

Learning through the Grapevine: the Impact of Message Mutation, Transmission Failure, and Deliberate Bias

Matthew O. Jackson, Suraj Malladi, and David McAdams *

Draft: June 2019

Abstract

We analyze whether learning is simply difficult or impossible when first hand information is relayed to agents through long chains of noisy person-to-person communication. Suppose noise only takes the form of random mutations and message dependent transmission failures. Agents can still learn but at the cost of accessing many more chains of communication. They learn fully when they have more than a certain threshold number of chains and learn nothing with fewer. On a positive note, this threshold remains unchanged when agents adopt a simpler, naive learning rule. In particular, both the frequency and content of their communications are informative about the state, but agents can safely restrict attention to the more important of these dimensions. Results change dramatically when there are biased agents who relay their preferred message despite what they hear. Learning is typically impossible in the presence of even an arbitrarily small fraction of biased agents. This can explain why people become stuck at potentially different priors, despite a significant body of primary evidence pointing to one answer (e.g., disagreement on the effects of vaccination or the reality of global warming). We show that a planner can recover partial learning by limiting the number of contacts to whom agents can pass along a given message, a policy some messaging platforms are starting to use.

JEL CLASSIFICATION CODES: D83, D85, L14, O12, Z13

KEYWORDS: Social Learning, Communication, Noise, Mutation, Bias, Fake News

*Jackson is from the Department of Economics, Stanford University, Stanford, California 94305-6072 USA, and is also an external faculty member at the Santa Fe Institute, and a fellow of CIFAR. Malladi is from the Graduate School of Business, Stanford University. McAdams is at the Fuqua School of Business and Economics Department, Duke University. Emails: jacksonm@stanford.edu, surajm@stanford.edu, david.mcadams@duke.edu. We gratefully acknowledge financial support under NSF grant SES-1629446 and from Microsoft Research New England. We thank Arun Chandrasekhar, Ben Golub, David Hirshleifer, Mallesh Pai, Evan Sadler, and Omer Tamuz for helpful conversations and suggestions.

1 Introduction

People rely on word-of-mouth learning when deciding whether to vaccinate their children, adopt a new diet, participate in a government program, adopt a new technology, support a challenger over an incumbent politician, etc. When does such word-of-mouth learning—interpreted as any form of relayed information—lead to efficient aggregation of information and informed decisions?

Francis Galton’s famous article “Vox Populi” (1907) showed that the information possessed by a group of people, when centrally aggregated, can be remarkably accurate. Galton examined 787 entries in a contest at the “West of England Fat Stock and Poultry Fair,” in which people guessed the dressed weight of an ox. Even though more than half of the guesses were off by more than 3 percent, the mean and median guesses were each within 0.1 percent of the ox’s true weight.

This sort of dispersed information can also be successfully aggregated in a decentralized manner through repeated communication in a social network, as people learn indirectly from neighbors of neighbors through the updated beliefs of their neighbors, and so on.¹

Under this view of the world, once sufficiently strong evidence is gathered in favor of a certain position, repeated communication should ensure that everyone’s beliefs converge to the truth. However, people remain strongly divided even on objective issues for which there is scientific consensus. For example, Largent (2012) estimates that 40 percent of American parents have delayed or denied recommended vaccinations for their children. Meanwhile, around twenty-seven percent of Americans do not accept that climate change is happening, and even more do not believe in human-caused climate change.² One explanation for this phenomenon is that primary evidence diffuses through the relevant population by means of noisy and potentially biased communication. This would dampen the informativeness of socially acquired information—people who start with different beliefs but do not consult primary sources of information may get stuck there.

Consider the children’s game of Telephone, in which a starting message is whispered from one player to the next. The final message typically bears little resemblance to the original message because of “mutations” that occur along the transmission chain.

Such mutations are not unique to a child’s game. Simmons, Adamic and Adar (2011) discuss a revealing example of message mutation. An initial tweet, “Street style shooting in Oxford Circus for ASOS and Diet Coke. Let me know if you’re around!”, was an invitation for people to join the crowd for a commercial being filmed in London. This was

¹Golub and Jackson (2010a) show that individuals can converge to accurate beliefs by repeatedly (weighted) averaging their beliefs with those of their neighbors, as long as the social network is balanced in a sense that no individual is unduly influential. Mueller-Frank (2014) shows that, if some individuals are more sophisticated, then learning can occur in a broader set of networks. For a review of the large literature on social learning that we do not survey here; see e.g., Golub and Sadler (2016).

²This figure is from the December 2018 “Climate Change in the American Mind” poll conducted jointly by the Yale Program on Climate Change Communication and the George Mason University Center for Climate Change Communication.

misunderstood and within minutes had mutated to “Shooting in progress in Oxford Circus? What?” and then retweeted as “Shooting in progress in Oxford Circus, stay safe people.” The informational content of the message completely changed.

Mutations can occur frequently. In Adamic, Lento, Adar and Ng (2016)’s study of online viral memes, one meme was reposted more than 470,000 times, with a mutation rate of around 11 percent and more than 100,000 variants. This was not an outlier in their analysis: 121 of the 123 most viral memes each had more than 100,000 variants. Other examples of mutating messages include mythology and the morphing of religious texts. Gurry (2016) estimates that there are around 500,000 textual variants of the Greek New Testament, not including spelling errors.

The information that reaches the receiver at the end of a transmission chain also depends on the likelihood that messages get transmitted, which itself may depend on message content. For instance, a person may be more likely to pass along information that they find surprising, or that is in line with their prior beliefs. Even in the publication of scientific articles, reviewers may be more likely to agree to publish (pass along) significant or surprising results than insignificant or expected ones.

In addition, some people may be ideologues, relaying messages that they prefer telling rather than what they heard (e.g., “fake news”).³ To give an example, Largent (2012) finds that parents who are skeptical of vaccinations have “tremendously high trust in medical communities.” But he writes “[who] don’t they trust? The feds, and pharma.”⁴ These are agents who participate in the process of diffusing scientific research and whom some parents worry may be ideologically motivated to misreport scientific research.

All three types of distortions—mutation, content-dependent transmission, and deliberate ideological bias—build up as relaying chains grow. Moreover, the paths that word-of-mouth communication follow can be quite long. Liben-Nowell and Kleinberg (2008) found instances of Internet chain letters that travelled median distances of over one hundred links.⁵ Adamic et al. (2016) examined hundreds of millions of instances of thousands of memes and found chains with lengths in the hundreds and typical distances well into the dozens.

Here, we investigate the consequences for learning when information reaches a learner over lengthy chains. In our model, information is relayed from its original source via a sequence of individuals to the eventual learner, who wishes to learn the state of the world. We code the state of the world and messages as either being in favor of an action (“1”, e.g., it is best to vaccinate a child) or against it (“0”, it is best not to vaccinate). With noiseless word-of-mouth communication and sufficiently many starting sources of conditionally independent information, the learner learns the true state. However, along each chain, the message may mutate, be dropped, or be deliberately biased—reducing the information content of the

³For other perspectives on the role of biased agents in the spread of false information, see Acemoglu, Ozdaglar and ParandehGheibi (2010) and Bloch, Demange and Kranton (2018).

⁴This was quoted from Mark Largent’s book, *Vaccine: The Debate in Modern America*, in *The Atlantic* article “Anti-Vaxers Aren’t Stupid” by Emma Green.

⁵ Golub and Jackson (2010b) explain why the resulting trees can be much longer than they are wide.

signals that reach the learner.

We characterize the (sharp) threshold number of independent chains that the learner needs in order to learn the true state, as chains grow long. We also contrast the implications of different sorts of transmission distortion on the difficulty in learning. And finally, we discuss a form of communication platform design that can partly overcome the difficulties presented by these distortions.

First, in a baseline model in which mutation and message dropping are the only sources of noise, we show that the number of independent original sources of information that a learner needs grows exponentially in the mutation rate, the dropping rate, and the length of chains over which messages travel. Thus, given the nontrivial empirically-observed mutation rates and chain lengths mentioned earlier, accurate learning may require observing many chains.

Second, when transmission/dropping likelihood depends on message content,⁶ we show that the learner can draw meaningful inferences from the number of messages received. To gain intuition, suppose that “surprising” messages are more likely to be shared. If the state is surprising, people will be more likely to share a message that matches the state, until a mutation occurs that changes the message’s surprisingness. This *relative* advantage from early stages is never erased: the *relative* likelihood a chain delivers a message to the learner is higher by a non-vanishing amount when the true state is surprising. This allows the learner to make inferences about the state from the number of messages received—even if she is unable to view or trust message content (which is likely to have mutated along the way).

We also bound how much more likely a Bayesian agent—who updates based on both message survival and content—is to guess the state compared to someone who looks only at message content or only at message survival. The Bayesian’s advantage vanishes as chain length increases; so, full Bayesian learning can be well-approximated by simple rules of thumb conditioned on just message survival or just message content. Thus, learning in the presence of mutations and transmission failures is not only possible, but may be easy as well—provided sufficiently many initial sources of information.

Third, we uncover a crucial difference between mutation and ideological bias. Mutations noise up a message sequence but even when chains are long enough that mutations are very likely to occur at some point, one can make probabilistic inferences about a sequence’s starting state and learn well from large numbers of sequences. In contrast, the information in a sequence that contains an ideologue is lost and cannot be recovered, even probabilistically. Learning from information passed along long chains is therefore only possible if ideologues can be identified and/or if their exact frequency is known. With even a small amount of uncertainty about ideologues’ location and relative bias for/against each state, learning is precluded. Indeed, uncertainty about the other parameters of the model, such as mutation

⁶In Banerjee (1993), a rumor’s survival rate increases with the number of people who have made an investment. Hearing a rumor therefore conveys information about how many people invested. The survival bias here comes from a different source, but there is still an inference to be made purely from information survival regardless of content.

rate or number of starting sources, would not produce the same effect. In the absence of biased agents, there would still be some partial learning in these cases.

Finally, we examine policies that offset the impact of ideologues on learning and reduce the frequency with which false messages are sent. When firsthand information can originate from anywhere in a network, one such policy is to limit the number of steps that a message can travel, which limits the chances it is distorted and allows partial learning for those near the original information. People see fewer messages, but ones that are more likely to be informative - and this increases the overall signal to noise ratio. If the distances that a message travels cannot be tracked, a cruder but still partially effective policy is to limit the number of other people to whom any given agent can forward a message. This also limits the number of messages flooding a network, and increases the relative frequency of messages that traveled short distances to longer distances. Thus received messages are less distorted. Interestingly, such policies have been adopted by online messaging platforms, like WhatsApp, for the express purpose of curbing the spread of false information.

2 The Base Model of Noisy Information Transmission

Information passes by “word-of-mouth” (oral or written) from an original source to “the learner”.

There are two possible states of the world, $\omega \in \{0, 1\}$ and the prior probability that the state is 1 is $\theta \in (0, 1)$.

A sequence of agents $\{1, 2, \dots, t\}$, referred to as a “chain,” successively relays a signal of the state via word of mouth, terminating with the learner at $t \geq 1$.

We do not model what the learner does with this information, but one can think of the learner preferring to match her action with the state. For instance, the learner may hear from friends about whether there is a link between vaccines and autism and then decide whether to vaccinate her child.

A first agent in a chain, interpreted as “an original source,” observes a noisy signal of the state, $s_1 \in \{0, 1, \emptyset\}$.⁷ That signal is transmitted with noise becoming $s_2 \in \{0, 1, \emptyset\}$, and so on until signal s_t reaches the learner.

The “null signal,” $s_\tau = \emptyset$, indicates that no signal was received, in which case no signal is transmitted. In particular, if agent $\tau \geq 1$ receives the null signal $s_\tau = \emptyset$, then all subsequent agents (including the learner) also receive the null signal.

If instead agent $\tau \geq 1$ receives a signal $s_\tau \in \{0, 1\}$, then that agent passes a signal along ($s_{\tau+1} \neq \emptyset$) with probability p if $s_\tau = 1$, and with probability q if $s_\tau = 0$. Thus, when $p > q$, agents are more likely to transmit a signal if they heard a 1, and vice versa if $p < q$. With the remaining probabilities of $1 - p$ and $1 - q$, respectively, the signal is dropped and $s_{\tau+1} = \emptyset$.

⁷We focus on a binary world to crystallize the main ideas. Extensions to richer state spaces and signal structures are left for future research.

Each time a non-null signal is transmitted, a mutation occurs with probability $\mu \in [0, 1/2]$, changing a non-empty signal from 0 to 1 or from 1 to 0. Bearing in mind that imperfections in “telling and listening” apply equally to any message, we assume that the likelihood of mutation $\mu \in [0, 1/2]$ does not depend on signal content (and defer bias in transmission to Section 5). Thus, conditional upon $s_{\tau+1} \neq \emptyset$, the probability that $s_\tau = s_{\tau+1}$ is $1 - \mu$ and the chance that $s_\tau \neq s_{\tau+1}$ is μ .

To keep equations as simple as possible, we focus on a stationary setting in which the initial signal s_1 is derived from the state in the same way, as if nature were “agent 0” in the chain with signal s_0 equal to the state.⁸

In summary, if $s_{t-1} = 1$, the next agent (including $t = 1$) hears $s_t = 1$ with probability $p(1 - \mu)$, $s_t = 0$ with probability $p\mu$, and $s_t = \emptyset$ with probability $1 - p$. Similarly, conditional on $s_{t-1} = 0$, $s_t = 1$ with probability $q\mu$, $s_t = 0$ with probability $q(1 - \mu)$, and $s_t = \emptyset$ with probability $1 - q$. If $s_\tau = \emptyset$ for some τ , then that is true for all subsequent signals. This defines a 3×3 Markov chain in which \emptyset is an absorbing state.

Most of our analysis presumes that the learner has access to some number $n \geq 1$ of chains, each of length t and along each of which conditionally independent signals of the state are independently relayed via the same noisy process to the same learner.

The base model described here allows for signal mutation and content-dependent signal dropping. In Section 5, we extend the model to include biased agents who distort signals.

3 Learning from Chains of Noisy Transmission

In Sections 3.1-3.3 we consider the case in which signals mutate, but their content does not affect the likelihood of further transmission ($p = q$). In Section 3.4 we allow for content-dependent signal survival ($p \neq q$).

3.1 The Rate of Decay due to Mutation

Lemma 1 characterizes the rate of information decay as a signal is passed along a single chain, which has a remarkably simple formula.

LEMMA 1 *Suppose that $p = q > 0$ and consider any mutation rate $\mu \in (0, 1/2)$. If agent $t \geq 1$ receives a non-null signal, it matches the true state with probability*

$$X(t) = \frac{1}{2} (1 + (1 - 2\mu)^t). \tag{1}$$

Note that $X(t) > 1/2$ for all t , $\lim_{t \rightarrow \infty} X(t) = 1/2$, and $X(t) - 1/2$ decreases exponentially at rate $1 - 2\mu$. Intuitively, the rate of decay, $1 - 2\mu = (1 - \mu) - \mu$, is how much more likely one is to get an unmutated signal than a mutated one from one period to the next.

⁸Our analysis is easily extended to settings in which first-signal accuracies and dropping rates differ from subsequent ones.

3.2 Learning (or Not) from Many Mutated Messages

We now characterize the threshold number of independent word-of-mouth chains that a Bayesian learner needs to access in order to have an accurate view of the true state.

Suppose that the learner has access to $n(t)$ independent chains, let $I_{n(t)}$ be the vector of (potentially null) random signals that the learner receives the chains, and let the random variable $b(t) = \Pr(\omega = 1 | I_{n(t)})$ be the posterior probability that the state equals 1 conditional on the signals. We index n by t since we wish to characterize how many chains are needed as a function of their length. Longer chains are more likely to be null or to have an incorrect signal and so more are needed to deliver an equivalent amount of information.

DEFINITION 1 (Threshold for learning) $\tau(t)$ is a threshold for learning if (i) $\text{Plim } b(t) = 1$ or 0 whenever $n(t)/\tau(t) \rightarrow \infty$ and (ii) $\text{Plim } b(t) = \theta$ whenever $n(t)/\tau(t) \rightarrow 0$.

Note that if $\text{Plim } b(t) = 1$ or 0, then Bayesian-updated beliefs are correct with a probability going to 1. Thus, a threshold for learning is sharp in that if the number of chains of signals is of higher order, then the receiver learns the true state with a probability going to one, while if it is of lower order, the receiver learns *nothing*.

PROPOSITION 1 If $p = q > 0$ and $\mu \in [0, 1/2)$,⁹ then $\frac{1}{p^t(1-2\mu)^{2t}}$ is a threshold for learning.

The proof of Proposition 1 is provided in the appendix but to build intuition we provide an informal argument here.

Each chain delivers a non-null message to the learner with probability p^t and, conditional on a signal being received, it is a Bernoulli random variable that matches the true state – is “true” – with probability $X(t) = \frac{1}{2}(1 + (1 - 2\mu)^t)$ by equation (1). If $m(t)$ non-null signals are received, the standard deviation of the fraction of true signals is

$$\left(\frac{X(t)(1 - X(t))}{m(t)} \right)^{1/2}.$$

The expected fraction of true signals is more than k standard deviations away from $1/2$ if

$$X(t) - k \left(\frac{X(t)(1 - X(t))}{m(t)} \right)^{1/2} \geq 1/2. \quad (2)$$

Inequality (2) – which roughly speaking ensures that the difference between the average received message and $1/2$ can be used to confidently infer the true state if k is large – is equivalent to

$$(1 - 2\mu)^{2t} m(t) \geq 4k^2 X(t)(1 - X(t)). \quad (3)$$

⁹If $\mu = 1/2$ then learning becomes impossible since all messages are independent of the starting state and the threshold diverges.

The right hand side in (3) converges to k^2 since $X(t)$ converges to $1/2$; so, for large t , (2) is approximately the same as

$$m(t) \geq \frac{k^2}{(1 - 2\mu)^{2t}}. \quad (4)$$

Finally, since a message is received from each chain with probability p^t ,

$$n(t) \geq \frac{k^2}{p^t(1 - 2\mu)^{2t}} \quad (5)$$

chains are needed in order to generate, on average, enough messages to satisfy (4). This is not a formal proof, but inequality (5) turns out to be a precise expression for the threshold for learning, as shown in the Appendix.

When inequality (5) holds for arbitrarily large k , the information contained in the set of received messages swamps the prior and allows the learner to perfectly discern the true state. If (5) does not hold for any positive k , then nothing can be learned.

3.3 Trees

The threshold number of word-of-mouth chains needed for learning translates into an exact degree cutoff when chains correspond to paths through a tree.

Suppose that the learner receives word-of-mouth signals through a random network; in particular, a random tree generated by a Galton-Watson branching process in which the ‘offspring’ distribution has strictly positive support.¹⁰ Initial signals about the state are received by leaf nodes, at distance t from the learner, and propagate toward the learner through the tree as independent transmission chains, as modeled in Section 2. An intermediate node hears multiple signals and transmits those as a vector to the next node, which each entry of the vector independently being subjected to the noise process described above.

For instance, leaf nodes could be a sales force in a firm who relay reports to their managers, who then collect the reports from their team and send them on to their managers. As each report flows up the organizational ladder, it becomes more likely to have mutated at least once along the way. The length of the chains over which word-of-mouth reports have to pass depends on the “flatness” of the firm’s organizational structure. See Figure 1, in which there are eight sources of information (and hence eight paths toward the learner) and three levels to the tree. In this context, Proposition 1 implies that, in a large firm, the top manager’s ability to learn from word-of-mouth messages depends on the expected degree of the graph corresponding to the firm’s organizational structure.

COROLLARY 1 *Suppose that word-of-mouth signals pass through a tree generated by a Galton-Watson branching process having a degree distribution with mean d , finite variance¹¹, and*

¹⁰This condition ensures that the tree does not die out and so has at least some paths of depth t with probability one. The analysis can be adapted to allow for extinction, but no new insight emerges.

¹¹The assumption of finite variance can be relaxed. If X is the degree of the offspring distribution, the finite variance assumption can be replaced with $E[X \log(X)] < \infty$.




Figure 1: The root node (“learner”) receives messages passed through eight paths, each starting from a different “information source” distance three from the learner. The absence of an arrow from one node to the one below it indicates that no message was delivered, a dashed arrow indicates the message was delivered but mutated, and a solid arrow indicates that the message was delivered unmutated. In the instance illustrated here: The true state is 1 and paths 1-3 and 6-8 begin with a correct initial signal, while path 4 begins with an incorrect initial signal and path 5 begins with no signal received. Initial messages are delivered on paths 1,2,4, and 6-8, mutating from 1 to 0 on path 1 and from 0 to 1 on path 4, and undelivered on paths 3 and 5. Messages are then relayed on path 2,4, and 6-8, mutating from 1 to 0 on path 6, but dropped on path 1. Finally, messages are re-relayed on paths 2 and 6-8, mutating from 0 to 1 on path 6, but dropped on path 4. Overall, the learner hears four messages, of which two never mutated, one mutated once, and one mutated twice.

strictly positive support. The root node learns perfectly ($\text{Plim } b(t) = 1$ or 0) if

$$d > \frac{1}{p(1-2\mu)^2}; \tag{6}$$

but the root node learns nothing ($\text{Plim } b(t) = \theta$) if

$$d < \frac{1}{p(1-2\mu)^2}.$$

To build intuition for this result, consider a (non-random) regular tree with degree d at each step. Since the number $n(t)$ of nodes distance t from the learner equals d^t , $n(t)$ exceeds the threshold for learning from Proposition 1 if and only if $d > \frac{1}{p(1-2\mu)^2}$. The proof handles the details associated with this being a random tree from a Galton-Watson process.

3.4 Learning from Message Survival

We now consider the case in which $p \neq q$, so that current signal content affects the likelihood that the signal is dropped along a word-of-mouth transmission chain. Without loss of generality, we focus on the case in which $p > q$, meaning that people are more likely to pass along signal 1 than 0.

The fact that a message has survived to reach the learner carries information about the original state. Moreover, unlike the content of a single message, which becomes nearly meaningless as chains grow long (due to mutation), the information conveyed by a single message's survival does not vanish in the long-chain limit.¹²

Let

$$z \equiv \frac{p}{q} \left(1 + (1-2\mu) \frac{(p-q)}{q + \mu(p-q)} \right).$$

When $p > q$ and $0 < \mu < 1/2$, this is strictly greater than $1/2$.

PROPOSITION 2 *Suppose that $1 \geq p > q > 0$ and $\mu \in (0, 1/2]$.*¹³

1. *The relative probability of message survival over a chain of length t conditional on state 1 versus state 0 is uniformly bounded away from $\frac{p}{q}$:*

$$\frac{\text{Pr}(s_t \neq \emptyset | \omega = 1)}{\text{Pr}(s_t \neq \emptyset | \omega = 0)} \geq z \geq \frac{p}{q} \text{ for all } t \geq 1, \tag{7}$$

with strict inequalities when $\mu < 1/2$.

¹²The amount of belief updating due to a single message's *non-survival* vanishes in the long-chain limit, but the amount due to message survival does not. What holds this together, of course, is the fact that the probability of survival itself vanishes in the long-chain limit.

¹³If $\mu = 0$ then $\frac{\text{Pr}(s_t \neq \emptyset | \omega = 1)}{\text{Pr}(s_t \neq \emptyset | \omega = 0)} = (p/q)^t$, which diverges, and the problem becomes trivial. Similarly, if $q = 0$ then $\text{Pr}(s_t \neq \emptyset | \omega = 0) = 0$ and the problem becomes trivial. Note that here we do not require that $\mu < 1/2$ since survival at the first step contains information, even if subsequent steps are completely random. This contrasts with the case in which $p = q$, in which learning is precluded when $\mu = 1/2$.

2. The ratio in (7) converges as chain-length grows: $y \equiv \lim_{t \rightarrow \infty} \frac{\Pr(s_t \neq \emptyset | \omega = 1)}{\Pr(s_t \neq \emptyset | \omega = 0)}$ exists.
3. Upon seeing a surviving message, the learner's updated belief $\Pr(\omega = 1 | s_t \neq \emptyset)$ is uniformly bounded below by $\frac{\theta}{\theta + (1-\theta)/z} > \theta$ and bounded above in the limit by $\frac{\theta}{\theta + (1-\theta)/y} < 1$.
4. In the limit, updating is entirely due to signal survival and not content: $\lim_{t \rightarrow \infty} \Pr(\omega = 1 | s_t = 1) = \lim_{t \rightarrow \infty} \Pr(\omega = 1 | s_t = 0) = \lim_{t \rightarrow \infty} \Pr(\omega = 1 | s_t \neq \emptyset)$.

The reason that signal survival is informative is that early survival affects the *relative* survival of a whole chain. To gain intuition, suppose for a moment that only the first agent in each chain was biased in favor of message 1, with other agents transmitting with probability \hat{p} regardless of signal content. The likelihood of survival to t is $p(\hat{p})^{t-1}$ if the first agent saw signal 1 or $q(\hat{p})^{t-1}$ if the first agent saw signal 0. Thus, the relative likelihood of survival equals $p/q > 1$ (favoring signal 1) no matter how long the chain. Moreover, biasing all agents in favor of transmitting message 1 further increases the relative likelihood of survival from state 1 since, by an extension of the reasoning in Lemma 1, signal 1 is more likely to be received at each step along the chain when the true state is 1 rather than 0.

As with learning from signal content only, the learner can discern the state from signal survival only, as long as transmission is more likely after one signal than the other ($p \neq q$) and there are sufficiently many starting sources of information.

PROPOSITION 3 *Suppose that the learner receives $n(t)$ signals along independent chains of length t , and that $\mu \in (0, 1/2]$ and $1 > p > q > 0$.¹⁴ There exists $\lambda(t) = c + o(1)$ for some $c \in (0, 1)$, such that a threshold for learning when conditioning only upon signal survival is*

$$\frac{1}{(p\lambda(t) + (1 - \lambda(t))q)^t}.$$

Given that messages mutate, the probability that any agent transmits a message lies somewhere between p and q . Conditional on the initial message being 1, the overall probability that a message is transmitted all the way to the end of a length- t chain must therefore take the form $(p\lambda(t) + (1 - \lambda(t))q)^t$ for some $\lambda(t) \in (0, 1)$. If the number of sequences $n(t)$ grows faster than this conditional survival likelihood decreases, then a growing number of signals survive conditional upon starting out as a 1. The learner can then discern the state (perfectly in the limit) based on the actual number of signals that survive. On the other hand, if the number of chains grows more slowly than $(p\lambda(t) + (1 - \lambda(t))q)^t$ decreases, then approximately zero messages are expected to survive in either state over long chains and approximately nothing is learned with probability going to one.

¹⁴If $\mu = 0$, then it is easy to check that the threshold is $1/p^t$, which is then the threshold for messages to survive conditional upon state $\omega = 1$, which are the more likely to survive.

4 Full Bayesian Learning vs Learning Only from Survival or Only from Content

In this section we provide a bound on how much more likely a Bayesian agent using both signal survival and message content is to guess the true state compared to someone who pays attention only to signal survival or only to signal content.

Without loss of generality, we focus on the case in which $p \geq q$.

Suppose for a moment that the learner has access to a single chain and needs to guess the state based on the signal $s_t \in \{0, 1, \emptyset\}$. We consider four different ways in which the learner might guess.

- A “Bayesian agent,” B , guesses the most likely state conditional on both signal survival and signal content.
- A “survival rule-of-thumb agent,” S , guesses 1 if a signal is received and guesses 0 if no signal is received.
- A “content rule-of-thumb agent,” C , guesses 1 if signal 1 is received, 0 if signal 0 is received, and guesses in favor of the prior if no message is received (flipping a coin if $\theta = 1/2$).
- A “naive agent,” N , always guesses in favor of the prior.

S, C, N are collectively referred to as “limited learners” since they make their guess based on less information than is available.

PROPOSITION 4 *Suppose that $1 \geq p \geq q \geq 0$ and $\mu \in [0, 1/2]$. The probability that a Bayesian agent is correct in guessing the state is at most $\frac{4}{3}$ higher than the best limited learner when $t = 1$, and at most $\frac{3}{2}$ higher than the best limited learner for all $t > 1$.¹⁵ Moreover, as t grows, this upper bound converges to 1.*

Proposition 4 implies that, when word-of-mouth chains are long, a belief-updating strategy that uses only message survival or only message content is approximately equivalent to one that uses all available information, no matter what the parameters and no matter what the realized state.¹⁶

Suppose now that the learner observes multiple chains. In this context, define “C” to be an agent who guesses 1 whenever the fraction of 1 messages compared to 0 messages is above a threshold, and define “S” to be the an agent who guesses 1 whenever the number of messages that survive is above or below a threshold. These thresholds are the conditional Bayesian ones, but these agents only consider one aspect of the information available.

¹⁵We conjecture that the bound is $\frac{4}{3}$ for any $t > 1$.

¹⁶This is obvious when $\mu = 1/2$, in which case message content contains no information, or when $p = q$, in which case message survival contains no information.

It is difficult to give tight bounds on the relative performance of the Bayesian agent and the best of the limited learners when there are many sequences. However, we establish limiting results. In particular, everywhere in the parameter space, the threshold for learning for agent B is the same as for one of the limited learners. Thus, there is no number of starting messages for which a Bayesian agent can learn but none of the naive agents can. Indeed, for large t full learning can be obtained from just one dimension, and we get the following result.

PROPOSITION 5 *For any $\theta, p, q \in [0, 1]$ and $\mu \in [0, \frac{1}{2}]$, the threshold for learning is the same for B as it is for the better of C or S.*

5 The Impossibility of Learning with Mutation and Deliberately Biased Agents

Finally, we consider the impact of ideologically-biased (“extremist”) agents on learning, starting, for simplicity, with the case in which $p = q$.

We consider situations in which each agent in the population is a “1-ideologue” (only sending message 1, no matter what message they received) with probability πf and a “0-ideologue” with probability $(1 - \pi)f$. $f < 1$ is each agent’s likelihood of being an ideologue and $\pi \in [0, 1]$ is the expected fraction of ideologues biased in favor of message 1.

Note that the presence of an ideologue kills the information content of a chain, as the same signal is forwarded regardless of the incoming signal.

If the learner knows ideologues’ identities, chains containing an ideologue can be ignored—resulting in the same amount of learning as if ideologues always drop their messages, i.e., as if each message is passed along with probability $\hat{p} = p(1 - f)$ rather than p . Alternatively, if messages are passed along a tree, one can think of removing ideologues entirely, reducing the expected degree of the tree from d to $\hat{d} = d(1 - f)$ and thereby eliminating all chains to which any ideologue belongs. The new threshold for learning is that a tree have expected degree of more than

$$\frac{1}{p(1 - f)(1 - 2\mu)^2}.$$

If the learner knows f and π exactly, but not the particular identities, the true state can still be learned in the limit when the learner has access to sufficiently many chains, with the above threshold being a necessary condition.

A much more challenging case is one in which there is some uncertainty about the relative frequency of the different types of ideologues. Even with arbitrarily small, but nonzero, uncertainty about π , all learning is precluded in the limit *regardless of the number of chains*.

PROPOSITION 6 *Consider a learner getting signals from $n(t)$ independent chains¹⁷ of length t , with dropping rates $p = q$, mutation rate $\mu \in [0, 1/2)$, and a fraction of ideologues f , with proportion π biased to 1 and $1 - \pi$ biased to 0. If the learner does not know anything about the location of ideologues¹⁸ but knows f exactly¹⁹ and has a prior on π that has an atomless continuous distribution with convex support (however concentrated), then for any $n(t)$ the receiver learns nothing in the limit; i.e., the learner’s posterior converges in probability to θ as t grows large.*

To gain intuition for Proposition 6, consider the special case with no mutation, $\mu = 0$, and no message dropping, $p = q = 1$. The final signal in a chain will match the state with probability $(1 - f)^t$, which tends to 0. Each contaminated signal is determined by the last ideologue in the chain, and hence equals 1 with probability π and 0 with probability $1 - \pi$. When the state is 1, a fraction $(1 - f)^t + \pi(1 - (1 - f)^t)$ of all signals equal 1, while if the true state is 0, a fraction $\pi(1 - (1 - f)^t)$ of signals equals 1. When π and f are exactly known, the vanishing difference $(1 - f)^t$ in expected signal content depending on the state allows the learner to discern the true state given enough signals. However, with any uncertainty about π , the vanishing difference of $(1 - f)^t$ is completely obscured by the non-vanishing overall base level of $\pi(1 - (1 - f)^t)$, which varies nontrivially. This logic holds no matter how many signals are received.

Proposition 6 can be viewed as an impossibility result, since it shows how learning may be impossible even when there is no signal mutation and the learner has good (but imperfect) knowledge of f and π . Even an arbitrarily small amount of uncertainty about how the relative frequency of different types of ideologues drowns the information in the noise.

Returning to the example of anti-vaxers given in the introduction, their concern about messages being deliberately distorted by biased parties like the feds or pharma reduces what they infer from messages. But the presence of such parties doesn’t necessarily preclude learning in our model, unless the parents are also unsure of how many messages were distorted by these groups.

This stands in stark contrast to unintentional mutations, which can still be unraveled probabilistically. Even though one does not know where mutations lie, one can figure out their frequency and make inferences about starting states.²⁰ In contrast, information about

¹⁷If messages are passed through a tree as modeled in Section 3.3, the learner can make inferences about which agents are biased based on the messages passed through them. We abstract from this inference problem here, implicitly assuming that each agent is part of only one word-of-mouth transmission chain. This simplifies the exposition and the proof of Proposition 6, but is not essential. As the proof makes clear, our main qualitative findings continue to hold if messages are passed through a tree.

¹⁸The result generalizes to a setting in which the learner knows the identity of some but not all ideologues, as long as the fraction of unknown ideologues does not go to zero.

¹⁹If there is uncertainty about both f and π , the result clearly extends.

²⁰If mutations were biased in one direction or the other - for instance more likely to turn a 1 into a 0 than the reverse, and that bias were unknown, then those inferences would also be clouded. Thus, one can also interpret “mutations” as unbiased distortions, and the deliberate changes as biased distortions.

the starting state in a sequence that contains at least one extremist is irrecoverable. Extremists render learning impossible, unless they can be identified and their sequences ignored.

Next, we comment on the case in which $p \neq q$ and both lie between 0 and 1. As t grows, then only a vanishing fraction of signals survive. When $p \neq q$, changing π slightly changes that fraction by orders of magnitude even though it will still be vanishing. Again, this crowds out the information about the original state that can be gleaned from survival, which dies out over the sequence.²¹ Thus, the difficulty with learning in the face of uncertain bias is not overcome by considering survival when $p \neq q$.

Finally, we note that other forms of uncertainty can also hamper learning. For example, agents may have uncertainty about the average degree in the network (tree). Again, as the observer is learning from a fractional difference in survival rates, having a larger order uncertainty over the overall number of sequences obscures the finer information needed to infer the state. This happens whether or not there are biased agents present. Interestingly, without any biased agents, uncertainty about the degree does not preclude learning from signal content, since that is based on relative frequencies and not how many sequences survive. Indeed, without biased agents, uncertainty over the rate of mutation does not completely crowd out learning either.

PROPOSITION 7 *Suppose the learner receives signals from $n(t)$ independent chains of length t , with dropping rates $p, q \in [0, 1)$ and mutation rate $\mu \in [0, \frac{1}{2})$. Moreover, suppose the learner knows p and q exactly, but only has a prior on μ with support bounded away from $\frac{1}{2}$. Then there exists $f(t)$ such that if $n(t)/f(t) \rightarrow \infty$, $\text{Plim } b(t) = 0$ or 1 .*

Unlike in the case of uncertainty over the composition of biased agents, uncertainty over the mutation rate does not preclude the possibility of learning as t grows large. Thus, there are subtle differences in which sorts of uncertainty crowd out which sorts of inference.

6 Overcoming Deliberate Bias by Restricting Communication

We have so far assumed that agents learn from messages that originated from distant sources in the network. However, agents were assumed to know how many times a message was

²¹In particular, to see this (wlog) consider a case in which $1 > p > q > 0$. Let $Z_\pi(t, s)$ be the probability that a signal survives t periods conditional on starting out as a signal s , given π . The key observation is that $Z_\pi(t, 0)/Z_{\pi-\varepsilon}(t, 1)$ grows without bound as t grows, for any ε . Both probabilities are tending to 0, but eventually they mix and the probabilities are strings of products of combinations of ps and qs , and tilting that combination one way or the other eventually accumulates arbitrarily in terms of *relative* probabilities as things are exponentiated. Even a small shift in the fraction of extremists completely overturns the advantage of the starting state. Then tiny uncertainty about π introduces much larger swings in the survival rates than the starting states.

relayed before reaching them, which helps them more accurately update beliefs. Here, we relax both of these assumptions and address an aspect of communication platform design.

In reality, people hear information originating close to them, such as from friends, friends of a friend. However they may not be able to distinguish whether messages originated from someone close by or far away. Although news articles or primary research papers identify their sources and authors, messages passed through more informal channels of communication rarely come with attributions to the originators. Moreover, even learning the identity of the sources of a message may not fully reveal how many times it was relayed before reaching people’s ears.

Consider a network in which information can originate from any node, and there are biased agents. Increasing the size of the network, for instance by increasing the average degree, can actually worsen the learning. As the degree grows, more messages come to a receiver from a greater distance. In the limit, this leads to almost all messages coming from an arbitrary distance, and the logic of Theorem 6 applies and learning is precluded. Limiting the distance that messages can travel can enable partial learning.

In particular, if distances that messages travel are unlimited and degree is large, then learning is precluded. At the other extreme in which distances are limited to be very short, then an agent does not have much information. If one can limit distances, then optimal learning occurs at in intermediate range. The formula of the precise distance that maximizes learning appears to be intractable as there is also learning about the distance a message traveled from its survival, and this depends upon the prior beliefs about π . However, it is possible to characterize a closely related policy that aims to choose the maximal distance that still ensures that messages are more likely to be true than false. This would be the policy of a platform that was interested in maximizing traffic, but subject to having a reputation that a majority of its information was correct. We explore this, and another related policy of limiting degree in this section.

We extend our model to allow any node to be an original source with probability $r > 0$. For simplicity, suppose that 0 and 1 messages survive with the same probability ($p = q > 0$) and that messages do not mutate ($\mu = 0$) (the arguments below extend to the more general case with more complex expressions). Agents are biased with known probability $f > 0$, and there is a fraction π of biased agents who are biased towards always forwarding 1’s, while the remainder always forward 0’s. The receiver does not know π and has a symmetric prior around $\pi = 1/2$. Without loss of generality, consider a realized $\pi \geq 1/2$.

In this environment, since (weakly) more biased agents tend to forward 1’s, received messages will always be more likely to reflect the true state when the state is 1. Messages are more likely to be incorrect in the other state, 0. For large t , messages which have only passed through unbiased agents are slightly more likely to match the state. However, conditioning on the messages that pass through at least one biased agent, the *last* biased agent through which the message passed is more likely to be biased towards 1 than 0 (since $\pi \geq 1/2$). Since there are no further mutations, the correctness of the message only depends on the direction of bias of this last biased agent along the chain.

By this argument, the probability that in state 0, a message that originated t steps away and reached the learner matches the true state is $(1 - f)^t + (1 - \pi)(1 - (1 - f)^t)$. With probability $\pi(1 - (1 - f)^t)$, the message will be false. Indeed, messages in state 0 will more often be true precisely when

$$\begin{aligned} (1 - f)^t + (1 - \pi)(1 - (1 - f)^t) &\geq \pi(1 - (1 - f)^t) \\ \iff \frac{1}{2\pi - 1}(1 - f)^t &\geq 1 - (1 - f)^t \\ \iff \frac{2\pi}{2\pi - 1}(1 - f)^t &\geq 1, \end{aligned}$$

and this last inequality holds precisely when

$$t \leq t^* \equiv \frac{\log(2\pi - 1) - \log(2\pi)}{\log(1 - f)} \quad (8)$$

Thus, only messages that originate sufficiently close to the learner are more likely to be true when the state is 0.

The example above suggests that one policy to reduce the frequency of false messages in a network is to limit the number of times a message (and its variants) can be forwarded. Indeed, social media companies such as WhatsApp and Facebook have recently been concerned with curbing the spread of misinformation over their platforms.

But following the life-cycle of a message can be difficult or infeasible to achieve in practice (e.g., for reasons of privacy, a designer may not be allowed to observe the messages being forwarded). In such cases, a designer can still cut down the number of long chains by restricting the number of others to whom any given agent can forward messages. We show how this works as follows.

Suppose that the receiver is in a tree of degree d . Let T be the largest integer smaller than t^* (which we take not to be an integer, which is generically true). Then the expected number of messages that come from nodes at a distance of less than t^* is

$$r \sum_{t=1}^T (pd)^t = r \frac{pd(1 - (pd)^T)}{1 - pd},$$

while the expected number of messages that come from nodes at a distance of more than t^* is

$$r \sum_{t=T+1}^{\infty} (pd)^t = r \frac{(pd)^{T+1}}{1 - pd}.$$

The closer messages are more likely to be correct than false, while the latter are all more likely to be false (conditional on state 0).

More messages are expected to be true than are expected to be false if and only

$$r \frac{pd(1 - (pd)^T)}{1 - pd} > r \frac{(pd)^{T+1}}{1 - pd}$$

which simplifies to

$$1 - (pd)^T > (pd)^T \quad \text{or} \quad pd < \left(\frac{1}{2}\right)^{1/T}.$$

Putting this all together, it is enough to limit d to

$$d < \frac{1}{p} \left(\frac{1}{2}\right)^{\frac{\log(1-f)}{\log(2\pi-1)-\log(2\pi)}}.$$

Our result above suggests that platforms can reduce the frequency of false communications by limiting how many people each person can spread information to. One effect of such a policy is that it enable agents to partially learn the true state in the presence of biased agents. Restricting communication creates a positive selection whereby relatively more of an agent’s received messages will have originated from nearby sources. Nearby messages have fewer distortions, and so this increases the signal to noise ratio even though it lowers the total number of messages that agents see.

Another very attractive feature of this policy is that it does not violate privacy and is more content neutral than having the planner directly censor messages agents send each other.

Indeed, cutting down on “fake news” has recently become an express objective for social media platforms who have been blamed for facilitating the spread of misinformation. For example, the messaging platform WhatsApp was criticized for allowing the spread of incendiary rumors leading to mob killings in India and for serving as a vehicle of misinformation in 2018 Brazilian elections. In line with the results of this section, WhatsApp, placed a cap of 5 on the number of people to whom a message can be forwarded in an attempt to curtail the spread of false information (see Hern and Safi (2019)).

7 Concluding Remarks

We introduced a benchmark model of social learning via relayed signals in the presence of signal mutation, dropping, and biasing. We showed that learning is challenging in the presence of mutations and dropped messages, in that it requires an exponentially growing number of original sources as the length of the chains over which information is relayed grows. We also showed that simpler information processing rules than full Bayesian updating suffice for making correct inferences. The presence of ideologues, however, renders learning impossible regardless of the number of chains observed, as it is difficult to imagine contexts where learners know the exact relative fraction of ideologues taking one or another side. In some cases, a central planner can partly overcome this issue and restore some learning by either optimally limiting the diameter of networks of information transmission (the number of steps that information can be relayed), or by implementing a second-best policy of restricting the number of people to whom someone can relay a message.

These challenges also naturally motivate learners to seek out information from closer, trusted contacts, and to down-weight or ignore more distantly-sourced information. However, learning is challenging in a network with varying lengths of sequences and with cycles that introduce interdependencies. Still, there may be partial learning since news sourced from nearby in the graph has lower chances of mutation, deliberate bias, and dropping. This gives a rationale for paying attention only to things for which one can directly trace and fact-check — short distances where the source of information has been retained.

References

- Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi**, “Spread of (mis) information in social networks,” *Games and Economic Behavior*, 2010, *70* (2), 194–227.
- Adamic, Lada A., Thomas M. Lento, Eytan Adar, and Pauline C. Ng**, “Information evolution in social networks,” in “Proceedings of the ninth ACM international conference on web search and data mining” ACM 2016, pp. 473–482.
- Banerjee, Abhijit V.**, “The economics of rumours,” *The Review of Economic Studies*, 1993, *60* (2), 309–327.
- Bloch, Francis, Gabrielle Demange, and Rachel Kranton**, “Rumors and social networks,” *International Economic Review*, 2018, *59* (2), 421–448.
- Durrett, Richard**, *Random graph dynamics*, Vol. 200:7, Cambridge university press Cambridge, 2007.
- Galton, Francis**, “Vox Populi,” *Nature*, 1907, *75*, 450–451.
- Golub, Benjamin and Evan Sadler**, “Learning in social networks,” *The Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. Rogers, Oxford University Press, 2016, pp. 504–542.
- Golub, Benjamin and Matthew O. Jackson**, “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2010, *2* (1), 112–149.
- and —, “Using selection bias to explain the observed structure of internet diffusions,” *Proceedings of the National Academy of Sciences*, 2010, *107* (24), 10833–10836.
- Gurry, Peter J.**, “The Number of Variants in the Greek New Testament: A Proposed Estimate,” *New Testament Studies*, 2016, *62* (1), 97–121.
- Harris, Theodore E.**, “Branching processes,” *The Annals of Mathematical Statistics*, 1948, pp. 474–494.

Hern, Alex and Michael Safi, “WhatsApp puts limit on message forwarding to fight fake news,” Jan 2019.

Largent, Mark A, *Vaccine: The debate in modern America*, JHU Press, 2012.

Liben-Nowell, David and Jon Kleinberg, “Tracing information flow on a global scale using Internet chain-letter data,” *Proceedings of the national academy of sciences*, 2008, 105 (12), 4633–4638.

Mueller-Frank, Manuel, “Does one Bayesian make a Difference?,” *Journal of Economic Theory*, 2014, 154, 423–452.

Simmons, Matthew P., Lada A. Adamic, and Eytan Adar, “Memes Online: Extracted, Subtracted, Injected, and Recollected.,” *icwsm (International Conference on Web and Social Media Proceedings)*, 2011, 11, 17–21.

Appendix

Proof of Lemma 1: The proof is by induction.

First, note that if $t = 1$ then this expression simplifies to $1 - \mu$, which is exactly the probability that the signal has not mutated, and so this holds for $t = 1$.

Then for the induction step, suppose that the claimed expression is correct for $t - 1$, we show it is correct for t .

The probability of an even number of mutations at t is the probability of an odd number at $t - 1$ times μ plus the probability of an even number at $t - 1$ times $1 - \mu$, which by the induction assumption can be written as

$$\left[1 - \frac{1}{2} (1 + (1 - 2\mu)^{t-1})\right] \mu + \left[\frac{1}{2} (1 + (1 - 2\mu)^{t-1})\right] (1 - \mu).$$

This can be rewritten as

$$\mu + \left[\frac{1}{2} (1 + (1 - 2\mu)^{t-1})\right] (1 - 2\mu),$$

or

$$\frac{1}{2} (1 + (1 - 2\mu)^t),$$

as claimed. ■

LEMMA 2 *Let $p = q$ and suppose a Bayesian agent with prior θ on the state being 1 observes $k > 0$ more of one type of signal than the other out of $n(t)$ independent sequences.*

If there are k more 1 signals, then the agent’s posterior that the state is 1 is

$$\frac{\theta X(t)^k}{\theta X(t)^k + (1 - \theta)(1 - X(t))^k}.$$

If there are k more 0 signals, then the agent's posterior that the state is 1 is

$$\frac{\theta(1 - X(t))^k}{\theta(1 - X(t))^k + (1 - \theta)X(t)^k}.$$

The proof is direct and omitted.

Proof of Proposition 1:

Case 1: $\mu = 0$. As long as any signal survives, it will be perfectly informative of the state. It is easily checked that $\frac{1}{p^t}$ is the threshold for survival of a signal, and so the result follows.

Case 2: $\mu \in (0, \frac{1}{2})$.

We first show that if $\frac{1}{p^t(1-2\mu)^{2t}} = o(n(t))$ then a Bayesian learner's beliefs will converge to the correct posterior of either 1 or 0 with a probability going to 1.

We first develop an expression for the variance of the fraction of correct signals, conditional on observing some signals.

Let $r(t, m) \sim \text{Bin}(m, X(t))$ be the number of correct signals given m received signals. Let $s(t, m) = \frac{r(t, m)}{m}$ be the share of signals received which are correct. Then

$$\begin{aligned} \text{var}(s(t, m)|m > 0) &= \text{var}\left(\frac{r(t, m)}{m} | m > 0\right) \\ &= E[\text{var}\left(\frac{r(t, m)}{m} | m\right) | m > 0] \\ &= E\left[\frac{1}{m^2} X(t)(1 - X(t))m | m > 0\right] \\ &= X(t)(1 - X(t))E\left[\frac{1}{m} | m > 0\right]. \end{aligned}$$

Next, we show that a Bayesian agent can infer the correct state with high probability as t grows large. It is sufficient to show that (1) $Pr(m > 0) \rightarrow 1$, and (2) for every k ,

$$E[s(t, m)|m > 0] - k\sqrt{\text{var}(s(t, m)|m > 0)} > \frac{1}{2}$$

for large enough t .

(2) implies, by Chebyshev's inequality, that the probability that the majority signal is correct approaches 1, conditional some signal surviving to the receiver. (1) says that at least some signal survives with probability approaching 1, so conditioning on this event poses no obstacle for showing the claim.

Note that (1) follows by standard arguments given that $\frac{1}{p^t(1-2\mu)^{2t}} = o(n(t))$, which implies that $\frac{1}{p^t} = o(n(t))$ (e.g., then Chernoff Bounds imply that the probability that $m(t) < p^t n(t)/2$ is less than $e^{-p^t n(t)/8}$ which goes to 0).

To show (2), first note that

$$E[s(t, m)|m > 0] - k\sqrt{\text{var}(s(t, m)|m > 0)} > \frac{1}{2} \iff X(t) - k\sqrt{X(t)(1 - X(t))E\left[\frac{1}{m} | m > 0\right]} > \frac{1}{2}$$

$$\begin{aligned} &\iff \frac{(1-2\mu)^t}{2} > k\sqrt{X(t)(1-X(t))E[\frac{1}{m}|m>0]} \\ &\iff \frac{(1-2\mu)^{2t}}{4k^2X(t)(1-X(t))} > E[\frac{1}{m}|m>0]. \end{aligned}$$

Since $4X(t)(1-X(t)) \rightarrow 1$, it suffices to show that $\frac{(1-2\mu)^{2t}}{k^2} > E[\frac{1}{m}|m>0]$ for any k and large enough t .

Let $\tau(t) = (1-\delta)n(t)p^t$ for some $\delta \in (0, 1)$. Then

$$\begin{aligned} E[\frac{1}{m}|m>0] &= Pr(m < \tau(t)|m>0)E[\frac{1}{m}|0 < m < \tau(t)] + Pr(m \geq \tau(t)|m>0)E[\frac{1}{m}|m \geq \tau(t)] \\ &< Pr(m < \tau(t)) + Pr(m \geq \tau(t)|m>0)\frac{1}{\tau(t)} \\ &< Pr(m < \tau(t)) + \frac{1}{\tau(t)} \\ &= Pr(n(t)p^t - \tau(t) < n(t)p^t - m) + \frac{1}{\tau(t)} \\ &\leq Pr((n(t)p^t - \tau(t))^2 < (n(t)p^t - m)^2) + \frac{1}{\tau(t)}, \text{ since } \tau(t) < n(t)p^t \\ &\leq \frac{\text{var}(m)}{(n(t)p^t - \tau(t))^2} + \frac{1}{\tau(t)}, \text{ by Markov's inequality} \\ &= \frac{n(t)p^t(1-p^t)}{(\delta n(t)p^t)^2} + \frac{1}{(1-\delta)n(t)p^t} \\ &= \Theta\left(\frac{1}{n(t)p^t}\right) \\ &= o((1-2\mu)^{2t}), \end{aligned}$$

completing the proof of (2). ²²

It remains to show that receiving too few signals precludes learning. Let $k(t)$ be the number of true signals minus the number of false signals received by the agent when $n(t) = o(\frac{1}{p^t(1-2\mu)^t})$. Then $E[k(t)] = n(t)p^t(2X(t) - 1) = n(t)p^t(1 - 2\mu)^t \rightarrow 0$ as $t \rightarrow \infty$. Moreover, the variance of $k(t)(1-2\mu)^t$, by similar calculations, tends to 0 as $t \rightarrow \infty$. Therefore $-2k(t)(1-2\mu)^t \rightarrow 0$ in probability. This in turn means (using that $\ln(1+x) = x + o(x^2)$) that $k(t)[\ln(1 - (1-2\mu)^t) - \ln(1 + (1-2\mu)^t)] \rightarrow 0$ so $(\frac{1-(1-2\mu)^t}{1+(1-2\mu)^t})^{k(t)} \rightarrow 1$. By Lemma 2, the agent's belief converges in probability to the prior as $t \rightarrow \infty$. ■

Proof of Corollary 1:

If the offspring distribution is degenerate, i.e., the tree is d -regular, the proof follows directly from Proposition 1 by noting that d is the t -th root of $n(t)$.

So, consider the general case in which d is the mean of the offspring distribution and let $d^* \equiv \frac{1}{p(1-2\mu)^2} > 1$ (the case where $d^* = 1$ is degenerate in which all streams survive

²² If $f(t) = \Theta(g(t))$, then there exist constants $0 < c < C$ such that $cg(t) \leq f(t) \leq Cg(t)$ for all sufficiently large t .

and do not mutate, in which case the result is direct). By the previous proposition and Lemma 2, it suffices to show that when $d > d^*$, the number of source nodes asymptotically dominates $\frac{1}{p^t(1-2\mu)^{2t}}$ with probability going to 1; and when $d < d^*$, the number of source nodes is $o(\frac{1}{p^t(1-2\mu)^{2t}})$ with probability going to 1. We show the former claim and henceforth assume $d > d^*$, since the proof of the latter follows similar steps.

Let $g(t)$ be the number of information sources, or leaves, in a depth t tree. Let $\sigma^2 < 0$ be the variance of the offspring distribution. Given that $d > d^* > 1$, it follows from standard arguments that $E[g(t)] = d^t$ (e.g., this can be shown by induction as in Harris (1948)).

First, note that there exists a random variable W with c.d.f F such that $\frac{g(t)}{d^t}$ converges to W in distribution by the martingale convergence theorem.

Next, note that under the assumption that the offspring distribution has strictly positive support the tree does not have any extinction. Given this, it can be shown that W has weight 0 on 0. This can be deduced from Harris (1948).²³ Therefore, since $(d^*)^t = o(d^t)$, it follows that $(d^*)^t = \frac{1}{p^t(1-2\mu)^{2t}} = o_p(g(t))$.²⁴ That is to say, the number of source nodes asymptotically dominates $\frac{1}{p^t(1-2\mu)^{2t}}$ with high probability. ■

Proof of Proposition 2, Part 1:

For ease of notation, let $P_{1S}^t \equiv Pr(s_t \neq \emptyset | \omega = 1)$ and $P_{0S}^t \equiv Pr(s_t \neq \emptyset | \omega = 0)$. These are the probabilities of signal survival to time t conditional on the first signal.

First we prove that $\frac{P_{1S}^t}{P_{0S}^t} \geq \frac{p}{q}$, with strict inequality when $\mu < 1/2$ and $t > 1$.

This is proven by induction. First, $P_{1S}^1 = p > q = P_{0S}^1$. Next, let us show that $\frac{P_{1S}^t}{P_{0S}^t} \geq \frac{p}{q}$ given that $P_{1S}^{t-1} > P_{0S}^{t-1}$. Note that given the stationarity of the process, $P_{1S}^{t-1} = Pr(s_t \neq \emptyset | s_1 = 1)$ and $P_{0S}^{t-1} = Pr(s_t \neq \emptyset | s_1 = 0)$, and then we can write²⁵ The first part

$$P_{1S}^t = p[(1 - \mu)Pr(s_t \neq \emptyset | s_1 = 1) + \mu Pr(s_t \neq \emptyset | s_1 = 0)],$$

and so then it follows that

$$P_{1S}^t = p(1 - \mu)P_{1S}^{t-1} + p\mu P_{0S}^{t-1}.$$

Then by the inductive step ($P_{1S}^{t-1} > P_{0S}^{t-1}$) and so it follows that

$$P_{1S}^t \geq p(1 - \mu)P_{0S}^{t-1} + p\mu P_{1S}^{t-1},$$

with strict inequality when $\mu < 1/2$ and $t > 1$. Similarly,

$$P_{0S}^t = q(1 - \mu)P_{0S}^{t-1} + q\mu P_{1S}^{t-1}.$$

²³For instance, see the third paragraph on page 477; or see Theorem 2.1.7 in Durrett (2007).

²⁴For a sequence of random variables X_t , and a deterministic function f , $f(t) = o_p(X_t)$ means for any $\varepsilon > 0$, $Pr(\frac{f(t)}{X_t} > \varepsilon) \rightarrow 0$

²⁵The starting state s_0 is 1 in this calculation and so then there is a probability p that the signal survives to the first period, and then the calculation inside the $[\cdot]$ handles the two possible values of the first period signal and then the probability the signal survives to t if it has made it to the first period in the two possible values it could have in the first period.

Therefore

$$\frac{P_{1S}^t}{P_{0S}^t} \geq \left(\frac{p}{q}\right) \frac{(1-\mu)P_{0S}^{t-1} + \mu P_{1S}^{t-1}}{(1-\mu)P_{0S}^{t-1} + \mu P_{1S}^{t-1}} = \frac{p}{q},$$

with strict inequality when $\mu < 1/2$ and $t > 1$, as claimed.

Now we complete the proof of the first part of the proposition. Note that (from above)

$$\frac{P_{1S}^t}{P_{0S}^t} = \left(\frac{p}{q}\right) \frac{(1-\mu)P_{1S}^{t-1} + \mu P_{0S}^{t-1}}{(1-\mu)P_{0S}^{t-1} + \mu P_{1S}^{t-1}}.$$

Therefore,

$$\frac{P_{1S}^t}{P_{0S}^t} = \left(\frac{p}{q}\right) \left(\frac{(1-\mu)P_{0S}^{t-1} + \mu P_{1S}^{t-1} + (1-2\mu)(P_{1S}^{t-1} - P_{0S}^{t-1})}{(1-\mu)P_{0S}^{t-1} + \mu P_{1S}^{t-1}} \right),$$

and then since $p > q$ and $\frac{P_{1S}^{t-1}}{P_{0S}^{t-1}} \geq \frac{p}{q}$, with strict inequality when $\mu < 1/2$ and $t > 1$, it follows that

$$\frac{P_{1S}^t}{P_{0S}^t} = \left(\frac{p}{q}\right) \left(1 + (1-2\mu) \frac{P_{1S}^{t-1} - P_{0S}^{t-1}}{P_{0S}^{t-1} + \mu(P_{1S}^{t-1} - P_{0S}^{t-1})} \right) \geq \frac{p}{q} \left(1 + (1-2\mu) \frac{(p-q)}{q + \mu(p-q)} \right),$$

with strict inequality when $\mu < 1/2$ and $t > 1$ (and it directly follows that this expression (z) is strictly larger than p/q when $\mu < 1/2$), as claimed. ■

The following lemma is useful in the proofs of the remaining parts of the proposition.

LEMMA 3 Fix $\theta \in (0, 1)$, $\mu \in (0, 1/2]$, $0 < q \leq p \leq 1$. For all $t > 0$,

$$Pr(s_t = 1|\omega = 1) \geq Pr(s_t = 0|\omega = 1).$$

Moreover, either there exists T large enough such that

$$Pr(s_t = 1|\omega = 0) \geq Pr(s_t = 0|\omega = 0) \quad \text{for all } t \geq T,$$

or

$$Pr(s_t = 1|\omega = 0) < Pr(s_t = 0|\omega = 0) \quad \text{for all } t.$$

Finally, the sequence

$$\frac{Pr(s_t = 1|\omega = 1)}{\min\{Pr(s_t = 1|\omega = 0), Pr(s_t = 0|\omega = 0)\}}$$

is bounded above.

Proof of Lemma 3:

The first claim is proven by induction:

Since $\mu \leq 1/2$, $Pr(s_1 = 1|\omega = 1) \geq Pr(s_1 = 0|\omega = 1)$. Suppose $Pr(s_t = 1|\omega = 1) \geq Pr(s_t = 0|\omega = 1)$. Note that,

$$\begin{aligned} Pr(s_{t+1} = 1|\omega = 1) &= p(1 - \mu)Pr(s_t = 1|\omega = 1) + q\mu Pr(s_t = 0|\omega = 1) \\ Pr(s_{t+1} = 0|\omega = 1) &= p\mu Pr(s_t = 1|\omega = 1) + q(1 - \mu)Pr(s_t = 0|\omega = 1). \end{aligned}$$

The result then follows from the inductive hypothesis and the facts that $p \geq q$ and $\mu \leq 1/2$.

Next, to show the second claim in the lemma, note that

$$\begin{aligned} Pr(s_{t+1} = 1|\omega = 0) &= p(1 - \mu)Pr(s_t = 1|\omega = 0) + q\mu Pr(s_t = 0|\omega = 0) \\ Pr(s_{t+1} = 0|\omega = 0) &= p\mu Pr(s_t = 1|\omega = 0) + q(1 - \mu)Pr(s_t = 0|\omega = 0). \end{aligned}$$

Then if $Pr(s_t = 1|\omega = 0) \geq Pr(s_t = 0|\omega = 0)$ for some $t = T$, the same will hold for all $t > T$ by a similar inductive proof. Otherwise $Pr(s_t = 1|\omega = 0) < Pr(s_t = 0|\omega = 0)$ for all t , and then the result holds directly.

Finally, we show the third part of the claim. By the second part of this lemma, there are two cases to consider. If $Pr(s_t = 1|\omega = 0) < Pr(s_t = 0|\omega = 0)$ for all t . Then

$$\frac{Pr(s_t = 1|\omega = 1)}{\min\{Pr(s_t = 1|\omega = 0), Pr(s_t = 0|\omega = 0)\}} = \frac{Pr(s_t = 1|\omega = 1)}{Pr(s_t = 1|\omega = 0)}$$

If instead there is a T such that for all $t \geq T$, $Pr(s_t = 1|\omega = 0) \geq Pr(s_t = 0|\omega = 0)$, then

$$\begin{aligned} \frac{Pr(s_t = 1|\omega = 1)}{\min\{Pr(s_t = 1|\omega = 0), Pr(s_t = 0|\omega = 0)\}} &= \frac{Pr(s_t = 1|\omega = 1)}{Pr(s_t = 0|\omega = 0)} \\ &= \frac{p(1 - \mu)Pr(s_{t-1} = 1|\omega = 1) + q\mu Pr(s_{t-1} = 0|\omega = 1)}{p\mu Pr(s_{t-1} = 1|\omega = 0) + q(1 - \mu)Pr(s_{t-1} = 0|\omega = 0)} \\ &\leq \frac{(p(1 - \mu) + q\mu)Pr(s_{t-1} = 1|\omega = 1)}{p\mu Pr(s_{t-1} = 1|\omega = 0) + q(1 - \mu)Pr(s_{t-1} = 0|\omega = 0)} \\ &< \frac{p(1 - \mu) + q\mu}{p\mu} \frac{Pr(s_{t-1} = 1|\omega = 1)}{Pr(s_{t-1} = 1|\omega = 0)}, \end{aligned}$$

where the second to last inequality uses the first part of this lemma. We can therefore handle both cases simultaneously by showing that the sequence $\frac{Pr(s_t=1|\omega=1)}{Pr(s_t=1|\omega=0)}$ is bounded above.

To that end, note that

$$\begin{aligned} Pr(s_t = 1|\omega = 0) &\geq Pr(s_t = 1|\omega = 0, s_1 = 1)Pr(s_1 = 1|\omega = 0) \\ &= Pr(s_{t-1} = 1|\omega = 1)q\mu. \end{aligned}$$

So,

$$\frac{Pr(s_t = 1|\omega = 1)}{Pr(s_t = 1|\omega = 0)} \leq \frac{Pr(s_t = 1|\omega = 1)}{Pr(s_{t-1} = 1|\omega = 1)} \frac{1}{q\mu}.$$

It then suffices to show that $Pr(s_t = 1|\omega = 1) \leq Pr(s_{t-1} = 1|\omega = 1)$, since then from above

$$\frac{Pr(s_t = 1|\omega = 1)}{Pr(s_t = 1|\omega = 0)} \leq \frac{1}{q\mu},$$

which is finite given that $q > 0$ and $\mu > 0$. To see that $Pr(s_t = 1|\omega = 1) \leq Pr(s_{t-1} = 1|\omega = 1)$,

$$\begin{aligned} Pr(s_t = 1|\omega = 1) &= p(1 - \mu)Pr(s_t = 1|s_1 = 1) + p\mu Pr(s_t = 1|s_1 = 0) \\ &= p(1 - \mu)Pr(s_{t-1} = 1|\omega = 1) + p\mu Pr(s_{t-1} = 1|\omega = 0) \\ &\leq p(1 - \mu)Pr(s_{t-1} = 1|\omega = 1) + p\mu Pr(s_{t-1} = 1|\omega = 1) \\ &= pPr(s_{t-1} = 1|\omega = 1), \end{aligned}$$

where the inequality follows from the first part of the lemma, establishing the claim. ■

Proof of Proposition 2, Part 2:

We show that $\lim_{t \rightarrow \infty} \frac{P_{1S}^t}{P_{0S}^t} = \lim_{t \rightarrow \infty} \frac{pPr(s_{t-1}=1|\omega=1)+qPr(s_{t-1}=0|\omega=1)}{pPr(s_{t-1}=1|\omega=0)+qPr(s_{t-1}=0|\omega=0)}$ exists.

The sequence is bounded above by the first and last part of Lemma 3: it is bounded above by either $\frac{Pr(s_t=1|\omega=1)}{Pr(s_t=1|\omega=0)}$ or $\frac{Pr(s_t=1|\omega=1)}{Pr(s_t=0|\omega=0)}$, both of which are bounded above. Furthermore, the sequence is bounded below by the first part of Proposition 2.

To complete the proof that the limit exists, we show that the sequence is monotone. For this, we will start by writing, r_t , the t^{th} term in the sequence, as $\frac{Pr(s_{t-1}=1|\omega=1)+\ell_1 Pr(s_{t-1}=0|\omega=1)}{Pr(s_{t-1}=1|\omega=0)+\ell_1 Pr(s_{t-1}=0|\omega=0)}$, where $\ell_1 = q/p$. Now the $t + 1^{\text{st}}$ is

$$\begin{aligned} r_{t+1} &= \frac{Pr(s_t = 1|\omega = 1) + \ell_1 Pr(s_t = 0|\omega = 1)}{Pr(s_t = 1|\omega = 0) + \ell_1 Pr(s_t = 0|\omega = 0)} \\ &= \frac{(p(1 - \mu) + \ell_1 p\mu)Pr(s_t = 1|s_1 = 1) + (q\mu + \ell_1 q(1 - \mu))Pr(s_t = 0|s_1 = 1)}{(p(1 - \mu) + \ell_1 p\mu)Pr(s_t = 1|s_1 = 0) + (q\mu + \ell_1 q(1 - \mu))Pr(s_t = 0|s_1 = 0)} \\ &= \frac{Pr(s_{t-1} = 1|\omega = 1) + \ell_2 Pr(s_{t-1} = 0|\omega = 1)}{Pr(s_{t-1} = 1|\omega = 0) + \ell_2 Pr(s_{t-1} = 0|\omega = 0)}, \end{aligned}$$

where $\ell_2 = \frac{q\mu + \ell_1(1-\mu)}{p(1-\mu) + \ell_1\mu}$. Consider the sequence ℓ_t , where $\ell_{t+1} = \frac{q\mu + \ell_t(1-\mu)}{p(1-\mu) + \ell_t\mu}$ and $\ell_1 = \frac{q}{p}$. Note that ℓ_t is non-decreasing in t given that $\mu \leq 1/2$ and it is strictly increasing when $\mu < 1/2$. Iterating on the above logic

$$r_t = \frac{Pr(s_1 = 1|\omega = 1) + \ell_{t-1} Pr(s_1 = 0|\omega = 1)}{Pr(s_1 = 1|\omega = 0) + \ell_{t-1} Pr(s_1 = 0|\omega = 0)}.$$

To see that r_t is monotone in t , note that the sign of the derivative of r_t with respect to ℓ_t only depends on the sign of $Pr(s_1 = 0|\omega = 1)Pr(s_1 = 1|\omega = 0) - Pr(s_1 = 1|\omega = 1)Pr(s_1 = 0|\omega = 0)$, and so it is monotone given the monotonicity of ℓ_t in t . ■

Proof of Proposition 2, Part 3:

That $Pr(\omega = 1|s_t \neq \emptyset) \geq \frac{\theta z}{1 + \theta(z-1)}$ for any $t > 1$, with strict inequality when $\mu < 1/2$, follows from Part 1 and Bayes rule (and it is evident from the proof that this lower bound is not tight). Therefore, it remains to show that $\lim_{t \rightarrow \infty} Pr(\omega = 1|s_t \neq \emptyset)$ exists, a step which is deferred to the proof of Part 4.

The fact that $\lim_{t \rightarrow \infty} Pr(\omega = 1|s_t \neq \emptyset) = \frac{\theta}{\theta + (1-\theta)/y} < 1$ follows from Part 2 and Bayes' Rule. ■

Proof of Proposition 2, Part 4:

It suffices to show that $\lim_{t \rightarrow \infty} Pr(\omega = 1 | s_t = 1) = \lim_{t \rightarrow \infty} Pr(\omega = 1 | s_t = 0)$, as this implies that $\lim_{t \rightarrow \infty} Pr(\omega = 1 | s_t \neq \emptyset)$ exists and has the same value. This limiting equality between posterior distributions can equivalently be expressed in terms of likelihood ratios:

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{Pr(s_t = 1 | \omega = 0)}{Pr(s_t = 1 | \omega = 1)} &= \lim_{t \rightarrow \infty} \frac{Pr(s_t = 0 | \omega = 0)}{Pr(s_t = 0 | \omega = 1)} \\ \iff \lim_{t \rightarrow \infty} \frac{Pr(s_t = 1 | s_t \neq \emptyset, \omega = 0)}{Pr(s_t = 1 | s_t \neq \emptyset, \omega = 1)} &= \lim_{t \rightarrow \infty} \frac{Pr(s_t = 0 | s_t \neq \emptyset, \omega = 0)}{Pr(s_t = 0 | s_t \neq \emptyset, \omega = 1)}. \end{aligned} \quad (9)$$

We show that

$$\lim_{t \rightarrow \infty} Pr(s_t = 1 | s_t \neq \emptyset, \omega = 0) = \lim_{t \rightarrow \infty} Pr(s_t = 1 | s_t \neq \emptyset, \omega = 1), \quad (10)$$

since this implies that both sides of equation 9 are equal to 1.²⁶

Denote by S a sequence of signals that evolve according to our process, starting with $S_0 = 1$ and S' another (independent) sequence of signals with $S'_0 = 0$. Let $\tau = \min\{t | S'_t = 1\}$, where $\tau = \infty$ if S' is dropped at some step before mutating to signal 1, or if $S'_t = 0$ for all t .

In this notation, equation 10 can equivalently be expressed as: $\lim_{t \rightarrow \infty} Pr(S_t = 1 | S_t \neq \emptyset) = \lim_{t \rightarrow \infty} Pr(S'_t = 1 | S'_t \neq \emptyset)$. Note the following relationship between the two independent paths:²⁷

$$\begin{aligned} Pr(S'_t = 1 | S'_t \neq \emptyset) &= \sum_{i=1}^t Pr(S'_t = 1 | S'_t \neq \emptyset, \tau = i) Pr(\tau = i | S'_t \neq \emptyset) \\ &= \sum_{i=1}^t Pr(S_{t-i} = 1 | S_{t-i} \neq \emptyset) Pr(\tau = i | S'_t \neq \emptyset) \\ &\equiv \left(\sum_{i=1}^t Pr(S_{t-i} = 1 | S_{t-i} \neq \emptyset) w_i^t \right), \end{aligned} \quad (11)$$

where $w_i^t = Pr(\tau = i | S'_t \neq \emptyset)$.

The result then follows from the following three facts, to be proved:

1. For any $\varepsilon > 0$ and positive integer k , for all sufficiently large t , $\sum_{i=t-k}^t w_i^t < \varepsilon$.
2. $\lim_{t \rightarrow \infty} Pr(S_t = 1 | S_t \neq \emptyset)$ exists.
3. $\sum_{i=1}^t w_i^t + w_\infty^t = 1$. Moreover, $w_\infty^t \rightarrow 0$ as $t \rightarrow \infty$, i.e., the probability that the signal never mutated conditional on survival to t goes to 0 as t grows.

²⁶Subtract each side of equation 10 from 1 before taking ratios to see that the right side of equation 9 is also 1.

²⁷Note that if $\tau > t$, then the probability that $S'_t = 1$ is 0.

To see that these facts imply the result, note that by fact 1, most of the weight falls on the first $t - k$ terms of the sum in equation 11 for large enough t . By fact 2, for a large enough k (growing slower than t), these first $t - k$ terms will be close to $\lim_{t \rightarrow \infty} Pr(S_t = 1 | S_t \neq \emptyset)$, and therefore by fact 3 the limiting weighted sum of these terms converges to this value as well.

Fact 3 is clear, so we prove the other two.

First we prove fact 1. Note that $q^i(1 - \mu)^{i-1}\mu$ is the probability of survival with no mutation through $i - 1$ and then survival with mutation at $t = i$, i.e., $Pr(\tau = i) = q^i(1 - \mu)^{i-1}\mu$. Second, let m_i be number of mutations through time i . Obviously, $Pr(S'_i \neq \emptyset) > Pr(S'_i \neq \emptyset \text{ and } m_i = 1)$. Third, if survival were always at rate q , then $Pr(S'_i \neq \emptyset \text{ and } m_i = 1) = iq^i(1 - \mu)^{i-1}\mu$. However, since survival likelihood immediately after the first mutation, p , is strictly higher than q and mutations sometimes occur (note, we assume $\mu > 0$), $Pr(S'_i \neq \emptyset \text{ and } m_i = 1) > iq^i(1 - \mu)^{i-1}\mu$. Putting these observations together, we have

$$\lim_{i \rightarrow \infty} \frac{Pr(\tau = i)}{Pr(S'_i \neq \emptyset)} < \lim_{i \rightarrow \infty} \frac{q^i(1 - \mu)^{i-1}\mu}{iq^i(1 - \mu)^{i-1}\mu} = \lim_{i \rightarrow \infty} \frac{1}{i} = 0, \quad (12)$$

where, as noted earlier, the inequality arises from replacing $Pr(S'_i \neq \emptyset)$ with a lower bound on the probability of exactly one mutation occurring over the course of the first i periods, and all the ways this could happen, and then noting that $q < p$. Now

$$\begin{aligned} Pr(\tau = i | S'_t \neq \emptyset) &= \frac{Pr(S'_t \neq \emptyset | \tau = i) Pr(\tau = i)}{Pr(S'_t \neq \emptyset)} \\ &= \frac{Pr(S_{t-i} \neq \emptyset) Pr(\tau = i)}{Pr(S'_t \neq \emptyset)} \\ &= \frac{Pr(S_{t-i} \neq \emptyset) Pr(\tau = i)}{Pr(S'_{t-i} = 1) Pr(S_i \neq \emptyset) + Pr(S'_{t-i} = 0) Pr(S'_i \neq \emptyset)} \\ &< \frac{Pr(S_{t-i} \neq \emptyset) Pr(\tau = i)}{Pr(S'_{t-i} \neq \emptyset) Pr(S'_i \neq \emptyset)}, \end{aligned}$$

where the inequality follows from the fact that $Pr(S_i \neq \emptyset) > Pr(S'_i \neq \emptyset)$, by Proposition 2, Part 1. $\frac{Pr(S_{t-i} \neq \emptyset)}{Pr(S'_{t-i} \neq \emptyset)}$ is bounded by Proposition 2 Part 2 (as it has a limit), and $\frac{Pr(\tau = i)}{Pr(S'_i \neq \emptyset)}$ can be made arbitrarily small for large enough i by equation 12. Thus, for any δ and k we can find large enough t for which $w_i^t < \delta$ for $i > t - k$. Choosing $\delta = \varepsilon/k$ establishes fact 1.

Finally, we prove fact 2. The probability distribution of S_t is given by $e_1' A^t$, where

$$A = \begin{bmatrix} p(1 - \mu) & p\mu & 0 \\ q\mu & q(1 - \mu) & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

is the Markov transition matrix for S . Let B be the principal 2×2 submatrix of A . By the partitioned matrix multiplication formula, $Pr(S_t = 1 | S_t \neq \emptyset) = \frac{e_1' B^t e_1}{e_1' B^t \mathbf{1}}$. Since B is strictly positive, the Perron-Frobenius theorem implies that this expression converges to the first entry of eigenvector corresponding to the largest eigenvalue of B . ■

Proof of Proposition 3:

$\lim_{t \rightarrow \infty} \frac{P_{1s}^t}{P_{1S}^t} = r$ for some $r < 1$, by Proposition 2. Let r_t be the t^{th} term in the sequence. Let $m(t)$ be the number of surviving signals. By Chernoff bounds, it follows that

$$Pr(m(t) > n(t)P_{1s}^t(1 + r_t)/2 | \omega = 1) \rightarrow 1$$

and

$$Pr(m(t) < n(t)P_{1s}^t(1 + r_t)/2 | \omega = 0) \rightarrow 1$$

provided that $n(t)P_{1s}^t \rightarrow \infty$. Given this separation, it is easy to check that if $n(t)P_{1s}^t \rightarrow \infty$, the beliefs will converge to 0 or 1 in probability.

Next, note that if $n(t)P_{1s}^t \rightarrow 0$, then the expected number of surviving signals in either state is 0, and that happens with the probability going to 1 by Chebychev, and so there is no learning. So, the threshold is $1/P_{1s}^t$.

Note that survival lies between $1/p^t$ and $1/q^t$ and so

$$1/P_{1s}^t = \frac{1}{(p\lambda(t) + (1 - \lambda(t))q)^t}.$$

The fact that $\lambda(t)$ converges to some λ then follows since this is a Markov chain and the probability that it survives in any given period (the third state with $s_t = \emptyset$ is absorbing) converges to a steady state distribution, which in this case lies between p and q . ■

The following lemma is useful in the proof of Proposition 4.

Let P_1^t (P_0^t) denote the Bayesian posterior probability that the state is 1 conditional upon a signal being received at time t and being 1 (0). Similarly, let P_\emptyset^t (P_S^t) denote the Bayesian posterior probability that the state is 1 conditional upon no signal (some signal) being received at time t .

LEMMA 4 *If $p > q$, then $P_1^t \geq P_0^t$ and $P_1^t \geq P_\emptyset^t$.*

Proof of Lemma 4

Let s^t denote the state of the signal at period t . That $P_1^t \geq P_0^t$ holds when $t = 1$ is easy to check from Bayes rule, given that $p > q$ and $\mu \leq 1/2$. Now suppose $P_1^t \geq P_0^t$ for some t . Then by the law of total probability, it follows that

$$\begin{aligned} P_1^{t+1} &= Pr(s^t = 0 | s^{t+1} = 1)P_0^t + Pr(s^t = 1 | s^{t+1} = 1)P_1^t \\ &= \frac{q\mu Pr(s^t = 0)}{q\mu Pr(s^t = 0) + p(1 - \mu)Pr(s^t = 1)}P_0^t + \frac{p(1 - \mu)Pr(s^t = 1)}{q\mu Pr(s^t = 0) + p(1 - \mu)Pr(s^t = 1)}P_1^t \end{aligned}$$

Similarly,

$$P_0^{t+1} = \frac{q(1 - \mu)Pr(s^t = 0)}{q(1 - \mu)Pr(s^t = 0) + p\mu Pr(s^t = 1)}P_0^t + \frac{p\mu Pr(s^t = 1)}{q(1 - \mu)Pr(s^t = 0) + p\mu Pr(s^t = 1)}P_1^t$$

Since $P_1^t \geq P_0^t$ by the inductive hypothesis, it suffices to show that

$$\frac{p(1-\mu)Pr(s^t=1)}{q\mu Pr(s^t=0) + p(1-\mu)Pr(s^t=1)} \geq \frac{p\mu Pr(s^t=1)}{q(1-\mu)Pr(s^t=0) + p\mu Pr(s^t=1)}$$

i.e., that

$$\frac{1}{1 + \frac{q Pr(s^t=0) - \mu}{p Pr(s^t=1) 1-\mu}} \geq \frac{1}{1 + \frac{q Pr(s^t=0) 1-\mu}{p Pr(s^t=1) \mu}}$$

which follows, since $\mu \leq 1 - \mu$.

To see that $P_1^t \geq P_\emptyset^t$, note that it suffices to prove that $P_S^t \geq P_\emptyset^t$, since P_S^t is a convex combination of P_1^t and P_0^t , and we just proved $P_1^t \geq P_0^t$. Now the statement follows directly from part 1 of Proposition 2. ■

Proof of Proposition 4:

First, note that we can focus on the case in which $p \neq q$ as otherwise there is nothing to be learned from signal survival, and agent C does as well as B . Without loss of generality we take $p > q$. Similarly, if $\mu = 1/2$, then all learning is from survival and S does as well as B , and so we can take $\mu < 1/2$.

Note that by Lemma 4, $P_1^t \geq P_0^t$ and $P_1^t \geq P_\emptyset^t$. In order for B to do strictly better in expectation than the other agents, it must be that $P_1^t > 1/2$ and at least one of P_0^t and P_\emptyset^t are less than $1/2$. To see this note that if all three are on the same side of $1/2$, then they must lie on the same side as the prior.²⁸ If $\theta \neq 1/2$ then N gets the same payoff as B . If $\theta = 1/2$, then for all three to lie on the same side of the prior it must be that $p = q$, in which case there is nothing learned from survival and C does as well as B in expectation.

Thus, $P_1^t > 1/2$ and at least one of P_0^t and P_\emptyset^t are less than $1/2$. If it is just P_\emptyset^t that is less than $1/2$, then S guesses the same as B (or equivalently in expected payoff terms). Thus, we need $P_0^t < 1/2$ to have a difference.

If it is just P_0^t that is less than $1/2$, then C guesses the same as B except if $\theta \leq 1/2$. But for such a θ , it must be that $P_\emptyset^t \leq 1/2$ and so C guesses as well as B .

So, consider the case in which $P_1^t > 1/2$ and $P_0^t < 1/2$ and $P_\emptyset^t < 1/2$. For C to guess differently than B , it must be that $\theta \geq 1/2$.

We can compute the expected payoff's for the three most relevant agents for this remaining case (we ignore N now, since in these conditions it is dominated by one of the others) for a given (p, q, μ, θ) satisfying the above constraints.

Letting U_B, U_C, U_S be the expected payoffs of agents B, C ²⁹ and S respectively, it follows that

$$U_B = Pr(s_t = 1)P_1^t + Pr(s_t = 0)(1 - P_0^t) + (1 - Pr(s_t = 1) - Pr(s_t = 0))(1 - P_\emptyset^t)$$

²⁸They cover three disjoint events whose union is all possibilities, and so the overall probability of a 1 is a convex combination of these conditionals, and so it is impossible to have them all weakly and some strictly greater (or all weakly and some strictly less) than the prior.

²⁹ C has expected payoff $U_C = Pr(s_t = 1)P_1^t + Pr(s_t = 0)(1 - P_0^t) + (1 - Pr(s_t = 1) - Pr(s_t = 0))(I_{\theta > 1/2}P_\emptyset^t + I_{\theta = 1/2}1/2)$. The expression in the main text is obtained by noting that the worst ratio for this compared to B will be in cases for which $\theta > 1/2$

$$\begin{aligned}
U_C &= Pr(s_t = 1)P_1^t + Pr(s_t = 0)(1 - P_0^t) + (1 - Pr(s_t = 1) - Pr(s_t = 0))P_0^t \\
U_S &= Pr(s_t = 1)P_1^t + Pr(s_t = 0)P_0^t + (1 - Pr(s_t = 1) - Pr(s_t = 0))(1 - P_0^t)
\end{aligned}$$

First, note that if $q < 1$ and $\mu > 0$, then as $t \rightarrow \infty$, then $Pr(s_t = \emptyset) \rightarrow 1$ and $P_0 \rightarrow 1/2$, in which case the ratio of B to either of these goes to 1. If $q < 1$ and $\mu = 0$, then B does as well as S for every t . If $p = q = 1$, then B does as well as C for every t . These facts together establish the last claim in the proposition that as $t \rightarrow \infty$, the ratio $\frac{U_B}{\max\{U_S, U_C\}} \rightarrow 1$.

That the ratio is bounded above by $3/2$ can be seen as follows. Since $\theta \geq 1/2$ and $p > q$, it follows that

$$Pr(s_t = 1) \geq Pr(s_t = 0), \quad P_1^t \geq (1 - P_0^t), \quad \text{and so } Pr(s_t = 1)P_1^t \geq Pr(s_t = 0)(1 - P_0^t).$$

Then if $Pr(s_t = 0)(1 - P_0^t) \leq (1 - Pr(s_t = 1) - Pr(s_t = 0))(1 - P_0^t)$ it follows that $U_S \geq U_B 2/3$. If $Pr(s_t = 0)(1 - P_0^t) \geq (1 - Pr(s_t = 1) - Pr(s_t = 0))(1 - P_0^t)$ then it follows that $U_C \geq U_B 2/3$.

To complete the proof, we compute

$$\max_{p, q, \theta, \mu \in [0, 1]} \frac{U_B}{\max\{U_S, U_C\}}.$$

for $t = 1$. We can rewrite the payoffs of agents B , S and C in the case $P_1^1 > 1/2$ and $P_0^1 < 1/2$ and $P_0^1 < 1/2$ as follows:

$$\begin{aligned}
U_B &= \theta p(1 - \mu) + (1 - \theta)(1 - q\mu) \\
U_C &= \theta(1 - p\mu) + (1 - \theta)q(1 - \mu) \\
U_S &= \theta p + (1 - \theta)(1 - q)
\end{aligned}$$

where

$$\theta p \mu \leq q(1 - \theta)(1 - \mu) \tag{13}$$

$$\theta(1 - p) \leq (1 - \theta)(1 - q) \tag{14}$$

$$\theta \geq 1/2 \tag{15}$$

$$\mu \leq 1/2 \tag{16}$$

$$p \geq q \tag{17}$$

$$p, q, \mu, \theta \in [0, 1]. \tag{18}$$

Case 1: $U_S \leq U_C$.

This condition can be rewritten as

$$\theta(p\mu + (p - 1)) \leq (1 - \theta)(q(1 - \mu) + (q - 1)) \tag{19}$$

The program with this additional constraint can be written as

$$\max_{p, q, \theta, \mu \text{ satisfy 13-19}} \frac{U_B}{U_C} \equiv \max_{p, q, \theta, \mu \text{ satisfy 13-19}} \frac{\theta p(1 - \mu) + (1 - \theta)(1 - q\mu)}{\theta(1 - p\mu) + (1 - \theta)q(1 - \mu)}$$

$$\begin{aligned}
&= \max_{p,q,\theta,\mu \text{ satisfy 13-19}} \frac{\theta p + (1 - \theta) - \mu(\theta p + (1 - \theta)q)}{\theta + (1 - \theta)q - \mu(\theta p + (1 - \theta)q)} \\
&\leq \max_{p,q,\theta,\mu \text{ satisfy 13-19}} \frac{\theta + (1 - \theta)2q - 2\mu(\theta p + (1 - \theta)q)}{\theta + (1 - \theta)q - \mu(\theta p + (1 - \theta)q)}
\end{aligned}$$

where the inequality is from rearranging constraint 19, as $\theta p + (1 - \theta) \leq \theta + (1 - \theta)2q - \mu(\theta p + (1 - \theta)q)$, and plugging this into the numerator. It is easily verified that the above ratio is decreasing in μ for any values of the remaining parameters³⁰. Moreover, reducing μ to 0 only relaxes constraints 13, 16 and 19, and leaves the other constraints unaffected. Therefore,

$$\max_{p,q,\theta,\mu \text{ satisfy 13-19}} \frac{U_B}{U_C} \leq \max_{p,q,\theta \text{ satisfy 14-19}} \frac{\theta + (1 - \theta)2q}{\theta + (1 - \theta)q}$$

It is clear that smaller values of θ increase this ratio, and by constraint 15, the smallest value of θ is $\frac{1}{2}$. But while reducing θ down to $\frac{1}{2}$ for given p and q relaxes constraint 14, doing so may violate constraint 19. We therefore separately consider the cases where either 19 or 15 bind, since at least one of them must at the optimum.

Subcase 1: 19 is satisfied with equality, i.e., $\theta(1 - p) = (1 - \theta)(1 - 2q)$. Plugging this in, the objective then becomes $2\frac{1+\theta(p-1)}{1+\theta(p-1)+\theta}$, which is decreasing in θ , so it is optimal to set $\theta = \frac{1}{2}$. The objective is then $2\frac{1+p}{2+p} \leq 4/3$. Note that at $p = 1, q = \frac{1}{2}, \theta = \frac{1}{2}, \mu = 0, \frac{U_B}{U_C} = \frac{4}{3}$, so this upper bound is tight.

Subcase 2: $\theta = 1/2$. Then

$$\frac{U_B}{U_C} = \frac{p + 1 - \mu(p + q)}{q + 1 - \mu(p + q)},$$

which is weakly increasing in μ by constraint 17. Constraint 19 can be rearranged to be

$$\mu \leq \frac{2q - p}{p + q},$$

which, first, implies that

$$\frac{U_B}{U_C} \leq \frac{2(p - q) + 1}{(p - q) + 1},$$

and second, along with the condition that $\mu \geq 0$, implies that

$$q \geq \frac{p}{2}.$$

Since $\frac{2(p-q)+1}{(p-q)+1}$ is decreasing in q , this expression is maximized under the given constraints when $q = \frac{p}{2}$. Therefore, $\frac{U_B}{U_C} \leq \frac{p+1}{\frac{p}{2}+1}$, which is maximized when $p = 1$ and equals $4/3$.

Case 2: $U_S \geq U_C$. The new constraint is

$$\theta(p\mu + (p - 1)) \geq (1 - \theta)(q(1 - \mu) + (q - 1)), \tag{20}$$

³⁰ $\frac{d}{dx} \frac{A-2x}{B-x} \leq 0$ if $A \leq 2B$ and $A, B > 0$.

and the relevant maximization program is

$$\max_{p,q,\theta,\mu \text{ satisfy 14-18, 20}} \frac{U_B}{U_S} \equiv \max_{p,q,\theta,\mu \text{ satisfy 14-18, 20}} \frac{\theta p(1-\mu) + (1-\theta)(1-q\mu)}{\theta p + (1-\theta)(1-q)}.$$

Notice that the ratio $\frac{\theta p(1-\mu) + (1-\theta)(1-q\mu)}{\theta p + (1-\theta)(1-q)}$ is linear and decreasing in μ , and the constraints are linear in μ as well. Constraint 13 only places an upper bound on μ , so it is not relevant in pinning down this value at the optimum. On the other hand, constraint 20, which can be rewritten as

$$2q(1-\theta) - p\theta - (1-2\theta) \leq (p\theta + q(1-\theta))\mu$$

and the constraint that $\mu \geq 0$ are relevant. There are two cases:

Subcase 1: $2q(1-\theta) - p\theta - (1-2\theta) \geq 0$, $\mu = \frac{2q(1-\theta) - p\theta - (1-2\theta)}{p\theta + q(1-\theta)}$.

Then

$$\begin{aligned} \frac{U_B}{U_S} &= \frac{\theta p + (1-\theta) - \mu(\theta p + (1-\theta)q)}{\theta p + (1-\theta) - (1-\theta)q} \\ &= \frac{\theta p + (1-\theta) - 2(1-\theta)q + p\theta + (1-2\theta)}{\theta p + (1-\theta) - (1-\theta)q} \\ &= \frac{2\theta p + (2-3\theta) - 2(1-\theta)q}{\theta p + (1-\theta) - (1-\theta)q} \\ &= 2 \frac{\theta p + (1-\frac{3}{2}\theta) - (1-\theta)q}{\theta p + (1-\theta) - (1-\theta)q} \end{aligned}$$

Clearly, the ratio is decreasing in θ , and moreover, decreasing θ only relaxes constraints 13 and 14. Therefore, constraint 15 binds and $\theta = \frac{1}{2}$ at the optimum, so

$$\frac{U_B}{U_S} = 2 \frac{p + \frac{1}{2} - q}{p + 1 - q}$$

Since $\mu = \frac{2q-p}{p+q}$ at $\theta = \frac{1}{2}$, constraints 13 and 14 reduce to just $p \geq q$. Since the ratio is increasing in $p - q$, the only binding constraint is that $\mu \geq 0$, i.e., $2q \geq p$. Therefore at the optimum, $p = 1$, $q = \frac{1}{2}$, $\mu = 0$, $\theta = \frac{1}{2}$, and $\frac{U_B}{U_S} = \frac{4}{3}$.

Subcase 2: $2q(1-\theta) - p\theta - (1-2\theta) \leq 0$, $\mu = 0$. In this case, the problem reduces to

$$\max_{p,q,\theta, \text{ satisfy 14,15,17, 18, 20}} \frac{\theta p + (1-\theta)}{\theta p + (1-\theta)(1-q)},$$

which is decreasing in θ . Now

$$\begin{aligned} 2q(1-\theta) - p\theta - (1-2\theta) &\leq 0 \\ \iff \theta(2-2q-p) &\leq 1-2q \end{aligned}$$

Suppose $2-2q-p < 0$. Since $1-2q \leq 1-2q + (1-p) = 2-2q-p < 0$, it follows that $\theta(2-2q-p) \geq \theta(1-2q) \geq 1-2q$. Therefore, the only way to satisfy the constraint is if $p = 1$ and $\theta = 1$, in which case $\frac{U_B}{U_S} = 1$.

If $2 - 2q - p \geq 0$, then $\theta = \frac{1}{2}$ at the optimum, and so the constraint in this sub-subcase becomes $2q \leq p$, while the objective function is $\frac{p+1}{p+1-q}$. This constraint binds at the optimum and again the optimal value is $\frac{4}{3}$ at $p = 1$ and $q = \frac{1}{2}$. ■

Proof of Proposition 5:

We proceed by cases for different values of the parameters. We concentrate on situations in which $\mu < 1/2$ since if $\mu = 1/2$ then content is completely uninformative and the result is direct.

Case 1: $\mu = 0$. Suppose without loss of generality that $q \leq p$. Any signal that reaches the agent is perfectly informative of the state, so a threshold for learning for agents B and C is the threshold for at least one signal to survive, which (following the logic of the proofs above) is $\frac{1}{p^t}$.

Case 2: $p = q$ and $\mu > 0$. By Proposition 1, the threshold for learning for agent B is $\frac{1}{p^t(1-2\mu)^{2t}}$. In this case there is no information from signal survival, and by Lemma 2, agent B's posterior is the same as agent C's posterior. Therefore, agent C has the same threshold for learning as B.

Case 3: $q \neq p$ and $\mu > 0$. Without loss of generality let $p > q$. Then $\tau(t) = \frac{1}{P_{1S}^t}$ is a threshold for learning for an agent conditioning only on signal survival, as shown in the proof of Proposition 3. Let $b(t)$ denote the beliefs of agent B after observing the outcome of $n(t)$ original sources of information sent along chains of depth t . Since agent B conditions on survival and signal content, $\text{plim } b(t) \rightarrow 1$ or 0 whenever $n(t)/\tau(t) \rightarrow \infty$. When $n(t)/\tau(t) \rightarrow 0$, then the probability of even a single signal surviving to reach the agent approaches 0. This holds regardless of the starting state by Proposition 2 part 2, so $\text{plim } b(t) \rightarrow \theta$. Therefore, agent B and S have the same thresholds for learning in this case.³¹ ■

Proof of Proposition 6:

We give the proof for the case in which $p^t n(t) \rightarrow \infty$. (With fewer paths there are even fewer signals from which to learn.) The following lemma is straightforward (and so its proof is omitted) but it useful.

LEMMA 5 Consider a sequence of $k \leq m$ such that $m \rightarrow \infty$ and $\frac{k}{m} \rightarrow a$. The maximizer of $z^k(1-z)^{m-k}$ is $z(m, k) = \frac{k}{m}$, and

$$\frac{z(m, k)^k(1-z(m, k))^{m-k}}{z^k(1-z)^{m-k}} \rightarrow \infty$$

for any $z \neq a$, the size of this ratio increases with the distance of z from a (as $z^a(1-z)^{1-a}$ is strictly concave). Moreover, for any atomless and continuous probability measure G on z that has connected support and includes a in its interior

$$\frac{\int_{a-\varepsilon}^{a+\varepsilon} z^k(1-z)^{m-k} dG(z)}{\int_0^1 z^k(1-z)^{m-k} dG(z)} \rightarrow 1,$$

³¹Strictly speaking, we only showed that they share a common threshold, but it is easy to see that being a threshold for learning for B, for S or for neither partitions the space of functions on $\mathbb{N} \rightarrow \mathbb{N}$.

for any $\varepsilon > 0$.

Let $Y_\pi(t)$ be the probability that a sequence ends with a signal 1 at time t conditional upon there being at least one extremist in the sequence, the sequence reaching time t , and the fraction of extremists always sending 1 being given by π . Note that this is independent of the starting state. Then $1 - Y_\pi(t)$ is the probability that a sequence with at least one extremist ends with a signal 0 at time t .

Moreover, the process is a time-homogeneous irreducible and aperiodic Markov chain, and so it has a steady state distribution, and $Y_\pi(t)$ converges to that steady-state as t grows. It is direct to calculate the limit Y_π , and in particular, note that the probability that a signal that was received as a 1 in one period is received as a 0 in the next period (conditional upon being received) is

$$p_{10} = (1 - f)\mu + f\pi\mu + f(1 - \pi)(1 - \mu) = \mu + f(1 - \pi)(1 - 2\mu)$$

and similarly the probability that a signal that was received as a 0 is then received as a 1 in the next period is

$$p_{01} = \mu + f\pi(1 - 2\mu).$$

This means, via a standard calculation for the limiting distribution for a two-state Markov chain, that the steady state limit is

$$Y_\pi = \frac{p_{01}}{p_{01} + p_{10}} = \frac{\mu + f\pi(1 - 2\mu)}{2\mu + f(1 - 2\mu)},$$

which is linear in π . Also, note that

$$Y_\pi(1) = \pi(1 - \mu) + (1 - \pi)\mu = \mu + \pi(1 - 2\mu),$$

and for $t > 1$,

$$Y_\pi(t + 1) = (1 - (1 - f)^t) [Y_\pi(t)(1 - p_{10}) + (1 - Y_\pi(t))p_{01}] + (1 - f)^t f (\mu + \pi(1 - 2\mu)),$$

where the first expression captures the probability that the process has already had an extremist, and the second expression is the chance that this is the first period after an extremist. This becomes

$$Y_\pi(t+1) = (1 - (1 - f)^t) [Y_\pi(t)(1 - f)(1 - 2\mu) + \mu + f\pi(1 - 2\mu)] + (1 - f)^t f (\mu + \pi(1 - 2\mu)),$$

which is also linear and increasing in π , by induction (noting that $Y_\pi(1)$ is linear and increasing in π). From the above it is also clear that the slopes of $Y_\pi(t)$ converge to that of Y_π , and so for large enough t , $\frac{\partial Y_\pi(t)}{\partial \pi} > \delta > 0$, for some δ regardless of t .

Also note that since $Y_\pi(t)$ is increasing in π , it is invertible.

The probability that some sequence ends in a 1 conditional on π and starting in state $\omega = 1$ is

$$(1 - (1 - f)^t) Y_\pi(t) + (1 - f)^t X(t).$$

The probability that some sequence ends in a 1 conditional on π and starting in state $\omega = 0$ is

$$(1 - (1 - f)^t) Y_\pi(t) + (1 - f)^t(1 - X(t)).$$

Similar calculations provide probabilities of ending in a 0.

The chance of observing k 1s, conditional on m sequences reaching the receiver, on π and on the starting state being $\omega = 1$ is then

$$P_{k,m,t,\pi}(1) = \binom{m}{k} [(1 - (1 - f)^t) Y_\pi(t) + (1 - f)^t X(t)]^k [(1 - (1 - f)^t) (1 - Y_\pi(t)) + (1 - f)^t(1 - X(t))]^{m-k}.$$

Then the chance of observing k 1s out of m sequences that reach the receiver conditional the starting state being $\omega = 0$ is then

$$P_{k,m,t,\pi}(0) = \binom{m}{k} [(1 - (1 - f)^t) Y_\pi(t) + (1 - f)^t(1 - X(t))]^k [(1 - (1 - f)^t) (1 - Y_\pi(t)) + (1 - f)^t X(t)]^{m-k}.$$

First consider the case where π is known, and suppose the state is 1 (the argument for the case where the state is 0 is analogous). As the number of signals grows large (keeping t fixed), $\frac{k}{m-k} \rightarrow \frac{(1-(1-f)^t)Y_\pi(t)+(1-f)^tX(t)}{(1-(1-f)^t)(1-Y_\pi(t))+(1-f)^t(1-X(t))} = a_{t,1}$ in probability, and $a_{t,1} > 1$. Now, the Bayesian's posterior that the state is $\omega = 1$ conditional upon seeing k 1's out of m sequences that reached the receiver³²

$$\frac{\theta P_{k,m,t,\pi}(1)}{\theta P_{k,m,t,\pi}(1) + (1 - \theta) P_{k,m,t,\pi}(0)},$$

By Lemma 5, $\frac{P_{k,m,t}(0)}{P_{k,m,t}(1)} \rightarrow 0$ in probability as the number of signals grow large, so it follows that

$$\frac{\theta P_{k,m,t}(1)}{\theta P_{k,m,t}(1) + (1 - \theta) P_{k,m,t}(0)} \rightarrow 1.$$

Therefore, since the agent can learn the true state with sufficiently many paths for any given t , it follows that the agent can learn the true state as $t \rightarrow \infty$ if $n(t)$ grows quickly enough.

Now we consider the case when π is unknown but follows an atomless distribution F with connected support. A Bayesian's posterior that the state is $\omega = 1$ conditional upon seeing k 1's out of m sequences that reached the receiver is

$$\frac{\theta \int_\pi P_{k,m,t,\pi}(1) dF(\pi)}{\theta \int_\pi P_{k,m,t,\pi}(1) dF(\pi) + (1 - \theta) \int_\pi P_{k,m,t,\pi}(0) dF(\pi)},$$

³²Note that each sequence has an independent probability p^t of reaching the observer, so this there is nothing to update about which sequences reach the observer when $p = q$.

and so if we can show that $\int_{\pi} P_{k,m,t,\pi}(1)dF(\pi)/\int_{\pi} P_{k,m,t,\pi}(0)dF(\pi)$ converges to one in probability, then we conclude the proof.

Given a true π^* in the interior of the support of F , the realized k, m will be such that $\frac{k}{m-k} - \frac{Y_{\pi^*}(t)}{1-Y_{\pi^*}(t)}$ converges to 0 in probability, and $\frac{Y_{\pi^*}(t)}{1-Y_{\pi^*}(t)} \rightarrow \frac{Y_{\pi^*}}{1-Y_{\pi^*}} = a$.

By the first part of Lemma 5, for any k, m , $P_{k,m,t,\pi}(1)$ is maximized when

$$Y_{\pi}(t) = \frac{\frac{k}{m} - (1-f)^t X(t)}{1 - (1-f)^t},$$

and $P_{k,m,t,\pi}(0)$ is maximized when

$$Y_{\pi}(t) = \frac{\frac{k}{m} - (1-f)^t(1-X(t))}{1 - (1-f)^t}.$$

Under the true π^* , if $m(t) = p^t n(t) \rightarrow \infty$, then these two right hand sides converge to each other as t becomes large. Given that $\frac{\partial Y_{\pi}(t)}{\partial \pi} > \delta > 0$, it then follows that the $\pi_1(t, k, m)$ and $\pi_0(t, k, m)$ that are the corresponding maximizers, converge to each other, and to π^* , as well in probability.³³ It therefore follows from Lemma 5 that

$$plim \frac{\int_{\pi} P_{k,m,t,\pi}(1)dF(\pi)}{\int_{\pi} P_{k,m,t,\pi}(0)dF(\pi)} = plim \frac{\int_{\pi_1(t,k,m)-\varepsilon}^{\pi_1(t,k,m)+\varepsilon} P_{k,m,t,\pi}(1)dF(\pi)}{\int_{\pi_0(t,k,m)-\varepsilon}^{\pi_0(t,k,m)+\varepsilon} P_{k,m,t,\pi}(0)dF(\pi)}$$

for any $\varepsilon > 0$.

Letting $[l, h]$ be the support of π , and recalling that $Y_{\pi}(t)$ is linear in π and note that $P_{k,m,t,\pi_1(t,k,m)}(1) = P_{k,m,t,\pi_0(t,k,m)}(0)$. From this, it follows that

$$P_{k,m,t,\pi_1(t,k,m)+\delta}(1) = P_{k,m,t,\pi_0(t,k,m)+\delta}(0)$$

for any $\delta \in \mathbb{R}$ such that both $\pi_1(t, k, m) + \delta$ and $\pi_0(t, k, m) + \delta$ fall in (l, h) . In particular, if we let $\varepsilon_t = \frac{1}{2} \min\{\pi_1(t, k, m) - l, h - \pi_0(t, k, m)\}$ and if $\varepsilon_t > 0$, the intervals $[\pi_1(t, k, m) - \varepsilon_t, \pi_1(t, k, m) + \varepsilon_t]$ and $[\pi_0(t, k, m) - \varepsilon_t, \pi_0(t, k, m) + \varepsilon_t]$ strictly lie in (l, h) . So by the earlier observation,

$$\int_{\pi_1(t,k,m)-\varepsilon_t}^{\pi_1(t,k,m)+\varepsilon_t} P_{k,m,t,\pi}(1)dF(\pi) = \int_{\pi_0(t,k,m)-\varepsilon_t}^{\pi_0(t,k,m)+\varepsilon_t} P_{k,m,t,\pi}(0)dF(\pi)$$

Moreover $plim \varepsilon_t = \frac{1}{2} \min\{\pi^* - l, h - \pi^*\} > 0$. Therefore, by the continuous mapping theorem,

$$plim \frac{\int_{\pi} P_{k,m,t,\pi}(1)dF(\pi)}{\int_{\pi} P_{k,m,t,\pi}(0)dF(\pi)} = plim \frac{\int_{\pi_1(t,k,m)-\varepsilon_t}^{\pi_1(t,k,m)+\varepsilon_t} P_{k,m,t,\pi}(1)dF(\pi)}{\int_{\pi_0(t,k,m)-\varepsilon_t}^{\pi_0(t,k,m)+\varepsilon_t} P_{k,m,t,\pi}(0)dF(\pi)} = 1,$$

³³There will be realizations for which $\pi_1(t, k, m)$ and $\pi_0(t, k, m)$ that exactly solve the equations do not exist and then a corresponding corner solution of the extremes of the support of the prior can be used, but those will occur with vanishing probability.

which concludes the proof. ■

Proof of Proposition 7:

Let $\bar{\mu}$ be the supremum of the support of the agent's prior on μ . By assumption, $\bar{\mu} < 1/2$. By Propositions 1 and 3, we know there exists an $f(t)$ such that if the mutation rate is known to be $\bar{\mu}$, $\text{Plim } b(t) = 0$ or 1 whenever $n(t)/f(t) \rightarrow \infty$. With such a choice of $n(t)$, the probability of full learning approaches 1 in t conditional on any value of μ in the support of the agent's prior. Therefore the agent's unconditional posterior also converges to 0 or 1 in probability.