Learning in a Post-Truth World

Abstract

Misinformation has emerged as a major societal challenge in the wake of the 2016 U.S. elections, Brexit, and the COVID-19 pandemic. One of the most active areas of inquiry into misinformation examines how the cognitive sophistication of people impacts their ability to fall for misleading content. In this paper, we capture sophistication by studying how misinformation affects the two canonical models of the social learning literature: sophisticated (Bayesian) and naive (DeGroot) learning. We show that sophisticated agents can be *more* likely to fall for misinformation. Our model helps explain several experimental and empirical facts from cognitive science, psychology, and the social sciences. It also shows that the intuitions developed in a vast social learning literature should be approached with caution when making policy decisions in the presence of misinformation. We conclude by discussing the relationship between misinformation and increased partisanship, and provide an example of how our model can inform the actions of policymakers trying to contain the spread of misinformation.

1 Introduction

Misinformation has become an inextricable component of how people learn about the world and make decisions. Persistent disagreement over objective facts is becoming increasingly commonplace (Alesina et al., 2020; Bursztyn et al., 2020), leading to the assertion that we currently live in a post-truth, "alternative facts" world (McIntyre, 2018). This makes it particularly difficult when society faces a collective action problem whose outcome depends on a substantial proportion of the population agreeing to take a specific action, e.g., vaccination or wearing a mask. Indeed, Loomba et al. (2021) show that in the U.S. and U.K., an average decline of 6.2 percentage points in the acceptance of the COVID-19 vaccine is directly attributable to misinformation.

What makes people believe false claims? A growing empirical and experimental literature argues that it depends on their cognitive sophistication.[1] Pennycook and Rand (2018) show that more sophisticated agents are less likely to fall for misinformation. At the same time, there is

[1] From Tappin et al. (2020): "indicators of cognitive sophistication [can include] educational attainment, science literacy, numeracy, specific topic knowledge, and a propensity for analytic thinking." This is also correlated with performance on the Cognitive Reflection Test (Pennycook and Rand, 2018). Pennycook et al. (2021) show that getting people to think more carefully about the accuracy of the news they read can lead to better discernment of false information.

ample evidence that sophisticated agents are more likely to disagree over objective facts.[2] This presents an interesting puzzle: if sophisticated agents are more likely to learn the truth, then why do we observe more disagreement within that group over what the truth is? Competing explanations argue that this disagreement arises from political polarization and partisan bias (e.g., Taber and Lodge (2006); Taber et al. (2009); Kahan et al. (2017)), or can simply be explained as an outcome of unbiased Bayesian reasoning (e.g., the recent experiments in Tappin et al. (2020)).

In this paper, we present a simple model that helps explain and reconcile the above empirical findings. These findings cannot be explained using classic social learning results, which until recently,[3] were developed under the assumption that agents learn from *organic* news that can be inaccurate but contains no deliberate misinformation. Our results depart from the conventional wisdom of social learning and economic modeling in general — that the rational economic agent is the standard against which all agent models should be measured.[4] Our model therefore serves as a foundation for further theoretical and experimental work that integrates misinformation into social learning, and can also be used to inform policy decisions on these issues.

We capture sophistication heterogeneity by studying how misinformation affects the two canonical models of the social learning literature: Bayesian (e.g., Acemoglu et al. (2011, 2014)) and naive (DeGroot) learning (e.g., Golub and Jackson (2010)). In these models, agents try to uncover an underlying state of the world, e.g., whether a vaccine is safe or not, and do so by learning from news they receive as well as by exchanging opinions with each other. Bayesian agents employ their sophisticated reasoning abilities before they accept a piece of news as factual, whereas DeGroot agents simply aggregate the information they receive into their beliefs without much deliberation. In the presence of misinformation, which of these two agent types is more likely to mislearn?

Theorem 1 answers this question (and reconciles the empirical literature) by allowing us to think about learning as a function of agent sophistication and the amount of misinformation in the system. When there is little or no misinformation, we recover the classic result that sophisticated agents learn better than naive ones.[5] As the amount of misinformation increases, we enter

[2]This relationship between increased cognitive sophistication and disagreement is widely documented across multiple issues; see for example Drummond and Fischhoff (2017); Kahan et al. (2012); Hamilton et al. (2015).

[3]For recent work that incorporates misinformation into the study of learning and how it impacts beliefs of the true state, see Allon et al. (2019); Mostagir et al. (2021); Mostagir and Siderius (2021); for work that studies how misinformation impacts platform design and/or the strategic sharing aspect of misinformation (but *not* how misinformation impacts beliefs of the true state), see Candogan and Drakopoulos (2020); Papanastasiou (2020); Acemoglu et al. (2021).

[4]For a recent example in a social learning context, Dasaratha et al. (2020) measure how well naive learning performs by comparing it to the Bayesian benchmark.

[5]This agrees with the hypothesis that more sophisticated agents are more likely to learn in the presence of misin-

a relatively unexplored territory where naive agents outperform sophisticated agents, who can no longer agree on what the true state is.[6] This result is slightly ironic given that the interest in studying DeGroot agents stems from the critique that Bayesian agents possess reasoning abilities that are unrealistic to demand of the average person. Theorem 1 shows that this learning superiority falls apart once there is enough misinformation in the system.

Underlying these diverging outcomes is the fact that the role of misinformation extends beyond corrupting the news that agents receive. For the Bayesians agents, the mere knowledge that there is misinformation in the system allows them to spin narratives that justify their existing positions. A post-truth world emerges where the presence of evidence to the contrary of one's beliefs is *rationally* dismissed as misinformation, and learning is hampered as a result. This is not the case for the DeGroots, who simply include any new information into their beliefs without worrying about its accuracy. This makes them vulnerable to misleading information, but they also benefit from not being paranoid about whether every piece of information is correct or not. This "openness" can allow them to learn better than the Bayesians, who develop a form of *endogenous stubbornness* that stops them from learning. As a result, increased sophistication can lead to less learning and more disagreement as agents "dig in their heels" and refuse to move from their existing positions.

We remark that the above result is similar to the findings of Tappin et al. (2020), which suggest that sophisticated agents disagree as a result of unbiased Bayesian behavior and not partisan bias. Our model allows us to comment more broadly on the polarization debate (see Osmundsen et al. (2020); Tappin et al. (2020)): does an increase in polarization and partisan bias lead to more disagreement over objective facts? Theorem 3 shows that in the presence of enough misinformation, an increase in polarization breaks down Bayesian learning, but can leave DeGroot learning unaffected and in some cases even make these simpler agents better off. This means that studies that measure the impact of polarization on susceptibility to misinformation should account for agent sophistication and the amount of misinformation in the system, otherwise they risk leaving out important confounding factors. This extends to policies that target specific agents in the population (for example, through educational outreach) in order to stop misinformation. Proposition 3 shows that the question of whom to target has different answers depending on the sophistication level of society: a planner should target agents with moderate beliefs in a DeGroot society, but focus her efforts on the extremists in a Bayesian one.

formation, as documented empirically in Bronstein et al. (2019); Bago et al. (2020); Pennycook and Rand (2018) and a wealth of other literature.

[6] This in turn is consistent with the empirical findings that more sophisticated agents are more likely to disagree, as shown in Drummond and Fischhoff (2017); Kahan et al. (2012); Hamilton et al. (2015).

2 Model

There is a true and unknown state of the world $\theta \in \{L, R\}$ indicating whether a left-leaning or right-leaning idea is correct. There are N agents in the population who are trying to learn θ to make an informed binary decision. Agents receive independent information (messages) about θ and then interact with others to try and learn what the true value of θ is.

Timing. We present our model using a parsimonious three-period model $t = 0, 1, 2$:

(i) At $t = 0$, agents start out with heterogeneous ideological beliefs. The initial belief, $\pi_{i,0}$, of each agent i is drawn i.i.d. from a continuous distribution H, corresponding to her belief in $\theta = R$. Agents with beliefs $\pi_{i,0} < 1/2$ are (initially) left-leaning, while those with beliefs $\pi_{i,0} > 1/2$ are (initially) right-leaning.

(ii) At $t = 1$, agents receive (independent) messages advocating for either state L or state R. Formally, each agent i receives a single message $m_i \in \{L, R\}$.[7] Some of the messages come from *organic news*, which are correlated with the state; in particular, $\mathbb{P}[m_i = \theta] = p > 1/2$ for every agent i. We assume throughout that the population N is large (i.e., $N \to \infty$) so that in the presence of organic news only, the truth is discernible with high probability by aggregating all of the messages.[8]

The messages that agents receive could also be *misinformation* that is orthogonal to the state. Agents cannot discern whether a given message contains misinformation or not. The probability that a message contains misinformation is $q < 1/2$, i.e., most news is organic. Agents are aware of the existence of misinformation (they know q), but do not know how this misinformation is broken down along the two possible states, i.e., they do not know the proportion of misinformation arguing for L vs. the proportion of misinformation arguing for R. This follows the empirical observation in van der Linden et al. (2020) that shows that while people agree about the existence of misinformation and even its extent, they do not agree on whether this misinformation leans more left or right.

Let r denote the proportion of misinformation *on the right*. We assume that r is drawn from a differentiable distribution $r \sim F(\cdot)$ at $t = 0$ and is independent of θ. We assume that F has

[7] Receiving a single message is without loss of generality. Appendix B.4 considers a generalization where agents get multiple (independent) messages over time, but our results apply identically.

[8] In Appendix B.5.2, we measure the sensitivity of our results to this assumption via simulations with finite N populations.

full support[9] on some interval $[\underline{r}, \bar{r}]$ and no support[10] outside of this interval.[11] Similarly, we assume the distribution of prior beliefs H has full support over $[\underline{\pi}, \bar{\pi}]$ and no support outside of it. In other words, we assume the supports of all distributions are *convex*.

(iii) At $t = 2$, agents observe the broadcasted beliefs of other agents in periods $t = 0$ and $t = 1$ and use these beliefs (as well as their own message) to form a final belief $\pi_{i,2}$ (as described below). Following this, each agent i makes a binary decision $a_i \in \{L, R\}$ based on which state she believes is more likely.

We consider two types of populations: **Bayesian** and **DeGroot**. Bayesian agents learn about θ by updating their beliefs in a fully Bayesian way, whereas DeGroots use simple learning heuristics. We use $\mathbf{1}_{\theta=R}$ to denote the indicator function of $\theta = R$ (i.e., $\mathbf{1}_{\theta=R}$ is equal to 1 when $\theta = R$ and 0 when $\theta = L$). Recall $\pi_{i,0}, \pi_{i,1}$, and $\pi_{i,2}$ are the beliefs of agent i at times $t = 0, 1$, and 2, respectively.

(i) **Bayesian Society**: At $t = 1$, each Bayesian agent forms a posterior update about the state, $\pi_{i,1}$, given the article with message m_i and knowing content may contain misinformation:

$$\pi_{i,1}(m_i = R) = \mathbb{E}[\mathbf{1}_{\theta=R}|m_i = R] = \int_0^1 \frac{(p(1-q) + qr)\pi_{i,0}}{p(1-q)\pi_{i,0} + (1-p)(1-q)(1-\pi_{i,0}) + qr} f(r)\ dr$$

$$\pi_{i,1}(m_i = L) = \mathbb{E}[\mathbf{1}_{\theta=R}|m_i = L] = \int_0^1 \frac{((1-p)(1-q) + q(1-r))\pi_{i,0}}{(1-p)(1-q)\pi_{i,0} + p(1-q)(1-\pi_{i,0}) + q(1-r)} f(r)\ dr$$

At time $t = 2$, agents form Bayesian posterior estimates about the state, $\pi_{i,2}$, given their article with message m_i and the beliefs of agents in the population $\{\pi_{j,0}, \pi_{j,1}\}_{j \neq i}$, again, fully aware that there may be misinformation in the system. This is akin to the updating process in Acemoglu et al. (2016), where agents are uncertain about the underlying message distribution.

(ii) **DeGroot Society**: DeGroot agents are boundedly rational agents who use a learning heuristic to learn θ. At $t = 1$, each DeGroot agent updates her belief of the state using Bayes' rule taking the news at *face value* (i.e., assuming there is no misinformation in the system). This is similar to how these agents update their beliefs in Jadbabaie et al. (2012)):

$$\pi_{i,1}(m_i = R) = \mathbb{E}[\mathbf{1}_{\theta=R}|m_i = R,\ q = 0] = \frac{p\pi_{i,0}}{p\pi_{i,0} + (1-p)(1-\pi_{i,0})} \quad (1)$$

[9] We define full support of a distribution G on an interval $[a, b]$ as having its density g satisfy the following property: there exists $\mu > 0$ such that $g(\alpha) > \mu$ for all $\alpha \in [a, b]$.

[10] No support over a set $\mathcal{A}$ means the distribution G draws an element from $\mathcal{A}$ with probability 0.

[11] While this does not rule out full support of F on $[0, 1]$, this more general assumption allows us to capture the effect of relatively symmetric vs asymmetric prevalence of misinformation.

$$\pi_{i,1}(m_i = L) = \mathbb{E}[\mathbf{1}_{\theta=R}|m_i = L,\, q = 0] = \frac{(1-p)\pi_{i,0}}{(1-p)\pi_{i,0} + p(1-\pi_{i,0})} \tag{2}$$

Based on these observations, each DeGroot takes an average of all the time $t = 1$ beliefs of the agents in society to form their time $t = 2$ belief, i.e., $\pi_{i,2} = \frac{1}{N}\sum_{j=1}^{N} \pi_{j,1}$. In other words, DeGroot agents employ "rule-of-thumb" learning to update their beliefs instead of forming a Bayesian posterior belief.

Learning. At $t = 2$, agent i chooses a binary terminal action $a_i \in \{L, R\}$ that minimizes her quadratic loss $\mathbb{E}[(a_i - \mathbf{1}_{\theta=R})^2]$ given her belief, $\pi_{i,2}$. We follow the standard definition of learning (e.g., Acemoglu et al. (2011)) and say that society *learns* if all agents take the correct action ($a_i = \theta$); otherwise, society *mislearns*. In Appendix B.5.1, we explore how our results are affected when characterizing the expected proportion of agents who (mis)learn instead of the classical "all-or-nothing" measure of (mis)learning.

Remark — While we adopt the three-period learning model to most transparently demonstrate the main concepts, richer learning dynamics can be supported without compromising any of the key results. In particular, when agents learn from each others' beliefs over a social network (and thus do not observe all beliefs in the population), our findings generalize, provided there is a longer learning horizon and given some mild assumptions on the network structure. The details of the reduction from more general networked learning to the three-period model are supplied in Appendix B.4.

3 Illustrative Example

We present an example to show that agents who use simple learning heuristics can learn better than fully-rational agents in the presence of misinformation. For concreteness, we fix $\theta = L$ as the true state (which, by assumption, is unknown to the agents). We also assume the misinformation ideological split r is uniformly distributed on $[0, 1]$, i.e., the split is ex-ante symmetric for left-leaning and right-leaning misinformation. Agents do not know the exact value of r, but they know that it comes from the uniform distribution on $(0, 1)$. We compare the following two settings:

Setting A: Weak Organic Messages and No Misinformation. Consider the baseline case studied throughout the social learning literature. There is no misinformation, i.e., $q = 0$, and $p = 0.54$, so that 54% of organic messages align with $\theta = L$. In this setup, organic news is (weakly) correlated with the truth.

Do agents learn the correct state (almost surely) when the population is large? The short answer, as already developed in a vast literature, is yes. Both agent types correctly learn that the true state is L. This happens regardless of their initial prior beliefs (i.e., even those on the extreme right still learn that the correct state is L) and despite the fact that news is weakly correlated with θ. In this setup, both the Bayesian and DeGroot societies take the correct action.

Setting B: Misinformation with Stronger Organic Messages. In this setting, misinformation exists in the system, with $q = 0.25$. On the other hand, organic news is of higher quality, with $p = 0.6$. For this example, assume that $r = 0.64$, i.e., 64% of the misinformation advocates for $\theta = R$ and 36% advocates for $\theta = L$. This means that among the large collection of messages $\{m_i\}_{j=1}^N$, roughly 54% correspond to L (i.e., $p(1-q) + (1-r)q$) and 46% correspond to R (i.e., $(1-p)(1-q) + qr$). This distribution is *identical* to the one in Setting A, under which agents were able to correctly learn. In Setting B, if agents did not know misinformation exists, learning will proceed exactly as before and everyone will take the correct action despite the presence of misinformation. However, agents are now aware that there is misinformation in the system. Given this, we analyze how each society updates its beliefs.

Bayesian agents: Agents observe the initial beliefs $\pi_{j,0}$. At $t = 1$, each Bayesian agent i forms a posterior belief $\pi_{i,1}$ based on m_i as mentioned in Section 2. Observe that $\pi_{j,1} > \pi_{j,0}$ if and only if $m_i = R$ and $\pi_{j,1} < \pi_{j,0}$ if and only if $m_i = L$. Thus, by observing beliefs in the second period, every agent i can deduce the messages $\{m_i\}_{j=1}^N$. As noted, When N is large, the law of large numbers guarantees that roughly 54% of the messages that agent i observes will favor L.

Note that there are *exactly* two realizations of r for which 54% of the messages are L and 46% are R. The first is the true realization, where $\theta = L$ and $r = 0.64$. The other is where $\theta = R$ but $r = 0.04$ (i.e., only 4% of the misinformation is on the right). This situation is depicted in Figure 1. Because r is uniformly distributed, *both* of these scenarios are equally likely. This implies that $\pi_{j,2} = \pi_{j,0}$ because the messages provide no information about the state θ. In other words, right-leaning agents spin a narrative that the vast majority of misinformation is on the left, and cannot use the massive quantity of news to change their views. Similarly, left-leaning agents do the same, believing (correctly, in their case) that most of the misinformation must be on the right. The society of Bayesians does not learn and there is persistent disagreement about θ in the population.

DeGroot agents: DeGroot agents take messages at face value and use them to update their beliefs (via Equations (1) and (2)) using $p = 0.6$. Thus, noting that 54% of messages are L and 46% of

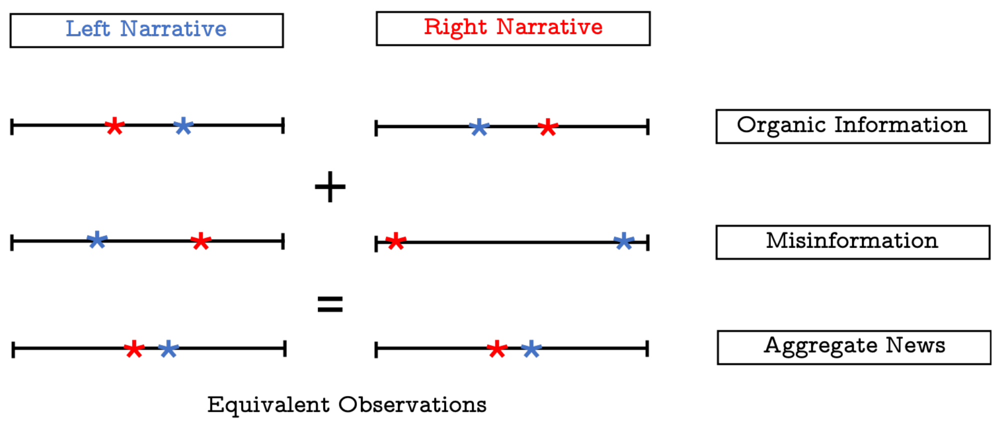

Figure 1. Two narratives that give rise to equivalent observations. In this example, the true state is L and most of the misinformation comes from the right with $r = 0.64$, i.e., the left narrative is the correct one. However, agents who hold right-leaning beliefs can rationalize their observations as coming from the right narrative, with the organic information arguing for R and most of the misinformation coming from the left, with $r = 0.04$.

messages are R, DeGroot agents hold beliefs $\pi_{j,t}$ for all $t \geq 2$:

$$\int_0^1 \left(.46 \cdot \frac{.6\alpha}{.6\alpha + .4(1-\alpha)} + .54 \cdot \frac{.4\alpha}{.4\alpha + .6(1-\alpha)} \right) \, d\alpha = 0.495 < 1/2$$

so all agents in society learn the correct state, in contrast to the Bayesians.

Likelihood of Mislearning. We used a specific misinformation split $r = 0.64$ in the above example for simplicity. Other values of r would give rise to different outcomes. Instead of focusing on a specific realization of r, we can instead look at the *likelihood* that society does not learn when the split r is randomly drawn from its true (uniform, in this example) distribution. How do Bayesian agents perform relative to DeGroot agents *on average*?

It turns out that in this case, Bayesian agents mislearn *twice as much* as DeGroots (probability that Bayesians mislearn is 40% vs. 20% for the DeGroots – see Appendix B.1 for calculations). We characterize this ratio as a function of the distribution of r in Theorem 2. Generally, while Bayesian agents thrive in environments where information is organic, they are much more vulnerable to mislearning and taking the wrong action in the presence of misinformation.

4 Learning in Bayesian vs. DeGroot Societies

In this section, we generalize the previous example and investigate the conditions under which learning breaks down in each society. We follow this up with two technical results. Theorem 1

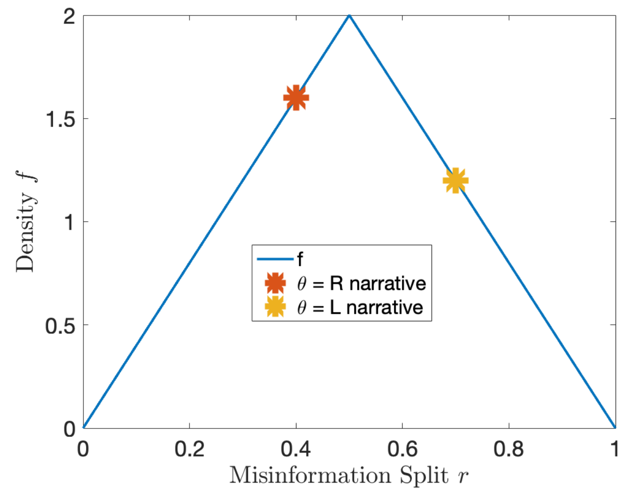

Figure 2. Two narratives that justify the same observed message distribution under a specific misinformation split r. Agents do not know r but they know that it follows a triangular distribution. This makes them believe that the more likely narrative is the one that corresponds to the true state being R. In this case, all agents move further to the right and away from the correct state.

characterizes *when* Bayesians learn worse than DeGroots as a function of the *amount* of misinformation in the system, and Theorem 2 quantifies *how much* worse they learn as a function of the *distribution* of that misinformation.

4.1 DeGroot (Mis)learning: Propaganda for the Incorrect State

DeGroot agents update their beliefs about the state by averaging the opinions of others. Because of this simple updating process, they always converge to a (possibly incorrect) consensus about what the true state is. When there is no misinformation (i.e., $q = 0$), agents *will* learn the correct state and choose $a_i = \theta$ (see Golub and Jackson (2010)). When misinformation is present, we can provide necessary and sufficient conditions for DeGroot learning in terms of the ideological split of misinformation r:

Proposition 1. *When $\theta = L$, there exists a threshold r^*_D such that if $r < r^*_D$, the DeGroot society learns and if $r > r^*_D$, the DeGroot society mislearns.*

Proposition 1 shows that failure of learning depends on whether propaganda for the incorrect state is sufficiently high to direct agents away from the belief that the organic news argues for. Agents learn as long as this propaganda does not overpower organic news.

4.2 Bayesian (Mis)learning: Rationalization and Competing Narratives

Bayesian agents make inferences about the state by observing the distribution of messages in the population. They then rationalize the values for the misinformation split r that gives rise to this message distribution. We refer to these values as *narratives*. There can be at most two narratives: one that corresponds to $\theta = L$ and one that corresponds to $\theta = R$. Only one of these narratives is correct, and agents will stick with the narrative that fits their beliefs.[12] However, there can also be a *single* correct narrative, in which case all agents learn the true state θ. This again depends on the extent of misinformation arguing for the incorrect state:

Proposition 2. *When $\theta = L$, there exists a threshold r^*_B such that if $r < r^*_B$, a single narrative exists and the Bayesian society learns.*

If $r > r^*_B$, then two narratives exist. Figure 2 shows this situation for a specific $r > r^*_B$ value. Recall that agents do not know r, but in this example they know that it follows a triangular distribution. This makes them believe that the more likely narrative is the one corresponding to R. In this case, *all agents* move further to the right (and away from the true state). Note that unlike the example in Section 3, where the two competing narratives were equally likely under the uniform distribution, the triangle distribution makes the incorrect narrative strictly more likely. Section 4.4 shows how to quantify the likelihood of mislearning as a function of the hazard rate of the distribution of r.

4.3 Low vs High Misinformation Regimes

Next, we present our main theorem, which analyzes the settings under which DeGroot or Bayesian agents mislearn more often:

Theorem 1. *Suppose H is symmetric about $1/2$.*[13] *Then there exists a threshold $q^* \in (0, 1)$ such that:*

(a) If $q < q^*$, the Bayesian society mislearns with lower probability than the DeGroot society;

(b) If $q > q^*$ and H and F have full support on $[0, 1]$, the Bayesian society mislearns with strictly higher probability than the DeGroot society.

[12]This is reminiscent of the effect informally described by the Today show co-host Al Roker, commenting on conflicting results of science experiments: "I think the way to live your life is you find the study that sounds best to you and you go with that."

[13]This corresponds to a society where every belief on the left is perfectly mirrored by a belief on the right of the same extremity. This allows us to analyze both the $\theta = L$ and $\theta = R$ cases identically and abstracts away from scenarios where society begins either initially biased toward or away from the correct state, in order to focus on the underlying learning mechanisms.

Theorem 1 establishes that when the amount of misinformation is not too large, Bayesian agents can still learn better than their DeGroot counterparts. This is the classic intuition from the social learning literature and the one empirically observed in Bronstein et al. (2019); Bago et al. (2020); Pennycook and Rand (2018). On the other hand, once misinformation becomes more rampant, DeGroot agents become *more* adept at aggregating the organic information compared to the Bayesians, who find themselves in consistent disagreement over what the truth is, as in the empirical work of Drummond and Fischhoff (2017); Kahan et al. (2012); Hamilton et al. (2015).

The intuition for Theorem 1 is as follows. When misinformation is relatively low, the organic news is enough for Bayesians to infer the true narrative and see the misinformation as purposefully deceptive. On the other hand, with DeGroot agents, this misinformation still has a chance of successfully steering the beliefs of the population away from the truth. Once the amount of misinformation becomes large, a *post-truth* effect kicks in for the Bayesians: any reasonable narrative can be told about the source of misinformation, and disagreement over the true state ensues. While DeGroots are not guaranteed to learn either, their simple updating tends to work well when misinformation is not too heavily skewed in one direction or another. Thus, in more instances, the DeGroot society is able to learn from the organic news while allowing the misinformation on both sides to nearly "wash out." We can quantify how much better the DeGroots do in this case, which is the subject of the next section.

4.4 High Misinformation and Mislearning Rates

When $q > q^*$, we observe in Theorem 1 that DeGroot agents learn more effectively than Bayesian agents. This section is for readers who are interested in quantifying the extent with which DeGroots outperform Bayesians. This depends on the misinformation split distribution $F(\cdot)$. We formalize this as follows. The exact difference in rates of mislearning between Bayesian and DeGroots in the high-misinformation regime (i.e., $q > q^*$) is determined by the *hazard rate* $\lambda_F(\alpha)$ of F. Recall that the hazard rate is given by $\lambda_F(\alpha) = \frac{f(\alpha)}{1-F(\alpha)}$, where $f(\cdot)$ is the density of cumulative distribution function F. Denote by μ the the relative frequency of DeGroot to Bayesian mislearning (which by Theorem 1 is less than 1), then we can characterize how μ changes as a function of the strength of organic signals p (i.e., whether agents can learn from strongly informative content):

Theorem 2. *Suppose that H and F have full support on $[0,1]$, H is symmetric, and $q > q^*$ as in Theorem 1(b). Consider $\alpha = \frac{1-2(1-p)(1-q)}{2q}$ and $\beta = p\left(1 - \frac{q^*}{q}\right)$.*[14] *The ratio μ is increasing in p if $\lambda_F(\alpha) < 2\lambda_F(\alpha-\beta)$, decreasing in p if $\lambda_F(\alpha) > 2\lambda_F(\alpha-\beta)$, and unchanging if $\lambda_F(\alpha) = 2\lambda_F(\alpha-\beta)$.*

[14] Note that it can be shown $0 \leq \alpha - \beta \leq \alpha \leq 1$, so the hazard rates are well-defined everywhere.

Informally, Theorem 2 states that Bayesians perform comparatively worse than DeGroots when misinformation is more evenly distributed. The reason is that while Bayesians are adept at making inferences about the possibility of strongly misleading misinformation, misinformation that is relatively balanced on both sides permits more rationalization of narratives and more disagreement. On the other hand, more balanced misinformation is always better for DeGroot learning because it permits greater likelihood of *overall balanced* news and less propaganda for the incorrect state. Appendix B.3 shows how to apply this result when r comes from skewed or unskewed distributions.

5 Polarization

Polarization and political partisanship have been steadily increasing (see Abramowitz (2010) and Pew Research Center (2014) for evidence from the United States). As noted in the introduction, a substantial literature advocates that increased polarization makes people more likely to disagree over objective facts because of politically-biased reasoning, whereas new findings (see Tappin et al. (2020)) suggest that this disagreement can be explained as a byproduct of sophisticated (Bayesian) reasoning and not partisan bias.

Our model lends support to the latter explanation by showing that disagreement can naturally arise from Bayesian updating. More generally, Theorem 3 highlights the fact that attempts to establish a connection between partisanship and disagreement should take into account the sophistication level of the agents and the amount of misinformation in the system. As misinformation becomes more prevalent, increased polarization leads to failure of learning in Bayesian societies, but leaves DeGroot societies relatively unaffected. Thus, measuring the effects of polarization on learning and consistent disagreement without incorporating these elements can lead to seemingly contradictory evidence.

We operationalize polarization as follows: consider some symmetric (about $1/2$) belief distribution H (density h) with support $[\underline{\pi}, \bar{\pi}]$ and a mean-preserving spread of H to some $\tilde{H}_\gamma$ (density h_γ) defined as:

$$\tilde{h}_\gamma(\pi + (\pi - 1/2)\gamma) = \frac{1}{1+\gamma} h(\pi)$$

where $\gamma \in [-1, \bar{\gamma}]$ where $\bar{\gamma} = \min\left\{\frac{1-\bar{\pi}}{\bar{\pi}-1/2}, \frac{\underline{\pi}}{1/2-\underline{\pi}}\right\}$. One can think of γ as a measure of belief polarization in society. For larger γ, $\pi + (\pi - 1/2)\gamma$ is closer to 0 when $\pi < 1/2$ and closer to 1 when $\pi > 1/2$; thus, the probability of realizing "tail" beliefs grows when γ increases. (Note that $1/(1+\gamma)$ is simply a scaling factor to guarantee the $\int_0^1 \tilde{h}_\gamma(\pi)\, d\pi = 1$.) A simple example of

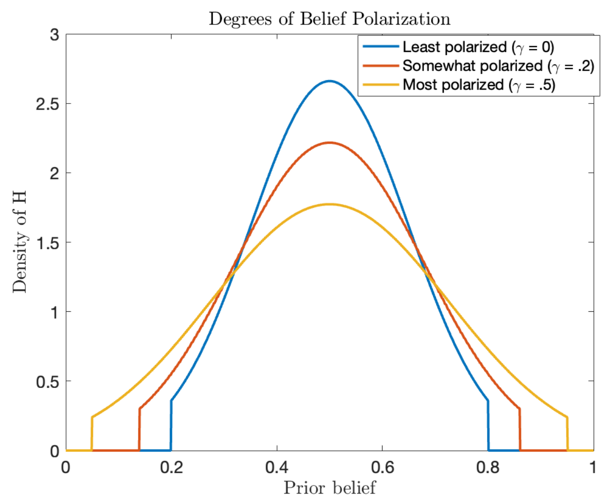

Figure 3. Polarization of beliefs captured through parameter γ.

increasing polarization for a truncated normal distribution of beliefs is given in Figure 3.

Our next result establishes a threshold result for polarization and its consequences to both Bayesian and DeGroot mislearning:

Theorem 3. *Let H be symmetric about $1/2$ and $q > q^*$. There exists a threshold γ^* such that if $\gamma < \gamma^*$, the DeGroot society mislearns more often than the Bayesian society, whereas if $\gamma > \gamma^*$, the DeGroot society mislearns less often than the Bayesian society.*

In Bayesian societies, an increase in belief polarization always hurts learning. Mean-preserving spreads necessarily create more tension in the learning process because agents cannot agree on what is likely to be the true narrative. The impact is more mild on DeGroot agents because they communicate and update beliefs by taking averages of their neighbors. In this case, mean-preserving spreads do not affect their general ability to aggregate organic information, even when they start from very different initial beliefs.

Remark — Note that when H is asymmetric, polarization still always hurts Bayesian societies (see the proof of Theorem 3 in Appendix A). However, it is possible that in this case polarization *improves* the chances of learning in a DeGroot society (see Appendix B.2 for an example). This occurs because when society is initially well-informed, additional polarization pulls initial opinions towards the correct and incorrect states. However, the convexity of how agents update their beliefs from news (see Equations (1) and (2)) leads to an overall movement towards the correct state. This fails in Bayesian societies because of persistent disagreement about the truth, and further accentuates the result of Theorem 3 that polarization is more damaging to a Bayesian society than a DeGroot one.

Targeting Policies. We now consider the problem of a planner who can target a subset of the population with information arguing for one state over the other. One can think of this policy as an educational outreach intervention. For example, governments have been ramping up their efforts to convince citizens to vaccinate against COVID-19, and a part of these efforts is targeted advertising. The question is which agents should the planner target? For example, should she target the most polarized agents? We show that the answer, again, crucially depends on the level of sophistication of the agents.

Formally, the planner informs a small but positive measure of agents that the correct state is either L or R. We assume the agents interpret this information in the same way they do news: combined with the message m_j that agent j receives, she also gets the planner's message $m_j^p \in \{L, R\}$ and updates using both messages. The next result describes the planner's targeting policy:

Proposition 3. *There exists* $1/2 < \pi^* < 1$ *such that the targeting policy for DeGroot agents is to target those agents whose beliefs lie in an open interval containing* π^* *(and bounded away from 1). The policy for Bayesians agents is to target those agents whose beliefs are farthest from the truth, i.e., the extremists.*

Proposition 3 states that in a DeGroot world with $\theta = L$, the planner wants to influence right-leaning *moderates*, as these are the agents who change their belief most when seeing message L. Extremists in this society mostly dismiss messages that don't agree with their priors, and the planner has little to gain by targeting them. However, with Bayesian agents, *extremists* are exactly the agents that the planner needs to target. While the efficacy of her work is limited, these are the agents who are most inclined to spin narratives that anchor them to the incorrect state.

Remark 1 — Proposition 3 recommends a policy when the planner knows the true state, but the insights generalize when the planner herself is uncertain about what the state actually is. In a DeGroot society, the planner tries to make relatively moderate agents (of both ideologies) more moderate. In Bayesian societies, the planner tries to make the extremists (of both ideologies) move toward the center. Both of these policies push toward decreasing the polarization of ideological beliefs (albeit in different ways).

Remark 2 — Under a different learning objective, such as minimizing the proportion of agents who mislearn the state, the DeGroot policy in Proposition 3 remains unaffected, but the Bayesian policy becomes more subtle. In many environments, the regulator should still target the most polarized Bayesian agents in society, but in some other instances, targeting mostly moderate

Bayesian dissenters can be more effective. The nuances of targeting interventions under this alternative objective are explored in detail in Appendix B.5.1.

6 Final Remarks

As argued recently in Watts et al. (2021), accurate information is very much a prerequisite for successful democratic discourse. The Internet and social media have made it easier to disseminate misinformation (Allcott and Gentzkow, 2017), with far-reaching consequences.[15] There are ongoing efforts across multiple disciplines to try and uncover the mechanisms by which misinformation spreads. In this paper, we contribute to these efforts by examining misinformation through the lens of social learning and focusing on agents' sophistication types. While the learning mechanisms of these types have been studied in a broad context, they have not been analyzed and compared when there is rampant misinformation. We show a reversal of results and intuitions that hold in many normative learning setups, but not in the presence of misinformation. We do this through a parsimonious framework whose results reconcile several empirical studies and whose predictions show the need for researchers and policymakers to jointly consider sophistication and social learning as integral components in studying the spread of misinformation.

[15]Examples range from individual actions along the lines of Pizzagate (Fisher et al., 2016), to belief in the "Death Panels" of the Affordable Care Act (Watts et al., 2021), to more collective action failures like the spread of measles in Eastern Europe as a result of Russian disinformation (Broniatowski et al., 2018).

References

Abramowitz, Alan I. (2010), *The Disappearing Center: Engaged Citizens, Polarization, and American Democracy*. Yale University Press.

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar (2014), "Dynamics of information exchange in endogenous social networks." *Theoretical Economics*, 9, 41–97.

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz (2016), "Fragility of asymptotic agreement under bayesian learning." *Theoretical Economics*, 11, 187–225.

Acemoglu, Daron, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), "Bayesian learning in social networks." *The Review of Economic Studies*, 78, 1201–1236.

Acemoglu, Daron, Asuman Ozdaglar, and James Siderius (2021), "Misinformation: Strategic sharing, homophily, and endogenous echo chambers." Technical report, National Bureau of Economic Research.

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva (2020), "The polarization of reality." In *AEA Papers and Proceedings*, volume 110, 324–28.

Allcott, Hunt and Matthew Gentzkow (2017), "Social media and fake news in the 2016 election." *Journal of economic perspectives*, 31, 211–36.

Allon, Gad, Kimon Drakopoulos, and Vahideh Manshadi (2019), "Information inundation on platforms and implications." In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 555–556.

Avella-Medina, Marco, Francesca Parise, Michael T. Schaub, and Santiago Segarra (2020), "Centrality measures for graphons: Accounting for uncertainty in networks." *IEEE Transactions on Network Science and Engineering*, 7, 520–537. ArXiv: 1707.09350.

Bago, Bence, David G Rand, and Gordon Pennycook (2020), "Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines." *Journal of experimental psychology: general*.

Broniatowski, David A, Amelia M Jamison, SiHua Qi, Lulwah AlKulaib, Tao Chen, Adrian Benton, Sandra C Quinn, and Mark Dredze (2018), "Weaponized health communication: Twitter bots and russian trolls amplify the vaccine debate." *American journal of public health*, 108, 1378–1384.

Bronstein, Michael V, Gordon Pennycook, Adam Bear, David G Rand, and Tyrone D Cannon (2019), "Belief in fake news is associated with delusionality, dogmatism, religious fundamentalism, and reduced analytic thinking." *Journal of applied research in memory and cognition*, 8, 108–117.

Bursztyn, Leonardo, Aakaash Rao, Christopher P Roth, and David H Yanagizawa-Drott (2020), "Misinformation during a pandemic." Technical report, National Bureau of Economic Research.

Candogan, Ozan and Kimon Drakopoulos (2020), "Optimal signaling of content accuracy: Engagement vs. misinformation." *Operations Research*, 68, 497–515.

Dasaratha, Krishna (2020), “Distributions of centrality on networks.” *Games and Economic Behavior*, 122, 1–27. Publisher: Elsevier.

Dasaratha, Krishna, Benjamin Golub, and Nir Hak (2020), “Learning from neighbors about a changing state.” *Available at SSRN 3097505*.

Degroot, Morris H. (1974), “Reaching a Consensus.” *Journal of the American Statistical Association*, 69, 118–121.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel (2003), “Persuasion Bias, Social Influence, and Unidimensional Opinions*.” *The Quarterly Journal of Economics*, 118, 909–968.

Drummond, Caitlin and Baruch Fischhoff (2017), “Individuals with greater science literacy and education have more polarized beliefs on controversial science topics.” *Proceedings of the National Academy of Sciences*, 114, 9587–9592.

Fisher, Marc, John Woodrow Cox, and Peter Hermann (2016), “Pizzagate: From rumor, to hashtag, to gunfire in dc.” *Washington Post*.

Gale, Douglas and Shachar Kariv (2003), “Bayesian learning in social networks.” *Games and Economic Behavior*, 45, 329–346.

Golub, Benjamin and Matthew O Jackson (2010), “Naive learning in social networks and the wisdom of crowds.” *American Economic Journal: Microeconomics*, 2, 112–49.

Golub, Benjamin and Evan Sadler (2017), “Learning in Social Networks.” SSRN Scholarly Paper ID 2919146, Social Science Research Network, Rochester, NY.

Hamilton, Lawrence C, Joel Hartter, and Kei Saito (2015), “Trust in scientists on climate change and vaccines.” *Sage Open*, 5, 2158244015602752.

Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi (2012), “Non-bayesian social learning.” *Games and Economic Behavior*, 76, 210–225.

Kahan, Dan M, Ellen Peters, Erica Cantrell Dawson, and Paul Slovic (2017), “Motivated numeracy and enlightened self-government.” *Behavioural public policy*, 1, 54–86.

Kahan, Dan M, Ellen Peters, Maggie Wittlin, Paul Slovic, Lisa Larrimore Ouellette, Donald Braman, and Gregory Mandel (2012), “The polarizing impact of science literacy and numeracy on perceived climate change risks.” *Nature climate change*, 2, 732–735.

Loomba, Sahil, Alexandre de Figueiredo, Simon J Piatek, Kristen de Graaf, and Heidi J Larson (2021), “Measuring the impact of covid-19 vaccine misinformation on vaccination intent in the uk and usa.” *Nature human behaviour*, 1–12.

McIntyre, Lee (2018), *Post-truth*. MIt Press.

Mossel, Elchanan, Allan Sly, and Omer Tamuz (2014), “Asymptotic learning on Bayesian social networks.” *Probability Theory and Related Fields*, 158, 127–157.

Mostagir, Mohamed, Asu Ozdaglar, and James Siderius (2021), “When is society susceptible to manipulation?” *Management Science, Forthcoming*.

Mostagir, Mohamed and James Siderius (2021), "Social inequality and the spread of misinformation." *Management Science, Forthcoming*.

Mueller-Frank, Manuel (2013), "A general framework for rational learning in social networks." *Theoretical Economics*, 8, 1–40.

Mueller-Frank, Manuel (2014), "Does one bayesian make a difference?" *Journal of Economic Theory*, 154, 423–452.

Osmundsen, Mathias, Alexander Bor, Peter Bjerregaard Vahlstrup, Anja Bechmann, and Michael Bang Petersen (2020), "Partisan polarization is the primary psychological motivation behind "fake news" sharing on twitter." *American Political Science Review*, 115.

Papanastasiou, Yiangos (2020), "Fake news propagation and detection: A sequential model." *Management Science*, 66, 1826–1846.

Pennycook, Gordon, Ziv Epstein, Mohsen Mosleh, Antonio A Arechar, Dean Eckles, and David G Rand (2021), "Shifting attention to accuracy can reduce misinformation online." *Nature*, 592, 590–595.

Pennycook, Gordon and David G Rand (2018), "Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning." *Cognition*.

Pew Research Center, C (2014), "Political polarization in the american public." *Retrieved September*, 2, 2019.

Taber, Charles S, Damon Cann, and Simona Kucsova (2009), "The motivated processing of political arguments." *Political Behavior*, 31, 137–155.

Taber, Charles S and Milton Lodge (2006), "Motivated skepticism in the evaluation of political beliefs." *American journal of political science*, 50, 755–769.

Tappin, Ben M, Gordon Pennycook, and David G Rand (2020), "Bayesian or biased? analytic thinking and political belief updating." *Cognition*, 204, 104375.

van der Linden, Sander, Costas Panagopoulos, and Jon Roozenbeek (2020), "You are fake news: political bias in perceptions of fake news." *Media, Culture & Society*, 42, 460–470.

Watts, Duncan J, David M Rothschild, and Markus Mobius (2021), "Measuring the news and its impact on democracy." *Proceedings of the National Academy of Sciences*, 118.

Online Appendix

A Proofs

We first provide two auxillary lemmas that we use in the proofs of our main results.

A.1 Auxiliary Lemmas

Recall that H is the distribution of prior beliefs in the population and p is the strength of organic news (likelihood a message corresponds to the true state θ). For misinformation, q denotes the amount of misinformation in the system, r denotes how much of this misinformation argues for state R, and F is the distribution of r.

Lemma 1. *Let H be symmetric. If $\theta = L$, the DeGroot society mislearns if and only if $r \geq (1-2(1-q)(1-p))/(2q)$; if $\theta = R$, the DeGroot society mislearns if and only if $r \leq (1-2(1-q)p)/(2q)$.*

Proof of Lemma 1. We prove this for $\theta = L$; the case of $\theta = R$ is similar. For a fixed realization $r \sim F(\cdot)$, every DeGroot agent i converges to belief π_∞ about $\theta = R$:

$$\pi_\infty = \int_0^1 \left(((1-p)(1-q)+qr)\frac{p\alpha}{p\alpha+(1-p)(1-\alpha)} + (p(1-q)+q(1-r))\frac{(1-p)\alpha}{(1-p)\alpha+p(1-\alpha)} \right) h(\alpha)\, d\alpha$$

Via the Leibniz integral rule, we see that:

$$\begin{aligned}
\frac{d\pi_\infty}{dr} &= \int_0^1 \frac{\partial}{\partial r} \left((1-p+qr)\frac{p\alpha}{p\alpha+(1-p)(1-\alpha)} + (p+q(1-r))\frac{(1-p)\alpha}{(1-p)\alpha+p(1-\alpha)} \right) h(\alpha)\, d\alpha \\
&= q\int_0^1 \left(\frac{p\alpha}{p\alpha+(1-p)(1-\alpha)} - \frac{(1-p)\alpha}{(1-p)\alpha+p(1-\alpha)} \right) h(\alpha)\, d(\alpha) \\
&= q\int_0^1 \frac{(1-\alpha)\alpha(2p-1)}{(p-\alpha(2p-1))((1-p)+\alpha(2p-1))}\, h(\alpha)\, d\alpha
\end{aligned}$$

Note that all expressions are positive because $p > 1/2$. The only non-trivial one to verify is the first expression in the denominator, which is linear in α and thus it is sufficient to verify it is non-negative for all $\alpha \in \{0,1\}$ to prove it is non-negative for all $\alpha \in [0,1]$. When $\alpha = 0$ it is equivalent to p and when $\alpha = 1$ it is equivalent to $1-p$.

Thus, $d\pi_\infty/dr > 0$ for all r. Consider the expression for $\pi_\infty(r)$ when $r = \tilde{r} \equiv (1-2(1-q)(1-p))/(2q)$:

$$\begin{aligned}
&\int_0^1 \left(((1-p)(1-q)+q\tilde{r})\frac{p\alpha}{p\alpha+(1-p)(1-\alpha)} + (p(1-q)+q(1-\tilde{r}))\frac{(1-p)\alpha}{(1-p)\alpha+p(1-\alpha)} \right) h(\alpha)\, d\alpha \\
&= \int_0^{1/2} \frac{p\alpha}{2(p\alpha+(1-p)(1-\alpha))}\, h(\alpha)\, d\alpha + \int_{1/2}^1 \frac{(1-p)\alpha}{2((1-p)\alpha+p(1-\alpha))}\, h(\alpha)\, d\alpha
\end{aligned}$$

For the second integral expression, we make the change of variables $\beta \equiv 1-\alpha$, which yields:

$$= \int_0^{1/2} \frac{p\alpha}{2(p\alpha+(1-p)(1-\alpha))}\, h(\alpha)\, d\alpha + \int_0^{1/2} \frac{(1-p)(1-\beta)}{2((1-p)(1-\beta)+p\beta)}\, h(1-\beta)\, d\beta$$

By symmetry of H, we know that $h(1-\beta) = h(\beta)$, thus the above expression simplifies to $\int_0^{1/2} h(\alpha)\, d\alpha = 1/2$ because H is symmetric and $\int_0^1 h(\alpha)\, d\alpha = 1$. Similarly, because $d\pi_\infty/dr > 0$, we know whenever $r > \tilde{r}$, $\pi_\infty(r) > 1/2$ and so agents elect action $a_i = R$, and the society mislearns. Whenever $r < \tilde{r}$, $\pi_\infty(r) < 1/2$ and so agents elect action $a_i = L$, and the society *does* learn. □

Lemma 2. *If H has full support and $\theta = L$, the Bayesian society mislearns if and only if $r \geq \underline{r} + \frac{(2p-1)(1-q)}{q}$; if $\theta = R$, the Bayesian society mislearns if and only if $r \leq \bar{r} - \frac{(2p-1)(1-q)}{q}$.*

Proof of Lemma 2. Once again, we prove this $\theta = L$ and remark that the case of $\theta = R$ is similar. Unlike Lemma 1, we prove both the "if" and "only if" parts separately:

(i) *If part*: Suppose $r \geq \underline{r} + \frac{(2p-1)(1-q)}{q}$. There is some proportion ρ_L of messages that advocate for L and some proportion $\rho_R = 1 - \rho_L$ of messages that advocate for R.

When $\theta = L$, $p(1-q)$ proportion of the messages are both organic and advocate for L, i.e., $m_i = L$, when the population is large. Similarly, $q(1-r)$ proportion of messages are inorganic ("misinformation") and advocate for L as well. Call this scenario 1.

When $\theta = R$, $(1-p)(1-q)$ proportion of the messages are both organic and advocate for L, with once again, $q(1-r)$ proportion of messages that are inorganic and advocate for L too. Call this scenario 2.

Both scenarios occur with positive probability if there exist some r_1 and r_2, both within $[\underline{r}, \bar{r}]$, where both scenarios yield the same realized distribution ρ_L and ρ_R. In scenario 1 we admit $\rho_L^1 = p(1-q) + q(1-r_1)$ and in scenario 2 we admit $\rho_L^2 = (1-p)(1-q) + q(1-r_2)$. Because $\theta = L$, we know that $r_1 = r$ and scenario 1 occurs with positive probability. To see if scenario 2 occurs with positive probability, one needs to find the existence of $r_2 \in [0,1]$ such that $p(1-q)+q(1-r) = (1-p)(1-q)+q(1-r_2)$. When $r_2 = 1$ and since $p(1-q)+q(1-r) > (1-p)(1-q)$, there exists a value for r_2 where the left-hand side is greater than the right-hand side. Moreover, the left-hand side is decreasing in r_2, so verify there exists some r_2 where equality can be obtained, it is sufficient to have $p(1-q)+q(1-r) \leq (1-p)(1-q)+q$. Rearranging gives the condition in the lemma.

Finally, we note that under this condition, both scenario 1 and scenario 2 occur with probability $\eta_1, \eta_2 > 0$. The probability that an agent with prior $\pi_{i,0}$ about scenario 2 (i.e., $\theta = R$) is:

$$\pi_{i,2} = \frac{\eta_2 \pi_{i,0}}{\eta_2 \pi_{i,0} + \eta_1(1-\pi_{i,0})}$$

Taking $\pi_{i,0}$ sufficiently close to 1 yields $\pi_{i,2} > 1/2$, and given H has full support, implies some positive fraction of the Bayesian population mislearns, so society fails to learn as well.

(ii) *Only if part*: Suppose $r < \underline{r} + \frac{(2p-1)(1-q)}{q}$. Then by the same argument in the "if" proof, there exists no value for r_2 such that $p(1-q)+q(1-r) = (1-p)(1-q)+q(1-r_2)$ because $p(1-q)+q(1-r) > (1-p)(1-q)+q(1-r_2)$ for all $r_2 \in [\underline{r}, \bar{r}]$. Thus, there exists a unique value for r that yields message distribution ρ_L (scenario 1) and it necessarily corresponds to $\theta = L$. Note that:

$$\pi_{i,2} = \frac{\eta_2 \pi_{i,0}}{\eta_2 \pi_{i,0} + \eta_1(1-\pi_{i,0})} = 0$$

given that $\eta_1 > 0$ and $\eta_2 = 0$, and $\pi_{i,0} \in (0,1)$ almost surely. Thus, all of the Bayesian agents learn the correct state $\theta = L$. □

A.2 Proofs of Section 4

Proof of Proposition 1. We showed in Lemma 1 that $d\pi_\infty/dr > 0$ without utilizing the symmetry assumption on H. The DeGroot society mislearns if and only if that $\pi_\infty(r) > 1/2$ when $\theta = L$. Thus, there is a unique cutoff r_D^* such that the DeGroot society mislearns if and only if $r > r_D^*$. □

Proof of Proposition 2. We showed in Lemma 2 there is a single narrative when $r < r_B^* \equiv \underline{r} + \frac{(2p-1)(1-q)}{q}$, which implies the Bayesian society must learn in this setting, as $\pi_{i,2} = 0$ when $\theta = L$. □

Proof of Theorem 1. Note by Lemma 1, the DeGroot society mislearns with positive probability when $q > \frac{2p-1}{2p}$. By the arguments in Lemma 2, given that H may or may not have full support, a necessary (but not necessarily sufficient) condition for the Bayesian society to mislearn is that $q > \frac{2p-1}{2p-\underline{r}} \geq \frac{2p-1}{2p}$. This establishes part (a) by taking $q^* = \frac{2p-1}{2p}$.

For part (b), when $q > q^*$ and H and F have full support (so $\underline{r} = 0$), the Bayesian society mislearns with probability $1 - F\left(\frac{(2p-1)(1-q)}{q}\right)$ and by Lemma 1 the DeGroot society mislearns with probability $1 - F\left(\frac{1-2(1-q)(1-p))}{2q}\right)$. Observe that:

$$\frac{(2p-1)(1-q)}{q} - \frac{1-2(1-q)(1-p))}{2q} = \frac{2p(1-q)-1}{2q}$$

which is a decreasing function in q and is exactly equal to 0 when $q = q^*$. Thus, by monotonicity of F, $F\left(\frac{(2p-1)(1-q)}{q}\right) < F\left(\frac{1-2(1-q)(1-p))}{2q}\right)$ and so $1 - F\left(\frac{(2p-1)(1-q)}{q}\right) > 1 - F\left(\frac{1-2(1-q)(1-p))}{2q}\right)$, implying the Bayesian society mislearns more often. □

Proof of Theorem 2. When $\theta = L$, note the ratio of DeGroot mislearning to the ratio of Bayesian mislearning is given by:

$$\mu = \frac{1 - F\left(\frac{1-2(1-q)(1-p)}{2q}\right)}{1 - F\left(\frac{(2p-1)(1-q)}{q}\right)}$$

by Lemma 1 and Lemma 2. Differentiating with respect to p, we get that

$$\frac{\partial \mu}{\partial p} = \frac{f\left(\frac{(2p-1)(1-q)}{q}\right)\frac{2(1-q)}{q}\left(1 - F\left(\frac{1-2(1-q)(1-p)}{2q}\right)\right) - f\left(\frac{1-2(1-q)(1-p)}{2q}\right)\frac{1-q}{q}\left(1 - F\left(\frac{(2p-1)(1-q)}{q}\right)\right)}{\left(1 - F\left(\frac{(2p-1)(1-q)}{q}\right)\right)^2}$$

Note that $\partial\mu/\partial p > 0$ if and only if

$$2f\left(\frac{(2p-1)(1-q)}{q}\right)\left(1 - F\left(\frac{1-2(1-q)(1-p)}{2q}\right)\right) > f\left(\frac{1-2(1-q)(1-p)}{2q}\right)\left(1 - F\left(\frac{(2p-1)(1-q)}{q}\right)\right)$$

It is easy to see that $q^* = \frac{2p-1}{2p}$ from the proof of Theorem 1, and thus $\alpha \equiv \frac{1-2(1-p)(1-q)}{2q}$ and $\alpha - \beta = \frac{(2p-1)(1-q)}{q}$ (given that $\beta \equiv p\left(1 - \frac{q^*}{q}\right)$). Substituting we have that $\partial\mu/\partial p > 0$ if and only if

$$2f(\alpha - \beta)(1 - F(\alpha)) > f(\alpha)(1 - F(\alpha - \beta))$$

or in other words, $2\lambda_F(\alpha-\beta) > \lambda_F(\alpha)$, which proves the claim. $\square$

A.3 Proofs of Section 5

Proof of Theorem 3. Observe that larger values of γ decrease the lower support of the distribution H: if $\bar{\pi} > 1/2$ is the upper support for h, then $\bar{\pi}_\gamma = \bar{\pi} + (\bar{\pi} - 1/2)\gamma$ is increasing in γ and when $\gamma = 0$, $\bar{\pi}_\gamma = \bar{\pi}$. Moreover, all values of γ preserve the symmetry of H.

By Lemma 1, and since H is always symmetric, the probability of DeGroot mislearning does not depend of γ. For Bayesian learning, there is always either one ($r < r_B^*$) or two narratives ($r > r_B^*$). In the former case, Bayesians always learn. In the latter case, let η_L be the likelihood of the $\theta = L$ narrative and η_R be the likelihood of the $\theta = R$ narrative. Then the Bayesian society mislearns if and only if:

$$\frac{\eta_R\bar{\pi}}{\eta_R\bar{\pi} + \eta_L(1-\bar{\pi})} > 1/2$$

But note that the left-hand side is increase $\bar{\pi}$, so for every realization of r, mislearning can only become "more likely" as $\bar{\pi}$ increases.[16] Integrating over all of r shows that the likelihood of Bayesian mislearning is increasing in $\bar{\pi}$ and, in particular, is increasing in γ.

Next observe that when $\gamma = -1$, the Bayesian society mislearns with lower probability than the DeGroot society. To show this, note that when $\gamma = -1$, the density h is a Dirac-delta function at belief $\pi = 1/2$. Thus, all Bayesian agents initially agree, so there is a homogenous prior. By the improvement principle (see, for instance, Golub and Sadler (2017)), Bayesian agents must be able to outperform the DeGroot heuristic.

Finally, when $\gamma = \bar{\gamma}$, then H has full support on $[0, 1]$, so by Theorem 1(b), we know the Bayesian society mislearns with higher probability than the DeGroots. By the previous paragraph, we know that when $\gamma = -1$, then the Bayesian society mislearns. Because the mislearning probability is increasing in γ for Bayesians but constant for DeGroots, there must be a unique single-crossing γ^* that determines the phase transition. $\square$

Proof of Proposition 3. We prove the two parts of the result, which separate into DeGroot and Bayesian societies. For both, we fix $\theta = L$ for concreteness.

(i) *DeGroot society*: Because we have fixed $\theta = L$, we know that $m^p = L$ for targeted agents. Thus, there are two cases for $\pi_{i,1}$ for the DeGroot agent that depend on whether (i) agent i receives $m_i = R$ and $m^p = L$ or (ii) agent i receives both $m_i = L$ and $m^p = L$. The former case occurs with probability $(1-p)(1-q) + qr$ whereas the latter occurs with probability $p(1-q) + q(1-r)$. In the former, it is easy to verify her belief remains unaffected, i.e., $\pi_{i,1} = \pi_{i,0}$. In the latter case, applying Bayes' rule we see:

$$\pi_{i,1} = \frac{(1-p)^2\pi_{i,0}}{(1-p)^2\pi_{i,0} + (1-(1-p)^2)(1-\pi_{i,0})}$$

Thus, the expected belief update of agent i is given by:

$$\mathbb{E}[\pi_{i,1}|\pi_{i,0}] = ((1-p)(1-q)+qr)\pi_{i,0} + (p(1-q)+q(1-r))\frac{(1-p)^2\pi_{i,0}}{(1-p)^2\pi_{i,0} + (1-(1-p)^2)(1-\pi_{i,0})}$$

[16] "More likely" is a slight abuse of terminology, because for a given realization of r, the society either learns or does not almost surely. Formally, we mean that an increase in $\bar{\pi}$ can not transition society from mislearning to learning for this value of r.

Note that the change in belief from targeting, given by $\Delta \equiv \mathbb{E}[\pi_{i,1}|\pi_{i,0}] - \pi_{i,0}$ is:

$$\frac{\partial \Delta}{\partial \pi_{i,0}} = p\left(\frac{(2-p)(1-p)^2(p(1-q)+q(1-r))}{(\pi_{i,0}(2p^2-4p+1)+(2-p)p)^2} - (1-q)\right) - q(1-r)$$

When $p > 1/2$, note that $2p^2 - 4p + 1 < 0$, so $\partial\Delta/\partial\pi_{i,0}$ is strictly increasing in $\pi_{i,0}$, and thus Δ is strictly convex is $\pi_{i,0}$. Moreover, $\Delta(\pi_{i,0} = 0) = 0$ and $\Delta(\pi_{i,0} = 1) = 0$, so Δ is maximized at some unique $\pi^*_{i,0} \in (0,1)$. Note that when $\pi_{i,0} = 1/2$, then

$$\frac{\partial \Delta}{\partial \pi_{i,0}}(\pi_{i,0} = 1/2) = -(1-2(2-p)p)^2(p(1-q)+q(1-r)) < 0$$

Because $\Delta'' > 0$ everywhere, this implies that $\Delta'(\pi^*_{i,0}(r)) = 0$ for some $\pi^*_{i,0}(r) > 1/2$, which might depend on r.

Consider the set of agents with belief $\alpha = \pi_{i,0}$, denoted by $\mathcal{A}$, who are targeted. Recall that all DeGroots converge to a consensus belief π_∞:

$$\pi_\infty(r) = \int_{\alpha\in\mathcal{A}} \mathbb{E}[\pi_{i,1} \,|\, \alpha\,,\, m^p = L]\, h(\alpha)\, d\alpha + \int_{\alpha\notin\mathcal{A}} \mathbb{E}[\pi_{i,1} \,|\, \alpha\,,\, m^p = \varnothing]\, h(\alpha)\, d\alpha$$

If $\mathcal{A}$ is restricted to have some small measure ν (in prior space h), and the objective is to minimize π_∞ it is clear that the optimal choice of $\mathcal{A}$ is top pick an open interval around $\pi^*_{i,0}(r)$ given that Δ is continuous in $\pi_{i,0}$.

Finally, note that $\pi_\infty(r)$, under the optimal choice of $\mathcal{A}$ for each r, is continuous due to Berge's theorem of the maximum. Finally, because of continuity and the fact $\pi_\infty(r) > 1/2$ for all r, we know there exists an interval $(\underline{\pi}^*_0, \bar{\pi}^*_0)$ such that for all r, π^*_0 lies in this interval, with $1/2 < \underline{\pi}^*_0 < \bar{\pi}^*_0 < 1$. The probability that DeGroot agents mislearn is given by $\mathbb{E}_r[\mathbf{1}_{\pi_\infty(r)>1/2}]$, and because $\mathbf{1}_{\pi_\infty(r)>1/2}$ is monotone in $\pi_\infty(r)$, this implies that the optimal choice of π^*_0 to maximize this expectation also satisfies $\pi^*_0 > 1/2$.

(ii) *Bayesian society*: We claim that reducing the highest belief $\pi_{i,0}$ is equivalent to decreasing the likelihood of mislearning. Observe that $\pi_{i,2}$ (which is equal to $\pi_{i,T}$ for large T) is equivalent to decreasing the likelihood of mislearning for Bayesian agents. To see this, note that for Bayesian agent i, $\pi_{i,2}$ is strictly increasing in $\pi_{i,0}$, fixing the messages $\{m_i\}_{i=1}^N$, which all Bayesians are able to deduce by period 2. Note that if there exists an open interval $(\pi^1, \pi^2) \subset [\underline{\pi}, \bar{\pi}]$ such that all agents with $\pi_{i,0} \in (\pi^1, \pi^2)$ mislearn, the Bayesian society mislearns. By the assumption that $[\underline{\pi}, \bar{\pi}]$ has full support, targeting some open interval nearest $\underline{\pi}$ maximizes the probability of all agents learning when $\theta = L$. Thus, the optimal policy targets an agent who is most extreme near $\underline{\pi}$. $\square$

B Supplemental Material

B.1 Likelihood of Mislearning

We show that Bayesian agents perform worse than DeGroot agents on average. We adopt the environment from Setting B in Section 3, but we do not focus on a specific realization of the misinformation split. Instead, we look at the *likelihood* that society does not learn when we draw the split r from its true (uniform, in this example) distribution. How do Bayesian agents perform relative to DeGroot agents *on average*?

We track the beliefs of both societies:

Bayesian Population. At $t = 1$, each Bayesian agent i forms a posterior belief $\pi_{i,1}$ based on m_i according to Section 2. Once again, we know that $\pi_{j,1} > \pi_{j,0}$ if and only if $m_i = R$ and $\pi_{j,1} < \pi_{j,0}$ if and only if $m_i = L$. So every agent i can deduce all of the messages $\{m_j\}_{j=1}^N$ by period 2.

Among the collection of messages $\{m_j\}_{j=1}^N$, if the state is $\theta = L$, there are (roughly) $p(1-q) + q(1-r)$ proportion of L messages and $(1-p)(1-q) + qr$ proportion of R messages, whereas if the state is $\theta = R$, there are $(1-p)(1-q) + q(1-r)$ proportion of L messages and $p(1-q) + qr$ proportion of R messages. The state cannot be pinned down if there exists a value $r' \in [0,1]$ such that $p(1-q) + q(1-r) = (1-p)(1-q) + q(1-r')$; in this case, there exist *exactly* two realizations of r (the true r and another r') for which the given distribution of messages can be explained under two different states, $\theta = L$ (correct) and $\theta = R$ (incorrect). Moreover, because r is uniformly distributed on $[0,1]$, *both* of these scenarios are equally likely. This implies that $\pi_{j,2} = \pi_{j,0}$ because the messages provide no information about the state θ; consequently, the society of Bayesian agents does not learn.

Note that $(1-p)(1-q) + q(1-r')$ is maximized when $r' = 0$; thus it is sufficient to consider for what values of r the inequality $p(1-q) + q(1-r) \le (1-p)(1-q) + q$ holds; this corresponds to the values of r for which there is mislearning. Rearranging, we see that $r \ge 1 - \frac{1-2p(1-q)}{q} = 0.6$ when $p = 0.6$ and $q = 0.25$. Hence, the Bayesian society mislearns with probability 40% (given that r is uniformly distributed on $[0,1]$) in the setting with 25% misinformation.

DeGroot Population. DeGroot agents update in the same way as before (i.e., via Equations (1) and (2)) using $p = 0.6$. Thus, noting that $p(1-q) + q(1-r)$ of the messages are L and $(1-p)(1-q) + qr$ of the messages are R, DeGroot agents hold beliefs $\pi_{j,t}$ for all $t \ge 2$:

$$\pi_\infty(r) \equiv \int_0^1 \left((.4(.75) + .25r) \cdot \frac{.6\alpha}{.6\alpha + .4(1-\alpha)} + (.6(.75) + .25(1-r)) \cdot \frac{.4\alpha}{.4\alpha + .6(1-\alpha)} \right) d\alpha$$

Note that $\pi_\infty(r)$ is monotonically increasing in r and it can be shown that $\pi_\infty(r) \le 1/2$ if and only if there are less than 50% R messages; thus, $\pi_\infty(r) \le 1/2$ if and only if $(1-p)(1-q) + qr \le 1/2$, or in other words, $r \le \frac{1/2-(1-p)(1-q)}{q} = 0.8$. Thus, since DeGroot agents mislearn the true state only when $r \ge r^*$, we see they mislearn 20% of the time, which *outperforms* the Bayesian population by a factor of two. Recall that quantifying how much better DeGroots do than Bayesians in general is formally analyzed in Theorem 2 in the paper.

B.2 Polarization

The next example shows that in some cases, mean-preserving spreads (i.e., more polarization) can *improve* the learning outcomes of DeGroot agents.

Example 1. Consider a world where $\theta = L$ and as in Section 3, the misinformation is $q = 0.25$, so DeGroot societies (as well as Bayesian societies) mislearn with positive probability.

First, suppose H is distributed with a small right bias, as demonstrated in Figure 4a and Figure 4b. In the more polarized society of Figure 4a, many prior opinions start off initially quite misinformed, so not much misinformation on the right-side can support learning (i.e., the realization of r must be lower); in particular, $r \leq .281 \equiv r^*$ is required to support learning. Next, consider a decrease in polarization to the Dirac-delta function on the average opinion of H, as shown in Figure 4b. This increases the threshold of right-leaning misinformation that can be tolerated for learning to $r^* = .319$, and the corresponding probability that learning occurs also increases. Because moderate right-leaning agents are the most likely to be influenced by organic left-leaning news, less polarization helps the DeGroot society learn, as is the case with the Bayesians.

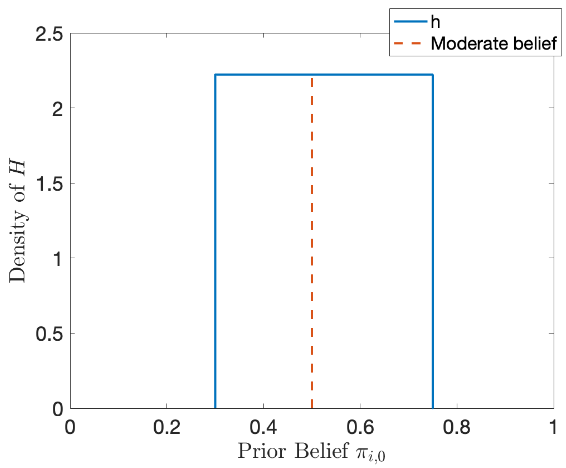

(a) Right-leaning but polarized society.

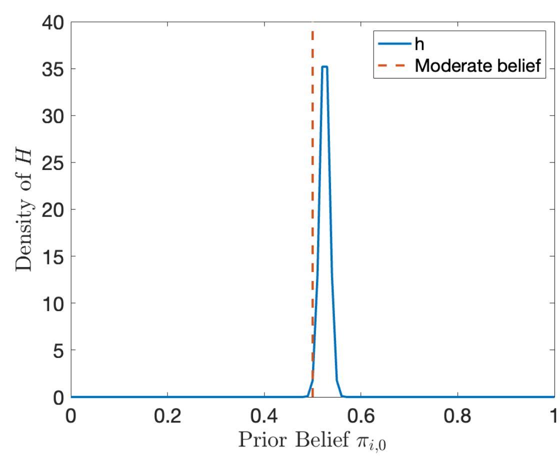

(b) Right-leaning but non-polarized society.

Figure 4. Two right-leaning distributions of prior beliefs (with the same mean belief), one of which is polarized and the other is not. The less polarized community mislearns less often because less evidence is needed to convince moderate right-leaning agents of $\theta = L$.

Conversely, suppose H is distributed with a small left bias, as demonstrated in Figure 5a and Figure 5b, so beliefs tend to support the correct state of the world. When (almost all) beliefs are concentrated just left of center (Figure 5b), there is positive probability of mislearning because there is some chance that right-leaning misinformation dominates and causes all agents to move closer to right-leaning ideas. However, these effects are mitigated when polarization increases, and learning occurs with probability 1 when it is sufficiently high, as shown in Figure 5a. This is because strong left-leaning believers are not swayed as much by right-leaning misinformation, but moderate right-leaning agents *can* be considerably convinced of left-leaning organic news. Thus, polarization generally helps when then initial belief distribution is already slanted toward the correct state. This is somewhat surprising given that additional polarization pushes more agents toward the incorrect belief of the world. Observe that this stands in contrast to polarization in Bayesian societies: whereas DeGroots always reach a consensus, and thus polarization can sometimes nudge the entire society closer to truth (in the aggregate), the ability of Bayesians to spin multiple narrative impedes consensus in the face of increasing polarization. □

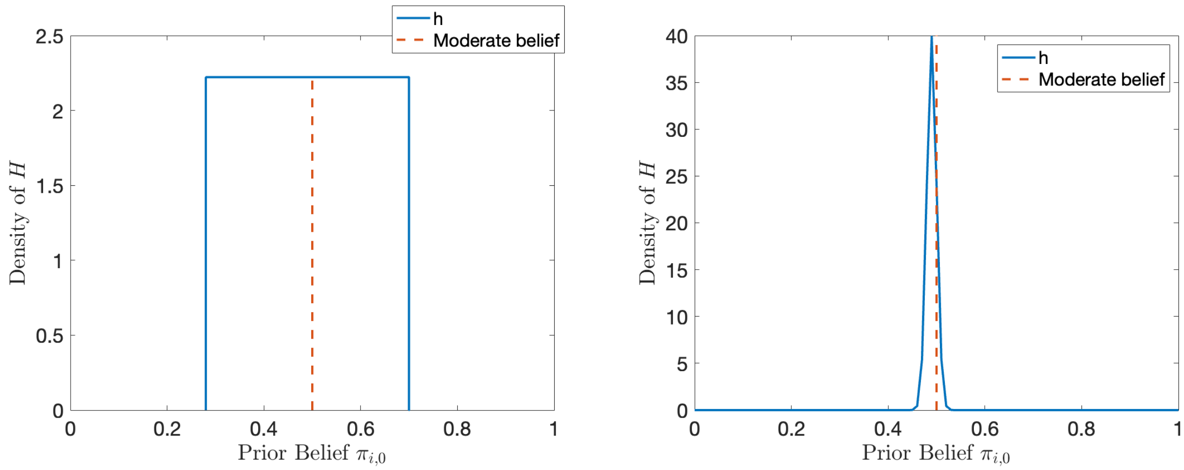

(a) Left-leaning but polarized society.

(b) Left-leaning but non-polarized society.

Figure 5. Two left-leaning distributions of prior beliefs (with the same mean belief), one of which is polarized and the other is not. The less polarized community mislearns *more often* because they are more susceptible to believing right-leaning ideas when right-leaning misinformation is more pervasive.

B.3 Mislearning Rates with High Misinformation

Let $q > q^*$. We denote by μ the ratio of the probability of DeGroot mislearning to the probability of Bayesian mislearning. Recall $\mu < 1$ by Theorem 1 (i.e., the "relative" frequency of DeGroot to Bayesian mislearning). Low values of μ indicate DeGroots do much better than Bayesians, whereas relatively larger values indicate Bayesians close the gap in mislearning more. Next, we try to reason about the conditions under which μ is increasing (i.e., Bayesian agents start picking up an advantage relative to DeGroot agents) or μ is decreasing (i.e., DeGroots learn more frequently relative to Bayesians) as a function of the strength of organic signals p and the misinformation in the system q. For, this we make the following standard definition of the *hazard rate* of a distribution:

Definition 1. The *hazard rate* $\lambda_G(\alpha)$ of a distribution G is given by $\lambda_G(\alpha) = \frac{g(\alpha)}{1-G(\alpha)}$ where $g(\cdot)$ is the density of cumulative distribution function G.

The hazard rate at point α^* corresponds the *likelihood* of the realization $\alpha \in (\alpha^*, \alpha^* + d\alpha)$ relative to the interval size $d\alpha$, conditional on $\alpha \geq \alpha^*$. Theorem 2 relates the hazard rate at specific points on the F distribution (local properties of F) to the sensitivity of μ to p (global property[17] of F). To gather some intuition for Theorem 2, let us look at three applications:

(i) *Uniform distribution*: Recall the uniform distribution we assumed for F in Section 3 showed that when $q = 0.25$, DeGroot agents mislearn half as often as the Bayesians did. How does this depend on p? The hazard rate is $\lambda_F(r) = \frac{1}{1-r}$, which is increasing in r because the likelihood of falling within an interval of fixed length dr is increasing as one conditions on higher values of r. Thus, while $\lambda_F(\alpha - \beta) < \lambda_F(\alpha)$, it is unclear its relation to $\lambda_F(\alpha)/2$. Some basic algebra reveals that $\lambda_F(\alpha - \beta) = \lambda_F(\alpha)/2$ for *all* values of α, β, so μ has no dependence on

[17] It might appear as though the hazard rate condition in Theorem 2 is still a local property, because we are looking at local changes in p, but recall that to characterize mislearning probability, one must consider all realizations r that come from F.

p. Thus, DeGroot societies always mislearn half as often as Bayesian ones on the uniform distribution.

(ii) *Unskewed misinformation (Figure 2 in the paper)*: This distribution of F is one where misinformation is likely to evenly balanced between L and R. Instead of computing the hazard rate for α and $\alpha - \beta$ explicitly, we will draw inferences by comparing it to the uniform distribution. When $r < 1/2$, the hazard rate is given by $\lambda_F(r) = \frac{4r}{1-2r^2}$ whereas when $r > 1/2$, the hazard rate is given by $\lambda_F(r) = \frac{4-4r}{2-4r+2r^2}$. It is easy to show the ratio of the hazard rate of this distribution to the uniform distribution is increasing on $r < 1/2$ and constant on $r > 1/2$. Thus, $2\lambda_F(\alpha - \beta) \leq \lambda_F(\alpha)$ and the ratio μ is decreasing in p, meaning that DeGroots do comparatively better with more precise organic information in an inverted V-distribution (i.e., the Bayesians are more than twice as likely to mislearn than the Bayesians). The intuition is simple: more moderate misinformation increases the likelihood that Bayesian agents can spin a narrative to their liking, whereas for DeGroots it corresponds to a greater likelihood of having balanced misinformation (that washes out), allowing the organic news to win out.

(iii) *Skewed misinformation (inverse of Figure 2 in the paper)*: Relative to application (ii), the opposite effect occurs here. When $r < 1/2$, the hazard rate is decreasing relative to the uniform distribution; when $r > 1/2$, the hazard rate is again constant. This means the opposite inequality holds (i.e., $\lambda_F(\alpha) \leq 2\lambda_F(\alpha - \beta)$) and by Theorem 2, the ratio μ is increasing in p. When the misinformation is more extreme, Bayesians are comparatively more resilient, and mislearn less than twice as often as their DeGroot counterparts. The high likelihood of very misleading misinformation is not as well-handled by the DeGroot agents relative to a Bayesian society. While the Bayesian society can use more extreme misinformation to dismiss an incorrect narrative, the DeGroot society falls victim to such misinformation.

B.4 Network Learning Dynamics and Multiple Messages

In Section 2, we considered a model of learning where all agents observe the beliefs of all other agents. However, this is often an unrealistic assumption, and there is a wide array of literature that considers the subtleties of learning when these observations are incomplete (see Golub and Sadler (2017) for a survey). The common approach to modeling this incompleteness is to assume there is a social network with pairwise connections that determines who can observe (or talk to) whom. In this context, the model in Section 2 assumes a *complete* social network, which simplifies the relevant dynamics to two periods.

In this section, we relax this assumption by considering arbitrary network architectures and the richer learning dynamics that occur over a longer time horizon. Under relatively mild conditions on the network structure, we show network learning leads to the same outcomes and insights found in the more parsimonious complete network setting, thereby rendering our assumption to be largely without loss of generality. We do this by building off of the previous literature on network learning in both Bayesian and DeGroot populations.

Network Preliminaries. We assume that all agents are arranged in an undirected social network $\mathbf{G}$. A link $i \leftrightarrow j$ denotes that agent i and agent j observe (or talk to) each other. We let $\mathcal{N}_i$ denote the neighborhood of agent i (i.e., the set of agents j with $i \leftrightarrow j$). The adjacency matrix $\mathbf{A}$ of $\mathbf{G}$ is a binary matrix with $[\mathbf{A}]_{ii} = 1$ and $[\mathbf{A}]_{ij} = 1$ if and only if $i \leftrightarrow j$. Let $d_i^{\mathbf{G}}$ be the degree of agent i in $\mathbf{G}$. We say a network $\mathbf{G}$ is k-regular if all agents have degree k.

We consider a discrete time model (as before) but with a much longer learning horizon T, $t = 0, 1, 2, \ldots, T$. We let $\tilde{\pi}_{i,t}$ denote the belief of agent i at time t under network learning, whereas $\pi_{i,0}, \pi_{i,1}$, and $\pi_{i,2}$ denote the beliefs of agent i at time $0, 1$, and 2, respectively, in the baseline model (i.e., a complete network).

Bayesian Population. Network learning in settings with fully rational (i.e., Bayesian) agents has been studied in many contexts, most notably in Acemoglu et al. (2011) and Gale and Kariv (2003). As is common in many models of Bayesian network learning,[18] we assume that the network $\mathbf{G}$ and initial priors $\pi_{i,0}$ are common knowledge.[19] Bayesian agents observe the beliefs of all agents in their neighborhoods for all $t \geq 1$ (i.e., agent i observes at time t the beliefs from $t-1$, $\{\pi_{j,t-1}\}_{j \in \mathcal{N}_i}$). Our next result shows that terminal beliefs in network learning indeed converge to the terminal beliefs of the baseline model:

Claim 1. *Suppose* $\mathbf{G}$ *is connected. Then as* $T \to \infty$, $\tilde{\pi}_{i,T} \to \pi_{i,2}$.

This claim follows directly from Mueller-Frank (2013). While agents do not hold a common prior about θ, common knowledge of the heterogenous priors $\{\pi_{j,0}\}_{j=1}^{N}$ allows agents to recalibrate the (updated) beliefs they see to their own prior. It is clear that the private information at $t = 1$ (i.e., the messages) are drawn from a finite partition of the θ state space (conditional on misinformation split r). Thus, by Theorem 4 of Mueller-Frank (2013), all Bayesian agents uncover the private information (i.e., $t = 1$ messages) of all other agents (including non-neighbors) in the network as $T \to \infty$, as is the case at $t = 2$ in the baseline model.

DeGroot Population. Due to demanding assumptions about the reasoning abilities of Bayesian agents, "rule-of-thumb" learning has become a popular alternative model. The most common model is that of Degroot (1974), and later expanded upon in works such as Golub and Jackson (2010) and DeMarzo et al. (2003). In these models, agents are assumed to update their beliefs using the simple heuristic of taking linear combinations of their neighbors' beliefs. Formally, agent i forms belief $\pi_{i,t+1}$ at each time t by computing:

$$\pi_{i,t+1} = \frac{1}{1 + d_i^{\mathbf{G}}} \left(\pi_{i,t} + \sum_{j \in \mathcal{N}_i} \pi_{j,t} \right)$$

Our next result provides conditions under which DeGroot learning over the network $\mathbf{G}$ leads to the same terminal beliefs as in our baseline model:

Claim 2. *Suppose* $\mathbf{G}$ *is a connected,* k*-regular network. Then as* $T \to \infty$, $\tilde{\pi}_{i,T} \to \pi_{i,2}$.

This claim follows directly from Golub and Jackson (2010). First, by Proposition 1 in Golub and Jackson (2010), observe that consensus is reached (as in the baseline model) because the normalized adjacency matrix $\mathbf{A}$ is irreducible and aperiodic, the former following from the connectedness assumption and the latter following from a positive diagonal on $\mathbf{A}$. Second, by Theorem 3 in Golub and Jackson (2010), the consensus belief of the agents as $T \to \infty$ is given by $\tilde{\pi}_{i,\infty} = \sum_{j=1}^{N} v_j^{\mathbf{G}} \pi_{j,1}$ for all agents i, where $v_j^{\mathbf{G}}$ is the (eigenvector) centrality of agent j (according to the row-stochastic normalized adjacency matrix $\mathbf{A}$). Because $v_j^{\mathbf{G}} = d_j^{\mathbf{G}} / \sum_{\ell=1}^{N} d_\ell^{\mathbf{G}}$, we obtain by k-regularity that $\tilde{\pi}_{i,\infty} = \frac{1}{N} \sum_{j=1}^{N} \pi_{j,1} = \pi_{i,2}$.

[18] In addition to Acemoglu et al. (2011) and Gale and Kariv (2003), see Mueller-Frank (2014) and Mossel et al. (2014).

[19] An alternative assumption, which does not require strong common knowledge assumptions of non-neighbor priors or the network structure, is that the size of the smallest neighborhood grows unboundedly as $N \to \infty$.

Observe that Claim 2 requires an additional condition not present in Claim 1, which is that no agent is more "influential" than any other agent in the network $\mathbf{G}$, as measured by her degree. This is easily satisfied by many network topologies, including several classes of random networks such as Erdos-Renyi networks (where links between agents occur uniformly at random).[20]

Multiple Messages. Let us consider the complete network setting of Section 2 for simplicity, but note that the reduction from arbitrary network learning discussed previously still applies.

In a Bayesian society with $N \to \infty$, by the strong law of large numbers, the first round of messages reveals the true fraction of R messages, ρ_R, and the true fraction of L messages, ρ_L, almost surely. Obtaining additional messages in subsequent rounds does not alter the (almost surely) known values of ρ_R or ρ_L, thus, learning is entirely unaffected by more incoming messages.

In a DeGroot society, after the first round of messages, agents converge to a consensus about θ which is a function of ρ_R (and ρ_L) alone. When H is symmetric, whether $\rho_R > 1/2$ or $\rho_L > 1/2$ determines if the consensus, call it π_2, lies more toward state R (i.e., $\pi_2 > 1/2$) or state L (i.e., $\pi_2 < 1/2$). By the martingale property of Bayesian updating, it is easy to see that $\mathbb{E}[\mathrm{BU}(\pi_2) \mid \rho_R > 1/2\,;\, \pi_2 > 1/2] > 1/2$ and $\mathbb{E}[\mathrm{BU}(\pi_2) \mid \rho_R < 1/2\,;\, \pi_2 < 1/2] < 1/2$ (where BU is the Bayesian update for DeGroot agents conditioning on the message, given by Equation (1) and Equation (2)). Therefore, one can show by induction that beliefs remain on the same side of belief $1/2$ as they are at $t = 2$ for all $t \geq T$, even with additional messages. Consequently, the likelihood of (mis)learning is unaffected by any further stream of messages.

B.5 Robustness: Learning Metric and Finite Populations

We conduct two robustness checks on our main learning results. First, we consider how our results change when looking at the *expected* fraction of mislearning agents, instead of the binary metric of whether the entire society learns or not. Second, we test the sensitivity of our large population assumption (i.e., $N \to \infty$) by simulating learning in settings with finite N.

B.5.1 Learning Metric

We consider the alternative learning metric of the expected proportion of the population that mislearns the true state. In particular, we look at (i) the environments where DeGroots or Bayesians perform better under this metric, and (ii) how the targeting policy changes under this other learning objective.

DeGroot vs Bayesian. How is learning affected when one evaluates the expected fraction of mislearning agents? Provided that H is sufficiently polarized (i.e., there are few moderate left or right-leaning agents and most agents have relatively strong opinions), the expected fraction of mislearning Bayesian agents is approximately half of the Bayesian mislearning rate; conversely, the expected fraction of mislearning DeGroot agents is exactly the mislearning rate. The former can be seen from Proposition 2: all agents in society learn when there is a single narrative, but when there are multiple narratives (and H is symmetric), only those who have priors that agree with the truth (i.e., 50%) will take the correct action. The latter can be seen from the fact that the DeGroot society always comes to consensus, so the two notions of learning coincide.

[20] k-regularity will hold approximately for large N in dense ER networks; see for instance, Avella-Medina et al. (2020) and Dasaratha (2020).

The comparison between Bayesian and DeGroot societies using this learning metric is best highlighted using the hazard rate analysis of Section 4.4 and Appendix B.3 in the high-misinformation regime (i.e., $q > q^*$), which depends on the distribution of misinformation F:

(i) *Uniform distribution*: When F has a uniform distribution, the Bayesian society mislearns twice as often as the DeGroot society. Thus, under the new learning metric, the expected fraction of mislearning agents is *the same* for both Bayesians and DeGroots.

(ii) *Unskewed distribution*: When F has an unskewed distribution (i.e., misinformation is more likely to be evenly split between ideologies), DeGroots perform better than Bayesians relative to the base case of the uniform distribution. Thus, under the new learning metric, DeGroots outperform Bayesian agents in the expected fraction of mislearning agents.

(iii) *Skewed distribution*: When F has a more skewed distribution (i.e., misinformation is more likely to come mostly from one ideology), Bayesians perform better than DeGroots relative to the base case of the uniform distribution. Thus, under the new learning metric, Bayesians would outperform DeGroot agents in the expected fraction of mislearning agents.

These cases make it clear that the main message of our paper —that reasoning abilities have an ambiguous effect on learning outcomes and that DeGroot agents can outperform Bayesian agents— is not an artifact of the learning definition but rather a fundamental property of learning in the presence of misinformation. Under this metric, the outcomes depend on both the level of misinformation (i.e., whether $q > q^*$ or $q < q^*$) and the concavity/convexity of f. An interesting implication of the above metric is that it suggests that DeGroot agents are better at learning the state when misinformation is likely to come from both sides of the spectrum, which is likely the case with most controversial political issues. On the other hand, Bayesian agents can be better at learning the state when misinformation is mostly one-sided, e.g., that it argues for the earth being flat.

Targeting Policies. We consider a targeting policy where the regulator wants to minimize the expected proportion of agents who mislearn. As before, we assume $q > q^*$ and the true state is $\theta = L$, which is known to the regulator. Because DeGroot agents always converge to a consensus, the regulator does not change her targeting policy because either all agents learn or none do; we let π_D^* denote the belief of the optimal DeGroot target which is the same as π^* in Proposition 3. However, the policy will change for the Bayesian society in a subtle way. For the most likely split of misinformation that generates two narratives, the regulator targets the agent whose posterior belief is barely to the right of belief $1/2$. We illustrate how this targeting policy works in practice for the three settings considered before:

(i) *Uniform distribution*: For the uniform distribution, when there are two narratives, no agents change their prior beliefs (and when there is only one, all agents learn anyway). Therefore, the optimal targeting policy is to target the most moderate right-leaning agent. This stands in contraposition to Proposition 3, as the regulator is better off targeting more moderate Bayesian agents and relatively more extreme DeGroots (i.e., $p_B^* < p_D^*$).

(ii) *Unskewed distribution*: When F follows the inverse-V distribution, there are two cases to consider. The first case is that the misinformation split $r = 1/2$ admits two narratives; this occurs when $q > \frac{2(2p-1)}{4p-1} > q^*$. In this case, there is some $r^* < 1/2$ that explains the $\theta = R$ narrative; however, this narrative is less compelling than the $\theta = L$ narrative (because

f is largest at $r = 1/2$). Thus, some moderate right-leaning agent with belief $\pi_B^* > 1/2$ will be the optimal target. As q increases, the r^* corresponding to the R narrative moves closer to $1/2$ and thus becomes more likely. Thus, for low q, π_B^* will be close to 1 (target the extremists, as in Proposition 3) whereas for high q, π_B^* will target the most right-leaning moderates (as in the uniform distribution).

The second case is that the misinformation split $r = 1/2$ admits only one narrative ($q^* < q < \frac{2(2p-1)}{4p-1}$). Then there are (barely) two narratives when $r = r^*$ satisfies $(1-p)(1-q) + qr^* = p(1-q)$, or $r^* = \frac{(2p-1)(1-q)}{q} > 1/2$, and this is the most likely split of misinformation conditional on the existence of two narratives. The other narrative (the $\theta = R$ narrative) has likelihood almost zero (this narrative occurs at $r = 0$), so only the most polarized right-leaning agents will prescribe to this narrative over the true narrative for $\theta = L$. Thus, the optimal policy is exactly the same as in Proposition 3: the regulator should target the most extreme Bayesian agents, which are strictly more extreme than the optimal DeGroot target.

(iii) *Skewed distribution*: When F follows the V-distribution, the most likely split of misinformation that admits two narratives is $r = 1$ (note that $r = 0$ is also the most likely, but when $\theta = L$, the L narrative is unique). There are always two narratives in this case and the $\theta = R$ narrative corresponds to $r^* = \frac{1-2p(1-q)}{q}$. When $q > \frac{2(2p-1)}{4p-1} > q^*$, $r^* > 1/2$, so increasing q increases r^* and makes the $\theta = R$ narrative more likely. Consequently, π_B^* decreases and the regulator targets more moderate right-leaning agents. When $q^* < q < \frac{2(2p-1)}{4p-1}$, $r^* < 1/2$, so increasing q increases r^* but makes the $\theta = R$ narrative *less* likely. Consequently, π_B^* *increases* and the regulator targets more extreme right-leaning agents. In particular, we recover the policy of Proposition 3 (target the extremists) when q has the intermediate value of $\frac{2(2p-1)}{4p-1}$.

Discussion — These cases highlight an interesting feature of the Bayesian targeting policy under this alternative learning metric: the optimal target depends on the quantity of misinformation q and can even be *non-monotone* in q.

In case (ii) (as with any concave distribution for F), with relatively low misinformation, the policy is the same as in Proposition 3: the regulator should target extremists because these are the most stubborn agents to convince. However, with relatively high misinformation, the policy changes to focus on more moderate right-leaning agents, because abundant misinformation will necessarily confound learning for the extremists.

In case (iii) (as with any convex distribution for F), the regulator employs a non-monotone targeting policy. When misinformation is small or large, both narratives conclude that misinformation is heavily skewed toward one ideology, which inhibits learning and admits an optimal policy of targeting moderate right-leaning agents. However, when misinformation is moderate, the incorrect narrative involves an even split of misinformation, which is unlikely; consequently, the regulator should target the most difficult agents to convince, the extremists, as in Proposition 3.

B.5.2 Finite Populations

Recall that Theorem 1 shows that when $N \to \infty$ and in settings where misinformation is high ($q > q^*$), DeGroot societies outperform Bayesian societies, whereas the converse holds when misinformation is low ($q < q^*$). We consider how robust these findings are to a *finite* population of agents.

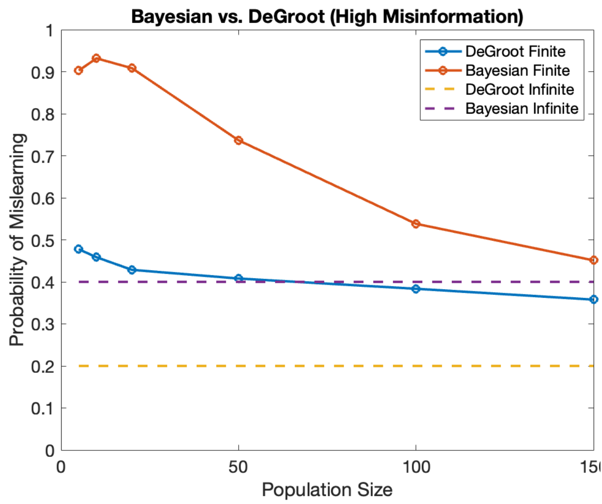

(a) High misinformation regime

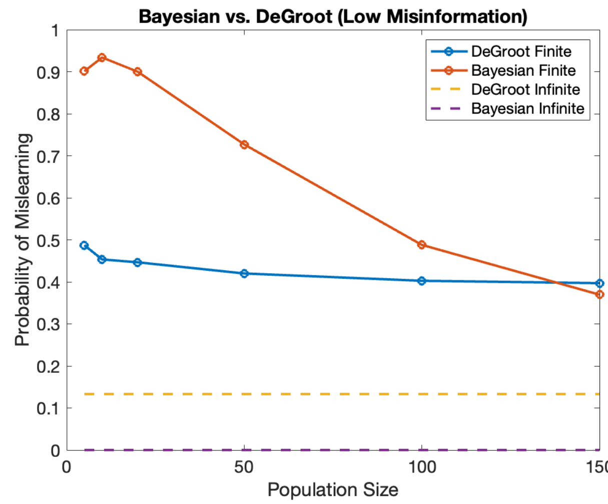

(b) Low misinformation regime

Figure 6. Beliefs in finite populations

High Misinformation. When there is high misinformation, DeGroot agents consistently outperform Bayesian agents for all finite populations, as shown in Figure 6a. Notably, Bayesian agents converge to their theoretical long-run mislearning average more quickly than DeGroot agents.[21] With few agents in the population, the small amount of information on the state θ allows for more wild narrative telling. We denote this kind of narrative telling as a *noise-based narrative* to distinguish it from the more subtle narrative telling identified in Section 4.2. Noise-based narratives arise from the agents believing that the existing (and limited) content is just by happenstance in opposition to one's priors, but is not indicative of θ. Once the population reaches a critical mass, the noise-based narrative dwindles and only the misinformation narrative spin identified in Section 4.2 persists, keeping Bayesian mislearning at or above 40%.

Low Misinformation. When there is low misinformation, DeGroot agents can still outperform Bayesian agents in small populations, as shown in Figure 6b. This is again related to the noise-based narrative that Bayesian agents can spin in finite populations. As seen in Figure 6b, with an infinite population, the Bayesian society learns almost surely, and so there are no misinformation narratives (of the form in Section 4.2) that can be told. As the population increases, the noise-based narratives attenuate and once the population is large enough (e.g., $N > 150$), the Bayesians begin to outperform the DeGroots, as predicted in Theorem 1 for the low misinformation regime.

[21]While both Figure 6a and Figure 6b show DeGroots bounded away from their long-run average, one can verify via simulation that after $N > 10000$, DeGroots are within 1% of their $N \to \infty$ learning rate (i.e., convergence is slower). The plots are capped at $N \leq 150$ because of the numerical issues associated with Bayesian learning when N is large but finite.