x}{\alpha}\right)^2}+1\right)} + \sqrt{\ln e^{\left(\frac{x}{\alpha}\right)^2}} < 4\frac{x}{\alpha}$$

Combining these inequalities, we find that for all sufficiently large x,

$$\frac{\alpha^2}{8}\left(\frac{e^{-\frac{1}{\alpha^2}x^2}}{x}\right) < f(x) < \alpha^2\left(\frac{e^{-\frac{1}{\alpha^2}x^2}}{x}\right)$$

In particular, if $\alpha < \frac{1}{\sqrt{k}}$, then $\lim_{x\to\infty} \frac{f(x)}{\frac{e^{-kx^2}}{x}} = 0$, and if $\alpha > \frac{1}{\sqrt{k}}$, then $\lim_{x\to\infty} \frac{\frac{e^{-kx^2}}{x}}{f(x)} = 0$.

Assume now that $\varepsilon < \frac{1}{2}$.

Taking $\alpha = \frac{1-\frac{\varepsilon}{3}}{\sqrt{k}}$, we have by lemma A.2 that for all sufficiently large t,

$$r_t > (1-\frac{\varepsilon}{3})\left(\frac{1-\frac{\varepsilon}{3}}{\sqrt{k}}\sqrt{\ln t}\right) > \frac{1-\varepsilon}{\sqrt{k}}\sqrt{\ln t}$$

and taking $\alpha = \frac{1+\varepsilon^{10}}{\sqrt{k}}$, we have by lemma A.2 that for all sufficiently large t,

$$(1-\varepsilon^5)r_t < \frac{1+\varepsilon^{10}}{\sqrt{k}}\sqrt{\ln t}$$

so

$$r_t < \frac{1+\varepsilon^{10}}{1-\varepsilon^5}\frac{1}{\sqrt{k}}\sqrt{\ln t} < \frac{1+\varepsilon}{\sqrt{k}}\sqrt{\ln t}$$

□

Proof of Theorem 4. Condition on $\theta = +1$. Then with probability 1, for all sufficiently large t, $a_t = 1$. Thus, with probability 1, for all sufficiently large t, $\ell_{t+1} = \ell_t + D_+(\ell_t)$.

Let the sequences $\bar{\ell}_t$ and $\underline{\ell}_t$ be as above. Then by Lemma A.3, with probability 1, for all t sufficiently large, $(1-\varepsilon)\underline{\ell}_t < \ell_t < (1+\varepsilon)\bar{\ell}_t$.

Finally, by Lemma A.4, there exist α_1 and α_2 such that for all sufficiently large t, $\alpha_1\sqrt{\ln t} < \underline{\ell}_t$ and $\bar{\ell}_t < \alpha_2\sqrt{\ln t}$.

Thus, with probability 1, $(1-\varepsilon)\alpha_1\sqrt{\ln t} < \ell_t < (1+\varepsilon)\alpha_2\sqrt{\ln t}$ for all sufficiently large t. □

Before proving Theorem 5 we state and prove the following lemma.

Lemma 11. *Conditioned on $\theta = +1$, all agents take action $+1$ with positive probability.*

Proof. The event that all agents take action $+1$ is equal to the event that $\ell_t > 0$ for all t. Since, conditioned on $\theta = +1$, almost surely $\ell_t > 0$ for all t large enough, and since ℓ_t is a Markov chain, there is some $x_0 > 0$ such that if, for some t, $\ell_t = x$ then with positive probability $\ell_{t'} > 0$ for all $t' > t$. A simple coupling argument shows that the same must then hold for any $x > x_0$. Finally, for any t_0, it holds that with positive probability the first t_0 agents take action $+1$. Therefore, if we take t_0 large enough, then $\ell_{t_0} \geq x_0$ with positive probability, with $\ell_t > 0$ for all $t \leq t_0$, and thus with positive probability $\ell_t > 0$ for all t. □

Proof of Theorem 5. Let E_t be the event that agent t takes action -1. Let A_t be the event that the first $t-1$ agents take action 1. Then $\mathbb{P}(E_t|\theta = 1) \geq \mathbb{P}(A_t \cap E_t|\theta = 1)$, so $\mathbb{P}(A_t \cap E_t|\theta = 1)$ provides a lower bound for $\mathbb{P}(a_t = -1|\theta = 1)$.

Now,

$$\mathbb{P}(A_t \cap E_t|\theta = 1) = \mathbb{P}(E_t|A_t, \theta = 1) \cdot \mathbb{P}(A_t|\theta = 1).$$

Let $p = \mathbb{P}(\{a_t = 1\}_{t\geq 0}|\theta = 1)$, the probability that all agents take the correct action when $\theta = 1$. By Lemma 11, $p > 0$. Thus, $\mathbb{P}(A_t|\theta = 1) \geq p$, so

$$\mathbb{P}(A_t \cap E_t|\theta = 1) \geq p \cdot \mathbb{P}(E_t|A_t, \theta = 1)$$

Putting these observations together, we find that

$$\mathbb{P}(a_t = -1|\theta = 1) \geq p \cdot \mathbb{P}(E_t|A_t, \theta = 1).$$

So what remains is to give a lower bound for $\mathbb{P}(E_t|A_t, \theta = 1)$.

Let ℓ_t^* denote the log-likelihood public belief given that the first $t-1$ agents take action 1. Then by the previous lemmas, for any $\varepsilon > 0$, $\ell_t^* < \beta\sqrt{\ln t}$ for all sufficiently large t, where $\beta = (1+\varepsilon)\frac{2\sqrt{2}}{\sigma}$. Further, by using L'Hôpital's rule, it can be shown that there exists $K > 0$ such that

$$\mathbb{P}(E_t|\ell_t^*, \theta = 1) > K\frac{e^{-(\sigma^2\ell_t^*/2+1)^2/(2\sigma^2)}}{\ell_t^*}$$

for all sufficiently large t.

Thus,

$$\mathbb{P}(E_n|\ell_t^*, \theta = 1) > K\frac{e^{-(\sigma^2\ell_t^*/2+1)^2/(2\sigma^2)}}{\ell_t^*} \geq K\frac{e^{-\left(\sigma^2\beta\sqrt{\ln t}/2+1\right)^2/(2\sigma^2)}}{\beta\sqrt{\ln t}}$$

for all sufficiently large t. Simplifying the right hand side, we find that for some $C > 0$ and all sufficiently large t,

$$\mathbb{P}(E_t|\ell_t^*, \theta = 1) > KCe^{-(1+\varepsilon)\ln t}.$$

Thus, we have that for all sufficiently large t,

$$\mathbb{P}(a_t = -1|\theta = 1) > pKCe^{-(1+\varepsilon)\ln t}.$$

□

Let us compare how the public belief changes when each agent has access to all of her predecessors' private signals. The probability of making a mistake in this case decays exponentially:

$$\mathbb{P}(\ell_t < 0|\theta = 1) = \Phi(\tfrac{-t}{\sqrt{t}\sigma}) = 1 - \Phi\left(\frac{\sqrt{t}}{\sigma}\right) = \frac{1}{2}\sigma\frac{e^{-t/\sigma^2}}{\sqrt{t}} + o\left(\frac{e^{-t/\sigma^2}}{t}\right).$$

References

[1] Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar, *Bayesian learning in social networks*, The Review of Economic Studies **78** (2011), no. 4, 1201–1236.

[2] Abhijit V Banerjee, *A simple model of herd behavior*, The Quarterly Journal of Economics (1992), 797–817.

[3] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch, *A theory of fads, fashion, custom, and cultural change as informational cascades*, Journal of political Economy (1992), 992–1026.

[4] Christophe Chamley, *Rational herds: Economic models of social learning*, Cambridge University Press, 2004.

[5] Lones Smith and Peter Sørensen, *Pathological outcomes of observational learning*, Econometrica **68** (2000), no. 2, 371–398.

[6] Peter Norman Sørensen, *Rational social learning*, Ph.D. Thesis, 1996.

California Institute of Technology