PNAS RESEARCH ARTICLE | ECONOMIC SCIENCES

Learning through the grapevine and the impact of the breadth and depth of social networks

Matthew O. Jackson[a,b,1], Suraj Malladi[c], and David McAdams[d]

Contributed by Matthew O. Jackson; received March 30, 2022; accepted July 8, 2022; reviewed by Manuel Mueller-Frank and Mallesh Pai

We study how communication platforms can improve social learning without censoring or fact-checking messages, when they have members who deliberately and/or inadvertently distort information. Message fidelity depends on social network depth (how many times information can be relayed) and breadth (the number of others with whom a typical user shares information). We characterize how the expected number of true minus false messages depends on breadth and depth of the network and the noise structure. Message fidelity can be improved by capping depth or, if that is not possible, limiting breadth, e.g., by capping the number of people to whom someone can forward a given message. Although caps reduce total communication, they increase the fraction of received messages that have traveled shorter distances and have had less opportunity to be altered, thereby increasing the signal-to-noise ratio.

social learning | network | misinformation

Significance

The online spread of misinformation has prompted debate about how social media platforms should police their content. A tacit assumption has been that censorship, fact-checking, and education are the only tools to fight misinformation. However, even well-intentioned censors may be biased, and fact-checking at the speed and scale of today's platforms is often impractical. We ask the policy-relevant question: can one improve the quality of information shared in networks without deciding what is true and false? We show that caps on either how many times messages can be forwarded or the number of others to whom messages can be forwarded increase the relative number of true versus false messages circulating in a network, regardless of whether messages are accidentally or deliberately distorted.

> The safety of our democracy is more important than shareholder dividends and CEO salaries ... That's why I'm calling on [tech companies] to take real steps right now to fight disinformation.
>
> Elizabeth Warren

> I don't think that Facebook or Internet platforms in general should be arbiters of truth.
>
> Mark Zuckerberg

Misinformation has always been a social concern, but there has recently been a resurgence of interest in addressing it. This is both because social media and messaging platforms make (mis)information spreading easier (e.g., see ref. 1) and because they have made its impact more transparent, with wide-ranging implications from voter attitudes toward election outcomes to their views on vaccination. The most common approach to regulate communication on online platforms involves some form of fact-checking and flagging or censorship. However, policing communication is challenging and problematic due to the enormous volume and fast pace of online communication and the potential for bias by whoever is "determining the truth," be it a government or private enterprise, an algorithm, or crowd. Moreover, communication on some platforms is intentionally encrypted to protect user privacy, making fact-checking impossible.

Given these challenges, we study policies that improve informational content and social learning without relying on anyone to know or decide what is true. In particular, we study how people learn as a function of their networks when information is subject to mutation, deliberate manipulation, and transmission failure. We characterize how learning depends on the depth and breadth of a person's network. Increasing depth and breadth of a network increases the number of messages passed and therefore received. However, it also increases the number of distant sources, from which surviving messages are less accurate, more than it increases the number of nearby sources. Our analysis shows how and when limiting the network can improve the accuracy of overall content without the need for censorship or message monitoring. Our results show how one can satisfy Warren's objectives outlined in the quote above, while respecting Zuckerberg's reticence for private enterprises to serve as arbiters of the truth.

In our model, information is relayed from original sources via sequences of individuals to an eventual "learner," who wishes to ascertain the state of the world. Our results apply whether the learner is a fully rational Bayesian facing substantial uncertainty or a naive learner who is simply influenced by the preponderance of messages. The state of the world and corresponding messages are 1 or 0 (e.g., the state can be "climate change is an urgent concern" or "climate change is not an urgent concern"). With

Author contributions: M.O.J., S.M., and D.M. designed research, performed research, and wrote the paper.

Reviewers: M.M.-F., IESE Business School; and M.P., Rice University Department of Economics.

The authors declare no competing interest.

Copyright © 2022 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

[1]To whom correspondence may be addressed. Email: jacksonm@stanford.edu.

Published August 15, 2022.

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.

noiseless word-of-mouth (oral or digital) communication and sufficiently many sources of conditionally independent information, the receiver learns the true state. However, along each transmission chain, the message may mutate (deliberately or inadvertently), become unintelligible, or be dropped, reducing the information content of the signals that reach the receiver.

Limiting how many times messages can be relayed (network depth) is key to overcoming the noise that builds up as messages are repeatedly relayed. People see fewer messages when depth is capped, but the messages that they do receive are more informative on average, thereby increasing the overall signal-to-noise ratio and enabling at least partial learning. We derive an optimal cap on depth and show that if depth can be capped, there is no need to otherwise restrict communication.

However, what if depth cannot be capped? For instance, some social media platforms do not track whether a message is new or forwarded. Moreover, tracking messages' trajectories can become particularly challenging when they mutate. Such platforms can instead limit network breadth by making it more difficult to forward a message to many people. Such a cap can be easier to implement even though it might be a cruder way of improving the ratio of true to false signals. In particular, breadth restrictions still improve learning since decreasing breadth increases the relative number of messages from closer sources nodes, as higher breadth increases the expansion properties of a network and thus the relative number of distant to nearby nodes.

Indeed, breadth limits have been adopted by online messaging platforms. For instance, WhatsApp has capped the number of people that someone can message, for the express purpose of curbing the spread of false information. Facebook has implemented a similar strategy, capping the number of people or groups to which a message can be forwarded (https://about.fb.com/news/2020/09/introducing-a-forwarding-limit-on-messenger/).

Some Background Observations

Consider the children's game of "Telephone," in which a message is whispered from one player to the next. The final message typically bears little resemblance to the original because of "mutations" along the transmission chain. Such mutations have been found to occur frequently. In ref. 2's study of online viral memes, one meme was reposted more than 470,000 times, with a mutation rate of around 11% and more than 100,000 variants. Indeed, 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 (3): estimates that there are around 500,000 textual variants of the Greek New Testament, not including spelling errors.] In particular, when relayed, a message can change—intentionally or inadvertently—in ways that alter its meaning (mutate) or render it unintelligible and/or irrelevant to the underlying state (drop), or simply not be forwarded (another way of being dropped). Ref. 4 discusses an interesting example of a message whose meaning was inadvertently lost in transmission. 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 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."

Tabloids are rife with examples of intentional message mutations. For example, the Yellowstone Volcano Observatory (YVO) discussed in an article how they repaired a monitoring station that was damaged during a storm. However, aspects of the article were quickly grabbed by a tabloid website, and within hours they had produced their own article that exaggerated and misconstrued the original information and implied Yellowstone was unmonitored. That story even failed to correctly copy and paste the original text. As an example, the word "borehole" in the original was misspelled as "boreal" in the tabloid version. The exact same misspelling was perpetuated in social media reports and by sources that repeatedly conduct misinformation campaigns about Yellowstone. They either did not know, or ignored, the fact that YVO was the original source of the article and that the tabloid changed the article's meaning to something ominous (https://www.usgs.gov/center-news/playing-telephone-miss-information).

In our model, both intentional and inadvertent changes in content are captured by the probability of nodes changing messages when they forward them, while failures to forward and garbling of content that renders a message uninterpretable are captured by the probability of a message being dropped. Importantly, we allow mutations to occur in different directions with different probabilities. For example, as we have noted, some mutations may be intentional as ideologues relay messages that they prefer telling rather than what they truly heard (e.g., "fake news"), and there may be more ideological pressure in one direction or another depending on the topic (5, 6).

In ref. 7, we show that in the presence of such mutations, slight uncertainty over the mutation rates severely limits what even a fully Bayesian receiver can learn from a distance, no matter how many independent chains of message relays they hear. Moreover, it appears that people are far from the ideal Bayesian social learner and tend to be swayed by the preponderance of messages (e.g., see ref. 8). Indeed, experimental evidence suggests that simply repeating falsehoods makes people more likely to believe them, even when no additional evidence is presented and they have prior knowledge to the contrary (see ref. 9 and the references therein). Also, ref. 10 shows that people are poor at distinguishing what is true and false but become more confident in false messages if shared by others. In light of these findings, we step back from particular specifications of sender intentions and receiver updating rules and focus directly on the problem of how platforms can increase message fidelity.

Several governments have increased regulations and fines for disinformation (11). This puts platforms that are hesitant to be arbiters of the truth in a difficult situation. Our analysis highlights a policy that such platforms could follow to ensure that most messages on their platforms are true, without needing to police message content. Limiting message passing does not eliminate false information but can improve the overall quality of what gets shared and allow people to learn more from the content they encounter.

A Model of Noisy Information Transmission

Information is relayed by word-of-mouth (oral, written, via social media, etc.) from original sources to a learner. For instance, the learner may hear from friends about whether there is a link between vaccines and autism and then decide whether to vaccinate their child.[*]

There are two possible states of the world, $\omega \in \{0, 1\}$. Let θ be the prior probability that the state is 1.[†]

[*] The learner may have information from sources outside of their network. If these sources are not direct, then they can be modeled as part of the network. Otherwise, we can think of this external information as being reflected in the prior.

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

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.

Agents are in an infinite network.

With probability $r > 0$ any given agent is an original source and gets information (exogenously from the network) about the state, and the agent is uninformed with probability $1 - r$. An original source observes a noisy signal of the state in $\{0, 1\}$, as described below. All original sources then transmit their information to all of their neighbors as a message. Neighbors relay each piece of information that they receive to their neighbors, and this process repeats.

Thus, in the first period, some random people hear direct (but noisy) information about the state. Then, in the next period, they broadcast that information to their neighbors via noisy messages. In each of the following periods, people broadcast the messages they have received in the last period, each time subject to noise.

To keep the analysis uncluttered, we abstract from cycles and multiple paths that the same information may follow. In particular, we study the perspective of some learner who is the root node of the subtree consisting of the learner's parent nodes, their parents, and so on. Each node in this subtree is a potential source of information that may then be relayed with noise to the learner.

We allow the tree to be random, following a standard Galton–Watson branching process: each node has an independent and identically drawn number of parent nodes. What is important in the analysis is the expected number of parent nodes at each level in the tree, which we denote by k. The exact distribution that gives that expected degree is not consequential, and one can think of a regular tree or some other distribution that has k parent nodes on average at each level.

Consider a path of length T that information travels from an original source to the learner. Label the source node 1 and the learner $T > 1$, with the message passing via the sequence (or chain) of agents $\{1, 2, \ldots, T\}$.

Messages are distributed as follows. Let $m_0 \in \{0, 1\}$ be the true state.

Messages m_t sent by agent $t \in \{1, 2, \ldots, T\}$ take on the values $\{0, 1, \emptyset\}$, with the null message, $m_t = \emptyset$, indicating that the message either was dropped at some earlier stage or was so garbled that it was not interpretable. In particular, if agent $t \geq 1$ receives a null message $m_t = \emptyset$, then all subsequent agents (including the learner) also receive a null message.

Noisy transmission along the chain proceeds according to the following procedure. If agent $t > 1$ receives a message $m_t \in \{0, 1\}$, then that agent passes along a nonnull message ($m_{t+1} \neq \emptyset$) with probability p, and the message is dropped/garbled ($m_{t+1} = \emptyset$) with probability $1 - p$.[‡]

Each time a nonnull message is transmitted, that message mutates from 0 to 1 with probability $\mu_{01} \in (0, 1/2)$ or from 1 to 0 with probability $\mu_{10} \in (0, 1/2)$.[§]

Thus, for each $t \geq 0$, if $m_t = 1$, then agent $t + 1$ passes along the message $m_{t+1} = 1$ with probability $p(1 - \mu_{10})$, $m_{t+1} = 0$ with probability $p\mu_{10}$, and $m_{t+1} = \emptyset$ with probability $1 - p$. Similarly, if $m_t = 0$, then $m_{t+1} = 1$ with probability $p\mu_{01}$, $m_{t+1} = 0$ with probability $p(1 - \mu_{01})$, and $m_{t+1} = \emptyset$ with probability $1 - p$. If $m_t = \emptyset$, then $m_{t+1} = \emptyset$. An agent in the tree may end up relaying multiple messages at once or at multiple times in the process. We treat each relay as independent.

We assume that the initial message m_1 depends on the true state m_0 in the same way that any other m_{t+1} depends on m_t, i.e., as if nature were agent 0 in the chain. This simplifies expressions, although our analysis easily extends to allow first-signal accuracies and dropping rates to differ from subsequent ones.[¶]

This defines a 3×3 Markov chain in which $\emptyset$ is an absorbing state.

The basic structure of the model is pictured in Fig. 1.

The interesting case is when mutation rates differ across states. If $\mu_{01} = \mu_{10}$, then messages are always (slightly) more likely to match the starting state. Even a naive learner who just follows the majority signal would make correct decisions on average by following messages originating from any distance. However, for instance, if $\mu_{01} < \mu_{10}$, then as T becomes large, any surviving message is more likely to be a 0 than a 1, regardless of the starting state.

In particular, letting $M = 1 - \mu_{01} - \mu_{10}$, if the state $\omega = 0$, the frequency of true (0) to false (1) messages originating T steps away is[#]

$$\frac{\mu_{10} + \mu_{01} M^T}{\mu_{01} - \mu_{01} M^T}. \quad \textbf{[1]}$$

If $\omega = 1$, the frequency of true (1) to false (0) messages originating T steps away is

$$\frac{\mu_{01} + \mu_{10} M^T}{\mu_{10} - \mu_{10} M^T}. \quad \textbf{[2]}$$

A Bayesian learner who knows the exact mutation rates (μ_{01}, μ_{10}) and the exact distance T that a message has traveled can leverage the difference between the frequency ratios in Eqs. **1** and **2** to glean information about the state from a message, even if it has traveled from a great distance. For such a learner, having more messages from any distance, and hence a broader and deeper network, is unambiguously good for learning.

However, perfect knowledge about message distances and mutation rates is unrealistic: people often cannot tell how far away the original source of a message is or how likely others are to alter messages. Thus, we analyze situations in which the learner cannot tell how far a message has traveled and faces uncertainty about mutation rates.[‖] With this in mind, we assume that the learner follows the majority of messages received. This may be seen as capturing learners who follow a simple rule of thumb; it can also be justified as a reduced-form model of Bayesian agents.[**] The analysis readily extends to nonsymmetric settings where there is some threshold applied that differs from a majority (e.g., the learner has a status quo bias and presumes it is state 0 unless messages of 1 outnumber messages of 0 by some number or by some ratio), but the symmetric case makes the expressions most transparent.

[‡]In another paper (7), we analyze social learning when the pass rate p depends on message content and characterize conditions under which learning can be enhanced by such dependence. However, those results shed little additional light on the role of network structure, and so we focus here on the case in which all messages have the same pass rate.

[§]In practice, agents pass along messages for a variety of reasons, be it to inform or persuade or mislead others, tell a joke, or even to just keep a record of what they heard. Since our goal is to understand the type of content that reaches a learner, we do not explicitly model agents' incentives for passing along messages. However, given any microfoundation of senders' incentives, the resulting average probabilities of mutations and message dropping at each step is what matters for the analysis.

[¶]Also, we ignore the possibility of the learner being an original source and just examine the marginal information that the learner gets from the network. Again, a direct extension addresses the case where the learner can be an original source, without any additional intuition but with some additional notation and two cases.

[#]See the proof of *Lemma* for a derivation.

[‖]In ref. 7, we characterize how agents learn with known mutation rates and known chain lengths. We show that learning from distant messages requires not only sufficient growth in the network but also that the learner has no uncertainty about the mutation rates. With any uncertainty about mutation rates, it is impossible to learn from any number of messages if they have traveled a sufficiently long distance.

[**]For instance, suppose a learner faces uncertainty about the mutation rates and has a symmetric prior on the state and on mutation rates. Moreover, suppose after receiving messages, the learner takes a binary action $a \in \{0, 1\}$ that yields payoff $+1$ if it matches the state or payoff -1 if not. Such an agent would, after Bayesian updating, find it optimal to choose the action corresponding to the most common message received.

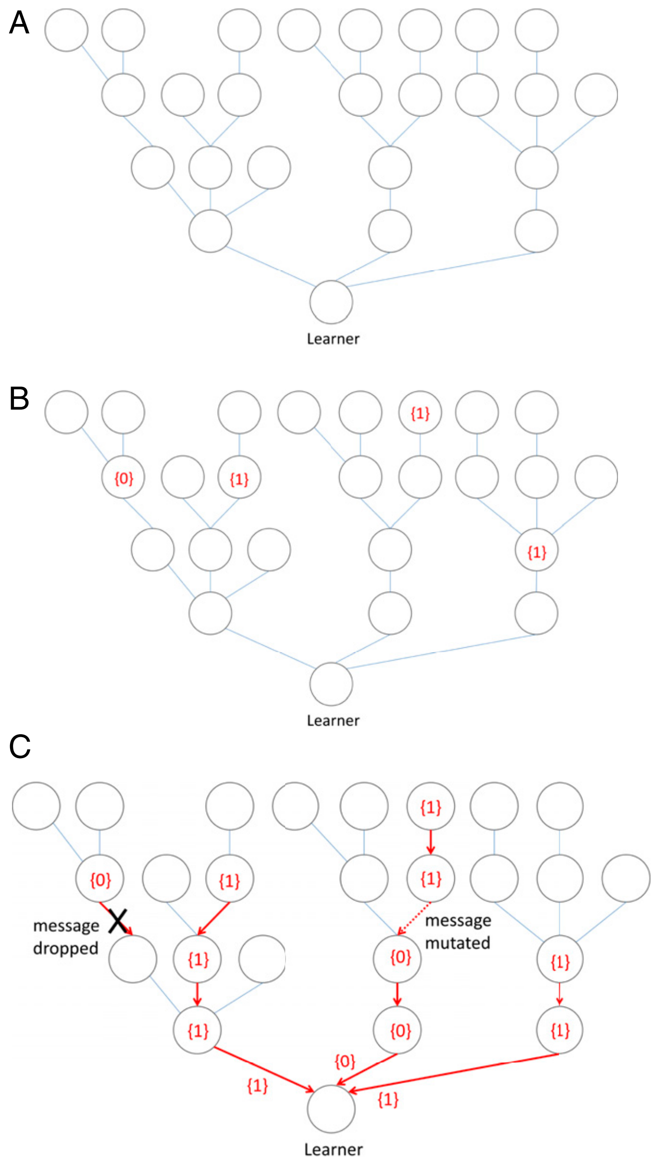

Fig. 1. An illustration of the model. The learner is at the root of a tree. From four original sources, messages are passed along with the possibility of dropping and mutation. The learner ends up hearing three messages, of which one mutated en route. (*A*) A random tree with a learner and 24 other agents who could pass information to the learner. (*B*) Randomly, with probability *r* some of the agents get information. Here three nodes get initial information of 1, and one gets initial information of 0, while others get no direct information. (*C*) Messages pass to the learner. One of the messages was randomly dropped, while three others made it to the learner. One message mutated along its path.

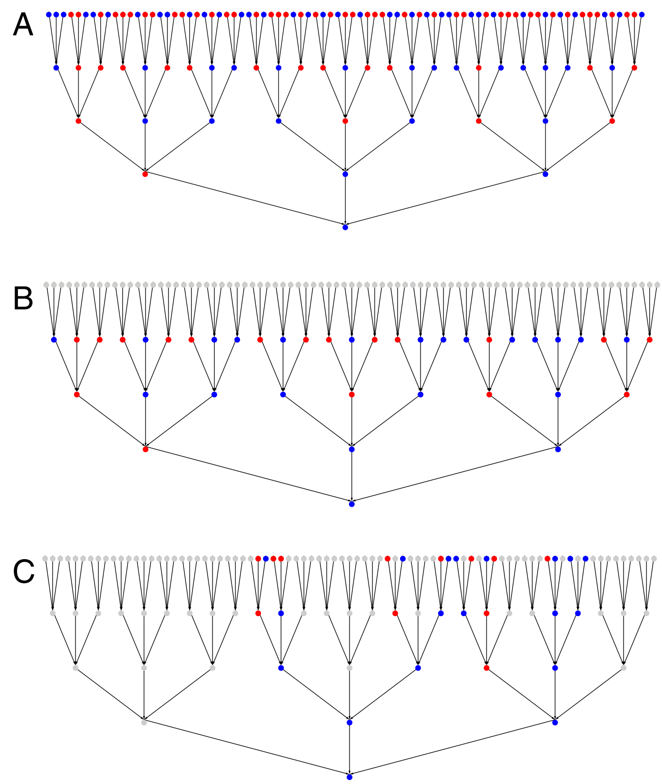

Fig. 2. A tree in which each node has three neighbors via whom messages are relayed and the learner is at the root. Blue nodes are those from which the signal gets to the learner in its true form: matching the state. Red nodes are those from which the signal is eventually corrupted and reach the learner in a false form: not matching the state. (This figure differs from Fig. 1 in that we do not show how messages mutate along the route but just color them at their point of origin based on whether they are true by the time they reach the learner.) Under the caps in *B* and *C*, the gray nodes are those whose messages no longer reach the learner. The parameter values are $p = 1$ so that there is no dropping, $\mu_{10} = 0.15$, $\mu_{01} = 0$, and the state is 1. (*A*) Without any cap, the learner hears 58 correct and 63 incorrect messages. More distant sources have a greater fraction of incorrect messages. (*B*) A depth cap blocks messages from nodes at a distance of 4 or greater, for which messages are more often incorrect than correct. The learner now hears 22 correct and 18 incorrect messages. (*C*) A breadth cap limits the extent to which messages are forwarded, so each person now only hears from two others on average. This increases the relative fraction of near to far sources. The learner hears 19 correct and 12 incorrect messages.

Improving Learning by Restricting Communication

We analyze how restricting the distance that messages can travel (depth) or reducing the network's degree (breadth) improves learning. To gain intuition, consider a situation in which $pk > 1$ so that the learner receives exponentially more messages from farther away. Whenever the mutation rates μ_{01}, μ_{10} are unequal, a majority of messages from sufficiently far away are more likely to match the mutation-favored state (state 0 if $\mu_{01} < \mu_{10}$ or state 1 if $\mu_{01} > \mu_{10}$), causing the learner to take the mutation-favored action regardless of the true state. The learner would be better off in a network with no retransmission at all: depth capped at $T = 1$. However, such a draconian restriction might not be optimal since the learner would receive too few messages.

We examine how to optimally restrict communication with respect to two related objectives. We first focus on the objective of maximizing the expected number of true minus false messages, which allows for a tractable analysis. We then turn to the objective of maximizing expected learner welfare, which simplifies to maximizing the probability that the majority of received messages match the true state. We show in a numerical example that the main qualitative findings continue to hold.[††]

Before offering results outlining breadth and depth caps that maximize the number of true minus false messages, we illustrate why such caps work in a simple example, as illustrated in Fig. 2. This is a simplified setting in which there is no dropping and the tree is regular and of depth four to begin with. Parameters are such

[††]Whenever the number of potential sources that are not too far from the learner is reasonably large, maximizing the expected number of true minus false messages also approximately maximizes the probability that true messages outnumber false ones. The reason, effectively, is that increasing how far the mean of the distribution of excess number of true messages is above zero also tends to reduce its tail probability of being below zero. Extra complications arise in our context, however, because depth and breadth caps impact the variance of this distribution, and these caps are themselves subject to an integer constraint.

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.

 https://doi.org/10.1073/pnas.2205549119

that messages are more likely to be true than false as long as they do not travel more than a distance of three. The optimal depth cap is to block messages traveling four or more, which changes the messages from being a majority false to be a majority true. Similarly, restricting breadth also increases the relative number of close to far messages and leads to a majority of true messages.

Say that a set of sources generates a "preponderance of true messages" if the expected number of true messages from these sources exceeds the expected number of false messages, regardless of the state. Whether the network generates a preponderance of true messages depends on the mutation rates (and other model parameters), which may differ depending on message topic, etc.

More generally, a platform may want there to be a preponderance of true messages for a set of possible mutation rates. The platform may be uncertain about which rates apply, may face varying conditions and mutation rates, or may have different rates across different topics that are transmitted across the same network. To deal with this, we can require a preponderance of true messages to hold with respect to a whole set of mutation rates. Let A be a compact set of pairs of mutation rates (μ_{01}, μ_{10}). Say that a set of sources generates a robust preponderance of true messages if there are more true than false stories from these sources, in expectation, for every pair of mutation rates $(\mu_{01}, \mu_{10}) \in A$ and regardless of the state. In a platform designed to generate a robust preponderance of true messages, a learner who simply follows the majority of messages expects that majority to be correct on average, for every pair of mutation rates in A and regardless of the state.

A depth (breadth) cap is undominated if there is no other depth (breadth) cap that generates a strictly higher expected number of true minus false messages in the mutation-disfavored state for some $(\mu'_{01}, \mu'_{10}) \in A$ and weakly higher for all $(\mu_{01}, \mu_{10}) \in A$.

A platform facing mutation rates in A with a goal of maximizing the expected number of true minus false messages would choose an undominated cap. Which one of the undominated caps they would choose would then depend on some more specific weightings on the possible mutation rates in A.

Capping Depth. *Proposition 1* establishes a critical distance such that messages traveling no more than that distance are always more likely to be true than false. This determines the set of caps that are undominated.

A rounding function is a mapping $r : \mathbb{R} \to \mathbb{Z}$ such that $r(x) \in \{\lfloor x \rfloor, \lceil x \rceil\}$ for all $x \in \mathbb{R}$.

Proposition 1. *Any given source generates a robust preponderance of true messages if and only if it is at distance from the learner of no more than* $T^* \equiv \min_{(\mu_{01},\mu_{10})\in A} T(\mu_{01}, \mu_{10})$, *where*

$$T(\mu_{01}, \mu_{10}) \equiv \frac{\log\left(\frac{1}{2}\right) + \log\left(1 - \min\left\{\frac{\mu_{01}}{\mu_{10}}, \frac{\mu_{10}}{\mu_{01}}\right\}\right)}{\log(1 - \mu_{01} - \mu_{10})}. \quad [3]$$

There exists a rounding function r such that a depth cap T is undominated if and only if $T \in [r(T^*), r(T^{**})]$,[‡‡] *where* $T^{**} \equiv \max_{(\mu_{01},\mu_{10})\in A} T(\mu_{01}, \mu_{10})$.

The proof of *Proposition 1* is relatively straightforward, and so we outline it here. The first part of *Proposition 1* follows from considering the values of T for which the ratios in Eqs. **1** or **2** are less than 1 for some pair of mutation rates. After some algebraic manipulations, that ratio is equal to 1 for the $T(\mu_{01}, \mu_{10})$ defined in Eq. **3**. Beyond that distance, the expected number of false messages exceeds true ones, and at less than that distance the reverse is true.

The cap of T^* is a conservative one, and it might be that a platform designer is willing to admit some additional false messages for some mutation rates in order to improve the expected number of true minus false for some other mutation rates. Beyond T^{**}, however, more false than true messages are then expected for every pair of mutation rates in A. This leads us to the latter statement of *Proposition 1*.

The magnitude of $T(\mu_{01}, \mu_{10})$, and hence the caps T^*, T^{**}, depends on how asymmetric the mutation rates are. On one extreme, when μ_{01} and μ_{10} are equal, it follows that $T(\mu_{01}, \mu_{10}) = \infty$. On the other extreme, when one of the mutation rates is near $1/2$ and the other is near zero, it follows that $T(\mu_{01}, \mu_{10})$ is close to 1. Thus, presuming that the designer is concerned about some asymmetric mutation rates—as those are the ones that upset learning—the designer will choose some finite $T \geq T^*$.

In principle, the set of possible mutation rates A can be determined by a platform that has data on different types of messages, what we will refer to as "topics." Setting a depth cap of T^* ensures that the expected number of true minus false messages is positive on every topic and in every state and is maximized for some topic. That said, if the platform designer can set different depth caps (and knows the mutation rates) for each topic, then setting a customized depth cap of $T(\mu_{01}, \mu_{10})$ for each topic maximizes the expected number of true minus false messages that users receive, topic by topic.

Capping Breadth. If the depth of the network has been capped at $T \leq T^*$, all received messages are more likely to be true than false. A platform designer intent on maximizing the expected number of true minus false messages would therefore not restrict breadth d.

However, what if the network's depth T cannot be capped? Capping T requires following the life cycle of a message, which can be infeasible in practice. For instance, it may be difficult or impossible to tell whether someone is mutating a previous message she heard or originating a new message. Alternatively, a designer maintaining user privacy may not even observe the messages being forwarded, which are encrypted on some platforms. When depth cannot be restricted, limiting breadth can help ensure that most messages are true. A designer can lower the relative number of long chains by restricting the number of others to whom any given agent can send or forward messages.

Consider again the infinite tree network with no cap on depth in which each node has average in-degree and out-degree $k > 0$. A cap on message forwards of $k^* \leq k$ limits each agent to passing along messages to at most k^* of its out-neighbors (Fig. 2). We suppose that agents follow such a cap by choosing their k^* out-neighbors uniformly at random. (The cap is not binding if the agent has k^* or fewer friends.)

Proposition 2. *An infinite-depth network generates a robust preponderance of true messages if and only if its degree is less than*

$$\overline{k} = \min_{(\mu_{01},\mu_{10})\in A} \frac{1 - 2\max\{\mu_{01}, \mu_{10}\}}{p(1 - \mu_{01} - \mu_{10})}.$$

The breadth cap

$$k(\mu_{01}, \mu_{10}) = \frac{1 - Z}{p\left(1 - (1 - \mu_{01} - \mu_{10})Z\right)}, \quad [4]$$

[‡‡] Caps in $[\lceil T^* \rceil, \lfloor T^{**} \rfloor]$ are always undominated; whether the caps $\lfloor T^* \rfloor$ or $\lceil T^{**} \rceil$ are undominated can depend on parameter values.

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.

where

$$Z = \left(\frac{\max(\mu_{10}, \mu_{01}) - \min(\mu_{10}, \mu_{01})}{2\max(\mu_{10}, \mu_{01})(1 - \mu_{01} - \mu_{10})} \right)^{1/2} \quad [5]$$

maximizes the expected number of true minus false messages in the unfavored state for mutation rates (μ_{01}, μ_{10}). *Let* $k^* \equiv \min_{(\mu_{01},\mu_{10})\in A} k(\mu_{01}, \mu_{10})$ *and* $k^{**} \equiv \max_{(\mu_{01},\mu_{10})\in A} k(\mu_{01}, \mu_{10})$. *It follows that there exists a rounding function r such that a breadth cap k is undominated if and only if* $k \in [r(k^*), r(k^{**})]$.

Proposition 2 is not an exact parallel to the result on depths, *Proposition 1*. The reason is that a depth cap produces an absolute cap on which nodes are in or out of consideration based on their distance from the learner. In contrast, a breadth cap does not rule out any distance absolutely but instead controls the relative frequency of nodes at various distances. Thus, the robust preponderance is stated relative to the tree in the case of a breadth cap, while it is stated for the source nodes in the case of a depth cap. This also makes *Proposition 2* harder to prove and leads to different expressions depending on whether we are considering a robust preponderance or undominated caps.

Note that pk captures the expansion property of a network with average degree k, with far-away messages dominating what the learner receives if $pk > 1$. Not surprisingly, the cap $k(\mu_{01}, \mu_{10})$ is thus always low enough so that $pk(\mu_{01}, \mu_{10}) \leq 1$, with $pk(\mu_{01}, \mu_{10}) = 1$ only in the special case with symmetric mutation rates, $\mu_{01} = \mu_{10}$.[§§] (To see why, note that $Z = 0$ if and only if $\mu_{01} = \mu_{10}$.)

As the likelihood p that messages are passed is decreased, fewer messages come from longer distances and the optimal cap is higher, meaning that people are less constrained in how widely they can forward any given message. In the extreme as p approaches 1, so messages are likely to travel far, the optimal cap is set such that the expected degree is less than 1.

If we increase the higher of the two mutation rates, then Z increases and $(1 - \mu_{01} - \mu_{10})$ decreases. This implies that $\frac{1-Z}{1-(1-\mu_{01}-\mu_{10})Z}$ decreases since the denominator shrinks more slowly, and a tighter cap is needed. In contrast, when the lower of the two mutation rates increases (but not so much to pass the other), a looser cap is possible. In other words, caps are less useful when mutation rates come more into balance.

As can be easily checked, the breadth cap $\overline{k}$ always exceeds the cap k^*. This difference highlights a potential conflict of interest between platforms and their users. In particular, if a platform seeks to maximize its overall volume of communication, subject to the constraint that there is a robust preponderance of the truth, then the platform might set a looser cap on network breadth than a social planner seeking to maximize learner welfare.

Probabilities that True Outnumber False Messages. The above results examine the expected number of true minus false messages, but more relevant for some learners is the probability that true messages exceed false ones. These objectives are closely related but do not always yield the same caps. The probability that true messages exceed false is not tractable in closed form, but we observe through numerical simulations that the comparative statics of varying network depth/breadth are similar to those of the caps above. In particular, the probabilities of having true messages exceed false ones for different combinations of d and T (for k-regular trees) are pictured in Fig. 3.

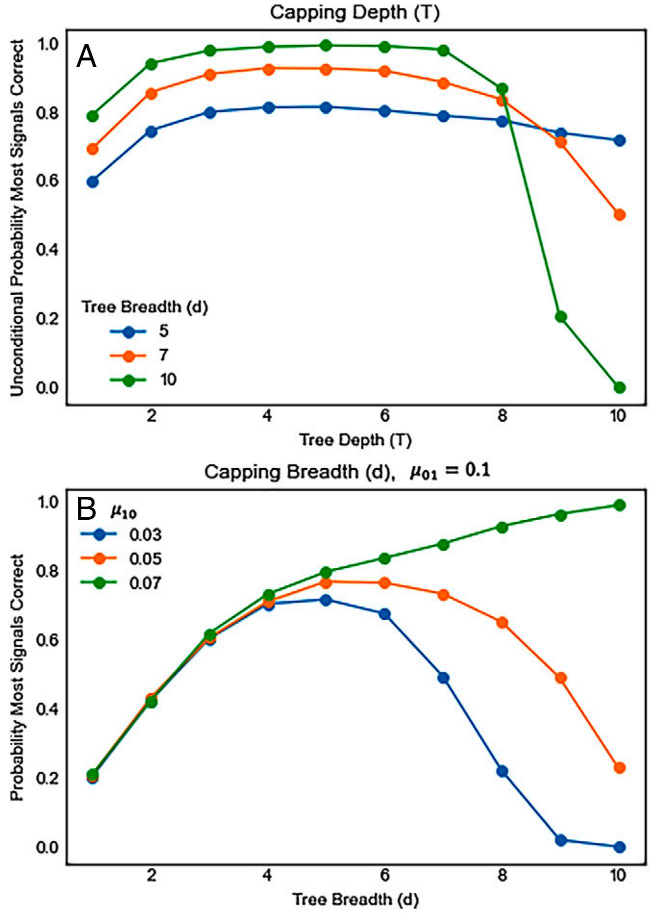

Fig. 3. These graphs show how the probability that most received messages are true in state 0 (the state more likely to have mutations) varies with the depth and breadth of the learner's social network. The parameters used in these simulations are $p = 0.2$, $\mu_{01} = 0.10$, and $\mu_{10} = 0.03$ for *A* and $p = 0.2$, $\mu_{01} = 0.10$, and $\mu_{10} \in \{0.03, 0.05, 0.07\}$ for *B*. (*A*) Probability that more messages are true than false in state 0, as a function of the cap on tree depth for three values of breadth. (*B*) Probability that more messages are true than false in state 0 for three relative mutation rates (all with depth 10).

In the numerical examples illustrated in Fig. 3, messages of any distance are more likely to be true than false in state 1, the mutation-favored state. However, once a message travels too far, it is more likely to be false than true in state 0. By *Proposition 1*, the threshold depth at which false messages on average overtake true ones is $T^* = 7.5$, regardless of network breadth. The breadth cap in *Proposition 2* that maximizes the extent to which true messages outnumber false ones is $k^* = 1.83$, while the cap in *Proposition 2* that ensures a robust preponderance of the truth is $\overline{k} = 4.6$.

Those thresholds do not maximize the probability that true messages outnumber false ones. Nonetheless, as is clear from Fig. 3, these are also close to the maximizers of probability.[¶¶]

Fig. 3 also shows that as the mutation rates become closer to each other, it could be optimal to avoid capping breadth. This illustrates how so the confound in learning originates from the asymmetry in mutation rates, which increasingly biases information the longer it travels.

[§§]In the special case with symmetric mutation rates and no dropping ($p = 1$), the optimal breadth limit $k^* = 1$, meaning that the optimally pruned tree is a single infinite chain. Otherwise, whenever $\mu_{01} \neq \mu_{10}$, the optimal breadth limit prunes the tree sufficiently that the expected number of sources is always finite.

[¶¶]The breadth figures are all done for a fixed depth of $T = 10$ rather than an infinite depth, for the sake of tractable simulations.

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.

Discussion

Cutting down on fake news has recently become an expressed objective for social media platforms that 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 the 2018 Brazilian elections. WhatsApp subsequently placed a cap of five on the number of people to whom a message can be forwarded in an attempt to curtail the spread of false information (12).

Our results above show platforms can reduce the relative frequency of false communications by limiting either the depth or breadth of people's communication networks. These policies enable agents to partially learn the true state in the presence of asymmetric mutations, by ensuring that relatively more of an agent's received messages originate from nearby.

Most importantly, such policies improve learning without requiring message content to be observed, an attractive feature for platforms intent on respecting user privacy. Even if messages are observable, a platform designer who does not know the true state of the world may not be able to discern true messages from false ones or may prefer to take a content-neutral stance and avoid censoring messages.

Of course, as with any model, there are important omissions, and we mention a couple here.

Tagging and Tracing. The challenges of learning from word-of-mouth communication can motivate learners to seek out information from closer, trusted contacts. Platforms differ in the ways that messages are forwarded and the extent to which a learner can trace the path back partway or all the way to the origin, as well as the extent to which this is easily disguised. To the extent that a platform can make it easy to trace the length of a chain and a message's history, that could also enhance learning. However, tracing information backward is difficult in many cases, especially for platforms with encrypted messages. In such cases, caps on the breadth of users' word-of-mouth networks can be powerful tools and relatively easily implemented. Nonetheless, studying how the ability to partially trace messages' transmission chains can help learners and the extent to which learners would take advantage of such an ability in practice are important topics for further research.

Homophily. A common feature of many social networks is that nodes tend to be more similar to neighbors than they are to those more distant from them. In communication networks, nearby nodes may share similar tendencies to mutate or drop messages. For example, one cluster of users may be primarily responsible for exaggerating information in one direction, while another group may be responsible for another.

Given our focus on robust bounds for a set of mutation rates, the caps we derive can account for this sort of heterogeneity. Nevertheless, there are interesting questions for further study. Optimizing learning within homophilous networks involves more than just trading off signal and noise (e.g., ref. 13). For instance, if clusters of nodes in different parts of the network tend to mutate messages in different directions, then loosening caps on depth and/or breadth could help receivers hear messages from other groups. Learning could then be nonmonotone in the caps.

Incentives. Our results imply that a profit-motivated platform might still prefer different looser caps than would be socially optimal if the platform profits from message traffic and/or conflict among its users. Studying how profit-motivated platforms would behave is an important future topic. (For some analysis of learning on strategic platforms, see refs. 14–17.)

Our model captures deliberate misinformation sources as appearing randomly in the network. They may have some discretion on their position. Fully analyzing sources' incentives when originating information is a topic for further study. Relatedly, if sources that wish to deliberately manipulate beliefs have choices of where to appear in the network or simply happen to appear more frequently in some parts of the network, that could introduce asymmetries in exposure to true versus false information across learners. This could result in more complex policies, which is also an interesting topic for further study.

Also, caps reduce the scope for false messages to go viral. To the extent that this diminishes ideological sources' incentive to originate false content more than it dissuades reliable sources from originating true content, the positive impact of caps is understated in our model, which treats information arrival as exogenous.

Proofs

The proof of *Proposition 1* is straightforward and outlined in the main text. Here we provide the proof of *Proposition 2*. We preface the proof with a lemma that provides some useful Markov chain formulas.

Lemma. *Suppose that $p > 0$ and consider any pair of mutation rates $\mu_{01}, \mu_{10} \in (0, 1/2)$. If the state is 0 and agent $t \geq 1$ receives a nonnull message, then the message is 0 (matching the true state) with probability*

$$X_0(t) = \Pr[m_t = 0 | m_t \neq \emptyset, \omega = 0] = \frac{\mu_{10} + \mu_{01} M^t}{\mu_{10} + \mu_{01}}.$$

If the state is 1 and agent $t \geq 1$ receives a nonnull message, the message is 1 (matching the true state) with probability

$$X_1(t) = \Pr[m_t = 1 | m_t \neq \emptyset, \omega = 1] = \frac{\mu_{01} + \mu_{10} M^t}{\mu_{01} + \mu_{10}}.$$

It follows that

$$X_0(t) - (1 - X_1(t)) = M^t.$$

As t grows, regardless of the starting state, the probabilities that a surviving message is a 0 and a 1, respectively, are

$$\pi_0 = \frac{\mu_{10}}{\mu_{10} + \mu_{01}} \text{ and } \pi_1 = \frac{\mu_{01}}{\mu_{10} + \mu_{01}}.$$

Finally, if $\mu_{01} = \mu_{10} = \mu$, then the message matches the true state with probability

$$X(t) = \frac{1 + M^t}{2}. \quad [6]$$

Proof of Lemma: We derive the expressions of X_0, X_1, which can also be deduced from standard Markov chain results, but it may be useful for the reader to see the derivation. The proof is by induction. We give the proof for X_0, when the state is 0. The proof for X_1 is symmetric, and the expression for X is a special case.

First, note that if $t = 1$, then this expression simplifies to $1 - \mu_{01}$, which is exactly the probability that the message has not mutated, and so this holds for $t = 1$.

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

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.

The probability of matching the true state at t is the probability of not matching at $t-1$ times μ_{10} plus the probability of matching at $t-1$ times $1-\mu_{01}$, which by the induction assumption can be written as

$$
\begin{aligned}
&\left[1-\frac{\mu_{10}+\mu_{01}M^{t-1}}{\mu_{10}+\mu_{01}}\right]\mu_{10}+\left[\frac{\mu_{10}+\mu_{01}M^{t-1}}{\mu_{10}+\mu_{01}}\right](1-\mu_{01})\\
&=\mu_{10}+\left[\frac{-\mu_{10}^2-\mu_{01}\mu_{10}M^{t-1}+\mu_{10}-\mu_{10}\mu_{01}+\mu_{01}M^{t-1}-\mu_{01}^2M^{t-1}}{\mu_{10}+\mu_{01}}\right]\\
&=\mu_{10}+\left[\frac{-\mu_{10}^2+\mu_{10}-\mu_{10}\mu_{01}+\mu_{01}M^{t-1}(1-\mu_{10}-\mu_{01})}{\mu_{10}+\mu_{01}}\right]\\
&=\frac{\mu_{10}+\mu_{01}M^t}{\mu_{10}+\mu_{01}}
\end{aligned}
$$

as claimed. ■

Proof of Proposition 2: Suppose without loss of generality that $\mu_{01}\le\mu_{10}$. Given this ordering, the worst case for the number of true minus false messages occurs when the true state is 1.

We first develop the expression for $k(\mu_{01},\mu_{10})$ and then the expression for $\overline{k}$.

When the average degree of the Galton–Watson tree is k, it follows by induction and iterated expectations that there are k^t nodes at a distance t away from the learner, in expectation. Therefore, the expected number of true messages received is $\sum_{t=1}^{\infty} r(pk)^t X_1(t)$, and the expected number of false messages received is $\sum_{t=1}^{\infty} r(pk)^t(1-X_1(t))$. The expected number of true minus false messages is

$$\sum_{t=1}^{\infty} r(pk)^t\left(X_1(t)-(1-X_1(t))\right).$$

Thus, we need to maximize $\sum_{t=1}^{\infty}(pk)^t(2X_1(t)-1)$, i.e.,

$$\sum_{t=1}^{\infty}(pk)^t\left(\frac{2\mu_{10}M^t+\mu_{01}-\mu_{10}}{\mu_{10}+\mu_{01}}\right).$$

The first-order conditions are

$$\sum_{t=1}^{\infty} tp^tk^{t-1}\left(\frac{2\mu_{10}M^t+\mu_{01}-\mu_{10}}{\mu_{10}+\mu_{01}}\right)=0. \qquad [7]$$

Multiplying this by k, this is equivalent to

$$\sum_{t=1}^{\infty} tp^tk^{t}\left(\frac{2\mu_{10}M^t+\mu_{01}-\mu_{10}}{\mu_{10}+\mu_{01}}\right)=0$$

or

$$\sum_{t=1}^{\infty} tp^tk^tM^t\left(\frac{2\mu_{10}}{\mu_{10}+\mu_{01}}\right)-\sum_{t=1}^{\infty} tp^tk^t\left(\frac{\mu_{10}-\mu_{01}}{\mu_{10}+\mu_{01}}\right)=0.$$

Noting that $\sum_{t=1}^{\infty} tz^t=z/(1-z)^2$, we rewrite the above as

$$\frac{pkM2\mu_{10}}{(1-pkM)^2}-\frac{pk(\mu_{10}-\mu_{01})}{(1-pk)^2}=0.$$

Rearranging terms to solve for k leads to the claimed expression expression for $k(\mu_{01},\mu_{10})$. To check the second-order conditions, note that the second derivative is

$$\sum_{t=2}^{\infty}(t-1)tp^tk^{t-2}\left(\frac{2\mu_{10}M^t+\mu_{01}-\mu_{10}}{\mu_{10}+\mu_{01}}\right).$$

Multiplying this by k does not change the sign and gives

$$\sum_{t=1}^{\infty}(t-1)tp^tk^{t-1}\left(\frac{2\mu_{10}M^t+\mu_{01}-\mu_{10}}{\mu_{10}+\mu_{01}}\right),$$

which is the same as Eq. 7 but where the weights are adjusted by $(t-1)$. Since the expression $\left(\frac{2\mu_{10}M^t+\mu_{01}-\mu_{10}}{\mu_{10}+\mu_{01}}\right)$ is decreasing in t (and eventually negative), then reweighting with $(t-1)$ tp^tk^{t-1} compared to tp^tk^{t-1} puts more weight on higher terms (in a dominance sense), and since Eq. 7 is 0, the second derivative expression is necessarily negative.

Next, we develop the expression for $\overline{k}$, again presuming that $\mu_{01}\le\mu_{10}$, as the expressions are symmetric.

By *Lemma*, the expected number of true messages exceed the expected number of false messages when the state is 0 since for every t, $X_0(t)=\frac{\mu_{10}+\mu_{01}M^t}{\mu_{10}+\mu_{01}}>\frac{\mu_{01}-\mu_{01}M^t}{\mu_{10}+\mu_{01}}=1-X_0(t)$.

It remains to be shown that the expected number of true messages also exceed the expected number of false messages when the state is 1. Using the expressions from the proof above, we need to characterize the conditions under which

$$\sum_{t=1}^{\infty} r(pk)^tX_1(t)>\sum_{t=1}^{\infty} r(pk)^t(1-X_1(t)).$$

Thus, we need $\sum_{t=1}^{\infty}(pk)^t(2X_1(t)-1)>0$, i.e., when

$$\sum_{t=1}^{\infty}(pk)^t\left(\frac{2\mu_{10}M^t+\mu_{01}-\mu_{10}}{\mu_{10}+\mu_{01}}\right)>0.$$

If $\mu_{01}=\mu_{10}$, then this holds regardless of k.

Next, suppose $\mu_{01}<\mu_{10}$. Then it follows that

$$
\begin{aligned}
&\sum_{t=1}^{\infty}(pk)^t\left(\frac{2\mu_{10}M^t+\mu_{01}-\mu_{10}}{\mu_{10}+\mu_{01}}\right)>0\\
&\iff \sum_{t=1}^{\infty}(pMk)^t(2\mu_{10})>\sum_{t=1}^{\infty}(pk)^t(\mu_{10}-\mu_{01})\\
&\iff \frac{pMk}{1-pMk}(2\mu_{10})>\frac{pk}{1-pk}(\mu_{10}-\mu_{01})\\
&\iff \frac{2\mu_{10}M}{1-pMk}>\frac{\mu_{10}-\mu_{01}}{1-pk}\\
&\iff \frac{1-pMk}{2\mu_{10}M}<\frac{1-pk}{\mu_{10}-\mu_{01}}\\
&\iff \frac{1}{2\mu_{10}M}-\frac{pk}{2\mu_{10}}<\frac{1}{\mu_{10}-\mu_{01}}-\frac{pk}{\mu_{10}-\mu_{01}}
\end{aligned}
$$

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.

$$
\begin{aligned}
&\iff \frac{1}{p}\left(\frac{1}{2\mu_{10}M}-\frac{1}{\mu_{10}-\mu_{01}}\right)\\
&\qquad < k\left(\frac{1}{2\mu_{10}}-\frac{1}{\mu_{10}-\mu_{01}}\right)\\
&\iff k < \frac{1}{p}\frac{\frac{\mu_{10}-\mu_{01}}{M}-2\mu_{10}}{\mu_{10}-\mu_{01}-2\mu_{10}}\\
&\iff k < \frac{1}{p}\frac{2\mu_{10}+\frac{\mu_{01}-\mu_{10}}{M}}{\mu_{10}+\mu_{01}}\\
&\iff k < \frac{1}{pM}\frac{\mu_{10}+\mu_{01}-2\mu_{10}(\mu_{10}+\mu_{01})}{\mu_{10}+\mu_{01}}\\
&\iff k < \frac{1-2\mu_{10}}{pM}.
\end{aligned}
$$

This is the condition in *Proposition 2* for $\mu_{01} \leq \mu_{10}$, which was without loss of generality and thus leads to the claimed expression for $\overline{k}$, concluding the proof. ■

Data, Materials, and Software Availability. There are no data underlying this work.

ACKNOWLEDGMENTS. We gratefully acknowledge financial support under NSF grants SES-1629446 and SES-2018554 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.

Author affiliations: [a]Department of Economics, Stanford University, Stanford CA, 94305; [b]Santa Fe Institute, Santa Fe, NM 87501; [c]Department of Economics, Cornell University, Ithaca, NY 14853; and [d]Fuqua School of Business and Economics Department, Duke University, Durham, NC 27708

1. S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online. *Science* **359**, 1146–1151 (2018).
2. L. A. Adamic, T. M. Lento, E. Adar, P. C. Ng, "Information evolution in social networks" in *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining* (Association for Computing Machinery, 2016), pp. 473–482.
3. P. J. Gurry, The number of variants in the Greek new testament: A proposed estimate. *New Testam. Stud* **62**, 97–121 (2016).
4. M. P. Simmons, L. A. Adamic, E. Adar, Memes online: Extracted, subtracted, injected, and recollected. *Proc. Annu. Int. Conf. Mob. Comput. Netw.* **11**, 17–21 (2011).
5. D. Acemoglu, A. Ozdaglar, A ParandehGheibi, Spread of (mis) information in social networks. *Games Econ. Behav.* **70**, 194–227 (2010).
6. F. Bloch, G. Demange, R. Kranton, Rumors and social networks. *Int. Econ. Rev.* **59**, 421–448 (2018).
7. M. O. Jackson, S. Malladi, D. McAdams, The fragility of social learning in the presence of noisy messages. arXiv:1812.03354v1 (2018).
8. A. G. Chandrasekhar, H. Larreguy, J. P. Xandri, Testing models of social learning on networks: Evidence from two experiments. *Econometrica* **88**, 1–32 (2020).
9. L. K. Fazio, N. M. Brashier, B. K. Payne, E. J. Marsh, Knowledge does not protect against illusory truth. *J. Exp. Psychol. Gen.* **144**, 993–1002 (2015).
10. M. Serra-Garcia, U. Gneezy, Mistakes, overconfidence, and the effect of sharing on detecting lies. *Am. Econ. Rev.* **111**, 3160–3183 (2021).
11. M. Posner, How social media companies need to address disinformation globally, Forbes, Forbes Magazine, 16 June 2019.
12. A. Hern, M. Safi, Whatsapp puts limit on message forwarding to fight fake news. *The Guardian*, GuardianNews and Media, 21 January 2019.
13. I. Lobel, E. Sadler, Preferences, homophily, and social learning. *Oper. Res.* **64**, 564–584 (2016).
14. M. Mueller-Frank, M. Pai, "Do online social networks increase welfare?" (IESE Business School Working Paper No. WP-1118-E, 2014).
15. O. Candogan, K. Drakopoulos, Optimal signaling of content accuracy: Engagement vs. misinformation. *Oper. Res.* **68**, 497–515 (2020).
16. D. Acemoglu, A. Ozdaglar, J. Siderius, A model of online misinformation. NBER working paper 28884. 10.3386/w28884 (2022).
17. O. Candogan, N. Immorlica, B. Light, J. Anunrojwong, Social learning under platform influence: Consensus and persistent disagreement. arXiv [Preprint] (1 March 2022). https://arxiv.org/abs/2202.12453.

Downloaded from https://www.pnas.org by STANFORD UNIVERSITY SERIALS DEPARTMENT on July 18, 2023 from IP address 171.66.162.18.