

A Model of Online Misinformation*

Daron Acemoglu[†]

Asuman Ozdaglar[‡]

James Siderius[§]

January 11, 2022

Abstract

We present a model of online content sharing where agents sequentially observe an article and must decide whether to share it with others. This content may or may not contain *misinformation*. Agents gain utility from positive social media interactions but do not want to be called out for propagating misinformation. We characterize the (Bayesian-Nash) equilibria of this social media game and show sharing exhibits strategic complementarity. Our first main result establishes that the impact of homophily on content virality is non-monotone: homophily reduces the broader circulation of an article, but it creates echo chambers that impose less discipline on the sharing of low-reliability content. This insight underpins our second main result, which demonstrates that social media platforms interested in maximizing engagement tend to design their algorithms to create more homophilic communication patterns (“filter bubbles”). We show that platform incentives to amplify misinformation are particularly pronounced when there is greater polarization and more divisive content. Finally, we discuss various regulatory solutions to such platform-manufactured misinformation.

Keywords: echo chambers, fake news, filter bubbles, homophily, misinformation, networks, social media

JEL classification: D83, D85, P16

*This paper builds on and replaces our earlier working paper on the same topic, entitled “Misinformation: Strategic Sharing, Homophily, and Endogenous Echo Chambers”. We are grateful to Hunt Allcott, Jackie Dewart, Matthew Gentzkow, Ali Makhdoumi, Azarakhsh Malekian, Mohamed Mostagir, Ro’ee Levy, Francesca Parise, Davy Perlman, Alexander Wolitzky, and numerous theory lunch seminar participants, as well as NSF and the Sloan Foundation.

[†]Massachusetts Institute of Technology, NBER, and CEPR, daron@mit.edu

[‡]Massachusetts Institute of Technology, asuman@mit.edu

[§]Massachusetts Institute of Technology, siderius@mit.edu

1 Introduction

Social media has become a major source of information for many Americans. Leading up to the 2016 US presidential election, around 14% of Americans indicated social media as their primary source of news (Allcott and Gentzkow (2017)), and by 2019, over 70% of Americans reported receiving at least *some* of their news from social media (Levy (2021)). At the same time, there is growing concern about viral misinformation in social media, including made-up news stories such as those claiming that there were no mass shootings under Donald Trump or that Hillary Clinton approved ISIS weapon sales.¹ Some recent evidence also suggests that misinformation on social media has impacted critical decisions such as vaccinations against COVID-19 (see Pennycook et al. (2018); Pennycook et al. (2020b)).

Although there is yet no consensus on what promotes the spread of falsehoods and misleading content on social media, two sets of factors have been emphasized. The first is the presence of echo chambers, which arise when individuals communicate and share content with like-minded users (Sunstein (2018), Lazer et al. (2018)). Törnberg (2018) and Vicario et al. (2016) show that echo chambers reinforce existing political viewpoints and tend to propagate misinformation. Echo chambers may be an unavoidable side effect of social media, which allow individuals to choose what and who to listen to a much greater extent than traditional media. However, there is also evidence that echo chambers are a result of the “filter bubbles” that platform algorithms create (see Levy (2021) on Facebook). The second factor conjectured to have fueled misinformation is the general political polarization in many countries, and especially the United States,² and there is some preliminary evidence suggesting that polarization has contributed to selective exposure to questionable content on social media as well (Guess et al. (2018)). Despite the importance and salience of these issues, we do not currently have a framework to understand how online interactions impact the spread of misinformation and what factors shape incentives for sharing low-reliability content.

In this paper, we develop a parsimonious model of online sharing behavior in the presence of misinformation, and as a first step, we focus on the behavior of fully Bayesian agents.³ Our model is inhabited by a set of N agents, each of which has a prior about the state of the world (“ideological bias”), and is connected to the rest of the users via a network, which is given by agents’ friends and acquaintances, and is also shaped by the algorithms of the social media platform. A news article, defined by an underlying type (truthful or containing misinformation), a message (right-wing or left-wing), and a level of reliability (which determines the likelihood of misinformation), is then seeded at

¹See <https://www.snopes.com/fact-check/mass-shootings-under-trump/> and <https://www.cnbc.com/2016/12/30/read-all-about-it-the-biggest-fake-news-stories-of-2016.html>, respectively.

²While there has been some debate about whether polarization has been mainly among politicians (see Fiorina et al. (2008) and Prior (2013)), there is considerable evidence that polarization has also risen among the general public (see Pew Research Center (2014) and Abramowitz (2010)).

³Myopic reactions and biased behavior appear to play some role, for example, via the “confirmation bias” in social media behavior (see, e.g., Buchanan (2020) and Pennycook and Rand (2018)), but we believe that the Bayesian benchmark we construct already generates a number of empirically-relevant and rich results. We view incorporating realistic and relevant behavioral biases as a next step in this research agenda.

one of the agents. The message and the level of reliability of the article are common knowledge, while whether it is truthful or contains misinformation is unobserved, and agents form beliefs about this component.

Given these beliefs, the agent in question decides whether to ignore, dislike, or share the news article. If it is shared, the article moves from the agent sequentially to her connections on social media, who are then faced with the same choices. If the article is ignored or disliked, it does not get past the agent. We assume that agents receive utility when their shared content is re-shared and incur a cost when it is disliked. The former aspect captures the role of positive engagement in social media, while the latter represents the reputation loss from being called out for sharing content containing misinformation. Agents additionally receive utility from disliking (or calling out) items that they believe contain misinformation.

We characterize the Bayesian-Nash equilibria of this sequential game and show that these equilibria always exist and are in cutoff strategies. In particular, our payoff structure implies that an individual will share any item that she believes is truthful with a high probability and will dislike articles that she believes to contain falsehoods. Items with intermediate beliefs will be ignored. Beliefs about the truthfulness of articles are formed on the basis of the article's reliability and message, and agents' prior beliefs/ideology. Moreover, we establish that ours is a game of strategic complements: when others are more likely to share an item, each agent also becomes more likely to do so. As a result, we show that the set of equilibria forms a lattice, with well-defined most-sharing and least-sharing equilibria. We present a number of comparative statics that describe how these extremal equilibria shift. For example, low-reliability articles are shared less, while articles that are more sensational are shared more. As expected, greater reputation costs reduce sharing uniformly.

We present two main results. First, we study the implications of the (social media) network structure. We establish non-monotone comparative statics with respect to the degree of homophily (which determines how likely agents are to be connected to others who are ideologically similar to them). Low levels of homophily ensure that agents are likely to be exposed to cross-cutting content, including "counter-attitudinal articles" that advocate views opposed to theirs. This in turn ensures that misinformation is unlikely to survive for very long. Perhaps paradoxically, for high-reliability articles, an increase in homophily reduces content virality. This is because greater homophily makes it less likely that an article escapes a given community, thus reducing its circulation throughout the network.

More interestingly and ominously, when relevant news items have low reliability, homophily increases virality. This is because, countering the circulation effect, high homophily also creates a perverse incentive effect: knowing that shared articles will be seen by like-minded individuals, agents become more likely to share questionable content. Strategic complementarities amplify this effect, because when others are expected to share, the benefits from sharing and reduces the likelihood of being called out for spreading misinformation are greater. It is particularly telling that homophily leads to the viral spread of low-reliability content, which are the ones more likely to contain misinformation.

We also show that political polarization and politically divisive articles are more likely to spread

virally when they are low-reliability and the level of homophily is already high, generating an echo chamber-like social media environment. Strategic complementarities tend to amplify these pernicious effects of political polarization and divisive content as well.

Our second main result turns to social media platforms' algorithm design choices and establishes a striking result: when the relevant articles have low reliability, social media platforms that aim to maximize engagement will tend to design algorithms that increase homophily and create filter bubbles, propagating misinformation (because virality means greater engagement). High-reliability content tends to spread anyway because most users recognize it as such and share it, and low homophily contributes to its spread by increasing its circulation throughout the network. In contrast, low-reliability content will be ignored or disliked by agents who disagree with its message and believe it to contain misinformation. The only way of increasing engagement with low-reliability content is thus to ensure that it remains among users ideologically aligned with its message, who would be willing to share it with like-minded others without fear of being called out by users with different ideologies. Hence, creating filter bubbles becomes an attractive strategy for engagement-maximizing platforms. It is particularly troublesome that such filter bubbles are created precisely when the relevant content is low-reliability and likely to contain misinformation.

If platform algorithms are propagating misinformation, can public policy counter and discourage this type of behavior?⁴ The answer is yes, but with some caveats. In the last part of the paper, we discuss four different types of regulatory policies, and in each case, we show how they may reduce misinformation but also point out the possibility that, if they are not designed well, they can backfire and exacerbate the problem.

First, we discuss regulations that force platforms to reveal the provenance of articles, making it easier for users to identify falsehoods (e.g., claims originating from less reputable sources, such as InfoWars). Though generally useful, such provenance information can backfire when it instills a sense of “false security” that the content has been verified by previous social media users. Second, we look at potential censorship of articles identified by a regulator as likely containing misinformation. While censorship can help reduce the viral spread of misinformation as well, it simultaneously generates an “implied truth” effect (Pennycook et al. (2020a)) that contributes to the viral spread of questionable content that escapes censorship. Third, we discuss (average) “performance targets”, where the regulator places limits on the amount of misinformation that circulates on the platform. Such targets tend to better align user and regulator preferences, but unless they appropriately target sensational content, platforms can meet the required performance limits by focusing on less sensational material that was unlikely to become viral in any case, and the least reliable content may still continue to circulate virally. Lastly, we show how direct regulation of platform algorithms and filter bubbles can reduce misinformation, but also point out that the non-monotone effects of homophily imply that such regulations need to be finely calibrated.

⁴As of August 2021, federal law protects social media platforms from being held responsible for content posted by its users (see Section 230 of the Communications Decency Act of 1996, discussed by <https://hbr.org/2021/08/its-time-to-update-section-230>).

Related Literature. Our paper builds on a large body of work on models of misinformation. In addition to the literature mentioned previously, several other papers in this literature are related to our findings.

Much previous work has focused on the susceptibility of boundedly-rational agents to engage with misinformation. In [Acemoglu et al. \(2010\)](#) and [Acemoglu et al. \(2013\)](#), the existence of persuasive agents can impede information aggregation and enable misinformed beliefs to survive, and sometimes even become dominant, in the population. In [Mostagir et al. \(2021\)](#) and [Mostagir and Siderius \(2021\)](#), a strategic principal who wants to persuade agents of an incorrect belief can distort the learning process by leveraging social connections and echo chambers to propagate misinformation. Similarly, models of misinformation “contagion”— without Bayesian agents or strategic decisions — have been studied in [Budak et al. \(2011\)](#), [Nguyen et al. \(2012\)](#), and [Törnberg \(2018\)](#). Our contribution relative to this literature is the possibility that misinformation spreads because of the strategic interactions of Bayesian agents and is exacerbated by profit-maximizing platform algorithms.

There is a growing literature on information design by platforms, building for the most part on the concept of Bayesian persuasion ([Kamenica and Gentzkow \(2011\)](#) and [Kamenica \(2019\)](#)). [Candogan and Drakopoulos \(2017\)](#) study how a platform with private knowledge of content’s accuracy should optimally signal to rational users whether to engage with it, while [Chen and Papanastasiou \(2021\)](#) and [Keppo et al. \(2019\)](#) consider more manipulative actions by platforms, including strategic seeding of information or “cheap talk” signals about quality. Also related are works on reputation and media bias. Motivated by the 2016 presidential election, [Allcott and Gentzkow \(2017\)](#) study the incentives of certain outlets to present misleading news, while [Gentzkow and Shapiro \(2006\)](#), [Hsu et al. \(2020\)](#) and [Allon et al. \(2019\)](#) explore other strategic reasons for media bias. Our paper contributes to this literature by highlighting the role of ideological leaning, strategic interactions, ideological homophily, and platform algorithms.

The most closely related work to ours is [Papanastasiou \(2020\)](#), who studies a model where agents hold heterogenous ideological beliefs and digest (and potentially share) a news article sequentially. Our work is different in three important dimensions. First, [Papanastasiou \(2020\)](#) focuses on costly inspection, which makes sharing decisions strategic substitutes, while our model generates strategic complementarities, because individuals care about further shares of the content they share.⁵ All of our results and formal analysis turn on strategic complementarities. Second, and relatedly, echo chambers play no role in [Papanastasiou \(2020\)](#).⁶ Third, our analysis of engagement-maximization by the platform and its implications for the spread of low-reliability content has no counterpart in [Papanastasiou \(2020\)](#) or any other work in this area we are aware of.⁷

⁵Our reading of the evidence is that strategic complementarities are more relevant for social media behavior than strategic substitutabilities. For example, [Eckles et al. \(2016\)](#) find evidence that feedback or “encouragement” from peers about Facebook posts have contributed significantly to future behavior and posting. See also [Taylor and Eckles \(2018\)](#) and [Aral and Dhillon \(2018\)](#).

⁶As already noted, echo chambers appear central to the spread of misinformation in practice. See, for example, [Lee et al. \(2011\)](#), [Törnberg \(2018\)](#), [Centola \(2010\)](#), and [Centola and Macy \(2007\)](#)

⁷[Papanastasiou \(2020\)](#) also discusses platform incentives, but assumes that the platform is interested in limiting misinformation. Our reading of the evidence in this instance, too, favors our interpretation, where platforms such as Facebook are (or at the very least used to be before regulatory pressure mounted) fairly indifferent to the presence of

The rest of the paper is organized as follows. The next section introduces our basic environment and describes the information structure and payoffs. Section 3 characterizes the (Bayesian-Nash) equilibria of this model and provides some basic comparative static results. Section 4 studies the effects of homophily by focusing on a special class of sharing networks that correspond to a set of “islands” of like-minded individuals who are less closely linked to those in other islands. Section 5 endogenizes the sharing network as a result of the algorithmic choices of the platform that aims to maximize engagement. Section 6 discusses a range of regulations aimed at containing misinformation. Section 7 concludes, while all proofs are provided in the Appendix.

2 Model

There is an underlying state of the world $\theta \in \{L, R\}$, for example, corresponding to whether the left-wing or the right-wing candidate is more qualified for political office. Agents have heterogeneous prior (ideological) beliefs about θ , and agent i 's prior that $\theta = R$ is denoted by b_i with an *ex ante* distribution $H_i(\cdot)$, which may or may not be the same across agents.

Sharing Network. We assume there are N agents in the population, who share a news item according to a *sharing network* defined by matrix \mathbf{P} of link probabilities:

$$\mathbf{P} \equiv \begin{pmatrix} 0 & p_{12} & \cdots & p_{1N} \\ p_{21} & 0 & \cdots & p_{2N} \\ \cdots & \cdots & \cdots & \cdots \\ p_{N1} & p_{N2} & \cdots & 0 \end{pmatrix}$$

where p_{ij} is the probability that agent i has a link to agent j . We define agent i 's neighborhood \mathcal{N}_i as the set of agents attached to her with an outgoing link, and denote her degree or the size of her neighborhood by $|\mathcal{N}_i|$. The sharing network reflects both an individual's social circle and the algorithms the platform uses for promoting shared content. The news item in question could be a news article or a post by one of the users, and throughout we refer to it as an “article”.

Misinformation and News Generation. Each article has a three-dimensional type (r, m, ν) . Here, $r \in [0, 1]$ indicates the *reliability* of the news, and $m \in \{L, R\}$ is the *message*, which corresponds to the article's viewpoint, for example, whether it argues for a left-wing or right-wing idea. Finally, ν is the article's *veracity*, which can either be \mathcal{T} , to indicate the article is truthful, or \mathcal{M} , to indicate the article contains *misinformation*.⁸

misinformation but strongly prioritize engagement maximization.

⁸Our focus in this paper is on misinformation, interpreted as items containing misleading information or arguments that can influence (a subset of) the public. Articles containing misinformation are in practice much more numerous than those that can be classified as “fake news”, which explicitly propagate demonstrably false information (e.g., Egelhofer and Lecheler (2019), Allen et al. (2020), Guess et al. (2019), Grinberg et al. (2019)). For example, according to this definition a news item that favorably describes a report denying climate change, without putting this in the context of hundreds of other reports reaching the opposite conclusion or mentioning the criticisms that it has received from experts, contains misinformation.

We assume that, at the beginning of the game, the type vector (r, ν, m) is drawn according to the following i.i.d. process:

- (i) The article's reliability $r \in [0, 1]$ is drawn from a continuous distribution F with density f .
- (ii) The veracity of the article is $\nu = \mathcal{T}$ (contains truthful content) with probability $\phi(r)$ or is $\nu = \mathcal{M}$ (contains misinformation) with probability $1 - \phi(r)$. We assume that ϕ is increasing and differentiable in r , and satisfies $\phi(0) = 0$ and $\phi(1) = 1$, so that the least reliable article always contains misinformation, and as the degree of reliability increases, the likelihood of misinformation monotonically declines and reaches zero.
- (iii) If $\nu = \mathcal{T}$ (the article is truthful), then its message is generated as $m = \theta$ with probability $p > 1/2$. Conversely, if $\nu = \mathcal{M}$ (the article contains misinformation), then its message is generated as $m = \theta$ with probability $q \leq 1/2$ and is weakly anti-correlated with the truth.

While m and r are common knowledge (for example, the message m is directly observed and reliability depends on certain commonly-observed characteristics such as source and headline), the third dimension, ν , is unknown to all agents. We assume that agents update their beliefs about ν using Bayes' rule given beliefs about the underlying state θ and the observables (r, m) of the article.

Social Media Behavior. Time is discrete $t = 1, 2, \dots$. Upon receipt of the article, an agent i can take one of three actions $a_i \in \{\mathcal{S}, \mathcal{I}, \mathcal{D}\}$, as described below:

- (i) Share (\mathcal{S}): The agent decides to *share* the article and passes it onto others after her.
- (ii) Ignore (\mathcal{I}): The agent decides to *ignore* the article and does not engage with it.
- (iii) Dislike (\mathcal{D}): The agent decides to *dislike* the article, which means expressing disagreement or contempt for the content contained in it.

The three possible actions are depicted in Figure 1 using Twitter as a sample social media platform. A given user sees the article and decides how to respond to it. She can (i) share it (\mathcal{S}), which actively puts it on other social media news feeds; (ii) ignore it (\mathcal{I}), where the user simply scrolls past the article; or (iii) actively dislike it (\mathcal{D}), expressing derision for the content.

At time $t = 1$, we assume that some initial seed agent i^* begins with the article. If the article is shared by agent i , it is passed to all $j \in \mathcal{N}_i$. In contrast, following ignore or dislike, the article does not propagate past agent i .

Payoffs. Let us define shares after i as $S_i = |\{j \in \mathcal{N}_i : a_j = \mathcal{S}\}|$ and dislikes after i as $D_i = |\{j \in \mathcal{N}_i : a_j = \mathcal{D}\}|$. Agent i 's utility can then be written as

$$U_i = \begin{cases} 0, & \text{if } a_i = \mathcal{I} \\ \tilde{u}\mathbf{1}_{\nu=\mathcal{M}} - \tilde{c}, & \text{if } a_i = \mathcal{D} \\ u\mathbf{1}_{\nu=\mathcal{T}} - c\mathbf{1}_{\nu=\mathcal{M}} + \kappa S_i - dD_i, & \text{if } a_i = \mathcal{S} \end{cases} \quad (1)$$

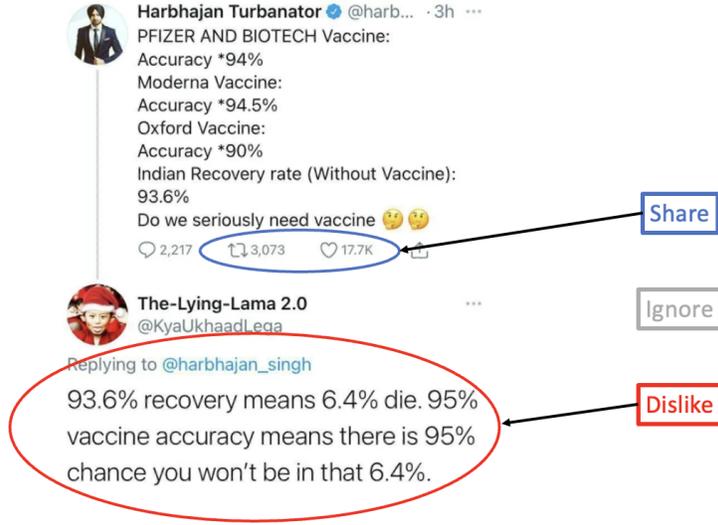


Figure 1. Sample Tweet.

where $\mathbf{1}$ is the indicator function (equal to 1 if true and 0 otherwise). Here, \tilde{u} , \tilde{c} , u , c , κ and d are strictly positive parameters, which we discuss below.

- (i) We normalize payoffs following ignore, \mathcal{I} , to $U_i = 0$.
- (ii) Payoffs from dislike, \mathcal{D} , depend on whether the article contains misinformation. We assume, in particular, that disliking has a cost of $\tilde{c} > 0$, regardless of whether the article is truthful (because of, say, the effort required to actively call out misinformation). In addition, disliking an article containing misinformation has a benefit of $\tilde{u} > \tilde{c}$, because individuals like calling out misleading articles. This formulation implies that disliking is never preferred to ignoring for an article that is truthful with probability 1, and is always preferred to ignoring for an article that contains misinformation with probability 1.
- (iii) Following a decision to share, \mathcal{S} , an agent receives utility from two sources. First, agents receive utility from sharing truthful content, but incur a cost from sharing misinformation. This explains the first component of utility following \mathcal{S} , $U_i^{(1)} = u\mathbf{1}_{\nu=\mathcal{T}} - c\mathbf{1}_{\nu=\mathcal{M}}$. Second, agents enjoy positive feedback from their peers (such as likes, or in our setting re-shares), but are negatively affected by dislikes. This is captured by the second component of utility,

$$U_i^{(2)} = \kappa S_i - dD_i$$

In this formulation, the parameter κ captures the importance of “popularity” for the agent’s sharing decision, while $-d$ represents the extent to which she cares about negative reactions.⁹

⁹Reputational loss associated with dislikes can be micro-founded straightforwardly by assuming that there are two types, normal and careless, and normal agents would like to distinguish themselves from careless ones who always share. We do not pursue this path, since imposing exogenous reputation costs simplifies the notation and analysis, with no major cost in terms of insights.

Information Structure and Solution Concept. Agents are not aware of when the article was first introduced onto social media, the prior sharing process, or the structure of the social network (though the link matrix \mathbf{P} is common knowledge).¹⁰ Moreover, while any agent i knows the distribution $\{H_i\}_{i=1}^N$ of beliefs in the population, she does not know any agent j 's belief (ideology) b_j . We focus on Bayesian-Nash equilibria, and refer to these as “equilibria” for short.

To eliminate trivial and unrealistic equilibria, we assume that the sensationalism of an article is upper bounded by $\bar{\kappa} = (c\bar{c} - u(\tilde{u} - \bar{c})) / (\tilde{u}N)$. This assumption guarantees that there is never an equilibrium where *every* agent *always* shares *all* articles. It also eliminates equilibria where agents may share and dislike, but never ignore.

Discussion — The basic assumptions introduced above are consistent with salient patterns of behavior and information structure in social media. First, as documented in studies such as [Pennycook et al. \(2021\)](#), users want to share content they believe to be truthful and not contain misinformation. Second, while users derive value from peer encouragement and re-shares on social media ([Eckles et al. \(2016\)](#)), they also suffer reputational costs when they get called out for sharing misinformation (see, for example, evidence from Facebook in [Altay et al. \(2020\)](#)). Finally, social media users often engage in criticisms of available content and inform others about misinformation (see, for example, [Kim et al. \(2020\)](#) for evidence in the context of 2018 midterm elections).

3 Equilibria in General Networks

In this section, we characterize the structure of equilibria for any sharing network structure \mathbf{P} and provide various comparative statics. Without loss of generality (and ease of exposition), we fix the article’s message as $m = R$ for the remainder of the paper.¹¹

3.1 Cutoff Strategies and Strategic Complementarities

When agent i receives an article with reliability r and message $m = R$, she updates her (*ex post*) belief, π_i , that the article is truthful according to Bayes’ rule:

$$\pi_i = \frac{(pb_i + (1 - p)(1 - b_i))\phi(r)}{(qb_i + (1 - q)(1 - b_i))(1 - \phi(r)) + (pb_i + (1 - p)(1 - b_i))\phi(r)}. \quad (2)$$

Clearly, π_i is increasing in b_i since an agent is more likely to believe in an article’s veracity when its message agrees with her prior. Moreover, π_i is increasing in r , as the agent updates more on the basis of more reliable articles.

¹⁰More formally, we assume that all agents have uniform priors over these objects. In this sense, agents do not know the exact interactions and sharing patterns outside their neighborhood, which is consistent with the evidence in [Breza et al. \(2018\)](#). That being said, because the equilibrium sharing process is Markovian, this assumption can be relaxed by replacing \mathbf{P} with the adjacency matrix.

¹¹To see this is without loss of generality, observe that the analysis applies identically with an $m = L$ message but with complementary priors $b'_i = 1 - b_i$.

We can also see that the payoff to sharing (\mathcal{S}) increases in π_i , since the first component of utility, $U_i^{(1)}$, is increasing in π_i (as the individual would like to share truthful articles), while $U_i^{(2)}$ is independent of π_i . With a similar reasoning, the payoff to disliking (\mathcal{D}) is decreasing in π_i , whereas the payoff to ignoring (\mathcal{I}) is independent of π_i . This monotone behavior of payoffs will lead to best-response decision rules for agents that take the form of cutoff strategies, as we explain next.

We say that agent i employs a *cutoff strategy* if there exists $b_i^*(r)$ and $b_i^{**}(r)$ such that agent i chooses \mathcal{S} when $b_i > b_i^{**}(r)$, chooses \mathcal{I} when $b_i^*(r) < b_i < b_i^{**}(r)$, and chooses \mathcal{D} when $b_i < b_i^*(r)$. Cutoff strategies in our context imply that agents who strongly agree with an article tend to share it, agents who strongly disagree with it tend to choose dislike, and those with intermediate beliefs typically ignore the article.

We will see in the next theorem that all equilibria will be in cutoff strategies. This means, in particular, that an equilibrium can be summarized by cutoff vectors $(\mathbf{b}^*, \mathbf{b}^{**}) = (b_1^*, b_1^{**}, \dots, b_N^*, b_N^{**})$. Furthermore, these cutoffs $b_i^*(r)$ and $b_i^{**}(r)$ will both be decreasing in r , so that as reliability increases, an article becomes more likely to be shared and less likely to be disliked.

We can also note that our social media game exhibits *strategic complementarities*. To see this, observe that when others share more, meaning that b_i^{**} (weakly) decreases for all i , this tends to encourage more sharing for every agent j , because the second component of utility, $U_j^{(2)}$, increases, raising the overall utility of sharing. Similarly, when others reduce their likelihood of disliking, meaning that now b_i^* (weakly) decreases for all i , this reduces the likely cost of sharing misinformation by mistake, also raising $U_j^{(2)}$. Strategic complementarities capture an important dimension of social media interactions — utility feedback from others' behavior tends to encourage agents to cohere with those behaviors.

Equilibrium Structure. The next theorem shows that an equilibrium always exists and is in cutoff strategies. At the same time, strategic complementarities ensure that there is a well-defined structure to the set of equilibria. To make this more concrete, we say that an equilibrium $(\mathbf{b}^*, \mathbf{b}^{**})$ has *uniformly more sharing* than other $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})$ if $\mathbf{b}^* \preceq \hat{\mathbf{b}}^*$ and $\mathbf{b}^{**} \preceq \hat{\mathbf{b}}^{**}$ (where \preceq is the component-wise order). This, in particular, means that all the thresholds for each agent is (weakly) lower in the former equilibrium (recall that lower thresholds mean more sharing).

Theorem 1.

- (i) *There exists a Bayesian-Nash equilibrium;*
- (ii) *All equilibria are in cutoff strategies;*
- (iii) *The set of cutoffs $(\mathbf{b}^*, \mathbf{b}^{**})$ forms a lattice, and thus there exists a least-sharing and most-sharing equilibrium.*

The structure of equilibria characterized in Theorem 1 facilitates our analysis, enabling us to focus on two sets of thresholds, $(\mathbf{b}^*, \mathbf{b}^{**})$, which are themselves monotonic in the reliability of the article in

question. Because of strategic complementarities, there can be multiple equilibria: when others are choosing to share an article with middling reliability, this further encourages sharing because one's own post will circulate more, increasing the utility from sharing. Conversely, if the same article with middling reliability is not shared by others, the payoff to sharing is reduced, while the cost of being found out to have circulated misinformation remains constant. This then discourages sharing.

Theorem 1 also shows that, despite this multiplicity, there are two focal equilibria on which we can concentrate: the equilibrium with the smallest vector of cutoffs (*most-sharing equilibrium*) and the equilibrium with the largest vector of cutoffs (*least-sharing equilibrium*).

Finally, the theorem's characterization provides an explicit measure of the amount of sharing. Recall that agent i 's prior belief b_i is drawn *ex ante* from the distribution H_i . Hence, in an equilibrium with cutoff b_i^{**} for agent i , the *ex ante* likelihood that this agent will share is $1 - H_i(b_i^{**})$. Thus, clearly, the most sharing equilibrium, which has the smallest equilibrium b_i^{**} for all i , has the highest likelihood that any agent i will share the article in question.

3.2 Comparative Statics (for Fixed P)

In this subsection we provide comparative statics for the most and least sharing equilibria as we change the parameters of the social media game, but holding the sharing network \mathbf{P} fixed. We discuss comparative statics with respect to the network in the next section. Throughout this subsection, we say that there is *uniformly more sharing* if sharing increases in both the most sharing and the least sharing equilibria.

Reputability and Quantity of Misinformation. The next proposition shows that greater reliability of articles and less misinformation lead to uniformly more sharing. We define greater reliability as a first-order stochastically dominant shift of the reliability distribution of articles from F to F' (where $F' \succeq_{FOSD} F$). We refer to less misinformation as a shift of the function ϕ to some $\phi' \geq \phi$.

Proposition 1. *More reliability of articles and/or less misinformation lead to uniformly more sharing.*

These results are intuitive. Holding constant the degree of misinformation, a shift towards more reliable articles reduces the relevant vector of thresholds ($\mathbf{b}^*, \mathbf{b}^{**}$), and this means that each agent becomes more likely to share. Holding constant the distributional reliability, less misinformation reduces the cost associated with sharing, triggering more aggressive sharing by all agents.

The predictions summarized in Proposition 1 are consistent with the existing evidence. Pennycook and Rand (2019) document that social media users recognize low-reliability content sources (e.g., Breitbart or Infowars), and this content is not typically shared by attentive social media users, regardless of partisanship (see Pennycook and Rand (2018)). This proposition also clarifies that viral spread of misinformation is not a mechanical effect in our model: if anything, less reliable articles that are more likely to contain misinformation are less likely to become viral. We will see that other aspects

of social media interactions—sensational content and the nature of the sharing network—are often responsible for viral spread of misinformation.

Sensational Content and Reputational Concerns. Let us next consider two key parameters of users’ utility. The first is κ , which we assume captures how *sensational* the article is. Higher κ implies that agents receive greater value from future shares, because these shares are associated with others paying more attention or perhaps being entertained more by the relevant posts. This greater utility is independent of the content’s veracity. As an example, headlines such as “Trump Pays \$2 Million to 8 Charities for Misuse of Foundation” are more sensational than headlines like: “108,000 people will get medical debt relief after Stacey Abrams’ PAC gifts \$1.34 million.”¹² We think of κ varying at the level of articles, so that some articles will be more sensational than others.

The second parameter is d , standing in for the importance of *reputational concerns*. Higher d means that dislikes are more damaging, which corresponds to the agent being more concerned about receiving many dislikes. We think of d as varying at the level of communities (certain communities of users, for example, academics, may have more reputational concerns).

Proposition 2. *More sensational articles and/or lower reputational concerns of users lead to uniformly more sharing.*

The results in this proposition are both intuitive and immediate. Nevertheless, they have important applications. First, the results of this proposition cohere with the patterns documented in [Duffy et al. \(2020\)](#), suggesting that social media participants often share a story that is “too good not to share”, and do so even when they realize it is also “too good to be true”. Second, the link between reputational concerns and misinformation is consistent with the evidence that in settings where reputation matters misinformation is less likely ([Altay et al. \(2020\)](#)), and when such reputational concerns are missing, even calling out individuals sharing misinformation is fairly ineffective ([Mosleh et al. \(2021a\)](#)).

Finally, this proposition provides a possible pathway for low-reliability content (often containing misinformation) to become viral. [Vosoughi et al. \(2018\)](#) argued that misinformation spreads farther, faster, deeper and more broadly than truthful news on social media. This evidence was criticized by [Grinberg et al. \(2019\)](#) who showed that, once the effects of sensational news items is controlled for, misinformation does not spread faster than truthful content. Propositions 1 and 2 provide a rationalization of these patterns. All else equal, misinformation does not spread faster than truthful content as in [Grinberg et al. \(2019\)](#). However, because, as observed in [Molina et al. \(2021\)](#) and [Kozyreva et al. \(2020\)](#), sensational content is often low-reliability, these two propositions together imply that misinformation may be more likely to become viral. Whether this happens or not depends on the boost from sensationalism. When this is limited, low-reliability articles containing misinformation spread less because of concerns of users that others will call them out for sharing this content. But when this sensationalism boost is high, misinformation can become viral. Additionally, the

¹²Sources are, respectively, <https://www.nytimes.com/2019/12/10/nyregion/trump-foundation-lawsuit-attorney-general.html> and <https://www.npr.org/2021/10/27/1049675704/stacey-abrams-medical-debt-relief-pac-donation>.

strategic complementarity in sharing decisions implies that sufficiently sensational misinformation can become viral very rapidly, because once an individual thinks others are going to share this sensational item, she becomes much more likely to share herself, even if she has doubts about its veracity.

4 Island Networks and the Implications of Homophily

In this section, we will conduct comparative statics with respect to the sharing network \mathbf{P} . Throughout this section, we take the sharing network as given, and then return to how it is shaped by the algorithms of social media platforms in Section 5.

The focus on comparative statics with respect to the sharing network necessitates two modifications from our analysis so far. *First*, we will restrict attention to *island networks* (or equivalently, the stochastic block model), which are lower-dimensional than general networks we have allowed so far. Namely, in an island network, agents are partitioned into k blocks of size N_1, N_2, \dots, N_k , called *islands* each with some constant (but not necessarily equal) share of the population N . Each agent i has a type $\ell_i \in \{1, \dots, k\}$ corresponding to which block (or “island”) she is in. Link probabilities are then given as:

$$p_{ij} = \begin{cases} p_s, & \text{if } \ell_i = \ell_j \\ p_d, & \text{if } \ell_i \neq \ell_j \end{cases}$$

where $p_s \geq p_d$. Without loss, we assume each of the islands is weakly connected.

Furthermore, we assume the prior distribution for agents on the same island ℓ is the same, and is denoted by H_ℓ . We also assume that islands are ranked according to their belief distributions. In particular, each island ℓ has distribution H_ℓ with support on $[b^{(\ell)}, b^{(\ell+1)}]$, where $1 \geq b^{(1)} > b^{(2)} > \dots > b^{(k)} > b^{(k+1)} \geq 0$.¹³ This implies that lower-indexed islands have stronger right-wing beliefs.

An important advantage of island networks, in addition to their lower-dimensional representation, is that, combined with this ranking assumption, they enable us to model the degree of *homophily*—the extent to which an individual interacts with others that have common characteristics as herself. Common characteristics for us are those that are relevant for prior beliefs, and therefore, by construction, individuals have more in common with those on the same island as themselves. As a result, homophily will be higher when most links are within islands and links between islands are sparse (high p_s and low p_d).¹⁴

Second, we shift our focus from *uniformly more sharing* to the weaker notion of *content virality* (which we define as more sharing, but not necessarily uniformly so). The reason for this is that, while

¹³This assumption is adopted for simplicity. Our results generalize if we instead assume that these distributions are ranked in terms of first-order stochastic dominance: $H_1 \succeq_{FOSD} H_2 \succeq_{FOSD} \dots \succeq_{FOSD} H_k$. However, this generalization requires considerably more formalism and notation, motivating our focus on disjoint supports.

¹⁴The homophilic structure and greater congruence of beliefs within islands are consistent with the evidence presented in [Bakshy et al. \(2015\)](#): “friend networks” on Facebook are ideologically segregated, with the median share of friends from the opposing ideology around only 20%. [Mosleh et al. \(2021b\)](#) provides evidence of similar homophily on Twitter.

our game continues to exhibit strategic complementarities, changes in the sharing network do not induce monotone shifts, and thus the most and least sharing equilibria do not necessarily shift in the same direction. We will see, nevertheless, that content virality will change in well behaved ways.

Formally, we define content virality as follows. We suppose that an article is seeded at some agent i^* at $t = 1$. We then define \mathbf{S}_{i^*} as the (random) proportion of the population that shares when agent i^* is the seed agent in the most-sharing equilibrium σ . We say σ_1 has more *content virality* than σ_2 if $\max_{i^*} \mathbb{E}_{\sigma_1}[\mathbf{S}_{i^*}] \geq \max_{i^*} \mathbb{E}_{\sigma_2}[\mathbf{S}_{i^*}]$. In words, our notion of content virality compares the spread of an article provided that it starts from the seed that is most favorable to its ultimate circulation. The reason we start from the most favorable seed is that, as we will see in the next section, social media platforms have an incentive to implement sharing algorithms that place articles in such favorable seeds. For future reference, we also note that content virality is also the same as expected overall engagement with an article, conditional on favorable seeding.

Finally, we define homophily more formally. We say that an island network with (p_s, p_d) has *more homophily* than an island network with (p'_s, p'_d) and the same expected degree of each agent, but where $p_s > p'_s$ and $p_d < p'_d$.¹⁵ From Theorem 1, We know that the equilibrium is in cutoff strategies of the form $(\mathbf{b}^*, \mathbf{b}^{**}) \equiv (b_1^*, b_1^{**}, \dots, b_N^*, b_N^{**})$. However, because there is symmetry within islands, equilibria now take a simpler, “semi-symmetric” form as shown in the next lemma.

Lemma 1. *All equilibria are semi-symmetric: for every equilibrium, there exist $\{(b_\ell^*, b_\ell^{**})\}_{\ell=1}^k$ such that $b_i^* = b_{\ell_i}^*$ and $b_i^{**} = b_{\ell_i}^{**}$ for all agents i in island ℓ .*

The simplification established in Lemma 1 will allow us to work with a lower dimensional cutoff vector (just two cutoffs for each island).

4.1 Comparative Statics: Homophily

The next theorem, characterizing the effects of homophily on the spread of misinformation, is our first main result:

Theorem 2. *There exist $0 < \underline{r} < \bar{r} < 1$ such that:*

- (a) *If $r < \underline{r}$, an increase in homophily increases the virality of content.*
- (b) *If $r > \bar{r}$, a decrease in homophily increases the virality of content.*

Theorem 2 shows how low-reliability content can spread virally in networks with high homophily. Intuitively, when content comes from a low-reliability source, only agents who (strongly) agree with the article’s message share it. However, as homophily increases, users know that they will mostly share with other like-minded people, who will also be inclined to share this content. This creates an echo chamber-like environment: the likelihood of being called out for spreading misinformation is

¹⁵Because network density raises connectivity and can directly increase virality, we hold network density fixed in order to isolate the effects from homophily.

now lower, making users “less disciplined” or more likely to share lower-reliability content. Strategic complementarities then extend these incentives throughout the network. In this way, homophily leads to the viral spread of low-reliability articles that likely contain misinformation.

However, Theorem 2 shows that homophily can have non-monotone effects. This is because greater homophily also keeps an article circulating among the same group of like-minded users and reduces the likelihood that it will reach other communities. Theorem 2(a) establishes that the first effect of homophily, working through incentives to share low-reliability content, is more powerful than the second, “circulation effect”, when we focus on particularly low-reliable content (with $r < \underline{r}$). This implies, in particular, that homophily’s impact is to increase the virality of especially low-quality content, which is of course relevant for public policy (as we discuss in Section 6).

The results in Theorem 2 are in line with recent evidence highlighting the importance of echo chambers for the spread of misinformation. Törnberg (2018) and Vicario et al. (2016), among others, show that homophily in sharing behavior propagates ideologically-congruent ideas, with little incentive to question the veracity of this information, while Quattrociocchi et al. (2016) document how echo chambers on Facebook fuel conspiracy theories and the popularization of incorrect scientific ideas, for example, on vaccines. Levy (2021) provides evidence that “filter bubbles” generated by Facebook’s algorithms are an important source of propagation of misinformation.

4.2 Comparative Statics: Divisive Content and Belief Polarization

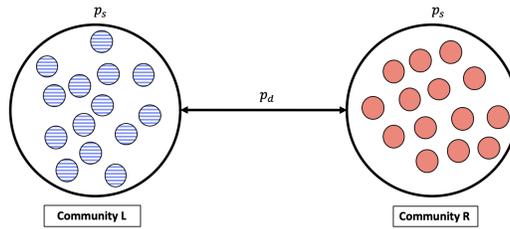


Figure 2. Two-Island Model.

For simplicity, for the next result, we focus on the case of just two islands, a left-wing and a right-wing one with prior distributions H_L and H_R , respectively, as pictured in Figure 2. Moreover, we suppose that there is disjoint support of prior beliefs across communities. Formally, we assume H_R has support on $[\underline{b}_R, \bar{b}_R]$ and H_L has support on $[\underline{b}_L, \bar{b}_L]$, with $\bar{b}_L < 1/2 < \underline{b}_R$.

Divisiveness of Content. We say content with parameters (p', q') is *more divisive* than content with parameters (p, q) if $p \geq p'$ and $q \leq q'$. Divisive content has a message that is more tethered to the true state θ when it is truthful (and more likely to argue against θ if it is misinformation). In our case, we think of state θ as related to political ideology. Therefore, non-political content, such as wedding photos or cat videos, has little divisiveness relative to more political ones, such as “Obama Signs Executive Order Banning The Pledge of Allegiance in Schools Nationwide” (Fourney et al. (2017)).

Note also that, from equation (2), prior beliefs about θ matter more for updating and the assessment of an article's veracity when content is more divisive.

Polarization of Beliefs. We say H_2 is *more polarized* than H_1 if it satisfies the following single crossing property: $H_2^{-1}(\alpha) - H_1^{-1}(\alpha)$ is a nondecreasing function in α , crossing zero at $\alpha^* = 1/2$ with $H_1(1/2) = H_2(1/2) = 1/2$.

An increase in polarization results in a “stretching” of the belief distribution around the most moderate user (i.e., $b = 1/2$) while preserving an equal distribution of left-wing and right-wing agents (meaning that $H(1/2) = H(1/2) = 1/2$, which is applied in the island model to the average distribution of beliefs, $H = \frac{1}{N} \sum_{\ell=1}^N N_{\ell} H_{\ell}$). The available evidence indicates that the US public has become more polarized (see [Pew Research Center \(2014\)](#) and [Abramowitz \(2010\)](#)), and an important question of debate has been whether this polarization has fueled the spread of misinformation on social media.

Virality of Content. The next theorem studies how political divisiveness and polarization impact social media behavior and the spread of misinformation. To highlight our main insights, we focus on the same two regimes identified in Theorem 2 (low reliability with $r < \underline{r}$ and high reliability with $r > \bar{r}$) and in addition distinguish between high and low levels of homophily ($p_s/p_d > \bar{p}$ and $p_s/p_d < \underline{p}$):

Proposition 3. *There exists $r^* \in (0, 1)$ and $p^* \in (0, 1)$ such that:*

- (a) If $r < r^*$ and $p_s/p_d > p^*$, then an *increase* in divisiveness or an increase in polarization leads to greater content virality.
- (b) If $r > r^*$ and $p_s/p_d < p^*$, then a *decrease* in divisiveness or a decrease in polarization leads to greater content virality.

Proposition 3 builds on and is complementary to Theorem 2. When the content in question has high reliability ($r > \bar{r}$), more divisive content or greater polarization do not accentuate virality (and in fact may reduce content virality when homophily is limited, because in a well-connected, non-homophilic network, controversial articles will solicit a wide range of reactions, disciplining those tempted to share misinformation). In contrast, when the article in question has low reliability and there is significant homophily, there are again echo chamber-like effects. More divisive content generates more divergent behavior from individuals with different ideologies, and greater polarization means there are sharper differences in terms of these ideologies. As a result, echo chambers matter especially for divisive content and in the presence of polarization. The intuition of this result builds on Theorem 2: greater polarization can trigger more negative reactions, especially for divisive content and especially when there more polarized opinions. But when there is a high degree of homophily, individuals are fairly certain that their shared content will be viewed by like-minded users, and hence become less disciplined in sharing low-reliability, divisive content. The presence of strategic complementarities once again amplifies this effect, as users recognize that others in their community will tend to share divisive content, and this makes them even more willing to share.

It is notable that Proposition 3, like Theorem 2, implies that greater divisiveness and polarization increase the virality of especially low-reliability content, which is most likely to contain misinformation. These two results together thus imply that echo chambers, greater political polarization and divisive content all exacerbate the circulation of misinformation on social media.

5 Platform Design and Filter Bubbles

We now turn to our second main result: how platform behavior affects misinformation. Consider a collection of social media users with beliefs distributed according to a distribution H . The platform can identify communities of users according to prior ideological beliefs, for example, based on content previously shared or affiliations with ideological groups. In particular, each user is binned into one of k communities, with each community ℓ having a belief distribution H_ℓ with support over $[b_\ell, b_{\ell+1}]$, and where $1 \geq b^{(1)} > b^{(2)} > \dots > b^{(k)} > b^{(k+1)} \geq 0$ (with at least one left-wing and one right-wing community). The size of these bins depends on the platform’s microtargeting technology at identifying users’ ideological beliefs (see, for example, Papakyriakopoulos et al. (2018)). Formally, we let $\varepsilon \equiv \max_\ell (b_{\ell+1} - b_\ell)$, with the interpretation that lower values of ε correspond to better platform technology for identifying ideology.

The platform’s objective is to maximize user engagement, which is equivalent to maximizing content virality (see the definition of content virality in the previous section).¹⁶ The platform does not directly care about whether the content users are engaging with is truthful or contains misinformation.

The platform chooses how content is shared across users. That is, for each article, the platform not only picks the seed agent at $t = 1$ to whom it recommends this article, but also chooses the sharing network—the matrix of link probabilities \mathbf{P} .¹⁷ The platform’s choice of \mathbf{P} can be interpreted as its “algorithm” to determine how users are exposed to content circulating in the social media site.

5.1 Optimality of Island Networks and Filter Bubbles

We remind the reader that the island networks of Section 4 are parameterized by three components: the within-island link probability (p_s), the across-island link probability (p_d), and the number of islands (k). As special cases, we have (i) an island model that has *maximal homophily*, where $p_s > 0$ but $p_d = 0$ (and thus there is extreme ideological segregation on the network); and (ii) an island model

¹⁶This objective is rooted in the fact that social media sites, like Facebook, primarily rely on advertising revenue, which becomes more valuable as users increase their activity on the site. For example, 85% of Facebook’s total revenue in 2011 was from advertising, and from 2017-2019, around 98% was (see Andrews (2012) and <https://www.nasdaq.com/articles/what-facebooks-revenue-breakdown-2019-03-28-0>).

Strictly speaking, we are modeling social media platform objectives before the more recent public backlash over misinformation. If the platform faces potential penalties from public backlash or regulators for spreading misinformation, its objective function will change, as we explore in greater detail in Section 6.

¹⁷Facebook’s algorithms may induce different sharing networks depending on features of the article, such as whether it contains cat videos, wedding photos, or political content. See <https://about.fb.com/news/2021/01/how-does-news-feed-predict-what-you-want-to-see/>.

with *maximal connectivity*, where $p_s = p_d$ (and there is minimal homophily and no segregation by ideology). Our second main result is:

Theorem 3. *There exists $\bar{\varepsilon} > 0$ such that if $\varepsilon < \bar{\varepsilon}$, the platform's profit-maximizing sharing network is determined by a reliability threshold $r_P \in (0, 1)$ such that:*

(i) *If $r < r_P$, the platform's profit-maximizing sharing network has maximal homophily.*

(ii) *If $r > r_P$, the platform's profit-maximizing sharing network has maximal connectivity.*

Part (i) of the theorem shows that when articles are mostly unreliable — and likely to contain misinformation — the platform creates an extreme filter by designing its algorithms to achieve a sharing network with the greatest homophily. In contrast, part (ii) demonstrates that when articles have higher reliability, the platform refrains from introducing algorithmic homophily. This result highlights an important channel by which misinformation spreads: it is precisely when articles are likely to contain misinformation that the platform seeks to maximize engagement by creating (endogenous) echo chambers, or filter bubbles, where these articles spread virally within like-minded communities. Put differently, with low reliability content, neither the platform nor the users are disciplined about sharing misinformation, and so these news items spread virtually uninhibited.

It is worth noting that this theorem builds on but also significantly strengthens Theorem 2. In Theorem 2, the effects of homophily are non-monotone and are ambiguous when an article is neither very low reliability nor very high reliability ($r \in (\underline{r}, \bar{r})$). In contrast, Theorem 3 gives a sharp characterization of the platform's algorithm: when $r > r_P$, the platform goes for maximal connectivity, and when $r < r_P$, it chooses maximal homophily, and in this case, misinformation spreads virally precisely because of the echo chambers that the platform has manufactured. In both cases, the island structure of the network considered in Section 4 arises endogenously as the profit-maximizing sharing network for the platform.

In addition, Theorem 3 shows that when a platform can shape the network topology through its recommendation algorithm, the echo chamber effect that arises from Theorem 2(a) is exactly what fuels misinformation (whereas in Theorem 2(b), echo chambers were harmless). This corroborates and extends much of the previous literature mentioned earlier on the interaction between recommender systems, echo chambers, and misinformation.

Finally, it is notable that the platform designs its algorithms to generate filter bubbles within ideologically-congruent groups precisely when the relevant content is low-reliability and likely to contain misinformation. This observation generalizes to cases where the platform has less precise microtargeting technology, albeit in a less sharp way than the result presented in Theorem 3.¹⁸

¹⁸In this case, the platform still induces echo chamber-like environments to generate viral spread of low-reliability content, even if its choice does not take the form of an island network. For instance, if there are only three communities (with broad ideology spectra), a misinformation right-wing article can still spread virally via the usage of filter bubbles. However, the platform may now prefer a sharing network that lies outside the class of island networks considered in Section 4. Specifically, user engagement may be maximized if the left-wing community is completely disconnected from all other communities,

Remark — In Theorem 3, we assume the platform can select any network \mathbf{P} it desires through its recommendation algorithm. This is without loss of generality. If we assume the social network originally begins as an arbitrary island network in Section 4, and the platform can hide and amplify content across different links, the same result readily follows.

5.2 Comparative Statics on r_P : Divisiveness and Polarization

In Theorem 3, the threshold r_P fully summarizes the extent to which misinformation will spread virally on social media. We next perform comparative statics for this threshold to understand the conditions under which the platform will create a filter bubble and propagate misinformation.

Proposition 4. *The optimal reliability threshold r_P increases as message divisiveness and/or belief polarization increases.*

Proposition 4 mimics the conclusions of Proposition 3(a). As divisiveness or polarization increases, content is consumed more aggressively within echo chambers and scrutinized more aggressively outside of them. In these settings, filter bubbles become more advantageous to the platform, especially when the relevant content has low reliability. This is because communities with more extreme beliefs now feel more strongly about news in general, rarely second-guess politically-congruent news, but often doubt and dislike counter-attitudinal news. As a result, low-reliability content would have been quickly disliked and stopped without echo chambers, which would have discouraged its sharing even by agents inclined to believe it. In contrast, inside of the platform’s filter bubbles, such content continues to spread virally.

Proposition 4 also provides a possible (albeit of course speculative) interpretation for why accelerating political polarization and identity politics in the last two decades may have come with more aggressive filter bubble algorithms from social media sites (Apprich et al. (2018)). As the recent documentary *The Social Dilemma* puts it: “The way to think about it is as 2.5 billion Truman Shows. Each person has their own reality with their own facts. Over time you have the false sense that everyone agrees with you because everyone in your news feed sounds just like you.” Tellingly in this context, while Facebook cracked down on misinformation prior to the 2020 election in part due to political pressure, its algorithms have resumed promotion of misinformation in November and December of 2020: “...the measures [Facebook] could take to limit harmful content on the platform might also limit its growth: In experiments Facebook conducted last month, posts users regarded in surveys as ‘bad for the world’ tended to have a greater reach — and algorithmic changes that reduced the visibility of those posts also reduced users’ engagement with the platform...”¹⁹

but the moderate (middle) community has sparse connections to the right-wing community. This facilitates a strong echo chamber within the right-wing community but also allows the article to spread to the more moderate community (while receiving no discipline from the left-wing community).

¹⁹See *Vanity Fair*:

<https://www.vanityfair.com/news/2020/12/with-the-election-over-facebook-gets-back-to-spreading-misinformation> and also <https://www.technologyreview.com/2021/03/11/1020600/facebook-responsible-ai-misinformation/>.

6 Regulation

Our analysis so far raises the natural question of what types of regulations might counter the viral spread of misinformation and platform choices leading to excessive ideological homophily. We now briefly discuss four distinct types of regulations that have been discussed in this context: (1) regulations that force platforms to reveal articles’ *provenance*; (2) *censorship* or tagging of misinformation; (3) *performance targets* that require the platform to keep misinformation below a given threshold; (4) *network regulations*, restricting the extent of ideological homophily or segregation introduced by platform algorithms intended to maximize engagement.

We consider the effects of these policies both on a fixed sharing network and when the platform can optimally choose the sharing network given public policy. For simplicity, we suppose the regulator’s objective is to decrease the virality of articles containing misinformation on the platform.²⁰ We will say the policy has an *ambiguous* effect when there exists a sharing network (and set of conditions) where the virality of misinformation *increases* following the regulatory policy.

6.1 Provenance

We consider a policy that requires the platform to reveal the original context or *provenance* of a piece of content. For example, provenance may point the user to a peer-reviewed medical study or the full discourse from which a quote was pulled. Such a policy allows users to verify (or “fact-check”) social media content easily and quickly.

We model a provenance policy by allowing some users to fact-check the article before making their sharing and dislike decisions. In particular, we assume revealing provenance allows each agent to check the veracity of the content instantly with probability $\rho > 0$, with higher values of ρ indicating more effective provenance policies.²¹ For simplicity, we assume fact-checks perfectly reveal whether the article has veracity $\nu = \mathcal{T}$ or $\nu = \mathcal{M}$. The user can then make a more informed (but still strategic) decision about whether to share the article or take one of the other two actions.

Proposition 5. *There exists $\rho^* < 1$ such that:*

- (a) *If $\rho > \rho^*$, for any fixed sharing network, the provenance policy reduces the virality of misinformation.*
- (b) *If $\rho < \rho^*$, the provenance policy has an ambiguous effect on the virality of misinformation.*
- (c) *If $\rho > \rho^*$, the provenance policy reduces the virality of misinformation even when the platform optimally chooses the sharing network conditional on the policy.*

²⁰Formally, the regulator minimizes the virality of content conditional on the article containing misinformation; in mathematical notation, the regulator minimizes $\mathbb{E}_{\max_{i^*}}[\mathbf{S}_{i^*} \mid \nu = \mathcal{M}]$.

²¹In practice, certain demographic groups, such as users over 65 years old, appear more likely to accept misinformation without fact-checking, perhaps because of poor media interpretation skills (see Grinberg et al. (2019) and Guess et al. (2019)). Thus, provenance policies which provide a less clear pathway to fact-checking may lead certain social media users to fact-check and others to not.

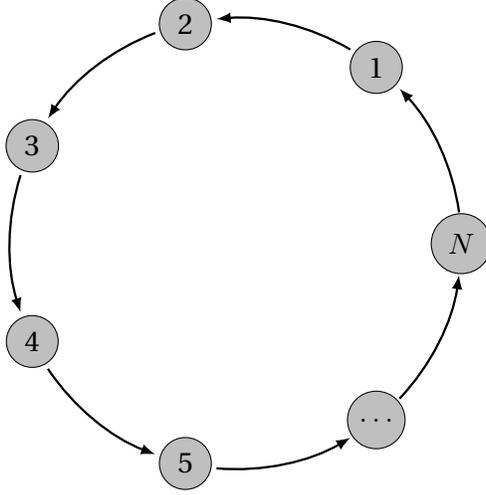


Figure 3. Ring network of Example 1.

We see that an effective provenance policy, which helps users identify misinformation with high likelihood, successfully reduces the circulation of this content. The reason is that when it is common knowledge that a piece of content contains misinformation, incentives toward sharing are uniformly diminished across all users on the platform. Strategic complementarities then further reduce sharing. However, less effective provenance policies ($\rho < \rho^*$) are *not* always beneficial and can sometimes backfire (part (b)), as illustrated in the next example.

Example 1. Without loss of generality, we fix an article that has message $m = R$ and some intermediate reliability r , so that ex-ante the article may or may not contain misinformation. We also condition on an article that turns out to contain misinformation. Moreover, we assume the regulator has only a semi-effective provenance policy ($\rho = 1/4$) and the sharing network is a ring, as pictured in Figure 3. Virality is therefore measured as the expected distance between the seed agent and the first agent who elects not to share.

For simplicity, we assume that strategic complementarities play a small role (i.e., $\kappa \approx 0$ and $d \approx 0$), so sharing decisions are based mostly on an agent’s assessment of the article’s veracity. Moreover, we assume all prior beliefs are drawn from the same distribution H . First, we consider the virality of content with no provenance policy. Per Theorem 1, there exists some b^{**} such that agents with $b_i > b^{**}$ share and agents with $b_i < b^{**}$ do not share. As a result, each agent shares the article with probability $\zeta \equiv 1 - H(b^{**})$, and the expected spread of the article is given by $\sum_{i=1}^N \zeta^i$. Note also that this probability applies regardless of whether the article contains misinformation, which is not observable to the agent without fact-checking.

Next, consider the introduction of a provenance policy with $\rho = 1/4$, so the likelihood that an agent fact-checks the article before choosing her action is $1/4$ and the likelihood that an agent takes her action without fact-checking is $3/4$. Because, by assumption, the article contains misinformation, a fact-check will set agent i ’s belief of veracity to zero, ensuring that she does not share this article. Therefore, the overall probability that the first agent will share the article is $(1 - \rho)\zeta = 3\zeta/4$. While

a fact-check immediately kills the article (which occurs with probability $1/4$), no fact-check leads to sharing with probability ζ (also occurring with probability $3/4$).

Given this provenance policy, we now consider agent 2, who next receives the article. Agent 2 does not know whether the article has already been fact-checked by the first agent, but knows the provenance policy is in place. Given the article was passed from the first agent to the second agent, she assesses that there is a higher probability that the article is truthful relative to the case with no provenance policy. This is because the act of agent 1 sharing the article with agent 2 may have been the result of a fact-check that came back positive (article is truthful), and so observing this action leads her to a higher posterior belief the article is truthful. Consequently, the prior cutoff for sharing will be below b^{**} (the cutoff without provenance), and if she does not fact-check herself, her likelihood of sharing will be some $\zeta_2 > \zeta \equiv \zeta_1$. As before, the probability that agent 2 will share is $(1 - \rho)\zeta_2 = 3\zeta_2/4$. Extending this reasoning, the probability agent i will share is $(1 - \rho)\zeta_i$, where $\zeta_i > \zeta_{i-1}$ for all i .

Thus, to determine whether the provenance policy is effective, one needs to compare $\sum_{i=1}^N \zeta^i$ and $\sum_{i=1}^N (1 - \rho)^i \left[\prod_{j=1}^i \zeta_j \right]$. In general, this comparison is ambiguous, but in the proof of Proposition 5(b), we show that the latter sum is greater because ζ_j grows rapidly, reflecting the fact that later users presume others before them have fact-checked and do not make independent judgments based on the reliability of the content. In summary, while the provenance policy effectively stops misinformation when identified quickly, it opens a potential gateway for misinformation to be much more widespread. ■

The observation in Example 1 is related to the literature on informational cascades and herding (for example, see Banerjee (1992), Bikhchandani et al. (2021), and Chen and Papanastasiou (2021)). When provenance policies are enacted, agents tend to follow the sharing decisions of those before them instead of making independent inferences about content veracity before sharing. This opens the door for share cascades based on what is sometimes called a “bystander effect”—sharing based on the belief that the content was already vetted by others.

Finally, part (c) of the proposition shows that when $\rho > \rho^*$, the provenance policy reduces misinformation even when the platform responds to this policy by changing the sharing network. This ensures that the platform cannot game the policy by choosing a different network structure that circumvents provenance regulation.

6.2 Censorship

Next, we consider a policy where the platform can censor misinformation that appears on the platform.²² Formally, we model this as the platform being able to adopt a censorship policy that removes $\delta \in (0, 1)$ fraction of the content containing misinformation (with each piece of

²²Alternately, it can “tag” the article in question as disputed by outside sources, with analogous implications. See, for example, Facebook’s policies leading up to the 2020 election on labeling suspected misinformation: <https://about.fb.com/news/2019/10/update-on-election-integrity-efforts/>.

misinformation removed with probability δ).²³ For simplicity, we assume this content is removed at $t = 0$ before it is even observed by any of the users.

Proposition 6. *There exists $\delta^* < 1$ such that:*

- (a) *If $\delta > \delta^*$, for any fixed sharing network, the censorship policy reduces the virality of misinformation.*
- (b) *If $\delta < \delta^*$, the censorship policy has an ambiguous effect on the virality of misinformation.*
- (c) *If $\delta > \delta^*$, the censorship policy reduces the virality of misinformation even when the platform optimally chooses the sharing network conditional on the policy.*

This result mimics Proposition 5: when censorship can adequately remove most misinformation, it is an effective policy both in given sharing networks (part (a)) and in sharing networks selected by the platform in response to the policy (part (c)). However, when the censorship policy fails to identify some of the misinformation, there is an “implied truth” effect that kicks in for uncensored articles (as empirically identified in Pennycook et al. (2020a)). Bayesian users believe, correctly, that articles are more likely to be truthful when there is censorship of misinformation. Under certain conditions, this policy can ultimately lead to the greater spread of misinformation (part (b)).

While provenance and censorship policies have similar features in that both are effective when done properly (and backfire when not), they possess different comparative advantages. For instance, provenance policies tend to perform better when most of the content on the platform is low reliability. In this case, multiple (independent) fact-checks by social media users typically remove this content more efficiently than just a single regulator can. In contrast, because users might share sensational content that contains misinformation even when there is an effective provenance policy, censorship may be more attractive relative to provenance policies in the presence of misinformation carried by sensationalist articles.

6.3 Performance Targets

Another possible regulation, which has recently been proposed by social media platforms (see Facebook (2020)), is to set *performance targets* for their content. As monitoring and removing misinformation is imperfect, a performance target allows the platform some leeway to have “bad” content on their site while still requiring a degree of accountability.

While in Section 6.2 we considered a policy where the regulator was responsible for removing content, a performance target transfers the burden of content removal to the platform itself. In this subsection, we consider how the platform would remove content when misinformation must be kept below a given threshold. We assume that the platform can discard content (by not recommending it to any user), and in doing so, forfeits any potential engagement this content may have with the users.

²³As in Section 6.1, we think of δ as being a technology parameter related to how effective the regulator is in identifying misinformation. The assumption that the regulator may make type-I errors but not type-II errors (truthful articles are never misidentified, but misinformation is identified with some probability less than one) is adopted for simplicity.

As in our previous discussion, we assume that sensational articles tend to be more associated with misinformation (see Section 3.2). Formally, we suppose the sensationalism of every article is drawn as $\kappa \sim G(\cdot|\nu)$ conditional on the article’s veracity ν , where $G(\cdot|\nu = \mathcal{M}) \succeq_{FOSD} G(\cdot|\nu = \mathcal{T})$. Given this, we suppose the regulator sets a performance target of λ , which requires the proportion of misinformation shares (to total shares) on the platform to fall below λ .²⁴ The regulator enforces this performance target by auditing the platform and sampling the content to ensure it meets the required standard. Formally, we suppose the regulator has an auditing technology $\alpha \in (0, 1)$ which captures the probability of catching a platform who violates its performance target. For simplicity, we assume that if the platform is caught violating the performance target, it incurs a cost due to regulatory fines.

Our next result establishes how stricter performance targets affect the spread of misinformation:

Proposition 7. *There exists an auditing technology cutoff $\alpha^* < 1$ such that:*

- (a) *If $\alpha > \alpha^*$, for a fixed sharing network, a stricter performance target (lower λ) reduces the virality of misinformation;*
- (b) *If $\alpha < \alpha^*$, a stricter performance target has an ambiguous effect on the virality of misinformation;*
- (c) *If $\alpha > \alpha^*$, a stricter performance target reduces the virality of misinformation even when the platform optimally chooses the sharing network conditional on the policy.*

This result captures similar ideas to those in Propositions 5 and 6. When performance targets are well-enforced through auditing, reducing the target (making it stricter) necessarily curbs the spread of misinformation. When the platform is held more accountable for misinformation shares, it is enticed to reduce the misinformation in circulation, even if it happens to be more sensational. Lower targets therefore align regulator and platform incentives to remove the least reliable content.

However, with lower auditing, the incentives of the regulator and platform become more divergent. With lax performance targets, the platform still has some incentive to comply: it can remove low-reliability, low-sensational content and meet the performance standard without much of a hit to engagement. However, as the performance target continues to tighten, the platform must start removing low-reliability sensationalist content. This comes at the cost of potentially removing some of the most viral content possible. When the auditing probability is not too high, the platform may prefer to violate the performance target, in which case more strict targets backfire.

6.4 Network Regulations

As we saw in Theorem 3, when unregulated, the platform (endogenously) chooses the island model of Section 4 with parameters (p_s, p_d) . Here, we consider limits on the ideological homophily induced by

²⁴This metric for performance comes from Facebook’s own statements on platform standards: “Regulators could say that internet platforms must publish annual data on the ‘prevalence’ of content that violates their policies, and that companies must make reasonable efforts to ensure that the prevalence of violating content remains below some standard threshold” (from Facebook (2020)), with the definition of prevalence being: “We care most about how often content that violates our standards is actually seen relative to the total amount of times any content is seen on Facebook” (from <https://about.fb.com/news/2019/05/measuring-prevalence/>).

the platform’s algorithm. In particular, we suppose the regulator can choose a homophily standard, based on the ratio between within-island links to across-island links.

Proposition 8. *There exists $\gamma < \infty$ such that for any $p^* \geq \gamma$, if the regulator imposes $p_s/p_d \leq p^*$ then (i) the platform chooses the island model with $p_s/p_d \leq p^*$; and (ii) the virality of misinformation is reduced.*

The regulator can thus achieve its objective by moderating the homophily that the platform can impose on the sharing network. In particular, algorithmic choices that create extreme homophily (as demonstrated in Theorem 3(a)), can be curbed by imposing a standard that requires content to be shared across ideological lines. This is closely related to the “ideological segregation standard” discussed in Sunstein (2018), which aims to restrict the extent to which content is curated specifically to the ideology or interests of the group of users. Such standards ensure that echo chambers are broken and users of differing ideology interact more frequently, limiting the spread of misinformation.

We note that the regulation in Proposition 8 is not always binding for the platform. In the case of highly reliable articles (as in Theorem 3(b)), the platform maximizes engagement by implementing a maximally-connected sharing network, so the homophily constraint is moot. However, when the article is less reliable, the regulation will bind and the platform will be forced to maintain a minimum level of connectivity.

7 Conclusion

This paper has developed a simple model of the spread of misinformation over social media platforms. A group of Bayesian agents with heterogeneous priors receive and share news items (articles) according to a stochastic sharing network, determined by the social media platform. Articles may be truthful and informative about an underlying state, or may contain misinformation, making them (weakly) anti-correlated with the underlying state. Upon receiving an article, an agent can decide to share it with others, ignore it, or actively call out another agent for propagating misinformation (“dislike”). Misinformation spreads when agents share articles expecting positive social media feedback and little negative reactions.

Though simple and parsimonious, the model encapsulates several rich strategic interactions. Agents receive utility from sharing truthful articles and not misinformation, but also enjoy peer engagement with shared content. The former force pushes agents toward sharing pro-attitudinal content, whereas the latter force pushes agents toward pleasing those in their sharing network. The ideological congruence between an agent and those in her sharing network, which we capture with the notion of homophily, is critical for sharing decisions. Because individuals are more likely to dissent against articles that disagree with their prior beliefs, an agent will be more cautious in sharing articles that disagree with the views of those in her sharing network.

We provide several comparative static results. Some of those are very intuitive, though still useful for interpreting a range of results in the emerging empirical literature on social media and

misinformation. For example, we find that while misinformation typically spreads less than truthful content (holding all else constant), more sensational content tends to be shared more. Moreover, when misinformation is correlated with sensationalism, the rapid spread misinformation can be problematic.

Of particular interest are the comparative statics with respect to homophily—measuring whether agents tend to share articles with others with similar or dissimilar beliefs. We show that when there is a highly-reliable article, an increase in homophily reduces the virality of content. Because this article is unlikely to contain misinformation, it is of broad appeal to a wide range of social media users, independent of ideology. An increase in homophily then reduces the extent to which this article can spread throughout the sharing network. The implications of homophily for low-reliability articles are very different, however: in a well-connected network, such articles will be disliked and stopped by users who disagree with their message, and anticipating this behavior and the loss of reputation they can suffer from spreading misinformation, even those who agree with their message would not share them widely. In contrast, high homophily creates echo chambers, where users share low-reliability messages aligned with their beliefs, because they understand that there are few negative reputational consequences from doing so. Misinformation contained in low-reliability articles can then spread virally in these echo chambers.

Our framework enables a tractable study of platform incentives in designing algorithms that determine who shares with whom. To do this, we assume that the platform aims to maximize user engagement (which is a good approximation to the objectives of major social media platforms such as Facebook or Twitter). Our main result is a striking one. When an article is highly reliable, the platform chooses a sharing network with minimal homophily to maximize the spread and appeal of the content throughout the user community. In contrast to this case, when the relevant articles have lower reliability, the platform chooses a network with maximal homophily and recommends articles to users with aligned beliefs. These articles then spread rapidly in the “filter bubble” the platform’s algorithms have created—because now ideologically like-minded individuals know that they are unlikely to be caught sharing misinformation in their extreme echo chambers.

We also study regulations aimed at minimizing the spread of misinformation. Revealing the provenance of a news item (for example, providing full context for a quote) can be useful, because this additional information allows users to more easily fact-check the content for veracity. However, we also show that this type of regulation may backfire because of a “bystander effect” whereby future users assume content has been previously fact-checked and put more faith in it. Similarly, tagging or censoring articles from less reliable sources may make agents more confident in those that remain on the site untagged—thus creating a Bayesian version of the “false sense of security”. Finally, we show that regulating platform algorithms that shape sharing network properties (in particular, homophily) can be effective.

Our framework was purposefully chosen to be simple and several generalizations would be interesting to consider in future work. Most importantly, our assumption that agents are Bayesian

rational should be viewed as a useful, albeit not fully realistic, benchmark. In our setting, it brought out certain new strategic forces—highlighting how social media actions exhibit strategic complementarity and how the degree of homophily alters agents' strategic behavior. Nevertheless, misinformation may be particularly damaging when agents have cognitive limitations and are only boundedly rational. Incorporating such considerations is one of the most important directions for future research. Interesting questions that emerge in this case relate to whether the platform, in addition to designing algorithms that create filter bubbles, may also choose strategies that exploit the cognitive limitations of users.

Other theoretical generalizations that might be interesting to consider include extensions to repeated interactions with incomplete information, which would enable agents to also update their beliefs about the ideological position of other agents in their sharing network. Fully endogenizing reputational concerns, taking into account the network position of agents, would be another interesting direction for future research. In this case, the existing reputational capital of an agent will determine how likely she is to risk sharing misinformation. We can also use this extended setup with repeated interactions to study how agents update their initial political views. When there is limited misinformation, agents will gradually learn the true state. In contrast, when there is a significant probability of misinformation, agents will be uncertain about how to interpret articles that disagree with their priors and this may place an upper bound on the speed and possibility of learning (see [Acemoglu et al. \(2016\)](#)).

Despite its simplicity, our model makes several new empirical predictions, most notably related to the non-monotonic effects of homophily and polarization and to platform incentives and algorithmic decisions. Investigating these predictions empirically as well as generating new stylized facts about patterns of these information cascades on social media, is another important and exciting area for future research.

References

- Abramowitz, Alan I. (2010), *The Disappearing Center: Engaged Citizens, Polarization, and American Democracy*. Yale University Press.
- Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz (2016), “Fragility of asymptotic agreement under bayesian learning.” *Theoretical Economics*, 11, 187–225.
- Acemoglu, Daron, Giacomo Como, Fabio Fagnani, and Asuman Ozdaglar (2013), “Opinion Fluctuations and Disagreement in Social Networks.” *Mathematics of Operations Research*, 38, 1–27. Publisher: INFORMS.
- Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi (2010), “Spread of (mis)information in social networks.” *Games and Economic Behavior*, 70, 194–227.
- Allcott, Hunt and Matthew Gentzkow (2017), “Social media and fake news in the 2016 election.” *Journal of economic perspectives*, 31, 211–36.
- Allen, Jennifer, Baird Howland, Markus Mobius, David Rothschild, and Duncan J. Watts (2020), “Evaluating the fake news problem at the scale of the information ecosystem.” *Science Advances*, 6, eaay3539. Publisher: American Association for the Advancement of Science Section: Research Article.
- 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*, EC '19, 555–556, Association for Computing Machinery, New York, NY, USA.
- Altay, Sacha, Anne-Sophie Hacquin, and Hugo Mercier (2020), “Why do so few people share fake news? It hurts their reputation.” *New Media & Society*, 1461444820969893. Publisher: SAGE Publications.
- Andrews, Lori (2012), “Facebook is using you.” *The New York Times*, 4.
- Apprich, Clemens, Florian Cramer, Wendy Hui Kyong Chun, and Hito Steyerl (2018), *Pattern Discrimination*. meson press. Accepted: 2020-05-04T14:50:47Z.
- Aral, Sinan and Paramveer S. Dhillon (2018), “Social influence maximization under empirical influence models.” *Nature Human Behaviour*, 2, 375–382. Number: 6 Publisher: Nature Publishing Group.
- Bakshy, Eytan, Solomon Messing, and Lada A. Adamic (2015), “Exposure to ideologically diverse news and opinion on Facebook.” *Science*, 348, 1130–1132. Publisher: American Association for the Advancement of Science Section: Report.
- Banerjee, Abhijit V. (1992), “A Simple Model of Herd Behavior*.” *The Quarterly Journal of Economics*, 107, 797–817.

- Bikhchandani, Sushil, David Hirshleifer, Omer Tamuz, and Ivo Welch (2021), “Information Cascades and Social Learning.” Working Paper 28887, National Bureau of Economic Research. Series: Working Paper Series.
- Breza, Emily, Arun G. Chandrasekhar, and Alireza Tahbaz-Salehi (2018), “Seeing the forest for the trees? An investigation of network knowledge.” *arXiv:1802.08194 [physics, stat]*. ArXiv: 1802.08194.
- Buchanan, Tom (2020), “Why do people spread false information online? The effects of message and viewer characteristics on self-reported likelihood of sharing social media disinformation.” *PLOS ONE*, 15, e0239666. Publisher: Public Library of Science.
- Budak, Ceren, Divyakant Agrawal, and Amr El Abbadi (2011), “Limiting the spread of misinformation in social networks.” In *Proceedings of the 20th international conference on World wide web, WWW '11*, 665–674, Association for Computing Machinery, New York, NY, USA.
- Candogan, Ozan and Kimon Drakopoulos (2017), “Optimal Signaling of Content Accuracy: Engagement vs. Misinformation.” SSRN Scholarly Paper ID 3051275, Social Science Research Network, Rochester, NY.
- Centola, Damon (2010), “The spread of behavior in an online social network experiment.” *science*, 329, 1194–1197. Publisher: American Association for the Advancement of Science.
- Centola, Damon and Michael Macy (2007), “Complex Contagions and the Weakness of Long Ties.” *American Journal of Sociology*, 113, 702–734. Publisher: The University of Chicago Press.
- Chen, Li and Yiangos Papanastasiou (2021), “Seeding the Herd: Pricing and Welfare Effects of Social Learning Manipulation.” *Management Science*, 67, 6734–6750. Publisher: INFORMS.
- Duffy, Andrew, Edson Tandoc, and Rich Ling (2020), “Too good to be true, too good not to share: the social utility of fake news.” *Information, Communication & Society*, 23, 1965–1979. Publisher: Routledge _eprint: <https://doi.org/10.1080/1369118X.2019.1623904>.
- Eckles, Dean, René F. Kizilcec, and Eytan Bakshy (2016), “Estimating peer effects in networks with peer encouragement designs.” *Proceedings of the National Academy of Sciences*, 113, 7316–7322. Publisher: National Academy of Sciences Section: Colloquium Paper.
- Egelhofer, Jana Laura and Sophie Lecheler (2019), “Fake news as a two-dimensional phenomenon: a framework and research agenda.” *Annals of the International Communication Association*, 43, 97–116. Publisher: Routledge _eprint: <https://doi.org/10.1080/23808985.2019.1602782>.
- Facebook (2020), “Charting a Way Forward on Online Content Regulation.”
- Fiorina, Morris P., Samuel A. Abrams, and Jeremy C. Pope (2008), “Polarization in the American Public: Misconceptions and Misreadings.” *The Journal of Politics*, 70, 556–560. Publisher: The University of Chicago Press.

- Fourney, Adam, Miklos Z. Racz, Gireeja Ranade, Markus Mobius, and Eric Horvitz (2017), “Geographic and Temporal Trends in Fake News Consumption During the 2016 US Presidential Election.” In *CIKM*, volume 17, 6–10.
- Gentzkow, Matthew and Jesse M. Shapiro (2006), “Media Bias and Reputation.” *Journal of Political Economy*, 114, 280–316. Publisher: The University of Chicago Press.
- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer (2019), “Fake news on Twitter during the 2016 U.S. presidential election.” *Science*, 363, 374–378. Publisher: American Association for the Advancement of Science.
- Guess, Andrew, Jonathan Nagler, and Joshua Tucker (2019), “Less than you think: Prevalence and predictors of fake news dissemination on Facebook.” *Science Advances*, 5, eaau4586. Publisher: American Association for the Advancement of Science Section: Research Article.
- Guess, Andrew, Brendan Nyhan, and Jason Reifler (2018), “Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign.” *European Research Council*, 9, 4.
- Hsu, Chin-Chia, Amir Ajorlou, and Ali Jadbabaie (2020), “News Sharing, Persuasion, and Spread of Misinformation on Social Networks.” SSRN Scholarly Paper ID 3391585, Social Science Research Network, Rochester, NY.
- Kamenica, Emir (2019), “Bayesian Persuasion and Information Design.” *Annual Review of Economics*, 11, 249–272. eprint: <https://doi.org/10.1146/annurev-economics-080218-025739>.
- Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian Persuasion.” *American Economic Review*, 101, 2590–2615.
- Keppo, Jussi, Michael Jong Kim, and Xinyuan Zhang (2019), “Learning Manipulation Through Information Dissemination.” *Available at SSRN 3465030*.
- Kim, Dam Hee, S. Mo Jones-Jang, and Kate Kenski (2020), “Why Do People Share Political Information on Social Media?” *Digital Journalism*, 0, 1–18. Publisher: Routledge eprint: <https://doi.org/10.1080/21670811.2020.1827966>.
- Kozyreva, Anastasia, Stephan Lewandowsky, and Ralph Hertwig (2020), “Citizens Versus the Internet: Confronting Digital Challenges With Cognitive Tools.” *Psychological Science in the Public Interest*, 21, 103–156.
- Lazer, David MJ, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, and David Rothschild (2018), “The science of fake news.” *Science*, 359, 1094–1096. Publisher: American Association for the Advancement of Science.

- Lee, Chei Sian, Long Ma, and Dion Hoe-Lian Goh (2011), “Why Do People Share News in Social Media?” In *Active Media Technology* (Ning Zhong, Vic Callaghan, Ali A. Ghorbani, and Bin Hu, eds.), Lecture Notes in Computer Science, 129–140, Springer, Berlin, Heidelberg.
- Levy, Ro’ee (2021), “Social Media, News Consumption, and Polarization: Evidence from a Field Experiment.” *American Economic Review*, 111, 831–870.
- Molina, Maria D., S. Shyam Sundar, Thai Le, and Dongwon Lee (2021), ““Fake News” Is Not Simply False Information: A Concept Explication and Taxonomy of Online Content.” *American Behavioral Scientist*, 65, 180–212. Publisher: SAGE Publications Inc.
- Mosleh, Mohsen, Cameron Martel, Dean Eckles, and David Rand (2021a), “Perverse Downstream Consequences of Debunking: Being Corrected by Another User for Posting False Political News Increases Subsequent Sharing of Low Quality, Partisan, and Toxic Content in a Twitter Field Experiment.” In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 182, 1–13, Association for Computing Machinery, New York, NY, USA.
- Mosleh, Mohsen, Cameron Martel, Dean Eckles, and David G. Rand (2021b), “Shared partisanship dramatically increases social tie formation in a twitter field experiment.” *Proceedings of the National Academy of Sciences*, 118.
- 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.
- Nguyen, Nam P., Guanhua Yan, My T. Thai, and Stephan Eidenbenz (2012), “Containment of misinformation spread in online social networks.” In *Proceedings of the 4th Annual ACM Web Science Conference*, WebSci ’12, 213–222, Association for Computing Machinery, New York, NY, USA.
- Papakyriakopoulos, Orestis, Simon Hegelich, Morteza Shahrezaye, and Juan Carlos Medina Serrano (2018), “Social media and microtargeting: Political data processing and the consequences for Germany.” *Big Data & Society*, 5. Publisher: SAGE Publications Ltd.
- Papanastasiou, Yiangos (2020), “Fake News Propagation and Detection: A Sequential Model.” *Management Science*, 66, 1826–1846. Publisher: INFORMS.
- Pennycook, Gordon, Adam Bear, Evan T. Collins, and David G. Rand (2020a), “The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings.” *Management Science*, 66, 4944–4957. Publisher: INFORMS.
- Pennycook, Gordon, Tyrone D. Cannon, and David G. Rand (2018), “Prior exposure increases perceived accuracy of fake news.” *Journal of Experimental Psychology. General*, 147, 1865–1880.

- 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. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7855 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Communication;Decision making;Human behaviour;Technology Subject_term.id: communication;decision-making;human-behaviour;technology.
- Pennycook, Gordon, Jonathon McPhetres, Yunhao Zhang, Jackson G. Lu, and David G. Rand (2020b), “Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for a Scalable Accuracy-Nudge Intervention.” *Psychological Science*, 31, 770–780. Publisher: SAGE Publications Inc.
- 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.” SSRN Scholarly Paper ID 3165567, Social Science Research Network, Rochester, NY.
- Pennycook, Gordon and David G. Rand (2019), “Fighting misinformation on social media using crowdsourced judgments of news source quality.” *Proceedings of the National Academy of Sciences*, 116, 2521–2526. Publisher: National Academy of Sciences Section: Social Sciences.
- Pew Research Center (2014), “Political Polarization in the American Public.”
- Prior, Markus (2013), “Media and Political Polarization.” *Annual Review of Political Science*, 16, 101–127. eprint: <https://doi.org/10.1146/annurev-polisci-100711-135242>.
- Quattrociocchi, Walter, Antonio Scala, and Cass R. Sunstein (2016), “Echo Chambers on Facebook.” SSRN Scholarly Paper ID 2795110, Social Science Research Network, Rochester, NY.
- Sunstein, Cass R. (2018), *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press. Google-Books-ID: oVBLDwAAQBAJ.
- Tarski, Alfred (1955), “A lattice-theoretical fixpoint theorem and its applications.” *Pacific Journal of Mathematics*, 5, 285–309. Publisher: Pacific Journal of Mathematics, A Non-profit Corporation.
- Taylor, Sean J. and Dean Eckles (2018), “Randomized Experiments to Detect and Estimate Social Influence in Networks.” In *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks* (Sune Lehmann and Yong-Yeol Ahn, eds.), Computational Social Sciences, 289–322, Springer International Publishing, Cham.
- Topkis, Donald M. (1998), *Supermodularity and Complementarity*, first edition edition. Princeton University Press, Princeton, N.J.
- Törnberg, Petter (2018), “Echo chambers and viral misinformation: Modeling fake news as complex contagion.” *PLOS ONE*, 13, e0203958. Publisher: Public Library of Science.

Vicario, Michela Del, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, and Walter Quattrociocchi (2016), “The spreading of misinformation online.” *Proceedings of the National Academy of Sciences*, 113, 554–559. Publisher: National Academy of Sciences Section: Physical Sciences.

Vosoughi, Soroush, Deb Roy, and Sinan Aral (2018), “The spread of true and false news online.” *Science*, 359, 1146–1151. Publisher: American Association for the Advancement of Science Section: Report.

A Proofs

A.1 Auxiliary Lemmas

We define a (mixed-strategy) strategy σ_i for agent i to be a map from priors b_i to elements of the simplex $\Delta(\{\mathcal{D}, \mathcal{I}, \mathcal{S}\})$. In others words, σ_i specifies for each ideological prior b_i of agent i the probability that she will play each of the three actions, \mathcal{D} , \mathcal{I} , and \mathcal{S} . We let σ_{-i} denote the (vector of) strategies of all agents other than agent i .

Lemma A.1. *Given any set of strategies σ_{-i} , agent i 's best response is a cutoff strategy with cutoffs (b_i^*, b_i^{**}) such that if $b_i < b_i^*$ agent i dislikes (\mathcal{D}), if $b_i^* < b_i < b_i^{**}$ agent i ignores (\mathcal{I}), and if $b_i > b_i^{**}$ agent i shares (\mathcal{S}).*

Proof of Lemma A.1. When agent i receives an article, she forms (ex-post) belief π_i about the article's veracity which depends only on the observables (r, m) . By Bayes' rule:

$$\pi_i \equiv \mathbb{P}[\nu = \mathcal{T} | r, m = R] = \frac{\mathbb{P}[m = R | r, \nu = \mathcal{T}] \mathbb{P}[\nu = \mathcal{T} | r]}{\mathbb{P}[m = R | r, \nu = \mathcal{M}] \mathbb{P}[\nu = \mathcal{M} | r] + \mathbb{P}[m = R | r, \nu = \mathcal{T}] \mathbb{P}[\nu = \mathcal{T} | r]} .$$

By the law of total probability, we have:

$$\begin{aligned} \mathbb{P}[m = R | r, \nu = \mathcal{T}] &= \mathbb{P}[m = R | \nu = \mathcal{T}] = \mathbb{P}[m = R | \theta = R, \nu = \mathcal{T}] \mathbb{P}[\theta = R] + \mathbb{P}[m = R | \theta = L, \nu = \mathcal{T}] \mathbb{P}[\theta = L] \\ &= pb_i + (1 - p)(1 - b_i) ; \end{aligned}$$

$$\begin{aligned} \mathbb{P}[m = R | r, \nu = \mathcal{M}] &= \mathbb{P}[m = R | \nu = \mathcal{M}] = \mathbb{P}[m = R | \theta = R, \nu = \mathcal{M}] \mathbb{P}[\theta = R] + \mathbb{P}[m = R | \theta = L, \nu = \mathcal{M}] \mathbb{P}[\theta = L] \\ &= qb_i + (1 - q)(1 - b_i) . \end{aligned}$$

Putting these together we obtain equation (2). Moreover, π_i is monotone in b_i since

$$\frac{\partial \pi_i}{\partial b_i} = \frac{(1 - \phi(r))\phi(r)(p - q)}{(1 - b_i + q(1 - \phi(r))(2b_i - 1) - p(\phi(r) - 2\phi(r)b_i))^2} > 0 .$$

Note that $U_i(\mathcal{I})$ and $U_i(\mathcal{D})$ is independent of σ_{-i} , and in particular \mathcal{D} is a better response to \mathcal{I} if and only if $\pi_i < (\tilde{u} - \tilde{c})/\tilde{u}$. Because π_i is monotone in b_i , this implies there exists some \tilde{b} where \mathcal{D} is a better response to \mathcal{I} if and only if $b_i < \tilde{b}$ (where $\tilde{b} = 1$ if disliking dominates ignoring and $\tilde{b} = 0$ if ignoring dominates disliking). Next, recall that the payoff to sharing is $U_i(\mathcal{S}) = U_i^{(1)} + U_i^{(2)}$, where $U_i^{(1)} = u\mathbf{1}_{\nu=\mathcal{T}} - c\mathbf{1}_{\nu=\mathcal{M}}$ and $U_i^{(2)} = \kappa S_i - dD_i$. Observe that, as before, $U_i^{(1)}$ is independent of σ_{-i} and has expected payoff $(u + c)\pi_i - c$, which is monotonically increasing in π_i . Moreover, $\mathbb{E}_{\mathbf{P}, \sigma_{-i}}[\kappa S_i - dD_i]$ does not depend on b_i . Because π_i is monotone in b_i , we see that $U_i(\mathcal{S})$ is increasing in b_i , $U_i(\mathcal{I})$ is constant in b_i (it is always zero), and $U_i(\mathcal{D})$ is decreasing in b_i (it is equal to $\tilde{u}(1 - \pi_i) - \tilde{c}$). This implies that either (i) ignoring dominates sharing, (ii) sharing dominates ignoring, or (iii) $U_i(\mathcal{S}) = 0$ for some prior b' :

- (i) If ignoring dominates sharing, we set $(b_i^*, b_i^{**}) = (\tilde{b}, 1)$.

- (ii) If sharing dominates ignoring, then either sharing dominates disliking (in which case set $(b_i^*, b_i^*) = (0, 0)$), disliking dominates sharing (in which case we set $(b_i^*, b_i^{**}) = (1, 1)$), or there exists some prior b'' where $U_i(\mathcal{S}) = U_i(\mathcal{D})$ (in which case set $(b_i^*, b_i^{**}) = (b'', b'')$).
- (iii) Otherwise, if $\tilde{b} < b'$, set $(b_i^*, b_i^{**}) = (\tilde{b}, b')$; however, if $\tilde{b} \geq b'$, then we set $(b_i^*, b_i^{**}) = (b', b')$.

This is of the cutoff form claimed in the lemma. ■

An immediate consequence of Lemma A.1 is that any Bayesian-Nash equilibrium must be in cutoff strategies for all agents. In particular, we can limit our attention to cutoff strategies (b_i^*, b_i^{**}) for every agent i , which can be represented as $(\mathbf{b}^*, \mathbf{b}^{**})$ in vector notation. This is a partially-ordered set using the component-wise order \succeq . Likewise, for any collection of cutoffs $\{(\mathbf{b}^{*,(1)}, \mathbf{b}^{**, (1)}), (\mathbf{b}^{*,(2)}, \mathbf{b}^{**, (2)}), \dots, \}$ in the cutoff space, there is a greatest lower bound given by the component-wise infimum and a least upper bound given by the component-wise supremum. Consequently, the cutoff space (which we denote by \mathbf{B}) forms a *complete lattice*.

Next, we define a map $\psi : \mathbf{B} \rightarrow \mathbf{B}$ that maps a set of cutoffs $(\mathbf{b}^*, \mathbf{b}^{**})$ to best-response cutoffs $(\mathbf{b}^{*,BR}, \mathbf{b}^{**,BR})$. We note this map is well-defined because (i) H is a continuous distribution, so we need not specify the strategies of agents precisely on the cutoffs, and (ii) by Lemma A.1, for any set of strategies σ_{-i} (including the cutoff strategies given by $(\mathbf{b}^*, \mathbf{b}^{**})$), all agents' best responses are in cutoff form.

Lemma A.2. *The map ψ preserves the component-wise order \succeq .*

Proof of Lemma A.2. Consider some $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**}) \succeq (\mathbf{b}^*, \mathbf{b}^{**})$. Fixing an article with observables (r, m) , $U_i(\mathcal{D})$, $U_i(\mathcal{I})$ and $U_i^{(1)}$ are independent of $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})$ and $(\mathbf{b}^*, \mathbf{b}^{**})$. However, for $U_i^{(2)}$ we have:

$$\begin{aligned}
\mathbb{E}_{\mathbf{P}, (\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})} [U_i^{(2)}] &= \mathbb{E}_{\mathbf{P}, (\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})} [\kappa S_i - dD_i] \\
&= \sum_{j=1}^N p_{ij} \left(\kappa \mathbb{P}_{(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})} [a_j = \mathcal{S}] - d \mathbb{P}_{(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})} [a_j = \mathcal{D}] \right) \\
&= \sum_{j=1}^N p_{ij} \left(\kappa \mathbb{P}_H [b_j > \hat{b}_j^{**}] - d \mathbb{P}_H [b_j < \hat{b}_j^*] \right) \\
&\leq \sum_{j=1}^N p_{ij} \left(\kappa \mathbb{P}_H [b_j > b_j^{**}] - d \mathbb{P}_H [b_j < b_j^*] \right) \\
&= \sum_{j=1}^N p_{ij} \left(\kappa \mathbb{P}_{(\mathbf{b}^*, \mathbf{b}^{**})} [a_j = \mathcal{S}] - d \mathbb{P}_{(\mathbf{b}^*, \mathbf{b}^{**})} [a_j = \mathcal{D}] \right) \\
&= \mathbb{E}_{\mathbf{P}, (\tilde{\mathbf{b}}^*, \tilde{\mathbf{b}}^{**})} [\kappa S_i - dD_i] = \mathbb{E}_{\mathbf{P}, (\mathbf{b}^*, \mathbf{b}^{**})} [U_i^{(2)}].
\end{aligned}$$

As a result, $U_i^{(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})}(\mathcal{S}) \leq U_i^{(\mathbf{b}^*, \mathbf{b}^{**})}(\mathcal{S})$. As in Lemma A.1, we define \tilde{b} as the prior where $U_i(\mathcal{D}) = 0$ if such a \tilde{b} exists, otherwise let $\tilde{b} = 0$ if ignoring dominates disliking and $\tilde{b} = 1$ if disliking dominates ignoring. Observe that \tilde{b} is the same for both $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})$ and $(\mathbf{b}^*, \mathbf{b}^{**})$. We have three cases for

the best-response cutoffs $(\mathbf{b}^{*,BR}, \mathbf{b}^{**,BR})$ given other agents' cutoffs $(\mathbf{b}^*, \mathbf{b}^{**})$ (which we compare to $(\hat{\mathbf{b}}^{*,BR}, \hat{\mathbf{b}}^{**,BR})$ given other agents' cutoffs $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})$):

- (i) Ignoring dominates sharing for agent i (for given cutoffs $(\mathbf{b}^*, \mathbf{b}^{**})$). Then by virtue of $U_i^{(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})}(\mathcal{S}) \leq U_i^{(\mathbf{b}^*, \mathbf{b}^{**})}(\mathcal{S})$, ignoring dominates sharing with $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})$ as well. Thus, $(b_i^{*,BR}, b_i^{**,BR}) = (\hat{b}_i^{*,BR}, \hat{b}_i^{**,BR}) = (\tilde{b}, 1)$.
- (ii) Sharing dominates ignoring for agent i (for given cutoffs $(\mathbf{b}^*, \mathbf{b}^{**})$). Then either sharing dominates disliking (in which case $(b_i^{*,BR}, b_i^{**,BR}) = (0, 0) \preceq (\hat{b}_i^{*,BR}, \hat{b}_i^{**,BR})$ trivially), or there exists some prior b'' where $U_i^{(\mathbf{b}^*, \mathbf{b}^{**})}(\mathcal{S}) = U_i(\mathcal{D})$ denoted by b'' and $(b_i^{*,BR}, b_i^{**,BR}) = (b'', b'')$. Moreover, because $U_i^{(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})}(\mathcal{S}) \leq U_i^{(\mathbf{b}^*, \mathbf{b}^{**})}(\mathcal{S}) = U_i(\mathcal{D})$ at prior b'' , for an agent with prior b'' , playing \mathcal{D} is a (weakly) better response than sharing when other agents play according to cutoffs $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})$. By monotonicity of $U_i(\mathcal{S})$ and $U_i(\mathcal{D})$ in prior b_i , this implies that $b_i^{**,BR} \leq \hat{b}_i^{**,BR}$. If ignoring is never a best response when $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})$, then $\hat{b}_i^{*,BR} = \hat{b}_i^{**,BR}$. Otherwise, $\hat{b}_i^{*,BR} = \tilde{b} \geq b_i^{**,BR} = b_i^{*,BR}$.
- (iii) $U_i^{(\mathbf{b}^*, \mathbf{b}^{**})}(\mathcal{S}) = 0$ for some prior b' for agent i . Then $U_i^{(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})}(\mathcal{S}) \leq U_i^{(\mathbf{b}^*, \mathbf{b}^{**})}(\mathcal{S}) = 0$ implies that for an agent with prior b' playing \mathcal{I} is a (weakly) better response than sharing when other agents play according to $(\hat{\mathbf{b}}^*, \hat{\mathbf{b}}^{**})$. By monotonicity of $U_i(\mathcal{S})$ in prior b_i , this implies that $b_i^{**,BR} \leq \hat{b}_i^{**,BR}$. If $\tilde{b} < b'$, then $b_i^{*,BR} = \hat{b}_i^{*,BR} = \tilde{b}$; otherwise, if $\tilde{b} \geq b'$, $b_i^{*,BR} = b_i^{**,BR} = b' \leq \hat{b}_i^{*,BR}$.

This establishes that $(\hat{\mathbf{b}}_i^{*,BR}, \hat{\mathbf{b}}_i^{**,BR}) \succeq (\mathbf{b}_i^{*,BR}, \mathbf{b}_i^{**,BR})$, so the order \succeq is preserved by ψ . ■

Lemma A.3. *An increase in polarization of beliefs can always be constructed via the following process: take every belief b_i and either (i) add some $\epsilon_i > 0$ to b_i if $b_i > 1/2$, or (ii) subtract some $\epsilon_i > 0$ to b_i if $b_i < 1/2$.*

Proof of Lemma A.3. Let H_2 be more polarized than H_1 . For part (i), note that $H_1(b_i^1) = \alpha > 1/2$, so by single-crossing at $H_1^{-1}(1/2) = H_2^{-1}(1/2)$, we know that $H_2^{-1}(\alpha) - H_1^{-1}(\alpha) > 0$. Thus, for some $b_i^2 > b_i^1$, we have $H_2^{-1}(\alpha) = b_i^2$, or in other words, $H_2(b_i^2) = \alpha$. Setting $\epsilon_i = b_i^2 - b_i^1 > 0$ in this fashion for all $b_i > 1/2$ accomplishes claim (i). For part (ii), note that $H_1(b_i^1) = \alpha < 1/2$, so by single-crossing at $H_1^{-1}(1/2) = H_2^{-1}(1/2)$, we know that $H_2^{-1}(\alpha) - H_1^{-1}(\alpha) < 0$. Thus, for some $b_i^2 < b_i^1$, we have $H_2^{-1}(\alpha) = b_i^2$, or in other words, $H_2(b_i^2) = \alpha$. Setting $\epsilon_i = b_i^1 - b_i^2 > 0$ in this fashion for all $b_i < 1/2$ accomplishes claim (ii). ■

Lemma A.4. *If $\kappa \leq \bar{\kappa} \equiv (c\bar{c} - u(\tilde{u} - \bar{c})) / (\tilde{u}N)$, then for any agent i :*

- (i) *If $b_i^* > 0$ and $b_i^{**} < 1$, then $b_i^{**} > b_i^*$;*
- (ii) *For all $\bar{b} < 1$, there exists $\tilde{r} > 0$ such that agent i plays \mathcal{D} in any equilibrium for an article with $r < \tilde{r}$ and on any sharing network \mathbf{P} , provided that $b_i < \bar{b}$.*

Proof of Lemma A.4. For part (i), by way of contradiction suppose that $b_i^* = b_i^{**}$. Then for an agent with prior b_i^* (and corresponding ex-post belief π_i^* that the article is truthful), it must be the case that:

$$\tilde{u}(1 - \pi_i) - \bar{c} = u\pi_i - c(1 - \pi_i) + \mathbb{E}[\kappa S_i - dD_i] \geq 0.$$

Re-arranging we see that

$$\pi_i = \frac{\tilde{u} - \tilde{c} + c - \mathbb{E}[\kappa S_i - dD_i]}{\tilde{u} + u + c}.$$

Plugging back into the payoff for action \mathcal{D} , we see that:

$$\begin{aligned} U_i(\mathcal{D}) &= \tilde{u} \left(\frac{u + \tilde{c} + \mathbb{E}[\kappa S_i - dD_i]}{\tilde{u} + u + c} \right) - \tilde{c} \\ &\leq \tilde{u} \left(\frac{u + \tilde{c} + \kappa N}{\tilde{u} + u + c} \right) - \tilde{c} \\ &< \tilde{u} \left(\frac{u + \tilde{c} + \bar{\kappa} N}{\tilde{u} + u + c} \right) - \tilde{c} \\ &\leq 0. \end{aligned}$$

By assumption, $U_i(\mathcal{S}) = U_i(\mathcal{D}) < 0$, but since $U_i(\mathcal{I}) = 0$, ignoring is the best response at prior b_i^* , which is a contradiction.

For part (ii), notice by equation (2), for a fixed $b < 1$, as $r \rightarrow 0$, $\pi_i \rightarrow 0$, and therefore:

$$\begin{aligned} U_i(\mathcal{S}) &= u\pi_i - c(1 - \pi_i) + \mathbb{E}[\kappa S_i - dD_i] \\ &< u\pi_i - c(1 - \pi_i) + \bar{\kappa} N \\ &\leq u\pi_i - c(1 - \pi_i) + \frac{c}{N} N \\ &\stackrel{r \rightarrow 0}{=} -c + c = 0. \end{aligned}$$

where the last inequality follows from the observation that:

$$\bar{\kappa} \equiv \frac{c\tilde{c} - u(\tilde{u} - \tilde{c})}{\tilde{u}N} < \frac{c\tilde{c}}{\tilde{u}N} < \frac{c}{N},$$

because $\tilde{u} > \tilde{c}$. Thus, as $r \rightarrow 0$, ignoring is a better response than sharing. But note that $U_i(\mathcal{D}) = \tilde{u}(1 - \pi_i) - \tilde{c} \stackrel{r \rightarrow 0}{=} \tilde{u} - \tilde{c} > 0$, so as $r \rightarrow 0$, disliking is a better response than ignoring. As a result, disliking is a best response for any fixed $b < 1$ as $r \rightarrow 0$. The claim in (ii) thus follows from continuity of equation (2).

A.2 Proofs from Section 3

Proof of Theorem 1. Claim (ii) follows directly from Lemma A.1 and establishes that the Bayesian-Nash equilibria are exactly the fixed points of the map ψ . Clearly the cutoff space \mathbf{B} is convex and compact (it is defined by $[0, 1]^{2N}$). To see that ψ is continuous, notice that for $\psi : (\mathbf{b}^*, \mathbf{b}^{**}) \mapsto (\mathbf{b}^{*,BR}, \mathbf{b}^{**,BR})$, $\mathbb{E}_{\mathbf{P},(\mathbf{b}^*, \mathbf{b}^{**})}[U_i^{(2)}]$ is continuous because H is continuous (and $U_i(\mathcal{D})$, $U_i(\mathcal{I})$, and $U_i^{(1)}$ do not depend on $(\mathbf{b}^*, \mathbf{b}^{**})$). Moreover, by the same reasoning as in Lemma A.2, $U_i^{\mathbf{P},(\mathbf{b}^*, \mathbf{b}^{**})}(\mathcal{S})$ and $U_i^{\mathbf{P},(\mathbf{b}^*, \mathbf{b}^{**})}(\mathcal{S}) - U_i(\mathcal{D})$ are monotone and continuous. Because these expressions are continuous in $(\mathbf{b}^*, \mathbf{b}^{**})$, the corresponding best-response cutoffs, $(\mathbf{b}^{*,BR}, \mathbf{b}^{**,BR})$ are also continuous in $(\mathbf{b}^*, \mathbf{b}^{**})$. By

Brouwer's fixed-point theorem, there exists a Bayesian-Nash equilibrium, proving (i).

Finally, noting that the cutoff space \mathbf{B} is a complete lattice and ψ preserves the component-wise order \succeq (by Lemma A.2), Tarski's fixed-point theorem establishes that the set of equilibrium cutoffs forms a lattice (see Tarski (1955)). By definition of a lattice order, there exists a least-sharing equilibrium (largest \mathbf{b}^{**}) and a most-sharing equilibrium (smallest \mathbf{b}^{**}). ■

Proof of Proposition 1. Recall that π_i is given by equation (2) and provides the (ex-post) belief of the article's veracity conditional on observables (r, m) . Also observe that:

$$\frac{\partial \pi_i}{\partial r} = \frac{(1 - b_i + p(2b_i - 1))(1 - b_i + q(2b_i - 1))}{(1 - b_i + q(1 - \phi(r))(2b_i - 1) - p(\phi(r) - 2\phi(r)b_i))^2} \phi'(r).$$

Because $\phi'(r) > 0$, it is clear that when $b_i > 1/2$, $\partial \pi_i / \partial r > 0$. When $b_i < 1/2$, $1 - b_i + p(2b_i - 1)$ is minimized when $p = 1$, in which case it is equal to $b_i \geq 0$ (and with this inequality strict whenever $p < 1$). Similarly, when $b_i < 1/2$, $1 - b_i + q(2b_i - 1)$ is minimized when $q = 1/2$, in which case it is equal to $1/2 > 0$. Thus, $\partial \pi_i / \partial r > 0$ for all b_i .

Similarly, when $\phi' \geq \phi$, a given reliability score r with misinformation structure ϕ' can be translated into a higher reliability score $r' \geq r$ but with misinformation structure ϕ (because both ϕ, ϕ' are monotonically increasing). As a consequence, a decrease in misinformation is isomorphic to greater reliability of the articles. It is thus sufficient to prove the latter leads to uniformly more sharing in both the least and the most sharing equilibria.

Note that the social media game is supermodular and has increasing differences in reputability. To see this, note that for all $r' \geq r$:

$$\begin{aligned} [U_i(\mathcal{S}, r') - U_i(\mathcal{I}, r')] - [U_i(\mathcal{S}, r) - U_i(\mathcal{I}, r)] &= U_i(\mathcal{S}, r') - U_i(\mathcal{S}, r) \\ &= U_i^{(1)}(r') - U_i^{(1)}(r) \\ &= (u + c)(\pi_i(r') - \pi_i(r)) \geq 0, \end{aligned}$$

via the above observation that $\frac{\partial \pi_i}{\partial r} > 0$. Similarly, for all $r' \geq r$:

$$\begin{aligned} [U_i(\mathcal{S}, r') - U_i(\mathcal{D}, r')] - [U_i(\mathcal{S}, r) - U_i(\mathcal{D}, r)] &= [U_i(\mathcal{S}, r') - U_i(\mathcal{S}, r)] + [U_i(\mathcal{D}, r') - U_i(\mathcal{D}, r)] \\ &= (u + c)(\pi_i(r') - \pi_i(r)) + \tilde{u}(\pi_i(r') - \pi_i(r)) \geq 0, \end{aligned}$$

via the same observation. Finally, for all $r' \geq r$:

$$\begin{aligned} [U_i(\mathcal{I}, r') - U_i(\mathcal{D}, r')] - [U_i(\mathcal{I}, r) - U_i(\mathcal{D}, r)] &= U_i(\mathcal{D}, r') - U_i(\mathcal{D}, r) \\ &= \tilde{u}(\pi_i(r') - \pi_i(r)) \geq 0, \end{aligned}$$

once again. Thus, via Topkis's monotone comparative statics theorem (see Topkis (1998)), there is uniformly more sharing. ■

Proof of Proposition 2. The social media game is supermodular and has increasing differences in sensationalism and (the negative of) reputational concerns. To see this, note that for all $\kappa' \geq \kappa$ and $d' \leq d$:

$$\begin{aligned} [U_i(\mathcal{S}, \kappa', d') - U_i(\mathcal{I}, \kappa', d')] - [U_i(\mathcal{S}, \kappa, d) - U_i(\mathcal{I}, \kappa, d)] &= U_i(\mathcal{S}, \kappa', d') - U_i(\mathcal{S}, \kappa, d) \\ &= U_i^{(2)}(\kappa', d') - U_i^{(2)}(\kappa, d) \\ &= (\kappa' - \kappa)S_i + (d - d')D_i \geq 0. \end{aligned}$$

Moreover, note for all $\kappa' \geq \kappa$ and $d' \leq d$:

$$\begin{aligned} [U_i(\mathcal{S}, \kappa', d') - U_i(\mathcal{D}, \kappa', d')] - [U_i(\mathcal{S}, \kappa, d) - U_i(\mathcal{D}, \kappa, d)] &= U_i(\mathcal{S}, \kappa', d') - U_i(\mathcal{S}, \kappa, d) \\ &= U_i^{(2)}(\kappa', d') - U_i^{(2)}(\kappa, d) \\ &= (\kappa' - \kappa)S_i + (d - d')D_i \geq 0. \end{aligned}$$

Finally, for all $\kappa' \geq \kappa$ and $d' \leq d$:

$$[U_i(\mathcal{I}, \kappa', d') - U_i(\mathcal{D}, \kappa', d')] - [U_i(\mathcal{I}, \kappa, d) - U_i(\mathcal{D}, \kappa, d)] = 0.$$

Thus, via Topkis's theorem (see [Topkis \(1998\)](#)), there is uniformly more sharing. ■

A.3 Proofs from Section 4

Proof of Lemma 1. To obtain a contradiction, suppose that there exists an agent i and an agent j with $\ell_i = \ell_j$ but either (i) $b_i^* \neq b_j^*$ or (ii) $b_i^{**} \neq b_j^{**}$.

Without loss of generality, suppose that $b_i^* < b_j^*$. By way of contradiction suppose $b_i^{**} > b_j^*$, and consider priors $\tilde{b} \in (b_i^*, \min\{b_j^*, b_i^{**}\})$ where agent i would ignore but agent j with that same prior would dislike. However, both agents with prior \tilde{b} receive payoff $\tilde{u}(1 - \pi(\tilde{b})) - \tilde{c}$ from disliking and payoff 0 from ignoring. Thus, one of them must not be playing a best response, and $b_i^{**} = b_j^*$.

Thus, when agents i and j both have some prior $b' \in (b_i^*, b_j^*)$, agent i shares and agent j dislikes. By symmetry of agent i and j 's network positions, it is clear that for agent i and agent j with prior b' that $U_j(\mathcal{S}) - U_i(\mathcal{S}) = p_s(\kappa + d)$. Similarly, $U_j(\mathcal{D}) - U_i(\mathcal{D}) = 0$. But in this case,

$$\begin{aligned} [U_j(\mathcal{S}) - U_i(\mathcal{S})] - [U_j(\mathcal{D}) - U_i(\mathcal{D})] &= [U_j(\mathcal{S}) - U_j(\mathcal{D})] + [U_i(\mathcal{D}) - U_i(\mathcal{S})] \\ &= p_s(\kappa + d) > 0. \end{aligned}$$

This implies that either $[U_j(\mathcal{S}) - U_j(\mathcal{D})] > 0$ or $[U_i(\mathcal{D}) - U_i(\mathcal{S})] > 0$ (or both). This yields a contradiction because at prior b' , it is supposed to be a best response for agent j to play \mathcal{D} and a best response for agent i to play \mathcal{S} . Thus, $b_i^* = b_j^*$.

Without loss of generality, suppose that $b_i^{**} < b_j^{**}$. If $b_i^{**} \leq b_j^*$, then for priors $b'' \in (b_i^{**}, b_j^*)$, agent i

shares and agent j dislikes. Via the same reasoning as in the previous paragraph, this is a contradiction, so $b_j^* < b_i^{**} < b_j^{**}$. Let us consider some prior $\hat{b} \in (b_i^{**}, b_j^{**})$, where agent i shares and agent j ignores. By symmetry of agent i and j 's network positions, it is clear that for agent i and agent j with prior \hat{b} that $U_j(\mathcal{S}) - U_i(\mathcal{S}) = p_s \kappa$. Similarly, $U_j(\mathcal{I}) - U_i(\mathcal{I}) = 0$. Then notice that:

$$\begin{aligned} [U_j(\mathcal{S}) - U_i(\mathcal{S})] - [U_j(\mathcal{I}) - U_i(\mathcal{I})] &= [U_j(\mathcal{S}) - U_j(\mathcal{I})] + [U_i(\mathcal{I}) - U_i(\mathcal{S})] \\ &= p_s \kappa > 0. \end{aligned}$$

This implies that either $[U_j(\mathcal{S}) - U_j(\mathcal{I})] > 0$ or $[U_i(\mathcal{I}) - U_i(\mathcal{S})] > 0$ (or both). However, this is a contradiction because at prior b' , it is supposed to be a best response for agent j to play \mathcal{I} and a best response for agent i to play \mathcal{S} . Thus, $b_i^{**} = b_j^{**}$. ■

Proof of Theorem 2. For part (a), let us consider belief $b^{(2)} < 1$ and a reliability threshold \underline{r} such that for all \mathbf{P} , all agents with $b < b^{(2)}$ choose \mathcal{D} in every equilibrium (including the most-sharing equilibrium) whenever the article has reliability $r < \underline{r}$. Such an \underline{r} exists by Lemma A.4(ii). Thus, for all $r < \underline{r}$, every agent on an island $\ell \geq 2$ dislikes in the most-sharing equilibrium, regardless of \mathbf{P} .

Next, we consider an increase in homophily (while holding expected degree fixed). By our choice of \underline{r} , all agents on islands $\ell \geq 2$ still dislike in the most-sharing equilibrium whenever $r < \underline{r}$. We can thus consider the social media game that only involves island 1, treating islands 2 through k as automata that always dislike. Before the shift in homophily, consider the equilibrium cutoffs (b_1^*, b_1^{**}) for island 1 in the most-sharing equilibrium (the same for all agents on island 1, per Lemma 1) and let \mathbf{B}_1 denote the modified cutoff space defined by all cutoffs $(\hat{b}_1^*, \hat{b}_1^{**}) \preceq (b_1^*, b_1^{**})$. Finally we define a map $\varphi : \mathbf{B}_1 \rightarrow \mathbf{B}_1$ that maps cutoffs in \mathbf{B}_1 , $(\hat{b}_1^*, \hat{b}_1^{**})$, to best-response cutoffs $(\hat{b}_1^{*,BR}, \hat{b}_1^{**,BR})$, given that agents on island 1 play according $(\hat{b}_1^*, \hat{b}_1^{**})$. By the arguments in Lemma A.2, φ preserves \succeq and \mathbf{B}_1 is a complete sublattice, provided that the map φ is well-defined in that it always maps to an element in \mathbf{B}_1 .

To establish this, consider the utility $U_1(\mathcal{S})$ of sharing on island 1 with homophily parameters (p_s, p_d) , holding fixed the cutoff strategy $(\hat{b}_1^*, \hat{b}_1^{**})$ and the expected degree of each agent on island 1, ζ_1 . Thus, we can write $p_d = (\zeta - N_1 p_s) / (N - N_1)$ and observe then that

$$U_1(\mathcal{S}) = U_1^{(1)} + \kappa N_1 p_s (1 - H(\hat{b}_1^{**})) - d \left(N_1 p_s H(\hat{b}_1^*) + \frac{\zeta - N_1 p_s}{N - N_1} \cdot (N - N_1) \right),$$

and in particular,

$$\frac{\partial U_1(\mathcal{S})}{\partial p_s} = \kappa N_1 (1 - H(\hat{b}_1^{**})) + d N_1 (1 - H(\hat{b}_1^*)) > 0.$$

Therefore, if we compare utility $U_1'(\mathcal{S})$ after the increase in homophily to $U_1(\mathcal{S})$ before the increase in homophily (leaving $(\hat{b}_1^*, \hat{b}_1^{**})$ fixed), we see that $U_1'(\mathcal{S}) \geq U_1(\mathcal{S})$. Hence, φ necessarily maps any cutoffs in \mathbf{B}_1 into \mathbf{B}_1 . Applying the same Tarski fixed-point theorem as before (see Tarski (1955)), the set of fixed points (and thus Bayesian-Nash equilibria) form a lattice within the space of cutoffs \mathbf{B}_1 . Moreover, there is a most-sharing equilibrium in \mathbf{B}_1 , which is also the most-sharing equilibrium in \mathbf{B} . We denote

this equilibrium by $(b_1^*, b_1^{**'})$ and note that $(b_1^*, b_1^{**'}) \preceq (b_1^*, b_1^{**})$ (because it lies in \mathbf{B}_1). In particular, this means $b_1^{**'} \leq b_1^{**}$, and more agents share on island 1 in the most-sharing equilibrium following the rise in homophily.

To measure the change in virality, we first observe that the seed agent i^* (that maximizes $\mathbb{E}[\mathbf{S}_{i^*}]$) is chosen from the agents on island 1. We consider the virality of the article when agents on island 1 share with probability $1 - H(b_1^{**})$ under the stronger homophily structure (p'_s, p'_d) versus (p_s, p_d) (and all other agents kill the article). This is sufficient to show that virality increases following the increase in homophily, because obviously virality with $b_1^{**'} < b_1^{**}$ (but the same network \mathbf{P}) is strictly higher, given that agents on island 1 share more often ($1 - H(b_1^{**'}) > 1 - H(b_1^{**})$).

We consider the diffusion process of an article on the (p'_s, p'_d) network that starts with an agent on island 1. Let us define a *path* of the diffusion process to be a chain $i^* \rightarrow i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_z$ representing a sequence of agents who receive the article in this process, with i^* being the seed agent, i_1 through i_{z-1} all being agents who shared it, and agent i_z being an agent who either ignored or disliked the article. There may be many such paths for the diffusion of the article (by assumption, all agents with the possible exception of agent i_z must be on island 1).

For each path, we define an alternative path (generated randomly) as follows. For any links to agents other than to agent i_z (i.e., links within island 1), with probability $(p'_s - p_s)/p'_s$, the link instead goes to one of islands $2, \dots, k$ (chosen in proportion to their population) and otherwise remains the same. Applying this to all paths, we define an isomorphic diffusion process to one on a sharing network with weaker homophily parameters (p_s, p_d) . However, note that the length of every path cannot increase following this transformation. Because any transition to islands $2, \dots, k$ is necessarily the end of the path, paths can only shorten. Moreover, the number of paths must weakly decrease. As a result, the fraction of agents who receive the article (i.e., \mathbf{S}_{i^*}) must be lower, and virality is less under the (p_s, p_d) sharing network. This establishes part (a).

For part (b), we first note that there exists \bar{r} such that the most-sharing equilibrium when $r > \bar{r}$ is all-share ($b_\ell^{**} = 0$ for all islands ℓ) regardless of \mathbf{P} . Notice that equation (2) is minimized when $b_i = 0$, and in particular, for all agents i (regardless of their prior):

$$\pi_i \geq \frac{(1-p)\phi(r)}{(1-q)(1-\phi(r)) + (1-p)\phi(r)}.$$

Then, letting $\bar{\pi} = \max \left\{ \frac{c}{u+c}, \frac{\tilde{u}-\tilde{c}}{\tilde{u}} \right\} < 1$, we note that whenever

$$r \geq \phi^{-1} \left(\frac{(1-q)\bar{\pi}}{(p-q)\bar{\pi} + (1-p)} \right) \equiv \bar{r} \in (0, 1),$$

then $\pi_i \geq \bar{\pi}$. Of course, when all other agents (other than i) share and $r > \bar{r}$, $U_i(\mathcal{S}) \geq u\pi_i - c(1 - \pi_i) \geq 0$ and $U_i(\mathcal{D}) = \tilde{u}(1 - \pi_i) - \tilde{c} \leq 0$, so $a_i = \mathcal{S}$ is a best response for agent i . Thus, the most-sharing equilibrium is all-share (because it is an equilibrium and no other strategy profile can have more sharing).

Observe that when $r > \bar{r}$, virality is measured simply by the expected size of the connected component (formed by \mathbf{P}) containing the seed agent i^* . Regardless of the homophily parameters, the seed agent i^* will be chosen from the largest island (call this island ℓ^*). To see this, consider a seed agent on some other island ℓ' . The probability of a connection between island ℓ^* and some other island ℓ'' is $1 - (1 - p_d)^{N_{\ell^*}N_{\ell''}} > 1 - (1 - p_d)^{N_{\ell'}N_{\ell''}}$. Thus, for every island ℓ'' , the likelihood that island ℓ'' gets the article (directly) from island ℓ^* is higher than (directly) from island ℓ' . Moreover, the likelihood that island ℓ'' gets the article *indirectly* from island ℓ^* is also higher than (indirectly) from island ℓ' , because the probability of connected islands outside of ℓ^* and ℓ' are independent of the seed (conditional on it being either island ℓ^* or ℓ'). Finally, the engagement within island ℓ^* is higher than ℓ' because $N_{\ell^*} > N_{\ell'}$ (and both are weakly connected by assumption). Thus, seeding an agent i^* on island ℓ^* maximizes S_{i^*} , and virality is measured by the expected number of islands connected to island ℓ^* , weighted by their populations.

Lastly, we note that the probability that island ℓ has any connections to island ℓ' is given by $\tilde{p}_{\ell,\ell'} = 1 - (1 - p_d)^{N_{\ell}N_{\ell'}}$ *before* the decrease in homophily and $\tilde{p}'_{\ell,\ell'} = 1 - (1 - p'_d)^{N_{\ell}N_{\ell'}}$ *after* the decrease in homophily, with $\tilde{p}'_{\ell,\ell'} > \tilde{p}_{\ell,\ell'}$ for all pairs of islands (ℓ, ℓ') because $p'_d > p_d$. Using the same terminology as in the argument for part (a), we map the diffusion paths of an article under the less homophilic sharing network with (p'_s, p'_d) . Consider cycles between islands $\ell^* \rightarrow \ell_1 \rightarrow \ell_2 \dots \rightarrow \ell_z$, where ℓ_z is the same island as one of $\ell^*, \ell_1, \dots, \ell_z$ (in which case, no additional engagement is obtained thereafter the article returns to island ℓ_z). Before the decrease in homophily (where $p_d < p'_d$), we can construct an isomorphic diffusion process where an article remains within the same island (instead of switching to a different one) with probability $(p'_d - p_d)/p_d$. By construction of the cycle, whenever such an event occurs, the cycle becomes complete and the islands reached thereafter in the (p'_s, p'_d) sharing network are not (for that given cycle). Measuring across all cycles that occur in the (p'_s, p'_d) model, (weakly) more islands are reached than under the more homophilic (p_s, p_d) model. Consequently, virality is higher under the (p'_s, p'_d) sharing network (with less homophily) than with the (p_s, p_d) sharing network (more homophily). This establishes part (b). ■

Proof of Proposition 3. Let us define r^* as

$$r^* \equiv \phi^{-1} \left(\max \left\{ \frac{(1-q)(\tilde{u} - \tilde{c})}{(p-q)(\tilde{u} - \tilde{c}) + (1-p)\tilde{u}}, \frac{c}{u+c} \right\} \right) \in (0, 1).$$

For part (a), first consider the case of $r < r^*$ and $p_d = 0$ (by continuity, the result extends to the case of sufficiently large p_s/p_d). In the most-sharing equilibrium, the seed agent most conducive to the article's spread is on the right-wing island, and given that $p_d = 0$, the equilibrium on the left-wing island is immaterial to its virality. Let us denote the right-wing island cutoffs by (b_R^*, b_R^{**}) (cutoffs are only island-dependent, per Lemma 1). Similar to the proof of Theorem 2(a), we define a cutoff space \mathbf{B}_R such that $(\hat{b}_R^*, \hat{b}_R^{**}) \in \mathbf{B}_R$ if and only if $(\hat{b}_R^*, \hat{b}_R^{**}) \preceq (b_R^*, b_R^{**})$. Similarly, we define the map $\varphi : \mathbf{B}_R \rightarrow \mathbf{B}_R$ which maps an arbitrary cutoff $(\hat{b}_R^*, \hat{b}_R^{**})$ to best-response cutoffs $(\hat{b}_R^{*,BR}, \hat{b}_R^{**,BR})$. To show the map is well-defined, consider $U_R(\mathcal{S})$ *before* the increase in divisiveness or polarization and $U'_R(\mathcal{S})$ *after* the increase

in divisiveness or polarization. Because the network structure is fixed, note that $U_R^{(2)}(\mathcal{S}) = U_R^{(2)'}(\mathcal{S})$ when the cutoffs $(\hat{b}_R^*, \hat{b}_R^{**})$ are taken as given, so the difference $U_R'(\mathcal{S}) - U_R(\mathcal{S})$ depends only on the difference between $U_R^{(1)}(\mathcal{S})$ and $U_R^{(1)'}(\mathcal{S})$. Specifically, the difference in share payoff depends only on the change in π_i following the increase in divisiveness or polarization. Observe that:

$$\begin{aligned}\frac{\partial \pi_i}{\partial p} &= \frac{(2b_i - 1)(1 - \phi(r))\phi(r)(1 - q - b_i(1 - 2q))}{(b_i(2p\phi(r) + 2q(1 - \phi(r)) - 1) - p\phi(r) - q(1 - \phi(r)) + 1)^2} > 0; \\ \frac{\partial \pi_i}{\partial q} &= \frac{(2b_i - 1)(1 - \phi(r))\phi(r)(1 - p + b_i(2p - 1))}{((2b_i - 1)\phi(r)(p - q) + 2b_iq - b_i - q + 1)^2} > 0,\end{aligned}$$

whenever $b_i > 1/2$. Likewise, as we showed in Lemma A.1, $\partial \pi_i / \partial b_i > 0$ for all b_i and greater polarization increases ideological priors for agents with $b_i > 1/2$ (by Lemma A.3). By virtue of $\underline{b}_R > 1/2$, we observe that $U_R^{(1)'}(\mathcal{S}) > U_R^{(1)}(\mathcal{S})$, and so $U_R'(\mathcal{S}) > U_R(\mathcal{S})$. Thus, as in the proof of Theorem 2(a), φ is well-defined. Applying the Tarski fixed-point theorem (see Tarski (1955)), we find that the most-sharing equilibrium leads to more sharing in the right-wing island. Because the network structure \mathbf{P} remains constant and there is a uniform shift in sharing, our weaker notion of virality also increases.

For part (b), consider $r \geq r^*$. Note that for $r \geq r^*$, ignoring is a better response to disliking for any agent (regardless of the agent's prior) and sharing is a better response to ignoring for all $b_i > 1/2$. The former follows from noting $\pi_i \geq \frac{\tilde{u}-\tilde{c}}{\tilde{u}}$ for an agent with prior $b_i = 0$ and the latter from noting $\pi_i \geq \frac{c}{u+c}$ for agents with $b_i > 1/2$ and observing that disliking is a dominated strategy. Therefore, the right-wing island always shares, whereas the left-wing island has equilibrium cutoffs $(0, b_L^{**})$. Using the same approach as in part (a), it is enough to show that $U_L(\mathcal{S})$ increases following a decrease in divisiveness or polarization. Of course, for $b_i < 1/2$, we see that $\partial \pi_i / \partial p < 0$ and $\partial \pi_i / \partial q < 0$, and moreover, by Lemma A.3, decreasing polarization means that all agents on the left-wing island also have an increase in b_i . Thus, there is more sharing in the most-sharing equilibrium following a decrease in divisiveness or polarization. By strategic complementarity, the right-wing island remains at all-share, and sharing uniformly increases (and so does virality, naturally). ■

A.4 Proofs from Section 5

Proof of Theorem 3. Consider the complete sharing network where $\mathbf{P} = \mathbf{1}_{N \times N} - \mathbf{I}$. We claim that if the most-sharing equilibrium involves all agents choosing \mathcal{S} or \mathcal{I} (with probability 1) and agents never choosing \mathcal{D} under this configuration, then this is the platform's profit-maximizing sharing network. By Lemma 1, all agents employ the same cutoffs $(b^*, b^{**}) = (0, b^{**})$ and $1 - H(b^{**})$ determines the proportion who share in the most-sharing equilibrium.

We focus on a modified social media game that only allows agents to ignore or share, which necessarily increases virality of content for any sharing network \mathbf{P}' but does not increase the virality for the complete sharing network, by assumption. We show that for any other sharing network \mathbf{P}' , the largest fixed point (the most-sharing equilibrium), must necessarily be above $b^{**}\mathbf{1}$ (in the order \preceq). To

do this, we consider the largest fixed point under \mathbf{P}' (call this $\mathbf{b}^{**'}$), and use the same mathematical arguments as before (although we disregard the dislike cutoff). Let \mathbf{B}' be the cutoff space where $\hat{\mathbf{b}}^{**}$ satisfies $\hat{\mathbf{b}}^{**} \preceq \mathbf{b}^{**'}$ and let the map $\varphi : \mathbf{B}' \rightarrow \mathbf{B}'$ map fixed cutoff strategies $\hat{\mathbf{b}}^{**}$ to best-response sharing cutoff strategies under the complete sharing network. It only remains to prove that φ indeed maps into \mathbf{B}' . To do this, let $U_j^c(\mathcal{S})$ be the utility from sharing under the complete network and $U_j'(\mathcal{S})$ as sharing under \mathbf{P}' , and note that

$$U_j^c(\mathcal{S}) - U_j'(\mathcal{S}) = \kappa \sum_{\tilde{j}=1}^N (1 - p'_{j\tilde{j}})(1 - H(\hat{b}_{\tilde{j}}^{**})) \geq 0.$$

Thus, by Tarski's fixed-point theorem (Tarski (1955)), we obtain that $b^{**}\mathbf{1} \preceq \mathbf{b}^{**'}$. Finally, observe that this necessarily implies that \mathbf{P}' is less viral, because for every prior realization and seed agent i^* , \mathbf{S}_{i^*} is larger in the complete network than in any other sharing network, provided that $b^{**}\mathbf{1} \preceq \mathbf{b}^{**'}$ and $b^* = 0$ (no agent dislikes). By Proposition 1, (uniformly more) sharing is monotone in reliability, so there exists some r_P such that for $r > r_P$, the complete sharing network admits only shares and ignores, whereas when $r < r_P$, agents dislike with positive probability. When $r > r_P$, the network takes the form of part (ii) by setting $p_s = p_d = 1$.

Next, we consider the case where $r < r_P$, and so (b^*, b^{**}) are the cutoffs in the most-sharing equilibrium with a complete sharing network, but where $b^* > 0$. First, notice that there must exist an open interval where agents with priors $b_i \in (0, \bar{b})$ will *never* share, regardless of the sharing network \mathbf{P} . To see this, suppose that there exists some \mathbf{P}' where *all* agents either share or ignore, so the equilibrium cutoffs are determined by $\mathbf{b}^{**'}$ (and $\mathbf{b}^{*'} = \mathbf{0}$). Using the reasoning as in the previous paragraph (but extending it to the full cutoff space $(b^*\mathbf{1}, b^{**}\mathbf{1})$), we conclude that the cutoffs in the most-sharing equilibrium of the complete sharing network must satisfy $(b^*\mathbf{1}, b^{**}\mathbf{1}) \preceq (\mathbf{b}^{*'}, \mathbf{b}^{**'})$, and in particular, $b^*\mathbf{1} \preceq \mathbf{b}^{*'} = \mathbf{0}$. This implies that all agents share or ignore in the complete network, yielding a contradiction. Therefore, such an interval $(0, \bar{b})$ must exist, and in particular, we choose the largest such \bar{b} (in the supremum sense) where agents with priors in $(0, \bar{b})$ *never* share in any sharing network \mathbf{P} in the most-sharing equilibrium.

Next, we consider disconnecting (and removing) all agents in any community ℓ with $b^{(\ell)} < \bar{b}$, but leaving all other communities connected in a (partial) complete network. We call this network the *active network*. We claim that when ε is sufficiently small, all of the remaining agents in the active network either share or ignore. By definition, an agent with \bar{b} would share under some sharing network \mathbf{P}^* but any agent with $\bar{b} - \varepsilon$ would ignore (for arbitrarily small ε) under \mathbf{P}^* (and by leveraging Lemma A.4(i), not all agents with $b < \bar{b}$ dislike). This implies that $U_i(\mathcal{D}) < 0$ for an agent with prior \bar{b} (by monotonicity), and thus an agent with this prior either shares or ignores in the active network. Moreover, for agents with priors in a small half-open neighborhood around \bar{b} (i.e., an interval $(\bar{b} - \eta, \bar{b}]$ for some $\eta > 0$) ignoring is a better response to disliking in the active network. Thus, for sufficiently small ε , we obtain a partial complete network (the active network) with agents who only share and

ignore (with probability 1) and never dislike (with probability 0).

Finally, with these two observations, we claim that the profit-maximizing sharing network takes the form of part (i). First, consider all communities who participated in the active network described above (call these the active communities) and suppose that the agents in communities outside of this active network are non-existent in our model (call these the inactive communities). When the active communities are arranged in a complete sharing network, we showed in the previous paragraph that all agents either share or ignore. By the exact argument in the first two paragraphs then, engagement (and virality) are maximized (amongst only the active communities) when these communities are arranged in a complete sharing network. Second, by construction of the active network (and the active communities), all agents in inactive communities *never* share under *any* sharing network \mathbf{P} . Therefore, removing these agents is without loss to potential virality. Hence, whenever virality is maximized amongst only agents in active communities, it is also maximized in general.

Lastly, we note that we can form a (partial) complete network among the inactive communities, but provide no connections to the (partial) complete network of active communities who only share and ignore (with probability 1). By our previous observations, this is a profit-maximizing sharing network for the platform. At the same time, it is exactly the form of a two-island model with $(p_s, p_d) = (1, 0)$, which has maximal homophily. ■

Proof of Proposition 4. Note by Theorem 3 that an agent i with prior $b_i = b^{(k+1)}$ is indifferent between ignoring and disliking when $r = r_P$ (but strictly prefers to either share or ignore for all $r > r_P$), so r_P increases if and only if this agent (strictly) prefers to dislike following a shift in parameters. Because $b^{(k+1)} < 1/2$, an increase in polarization means that agent i 's prior decreases (see Lemma A.3), and given that $\partial\pi_i/\partial b_i > 0$ (see Lemma A.1), π_i decreases for this agent. As a consequence $U_i(\mathcal{D})$ increases but $U_i(\mathcal{I})$ remains the same, so agent i (strictly) prefers to dislike. Similarly, because $\partial\pi_i/\partial p < 0$ and $\partial\pi_i/\partial q < 0$ for $b_i < 1/2$ (see Proposition 3), π_i decreases for this agent (making $a_i = \mathcal{D}$ a best response). In both cases, we see that r_P increases. ■

A.5 Proofs from Section 6

Proof of Proposition 5. For part (a), let $V > 0$ be the virality of any given article before the provenance policy. Note by Lemma A.4 that $\kappa \leq c/N$, so in particular it is never a best response for agent i to play $a_i = \mathcal{S}$ when $\pi_i = 0$. Conditional on an article containing misinformation and given provenance policy ρ , the seed agent's probability of sharing it is bounded above $1 - \rho$, and the virality of the misinformation article is thus bounded above by $(1 - \rho)N$. Setting $\rho^* = (N - V)/N$ guarantees that the provenance policy reduces the virality of misinformation.

For part (b), we expand on some of the details in Example 1, where we fix an article that necessarily contains misinformation. Consider the deterministic ring network where each user potentially shares the article with one other user or kills the article's propagation altogether, as shown in Figure 3. Note

the first ignore or dislike in the ring ends the lifetime of the article. All user beliefs are uniformly distributed on $[0, 2/3]$, the message is $m = R$, $\phi(r) = 1/2$, $p = 0.7$, and $q = 0.3$. We suppose that parameters are such that strategic complementarities play no role ($\kappa = d = 0$), and that agents share content they believe is more likely truthful (more than probability $1/2$ of $\nu = \mathcal{T}$) and ignore or dislike content they believe is more likely misinformation (more than probability $1/2$ of $\nu = \mathcal{M}$). This is satisfied when $u = c$.

With no provenance policy, we note that users with beliefs $b_i > 1/2$ share the article and those with beliefs $b_i < 1/2$ do not. By symmetry, note that the choice of seed agent does not affect the expected spread of the article, so without loss of generality we can assume the article starts at agent 1 when computing its virality. With each user i sharing the article with probability $\zeta = 1/4$, the total virality of the article can thus be computed simply as $\sum_{i=1}^N \zeta^i \approx \zeta/(1 - \zeta) = 1/3$ for large N .

Next, we consider a provenance policy with $\rho = 1/4$, so with probability $1/4$ each user (independently) will fact-check the content. Given our previous assumption, if a user fact-checks she will share when the fact-check comes back as $\nu = \mathcal{T}$ and will not share when the fact-check comes back $\nu = \mathcal{M}$. We once again suppose the article starts at seed agent 1.²⁵ User 1, similar to the no provenance policy, will share the article with probability $\zeta = 1/4$ if there is no fact-check. On the other hand, user 1 will share the article with certainty if there is a fact-check and it comes back $\nu = \mathcal{T}$, but will ignore the article with certainty if there is a fact-check and it comes back $\nu = \mathcal{M}$. Thus, user 1 shares the article with probability $\zeta(1 - \rho) = 3/16$. Any user $i \geq 2$ will know the article might have been previously fact-checked, and therefore will assign higher likelihood to its veracity than π_i defined in equation (2). In particular, user 2 will assign likelihood

$$\pi'_2 = \frac{(\rho + (1 - \rho)\zeta)\pi_2}{\rho\pi_2 + (1 - \rho)\zeta} > \pi_2,$$

that the article is truthful. For our parameter values, one can verify that $\pi'_2 > 1/2$ regardless of $b_i \in [0, 2/3]$ and so user 2 shares the article with probability 1 if she does not fact-check. Thus, the probability that agent 2 shares is $1 - \rho = 3/4$. Moreover, because user 2 (and future users) might fact-check with probability $\rho = 1/4$, it is easy to see via induction that all users after user 1 share with probability 1 if they do not fact-check. Conditional on user 1 sharing, the expected spread of the article is given by $\sum_{i=2}^N (1 - \rho)^{i-1} \approx (1 - \rho)/\rho = 3$ for large N . Finally, the virality of the article after the provenance policy is $3/16 + 3/16 \cdot 3 = 3/4 > 1/3$. Thus, the provenance policy increases the virality of a misinformation article.

For part (c), we note that the generous bound in part (a) still applies identically even if the platform chooses a different sharing network following the provenance policy. ■

²⁵We comment that the platform might, in fact, want to randomize over seed agents with a provenance policy in place because a deterministic seeding rule will induce lower sharing from the seed agent (as she knows no previous agents have fact-checked). For this reason, we instead establish a lower bound on the virality of content, assuming a deterministic seeding rule that is higher than virality without provenance. This is sufficient to show the provenance policy leads to more virality, because the platform can never do worse.

Proof of Proposition 6. For part (a), once again let $V > 0$ be the virality of any given article before the censorship policy. The virality of a misinformation article that is not detected (and censored) by the platform is upper bounded by N , so provided that $\delta > (N - V)/N \equiv \delta^*$ (and therefore, $(1 - \delta)N < V$), the censorship policy reduces the virality of misinformation.

For part (b), we modify Example 1 to incorporate a censorship policy instead of a provenance policy. Suppose the exact same parameters as in Example 1 and where $\delta = 1/4$. Notice that the virality of a misinformation article before the censorship policy is the same as before the provenance policy (i.e., it is equal to $1/3$). On the other hand, we find that after the censorship policy, agents with priors $b_i \in [1/3, 2/3]$ will share the article instead of ignore or dislike it, i.e., half of the population will share the article instead of just $1/4$. This occurs because of an “implied truth” effect from content that is not tagged as misinformation and censored by the regulator. Consequently, conditional on the article not being flagged by the regulator as misinformation, the virality will be $\sum_{i=1}^N (1/2)^i \approx 1$ for large N . Because $1 - \delta = 3/4$ of the misinformation is permitted, the virality of misinformation is $3/4 \cdot 1 = 3/4 > 1/3$.²⁶

For part (c), as in the proof of Proposition 5, the generous bound on content virality following the censorship policy still applies even if the platform elects a different sharing network. Thus, the censorship policy with $\delta > \delta^*$ is still effective. ■

Proof of Proposition 7. For part (a), first consider a new (stricter) performance target $\lambda^* < \lambda$ (where λ is the current performance target). The platform can meet the new performance standard λ^* by recommending any content with $r \geq \phi^{-1}(\lambda^*)$, which by continuity of ϕ and positive density of f over $r \in [\phi^{-1}(\lambda^*), 1]$ guarantees some level of expected engagement $E > 0$ whereas (platform-optimal) engagement while neglecting the performance target is some $\tilde{E} \geq E$. Setting $\alpha^* = (\tilde{E} - E)/\tilde{E}$ implies that for all $\alpha > \alpha^*$, the platform prefers to meet λ^* than neglect it. Finally, note that to meet a lower performance target λ^* , the total shares on the platform (in expectation) must decrease; otherwise, with the more relaxed $\lambda > \lambda^*$ target, the platform could have increased engagement with a different sharing network (and thus is not implementing a profit-maximizing sharing network). Given that the performance target is the ratio of misinformation shares to total shares, and total shares have decreased, the platform must decrease misinformation shares; in particular, the platform decreases the virality of misinformation following the stricter performance target when $\alpha > \alpha^*$.

For part (b), we consider a complete sharing network with uniformly distributed beliefs on $[0, 1]$, but otherwise we adopt the same parameter values as in Example 1. Moreover, the auditing technology to enforce performance targets is only $\alpha = 1/2$.

For simplicity, we suppose that there are just three types of articles: high-reliability and low-sensationalism (type HL) articles, low-reliability and low-sensationalism articles (type LL) articles,

²⁶Note that the virality of misinformation in the provenance policy example (Proposition 5) and in the censorship example (Proposition 6) are identical just by coincidence, and this is not a property that holds generally.

and low-reliability but high-sensationalism (type LH) articles. The type-HL articles never contain misinformation ($\phi(r) = 1$) and have no sensationalism ($\kappa = 0$). The type-LL articles contain misinformation with probability $2/3$ ($\phi(r) = 1/3$) and have no sensationalism ($\kappa = 0$). The type-LH articles also contain misinformation with probability $2/3$ ($\phi(r) = 1/3$), but have high sensationalism given by $\kappa > (2c - u)/(3N) > 0$. We assume that all three article types are equally likely. In the case of type-HL and type-LH articles, all agents necessarily share in the most-sharing equilibrium. On the other hand, only agents with priors $b_i \in [3/4, 1]$ share the type-LL articles.

When $\lambda > 34/99$, the platform can allow all content. In particular, the platform gets an (engagement) payoff proportional to $N/3 + N/3 + N/48 = 11N/16$ because it meets the performance target: total shares are given by $11N/16$ and misinformation shares are given by $N/72 + 2N/9 = 17N/72$, and this ratio is below λ . However, when $\lambda = 1/3$, the platform faces the following trade-off: it can still allow all content or remove just the type-LL articles. By removing the type-LL content, it can meet the performance target exactly and guarantee engagement of $2N/3$, but by allowing the type-LL content, can increase engagement to $11N/16$. However, the $\alpha = 1/2$ probability of auditing (so expected engagement payoff of only $11N/32 < 2N/3$) is sufficient to have the platform remove the type-LL content. Therefore, reducing the performance target reduces the virality of misinformation.

Finally, however, suppose that λ is set close to 0. Thus, to meet the performance target, the platform must remove type-LH and type-LL content and receive engagement of only $N/3$. However, by allowing all content, the platform receives $11N/16$, which in expectation, is $11N/32 > N/3$ even with auditing technology $\alpha = 1/2$. Consequently, the lower performance target leads the platform to introduce type-LL articles back onto the platform, *increasing* the virality of misinformation.

For part (c), simply setting $\tilde{E} = N$ and letting $\alpha^* = (N - E)/N$ ensures that even if the platform chooses a new sharing network following the performance target regulation, that it will necessarily choose to meet that target whenever $\alpha > \alpha^*$. By the same argument as in part (a), total shares must then (weakly) decrease following this stricter performance target, which requires misinformation shares (and the virality of misinformation) to decrease as well. ■

Proof of Proposition 8. The network regulation does not bind for an article with $r > r_P$, so we need only consider $r < r_P$. Take some agent i with prior $b_i \in (\bar{b}, \bar{b} + \eta)$ in a small neighborhood $\eta > 0$ of \bar{b} (where \bar{b} is the same \bar{b} constructed in Theorem 3). Following the same line of reasoning as in Theorem 2(a), agents with priors in this interval elect to ignore instead of share following the network regulation (and when η is sufficiently small), and this necessarily reduces the virality of misinformation, showing (ii). To prove (i), we note that agents in this neighborhood around \bar{b} also do not share in the most-sharing equilibrium under *any* sharing network P' (following the network regulation), per the construction of \bar{b} in Theorem 3. Therefore, the platform cannot generate additional engagement by departing from the class of island models (specifically, two-island models) while maintaining $p_s/p_d \leq p^*$. ■