

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 convergence occurs, and show that it 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.

1. 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 make decisions based on the observed actions of others. 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 [7]. 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 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?

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

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 speed 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 over time [7]; 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.

We next analyze the specific canonical case of Gaussian private signals. Here we calculate precisely the asymptotic behavior of the log-likelihood ratio of the public belief. We show that learning from actions is significantly slower than learning from signals: the log-likelihood ratio behaves asymptotically as $\sqrt{\log t}$. To calculate this we develop a technique that allows, much more generally, for the long-term evolution of the public belief to be calculated for a large class of signal distributions.

Since, in our setting of unbounded signals, agents eventually take the correct action, an additional, natural measure of the speed of learning is the expected time at which this happens: how long does it take until no more mistakes are made?

We show that the answer to this question crucially depends on the signal distributions. For distributions, such as the Gaussian, in which strong signals occur with very small probability, we show that this

expected time is infinite.¹ However, when strong signals are less rare, this expectation is finite. Intuitively, when strong signals are rare, agents are more likely to emulate their predecessors, and so it may take a long time for a mistake to be corrected.

Finally, in the Gaussian case, we study another measure of the speed of learning. Namely, we consider directly how the probability of choosing the incorrect action varies as agents see more and more of the other agents' decisions before making their own. We find that this probability is asymptotically no less than $1/t^{1+\varepsilon}$ for any $\varepsilon > 0$. In contrast, when agents 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.

1.1. Related literature. Several previous studies have considered the same question. Chamley [4] gives an estimate for the evolution of the public belief for a class of private signal distributions with fat tails. He also studies the speed of convergence in the Gaussian case using a computer simulation. Sørensen [8] has published a claim related to our Theorem 1, with an unfinished proof. A concurrent paper by Rosenberg and Vieille [6] studies related questions. In particular they study the time until the first correct action, as well as the number of incorrect actions—which are related to our time to learn—and characterize when they have finite expectations.

A related model is studied by Lobel, Acemoglu, Dahleh and Ozdaglar [5], who consider agents who also act sequentially, but do not observe all of their predecessors' actions. They study how the speed of learning varies with the network structure. Vives [9], in a paper with a very similar spirit to ours, studies the speed of sequential learning in a model with actions chosen from a continuum, and where agents observe a noisy signal about their predecessors' actions. He similarly shows that learning is significantly slower than in the benchmark case. An overview of this literature is given by Vives in his book [10, Chapter 6].

2. MODEL

Let $\theta \in \{-1, +1\}$ be the true state of the world, with each state a priori equally likely². 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 cumulative distribution function (CDF) F_+ and if $\theta = -1$ they

¹In the benchmark case of observed signals this time is finite, for any signal distribution.

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

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

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

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:

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

In equilibrium, agent t chooses $a_t = +1$ iff⁴

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

A simple calculation shows that this occurs iff

$$\ell_t + L_t > 0.$$

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

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

where

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

Likewise, when $a_t = -1$,

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

³One 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.

⁴For simplicity, we assume that agents choose action -1 when indifferent. This will have no impact on our results.

where

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

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

3. THE EVOLUTION OF PUBLIC BELIEF

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 $\tilde{\ell}_t$ would equal the sum

$$\tilde{\ell}_t = \sum_{\tau=1}^t L_\tau.$$

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 $\lim_t \tilde{\ell}_t/t$ would—conditioned on (say) $\theta = +1$ —equal the conditional expectation of L_t , which is positive.⁵

3.1. Sub-linear public beliefs. 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 .

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

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

with probability 1.

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

⁵In 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.

3.2. Long-term behavior of public beliefs. We next turn to estimating more precisely the long-term behavior of the public log-likelihood ratio ℓ_t . 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 1; all later agents choose the correct action.

We consider without loss of generality the case that $\theta = +1$, so that ℓ_t is positive from some t on. Thus, recalling (1), we have that from some t on,

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

This is the recurrence relation that we need to solve in order to understand the long term evolution of ℓ_t . To this end, we consider the corresponding differential equation:

$$\frac{df}{dt}(t) = D_+(f(t)).$$

Recall that G_- is the CDF of the private log-likelihood ratio L_t , conditioned on $\theta = -1$. We show (Lemma 8) that $D_+(x)$ is well approximated by $G_-(-x)$ for high x , in the sense that

$$\lim_{x \rightarrow \infty} \frac{D_+(x)}{G_-(-x)} = 1.$$

In some applications (including the Gaussian one, which we consider below), the expression for G_- is simpler than that for D_+ , and so one can instead consider the differential equation

$$(2) \quad \frac{df}{dt}(t) = G_-(-f(t)).$$

This equation can be solved analytically in many cases in which G_- has a simple form. For example, if $G_-(-x) = e^{-x}$ then $f(t) = \log(t + c)$, and if $G_-(-x) = x^{-k}$ then $f(t) = ((k + 1) \cdot t + c)^{1/(k+1)}$.

We show that solutions to this equation have the same long term behavior as ℓ_t , given that G_- satisfies some regularity conditions.

Theorem 3. *Suppose that G_- and G_+ are continuous, and that the left tail of G_- is convex and differentiable. Suppose also that $f: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ satisfies*

$$(3) \quad \frac{df}{dt}(t) = G_-(-f(t))$$

for all sufficiently large t . Then conditional on $\theta = +1$,

$$\lim_{t \rightarrow \infty} \frac{\ell_t}{f(t)} = 1$$

with probability 1.

The condition⁶ on G_- is satisfied when the random variables L_t (i.e., the log-likelihood ratios associated with the private signals), conditioned on $\theta = -1$, have a distribution with a probability density function that is monotone decreasing for all x less than some x_0 . This is the case for the normal distribution, and for practically every non-atomic distribution one may encounter in the standard probability and statistics literatures.

3.2.1. Gaussian signals. In the Gaussian case, F_+ is Normal with mean $+1$ and variance σ^2 , and F_- is Normal with mean -1 and the same variance. A simple calculation shows that G_- is the Gaussian cumulative distribution function, and so we cannot solve the differential equation (2) analytically. However, we can bound $G_-(x)$ from above and from below by functions of the form $e^{-c \cdot x^2}/x$. For these functions the solution to (2) is of the form $f(t) = \sqrt{\log t}$, and so we can use Theorem 3 to deduce the following.

Theorem 4. *When private signals are Gaussian, then conditioned on $\theta = +1$,*

$$\lim_{t \rightarrow \infty} \frac{\ell_t}{(2\sqrt{2}/\sigma) \cdot \sqrt{\log t}} = 1$$

with probability 1.

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

3.3. The expected time to learn. When private signals are unbounded then with probability 1 the agents eventually all choose the correct action $a_t = \theta$. A natural question is: how long does it take for that to happen? Formally, we define the *time to learn*

$$T_L = \min\{t : a_\tau = \theta \text{ for all } \tau \geq t\},$$

and study its expectation. Note that in the baseline case of observed signals T_L has finite expectation, since the probability of a mistake at time t decays exponentially with t .

We first study the expectation of T_L in the case of Gaussian signals. To this end we define the *time of first mistake* by

$$T_1 = \min\{t : a_t \neq \theta\}$$

⁶By “the left tail of G_- is convex and differentiable” we mean that there is some x_0 such that, restricted to $(-\infty, x_0)$, G_- is convex and differentiable.

if $a_t \neq \theta$ for some t , and by $T_1 = 0$ otherwise. We calculate a lower bound for the distribution of T_1 , showing that it decays at most as fast as $1/t$.

Theorem 5. *When private signals are Gaussian then for every $\varepsilon > 0$ there exists a $k > 0$ such that for all t*

$$\mathbb{P}(T_1 = t) \geq \frac{k}{t^{1+\varepsilon}}.$$

Thus T_1 has a very thick tail, decaying far slower than the exponential decay of the baseline case. In particular, T_1 has infinite expectation, and so, since $T_L > T_1$, the expectation of the time to learn T_L is also infinite.

In contrast, we show that when private signals have thick tails—that is, when the probability of a strong signal vanishes slowly enough—then the time to learn has finite expectation. In particular, we show this when the left tail of G_- and the right tail of G_+ are polynomial.⁷

Theorem 6. *Assume that $G_-(-x) = c \cdot x^{-k}$ and that $G_+(x) = 1 - c \cdot x^{-k}$ for some $c > 0$ and $k > 0$, and for all x greater than some x_0 . Then $\mathbb{E}(T_L) < \infty$.*

An example of private signal distributions F_+ and F_- for which G_- and G_+ have this form is given by the probability density functions

$$f_-(x) = \begin{cases} c \cdot e^{-x} x^{-k-1} & \text{when } 1 \leq x \\ 0 & \text{when } -1 < x < 1 \\ c \cdot (-x)^{-k-1} & \text{when } x \leq -1. \end{cases}$$

and $f_+(x) = f_-(-x)$, for an appropriate choice of normalizing constant $c > 0$. In this case $G_-(-x) = 1 - G_+(x) = \frac{c}{k} x^{-k}$ for all $x > 1$.⁸

The proof of Theorem 6 is rather technically involved, and we provide here a rough sketch of the ideas behind it.

We say that there is an *upset* at time t if $a_{t-1} \neq a_t$. We denote by Ξ the random variable which assigns to each outcome the total number of upsets

$$\Xi = |\{t : a_{t-1} \neq a_t\}|.$$

⁷Recall that G_- is the conditional cumulative distribution function of the private log-likelihood ratios L_t .

⁸Theorem 6 can be proved for other thick-tailed private signal distributions: for example, one could take different values of c and k for G_- and G_+ , or one could replace their thick polynomial tails by even thicker logarithmic tails. For the sake of readability we choose to focus on this case.

We say that there is a *run* of length m from time t if $a_t = a_{t+1} = \dots = a_{t+m-1}$. As we will condition on $\theta = +1$ in our analysis, we say that a run from time t is *good* if $a_t = 1$ and *bad* otherwise. A trivial but important observation is that the number of maximal finite runs is equal to the number of upsets, and so, if $\Xi = n$, and if $T_L = t$, then there is at least one run of length t/n before time t . Qualitatively, this implies that if the number of upsets is small, and if the time to learn is large, then there is at least one long run before the time to learn.

We show that it is indeed unlikely that Ξ is large: the distribution of Ξ has an exponential tail. Incidentally, this holds for *any* private signal distribution:

Proposition 7. *For every private signal distribution there exist $c > 0$ and $0 < \gamma < 1$ such that for all $n > 0$*

$$\mathbb{P}(\Xi \geq n) \leq c\gamma^n.$$

Intuitively, this holds because whenever an agent takes the correct action, there is a non-vanishing probability that all subsequent agents will also do so, and no more upsets will occur.

Thus, it is very unlikely that the number of upsets Ξ is large. As we observe above, when Ξ is small then the time to learn T_L can only be large if at least one of the runs is long. When G_- has a thin tail then this is possible; indeed, Theorem 5 shows that the first finite run has infinite expected length when private signals are Gaussian. However, when G_- has a thick, polynomial left tail of order x^{-k} , we show that it is very unlikely for any run to be long: the probability that there is a run of length n decays at least as fast as $\exp(-n^{k/(k+1)})$, and in particular runs have finite expected length. Intuitively, when strong signals are rare then runs tend to be long, as agents are likely to emulate their predecessor. Conversely, when strong signals are more likely then agents are more likely to break a run, and so runs tend to be shorter.

Putting together these insights, we conclude that it is unlikely that there are many runs, and, in the polynomial signal case, it is unlikely that runs are long. Thus T_L has finite expectation.

3.4. Probability of taking the wrong action. Yet another natural metric of the speed of learning is the probability of mistake

$$p_t = \mathbb{P}(a_t \neq \theta).$$

Calculating the asymptotic behavior of p_t seems harder to tackle.

For the Gaussian case, while we cannot estimate p_t precisely, Theorem 5 immediately implies a lower bound: p_t is at least $k/t^{1+\varepsilon}$, for

every $\varepsilon > 0$ and k that depends on ε . This is much larger than the exponentially vanishing probability of mistake in the revealed signal baseline case.

More generally, we can use Theorem 1 to show that p_t vanishes sub-exponentially for any signal distribution, in the sense that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log p_t = 0.$$

To see this, note that the probability of mistake at time $t - 1$, conditioned on the observed actions, is exactly equal to

$$\min\{\mu_t, 1 - \mu_t\};$$

where we recall that

$$\mu_t = \mathbb{P}(\theta = +1 | a_1, \dots, a_{t-1}) = \frac{e^{\ell_t}}{e^{\ell_t} + 1}$$

is the public belief. This is due to the fact that if the outside observer, who holds belief μ_t , had to choose an action, she would choose a_{t-1} , the action of the last player she observed, a player who has strictly more information than her. Thus

$$p_t = \mathbb{E}(\min\{\mu_t, 1 - \mu_t\}) = \mathbb{E}\left(\frac{1}{e^{|\ell_t|} + 1}\right),$$

and since, by Theorem 1, $|\ell_t|$ is sub-linear, it follows that p_t is sub-exponential.

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, in the sense that the public log-likelihood ratio tends more slowly to infinity. 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 (among a large class of private signal distributions) and show that learning is very slow. We also show that in the Gaussian case the expected time to learn is infinite, as opposed to cases of more thick-tailed distributions, in which it is finite.

For the Gaussian case we also provide a lower bound for the probability of mistake. Finding a matching upper bound seems beyond our reach at the moment, and provides a compelling open problem for further research.

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] Ilan Lobel, Daron Acemoglu, Munther Dahleh, and Asuman Ozdaglar, *Rate of convergence of learning in social networks*, Proceedings of the american control conference, 2009.
- [6] Dinah Rosenberg and Nicolas Vieille, *On the efficiency of social learning*, private communication (2017).
- [7] Lones Smith and Peter Sørensen, *Pathological outcomes of observational learning*, Econometrica **68** (2000), no. 2, 371–398.
- [8] Peter Norman Sørensen, *Rational social learning*, Ph.D. Thesis, 1996.
- [9] Xavier Vives, *How fast do rational agents learn?*, The Review of Economic Studies **60** (1993), no. 2, 329–347.
- [10] ———, *Information and learning in markets: the impact of market microstructure*, Princeton University Press, 2010.

APPENDIX A. SUB-LINEAR LEARNING

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 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

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

It follows that

$$(4) \quad G_+(x) = \int_{-\infty}^x d\nu_+(\zeta) = \int_{-\infty}^x e^\zeta d\nu_-(\zeta).$$

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

Lemma 8.

$$\lim_{x \rightarrow \infty} \frac{D_+(x)}{G_-(-x)} = 1 \text{ and } \lim_{x \rightarrow -\infty} \frac{D_-(x)}{G_+(-x) - 1} = 1.$$

Proof. By definition,

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

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

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

Applying (4) yields

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

and so for any ε and all x large enough,

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

Using the same approximation of the logarithm, we have that

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

The statement for D_+ now follows by taking ε to zero. 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 8, 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. \square

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 $Q: \mathbb{R} \rightarrow \mathbb{R}_{>0}$ with $\lim_{x \rightarrow \infty} Q(x) = 0$, we can find CDFs such that $D_+(x) = -D_-(-x)$ and $\liminf_{x \rightarrow \infty} D_+(x)/Q(x) > 0$.

Fix any Q such that $\lim_{x \rightarrow \infty} Q(x) = 0$, but assume without loss of generality that $Q(x)$ is monotone decreasing.⁹ 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 .

⁹If 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

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

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 $x > 0$

$$(5) \quad G_-(-x) = C^{-1} \cdot Q(\lceil x \rceil - 1).$$

Now, by Lemma 8, 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

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

which, by (5) equals

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

□

APPENDIX B. LONG-TERM BEHAVIOR OF PUBLIC BELIEF

The primary goal of this section is to prove Theorem 3, which states that public belief is asymptotically given by the solution to the differential equation (3). The proof of this theorem uses two general lemmas regarding recurrence relations. We state these lemmas now and prove them later. The first lemma states that two similar recurrence relations yield similar solutions. The second shows that the solution to a recurrence relation (of the type we are interested in) is well approximated by the solution to the corresponding differential equation.

Lemma 9. *Let $A, B: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ be continuous, eventually monotone decreasing, and tending to zero.*

Let (a_t) and (b_t) be sequences satisfying the recurrence relations

$$\begin{aligned} a_{t+1} &= a_t + A(a_t) \\ b_{t+1} &= b_t + B(b_t). \end{aligned}$$

Suppose

$$\lim_{x \rightarrow \infty} \frac{A(x)}{B(x)} = 1.$$

Then

$$\lim_{t \rightarrow \infty} \frac{a_t}{b_t} = 1.$$

Lemma 10. *Assume that $A: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ is a continuous function with a convex differentiable tail, and that $A(x)$ goes to 0 as x goes to ∞ . Let (a_t) be any sequence satisfying the recurrence equation $a_{t+1} = a_t + A(a_t)$, and suppose there is a function $f: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ with $f'(t) = A(f(t))$ for all sufficiently large t . Then*

$$\lim_{t \rightarrow \infty} \frac{f(t)}{a_t} = 1.$$

Given these lemmas, we are ready to prove our theorem.

Proof of Theorem 3. Let (a_t) be any sequence in $\mathbb{R}_{>0}$ satisfying:

$$a_{t+1} = a_t + G_-(-a_t).$$

Then by Lemma 10, the sequence (a_t) is well approximated by $f(t)$, the solution to the corresponding differential equation:

$$\lim_{t \rightarrow \infty} \frac{a_t}{f(t)} = 1.$$

Now, conditional on $\theta = +1$, all agents take action +1 from some point on with probability 1. Thus, with probability 1,

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

for all sufficiently large t . Further, by Lemma 8,

$$\lim_{x \rightarrow \infty} \frac{D_+(x)}{G_-(-x)} = 1.$$

So by Lemma 9,

$$\lim_{t \rightarrow \infty} \frac{\ell_t}{a_t} = 1$$

with probability 1. Thus, we have

$$\lim_{t \rightarrow \infty} \frac{\ell_t}{f(t)} = \lim_{t \rightarrow \infty} \frac{\ell_t}{a_t} \cdot \frac{a_t}{f(t)} = 1$$

with probability 1. □

B.1. Proofs of Lemmas 9 and 10.

Proof of Lemma 9. We prove the claim in two steps. First, we show that for every $\varepsilon > 0$ there are infinitely many times t such that

$$(6) \quad (1 - \varepsilon)a_t \leq b_t \leq (1 + \varepsilon)a_t.$$

Second, we show that if (6) holds for some t large enough, then it holds for all $t' > t$, proving the claim.

We start with step 1. Assume without loss of generality that $a_t \leq b_t$ for infinitely many values of t . Fix $\varepsilon > 0$. To show that $(1 - \varepsilon)a_t \leq b_t \leq (1 + \varepsilon)a_t$ holds for infinitely many values of t , let $x_0 > 1$ be such that for all $x > x_0$ it holds that A and B are monotone decreasing,

$$A(x), B(x) < \varepsilon < 1$$

and

$$(7) \quad (1 - \varepsilon/2)A(x) < B(x) < (1 + \varepsilon/2)A(x).$$

Assume that $a_t, b_t > x_0$; this will indeed be the case for t large enough, since A and B are positive and continuous, and so both a_t and b_t are monotone increasing and tend to infinity. So

$$B(b_t) < (1 + \varepsilon/2)A(b_t) \leq (1 + \varepsilon/2)A(a_t),$$

where the first inequality follows from (7), and the second follows from the fact that A is monotone decreasing and $a_t < b_t$. Since $B(b(t)) = b_{t+1} - b(t)$ and $A(a_t) = a_{t+1} - a(t)$ we have shown that

$$b_{t+1} - b_t < (1 + \varepsilon/2)(a_{t+1} - a_t),$$

and so eventually $b_t \leq (1 + \varepsilon)a_t$. Also, notice that the first time this obtains, we also have that the left inequality in (6) holds at the same moment:

$$b_t > b_{t-1} > a_{t-1} = a_t - (a_t - a_{t-1}) > a_t - \varepsilon > a_t - \varepsilon a_t = (1 - \varepsilon)a_t.$$

This completes the first step. Now we go to step 2. Here we show that if (6) holds for large enough t then it holds for all $t' > t$.

Fix $\varepsilon > 0$, and let x_0 be defined as above. Suppose that $(1 - \varepsilon)a_t < b_t < (1 + \varepsilon)a_t$, with $a_t, b_t > x_0$. Assume without loss of generality that $b_t \geq a_t$. Then our assumptions and (7) imply

$$\begin{aligned} b_{t+1} &= b_t + B(b_t) \\ &< (1 + \varepsilon)a_t + (1 + \varepsilon)A(b_t). \end{aligned}$$

Because $a_t \leq b_t$ and A is decreasing we have

$$\begin{aligned} b_{t+1} &< (1 + \varepsilon)a_t + (1 + \varepsilon)A(a_t) \\ &= (1 + \varepsilon)a_{t+1}. \end{aligned}$$

For the other direction, note first that

$$b_{t+1} > b_t \geq a_t,$$

by assumption. We can write $a_t = (1 - \varepsilon)a_t + \varepsilon a_t$, and since $a_t > x_0 > 1$, $\varepsilon a_t > (1 - \varepsilon)\varepsilon$, and so

$$b_{t+1} > (1 - \varepsilon)a_t + (1 - \varepsilon)\varepsilon.$$

Now, $\varepsilon > A(a_t)$ since $a_t > x_0$, and so

$$\begin{aligned} b_{t+1} &> (1 - \varepsilon)a_t + (1 - \varepsilon)A(a_t) \\ &= (1 - \varepsilon)a_{t+1}. \end{aligned}$$

Thus

$$(8) \quad (1 - \varepsilon)a_{t+1} < b_{t+1} < (1 + \varepsilon)a_{t+1},$$

as required. □

Proof of Lemma 10. We restrict the domain of f to the interval (t_0, ∞) such that for $t > t_0$ it already holds that $f'(t) = A(f(t))$. Since A is continuous, $\lim_{t \rightarrow \infty} f(t) = \infty$, and so we can also assume that in the interval $(f(t_0), \infty)$ it holds that A is convex and differentiable.

Since f is strictly increasing in (t_0, ∞) , it has an inverse f^{-1} . For x large enough define $B(x) = f(f^{-1}(x) + 1) - x$.

Now, let (b_t) be any sequence satisfying the recurrence relation

$$b_{t+1} = b_t + B(b_t).$$

In order to apply Lemma 9, we will first show that

$$\lim_{x \rightarrow \infty} \frac{B(x)}{A(x)} = 1.$$

Let $t = f^{-1}(x)$. Such a t exists and is unique for all sufficiently large x , because f is monotone. Notice that by the definitions of $B(x)$ and $f'(x)$

$$\begin{aligned} B(x) &= f(f^{-1}(x) + 1) - x \\ &= f(f^{-1}(x) + 1) - x - f'(f^{-1}(x)) + f'(f^{-1}(x)) \\ &= f(t + 1) - f(t) - f'(t) + A(f(t)), \end{aligned}$$

where in the last equality we substitute $t = f^{-1}(x)$. Because f' is positive and decreasing (f is concave) then $f(t+1) - f(t) \geq f'(t+1)$, and so

$$B(x) \geq f'(t+1) - f'(t) + A(f(t)).$$

By the definition of f , $f'(t) = A(f(t))$, and so

$$B(x) \geq A(f(t+1)) - A(f(t)) + A(f(t)) = A(f(t+1)).$$

Again, due to concavity of f we have $f(t+1) \leq f(t) + f'(t)$ and as A is decreasing and convex we get

$$\begin{aligned} B(x) &\geq A(f(t) + f'(t)) \\ &\geq A'(f(t))f'(t) + A(f(t)) \\ &= A'(f(t))A(f(t)) + A(f(t)). \end{aligned}$$

We now substitute back $x = f(t)$:

$$\begin{aligned} B(x) &\geq A'(x)A(x) + A(x) \\ &= A(x)(A'(x) + 1) \end{aligned}$$

so in particular, since $A'(x) \rightarrow 0$ as $x \rightarrow \infty$,

$$\liminf_{x \rightarrow \infty} \frac{B(x)}{A(x)} \geq 1.$$

Now we are going to show that $\limsup_{x \rightarrow \infty} \frac{B(x)}{A(x)} \leq 1$ which will conclude the proof. By the definitions of $f^{-1}(x)$ and $B(x)$

$$B(x) = B(f(t)) = f(t+1) - f(t) = \int_t^{t+1} f'(\zeta) d\zeta.$$

As f' is decreasing it follows that

$$B(x) \leq \int_t^{t+1} f'(t) d\zeta = f'(t) = A(f(t)) = A(x).$$

Therefore,

$$\limsup_{x \rightarrow \infty} \frac{B(x)}{A(x)} \leq 1.$$

Hence, from these two inequalities we get that

$$\lim_{x \rightarrow \infty} \frac{B(x)}{A(x)} = 1.$$

Now notice that, by construction, $f(t+1) = f(t) + B(f(t))$. Thus, by Lemma 9,

$$\lim_{n \rightarrow \infty} \frac{f(t)}{a_t} = 1.$$

□

B.2. Monotonicity of solutions to a differential equation. We now prove a general lemma regarding differential equations of the form $a'(t) = A(a(t))$. It shows that the solutions to this equation are monotone in A . This is useful for calculating approximate analytic solutions whenever it is impossible to find analytic exact solutions, as is the case of Gaussian signals, in which we use this lemma.

Lemma 11. *Let $A, B: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ be continuous, and let $a, b: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ satisfy $a'(t) = A(a(t))$ and $b'(t) = B(b(t))$ for all sufficiently large t .*

Suppose that

$$\liminf_{x \rightarrow \infty} \frac{A(x)}{B(x)} > 1.$$

Then $a(t) > b(t)$ for all sufficiently large t .

Proof. Notice that $a(t)$ and $b(t)$ are eventually monotone increasing and tend to infinity as t tends to infinity. Thus for all x greater than some $x_0 > 0$ large enough, a and b have inverses that satisfy the following differential equations:

$$\begin{aligned} \frac{d}{dx} a^{-1}(x) &= \frac{1}{A(x)} \\ \frac{d}{dx} b^{-1}(x) &= \frac{1}{B(x)}. \end{aligned}$$

Since $\liminf_x A(x)/B(x) > 1$, we can furthermore choose x_0 so that for all $x \geq x_0$, $A(x) > (1 + \varepsilon)B(x)$ for some $\varepsilon > 0$. Thus, for $x > x_0$

$$\begin{aligned} a^{-1}(x) &= a^{-1}(x_0) + \int_{x_0}^x \frac{1}{A(x)} dx \\ b^{-1}(x) &= b^{-1}(x_0) + \int_{x_0}^x \frac{1}{B(x)} dx \end{aligned}$$

and so

$$\begin{aligned} a^{-1}(x) &< a^{-1}(x_0) + \frac{1}{1 + \varepsilon} \int_{x_0}^x \frac{1}{B(x)} dx \\ &= a^{-1}(x_0) + \frac{1}{1 + \varepsilon} (b^{-1}(x) - b^{-1}(x_0)) \end{aligned}$$

and thus

$$a^{-1}(x) - b^{-1}(x) < -\frac{\varepsilon}{1+\varepsilon}b^{-1}(x) + \left[a^{-1}(x_0) - \frac{1}{1+\varepsilon}b^{-1}(x_0) \right].$$

Since $b^{-1}(x)$ tends to infinity as x tends to infinity, it follows that for all sufficiently large x , $a^{-1}(x) < b^{-1}(x)$. Thus, for all sufficiently large t

$$t = a^{-1}(a(t)) < b^{-1}(a(t)),$$

and so, since $b(t)$ is monotone increasing,

$$b(t) < a(t).$$

□

B.3. Eventual monotonicity of public belief update. We end this section with a lemma that shows that under some technical conditions on the left tail of G_- , the function $u_+(x) = x + D_+(x)$ (i.e., the function that determines how the public log-likelihood ratio is updated when the action +1 is taken) is eventually monotone increasing.

Lemma 12. *Suppose G_- has a convex and differentiable left tail. Then the map $u_+(x) = x + D_+(x)$ is monotone increasing for all sufficiently large x .*

Proof. Recall that

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

Since G_- has a differentiable left tail, it has a derivative $g_-(-x)$ for all x large enough. It then follows from (4) that G_+ also has a derivative in this domain, and

$$\begin{aligned} u'_+(x) &= 1 + \frac{g_+(-x)}{1 - G_+(-x)} - \frac{g_-(-x)}{1 - G_-(-x)} \\ &= 1 + \frac{e^{-x}g_-(-x)}{1 - G_+(-x)} - \frac{g_-(-x)}{1 - G_-(-x)}. \end{aligned}$$

Since $1 - G_-(-x)$ and $1 - G_+(-x)$ tend to 1 as x tends to infinity,

$$\lim_{x \rightarrow \infty} u'_+(x) = \lim_{x \rightarrow \infty} 1 + e^{-x}g_-(-x) - g_-(-x).$$

Since G_- is eventually convex, $g_-(-x)$ tends to zero, and therefore

$$\lim_{x \rightarrow \infty} u'_+(x) = 1.$$

In particular, $u'_+(x)$ is positive for x large enough, and hence $u_+(x)$ is eventually monotone increasing. □

APPENDIX C. GAUSSIAN PRIVATE SIGNALS

C.1. Preliminaries. We say that private signals are Gaussian when F_- is the normal distribution with mean -1 and variance σ^2 , and F_+ is the normal distribution with mean $+1$ and variance σ^2 . To calculate the evolution of ℓ_t , we need to calculate G_+ and G_- , the conditional distributions of the private log-likelihood ratio L_t . Notice that in this case

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

so that L_t is simply proportional to the signal s_t . It follows that L_t is also normally distributed, conditioned on the state θ , and that G_+ and G_- are cumulative distribution functions of Gaussians, with variance $4/\sigma^2$.

C.1.1. Notation. In this section and those that follow, we denote by ℓ_t^* the public log-likelihood ratio when all agents before agent t take the correct action. Formally,

$$\ell_t^* = \log \frac{\mathbb{P}(\theta = +1 \mid a_1 = \dots = a_{t-1} = +1)}{\mathbb{P}(\theta = -1 \mid a_1 = \dots = a_{t-1} = +1)}.$$

For convenience, we will also use the notation $\mathbb{P}_+(\cdot)$ as shorthand for $\mathbb{P}(\cdot \mid \theta = +1)$.

C.2. The evolution of public belief.

Proof of Theorem 4. Let $f: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ be any function such that $f'(t) = G_-(-f(t))$ for all sufficiently large t . Then by Theorem 3,

$$\lim_{t \rightarrow \infty} \frac{\ell_t}{f(t)} = 1$$

with probability 1.

Recall from above that L_t is distributed normally, and $G_-(-x)$ is the CDF of a normal distribution with variance $\tau^2 = 4/\sigma^2$.

For $1 > \eta \geq 0$, define

$$F_\eta(x) = \frac{e^{-\frac{1-\eta}{2\tau^2}x^2}}{x}$$

$$f_\eta(t) = \frac{\sqrt{2}\tau}{\sqrt{1-\eta}} \sqrt{\log(t) + \log \frac{(1-\eta)^2}{2\tau^2}}.$$

By a routine application of L'Hospital's rule, F_0 and F_η are lower and upper bounds for G_- , in the sense that

$$\begin{aligned}\lim_{x \rightarrow \infty} \frac{G_-(-x)}{F_0(x)} &= \infty \\ \lim_{x \rightarrow \infty} \frac{F_\eta(x)}{G_-(-x)} &= \infty, \eta > 0.\end{aligned}$$

Since $f'_\eta(t) = F_\eta(f_\eta(t))$ for all sufficiently large t , we have by Lemma 11 that for any $\eta > 0$,

$$f_0(t) < f(t) < f_\eta(t)$$

for all sufficiently large t . So

$$\liminf_{t \rightarrow \infty} \frac{f(t)}{\sqrt{2\tau}\sqrt{\log t}} = \liminf_{t \rightarrow \infty} \frac{f(t)}{f_0(t)} \geq 1$$

and for any $\eta > 0$,

$$\limsup_{t \rightarrow \infty} \frac{f(t)}{\sqrt{2\tau}\sqrt{\log t}} = \frac{1}{\sqrt{1-\eta}} \cdot \limsup_{t \rightarrow \infty} \frac{f(t)}{f_\eta(t)} \leq \frac{1}{\sqrt{1-\eta}}.$$

Thus,

$$\lim_{t \rightarrow \infty} \frac{f(t)}{\sqrt{2\tau}\sqrt{\log t}} = \lim_{t \rightarrow \infty} \frac{f(t)}{(2\sqrt{2}/\sigma)\sqrt{\log t}} = 1$$

so with probability 1,

$$\lim_{t \rightarrow \infty} \frac{\ell_t}{(2\sqrt{2}/\sigma)\sqrt{\log t}} = \lim_{t \rightarrow \infty} \frac{\ell_t}{f(t)} \cdot \frac{f(t)}{(2\sqrt{2}/\sigma)\sqrt{\log t}} = 1.$$

□

To prove Theorem 5, we will need two lemmas. The first is general, and will be used several times in the sequel, while the second deals exclusively with the Gaussian case.

Denote by E_t the event that $a_\tau = +1$ for all $\tau \geq t$; that is, that there are no more mistakes after time t . The next lemma provides a uniform bound for the probability of E_t , conditioned on the public belief. It implies, in particular, that the probability of E_1 is positive, which we will use in the proof of Theorem 5.

Lemma 13. *Suppose G_- and G_+ are continuous, and G_- has a convex and differentiable left tail. Then for every $L \in \mathbb{R}$, there is some $m_L > 0$ such that for any t , $x \geq L$ implies $\mathbb{P}_+(E_t | \ell_t = x) \geq m_L$.*

Proof. Recall the definition of the public belief $\mu_t = \mathbb{P}(\theta = +1 | a_1, \dots, a_{t-1})$. The process (μ_1, μ_2, \dots) is a bounded martingale, and therefore, by a standard argument on bounded martingales, if we condition on $\mu_t = q$, then the probability that $\mu_\tau \leq 1/2$ for some $\tau > t$ is at most $2(1-q)$.¹⁰ This event is precisely the complement of E_t , and therefore we have that $\mathbb{P}(E_t | \mu_t = q)$ is at least $2q - 1$. Hence, conditioning on $\theta = +1$, we have that $\mathbb{P}_+(E_t | \mu_t = 1 - q) \geq (2q - 1)/q$, which is positive for all $q > 1/2$.

Since $\mu_t = q$ is equivalent to $\ell_t = \log q / (1 - q)$, what we have shown implies that there is an $\varepsilon > 0$ such that for all $x \geq 1$ (here the choice of 1 is arbitrary and can be replaced with any positive number)

$$\mathbb{P}_+(E_t | \ell_t = x) > \varepsilon.$$

Now, for any $L < 1$, the compactness of the interval $[L, 1]$, together with the continuity of G_- and G_+ , implies that there is an n_L such that if $\ell_t \geq L$, and if agents t through $t + n_L - 1$ take action +1, then $\ell_{t+n_L} > 1$. Further, since the probability of agents t through $t + n_L - 1$ all taking action +1 conditional on $\ell_t = x$ is continuous in x , there is a $p_L > 0$ such that

$$\mathbb{P}_+(E_t | \ell_t = x) \geq p_L \cdot \varepsilon$$

since with probability at least p_L there are no mistakes up to time $t + n_L$, and thence there are no mistakes with probability at least ε . \square

Lemma 14. *Assume private signals are Gaussian. For every $\varepsilon > 0$ there exists some $k > 0$ such that for all t ,*

$$\mathbb{P}_+(a_t = -1 | a_\tau = +1 \text{ for all } \tau < t) > \frac{k}{t^{1+\varepsilon}}.$$

Proof. By the definitions of ℓ_t^* and G_+ ,

$$\begin{aligned} \mathbb{P}_+(a_t = -1 | a_\tau = +1 \text{ for all } \tau < t) &= \mathbb{P}_+(a_t = -1 | \ell_t = \ell_t^*) \\ &= G_+(-\ell_t^*). \end{aligned}$$

Now, by Theorem 4, for every $\beta > 0$, $\ell_t^* < (1 + \beta) \frac{2\sqrt{2}}{\sigma} \sqrt{\log t}$ for all sufficiently large t . Further, it follows from a routine application of L'Hopital's rule (or from the standard asymptotic expansion for the CDF of a normal distribution) that for all sufficiently large x ,

$$G_+(-x) > \frac{e^{-(\sigma^2/8)x^2}}{x}.$$

¹⁰Intuitively, if I assign high belief now to the event $\theta = +1$, then the probability that I assign this event low belief in the future must be small.

Let $\varepsilon > 0$, and take $\beta < \sqrt{1 + \varepsilon} - 1$. Then by monotonicity of $G_+(-x)$ and a straightforward calculation,

$$\begin{aligned} G_+(-\ell_t^*) &> G_+(-(1 + \beta) \frac{2\sqrt{2}}{\sigma} \sqrt{\log t}) \\ &> \left[\frac{1}{(1 + \beta) \frac{2\sqrt{2}}{\sigma}} \right] \cdot \frac{t^{(1+\varepsilon)-(1+\beta)^2}}{\sqrt{\log t}} \cdot \frac{1}{t^{1+\varepsilon}} \\ &> \frac{1}{t^{1+\varepsilon}} \end{aligned}$$

for all sufficiently large t . From this, the claim follows immediately. \square

Proof of Theorem 5. Denote by C_t be the event that $a_\tau = +1$ for all $\tau < t$, and note that the event $T_1 = t$ is simply the intersection of C_t with the event that $a_t = -1$.

Let $\varepsilon > 0$. By Lemma 14 there is some $k' > 0$ such that for all t ,

$$\mathbb{P}_+(a_t = -1 \mid C_t) > \frac{k'}{t^{1+\varepsilon}}.$$

Now, put $\gamma = \mathbb{P}_+(a_\tau = +1 \text{ for all } \tau \geq 1)$, the probability that all agents take the correct action. By Lemma 13, $\gamma > 0$, so this provides a lower bound on the probability of the first $t - 1$ agents taking the correct action. Formally,

$$\mathbb{P}_+(C_t) \geq \mathbb{P}_+(a_\tau = +1 \text{ for all } \tau \geq 1) = \gamma.$$

Thus,

$$\begin{aligned} \mathbb{P}_+(T_1 = t) &= \mathbb{P}_+(a_t = -1, C_t) \\ &= \mathbb{P}_+(a_t = -1 \mid C_t) \cdot \mathbb{P}_+(C_t) \\ &\geq \frac{\gamma k'}{t^{1+\varepsilon}} \end{aligned}$$

for all t . \square

APPENDIX D. UPSETS AND RUNS

We recall a few definitions from Section 3.3. We say that there is an *upset* at time t if $a_{t-1} \neq a_t$. We denote by Ξ the random variable which assigns to each outcome the total number of upsets, and by Ξ_t the total number of upsets at times up to and including t . We say that there is a *run* of length m from t if $a_t = a_{t+1} = \dots = a_{t+m-1}$. Note that this definition does not preclude a run from being part of a longer run; we will refer to a run of finite length which is not strictly contained in

any other run as *maximal*. We say that a run from t is *good* if $a_t = +1$ and *bad* otherwise.

Notice that the number of maximal runs is exactly equal to the number of upsets. We use this observation now to show that the probability of having many maximal runs is very small, so that most of the probability is concentrated in the outcomes with few maximal runs.

Proof of Proposition 7. Denote by Υ the random variable which assigns to each outcome the number of finite maximal good runs it contains; note that with probability 1, Υ is finite.

By Lemma 13, there is a $\beta > 0$ such that for any $x \geq 0$, if $\ell_t = x$, then the probability that all agents from t on take the correct action is at least β . Formally,¹¹

$$\mathbb{P}_+(a_\tau = +1 \text{ for all } \tau \geq t \mid \ell_t = x) \geq \beta.$$

Thus, whenever $a_{t-1} = -1$ and $a_t = +1$ (or $t = 1$), the probability that there is exactly one more maximal good run is at most $1 - \beta$. It follows that for $n \geq 0$,

$$\mathbb{P}_+(\Upsilon = n + 1) \leq (1 - \beta)\mathbb{P}_+(\Upsilon = n)$$

and thus, for any $n \geq 0$,

$$\mathbb{P}_+(\Upsilon = n) \leq (1 - \beta)^n \mathbb{P}_+(\Upsilon = 0)$$

and so

$$\mathbb{P}_+(\Upsilon \geq n) \leq \frac{\mathbb{P}_+(\Upsilon = 0)}{\beta} \cdot (1 - \beta)^n.$$

Finally, since $\Upsilon = \lfloor \Xi/2 \rfloor$, we have for any n :

$$\mathbb{P}_+(\Xi \geq n) \leq \mathbb{P}_+(\Upsilon \geq \lfloor n/2 \rfloor) \leq c \cdot \gamma^n$$

where $c = \mathbb{P}_+(\Upsilon = 0)/\beta$ and $\gamma = (1 - \beta)^{\frac{1}{3}}$.

□

Whenever asymptotic learning occurs (that is, whenever the probability that all agents take the correct action from some point on is equal to 1), the total number of upsets is almost surely finite. In particular, the probability that Ξ_t is logarithmic in t tends to zero as t tends to infinity. Using Proposition 7, we can show that in fact this probability tends to 0 quickly:

Corollary 15. *Let c, γ be as in Proposition 7. Then*

$$\mathbb{P}(\Xi_t \geq -\frac{2.1}{\log \gamma} \log t) \leq c \cdot \frac{1}{t^{2.1}}.$$

¹¹We remind the reader that $\mathbb{P}_+(\cdot)$ is shorthand for $\mathbb{P}(\cdot \mid \theta = +1)$.

Proof.

$$\begin{aligned} \mathbb{P}(\Xi_t \geq -\frac{2.1}{\log \gamma} \log t) &\leq \mathbb{P}(\Xi \geq -\frac{2.1}{\log \gamma} \log t) \\ &\leq c \cdot \gamma^{-\frac{2.1}{\log \gamma} \log t} \\ &= c \cdot \frac{1}{t^{2.1}}. \end{aligned}$$

□

In fact, it is equally easy to show the same statement for exponents larger than 2.1, but this will suffice for our purposes.

One important consequence of Corollary 15 is that with high probability, there is at least one maximal run before time t which is long relative to t . Thus, much of the dynamics is controlled by what happens during long runs.

We previously analyzed only long runs that start at time 1, when the public log-likelihood ratio is equal to 0. If a long run starts at some public belief $\ell_t \neq 0$ then its evolution is different from the former case. However, if the run is long enough then the analysis above can still be applied. The following lemma states that if a run starts at some $\ell_t > 0$ then we can bound the future public belief from below using ℓ^* .

Lemma 16. *Suppose that G_- has a convex and differentiable left tail. Then there exists a $z > 0$ such that, if there is a good run of length s from t , then $\ell_{t+s} \geq \ell_{s-z}^*$.*

Proof. Let $u_+(x) = x + D_+(x)$. Then by (1), whenever agent t takes action $+1$, $\ell_{t+1} = u_+(\ell_t)$.

Since G_- is eventually convex and differentiable, $u_+(x)$ is monotone increasing for sufficiently large x , by Lemma 12. Take

$$z = \min \{t \in \mathbb{N} : u_+(x) \text{ is monotone on } (\ell_t^* - 1, \infty)\}.$$

Now, let $\mu = \inf_{x \in [0, \ell_z^*]} D_+(x)$. By continuity of $D_+(x)$ and compactness of $[0, \ell_z^*]$, $\mu > 0$, since $D_+(x) > 0$ for all x . Put $N = \lceil \frac{\ell_z^*}{\mu} \rceil$. Then for all $x \in [0, \ell_z^*]$, $u_+^N(x) \geq \mu \cdot N \geq \ell_z^*$. Further, since $u_+(x) > x$ for all x , it follows that whenever there is a run of length N from t , $\ell_{t+N} > \ell_z^*$.

This implies that if there is a good run from t of length $s \geq N$, then $\ell_{t+s} \geq \ell_{s-z}^*$.

□

APPENDIX E. DISTRIBUTIONS WITH POLYNOMIAL TAILS

In this appendix we prove Theorem 6, showing that for private log-likelihood distributions with polynomial tails, the expected time to learn is finite.

As in the setting of Theorem 6, assume that the conditional distributions of the private log-likelihood ratio satisfy

$$(9) \quad G_+(x) = 1 - \frac{c}{x^k} \text{ for all } x > x_0$$

$$(10) \quad G_-(x) = \frac{c}{(-x)^k} \text{ for all } x < -x_0$$

for some $x_0 > 0$.

We remind the reader that we denote by ℓ_t^* the log-likelihood ratio of the public belief that results when the first $t - 1$ agents take action +1. It follows from Theorem 3 that in this setting, ℓ_t^* behaves asymptotically as $t^{1/(k+1)}$. Notice also that, by the symmetry of the model, the log-likelihood ratio of the public belief that results when the first $t - 1$ agents take action -1 is $-\ell_t^*$.

We begin with the simple observation that a strong enough bound on the probability of mistake is sufficient to show that the expected time to learn is finite. Formally, we have the following lemma. We remind the reader that $\mathbb{P}_+(\cdot)$ is shorthand for $\mathbb{P}(\cdot | \theta = +1)$.

Lemma 17. *Suppose there exist $k, \varepsilon > 0$ such that for all $t \geq 1$, $\mathbb{P}_+(a_t = -1) < k \cdot \frac{1}{t^{2+\varepsilon}}$. Then $\mathbb{E}_+(T_L)$ is finite.*

Proof. Since $T_L = t$ only if $a_{t-1} = -1$, $\mathbb{P}_+(T_L = t) \leq \mathbb{P}_+(a_{t-1} = -1)$. Thus

$$\begin{aligned} \mathbb{E}_+(T_L) &= \sum_{t=1}^{\infty} t \cdot \mathbb{P}_+(T_L = t) \\ &\leq \mathbb{P}_+(T_L = 1) + \sum_{t=2}^{\infty} t \cdot \mathbb{P}_+(a_{t-1} = -1) \\ &\leq 1 + k \sum_{i=2}^{\infty} \frac{t}{(t-1)^{2+\varepsilon}} \\ &< \infty. \end{aligned}$$

□

Accordingly, this section will be primarily devoted to studying the rate of decay of the probability of mistake, $\mathbb{P}_+(a_t = -1)$. In order to bound this probability, we will need to make use of the following

lemmas, which give some control over how the public belief is updated following an upset.

Lemma 18. *For G_+ and G_- as in (9) and (10), $|\ell_{t+1}| \leq |\ell_t|$ whenever $|\ell_t|$ is sufficiently large and $a_t \neq a_{t+1}$.*

Proof. Assume without loss of generality that $a_t = +1$ and $a_{t+1} = -1$, so that

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

Thus, to prove the claim we compute a bound for D_- . To do so we first obtain a bound for the left tail of G_+ . By assumption, for $x > x_0$ (with x_0 as in (9) and (10)),

$$g_-(-x) = G'_-(-x) = \frac{ck}{x^{k+1}}$$

and so by (4),

$$g_+(-x) = e^{-x}g_-(-x) = ck \frac{e^{-x}}{x^{k+1}}.$$

Hence,

$$G_+(-x) = \int_{-\infty}^{-x} g_+(\zeta) d\zeta = \int_{-\infty}^{-x} ck \frac{e^\zeta}{(-\zeta)^{k+1}} d\zeta = ck \int_x^\infty \zeta^{-k-1} e^{-\zeta} d\zeta.$$

For ζ sufficiently large, ζ^{-k-1} is at least, say, $e^{-.1\zeta}$. Thus, for x sufficiently large,

$$G_+(-x) \geq ck \int_x^\infty e^{-.1\zeta} d\zeta = \frac{ck}{1.1} e^{-.1x}.$$

It follows that for x sufficiently large,

$$D_-(x) = \log \frac{G_+(-x)}{G_-(-x)} \geq \log \frac{ck}{1.1} - 1.1x + k \log x \geq -1.2x.$$

Thus, for ℓ_t sufficiently large,

$$\ell_{t+1} = \ell_t + D_-(\ell_t) = \ell_t + \log \frac{G_+(-\ell_t)}{G_-(-\ell_t)} \geq \ell_t + 1.2(-\ell_t) = -.2\ell_t$$

so in particular, $|\ell_{t+1}| < |\ell_t|$. □

We will make use of the following lemma, which bounds the range of possible values that ℓ_t can take.

Lemma 19. *For G_+ and G_- as in (9) and (10), there exists an $M > 0$ such that for all $t \geq 0$, $|\ell_s| \leq M \cdot \ell_t^*$ for all $s \leq t$.*

Proof. For each $\tau \geq 0$, define

$$M_\tau = \max \frac{|\ell_\tau|}{\ell_\tau^*}$$

where the maximum is taken over all outcomes. Note that there are at most 2^τ possible values for this expression, so M_τ is well-defined and finite. Put

$$M = \sup_{\tau \geq 0} M_\tau.$$

To establish the claim, we must show that M is finite. To do this, it suffices to show that for τ sufficiently large, $M_{\tau+1} \leq M_\tau$.

Now, let $u_+(x) = x + D_+(x)$ and $u_-(x) = x + D_-(x)$. Then as shown in the section about the model, whenever agent τ takes action $+1$, $\ell_{\tau+1} = u_+(\ell_\tau)$, and whenever agent τ takes action -1 , $\ell_{\tau+1} = u_-(\ell_\tau)$.

By Lemma 12, u_+ and u_- are eventually monotonic. Thus, there exists $x_0 > 0$ such that u_+ is monotone increasing on (x_0, ∞) and u_- is monotone decreasing on $(-\infty, -x_0)$.

For τ sufficiently large, $\ell_\tau^* > x_0$. Further, it follows from Lemma 18 that for τ sufficiently large, $|\ell_{\tau+1}| < |\ell_\tau|$ whenever $a_\tau \neq a_{\tau+1}$ and $|\ell_\tau| > |\ell_\tau^*|$. Let (a_τ) be any sequence of actions with $\frac{|\ell_{\tau+1}|}{\ell_{\tau+1}^*} = M_{\tau+1}$. If $a_\tau \neq a_{\tau+1}$

$$M_{\tau+1} = \frac{|\ell_{\tau+1}|}{\ell_{\tau+1}^*} \leq \frac{|\ell_\tau|}{\ell_\tau^*} \leq M_\tau.$$

If $a_\tau = a_{\tau+1}$, then either $M_{\tau+1} = 1$, in which case $M_{\tau+1} \leq M_\tau$, or $M_{\tau+1} > 1$. If $M_{\tau+1} > 1$, then since $|D_+|$ and $|D_-|$ are decreasing on (x_0, ∞) and $(-\infty, -x_0)$ respectively, $|\ell_{\tau+1} - \ell_\tau|/|\ell_\tau| \leq |\ell_{\tau+1}^* - \ell_\tau^*|/|\ell_\tau^*|$. So

$$M_{\tau+1} = \frac{|\ell_{\tau+1}|}{\ell_{\tau+1}^*} = \frac{|\ell_\tau| + |\ell_{\tau+1} - \ell_\tau|}{\ell_\tau^* + |\ell_{\tau+1}^* - \ell_\tau^*|}$$

where the second equality follows from the fact that ℓ_τ and $\ell_{\tau+1}$ have the same sign. Finally,

$$M_{\tau+1} = \frac{|\ell_\tau|}{\ell_\tau^*} \cdot \frac{1 + |\ell_{\tau+1} - \ell_\tau|/|\ell_\tau|}{1 + |\ell_{\tau+1}^* - \ell_\tau^*|/\ell_\tau^*} \leq \frac{|\ell_\tau|}{\ell_\tau^*} \leq M_\tau.$$

Thus, for all sufficiently large τ , $M_{\tau+1} \leq M_\tau$. □

Proposition 20. *There exists $\kappa > 0$ such that $\mathbb{P}_+(a_t = -1) < \kappa t^{-2.1}$ for all $t > 0$.*

Proof. Let $\beta = -2.1/\log \gamma$, where γ is as in Proposition 7. To carry out our analysis, we will divide the event that $a_t = -1$ into three disjoint events and bound each of them separately:

$$A = (a_t = -1) \text{ and } (\Xi_t > \beta \log t)$$

$$B_1 = (a_t = -1) \text{ and } (\Xi_t \leq \beta \log t) \text{ and } (|\{s: s < t, a_s = +1\}| \geq \frac{1}{2}t)$$

$$B_2 = (a_t = -1) \text{ and } (\Xi_t \leq \beta \log t) \text{ and } (|\{s: s < t, a_s = +1\}| < \frac{1}{2}t).$$

First, by Corollary 15 we have a bound for $\mathbb{P}_+(A)$

$$\mathbb{P}_+(A) \leq c \cdot \frac{1}{t^{2.1}}.$$

Next, we bound $\mathbb{P}_+(B_1)$. This is the event that the number of upsets so far is small and the majority of agents so far have taken the correct action.

Since there are at most $\beta \log t$ upsets, there are at most $\frac{1}{2}\beta \log t$ maximal good runs. Since, furthermore, there are at least $\frac{1}{2}t$ agents who take action +1, there is at least one maximal good run of length at least $t/(\beta \log t)$.

Thus, $\mathbb{P}_+(B_1)$ is bounded from above by the probability that there are some $s_1 < s_2 < t$ such that there is a good run of length $s_2 - s_1 \geq t/(\beta \log t)$ from s_1 and $a_{s_2} = -1$.

For fixed s_1, s_2 , denote by E_{s_1, s_2} the event that there is a good run of length $s_2 - s_1$ from s_1 . Denote by Γ_{s_1, s_2} the event $(E_{s_1, s_2}, a_{s_2} = -1)$. Then

$$\begin{aligned} \mathbb{P}_+(\Gamma_{s_1, s_2}) &= \mathbb{P}_+(a_{s_2} = -1 | E_{s_1, s_2}) \cdot \mathbb{P}_+(E_{s_1, s_2}) \\ &\leq \mathbb{P}_+(a_{s_2} = -1 | E_{s_1, s_2}). \end{aligned}$$

By Lemma 16, there exists a $z > 0$ such that E_{s_1, s_2} implies that $\ell_{s_2} \geq \ell_{s_2 - s_1 - z}^*$. Therefore,

$$\mathbb{P}_+(\Gamma_{s_1, s_2}) \leq G_+(-\ell_{s_2 - s_1 - z}^*).$$

Since for t sufficiently large $\ell_t^* > t^{\frac{1}{k+2}}$ and since $G_+(-x) \leq e^{-x}$ by (4),

$$\mathbb{P}_+(\Gamma_{s_1, s_2}) \leq e^{-\alpha(s_2 - s_1 - z)^{\frac{1}{k+2}}} \leq e^{-\alpha(t/(\beta \log t) - z)^{\frac{1}{k+2}}}.$$

To simplify, we further bound this last expression to arrive at, for some $c > 0$,

$$\mathbb{P}_+(\Gamma_{s_1, s_2}) \leq ce^{-t^{\frac{1}{k+3}}}$$

for all t . Since B_1 is covered by fewer than t^2 events of the form Γ_{s_1, s_2} (as s_1 and s_2 are less than t), it follows that

$$\mathbb{P}_+(B_1) < ct^2 e^{-t^{\frac{1}{k+3}}} < \frac{1}{t^{2.1}}$$

for all t large enough.

Finally we bound $\mathbb{P}_+(B_2)$. This is the event that the number of upsets so far is small and the majority of agents so far have taken the wrong action. As in B_1 , there is a maximal bad run of length at least $t/(\beta \log(t))$.

Denote by R the event that there is at least one bad run of length $t/(\beta \log(t))$ before time t and by R_s the event that agents s through $s+t/(\beta \log t) - 1$ take action -1 . Since B_2 is contained in R , and since R is contained in the union $\cup_{s=1}^t R_s$, we have that

$$\mathbb{P}_+(B_2) \leq \mathbb{P}_+(R) \leq \sum_{s=1}^t \mathbb{P}_+(R_s).$$

Taking the maximum of all the addends in the right hand side, we can further bound the probability of B_2 :

$$\mathbb{P}_+(B_2) \leq t \cdot \max_{1 \leq s \leq t} \mathbb{P}_+(R_s).$$

Conditioned on ℓ_s , the probability of R_s is

$$\mathbb{P}_+(R_s | \ell_s) = \prod_{r=s}^{s+t/(\beta \log t) - 1} G_+(-\ell_r).$$

By Lemma 19, there exists $M > 0$ such that $|\ell_r| \leq M\ell_t^*$, for all $r \leq t$. Therefore, since G_+ is monotone,

$$\mathbb{P}_+(R_s) \leq G_+(M\ell_t^*)^{t/(\beta \log t)}.$$

It follows that

$$\mathbb{P}_+(B_2) \leq t \cdot G_+(M\ell_t^*)^{t/(\beta \log t)}.$$

Since $G_+(x) = 1 - c \cdot x^{-k}$ for x large enough, and since ℓ_t^* is asymptotically at most $t^{1/(k+0.5)}$, we have that

$$\log G_+(M\ell_t^*) \leq -cM^{-k} \cdot t^{-k/(k+0.5)}.$$

Thus

$$\mathbb{P}_+(B_2) \leq t \cdot \exp(-cM^{-k} \cdot t^{1/(2k+1)}/(\beta \log t)) \leq t^{-2.1},$$

for all t large enough. This concludes the proof, because $\mathbb{P}_+(a_t = -1) = \mathbb{P}_+(A) + \mathbb{P}_+(B_1) + \mathbb{P}_+(B_2) \leq \kappa \frac{1}{t^{2.1}}$ for some constant κ . \square

Given this bound on the probability of mistakes, the proof of the main theorem of this section follows easily from Lemma 17.

Proof of Theorem 6. By Proposition 20, there exists $\kappa > 0$ such that $\mathbb{P}(a_t = -1 | \theta = +1) < \kappa \frac{1}{t^{2.1}}$ for all $t \geq 1$. Hence, by Lemma 17 $\mathbb{E}(T_L | \theta = +1) < \infty$. By a symmetric argument the same holds conditioned on $\theta = -1$. Thus, the expected time to learn is finite. \square

CALIFORNIA INSTITUTE OF TECHNOLOGY