Social Networks and the Market for News

Rachel Kranton and David McAdams

April 20, 2020

Abstract: This paper introduces a simple market model for news: consumers benefit from and want to share true news and producers incur costs to produce true news. News veracity is endogenous, shaped by the social network. When producer revenues derive from consumers' viewing stories (e.g., advertising revenue), veracity is low in dense networks, since even false news spreads widely. With revenues from consumers' actions based on stories (e.g, voting), veracity is higher in dense networks, since consumers make better inferences about news truth. Adding third-party misinformation can increase equilibrium true-news production as consumers respond by being more judicious when sharing stories.

Keywords: social networks, news veracity, misinformation

Kranton: Economics Department, Duke University (email: rachel.kranton@duke.edu). McAdams: Fuqua School of Business and Economics Department, Duke University (email: david.mcadams@duke.edu). We thank seminar participants at Cornell, Dartmouth, Duke, EUI, GSE Barcelona, Johns Hopkins, Penn State, the 6th European Conference on Networks, the 17th IO Theory Conference, and the 2019 Cowles Conference in Economic Theory for helpful comments. We also thank Guglhupf Bakery, Cafe, & Biergarten in Durham, NC for its hospitality while this research was conducted.

The 2016 Presidential election in the United States and the subsequent media environment have raised both public and academic interest in "fake news" and overall news quality. "Fake news" often refers to information that the provider knows to be false. As such, fake news is not new. Tabloid newspapers have long published questionable stories about celebrities. Governments have used false information to influence public opinion at home and abroad (elaborated below), and one objective of disinformation campaigns is to undermine overall trust in the news. At the same time, bona fide news producers make decisions about the quality of the stories they broadcast to the public. This paper studies news quality in a stylized model of the current news market, distinguished by providers who can reach consumers through online distribution channels and consumers who share stories through social media.

Our innovation is to endogenize the veracity of the news. Producers decide whether to incur costs to produce "high-quality," true stories. Consumers evaluate the news that they receive and desire to share and act only on true news. We characterize outcomes first in the baseline case when producers are paid per consumer who views their story, such as through accompanying advertising. We consider the impact of network structure on news veracity as well as the effect of third-party misinformation. We then study news producers with more partisan motives, where revenues derive from the number of consumers who take action based on their stories, such as voting.

The model applies to any decision-relevant information shared socially and specifically captures the spectrum of provider motivations in the contemporary media market. Traditional brick and mortar newspapers, such as *The New York Times*, or fictitious-news websites, such as denverguardian.com,[1] earn revenues from advertising accompanying their articles. While consumers might base decisions on their stories, these outlets do

[1]Fake-news website denverguardian.com famously published a false story linking Hillary Clinton to the death of an FBI agent (Borchers, 2016).

not typically directly earn revenue from those decisions. News outlets such as Fox News, MSNBC, Breitbart, and Sinclair Media earn revenues from advertising but can also be supported by owners who care about advancing their own political views.[2] Finally, government-sponsored media, such as the British propagandists of 1940 or the Russian troll factories of today, seek to induce people to take some desired action or to disrupt the media market altogether.[3]

Our results emphasize how network structure and the way in which producers earn revenue impact equilibrium news veracity. In a sparsely connected network when producers' revenues are based on views, more links increase equilibrium quality; producers correctly anticipate that true stories are shared and hence viewed more frequently. But as the network becomes very dense, producers have little incentive to invest in story quality since even false stories spread widely. By contrast, when producers' revenues derive from actions based on their stories and consumers are highly connected, producers have a strong incentive to invest in story quality. Consumers' inferences about the truth of a story become more precise as they are able to observe more neighbors' sharing decisions. However, even in this case, consumers' ability to discern the truth may be limited, i.e., there may not be a "wisdom of the crowd."

Misinformation from outside sources, such as from government agencies that publish false news stories, alter equilibrium veracity in different ways. As might be expected, a large quantity of misinformation leads to a breakdown of the news market. Consumers do not believe or share any stories, and bona fide producers do not invest in quality content. However, a small quantity of misinformation can in some cases increase true-

[2]For an exposé of Sinclair Media and its CEO David Smith, see Kroll (2017).

[3]Britian deployed three thousand operatives to the United Stated in 1940 to spread (sometimes false) stories under the guise of news reports to raise American popular support for entering the war effort against Nazi Germany (Cull 1995). More recently, the Russian-based Internet Research Agency created *Heart of Texas*, a fictitious advocacy group that promoted Texas secession from the United States and other provocative positions. When its Facebook page called for a protest against "the Islamification of Texas" in 2017, real people showed up to protest and counterprotest (Allbright, 2017).

news production. Knowing that false stories are being injected into the market from outside sources, consumers become more judicious when deciding which stories to share, which in turn gives bona fide producers more incentive to invest in publishing high-quality news.

The paper contributes to three distinct literatures: (i) social learning, information transmission, and networks, (ii) media markets, and (iii) misinformation.

Social learning, information transmission, and networks. The demand side of our market specifies information transmission and social learning that is both similar to and different from other models. Consumers here receive private signals and rationally update beliefs about each news item based on others' sharing decisions. However, unlike in the cascades literature (e.g., Banerjee 1992 and Bikhchandani, Hirshleifer, and Welch 1992), consumers observe multiple neighbors' independent sharing decisions in one round of social learning. As in Bloch, Demange, and Kranton (2018) and Chatterjee and Dutta (2016), but unlike much of the network literature on information diffusion (e.g., Acemoglu, Ozdaglar, and ParandehGheibi 2010 and Banerjee et al. 2013), consumers in our model choose whether or not to pass on information to their neighbors. These decisions ultimately determine the network's role in spreading and filtering the news. To the best of our knowledge, this is the first paper to endogenize the product which spreads in a network setting—how the product itself is shaped by the network.[4]

Media markets. Much previous work on news markets studies possible sources of media

[4]Many papers in diverse fields have examined how network structure impacts the decisions of a third party who cares about outcomes, e.g., a health authority deciding how best to control an epidemic (Peng et al. 2013) or a supply-chain manager deciding how best to operate its warehouses (Beamon and Fernandes 2004). The idea of endogenizing what passes through the network is rarely explored in these literatures, but there are exceptions, e.g., Read et al. (2015) on endogenous pathogen virulence and Bimpikis, Fearing, and Tahbaz-Salehi (2018) on upstream sourcing in a supply chain. Previous research studies the effect of social-network structure on other producer decisions for a given product, such as relying on traditional versus word-of-mouth advertising (Galeotti and Goyal 2009) or targeting consumers when launching a new product (Chatterjee and Dutta 2016, Bimpikis, Ozdaglar, and Yildiz 2016).

bias. In Gentzkow and Shapiro (2006), news producers earn revenues based on their reputation for accuracy and thus have an incentive to slant their news towards consumers' initial beliefs. In Besley and Prat (2006) and Gentzkow, Glaeser, and Goldin (2006), earning revenue from advertising, rather than a sponsor, reduces bias. In Ellman and Germano (2009), however, newspapers bias their news towards their advertisers. In the present paper, consumers care about the veracity of news. A key insight from the analysis is that news veracity is lower when producers' revenues depend only on advertising. In that case, producers only care about how many consumers view their stories and, in dense networks, even false news is widely viewed. When producers earn revenues from consumers' actions, in contrast, their incentive to produce true news is based on consumers' inferences, which improve in dense networks.[5]

Misinformation and gaslighting. In 1923, the Soviet Union launched the first modern black-propaganda office, with the aim of "manipulating a nation's intelligence system through the injection of credible but misleading data" (Safire 1989), a tactic Joseph Stalin dubbed "dezinformatsiya (disinformation)" (Manning and Romerstein 2004). State-sponsored disinformation efforts now abound[6] and are often online.[7] Consumers encounter false news from other sources as well, including individuals and social bots who

[5]Several recent papers study other features of contemporary media markets, such as competition for consumers' limited attention (Chen and Suen 2018), media bias when consumers have heterogeneous preferences and pass on news to like-minded individuals (Redlicki 2017), and competition to break a story that leads to lower-quality news (Andreottola and de Moragas 2018).

[6]To give some examples: In 2016, an Iranian operation published over one hundred fake articles on websites posing as legitimate news outlets, including a story apparently from the Belgian newspaper *Le Soir* claiming that Emmanuel Macron's campaign was financed by Saudi Arabia (Lim et al. 2019). In 2014, Russia spread false stories about the downing of a civilian airliner, attempting to implicate Ukrainian forces (Mills 2014). In 1985, the Soviets conducted "Operation INFEKTION" to drive world opinion that the United States had invented AIDS to kill black people (Boghardt 2009), a falsehood still believed by nearly one in five young black South Africans as late as 2009 (Grebe and Nattrass 2012). In 1978, a Soviet-controlled newspaper in San Francisco published a story falsely claiming that the Carter administration supported the apartheid government of South Africa (Romerstein 2001).

[7]With the rise of "deep fake" video technology, it will become even harder for news consumers to distinguish true from false sources. Even seeing may no longer be enough to believe (https://www.cnet.com/videos/were-not-ready-for-the-deepfake-revolution/).

spread conspiracy theories on social media and in memes.[8] The problem is so severe that, even seven years ago, the World Economic Forum listed digital misinformation in online social media as one of the main threats to our society (Howell 2013a, b). A large and varied literature studies misinformation, examining how falsehoods and conspiracy theories spread differently than fact-based information on the Internet (del Vicario et al. 2016 and Vosoughi, Roy, and Aral 2018).[9] The psychology literature on "gaslighting" also studies how an abuser, through persistent lying and contradiction, seeks to cause targeted individuals to doubt their own memory and perception (Dorpat 1996).

We evaluate the impact of misinformation, interpreted as the third-party broadcast of false stories, and gaslighting, interpreted as lowered consumer confidence in their private signals, on bona fide news producers' equilibrium investment in news truth. While large misinformation campaigns and highly-effective gaslighting unambiguously undermine the news market, smaller efforts can in some cases lead to increased true-news production. Consumers pass on fewer stories, but producers' incentive to broadcast true stories increases, as false stories are less likely to get through consumers' more stringent filter.

The paper proceeds as follows. Section 1 presents the basic news-market model. Section 2 characterizes equilibrium outcomes when producers are paid for views, and Section 3 considers how expanding social networks affects news veracity. In Section 4, we study large markets in which producers' revenues derive from consumers taking actions based on their stories. The Conclusion outlines directions for future research.

[8] A recent trending example is the meme "Epstein didn't kill himself" (Ellis 2019).

[9] Specific studies include the effect of misinformation on an Ebola outbreak in West Africa (Oyeyemi, Gabarron, and Wynn 2014) and on a French presidential election (Ferrera 2017); how exposure to misinformation can shape memory (Loftus 2005 and Zhu et al. 2010); and how to identify misinformation and reduce its harmful impact (Qazvinian et al. 2011 and Shao et al. 2016).

1 Model: The Market for News

The market for decision-relevant information, which we refer to as "news," consists of a large finite number N of consumers, of whom M generate revenue for producers.[10] Producers are modeled as a unit-mass continuum of agents, but the analysis applies equally to a setting with finitely-many producers or even a single identifiable producer, as long as producers lack commitment power.[11] Low quality stories are costless to produce and are false with probability one; high-quality stories entail a "reporting cost" $c_R > 0$ and are true with probability one.[12] (Our analysis extends easily to the case when low-quality news is sometimes true and high-quality news is sometimes false.) The cost c_R is an i.i.d. random variable across stories with support $(0, \infty)$ and continuous distribution $H(c_R)$.

Each consumer is linked to others in a directed social network, with a link from consumer i to consumer j indicating that i can observe whatever news j decides to share, i.e., i "follows" j. We use the word "neighbors" to describe consumers who are linked, with the context indicating the link's direction. For simplicity, we focus on networks in which each consumer follows d others and refer to d as "social connectedness;" networks with higher d are then "more connected."

---

[10]The distinction between revenue-generating and non-revenue-generating consumers allows us to study the impact of increasing the number of social links while holding producers' revenue base fixed; see Section 3. The model also encompasses situations in which producers only care about reaching (say) a single consumer. For example, some stories aired on Fox News in 2019 were aimed specifically at Donald Trump, reaching him while he watched the channel and indirectly through related social-media activity (Shields and Dlouhy 2019).

[11]The model thus applies to particular news providers or reporters, each interacting with consumers in its own "news market." In this setting, there could be a reputational cost associated with publishing a false story, which can be incorporated into the model by adjusting the support of the cost distribution (specified below) to allow for negative reporting costs.

[12]While we focus on a context in which the thing being produced is a *factual claim*, our analysis applies more broadly to settings where consumers care about any unobservable product characteristic, e.g., the entertainment value of a new movie, the effectiveness of a new scientific practice (with "consumers" being scientists), or the viability of a potential political candidate (with "consumers" being political donors).

The news-market game proceeds in three phases $t = \{0, 1, 2\}$. At $t = 0$, each producer sees the realization of his reporting cost c_R and decides whether to produce a high- or low-quality story. All stories are then "broadcast," seen by each consumer with independent probability $b \in (0, 1]$.[13] Let p_0 be the ex ante likelihood that stories are true, referred to as "news veracity."

At $t = 1$, each consumer who saw a story's broadcast decides whether to share the story with her neighbors. By assumption, consumers cannot directly observe story quality but they can evaluate stories based on their own expertise, personal experience, or access to other information, modeled as a private signal $s_i \in \{T, F\}$ about news truth. These signals are informative, with $\Pr(s_i = T|\text{true}) = \Pr(s_i = F|\text{false}) = \rho \in \left(\frac{1}{2}, 1\right)$, conditionally i.i.d. across consumers and i.i.d. across stories.[14] Consumers prefer to share true stories[15] but not false stories.[16] A consumer earns "sharing payoff" $\pi_T^S > 0$ from sharing a true story, $-\pi_F^S < 0$ from sharing a false story, and zero payoff from not sharing. Consumers therefore prefer to share whenever they believe that a story's likelihood of being true exceeds "sharing threshold" $p^S = \frac{\pi_F^S}{\pi_T^S + \pi_F^S} \in (0, 1)$. For notational simplicity, we normalize $\pi_T^S = \pi_F^S$ so that $p^S = \frac{1}{2}$.

At $t = 2$, consumers view the stories shared by their neighbors and each consumer who has seen a story decides whether to take an "action" based on it, earning $\pi_T^A > 0$ when acting on a true story, $-\pi_F^A < 0$ when acting on a false story, and zero payoff when not acting. Consumers therefore prefer to act on a story when its likelihood of being true

[13] In Appendix B, we extend the analysis to allow consumers to have different likelihoods of seeing the broadcast, b_i, different numbers of neighbors, d_i, and different private-signal precisions, ρ_i, among other asymmetries.

[14] Because private signals are i.i.d. across stories, signals and sharing behavior about one story are uninformative about other stories. We can therefore consider each story in isolation.

[15] A consumer's overall incentive to share a story could depend on observable characteristics, such as novelty, as well as on unobservable quality. Our analysis focuses on the strategic issues created by the presence of an unobservable characteristic, namely "truth," holding observable characteristics fixed.

[16] The analysis becomes trivial if consumers prefer to share false stories, since then nothing is learned from others' sharing behavior.

exceeds "action threshold" $p^A = \frac{\pi_F^A}{\pi_T^A + \pi_F^A} \in (0,1)$. For expositional ease, we assume in the main text that $\pi_T^A = \pi_F^A$ so that $p^A = \frac{1}{2}$.[17]

The truth of the story is revealed at the end of $t = 2$, at which point consumers' sharing and action payoffs are realized. In the baseline model studied next, producers earn a unit of revenue for each consumer that has viewed its story. We then study markets where producers earn a unit of revenue for each consumer that acts on its story.[18]

2 News Markets with Revenues from Views

In this section, we characterize equilibrium outcomes in the news market when each producer earns one unit of revenue per consumer who views their story.[19]

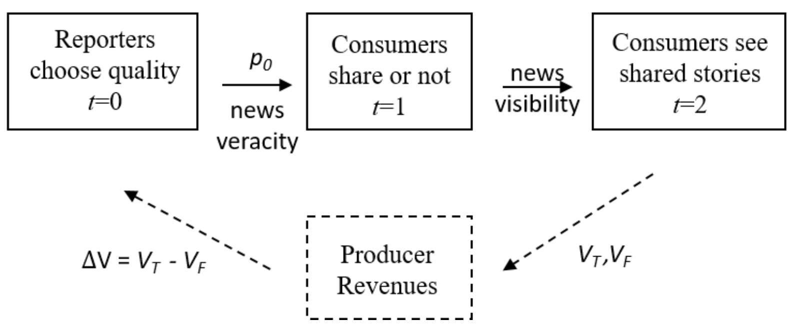

Figure 1: Schematic of a views-supported news market.

Figure 1 illustrates this news-market game. Producers decide whether to produce high- or low-quality news, which determines p_0. This news veracity informs consumers' sharing decisions, which in turn determine the likelihood that any given consumer sees a

[17]This assumption that consumers have the same threshold belief for action as for sharing ($p^A = p^S = \frac{1}{2}$) simplifies the analysis but also rules out some interesting possibilities. In Appendix C, we extend our analysis to allow consumers to have a higher or a lower standard for action than for sharing.

[18]Appendix C analyzes an extension in which producers earn revenue from both views and actions.

[19]Consumers may encounter the same news item multiple times but, by assumption, the producer is only paid once per consumer who sees the news. The analysis can be extended in a straightforward way to allow for non-linear producer revenue.

story, that story's "visibility," denoted V_T and V_F for true and false news, respectively. Producers' incentive to invest depends on the extra visibility of true news, denoted $\Delta V \equiv V_T - V_F$.

We focus on dynamically-stable Bayesian Nash equilibria in which consumers use the same sharing strategy ("equilibria," for short). (For details on dynamic stability, see Appendix A. For characterization of all Bayesian Nash equilibria, see Appendix B.) We solve for equilibria by working backward, first considering consumers' incentives to share and then producers' incentive to invest.

Optimal consumer sharing. Suppose that consumer i has seen a story's broadcast. Given private signal $s_i = T$ or $s_i = F$, i's updated beliefs, denoted $p_1(s_i; p_0)$, are

$$p_1(T; p_0) = \frac{p_0\rho}{p_0\rho + (1-p_0)(1-\rho)} \quad \text{and} \quad p_1(F; p_0) = \frac{p_0(1-\rho)}{p_0(1-\rho) + (1-p_0)\rho}$$

by Bayes' rule. Let z_T and z_F denote each consumer's likelihood of sharing after private signal $s_i = T$ and $s_i = F$, respectively. $\mathbf{z} = (z_T, z_F)$ is called the "sharing rule." Optimal consumer sharing depends on the prior p_0, as illustrated in Figure 2:

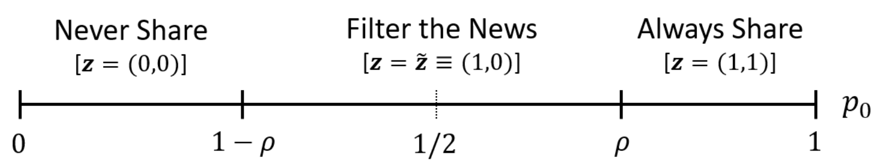

Figure 2: News-veracity regions and optimal consumer sharing.

- *Always-share region* $p_0 \in (\rho, 1]$*:* If news veracity is high enough, consumers find it optimal to share news after a good and after a bad signal, since both $p_1(T; p_0) > \frac{1}{2}$ and $p_1(F; p_0) > \frac{1}{2}$. In this case, sharing is uninformative.

- *Never-share region $p_0 \in [0, 1-\rho)$:* If news veracity is low enough, consumers find it optimal never to share news, since both $p_1(T; p_0) < \frac{1}{2}$ and $p_1(F; p_0) < \frac{1}{2}$.

- *Filtering region $p_0 \in (1-\rho, \rho)$:* If news veracity is in this intermediate range, consumers find it optimal to share after a good signal because $p_1(T; p_0) > \frac{1}{2}$, but find it optimal not to share after a bad signal because $p_1(F; p_0) > \frac{1}{2}$. Sharing here is informative and we say that consumers "filter" the news. Let $\tilde{\mathbf{z}} \equiv (1, 0)$ be shorthand for the optimal sharing rule in this case.

- *Thresholds $p_0 \in \{1-\rho, \rho\}$:* If news veracity is exactly $p_0 = \rho$, what we call the "always-share threshold," consumers are indifferent whether to share after seeing a bad signal ($p_1(F; p_0) = \frac{1}{2}$) and hence use a sharing rule of the form $\mathbf{z} = (1, z_F)$. Similarly, if news veracity is $p_0 = 1-\rho$, the "never-share threshold," consumers are indifferent after seeing a good signal ($p_1(T; p_0) = \frac{1}{2}$) and use a sharing rule of the form $\mathbf{z} = (z_T, 0)$.

Lemma 1 describes consumers' best-response correspondence, denoted $Z(p_0)$:

Lemma 1. *(i) If $p_0 < 1-\rho$, then $Z(p_0) = (0, 0)$. (ii) If $p_0 > \rho$, then $Z(p_0) = (1, 1)$. (iii) If $p_0 \in (1-\rho, \rho)$, then $Z(p_0) = \tilde{\mathbf{z}} \equiv (1, 0)$. (iv) $Z(1-\rho) = \{(z_T, 0) : z_T \in [0, 1]\}$ and $Z(\rho) = \{(1, z_F) : z_F \in [0, 1]\}$.*

Consumer sharing determines the visibility of news stories. Since each neighbor shares true stories with probability $b(\rho z_T + (1-\rho) z_F)$, we have

$$V_T(\mathbf{z}) = 1 - (1-b)(1 - b(\rho z_T + (1-\rho) z_F))^d. \tag{1}$$

Similarly, because each neighbor shares false stories with probability $b((1-\rho) z_T + \rho z_F)$, we have

$$V_F(\mathbf{z}) = 1 - (1-b)(1 - b((1-\rho) z_T + \rho z_F))^d. \tag{2}$$

Optimal producer investment. Next, we turn to each producer's decision whether to incur the reporting cost c_R to produce a true story. Let $R_T(\mathbf{z})$ and $R_F(\mathbf{z})$ be the expected revenue of true and false stories, respectively. With M revenue-generating consumers, $R_T(\mathbf{z}) = MV_T(\mathbf{z})$ and $R_F(\mathbf{z}) = MV_F(\mathbf{z})$. True stories earn a "revenue premium" $\Delta R(\mathbf{z}) = M\Delta V(\mathbf{z})$, where $\Delta V(\mathbf{z}) \equiv V_T(\mathbf{z}) - V_F(\mathbf{z})$.

Producers maximize expected profit by producing high quality whenever $c_R < M\Delta V(\mathbf{z})$, which occurs with ex ante probability $H(M\Delta V(\mathbf{z}))$. The resulting news veracity is denoted $p_0(\mathbf{z})$ and referred to as the "best-response news veracity:"

$$p_0(\mathbf{z}) = H\left(M\Delta V(\mathbf{z})\right). \tag{3}$$

Since $c_R = M\Delta V(\mathbf{z})$ occurs with probability zero, the producer has an essentially-unique best response to any sharing rule.

Figure 3 illustrates the best-response news veracity.[20] First, producers never invest if consumers always share or never share, i.e., $p_0(1,1) = p_0(0,0) = 0$, but do sometimes invest whenever consumer sharing is informative, i.e., $p_0(\mathbf{z}) > 0$ whenever $z_T > z_F$. Then, starting from the left, $p_0(z_T, 0)$ is increasing in z_T below a threshold $\overline{z}_T \leq 1$ and, if $\overline{z}_T < 1$ (as shown), decreasing in z_T above that threshold. Intuitively, sharing after a good private signal magnifies the visibility of true stories more than false stories when sharing is sufficiently rare. However, as z_T increases, false stories are increasingly shared as well, causing $\Delta V(\mathbf{z})$ to fall when d is sufficiently large.[21] Second, $p_0(1, z_F)$ is decreasing in z_F, since sharing after a bad signal spreads false stories more than true stories.

Lemma 2 gathers together these facts about best-response news veracity $p_0(\mathbf{z})$.

[20] To avoid confusion, note that the x-axis of this figure consists of all sharing rules that could potentially be a best response for consumers, i.e., those of the form $(z_T, 0)$ or $(1, z_F)$.

[21] The critical value $\overline{z}_T$ is strictly less then one if d is sufficiently large. However, the magnitudes involved are not large. For example, with $b = \frac{1}{2}$ and $\rho = \frac{2}{3}$, $\overline{z}_T \approx 0.823$ if $d = 5$ and $\overline{z}_T \approx 0.53$ if $d = 9$.

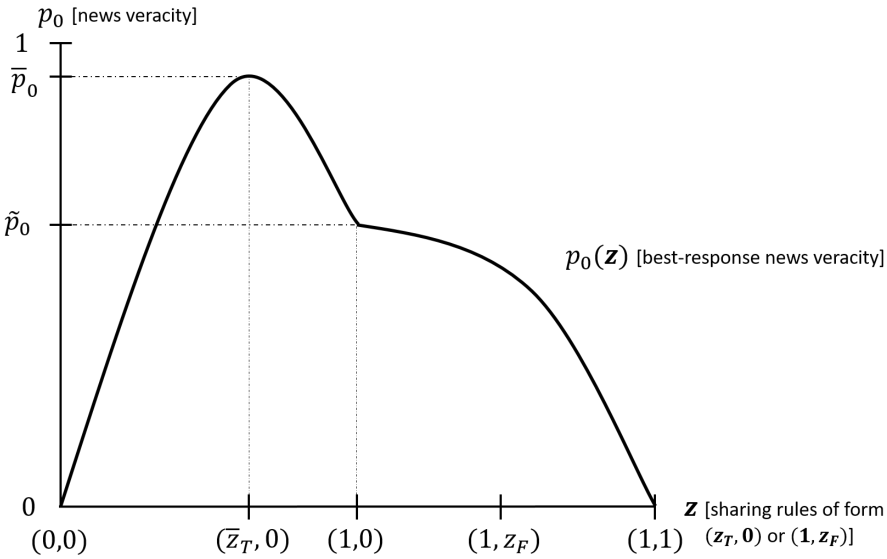

Figure 3: The best-response news veracity function $p_0(\mathbf{z})$, the maximal news veracity $\overline{p}_0 \equiv \max_{\mathbf{z}} p_0(\mathbf{z})$, and the filtering news veracity $\tilde{p}_0 \equiv p_0(1, 0)$.

Lemma 2. *(i) $p_0(z_T, z_F) > 0$ if $z_T > z_F$ and $p_0(z_T, z_F) = 0$ if $z_T = z_F$. (ii) $p_0(z_T, 0)$ is strictly increasing in z_T over the interval $[0, \overline{z}_T]$ and strictly decreasing in z_T over the interval $[\overline{z}_T, 1]$ for some $\overline{z}_T \in (0, 1]$. (iii) $p_0(1, z_F)$ is strictly decreasing in z_F. (iv) $\arg\max_{\mathbf{z}} p_0(\mathbf{z}) = (\overline{z}_T, 0)$. (v) $p_0(\mathbf{z})$ is continuous in $\mathbf{z}$.*

Two specific news-veracity levels are key to the equilibrium analysis below and also illustrated in Figure 3. First, the *maximal news veracity*, denoted $\overline{p}_0$, is the highest veracity that can be achieved when producers invest optimally, given any sharing rule:

$$\overline{p}_0 \equiv \max_{\mathbf{z}} p_0(\mathbf{z}) = p_0(\overline{z}_T, 0) = H\left(M\Delta V(\overline{z}_T, 0)\right),$$

with $(\overline{z}_T, 0)$ being the sharing rule that maximizes producers' incentive to invest (Lemma 2(iv)). Second, the *filtering news veracity*, denoted $\tilde{p}_0$, is the producers' best reply when consumers filter the news:

$$\tilde{p}_0 \equiv p_0(1, 0) = H\left(M \Delta V(1, 0)\right).$$

2.1 Equilibrium characterization

This section characterizes the equilibria of the news-market game. In a "dysfunctional equilibrium," producers never invest, all stories are false, and consumers never share. By contrast, in a "functional equilibrium," producers sometimes invest, consumers sometimes share, and some stories are true. We find, first, that a dysfunctional equilibrium always exists. Second, a functional equilibrium exists if and only if the maximal news veracity exceeds the never-share threshold, i.e., $\overline{p}_0 > 1 - \rho$. Third, when a functional equilibrium exists, it is essentially unique with news veracity, denoted p_0^*, equal to $\max\{1 - \rho, \min\{\tilde{p}_0, \rho\}\}$. In particular, p_0^* must equal either ρ (always-share threshold), $\tilde{p}_0$ (filtering news veracity), or $1 - \rho$ (never-share threshold). Since news veracity cannot exceed ρ, some false news circulates in any equilibrium.

Formally, let $H(\cdot; \gamma)$ denote the distribution of producers' reporting costs when scaled by a parameter $\gamma > 0$, so that $H(c_R; \gamma) = H(c_R/\gamma)$ for all $c_R > 0$. Viewed as functions of the cost-scaling parameter γ, the filtering news veracity is $\tilde{p}_0(\gamma) \equiv H(M\Delta V(1, 0)/\gamma)$ and the maximal news veracity is $\overline{p}_0(\gamma) \equiv H(M\Delta V(\overline{z}_T, 0)/\gamma)$. Since $\tilde{p}_0(\gamma)$ and $\overline{p}_0(\gamma)$ are each continuous and strictly decreasing functions of γ, with $\tilde{p}_0(\gamma) \leq \overline{p}_0(\gamma)$ for all γ, we can define thresholds $0 < \gamma_1 < \gamma_2 \leq \gamma_3 < \infty$ implicitly by the conditions $\tilde{p}_0(\gamma_1) = \rho$, $\tilde{p}_0(\gamma_2) = 1 - \rho$, and $\overline{p}_0(\gamma_3) = 1 - \rho$.

The equilibria depend on the distribution of producers' costs, as illustrated in Figure 4 for four cases, when γ is: (a) less than γ_1, so that $\tilde{p}_0(\gamma) > \rho$; (b) between γ_1 and γ_2, so that $\tilde{p}_0(\gamma) \in (1 - \rho, \rho)$; (c) between γ_2 and γ_3, so that $\tilde{p}_0(\gamma) < 1 - \rho < \overline{p}_0$; or

(d) greater than γ_3, so that $\overline{p}_0(\gamma) \leq 1 - \rho$. In each of the four panels of Figure 4: the x-axis depicts consumers' sharing rule; the y-axis is news veracity; the thick line depicts consumers' best-response correspondence (given by Lemma 1); and the thin line depicts the best-response news veracity (as in Figure 3).

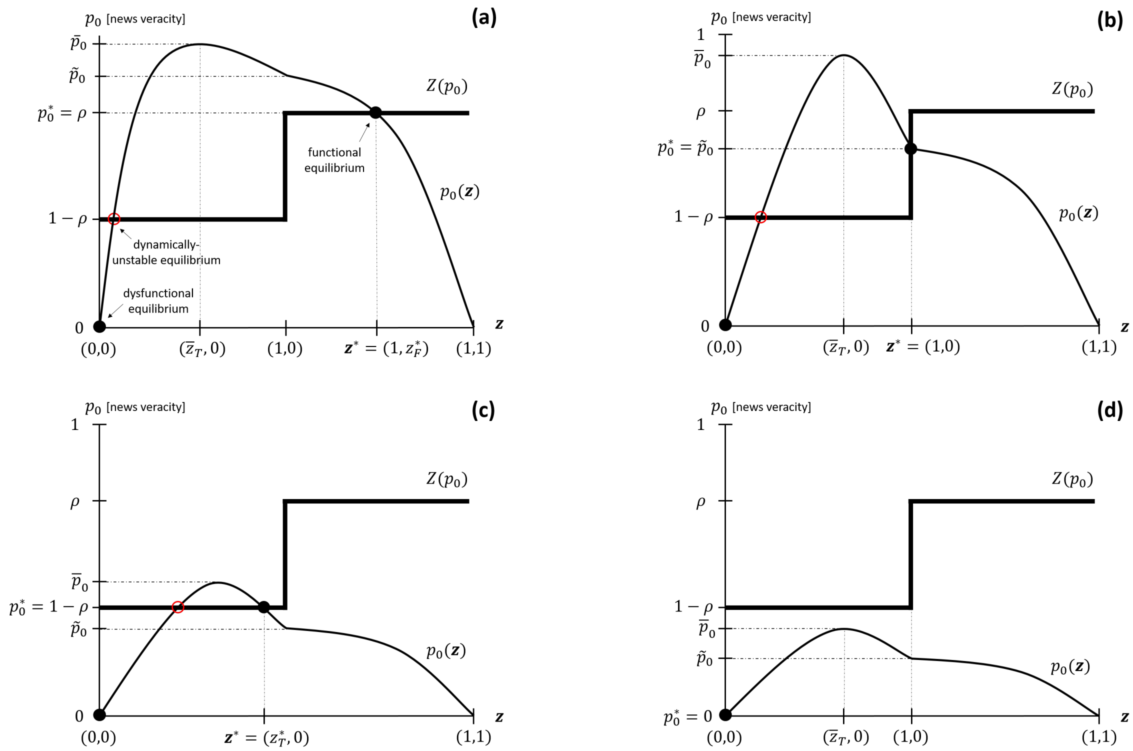

Figure 4: Illustration of the equilibria when (a) $\tilde{p}_0 > \rho$; (b) $\tilde{p}_0 \in (1-\rho, \rho)$; (c) $\tilde{p}_0 < 1 - \rho < \overline{p}_0$; and (d) $\overline{p}_0 < 1 - \rho$.

Each crossing point of these best-reply curves corresponds to a Bayesian Nash equilibrium. In (a-c), there are three such crossing points. However, as shown in Appendix A, only the highest and the lowest of these crossing points are dynamically stable,[22] corresponding to the unique functional equilibrium and the dysfunctional equilibrium, respectively.

[22] The intermediate-veracity equilibrium always has news veracity equal to $1 - \rho$. To see why this equilibrium is dynamically unstable, suppose that news veracity were perturbed to be slightly lower than $1 - \rho$. Consumer adaptation would lead them to share less frequently, inducing producers to invest less frequently, reinforcing the original perturbation.

(a) High equilibrium news veracity (Fig 4(a)). If $\gamma < \gamma_1$, low reporting costs are sufficiently likely that the filtering news veracity $\tilde{p}_0(\gamma)$ exceeds ρ. In the unique functional equilibrium, $p_0^* = \rho$ and the consumer sharing rule is of the form $\mathbf{z}^*(\gamma) = (1, z_F^*(\gamma))$. As γ increases over this range, news veracity remains equal to ρ but consumers' likelihood $z_F^*(\gamma)$ of sharing after a bad signal decreases from 1 to 0.

(b) Intermediate equilibrium news veracity (Fig 4(b)). If $\gamma \in (\gamma_1, \gamma_2)$, then $\tilde{p}_0(\gamma) \in (1-\rho, \rho)$. In the unique functional equilibrium, $p_0^* = \tilde{p}_0(\gamma)$ and consumers filter the news, using sharing rule $\mathbf{z}^*(\gamma) = \tilde{\mathbf{z}} \equiv (1, 0)$. As γ increases over this range, news veracity decreases from ρ to $1 - \rho$.

(c) Low equilibrium news veracity (Fig 4(c)). If $\gamma \in (\gamma_2, \gamma_3)$, then $\tilde{p}_0(\gamma) < 1 - \rho$ and $\overline{p}_0(\gamma) > 1 - \rho$. In the unique functional equilibrium, $p_0^*(\gamma) = 1 - \rho$ and consumers' sharing rule is of the form $\mathbf{z}^*(\gamma) = (z_T^*(\gamma), 0)$. As γ increases over this range, news veracity remains equal to $1 - \rho$ but consumers' likelihood $z_T^*(\gamma)$ of sharing after a good signal decreases from 1 to $\overline{z}_T$ (defined in Lemma 2).

(d) Dysfunctional news market (Fig 4(d)). If $\gamma > \gamma_3$, then $\overline{p}_0(\gamma) < 1 - \rho$ and the dysfunctional equilibrium is the unique equilibrium.

Theorem 1 summarizes these findings, including as well the boundary cases when the cost-scaling parameter γ equals one of the three thresholds.

Theorem 1. *For any γ, there exists a dysfunctional equilibrium in which no consumer shares and no producer invests. (a) If $\gamma \leq \gamma_1$, then a unique functional equilibrium exists and $p_0^*(\gamma) = \rho$. (b) If $\gamma \in (\gamma_1, \gamma_2)$, then a unique functional equilibrium exists, $p_0^*(\gamma) = \tilde{p}_0(\gamma) \in (1-\rho, \rho)$, and $p_0^*(\gamma)$ is strictly decreasing in γ. (c) If $\gamma \in [\gamma_2, \gamma_3)$, then a unique functional equilibrium exists and $p_0^*(\gamma) = 1 - \rho$. (d) If $\gamma \geq \gamma_3$, then no functional equilibrium exists.*

2.2 Misinformation and Equilibrium News Truth

This section examines the impact of misinformation injected into the news market. In addition to the unit mass of "bona fide producers" with reporting costs drawn from distribution $H(\cdot)$, suppose that there is a mass $m \geq 0$ of "misinformation agents" who only produce false stories (as if they have infinite reporting cost). Total quantity $1 + m$ of stories is produced, fraction $\frac{1}{1+m}$ by bona fide producers and fraction $\frac{m}{1+m}$ by misinformation agents.[23] As one might expect, if m is sufficiently large, then the dysfunctional equilibrium is the unique equilibrium. Moreover, when a functional equilibrium exists, adding more misinformation never increases equilibrium news veracity.

However, perhaps surprisingly, more misinformation can induce bona fide producers to invest more, resulting in a greater quantity of true news. Consumers share more judiciously when there is more misinformation in circulation, which in turn can increase the incentive to produce true news.

The analysis mirrors that of the previous section: Let $\tilde{p}_0(m)$, $\overline{p}_0(m)$, and $p_0^*(m)$ denote the filtering news veracity, maximal news veracity, and equilibrium news veracity, as functions of the quantity of misinformation, with shorthand $\tilde{p}_0 = \tilde{p}_0(0)$ and $\overline{p}_0 = \overline{p}_0(0)$ for the baseline case when $m = 0$. Given any sharing rule $\mathbf{z}$, bona fide producers optimally produce quantity $p_0(\mathbf{z})$ of true news. The share of news that is true is therefore $p_0(\mathbf{z}; m) = \frac{p_0(\mathbf{z})}{1+m}$; in particular, the filtering news veracity is $\tilde{p}_0(m) = \frac{\tilde{p}_0}{1+m}$ and maximal news veracity is $\overline{p}_0(m) = \frac{\overline{p}_0}{1+m}$. Define three key thresholds $m_1 < m_2 \leq \overline{m}$:

$$m_1 = \frac{\tilde{p}_0}{\rho} - 1; m_2 = \frac{\tilde{p}_0}{1-\rho} - 1;\ \text{and } \overline{m} = \frac{\overline{p}_0}{1-\rho} - 1. \tag{4}$$

[23] As discussed in footnote 11, our analysis also applies to identifiable individual producers. In that context, "misinformation" corresponds to false stories produced by a third party but made to appear as if produced by that individual, such as impersonation accounts (Goga, Venkatadri, and Gummadi 2015) and social bots (Shao et al. 2018).

By construction, $\tilde{p}_0(m_1) = \rho$, $\tilde{p}_0(m_2) = 1 - \rho$ and $\overline{p}_0(\overline{m}) = 1 - \rho$. Thus, $\tilde{p}_0(m) > \rho$ if and only if $m < m_1$, $\tilde{p}_0(m) > 1 - \rho$ if and only if $m < m_2$, and $\overline{p}_0(m) > 1 - \rho$ if and only if $m < \overline{m}$.

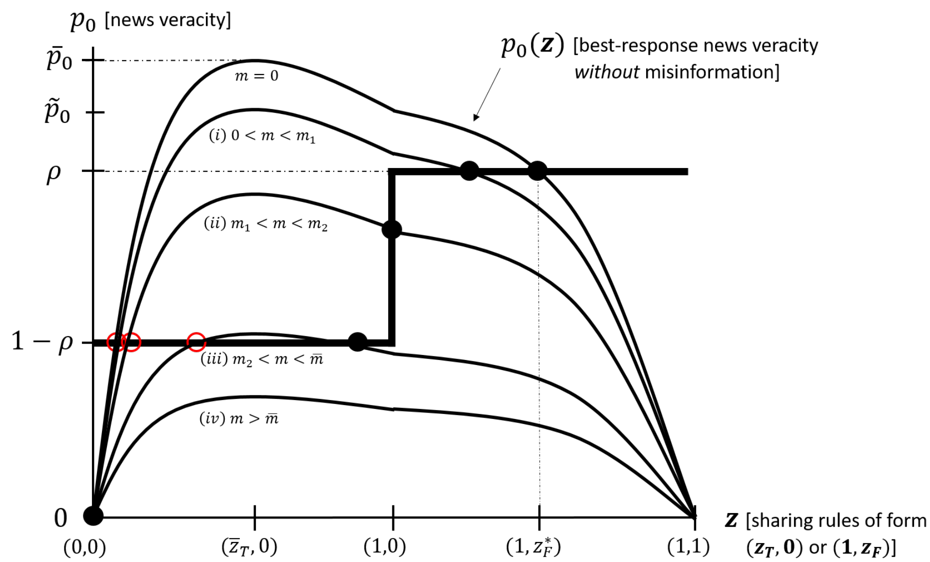

Figure 5: The impact of different quantities of misinformation.

Figure 5 shows how misinformation affects equilibrium outcomes in a scenario with low producer costs, $\gamma < \gamma_1$, so that $\tilde{p}_0 > \rho$ and hence $m_1 > 0$. The thick line shows consumers' best-response correspondence and the highest thin line shows the bona fide producers' best reply $p_0(\mathbf{z})$, as in Figure 4(a). The other four thin lines show the share of news that is true, $p_0(\mathbf{z}; m) = \frac{p_0(\mathbf{z})}{1+m}$, given different quantities of misinformation. The highest crossing-point of each thin line with the thick line corresponds to the functional equilibrium for that quantity of misinformation.

(i) Small quantity of misinformation. Suppose first that $m \in (0, m_1)$. The presence of misinformation shifts down the best-response news veracity curve, from the highest line ($m = 0$) to the second-highest line in Figure 5. However, because $m < m_1$, the filtering news veracity $\tilde{p}_0(m)$ remains above ρ; so, $p_0^*(m)$ remains equal to ρ. Since the overall fraction of news that is true remains the same, even though misinformation is

being injected into the market, bona fide news producers must be increasing true-news production from $p_0^* = \rho$ to $(1+m)p_0^*(m) = (1+m)\rho$. A sufficiently small amount of misinformation therefore *increases* the equilibrium quantity of true stories in circulation. As misinformation quantity increases from 0 to m_1, consumers become more judicious, reducing their probability of sharing after a false signal (z_F^* falls). This in turn increases the relative visibility of true news, giving bona fide producers more incentive to invest in news truth. This can be seen visually in Figure 5, as best-response news veracity is decreasing to the right of $(1, 0)$.

(ii) Intermediate quantity of misinformation. Suppose next that $m \in [m_1, m_2]$. The best-response news veracity curve shifts down from the second-highest to the third-highest line in Figure 5. Because $m \in [m_1, m_2]$, the filtering news veracity $\tilde{p}_0(m) \in [1-\rho, \rho]$. The unique functional equilibrium therefore has news veracity $p_0^*(m) = \tilde{p}_0(m) = \frac{\tilde{p}_0}{1+m}$ and consumers use the filtering rule $\tilde{\mathbf{z}} = (1, 0)$. Bona fide producers optimally respond by producing a constant quantity $\tilde{p}_0$ of true news, and overall equilibrium news veracity falls as misinformation increases over this range.

(iii) Large quantity of misinformation. Suppose next that $m \in (m_2, \overline{m})$, corresponding to the second-lowest line in Figure 5. In this case, filtering news veracity $\tilde{p}_0(m) < 1-\rho$ because $m > m_2$, but the maximal news veracity $\overline{p}_0(m) > 1-\rho$ because $m < \overline{m}$. The functional equilibrium now has news veracity $p_0^*(m) = 1-\rho$ and quantity $(1+m)(1-\rho)$ of true news, which is increasing in m. So, once again, increasing the quantity of misinformation can lead to more true news being produced.

(iv) Overwhelming quantity of misinformation. In the last case when $m > \overline{m}$, $\overline{p}_0(m) < 1-\rho$ and the dysfunctional equilibrium is the unique equilibrium.

Proposition 1 summarizes these findings:

Proposition 1. *(i) A functional equilibrium exists if and only if* $m < \overline{m}$*. (ii) Equilibrium news veracity,* $p_0^*(m)$*, is non-increasing over the range* $m < \overline{m}$ *and strictly decreasing over the range* $m \in (m_1, m_2)$*. (iii) The quantity of true news,* $(1+m)p_0^*(m)$*, is non-decreasing over the range* $m < \overline{m}$ *and strictly increasing over the ranges* $m \in (0, m_1)$ *and* $m \in (m_2, \overline{m})$*.*

2.3 Consumer Evaluation of News and Equilibrium Veracity

Here we study the impact of deep fake technology and efforts to "gaslight" consumers, reducing their ability to assess what is true. Decreasing signal precision from ρ to ρ' has three direct effects. First, holding fixed p_0, consumers switch from always sharing to filtering if $p_0 \in (\rho', \rho)$ or switch from filtering to never sharing if $p_0 \in (1-\rho, 1-\rho')$. Second, again holding fixed p_0, a consumer who only shares after a good signal will share more false stories and fewer true stories. Third, holding consumers' sharing rule fixed, the drop from ρ to ρ' leads to less true news shared and viewed, giving producers less incentive to produce true news.

The interaction of these effects determines how equilibrium news veracity, now denoted $p_0^*(\rho)$, varies with ρ. If $p_0^*(\rho) > 1/2$ so that most news is true, then all three effects go in the same direction and decreasing ρ unambiguously decreases $p_0^*(\rho)$. However, when $p_0^*(\rho) < 1/2$, decreasing ρ may actually increase equilibrium news veracity.

More precisely, the effect of decreasing ρ depends on how ρ compares to thresholds $\underline{\rho}$ and $\overline{\rho}$ implicitly defined by the conditions $\overline{p}_0(\underline{\rho}) = 1 - \underline{\rho}$ and $\tilde{p}_0(\overline{\rho}) = 1 - \overline{\rho}$, respectively. Intuitively, $\underline{\rho}$ is the lowest possible signal precision that supports a functional equilibrium, and $\overline{\rho}$ is the highest signal precision for which equilibrium news veracity is at the never-share threshold. Referring back to Figure 4:

(i) Low signal precision. If $\rho \leq \underline{\rho}$, then $\overline{p}_0(\rho) \leq 1 - \rho$ and the unique equilibrium is

the dysfunctional equilibrium. This parameter range is depicted in Figure 4(d). Any decrease in ρ has no effect.

(ii) Intermediate signal precision. If $\underline{\rho} < \rho \leq \overline{\rho}$, then $\overline{p}_0(\rho) > 1 - \rho \geq \tilde{p}_0(\rho)$. There is a unique functional equilibrium, with news veracity $p_0^*(\rho) = 1 - \rho$. Over this parameter range, then, $p_0^*(\rho)$ is strictly decreasing in ρ. This counter-intuitive effect can be seen graphically in Figure 4(c). As ρ decreases, producers' best-response news veracity (the thin line) moves down but the bottom step of consumers' best-response correspondence moves up. The crossing-point of these curves therefore moves up and to the left; producers invest more in quality and consumers' share less often after a good signal.

(iii) High signal precision. If $\rho > \overline{\rho}$, then $\tilde{p}_0(\rho) > 1 - \rho$. A unique functional equilibrium exists with news veracity $p_0^*(\rho) = \min\{\tilde{p}_0(\rho), \rho\}$. Since $\tilde{p}_0(\rho)$ and ρ are each greater than $1 - \rho$ and strictly increasing in ρ, so is $p_0^*(\rho)$ over this parameter range. Graphically, in Figure 4(a-b), as ρ decreases, producers' best-response news veracity (the thin line) and the top step of consumers' best-response correspondence (the thick line) both decrease. The crossing-point corresponding to the functional equilibrium must therefore also decrease.

Proposition 2 summarizes:

Proposition 2. *Signal-precision thresholds $\frac{1}{2} < \underline{\rho} \leq \overline{\rho} \leq 1$ exist such that: (i) if $\rho \in (1/2, \underline{\rho}]$, then the dysfunctional equilibrium is the unique equilibrium; (ii) if $\rho \in (\underline{\rho}, \overline{\rho}]$, then there is a unique functional equilibrium, $p_0^*(\rho) = 1 - \rho$, and $p_0^*(\rho)$ is decreasing in ρ over this range; and (iii) if $\rho \in (\overline{\rho}, 1]$, then there is a unique functional equilibrium, $p_0^*(\rho) = \min\{\tilde{p}_0(\rho), \rho\} > 1 - \rho$, and $p_0^*(\rho)$ is increasing in ρ over this range.*

Thus, if equilibrium news veracity is high enough that consumers always share after a good private signal, i.e., $p_0^*(\rho) > 1 - \rho$, then reducing consumers' ability to discern which stories are true causes news veracity to fall. On the other hand, if $p_0^*(\rho) = 1 - \rho$

so that consumers are indifferent whether to share after a good signal, slightly reducing consumers' ability to discern the truth causes news veracity to increase. As consumers become more cautious and share stories less often, the visibility of true and false stories declines, but more so for false stories. As a result, the extra visibility of true stories increases, giving producers greater incentive to invest.

3 Social Connectedness and News Veracity

This section examines how more links among consumers impacts equilibrium news veracity when producers are paid for views. We find that adding links to a sparse network leads to more true news, while adding links to a dense network leads to less true news. Moreover, when social connectedness is sufficiently high, equilibrium news veracity cannot exceed $1 - \rho$, the lowest level consistent with any consumer sharing.

3.1 Finitely Dense Networks

As a first step, consider how the filtering news veracity, denoted here as $\tilde{p}_0(d)$, varies with social connectedness d. In the Appendix, we show that $\tilde{p}_0(d)$ is single-peaked in d, rising to a maximum at a level denoted $\tilde{d}$ and then declining to zero as d grows large.

For intuition, consider the specific impact of increasing $d = 0$ to $d = 1$ and from d to $d + 1$ large. For $d = 0$, all stories are seen with broadcast probability b; true stories have no extra visibility. Producers therefore have no incentive to invest, and $\tilde{p}_0(0) = 0$. When a consumer follows one person, $d = 1$, that link increases the consumer's likelihood of viewing any given story by $(1 - b)b\rho$ if the story is true or by $(1 - b)b(1 - \rho)$ if the story is false. The extra visibility of true stories therefore increases from 0 to $(1 - b)b(2\rho - 1)$, inducing producers sometimes to invest. In particular, $\tilde{p}_0(1) = H(M(1-b)b(2\rho-1)) > 0$.

When a consumer follows d others, where d is large, almost all stories are false *condi-*

tional on not being shared by any of those d neighbors. The stories shared only by the $(d+1)$-st neighbor are therefore almost all false. Adding a $(d+1)$-st neighbor therefore increases false-news visibility more than it increases true-news visibility, reducing the extra visibility of true stories and thereby reducing producers' incentive to invest.

Given that $\tilde{p}_0(d)$ is single-peaked in d, there exist thresholds $\underline{d}$ and $\overline{d}$ such that $\tilde{p}_0(d) > 1-\rho$ if and only if $d \in \left(\underline{d}, \overline{d}\right)$. By Thm 1(a-b), a unique functional equilibrium therefore exists whenever $d \in \left(\underline{d}, \overline{d}\right)$, with news veracity $p_0^*(d) = \min\{\tilde{p}_0(d), \rho\}$. Note that $p_0^*(d)$ only depends on d through its impact on the filtering news veracity $\tilde{p}_0(d)$, and indeed that $p_0^*(d)$ is a non-decreasing function of $\tilde{p}_0(d)$. The fact that $\tilde{p}_0(d)$ is single-peaked and maximized at $\tilde{d}$ therefore implies that $p_0^*(d)$ is single-peaked over the range $d \in \left(\underline{d}, \overline{d}\right)$ and also maximized at $\tilde{d}$.

The flip side of this observation is that equilibrium news veracity cannot exceed $1-\rho$ whenever $d \leq \underline{d}$ or $d \geq \overline{d}$. In those cases, there are two possibilities: *either* only the dysfunctional equilibrium exists (as when $d = 0$), so that all stories are false, *or* a functional equilibrium exists with news veracity equal to $1-\rho$, so that consumers are indifferent whether to share after a good private signal. (We examine what happens in markets with large d in more detail in Section 3.2.)

Proposition 3 summarizes:

Proposition 3. *There exist* $0 \leq \underline{d} \leq \overline{d} < \infty$ *such that:* $p_0^*(d) > 1-\rho$ *if and only if* $d \in \left(\underline{d}, \overline{d}\right)$*;* $p_0^*(d)$ *is single-peaked in* d *over this range; and* $p_0^*(d)$ *is maximized at* $d = \tilde{d}$*. Moreover, for all* $d \leq \underline{d}$ *and* $d \geq \overline{d}$*,* either *the dysfunctional equilibrium is the unique equilibrium* or *a unique functional equilibrium exists with news veracity equal to* $1-\rho$*.*

Proposition 3 highlights a weakness of densely connected news markets. When producers are paid for views, as assumed in this section, most news is false in any equilibrium. To see why, suppose that news veracity were higher than $1-\rho$. Everyone with a positive private signal would then share, causing true and false stories all to be very likely seen. But this would give producers approximately zero incentive to invest, causing best-response news veracity to be approximately zero, a contradiction. However, equilibria can exist with news veracity equal to $1-\rho$ in which stories are shared—but not so widely as to undermine producers' incentive to invest. In such equilibria, consumers use a sharing rule of the form $\mathbf{z}^*(d) = (z_T^*(d), 0)$, where $z_T^*(d)$ grows small as d grows large.[24]

3.2 Infinitely Dense Networks

Consider next the limit of a sequence of news markets, each having the same number M of revenue-generating consumers but with social connectedness d going to infinity, what we refer to as the "limit-market." If news quality and consumers' sharing rule were fixed, consumers in the limit-market would, by the Law of Large Numbers, be able to discern perfectly which stories are true based on the (non-zero) fraction of their infinitely many neighbors who share each piece of news. There would be a "wisdom of the crowd" (Galton 1907). However, news quality and consumers' sharing behavior is not fixed but rather determined in equilibrium.

The equilibria in the limit-market depend on the precision of consumers' private signals. For each d, let $\underline{\rho}(d)$ be the lower signal-precision threshold defined in Section 2.3, and let $\underline{\rho}^\infty \equiv \lim_{d\to\infty} \underline{\rho}(d)$. Similarly, let $p_0^*(d)$ denote the news veracity in the unique functional equilibrium for any given d, if it exists, or $p_0^*(d) = 0$ if no functional equilibrium exists, and let $p_0^{*\infty} = \lim_{d\to\infty} p_0^*(d)$. We have:

Proposition 4. *(i) If $\rho < \underline{\rho}^\infty$, then $p_0^{*\infty} = 0$ and the limit-market is dysfunctional. (ii) If $\rho > \underline{\rho}^\infty$, then the limit market has a unique functional equilibrium and $p_0^{*\infty} = 1-\rho$.*

There is no "wisdom of the crowd" in equilibrium. In the only functional equilibrium,

[24] As established in Appendix B, $\lim_{d\to\infty} dz_T^*(d) \in (0,\infty)$, i.e., $z_T^*(d)$ goes to zero as rate $1/d$.

veracity is equal to $1-\rho$, and consumers are unable to infer for sure which stories are true based on others' sharing behavior. Mathematical details are in Appendix B but, to gain intuition, note that news veracity being equal to $1-\rho$ means that true stories must generate $H^{-1}(1-\rho)$ more revenue than false stories. This in turn requires that true-story visibility V_T exceed false-story visibility V_F by $H^{-1}(1-\rho)/M > 0$. Hence, a significant fraction of stories must go unseen by each consumer, even though they have infinitely many neighbors. A consumer who sees a story's broadcast but sees no neighbor sharing infers that others' failure to share the story is "bad news" but, since some true stories are also never shared, cannot determine for sure whether the story is true or false. Similarly, if a consumer were to see several neighbors sharing, that would be "good news" but also not definitive. Consumers will not know for sure which stories are true and which are false.

4 News Markets with Revenues from Actions

In this section, we study markets where producers earn revenue for each consumer who takes a specific action based on their story, such as voting (partisan news) or buying a product (sponsor-supported news).

Figure 6 illustrates this news-market game. Producers choose quality, and consumers then decide whether to share stories. Observing the number of neighbors that share a story, consumers make inferences about the story's truth, which inform their decisions whether or not to act on the story. Consumers' sharing decisions thus impact producers' incentives in two ways, affecting how widely true and false news spreads and consumers' inferences about news truth.

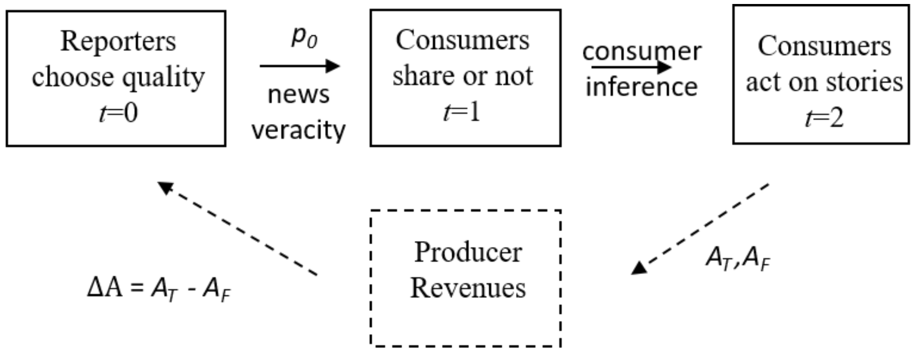

Figure 6: Schematic of an actions-supported news market.

4.1 Limit-market equilibrium characterization

For clarity and brevity, we focus on the "limit-market" introduced in Section 3.2.[25] We characterize the limit-market equilibria and the corresponding news veracity, denoted $p_0^{*A\infty}$. As in a views-supported market, there is always a dysfunctional equilibrium in which producers never invest, so that all stories are false, and consumers never share nor act on any stories. Moreover, no functional equilibrium exists with veracity higher than ρ or lower than $1-\rho$. If the former, consumers would share and act on all stories, giving producers no incentive to invest and causing all stories to be false, a contradiction. If the latter, consumers would never share nor act on any stories, and producers would again have no incentive to invest, another contradiction. All functional equilibria must therefore have news veracity in the range $[1-\rho, \rho]$.

Suppose for a moment that consumers were able to perfectly discern which stories are true. Consumers would then act on all true stories and not act on all false ones, generating the greatest possible revenue (M) for true stories and the least possible revenue (zero) for false ones. Producers would then have the greatest possible incentive to invest, whenever reporting cost $c_R < M$. We refer to the resulting news veracity, $H(M/\gamma)$, as the "maximal conceivable news veracity." Next, define cost-parameter thresholds $\underline{\gamma}$ and $\overline{\gamma}$ implicitly by

[25] Appendix C provides a full analysis of actions-supported news markets with finite social networks, as well as an extension allowing producers to be paid for both views and for actions.

the conditions $H(M/\overline{\gamma}) = 1 - \rho$ and $H(M/\underline{\gamma}) = \rho$. The thresholds $\underline{\gamma}$ and $\overline{\gamma}$ capture how high news veracity could conceivably be: greater than ρ if $\gamma < \underline{\gamma}$ ("low costs"); in the interval $(1 - \rho, \rho]$ if $\gamma \in [\underline{\gamma}, \overline{\gamma})$ ("intermediate costs"); or less than or equal to $1 - \rho$ if $\gamma \geq \overline{\gamma}$ ("high costs"). Figure 7 illustrates the equilibrium outcomes for costs in different ranges.

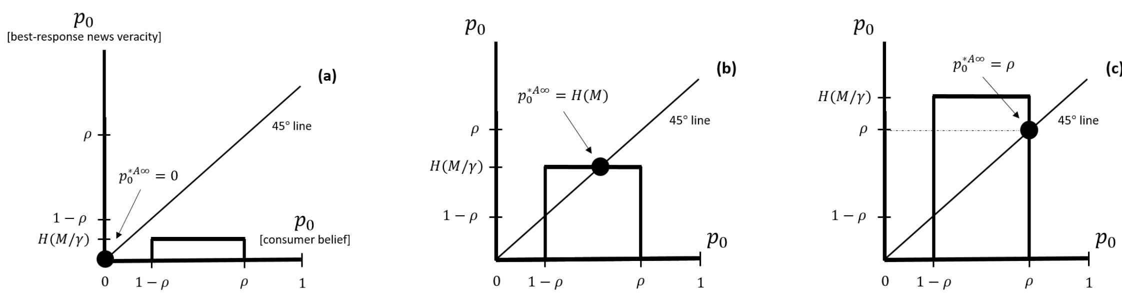

Figure 7: Illustration of limit-market equilibria when revenue is based on consumer actions, in three cases: (a) when $\gamma \geq \overline{\gamma}$ ("high costs"); (b) when $\gamma \in [\underline{\gamma}, \overline{\gamma})$ ("intermediate costs"); and (c) when $\gamma < \underline{\gamma}$ ("low costs").

Case #1: High costs and the dysfunctional limit-market. Suppose first that $\gamma \geq \overline{\gamma}$, pictured in Figure 7(a). Because $H(M/\gamma) \leq 1 - \rho$, news veracity must be strictly less than $1 - \rho$ in any equilibrium. Consumers never share in any equilibrium and, learning nothing from others, never act. True and false stories therefore all earn zero revenue, giving producers zero incentive to invest. Thus, the dysfunctional equilibrium is the unique equilibrium.

Case #2: Intermediate costs and the wise limit-market. Suppose that $\gamma \in [\underline{\gamma}, \overline{\gamma})$, pictured in Figure 7(b). The limit-market then has a functional equilibrium in which consumers use the filtering sharing rule $\tilde{\mathbf{z}} = (1, 0)$ and the maximal conceivable news veracity is realized, i.e., $p_0^{*A\infty} = H(M/\gamma)$. To see why, consider the case when $\gamma = \overline{\gamma} - \epsilon$ for some small ϵ, and suppose for a moment that consumers use the filtering sharing rule. Each neighbor's observed sharing behavior is then a binary random variable ("share/not share") with "share" occurring with probability $b\rho$ and $b(1-\rho)$, respectively, for true and false stories. By the Law of Large Numbers, a consumer who follows many others can therefore infer with high confidence whether any given story is true or false.

In the limit-market when following infinitely many others, consumers can perfectly discern which stories are true–and hence only act on true stories. True stories then earn M revenue while false stories earn zero revenue, causing producers to invest whenever $c_R < M$, the most possible, and inducing news veracity $H(M/\overline{\gamma})$. Finally, because γ is slightly below $\overline{\gamma}$, $H(M/\overline{\gamma})$ is slightly higher than $1-\rho$ and consumers find it optimal to use the filtering sharing rule; so, there is indeed an equilibrium in the limit-market in which producers invest maximally and consumers enjoy a "wisdom of the crowd."

Case #3: low costs and the high-veracity limit-market. Suppose that $\gamma < \underline{\gamma}$, pictured in Figure 7(c). In this case, the limit-market has a functional equilibrium with news veracity $p_0^{*A\infty} = \rho$.[26] In the sequence of finite-market equilibria converging to this limit-market equilibrium (as $d \to \infty$), consumers use a sharing rule of the form $\mathbf{z}^{*A}(d) = (1, z_F^{*A}(d))$ where $\lim_{d\to\infty} z_F^{*A}(d) = 1$. Because each consumer almost always shares, even after a bad private signal, each consumer in a large finite market views approximately fraction b of their neighbors sharing, for true and false stories. The exact number of sharing neighbors is informative, since true stories are slightly more likely to be shared, but such learning is limited even in the limit as consumers follow infinitely many others.

The amount that consumers are able to learn in equilibrium is determined by the need to provide producers with just enough incentive to invest that exactly fraction ρ of stories are true. In particular, for all large d, true-story visibility $V_T^*(d)$ must be exactly $H^{-1}(\rho)/M > 0$ more than false-story visibility $V_F^*(d)$. This requires each consumer not to encounter each true and false story with probability $1-V_T^*(d)$ and $1-V_F^*(d)$, respectively.

[26] An actions-supported limit market may have multiple functional equilibria. See Appendix C for a complete (implicit) characterization of all equilibria.

Theorem 2 summarizes:

Theorem 2. *In an actions-supported limit-market: (i) If $\gamma \geq \overline{\gamma}$, then the unique equilibrium is the dysfunctional equilibrium. (ii) If $\gamma \in [\underline{\gamma}, \overline{\gamma})$, then there is a unique functional equilibrium with $p_0^{*A\infty} = H(M/\gamma)$ and perfect learning by consumers. (iii) If $\gamma < \underline{\gamma}$, then there exists a functional equlibrium with news veracity $p_0^{*A\infty} = \rho$ and some (imperfect) learning by consumers.*

Impact of misinformation. Suppose that quantity $m \geq 0$ of misinformation is injected into a revenue-from-actions news market. The maximal conceivable news veracity falls from $H(M/\gamma)$ to $\frac{H(M/\gamma)}{1+m}$. Theorem 2 then applies with thresholds $\underline{\gamma}(m)$ and $\overline{\gamma}(m)$ defined implicitly by $\frac{H(M/\underline{\gamma}(m))}{1+m} = \rho$ and $\frac{H(M/\overline{\gamma}(m))}{1+m} = 1 - \rho$.

To visualize the impact of increased misinformation, suppose that bona fide producers have "low costs" ($H(M/\gamma) > \rho$) so that, absent any misinformation, limit-market equilibrium news veracity would be equal to ρ, as shown in Figure 7(c). Increasing the quantity of misinformation shifts down the news veracity that can be supported when consumers filter the news, so that the picture eventually shifts into the "wise market" case of Figure 7(b), and then into the "dysfunctional market" case of Figure 7(a). When there is an intermediate amount of misinformation, so that a "wise market" emerges, more true news is produced than when there was no misinformation at all and consumers are better able to discern which stories are true. In this way, misinformation campaigns can sometimes backfire and actually improve the performance of a news market.

Impact of gaslighting. Suppose that the precision of consumers' private signals falls from ρ to $\rho' \in (1/2, \rho)$, as in Section 2.3. Such a change shrinks the filtering region from $(1 - \rho, \rho)$ to $(1 - \rho', \rho')$, which can be seen visually as reducing the *width* of the step in Figure 7 but not changing its height. In the limit-market given any news veracity in $(1-\rho', \rho')$, consumers will still be able to perfectly discern which stores are true, maintaining producers' incentive to produce quantity $H(M/\gamma)$ of true news. Thus, if $H(M/\gamma) \in (1-\rho', \rho')$, then decreasing consumers' private-signal precision has no effect on equilibrium outcomes in the limit-market. By contrast, if $H(M/\gamma) \in (1-\rho', 1-\rho)$, then equilibrium outcomes will shift from the "wise market" case of Figure 7(b) to the "dysfunctional market" case of Figure 7(a), with equilibrium news veracity falling from $H(M/\gamma)$ to zero. Similarly, if $H(M/\gamma) \in (\rho', \rho)$, then equilibrium outcomes will shift from the "wise market" case to the "high-veracity market" case of Figure 7(c), with equilibrium news veracity falling from $H(M/\gamma)$ to ρ'.

4.2 Veracity Comparison: Revenue from Views vs. Actions

We find that whenever any true news is produced in a revenue-from-views limit-market, strictly more true news is produced in an otherwise-identical revenue-from-actions limit-market. Given our previous results, the argument is straightforward: Equilibrium news veracity in a views-supported limit market, denoted $p_0^{*V\infty}$, is anemic at best; by Proposition 4, either $p_0^{*V\infty} = 1-\rho$ or $p_0^{*V\infty} = 0$. Moreover, whenever $p_0^{*V\infty} = 1-\rho$, it must be the case that $H(M/\gamma)$, the maximal conceivable news veracity exceeds $1-\rho$. And then Theorem 2 implies $p_0^{*A\infty} = \min\{H(M/\gamma), \rho\} > 1-\rho$. [27]

Proposition 5. Either $p_0^{*V\infty} = p_0^{*A\infty} = 0$ or $p_0^{*A\infty} > p_0^{*V\infty}$.

Hence, when producers want consumers to actually believe their stories, social networks enhance producers' incentives to invest in quality news. Network connections both spread stories and allow followers to distinguish which news is true and which is false.

[27] Because all stories have visibility of at least b, the extra visibility of true stories $\Delta V \leq 1-b < 1$. Thus, even if consumers' signals were perfectly informative and only true stories were shared, producers' best-response news veracity would only be $H(M(1-b)/\gamma)$. The fact that $p_0^{*V\infty} = 1-\rho$ therefore implies $H(M(1-b)/\gamma) \geq 1-\rho$ and hence $H(M/\gamma) > 1-\rho$.

5 Conclusion

This paper analyzes the supply and demand for decision-relevant information, referred to as "news." Our model captures contemporary media markets in which consumers share stories over social networks, news producers cannot commit to quality, and news producers are paid when consumers view their stories or when consumers act on the basis of their stories. In each case, a dysfunctional equilibrium always exists in which no consumer shares any story and no producer invests in news truth. A unique functional equilibrium exists when producer costs are sufficiently low and/or consumers' private signals are sufficiently precise.

Social-network density has a non-monotonic effect on true-news production in views-supported news markets. Adding more links to a sparse network induces producers to invest more in news truth, since true stories are more likely to be shared. However, if consumers already follow many others, additional links favor the spread of false stories and hence reduce investment. Indeed, in the limit as consumers follow infinitely-many others, a views-supported news market either is dysfunctional or provides the bare minimum of true stories, just enough to induce consumers sometimes to share.

In contrast, when producers are paid for each consumer who acts on their story, dense social networks can support high news veracity. As long as the distribution of producer costs is neither too high nor too low, a "wisdom of the crowd" emerges when consumers follow infinitely-many others, allowing consumers to perfectly infer which stories are true and hence avoid acting on false information.

Any news market is vulnerable to misinformation, deep-fake technology, and gaslighting in the sense that, at sufficiently high levels, no functional equilibrium exists. However, smaller amounts of misinformation and smaller decreases in consumers' private abilities to discern true from false can actually prompt bona fide news producers to invest more in news production, since consumers become more cautious when sharing news stories.

This paper serves as a jumping-off point for the study of supply and demand in networked news markets. Several directions for future work could build on our analysis.

A natural next step would be to endogenize the social network, allowing consumers to decide how many people to follow, paying a cost for each social link. This possibility would have implications for the efficiency of media platforms. In an actions-supported news market, high-veracity news can arise in equilibrium if consumers are densely connected (Thm 2(iii)). But if stories are very likely to be true, consumers have little to gain by following others and hence little incentive to invest in social connections, potentially resulting in a sparse network that cannot support high news veracity. In this context, platforms such as Facebook and Twitter that make it easier for consumers to follow one another (reducing link costs) might indirectly promote higher-quality journalism.

On the supply side, natural next steps are to consider the industrial organization of news production and different business models. News producers recently have shifted toward subscription-based revenue (see e.g., New York Times 2015). Whoever controls a subscription channel then has an incentive to maximize its overall value to consumers, to increase subscribers' willingness to pay for channel access. However, subscriber engagement also drives advertiser and sponsor revenue[28] and these different revenue sources generate potentially competing incentives, in ways that deserve further study. For instance, a channel that earns its revenue only from subscribers might have an incentive to block readers from sharing content outside of its own walled garden, while one that also earns advertising and/or content-sponsor revenue might prefer to enable stories to be more widely shared by subscribers.

Newspapers and other news-distribution platforms can serve as intermediaries, screen-

[28]As the *New York Times* explained: "By focusing on subscribers, *The Times* will also maintain a stronger advertising business than many other publications. Advertisers crave engagement: readers who linger on content and who return repeatedly" (New York Times 2017).

ing the stories that consumers encounter and/or lending credibility to news producers. For instance, a politician with an opinion might post it directly on Twitter or some other online channel such as Medium that does not fact-check content if the goal is just to grab attention,[29] but submit it to a reputable paper such as the *Washington Post* for editorial review if the goal is to change minds.

News-distribution platforms can also create their own "news markets," by distinctively identifying the stories that consumers discover through their channels. For instance, at Facebook, a "curation team" consisting of journalists from partner news organizations decides which stories to highlight under the banner of "Today's Stories," creating a distinct news market with material re-published from original sources. Such curated channels could benefit consumers, by highlighting high-quality stories by high-quality producers.[30] However, consumers of such news might also share less judiciously, limiting how much others can learn from their sharing choices. In addition, a *dominant* curated channel might have anti-competitive and/or anti-democratic effects, if those curating the news seek to enhance the market power of existing producers and/or promote an ideological or partisan agenda.

Finally, future work could extend our analysis to allow for multidimensional investment. Producers in this present paper invest in a single unobservable characteristic ("truth") but, of course, producers also invest heavily in observable characteristics. Such investments directly affect consumers' incentives to share (e.g., consumers may want to

[29] Through a partnership with PolitiFact, Medium adds fact-check annotations to some posts after publication (PolitiFact 2015). This allows readers who encounter such stories *on Medium* to better assess which factual claims are true, akin in our model to providing an extra signal about news truth to all those who see the original broadcast, and may give politicians more incentive not to lie. However, to the extent that such claims are re-reported or spread by word of mouth without the extra annotations, falsehoods may still find their audience.

[30] Allcott, Gentzkow, and Yu (2019) found that, throughout 2017, user engagement with false content fell sharply on Facebook but continued rising on Twitter, suggesting that Facebook's efforts to combat misinformation after the 2016 election were effective. However, in September 2019, Facebook announced that it would not fact-check politicians' speech, exempting politicians' content and ads from a third-party fact-checking program used to assess other content (Constine 2019).

share funny or shocking content, even if they suspect it is untrue) and act, while also indirectly affecting those decisions by shaping consumers' beliefs about unobservable characteristics. For instance, suppose that a news producer can invest in the (unobservable) truth and/or (observable) appeal of its stories, where "appeal" increases consumers' payoff when sharing a story but has no effect on their action payoff. If the cost of increasing news appeal is small relative to the cost of news truth, producers paid for views may find it optimal to invest only in news appeal, leading to equilibrium outcomes in which all stories are false but appealing: widely shared because of their appeal but ineffective at driving action because no one believes them. By contrast, producers paid for actions may find it optimal to *disinvest* in news appeal as a way of increasing the return to their investments in news truth, as doing so can cause consumers to filter the news and thereby activate a "wisdom of the crowd."

References

Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi, "Spread of (Mis)information in Social Networks," *Games and Economics Behavior*, 2010, *70* (2), 194–227.

Allbright, Claire, "A Russian Facebook page organized a protest in Texas. A different Russian page launched the counterprotest," *Texas Tribune*, Nov 1 2017.

Allcott, Hunt, Matthew Gentzkow, and Chuan Yu, "Trends in the diffusion of misinformation on social media," *Research & Politics*, 2019, *6* (2), 2053168019848554.

Andreottola, Giovanni and Antoni-Italo de Moragas, "Scandals, Media Competition and Political Accountability," working paper 2018.

Banerjee, Abhijit, "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

_, **Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson**, "The Diffusion of Microfinance," *Science*, 2013, *341* (6144), 363–370.

Beamon, Benita M. and Clara Fernandes, "Supply-chain network configuration for product recovery," *Production Planning & Control*, 2004, *15* (3), 270–281.

Besley, Timothy and Andrea Prat, "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability," *American Economic Review*, 2006, *96* (3), 720–736.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Bimpikis, Kostas, Asuman Ozdaglar, and Ercan Yildiz, "Competitive targeted advertising over networks," *Operations Research*, 2016, *64* (3), 705–720.

_, **Douglas Fearing, and Alireza Tahbaz-Salehi**, "Multisourcing and Miscoordination in Supply Chain Networks," *Operations Research*, 2018, *66* (4), 1023–1039.

Bloch, Francis, Gabrielle Demange, and Rachel Kranton, "Rumors and Social Networks," *International Economic Review*, 2018, *59* (2), 421–448.

Boghardt, Thomas, "Soviet Bloc Intelligence and its AIDS disinformation campaign," *Studies in Intelligence*, 2009, *53* (4), 1–24.

Borchers, Callum, "This is a real news story about fake news stories," *Washington Post*, Nov 7 2016.

Chatterjee, Kalyan and Bhaskar Dutta, "Credibility and Strategic Learning in Networks," *International Economic Review*, 2016, *57* (3), 759–786.

Chen, Heng and Wing Suen, "Competition for Attention and News Quality," working paper 2019.

Constine, John, "Facebook promises not to stop politicians' lies & hate," *Techcrunch*, Sept 24 2019.

Del Vicario, Michela, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, , and Walter Quattrociocchi, "The spreading of misinformation online," *PNAS*, 2016, *113* (3), 554–559.

Dorpat, Theodore, *Gaslighting, the Double Whammy, Interrogation and Other Methods of Covert Control in Psychotherapy and Analysis*, Jason Aronson, Inc., 1996.

Ellis, Emma Grey, ""Epstein Didn't Kill Himself" and the Meme-ing of Conspiracy," *Wired*, Nov 15 2019.

Ellman, Matthew and Fabrizio Germano, "What do the Papers Sell? A Model of Advertising and Media Bias," *Economic Journal*, 2009, *119*, 680–704.

Ferrara, Emilio, "Disinformation and social bot operations in the run up to the 2017 French presidential election," *First Monday*, 2017, *22* (8).

Galeotti, Andrea and Sanjeev Goyal, "Influencing the Influencers: A Theory of Strategic Diffusion," *RAND Journal of Economics*, 2009, *40* (3), 509–532.

Galton, Francis, "Vox Populi," *Nature*, 1907, *75*, 450–451.

Gentzkow, Matthew and Jesse Shapiro, "Media Bias and Reputation," *Journal of Political Economy*, 2006, *114* (2), 280–316.

__ , **Edward Glaeser, and Claudia Goldin**, "The Rise of the Fourth Estate. How Newspapers Became Informative and Why It Mattered," in Edward Glaeser and Claudia Goldin, eds., *Corruption and Reform: Lessons from America's Economic History*, University of Chicago Press, 2006.

Goga, Oana, Giridhari Venkatadri, and Krishna P. Gummadi, "The doppelganger bot attack: Exploring identity impersonation in online social networks," *Proceedings of the 2015 Internet Measurement Conference*, 2015, pp. 141–153.

Grebe, Eduard and Nicoli Nattrass, "AIDS conspiracy beliefs and unsafe sex in Cape Town," *AIDS and Behavior*, 2012, *16* (3), 761–773.

Howell, Lee, "Global Risks 2013: Eighth Edition," World Economic Forum report 2013.

__ , "Only You Can Prevent Digital Wildfires," *New York Times*, Jan 8 2013.

Kroll, Andy, "Ready for Trump TV? Inside Sinclair Broadcasting's Plot to Take Over Your Local News," *Mother Jones*, Nov/Dec 2017.

Lim, Gabrielle, Etienne Maynier, John Scott-Railton, Alberto Fittarelli, Ned Moran, and Ron Deibert, "Burned After Reading: Endless Mayfly's Ephemeral Disinformation Campaign," U. Toronto working paper 2019.

Loftus, Elizabeth F., "Planting misinformation in the human mind: A 30-year investigation of the malleability of memory," *Learning & memory*, 2005, *12* (4), 361–366.

Manning, Martin J. and Herbert Romerstein, *Historical Dictionary of American Propaganda*, Greenwood Publishing Group, 2004.

Mills, Laura, "Russians fed conspiracy theories on Ukraine crash," *Associated Press*, July 22 2014.

New York Times, "Our Path Forward," *New York Times*, Oct 7 2015.

_ , "Journalism That Stands Apart: The Report of the 2020 Group," *New York Times*, Jan 2017.

Oyeyemi, Sunday Oluwafemi, Elia Gabarron, and Rolf Wynn, "Ebola, Twitter, and misinformation: a dangerous combination?," *BMJ*, 2014, *349*, g6178.

Peng, Xiao-Long, Xin-Jian Xu, Xinchu Fu, and Tao Zhou, "Vaccination intervention on epidemic dynamics in networks," *Physical Review E*, 2013, *87* (2), 022813.

PolitiFact, "PolitiFact, Knight team up to fact-check the 2016 candidates on Medium," press release Dec 18 2015. available at https://www.politifact.com/truth-o-meter/article/2015/dec/18/politifact-knight-team-fact-check-2016-candidates-/.

Qazvinian, Vahed, Emily Rosengren, Dragomir R. Radev, and Qiaozhu Mei, "Rumor has it: Identifying misinformation in microblogs," *Proceedings of the conference on empirical methods in natural language processing*, 2011, pp. 1589–1599.

Read, Andrew F., Susan J. Baigent, Claire Powers, Lydia B. Kgosana, Luke Blackwell, Lorraine P. Smith, David A. Kennedy, Stephen W. Walkden-Brown, , and Venugopal K. Nair, "Imperfect vaccination can enhance the transmission of highly virulent pathogens," *PLoS Biology*, 2015, *13* (7), e1002198.

Redilicki, Bartosz, "Spreading Lies," Cambridge Working Paper Economics: 1747 2017.

Romerstein, Herbert, "Disinformation as a KGB Weapon in the Cold War," *Journal of Intelligence History*, 2001, *1* (1), 54–67.

Safire, William, "ON LANGUAGE; The Glasnost Dangle," *New York Times Magazine*, Apr 30 1989.

Samuelson, Larry and Jianbo Zhang, "Evolutionary stability in asymmetric games," *Journal of Economic Theory*, 1992, *57* (2), 363–391.

Shao, Chengcheng, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer, "Hoaxy: A platform for tracking online misinformation," *Proceedings of the 25th international conference companion on world wide web*, 2016, pp. 745–750.

_ , _ , Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer, "The spread of low-credibility content by social bots," *Nature Communications*, 2018, *9* (1), 4787.

Shields, Todd and Jennifer A. Dlouhy, "Fox News advertisers get a direct line to the viewer in chief," *Bloomberg News*, Aug 28 2019.

Vosoughi, Soroush, Deb Roy, and Sinan Aral, "The spread of true and false news online," *Science*, 2018, *359* (6380), 1146–1151.

Zhu, Bi, Chuansheng Chen, Elizabeth F. Loftus, Chongde Lin, Qinghua He, Chunhui Chen, He Li, Gui Xue, Zhonglin Lu, and Qi Dong, "Individual differences in false memory from misinformation: Cognitive factors," *Memory*, 2010, *18* (5), 543–555.

Online Appendices

Appendix A provides details on dynamic stability of equilibria. Appendix B contains proofs and extended analysis for the market in which producers are paid for views of their news stories (Sections 2-3. Appendix C contains proofs and extended analysis for the market in which producers are paid for consumers' actions (Section 4).

A Dynamic stability

Our analysis focuses on equilibria that are "dynamically stable." In this appendix, we define our stability concept and identify conditions under which equilibria in a views-supported news market are dynamically stable or unstable.

Definition 1 (Perturbed best-response news veracity). *For any pair of sharing rules* $\mathbf{z}, \hat{\mathbf{z}}$*, let* $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}}) = p_0(\mathbf{z}(1-\epsilon) + \hat{\mathbf{z}}\epsilon)$ *denote the "perturbed best-response news veracity" when consumers use sharing rule* $\mathbf{z}(1-\epsilon) + \hat{\mathbf{z}}\epsilon$.

Definition 2 (Dynamic stability). *A sharing rule* $\mathbf{z}$ *is "dynamically stable" (or simply "stable") if, for all* $\hat{\mathbf{z}}$ *and all* $\epsilon \approx 0$, $\mathbf{z}$ *is a strictly better response for consumers than* $\hat{\mathbf{z}}$ *given news veracity* $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}})$.[31] *Similarly,* $\mathbf{z}$ *is "dynamically unstable" (or simply "unstable") if there exists* $\hat{\mathbf{z}}$ *such that, for all* $\epsilon \approx 0$, $\hat{\mathbf{z}}$ *is a strictly better response for consumers than* $\mathbf{z}$ *given news veracity* $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}})$.

Lemma A1. *Consider a views-supported news market. (i) Any equilibrium with news veracity* $p_0 \notin \{1-\rho, \rho\}$ *is dynamically stable. (ii) Any equilibrium with news veracity* $p_0 = \rho$ *is dynamically stable. (iii) Any equilibrium with news veracity* $p_0 = 1-\rho$ *and*

[31] Implicit in this definition is a simplifying assumption that producers adapt immediately to any change in consumers' sharing strategies while consumers adapt gradually over time to changes in producers' investment strategies. However, this is not essential. Our results hold under any monotone co-adaptation dynamics (Samuelson and Zhang 1992); straightforward details omitted to save space.

sharing rule of the form $(z_T, 0)$ *is dynamically stable or unstable if* $z_T^* > \overline{z}_T$ *or* $z_T^* \leq \overline{z}_T$*, respectively, where* $(\overline{z}_T, 0)$ *is the sharing rule that maximizes producers' incentive to invest (Lemma 2(iv)).*

Proof: Consider any equilibrium with news veracity $p_0 \notin \{1-\rho, \rho\}$. Consumers have a unique best response, $\mathbf{z}(p_0)$, equal to $(1,1)$ if $p_0 > \rho$, $(1,0)$ if $p_0 > (1-\rho, \rho)$, or $(0,0)$ if $p_0 < 1-\rho$. For any given $\hat{\mathbf{z}} \neq \mathbf{z}(p_0)$ and $\epsilon \approx 0$, $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}}) \approx p_0$. (By equation (3), $p_0(\mathbf{z})$ is continuous in $\mathbf{z}$, a fact we use repeatedly throughout the proof.) Thus, for all $\epsilon \approx 0$, $\mathbf{z}(p_0)$ continues to be consumers' unique best response; in particular, $\mathbf{z}(p_0)$ is a better reply than $\hat{\mathbf{z}}$ and hence $\hat{\mathbf{z}}$ cannot successfully invade. This completes the proof of (i).

Next, consider any equilibrium with news veracity equal to ρ. Consumers strictly prefer to share given signal $s_i = T$ and are indifferent whether to share given signal $s_i = F$; the equilibrium sharing rule must be $\mathbf{z} = (1, z_F)$ for some $z_F \in [0,1]$. For any $\hat{\mathbf{z}} \neq \mathbf{z}$, perturbed news veracity $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}}) \approx \rho$, given which consumers still strictly prefer to share when $s_i = T$ and are approximately indifferent whether to share when $s_i = F$. The rest of the proof that $\hat{\mathbf{z}}$ cannot successfully emerge has three steps. First, consider any $\hat{\mathbf{z}}$ with $\hat{z}_T < 1$. After receiving signal $s_i = T$ (probability $\Pr(s_i = T | p_0 = \rho) > 0$), a consumer who shares with probability $\hat{z}_T$ loses approximately $(1-\hat{z}_T)\pi^S(2\rho - 1) > 0$ relative to the best response of always sharing. By contrast, after receiving signal $s_i = F$, the benefit (if any) that a consumer gets by sharing with probability $\hat{z}_F$ rather than probability z_F goes to zero as ϵ goes to zero. Overall, then, $\hat{\mathbf{z}}$ is a worse reply than $\mathbf{z}$ for all small enough ϵ. Second, consider any $\hat{\mathbf{z}} = (1, \hat{z}_F)$ with $\hat{z}_F < z_F$, inducing perturbed news veracity $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}}) = p_0(1, z_F - \epsilon(z_F - \hat{z}_F))$. Because $p_0 = (1, z_F)$ is strictly decreasing in z_F (Lemma 2(ii)), $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}}) > \rho$ and consumers have a strict incentive to share after signal $s_i = F$. Since $\hat{z}_F < z_F$, $\hat{\mathbf{z}}$ is therefore a worse reply than $\mathbf{z}$. Third and finally, consider any $\hat{\mathbf{z}} = (1, \hat{z}_F)$ with $\hat{z}_F > z_F$, inducing perturbed news veracity

$p_0^\epsilon(\mathbf{z};\hat{\mathbf{z}}) = p_0(1, z_F+\epsilon(\hat{z}_F-z_F))$. Because $p_0(1, z_F)$ is strictly decreasing in z_F, $p_0^\epsilon(\mathbf{z};\hat{\mathbf{z}}) < \rho$, giving consumers a strict incentive not to share after signal $s_i = F$ and making $\hat{\mathbf{z}}$ a worse reply than $\mathbf{z}$ since $\hat{z}_F > z_F$. This completes the proof of (ii).

Finally, consider any equilibrium with news veracity equal to $1-\rho$. Consumers are indifferent whether to share after getting a positive signal $s_i = T$ and strictly prefer not to share after a negative signal $s_i = F$; the equilibrium sharing rule must be $\mathbf{z} = (z_T, 0)$ for some $z_T \in [0, 1]$. One can easily show that any $\hat{\mathbf{z}}$ with $\hat{z}_F > 0$ cannot emerge; so, we will only consider potential strategies of the form $\hat{\mathbf{z}} = (\hat{z}_T, 0)$. Suppose first that $z_T \leq \overline{z}_T$ and consider the perturbing sharing rule $\hat{\mathbf{z}} = (0, 0)$, inducing news veracity $p_0(z_T - \epsilon z_T, 0)$. By Lemma 2(iii), $p_0(z_T, 0)$ is strictly increasing over the range $[0, \overline{z}_T)$; so, $p_0(z_T - \epsilon z_T, 0) < p_0(z_T, 0) = p_0 = 1 - \rho$. Since consumers have a strict incentive *not* to share given private signal $s_i = F$ after the perturbation, the equilibrium is dynamically unstable. Suppose next that $z_T > \overline{z}_T$ and consider any $\hat{\mathbf{z}} = (\hat{z}_T, 0)$. By Lemma 2(iii), $p_0(z_T, 0)$ is strictly decreasing over the range $(\overline{z}_T, 1]$; so, $p_0^\epsilon(\mathbf{z};\hat{\mathbf{z}}) > 1-\rho$ whenever $\hat{z}_T < z_T$ (making $\mathbf{z}$ a better reply than $\hat{\mathbf{z}}$) and $p_0^\epsilon(\mathbf{z};\hat{\mathbf{z}}) < 1 - \rho$ whenever $\hat{z}_T > z_T$ (again making $\mathbf{z}$ a better reply than $\hat{\mathbf{z}}$). We conclude in this case that the equilibrium is dynamically stable. This completes the proof of (iii). $\square$

B Proofs and extensions: revenue from views

This appendix provides formal proofs for the results in Sections 2-3 and analyzes several extensions in the case when producers are paid for views.

Lemma 1 follows directly arguments provided in the main text. Its proof is omitted to save space.

B.1 Proofs

B.1.1 Proof of Lemma 2

Parts (i,v). By equation (3), $p_0(z_T, z_F) = H(M\Delta V(z_T, z_F))$, where

$$\Delta V(z_T, z_F) = (1-b)\left((1-b((1-\rho)z_T + \rho z_F))^d - (1-b(\rho z_T + (1-\rho)z_F))^d\right) \quad \text{(A1)}$$

by equations (1,2). Recall that, by assumption, the producer's reporting cost has c.d.f. $H(\cdot)$ and atomless support; so, $H(0) = 0$, $H(c_R) > 0$ for all $c_R > 0$, and $H(c_R)$ is continuous in c_R. Part (i) now follows immediately from the fact that $\Delta V(z_T, z_F) = 0$ when $z_T = z_F$ and $\Delta V(z_T, z_F) > 0$ when $z_T > z_F$. Part (v) is also immediate, following from the continuity of $H(\cdot)$ and the easily-checked continuity of $\Delta V(\cdot)$.

Part (ii). Define $x(z_T) = \frac{\Delta V(z_T,0)}{1-b}$. To prove part (ii), it suffices to show that $x(z_T)$ is strictly increasing in z_T over the interval $[0, \overline{z}_T]$ and strictly decreasing in z_T over $[\overline{z}_T, 1]$ for some $\overline{z}_T \in (0, 1]$. Note that

$$x'(z_T) = db\left(\rho(1-b\rho z_T)^{d-1} - (1-\rho)(1-b(1-\rho)z_T)^{d-1}\right). \quad \text{(A2)}$$

Suppose first that $d = 1$. Since $x'(z_T) = b(2\rho - 1) > 0$, $x(z_T)$ is strictly increasing over the whole interval $z_T \in [0, 1]$, establishing the desired result with respect to $\overline{z}_T = 1$. Suppose next that $d \geq 2$. $x'(z_T) > 0$ if and only if $\frac{\rho}{1-\rho} > \left(\frac{1-b(1-\rho)z_T}{1-b\rho z_T}\right)^{d-1}$ which, after re-arranging, can be written as $z_T < \hat{z}_T \equiv \frac{\left(\frac{\rho}{1-\rho}\right)^{\frac{1}{d-1}} - 1}{b\left(\rho\left(\frac{\rho}{1-\rho}\right)^{\frac{1}{d-1}} - (1-\rho)\right)}$. So, $x(z_T)$ is strictly increasing in z_T over the interval $[0, \min\{\hat{z}_T, 1\}]$ and, if $\hat{z}_T < 1$, strictly decreasing over the interval $[\hat{z}_T, 1]$, establishing the desired result with respect to $\overline{z}_T \equiv \min\{\hat{z}_T, 1\}$.

Part (iii). Define $w(z_F) = \frac{\Delta V(1,z_F)}{1-b}$. To prove part (iii), it suffices to show that $w(z_F)$ is strictly decreasing in z_F. Note that

$$w'(z_F) = db\left((1-\rho)(1-b(\rho+(1-\rho)z_F))^{d-1} - \rho(1-b(1-\rho+\rho z_F))^{d-1}\right). \tag{A3}$$

Since $\rho > \frac{1}{2}$ and $z_F \leq 1$, $1-\rho < \rho$ and $\rho + (1-\rho)z_F \geq 1-\rho+\rho z_F$; thus, $w'(z_F) < 0$ for all $z_F \in [0,1]$, as desired.

Part (iv). Define $y(z_T, z_F) = \frac{\Delta V(z_T, z_F)}{1-b}$. To prove part (iv), it suffices to show that $y(\overline{z}_T, 0) \geq y(z_T, z_F)$ for all $z_T, z_F \in [0,1]$. First, note that $y(z_T, z_F) \leq 0$ whenever $z_T \leq z_F$ but $y(\overline{z}_T, 0) > 0$; so, we may restrict attention to sharing rules with $z_T > z_F$. Next, note that $\rho > \frac{1}{2}$ and $z_T > z_F$ implies $1-\rho < \rho$ and $\rho z_T + (1-\rho)z_F > (1-\rho)z_T + \rho z_F$; thus,

$$\frac{\partial y(z_T, z_F)}{\partial z_F} = db\left((1-\rho)(1-b(\rho z_T + (1-\rho)z_F)^{d-1} - \rho(1-b((1-\rho)z_T + \rho z_F)^{d-1}\right) < 0 \tag{A4}$$

Finally, $y(\overline{z}_T, 0) \geq y(z_T, 0)$ for all $z_T \in [0,1]$ by definition of $\overline{z}_T$. We conclude $y(\overline{z}_T, 0) \geq y(z_T, z_F)$ for all z_T, z_F, as desired. □

B.1.2 Proof of Theorem 1

Here we characterize all Bayesian Nash equilibria (BNE) in which consumers use the same sharing rule.[32] These "symmetric BNE" may or may not be dynamically stable, and may or may not exhibit positive investment by producers. In the main text and here, we refer to any symmetric BNE that is dynamically stable and has zero producer investment as "dysfunctional," and any symmetric BNE that is dynamically stable and has positive producer investment as "functional."

[32] All BNE in which consumers use different sharing rules are characterized in Section C.1, when we consider the most general model in which producers may be paid for views and/or for actions.

Dysfunctional equilibrium exists. Suppose for a moment that all stories were false. Consumers would find it optimal to never share, causing all stories to earn the same revenue (Mb) and giving producers zero incentive to invest in news truth. Because producing true stories is costly (by assumption[33]), producers then find it optimal only to produce false stories. Thus, a BNE always exists in which consumers never share and producers never invest. Moreover, this equilibrium is dynamically stable by Lemma A1(i).

Equilibrium news veracity $p_0^* \not> \rho$. BNE do not exist with news veracity greater than ρ. In such an equilibrium, consumers would find it optimal to always share, causing all stories to generate equal revenue. But then producers would have zero incentive to invest, so that news veracity must be zero, a contradiction.

Equilibrium news veracity $p_0^* \notin (0, 1-\rho)$. BNE do not exist with news veracity between 0 and $1-\rho$. In such an equilibrium, consumers would find it optimal to never share, causing all stories to generate equal revenue and giving producers zero incentive to invest, a contradiction.

Proof of part (a). Suppose that $\gamma \leq \gamma_1$. Since $\tilde{p}_0(\gamma) \geq \rho$, no BNE exists with news veracity between $1-\rho$ and ρ. In such an equilibrium, consumers would find it optimal to filter the news, causing true stories to have extra visibility $\Delta V(\tilde{\mathbf{z}})$ and giving producers incentive to invest with probability $\tilde{p}_0 \geq \rho$, a contradiction. However, BNE do exist with news veracity equal to ρ and equal to $1-\rho$.

Given news veracity $p_0 = \rho$, consumers find it optimal to use sharing rules of the form $(1, z_F)$ for all $z_F \in [0,1]$. If $z_F = 0$, then producers will respond by investing with probability $p_0(1,0) = \tilde{p}_0 \geq \rho$ while, if $z_F = 1$, then best-response news veracity

[33] Our analysis and results can be easily extended to a setting in which producers' reporting cost is zero (or negative) with probability $\alpha > 0$. Results change in a qualitative way if $\alpha > 1-\rho$ since, in that case, enough true news is always produced to induce consumers to share after a good private signal. On the other hand, so long as $\alpha \leq 1-\rho$, a "minimally-functional equilibrium" continues to exist in which consumers never share any stories and producers only publish true stories when doing so is costless.

$p_0(0,0) = 0 < \rho$. Since $p_0(1, z_F)$ is continuous and strictly decreasing in z_F (Lemma 2), there exists a unique $z_F^* \in [0,1)$ such that $p_0(1, z_F^*) = \rho$. We conclude that there exists a unique symmetric BNE with news veracity ρ. Moreover, this equilibrium is dynamically stable by Lemma 2(i) and hence a functional equilibrium.

Given news veracity $p_0 = 1 - \rho$, consumers find it optimal to use sharing rules of the form $(z_T, 0)$ for all $z_T \in [0,1]$. Since $p(1,0) > 1 - \rho$ and $p_0(z_T, 0)$ is strictly decreasing over $z_T \in [\overline{z}_T, 1]$, $p_0(z_T, 0) > 1 - \rho$ for all $z_T \in [\overline{z}_T, 1]$. However, because $p_0(0,0) = 0$ and $p_0(z_T, 0)$ is continuous and strictly increasing over $z_T \in [0, \overline{z}_T]$ (Lemma 2), there exists a unique $z_T^* \in (0, \overline{z}_T)$ such that $p_0(1, z_F^*) = 1 - \rho$. We conclude that there exists a unique symmetric BNE with this news veracity level. However, because $0 < z_T^* < \overline{z}$, this equilibrium is dynamically unstable by Lemma 2(i) and hence is not a functional equilibrium.

Proof of part (b). Suppose that $\gamma \in (\gamma_1, \gamma_2)$ so that $\tilde{p}_0(\gamma) \in (1 - \rho, \rho)$. If consumers filter the news, then optimal investment induces best-response news veracity $\tilde{p}_0(\gamma)$, given which consumers find it optimal to filter the news. Hence, a symmetric BNE exists in which consumers use sharing rule $\tilde{\mathbf{z}} = (1, 0)$ and news veracity $p_0^* = \tilde{p}_0$. This equilibrium is dynamically stable by Lemma A1(i) and hence is a functional equilibrium.

To establish that this is the unique functional equilibrium, we rule out all other possibilities. First, no symmetric BNE exists in which consumers use sharing rule of the form $(1, z_F^*)$ for some $z_F^* > 0$. In such an equilibrium, it must be that $p_0^* \geq \rho$ but, since $p_0(1, z_F)$ is strictly decreasing in z_F (Lemma 2), producers' best response results in news veracity less than $\tilde{p}_0$, a contradiction. Second, consider any symmetric BNE in which consumers use a sharing rule of the form $(z_T^*, 0)$ for some $z_T^* \in (0,1)$. In such an equilibrium, it must be that news veracity $p_0^* = 1 - \rho$, so that consumers are indifferent whether to share after a good private signal. But since $p_0(1,0) > 1 - \rho$ and $p_0(1, z_F)$

is strictly increasing over $z_F \in [\overline{z}_T, 1]$, it must be that $z_F^* < \overline{z}_T$. Any such equilibrium must therefore be dynamically unstable (Lemma A1(iii)) and hence not a functional equilibrium.

Proof of part (c). Suppose that $\gamma \in [\gamma_2, \gamma_3)$ so that $\tilde{p}_0(\gamma) \leq 1 - \rho < \overline{p}_0(\gamma)$. No BNE exists with news veracity $p_0^* > 1 - \rho$ since, if it did, consumers who find it optimal to use a sharing rule of the form $(1, z_F^*)$ for some $z_F^* \in [0, 1]$ and best-response news veracity $p_0(1, z_F^*) \leq \tilde{p}_0 \leq 1 - \rho$, a contradiction. However, symmetric BNE do exist with news veracity $p_0^* = 1 - \rho$, and one of these is a functional equilibrium.

Recall that $p_0(z_T, 0)$ is strictly increasing in z_T over the range $[0, \overline{z}_T]$, rising from $p_0(0, 0) = 0$ to $p_0(\overline{z}_T, 0) = \overline{p}_0 > 1 - \rho$, and then strictly decreasing in z_T over the range $[\overline{z}_T, 1]$, falling from $\overline{p}_0$ to $p_0(1, 0) = \tilde{p}_0 \leq 1 - \rho$. Thus, there exist exactly *two* levels $z_T^1 \in (0, \overline{z}_T)$ and $z_T^2 \in (\overline{z}_T, 1]$ such that $p_0(z_T^1, 0) = p_0(z_T^2, 0) = 1 - \rho$, and hence two symmetric BNE with news veracity equal to $1 - \rho$. Of these two equilibria, Lemma A1(iii) implies that the one with more sharing is dynamically stable, while the one with less sharing is dynamically unstable. We conclude that there is a unique functional equilibrium, in which consumers use sharing rule $(z_T^1, 0)$ and news veracity equals $1 - \rho$.

Proof of part (d). When $\gamma > \gamma_3$, $\overline{p}_0(\gamma) < 1 - \rho$ and news veracity must be less than $1 - \rho$ in any BNE. But then consumers must never share, implying that producers must never invest; so, the dysfunctional equilibrium is the unique BNE. When $\gamma = \gamma_3$, $\overline{p}_0(\gamma) = 1 - \rho$, meaning that the best-response news veracity is strictly less than $1 - \rho$ unless consumers use the sharing rule $(\overline{z}_T, 0)$, in which case producers will invest just often enough to support news veracity equal to $1 - \rho$. We conclude that no symmetric BNE exists with news veracity greater than $1 - \rho$ but that a symmetric BNE *does* exist in which consumers use sharing rule $\mathbf{z}^* = (\overline{z}_T, 0)$ and news veracity $p_0^* = 1 - \rho$. However, this equilibrium is dynamically unstable by Lemma A1(iii); so, no functional equilibrium exists. □

B.1.3 Proof of Proposition 1

Equilibrium outcomes in a market with quantity m of misinformation and c.d.f. $H(c_R)$ for bona fide reporters' cost are identical to those in a market without misinformation and c.d.f. $H(c_R; m) \equiv \frac{H(c_R)}{1+m}$. Consequently, the maximal news veracity $\overline{p}_0(m) = \frac{\overline{p}_0}{1+m}$ and filtering news veracity $\tilde{p}_0(m) = \frac{\tilde{p}_0}{1+m}$. Under this interpretation, Proposition 1 follows fairly immediately from Theorem 1.

Proof of (i). If $m \geq \overline{m}$, then $\overline{p}_0(m) \leq 1 - \rho$. By the proof of Theorem 1(d), the dysfunctional equilibrium is the unique BNE. On the other hand, if $m < \overline{m}$, then $\overline{p}_0(m) > 1 - \rho$ and a unique functional equilibrium exists by Theorem 1(a-c).

Proof of (ii-iii). When $m \in [0, m_1]$, $\tilde{p}_0(m) \geq \rho$ and so news veracity $p_0^*(m) = \rho$ by Theorem 1(a); over this range, news veracity $p_0^*(m)$ is constant while true-news volume $(1+m)p_0^*(m) = (1+m)\rho$ is strictly increasing in m. When $m \in [m_1, m_2]$, $\tilde{p}_0(m) \in [1-\rho, \rho]$ and so news veracity $p_0^*(m) = \tilde{p}_0(m) = \frac{\tilde{p}_0}{1+m}$ by Theorem 1(b); over this range, news veracity $p_0^*(m)$ is strictly decreasing in m while true-news volume $(1+m)p_0^*(m) = \tilde{p}_0$ is constant. When $m \in [m_2, \overline{m})$, $\tilde{p}_0(m) \leq 1 - \rho M \overline{p}_0(m)$ and so news veracity $p_0^*(m) = 1 - \rho$ by Theorem 1(c); over this range, news veracity $p_0^*(m)$ is constant in m while true-news volume $(1+m)p_0^*(m) = (1+m)(1-\rho)$ is constant.

Putting this all together, news veracity $p_0^*(m)$ is non-increasing over the range $m < \overline{m}$ and strictly decreasing for $m \in [m_1, m_2]$, while true-news volume $(1+m)p_0^*(m)$ is non-decreasing over the range $m < \overline{m}$ and strictly increasing for $m \in [0, m_1] \cup [m_2, \overline{m})$, as desired. □

B.1.4 Proof of Proposition 2

Let $V_T(\mathbf{z};\rho)$ and $V_F(\mathbf{z};\rho)$ denote the visibility of true and false stories, respectively, given sharing rule $\mathbf{z}$ and signal precision ρ. By equations (1-2), $V_T(\mathbf{z};1/2) = V_F(\mathbf{z};1/2)$ and

$$\frac{\mathrm{d}V_T(\mathbf{z})}{\mathrm{d}\rho} = db(1-b)(z_T - z_F)\left(1 - b(\rho z_T + (1-\rho)z_F)\right)^{d-1} \tag{A5}$$

$$\frac{\mathrm{d}V_F(\mathbf{z})}{\mathrm{d}\rho} = -db(1-b)(z_T - z_F)\left(1 - b((1-\rho)z_T + \rho z_F)\right)^{d-1} \tag{A6}$$

Note that, so long as consumers are more likely to share after a good private signal ($z_T > z_F$), $V_T(\mathbf{z};\rho)$ is strictly increasing in ρ while $V_F(\mathbf{z};\rho)$ is strictly decreasing; thus, best-response news veracity $p_0(z_T, z_F;\rho) = H\left(M(V_T(\mathbf{z};\rho) - V_F(\mathbf{z};\rho))\right)$ is also strictly increasing in ρ for all such sharing rules. In particular, $\tilde{p}_0(\rho) \equiv p_0(1,0;\rho)$ and $\overline{p}_0(\rho) \equiv \max_{z_T\in(0,1]} p_0(z_T,0;\rho)$ are each strictly increasing in ρ.

Define signal-precision threshold $\overline{\rho}$ implicitly by $\tilde{p}_0(\overline{\rho}) = 1-\rho$, or $\overline{\rho} = 1$ if $\tilde{p}_0(1) \leq 1-\rho$. Similarly, define $\underline{\rho}$ implicitly by $\overline{p}_0(\underline{\rho}) = 1-\rho$, or $\underline{\rho} = 1$ if $\overline{p}_0(1) \leq 1-\rho$.

Proposition 2 now follows directly from Thm 1. *Part (i):* If $\rho \leq \underline{\rho}$, then $\overline{p}_0(\rho) \leq 1-\rho$ and hence $p_0^*(\rho) = 0$ by Thm 1(iv). *Part (ii):* If $\underline{\rho} < \rho \leq \overline{\rho}$, then $\overline{p}_0(\rho) > 1-\rho \geq \tilde{p}_0(\rho)$ and hence $p_0^*(\rho) = 1-\rho$ by Thm 1(iii). Note that $p_0^*(\rho)$ is strictly decreasing in ρ over this range. *Part (iii):* If $\rho > \overline{\rho}$, then $\tilde{p}_0(\rho) > 1-\rho$ and hence $p_0^*(\rho) = \min\{\tilde{p}_0(\rho),\rho\} > 1-\rho$ by Thm 1(i-ii). Note that $p_0^*(\rho)$ is strictly increasing in ρ over this range, since both $\tilde{p}_0(\rho)$ and ρ are strictly increasing in ρ. $\square$

B.1.5 Proof of Proposition 3

It suffices to show that the filtering news veracity $\tilde{p}_0(d)$ is single-peaked in d. To see why, recall by Thm 1 that $p_0^*(d) > 1-\rho$ if and only if $\tilde{p}_0(d) > 1-\rho$ and that, when this is the case, $p_0^*(d) = \min\{\tilde{p}_0(d),\rho\}$. $\tilde{p}_0(d)$ being single-peaked means that $\tilde{p}_0(d) > 1-\rho$

if and only if $\underline{d} < d < \overline{d}$; thus, $p_0^*(d) > 1-\rho$ if and only if $\underline{d} < d < \overline{d}$, establishing Proposition 3(i). Moreover, $p_0^*(d) = \min\{\tilde{p}_0(d), \rho\}$ when $d \in \left(\underline{d}, \overline{d}\right)$ means that $p_0^*(d)$ is also single-peaked over this range and maximized at $d = \tilde{d}$, establishing Proposition 3(ii-iii).[34]

By equation (3), $\tilde{p}_0(d) = H(M\Delta V(\tilde{\mathbf{z}}; d))$. Note that $\tilde{p}_0(d)$ depends on d only through $\Delta V(\tilde{\mathbf{z}}; d)$ and that $\tilde{p}_0(d)$ is strictly increasing in $\Delta V(\tilde{\mathbf{z}}; d)$. Thus, it suffices to show that $\Delta V(\tilde{\mathbf{z}}; d)$, the extra visibility of true stories when consumers filter the news, is itself single-peaked. By equations (1,2), $\Delta V(\tilde{\mathbf{z}}; d) = (1-b)\left((1-b(1-\rho))^d - (1-b\rho)^d\right)$; so,

$$\frac{\mathrm{d}\Delta V(\tilde{\mathbf{z}}; d)}{\mathrm{d}d} = (1-b)\left(\ln(1-b(1-\rho))(1-b(1-\rho))^d - \ln(1-b\rho)(1-b\rho)^d\right). \qquad \text{(A7)}$$

Re-arranging terms, we conclude that $\frac{\mathrm{d}\Delta V(\tilde{\mathbf{z}};d)}{\mathrm{d}d} \gtreqless 0$ if and only if

$$\left(\frac{1-b\rho}{1-b(1-\rho)}\right)^d \gtreqless \frac{\ln(1-b(1-\rho))}{\ln(1-b\rho)} \in (0,1). \qquad \text{(A8)}$$

Since $\rho > \frac{1}{2}$, $\frac{1-b\rho}{1-b(1-\rho)} < 1$ and the left-hand-side of (A8) is exponentially decreasing in d, while the right-hand-side of (A8) does not depend on d. We conclude that $\Delta V(\tilde{\mathbf{z}}; d)$ is strictly increasing in d up to some critical level and strictly decreasing thereafter; so, $\Delta V(\tilde{\mathbf{z}}; d)$ is single-peaked in d, as desired. □

B.1.6 Proof of Proposition 4

Let $\underline{\rho}^\infty = \lim_{d\to\infty} \underline{\rho}(d)$ where, as in Proposition 2, $\underline{\rho}(d)$ is defined implicitly by the condition $\overline{p}_0(d, \underline{\rho}(d)) = 1 - \underline{\rho}(d)$ and $\overline{p}_0(d,\rho) = H\left(M\Delta V(\overline{z}_T(d,\rho), 0; d, \rho)\right)$ is the maximal best-response news veracity given social connectedness d and signal precision ρ.

[34] The minimum of any two single-peaked functions is single-peaked; so, $\tilde{p}_0(d)$ being single-peaked implies that $p_0^*(d) = \min\{\tilde{p}_0(d), \rho\}$ is single-peaked. Similarly, because $\max_d \min\{\tilde{p}_0(d), \rho\} = \min\{\max_d \tilde{p}_0(d), \rho\}$, $\tilde{p}_0(d)$ being maximized at $\tilde{d}$ implies that $p_0^*(d)$ is maximized at $\tilde{d}$.

Proof of part (i): Suppose that $\rho < \underline{\rho}^{\infty}$. For all sufficiently large d, $\overline{p}_0(d,\rho) < 1-\rho$ (since $\rho < \underline{\rho}(d)$) and hence $p_0^*(d,\rho) = 0$ (by Thm 1(iv)); so, $p_0^{*\infty} = 0$ and the limit-market is dysfunctional.

Proof of part (ii): Suppose now that $\rho > \underline{\rho}^{\infty}$, so that $\overline{p}_0(d,\rho) > 1-\rho$ for all sufficiently large d. Since $\tilde{p}_0(d,\rho)$ is less than $1-\rho$ for all sufficiently large d (by Proposition 3), $p_0^*(d,\rho) = 1-\rho$ (by Thm 1(iii)) for all sufficiently large d; so, $p_0^{*\infty} = 1-\rho$.

To complete the proof, it remains to show that consumers cannot perfectly discern which stories are true from others' equilibrium sharing behavior in the $d \to \infty$ limit. For each large d, let $(z_T^*(d,\rho), 0)$ denote consumers' sharing rule in the dynamically-stable symmetric BNE with news veracity $1-\rho$. Define shorthand $V_T^*(d,\rho)$ and $V_F^*(d,\rho)$ for the equilibrium visibility of true and false stories, respectively, and $\Delta V^*(d,\rho) = V_T^*(d,\rho) - V_F^*(d,\rho)$ for the extra visibility of true stories. Similarly, define $V_T^{*\infty}(\rho) = \lim_{d\to\infty} V_T^*(d,\rho)$, $V_F^{*\infty}(\rho) = \lim_{d\to\infty} V_F^*(d,\rho)$, and $\Delta V^{*\infty}(\rho) = \lim_{d\to\infty} \Delta V^*(d,\rho)$.

It suffices to show that $b < V_F^{*\infty}(\rho) < V_T^{*\infty}(\rho) < 1$. Why? If this is the case, then each consumer i in the limit-market who sees the broadcast will have zero sharing neighbors with probability $\frac{1-V_T^{*\infty}(\rho)}{1-b}$ for true stories or $\frac{1-V_F^{*\infty}(\rho)}{1-b}$ for false stories. In this positive probability event, consumer i will update her belief by Bayes' Rules, from prior p_0 to updated belief $\hat{p}_0$, where $\frac{\hat{p}_0}{1-\hat{p}_0} = \frac{p_0}{1-p_0} \times \frac{1-V_T^{*\infty}(\rho)}{1-V_F^{*\infty}(\rho)} \in (0,1)$. So, in the limit when following infinitely-many others, consumers learn from others' sharing behavior but not perfectly; they continue to face uncertainty about which stories are true.

Next, we compute $V_T^{*\infty}(\rho)$ and $V_F^{*\infty}(\rho)$. Consider any large d. By equations (1,2), true stories have extra visibility

$$\Delta V^*(d,\rho) = (1-b)\left((1-b(1-\rho)z_T^*(d))^d - (1-b\rho z_T^*(d))^d\right) \tag{A9}$$

Using the basic mathematical fact that $\lim_{d\to\infty}(1 - X/d)^d = e^{-X}$,

$$\Delta V^{*\infty}(\rho) = (1-b)\left(e^{-b(1-\rho)\lim_{d\to\infty} dz_T^*(d)} - e^{-b\rho \lim_{d\to\infty} dz_T^*(d)}\right). \tag{A10}$$

Because $p_0^*(d,\rho) = 1-\rho$, it must be that $\Delta V^*(d,\rho) = H^{-1}(1-\rho)/M$ and hence that $\Delta V^{*\infty}(\rho) = H^{-1}(1-\rho)/M > 0$. Considering equation (A10), we conclude that $dz_T^*(d,\rho)$ must converge to a positive finite number $C^*(\rho)$ such that $e^{-b(1-\rho)C^*(\rho)} - e^{-b\rho C^*(\rho)} = \frac{H^{-1}(1-\rho)}{M(1-b)}$.[35] So, $V_T^{*\infty}(\rho) = 1-(1-b)e^{-b\rho C^*(\rho)}$ and $V_F^{*\infty}(\rho) = 1-(1-b)e^{-b(1-\rho)C^*(\rho)}$, implying as desired that $b < V_F^{*\infty}(\rho) < V_T^{*\infty}(\rho) < 1$. □

B.2 Extensions and supplementary analysis

In this section, we extend the revenue-from-views analysis of Section 2 to richer settings in which the social network is an arbitrary directed graph G and consumers may differ in their likelihood of seeing the broadcast (b_i), the precision of their private signals (ρ_i), and their belief-thresholds for sharing (p_i^S). In this enhanced model, it is assumed that all these individual parameters are common knowledge.

B.2.1 Characterization of all Bayesian Nash equilibria

Much as in the baseline model, the set of Bayesian Nash equilibria (BNE) can be easily characterized in terms of (i) how optimal consumer sharing varies with news veracity and (ii) how the relative visibility of true and false stories (which drives optimal producer investment) varies with consumer sharing.

[35] Note that $\frac{\mathrm{d}\left(e^{-b(1-\rho)C} - e^{-b\rho C}\right)}{\mathrm{d}C} = -b(1-\rho)e^{-b(1-\rho)C} + b\rho e^{-b\rho C} \gtrless$ if and only if $C \lessgtr \frac{\log(\rho/(1-\rho))}{b(2\rho-1)}$. Thus, the function $e^{-b(1-\rho)C} - e^{-b\rho C}$ is single-peaked and there are *two* solutions to the equation $e^{-b(1-\rho)C^*} - e^{-b\rho C^*} = \frac{H^{-1}(1-\rho)}{M(1-b)}$ when $\rho > \underline{\rho}^\infty$ (and zero solutions when $\rho < \underline{\rho}^\infty$). The higher solution corresponds to the dynamically-stable symmetric BNE with news veracity $1-\rho$, which is our focus here; the lower solution corresponds to the dynamically-unstable symmetric BNE.

Optimal consumer sharing. Consumer i's sharing incentives depend on news veracity ($p_0 \in [0,1]$), the precision of her own private signal ($\rho_i \in (1/2, 1)$), and her own belief threshold for sharing ($p_i^S \in [0,1]$).[36] Given news veracity p_0 and private signal $s_i \in \{T, F\}$, consumer i updates her belief to $p_{i1}(s_i; p_0)$ where, by Bayes' Rule,

$$\frac{p_{i1}(T;p_0)}{1-p_{i1}(T;p_0)} = \frac{\rho_i}{1-\rho_i} \times \frac{p_0}{1-p_0} \qquad \text{and} \qquad \frac{p_{i1}(F;p_0)}{1-p_{i1}(F;p_0)} = \frac{1-\rho_i}{\rho_i} \times \frac{p_0}{1-p_0} \tag{A11}$$

If $p_0 > p_i^{AS} \equiv \frac{(1-\rho_i)(1-p_i^S)+\rho_i p_i^S}{\rho_i p_i^S}$, then consumer i strictly prefers to share even after a negative private signal; p_i^{AS} is consumer i's "always-share threshold." Similarly, if $p_0 < p_i^{NS} \equiv \frac{\rho_i(1-p_i^S)+(1-\rho_i)p_i^S}{(1-\rho_i)p_i^S}$, then consumer i strictly prefers not to share even after a positive private signal; p_i^{NS} is consumer i's "never-share threshold." (In the baseline model in the main text, $p_i^S = \frac{1}{2}$, $p_i^{AS} = \rho$, and $p_i^{NS} = 1-\rho$ for all i.) Note that, in general, $p_i^{AS} > p_i^S > p_i^{NS}$, with p_i^{AS} increasing in ρ_i and p_i^{NS} decreasing in ρ_i.

Extra visibility of true news. Let $\mathcal{N}$ be the set of all consumers and let $\mathcal{N}_i \subset \mathcal{N}$ be the subset that consumer i follows. When others use sharing rules $\vec{\mathbf{z}}_{-i}$, each consumer $j \in \mathcal{N}_i$ shares with ex ante probability $b(\rho z_{jT} + (1-\rho)z_{jF})$ or $b((1-\rho)z_{jT} + \rho z_{jF})$ when the story is true or false, respectively. Consumer i's overall likelihood of viewing a true or false story is therefore

$$V_{iT}(\vec{\mathbf{z}}_{-i}) = 1 - (1-b)\Pi_{j\in\mathcal{N}_i}(1 - b(\rho z_{jT} + (1-\rho)z_{jF})) \tag{A12}$$

$$V_{iF}(\vec{\mathbf{z}}_{-i}) = 1 - (1-b)\Pi_{j\in\mathcal{N}_i}(1 - b((1-\rho)z_{jT} + \rho z_{jF})) \tag{A13}$$

[36] Consumer i's belief threshold for sharing depends on how her benefit when sharing true stories, π_{iT}^S, compares to her cost when sharing false stories, $-\pi_{iF}^S$, namely: $p_i^S = \frac{\pi_{iF}^S}{\pi_{iT}^S + \pi_{iF}^S} \in (0,1)$ when consumer i prefers to share only true stories ($\pi_{iT}^S > 0$ and $\pi_{iF}^S < 0$); $p_i^S = 0$ when consumer i prefers to share all stories ($\pi_{iT}^S \geq 0$ and $\pi_{iF}^S \leq 0$); or $p_i^S = 1$ when consumer i never benefits from sharing ($\pi_{iT}^S \leq 0$ and $\pi_{iF}^S \geq 0$).

and true stories enjoy extra visibility

$$\Delta V_i(\vec{\mathbf{z}}_{-i}) = (1-b)\left(\Pi_{j\in\mathcal{N}_i}(1-b((1-\rho)z_{jT}+\rho z_{jF})) - \Pi_{j\in\mathcal{N}_i}(1-b(\rho z_{jT}+(1-\rho)z_{jF}))\right) \tag{A14}$$

Best-response news veracity. Let $\mathcal{M} \subseteq \mathcal{N}$ be the set of all revenue-generating consumers. Given sharing-rule profile $\vec{\mathbf{z}}$, true stories earn additional expected revenue $\Delta R(\vec{\mathbf{z}}) \equiv \sum_{i\in\mathcal{M}} \Delta V_i(\vec{\mathbf{z}}_{-i})$, inducing best-response news veracity $H(\Delta R(\vec{\mathbf{z}}))$.

Equilibrium characterization. A BNE exists with news veracity p_0 and sharing-rule profile $\vec{\mathbf{z}}$ if and only if (i) $p_0 = H(\Delta R(\vec{\mathbf{z}}))$ and (ii) $\mathbf{z}_i \in Z_i(p_0)$ for all $i \in \mathcal{N}$.

Numerical example with differing private-signal precisions. Consider a simple example with just two consumers, where $N = M = 2$, each of whom follows one another ($d = 1$), but consumer 1 has a less precise signal than consumer 2, with $\rho_1 = \frac{4}{7}$ and $\rho_2 = \frac{6}{7}$. Moreover, suppose that broadcast reach $b = \frac{1}{2}$ and reporting cost c_R is uniformly distributed on $\left[0, \frac{3}{7}\right]$.

Visibility for consumer 1. When $p_0 > \frac{6}{7}$, consumer 2 always shares; so, consumer 1 views a story as long as anyone sees the broadcast: $V_{1T}(p_0) = V_{1F}(p_0) = 1-(1-b)^2 = \frac{3}{4}$ and $\Delta V_1(p_0) = 0$. When $\frac{1}{7} < p_0 < \frac{6}{7}$, consumer 2 shares only after a good signal; so, consumer 1 views a story when seeing the broadcast or when consumer 2 sees the broadcast and receives signal $s_2 = T$: $V_{1T}(p_0) = 1-(1-b)(1-b\rho_2) = \frac{5}{7}$, $V_{1F}(p_0) = 1-(1-b)(1-b(1-\rho_2)) = \frac{15}{28}$, and $\Delta V_1(p_0) = \frac{5}{28}$. Finally, when $p_0 < \frac{1}{7}$, consumer 2 never shares and consumer 1 views a story only when seeing the broadcast: $V_{1T}(p_0) = V_{1F}(p_0) = b = \frac{1}{2}$ and $\Delta V_1(p_0) = 0$.

Visibility for consumer 2. When $p_0 > \frac{4}{7}$, consumer 1 always shares; so, $V_{2T}(p_0) = V_{2F}(p_0) = 1-(1-b)^2 = \frac{3}{4}$ and $\Delta V_2(p_0) = 0$. When $\frac{3}{7} < p_0 < \frac{4}{7}$, consumer 1 shares only after a good signal; so, $V_{2T}(p_0) = 1-(1-b)(1-b\rho_1) = \frac{9}{14}$, $V_{1F}(p_0) = 1-(1-b)(1-b(1-\rho_1)) = \frac{17}{28}$, and $\Delta V_1(p_0) = \frac{1}{28}$. Finally, when $p_0 < \frac{3}{7}$, consumer 1 never shares; so, $V_{2T}(p_0) = V_{2F}(p_0) = b = \frac{1}{2}$ and $\Delta V_2(p_0) = 0$.

Best-response news veracity. True stories enjoy extra expected revenue $\Delta R(p_0) = \Delta V_1(p_0) + \Delta V_2(p_0)$. By the previous analysis: $\Delta R(p_0) = 0$ for all $p_0 > \frac{6}{7}$ and $p_0 < \frac{1}{7}$; $\Delta R(p_0) = \frac{5}{28}$ for all $p_0 \in \left(\frac{1}{7}, \frac{3}{7}\right)$ and $p_0 \in \left(\frac{4}{7}, \frac{6}{7}\right)$; and $\Delta R(p_0) = \frac{3}{14}$ for all $p_0 \in \left(\frac{3}{7}, \frac{4}{7}\right)$.

Maximal equilibrium news veracity. As shown in Figure A1, three news-veracity levels can be sustained in a dynamically-stable BNE: $p_0^* = \frac{1}{2}$, the maximal one, with both consumers filtering the news; $p_0 = \frac{5}{12}$, in which consumer 1 never shares and consumer 2 filters the news; and $p_0 = 0$, in the dysfunctional equilibrium in which both consumers never share. (In addition, news-veracity levels $p_0 = \frac{3}{7}$ and $p_0 = \frac{1}{7}$ can arise in dynamically-unstable BNE.)

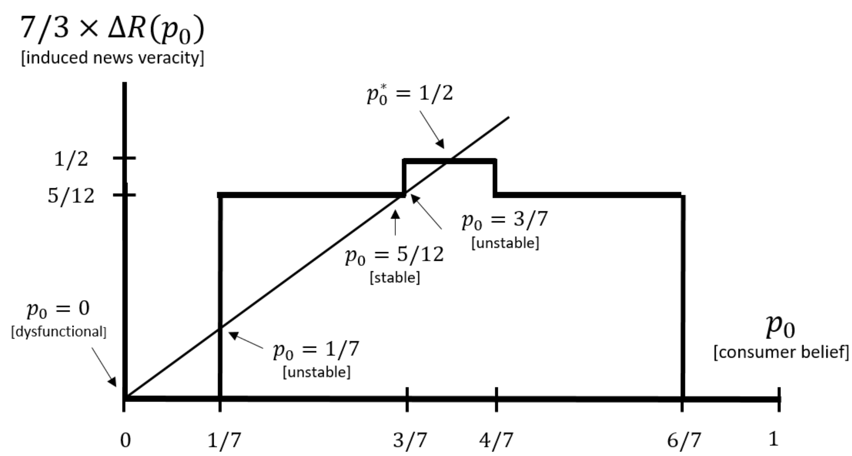

Figure B1: Illustration of equilibrium news-veracity levels in a simple numerical example with asymmetric consumers.

Discussion: impact of less-dispersed signal precisions. Consider changing consumers' precision-signal precisions from $(\rho_1, \rho_2) = \left(\frac{4}{7}, \frac{6}{7}\right)$ to $(\rho_1, \rho_2) = \left(\frac{5}{7}, \frac{5}{7}\right)$, making them less dispersed but with the same mean. Because of a special features of this numerical example (that each consumer only follows one other person), true stories' extra expected revenue when both filter the news does not change. In particular, $\Delta R\left(\frac{1}{2}\right) = M(1-b)b(2\rho-1) = \frac{3}{14}$ in either case and, given that $c_R \sim U\left[0, \frac{3}{7}\right]$, the maximal equilibrium news veracity in this example remains $p_0^* = \frac{1}{2}$. Although $\Delta R(p_0)$ remains unchanged at $\frac{3}{14}$ for $p_0 \in \left(\frac{3}{7}, \frac{4}{7}\right)$, it *increases* from $\frac{5}{28}$ to $\frac{3}{14}$ for $p_0 \in \left(\frac{2}{7}, \frac{3}{7}\right) \cup \left(\frac{4}{7}, \frac{5}{7}\right)$ and *decreases* from $\frac{5}{28}$ to 0 for $p_0 \in \left(\frac{1}{7}, \frac{2}{7}\right) \cup \left(\frac{5}{7}, \frac{6}{7}\right)$. Consequently, depending on the specific cost distribution $H(\cdot)$, the effect of this change could be to increase or to decrease the maximal equilibrium news veracity.

C Proofs and extensions: revenue from actions

This appendix provides all omitted proofs for Section 4 and generalizes the model to a setting where producers' revenues derive from both consumers' views of news stories and their actions based on news stories. The appendix first, in section C.1, characterizes equilibrium outcomes in the most general case, allowing for asymmetric consumers, an arbitrary finite social network, and a mix of revenue from both views and actions. Section C.2 then specializes that analysis to the case of symmetric consumers and revenue only from actions, and provides proofs for results in Section 4).

C.1 Extensions and supplementary analysis

This section builds on and generalizes the analysis in Appendix B.2, in two ways. First, producers may be paid for views and/or actions. In particular, each producer receives $\alpha^V \in [0,1]$ units of revenue for each view and $\alpha^A \in [0,1]$ from each action by a revenue-generating consumer, where $\alpha^V + \alpha^A = 1$. Second, consumers may have arbitrary (and asymmetric) sharing and action payoffs, parametrized by π^S_{iT}, π^S_{iF}, π^A_{iT}, and π^A_{iF}. Consequently, consumers may have different belief-thresholds for sharing, $p^S_i = \frac{\pi^S_{iF}}{\pi^S_{iT}+\pi^S_{iF}} \in (0,1)$, and for action, $p^A_i = \frac{\pi^A_{iF}}{\pi^A_{iT}+\pi^A_{iF}} \in (0,1)$. All individual parameters are common knowledge.[37]

Consumer strategies. Each consumer's strategy consists of a (i) a *sharing rule* $\mathbf{z}_i = (z_{is_i} : s_i \in \{T, F\})$ and (i) an *action rule* $\mathbf{a}_i = \left(a_{io_i} : o_i = (s_i, S_i) \in \{T, F\} \times 2^{\mathcal{N}_i}\right)$, where s_i is consumer i's private signal and $S_i \subseteq \mathcal{N}_i$ is the subset of i's neighbors who shared at time $t = 1$. (2^X denotes the set of all subsets of X.)

[37] Some consumers may always (or never) prefer to share or act, no matter whether the news is true, corresponding to belief thresholds $p^S_i = 0$ and $p^A_i = 0$ (or $p^S_i = 1$ and $p^A_i = 1$), respectively. Such consumers' behavior is trivial to describe and lacks information content; we therefore focus on the more interesting case in which $\pi^S_{iT}, \pi^S_{iF}, \pi^A_{iT}, \pi^A_{iF} > 0$, so that each consumer prefers to share and act on true stories but not on false ones.

Optimal sharing. Let $p_{i1}(s_i; p_0)$ denote consumer i's time-1 belief about the likelihood of news truth, depending on private signal $s_i \in \{T, F\}$. By Bayes' Rule,

$$\frac{p_{i1}(T; p_0)}{1 - p_{i1}(T; p_0)} = \frac{\rho_i}{1 - \rho_i} \times \frac{p_0}{1 - p_0} \quad \text{and} \quad \frac{p_{i1}(F; p_0)}{1 - p_{i1}(F; p_0)} = \frac{1 - \rho_i}{\rho_i} \times \frac{p_0}{1 - p_0} \tag{A15}$$

Define consumer i's "never-share threshold" p_i^{NS} and "always-share threshold" p_i^{AS} by the conditions $p_{i1}(T; p_i^{NS}) = p_i^S$ and $p_{i1}(F; p_i^{AS}) = p_i^S$. (When $p_i^S = \frac{1}{2}$, as in the main text, $p_i^{NS} = 1 - \rho_i$ and $p_i^{AS} = \rho_i$; more generally, $0 < p_i^{NS} < p_i^S < p_i^{AS} < 1$.) We can now characterize consumer i's set of optimal sharing rules, denoted $Z_i(p_0)$, much as in Lemma 1: $Z_i(p_0) = (0, 0)$ for all $p_0 < p_i^{NS}$; $Z_i(p_i^{NS}) = \{(z_{iT}, 0) : z_{iT} \in [0, 1]\}$; $Z_i(p_0) = (1, 0)$ for all $p_i^{NS} < p_0 < p_i^{AS}$; $Z_i(p_i^{AS}) = \{(1, z_{iF}) : z_{iF} \in [0, 1]\}$; and $Z_i(p_0) = (1, 1)$ for all $p_0 > p_i^{AS}$.

Optimal action. Let $p_{i2}(o_i; \vec{\mathbf{z}}_{-i}, p_0)$ denote consumer i's time-2 belief conditional on "observation" $o_i = (s_i, S_i)$, given ex ante belief p_0 and others' sharing-rule profile $\vec{\mathbf{z}}_{-i}$. Let $\mathcal{O}_i^>(\vec{\mathbf{z}}_{-i}, p_0) = \{o_i : p_{i2}(o_i; \vec{\mathbf{z}}_{-i}, p_0) > p^A\}$ and $\mathcal{O}_i^=(\vec{\mathbf{z}}_{-i}, p_0) = \{o_i : p_{i2}(o_i; \vec{\mathbf{z}}_{-i}, p_0) = p^A\}$ denote the subsets of observations given which, respectively, consumer i strictly prefers to act and is indifferent whether to act. Let $\mathcal{A}_i(\vec{\mathbf{z}}_{-i}, p_0)$ denote the set of "best-response action rules:" $\mathbf{a}_i \in \mathcal{A}_i(\vec{\mathbf{z}}_{-i}, p_0)$ if and only if $a_{io_i} = 1$ for all $o_i \in \mathcal{O}_i^>(\vec{\mathbf{z}}_{-i}, p_0)$, $a_{io_i} \in [0, 1]$ for all $o_i \in \mathcal{O}_i^=(\vec{\mathbf{z}}_{-i}, p_0)$, and $a_{io_i} = 0$ for all $o_i \notin \mathcal{O}_i^>(\vec{\mathbf{z}}_{-i}, p_0) \cup \mathcal{O}_i^=(\vec{\mathbf{z}}_{-i}, p_0)$.

Equilibrium characterization. An equilibrium exists with news veracity p_0, sharing-rule profile $\vec{\mathbf{z}}$, and action-rule profile $\vec{\mathbf{a}}$ if and only if (i) $\mathbf{z}_i$ is an optimal sharing rule for all i, i.e., $\mathbf{z}_i \in Z_i(p_0)$; (ii) $\mathbf{a}_i$ is an optimal action rule for all i, i.e., $\mathbf{a}_i \in \mathcal{A}_i(\vec{\mathbf{z}}_{-i}, p_0)$; and (iii) optimal producer investment generates news veracity p_0, i.e., $p_0 = H\left(\sum_{i \in \mathcal{M}} \Delta R_i(\vec{\mathbf{z}}_{-i}, \vec{\mathbf{a}})\right)$, where $\Delta R_i(\vec{\mathbf{z}}_{-i}, \vec{\mathbf{a}})$ is the extra expected revenue generated by true stories due to consumer i. Note that an equilibrium always exists in which all stories are false ($p_0 = 0$), consumers never share ($\vec{\mathbf{z}} = \mathbf{0}$), and consumers never act ($\vec{\mathbf{a}} = \mathbf{0}$). $Z_i(p_0)$ and $\mathcal{A}_i(\vec{\mathbf{z}}_{-i}, p_0)$ were characterized previously; in order to characterize all Nash equilibria, it remains for us to characterize $\Delta R_i(\vec{\mathbf{z}}_{-i}, \vec{\mathbf{a}})$ for each i.

Expected revenue of true and false stories. True and false stories earn expected revenue $R_{iT} = \alpha_i^V V_{iT} + \alpha_i^A A_{iT}$ and $R_F = \alpha_i^V V_{iF} + \alpha_i^A A_{iF}$, respectively, where (α_i^V, α_i^A) are the revenue weights on views and actions for consumer i and (V_{iT}, A_{iT}) and (V_{iF}, A_{iF}) are the ex ante likelihoods that consumer i will view and act conditional on a story being true or false, respectively.

V_{iT}, V_{iF} depend on others' sharing behavior but not directly on news veracity, and hence are functions (only) of $\vec{\mathbf{z}}_{-i}$. We characterized $V_{iT}(\vec{\mathbf{z}}_{-i})$ and $V_{iF}(\vec{\mathbf{z}}_{-i})$ earlier, in equations (A12, A13).

A_{iT}, A_{iF} also depend on others' sharing behavior, but now for two reasons, as sharing impacts both (i) the likelihood that consumers view the story and hence have the opportunity to act and (ii) how much consumers update their beliefs about whether the story is true based on neighbors' sharing behavior. Since consumers only act when their updated belief meets or exceeds the action threshold p^A, the ex ante likelihood p_0 that stories are true also impacts each consumer's likelihood of acting. Consequently, A_{iT}, A_{iF} are functions of both $\vec{\mathbf{z}}_{-i}$ and p_0, as well as consumer i's action rule $\mathbf{a}_i$. It remains for us to characterize $A_{iT}(\mathbf{a}_i, \vec{\mathbf{z}}_{-i}, p_0)$ and $A_{iT}(\mathbf{a}_i, \vec{\mathbf{z}}_{-i}, p_0)$ for any $(\mathbf{a}_i, \vec{\mathbf{z}}_{-i}, p_0)$.

Action likelihoods. Conditional on the story being true or false, the likelihood of any given observation $o_i = (s_i, S_i)$ depends (only) on others' sharing rules. Let $L_{iT}(o_i; \vec{\mathbf{z}}_{-i})$ and $L_{iF}(o_i; \vec{\mathbf{z}}_{-i})$ denote the ex ante likelihoods that consumer i views a story and observes o_i conditional on the story being true or false, respectively. Upon observing o_i, consumer

i updates her belief to $p_{i2}(o_i; \vec{\mathbf{z}}_{-i}, p_0)$ where, by Bayes' Rule,

$$\frac{p_{i2}(o_i; \vec{\mathbf{z}}_{-i}, p_0)}{1 - p_{i2}(o_i; \vec{\mathbf{z}}_{-i}, p_0)} = \frac{p_0}{1 - p_0} \times \frac{L_{iT}(o_i; \vec{\mathbf{z}}_{-i})}{L_{iF}(o_i; \vec{\mathbf{z}}_{-i})} \tag{A16}$$

We can now characterize the set of action-inducing observations in terms of their relative likelihood of arising when a story is true versus false. In particular: $o_i \in \mathcal{O}_i^{>}(\vec{\mathbf{z}}_{-i}, p_0)$ if and only if $\frac{L_{iT}(o_i;\vec{\mathbf{z}}_{-i})}{L_{iF}(o_i;\vec{\mathbf{z}}_{-i})} > \frac{p^A(1-p_0)}{(1-p^A)p_0}$ and $o_i \in \mathcal{O}_i^{=}(\vec{\mathbf{z}}_{-i}, p_0)$ if and only if $\frac{L_{iT}(o_i;\vec{\mathbf{z}}_{-i})}{L_{iF}(o_i;\vec{\mathbf{z}}_{-i})} = \frac{p^A(1-p_0)}{(1-p^A)p_0}$. Since consumer i chooses to act after observations in $\mathcal{O}_i^{>}(\vec{\mathbf{z}}_{-i}, p_0)$ and may randomize after observations in $\mathcal{O}_i^{=}(\vec{\mathbf{z}}_{-i}, p_0)$, any optimal action rule $\mathbf{a}_i \in \mathcal{A}_i(\vec{\mathbf{z}}_{-i}, p_0)$ induces action likelihoods of the form:

$$A_{iT}(\mathbf{a}_i, \vec{\mathbf{z}}_{-i}, p_0) = \sum_{o_i \in \mathcal{O}_i(\vec{\mathbf{z}}_{-i}, p_0)} L_{iT}(o_i; \vec{\mathbf{z}}_{-i}) + \sum_{o_i \in \mathcal{O}_i^{=}(\vec{\mathbf{z}}_{-i}, p_0)} a_{io_i} L_{iT}(o_i; \vec{\mathbf{z}}_{-i}) \tag{A17}$$

$$A_{iF}(\mathbf{a}_i, \vec{\mathbf{z}}_{-i}, p_0) = \sum_{o_i \in \mathcal{O}_i(\vec{\mathbf{z}}_{-i}, p_0)} L_{iF}(o_i; \vec{\mathbf{z}}_{-i}) + \sum_{o_i \in \mathcal{O}_i^{=}(\vec{\mathbf{z}}_{-i}, p_0)} a_{io_i} L_{iF}(o_i; \vec{\mathbf{z}}_{-i}) \tag{A18}$$

The extra likelihood that consumer i acts based on true stories is then simply

$$\Delta A_i(\mathbf{a}_i, \vec{\mathbf{z}}_{-i}, p_0) = \sum_{o_i \in \mathcal{O}_i(\vec{\mathbf{z}}_{-i}, p_0)} \Delta L_i(o_i; \vec{\mathbf{z}}_{-i}) + \sum_{o_i \in \mathcal{O}_i^{=}(\vec{\mathbf{z}}_{-i}, p_0)} a_{io_i} \Delta L_i(o_i; \vec{\mathbf{z}}_{-i}) \tag{A19}$$

where $\Delta L_i(o_i; \vec{\mathbf{z}}_{-i}) = L_{iT}(o_i; \vec{\mathbf{z}}_{-i}) - L_{iF}(o_i; \vec{\mathbf{z}}_{-i})$. It remains to characterize $L_{iT}(o_i; \vec{\mathbf{z}}_{-i})$ and $L_{iF}(o_i; \vec{\mathbf{z}}_{-i})$ for all $(o_i; \vec{\mathbf{z}}_{-i})$.

Observation likelihoods. Consider any consumer $j \in \mathcal{N}_i$ followed by consumer i. Consumer j's ex ante likelihood of sharing is $b_j(\rho_j z_{jT} + (1 - \rho_j) z_{jF})$ when a story is true or $b_j((1-\rho_j) z_{jT} + \rho_j z_{jF})$ when it is false. So long as at least one neighbor shares, consumer i is sure to view the story and hence receive a private signal s_i; alternatively, if no neighbor shares, consumer i will only view the story with probability b_i (by seeing the broadcast).

Overall, then, conditional on the story being true or false, any observation $o_i = (s_i, S_i)$ with $S_i \neq \emptyset$ has likelihood

$$L_{iT}(s_i, S_i; \vec{\mathbf{z}}_{-i}) = \Pr(s_i|\text{true}) \prod_{j \in S_i} \big(b_j(\rho_j z_{jT} + (1-\rho_j) z_{jF})\big) \prod_{j \in \mathcal{N}_i \setminus S_i} \big(1 - b_j(\rho_j z_{jT} + (1-\rho_j) z_{jF})\big)$$
$$L_{iF}(s_i, S_i; \vec{\mathbf{z}}_{-i}) = \Pr(s_i|\text{false}) \prod_{j \in S_i} \big(b_j((1-\rho_j) z_{jT} + \rho_j z_{jF})\big) \prod_{j \in \mathcal{N}_i \setminus S_i} \big(1 - b_j((1-\rho_j) z_{jT} + \rho_j z_{jF})\big)$$

Similarly, any observation $o_i = (s_i, \emptyset)$ has conditional likelihoods

$$L_{iT}(s_i, \emptyset; \vec{\mathbf{z}}_{-i}) = b_i \Pr(s_i|\text{true}) \prod_{j \in \mathcal{N}_i} \big(1 - b_j(\rho_j z_{jT} + (1-\rho_j) z_{jF})\big)$$
$$L_{iF}(s_i, \emptyset; \vec{\mathbf{z}}_{-i}) = b_i \Pr(s_i|\text{true}) \prod_{j \in \mathcal{N}_i} \big(1 - b_j((1-\rho_j) z_{jT} + \rho_j z_{jF})\big)$$

This completes our implicit characterization of all Nash equilibrium in the general case, allowing for an arbitrary social network and arbitrary belief thresholds for sharing and action, among other extensions of the baseline revenue-for-views and revenue-for-actions models analyzed in the main text.

C.2 Special case considered in Section 4

Here we specialize the previous analysis to the case examined in Section 4 of the text: (i) consumers are symmetric, i.e., $b_i = b$, $\rho_i = \rho$, and $\#(\mathcal{N}_i) = d$ for all i, (ii) producers earn revenue from actions only, i.e., $\alpha^A = 1$ and $\alpha^V = 0$, and (iii) consumers have the same belief-threshold for action as for sharing, i.e., $p^A = p^S = \frac{1}{2}$, and consumers follow the same strategy. Consumers' belief updating then depends only on *how many* neighbors share, not their particular identities. We can therefore simplify each consumer's "observation" o_i to consist of only her private signal $s_i \in \{T, F\}$ and the number of sharing neighbors $\sigma_i \in \{0, 1, ..., d\}$. For further simplicity, we focus on equilibria in which all consumers use the same action rule $\mathbf{a}$.

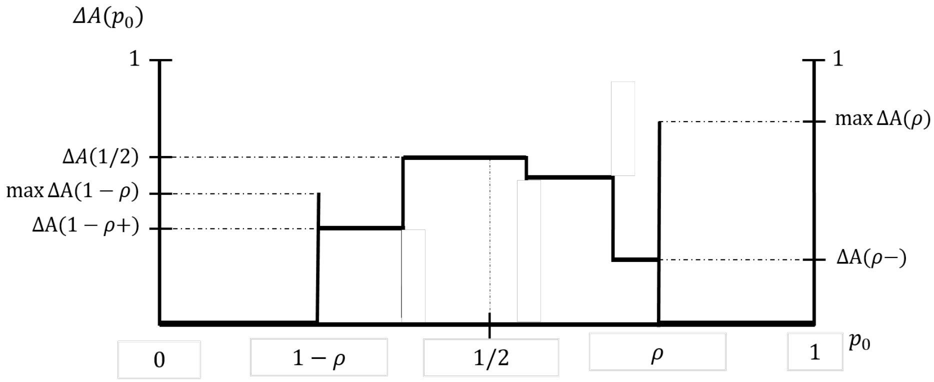

Figure 8: Illustration of extra-action-likelihood correspondence $\Delta A(p_0)$ in a finite network (Lemma A2).

Let $\Delta A(p_0) = \{\Delta A(\mathbf{a}, \mathbf{z}, p_0) : \mathbf{z} \in Z(p_0) \text{ and } \mathbf{a} \in \mathcal{A}(\mathbf{z}, p_0)\}$ be the *range* of possible values that the extra likelihood of true-story action (ΔA) can take when consumers share and act optimally given news veracity $p_0 \in [0, 1]$, and let $\mathcal{P}(p_0)$ denote the resulting news veracity when producers invest optimally. A (symmetric) equilibrium exists with news veracity p_0 if and only if $p_0 \in \mathcal{P}(p_0)$.

Never-share region ($p_0 < 1 - \rho$). In any equilibrium with news veracity $p_0 < 1 - \rho$, consumers must never share ($\mathbf{z} = \mathbf{0}$) and, because no stories are ever shared, consumers must learn nothing from others' decisions not to share. Since by assumption consumers' action threshold $p^A = \frac{1}{2}$, consumers must therefore never act ($\mathbf{a} = \mathbf{0}$); so, $\Delta A(p_0) = 0$ and hence $\mathcal{P}(p_0) = 0$ for all $p_0 < 1 - \rho$. We conclude that a dysfunctional equilibrium exists with zero news veracity, but no equilibria exist with news veracity $p_0 \in (0, 1 - \rho)$.

Always-share region ($p_0 > \rho$). In any equilibrium with news veracity $p_0 > \rho$, consumers must always share ($\mathbf{z} = \mathbf{1}$) and, because all stories are shared, consumers learn nothing from others' sharing behavior. Since consumers' action threshold $p^A = \frac{1}{2}$, consumers must always act ($\mathbf{a} = \mathbf{1}$), inducing equal action for true and false stories; so,

$\Delta A(p_0) = 0$ and hence $\mathcal{P}(p_0) = 0$ for all $p_0 > \rho$. We conclude that no equilibrium exists with news veracity greater than ρ.

Filtering region ($1 - \rho < p_0 < \rho$). When consumers filter the news, seeing someone share (or not share) is a positive (or negative) signal about news truth, and consumers find it optimal to act whenever enough of their neighbors share. In particular, there exist real-valued functions $(\tilde{\sigma}_T(p_0), \tilde{\sigma}_F(p_0))$, referred to as "neighbor-sharing thresholds," such that consumer i strictly prefers to act after observing $o_i = (s_i, \sigma)$ for any $\sigma > \tilde{\sigma}_{s_i}(p_0)$ and strictly prefers not to act after observing $o_i = (s_i, \sigma)$ for any $\sigma < \tilde{\sigma}_{s_i}(p_0)$. These neighbor-sharing thresholds, and the resulting action likelihoods for true and false stories, are explicitly derived in the proof of Lemma A2 below.

Lemma A2. *(i) For all $p_0 < 1-\rho$ and all $p_0 > \rho$, $\Delta A(p_0) = 0$. (ii) For all $p_0 \in (1-\rho, \rho)$, $\Delta A(p_0)$ is a continuous and interval-valued correspondence, with*

$$\min \Delta A(p_0) = \sum_{\sigma > \tilde{\sigma}_T(p_0)} \Delta L(T, \sigma; \mathbf{z}) + \sum_{\sigma > \tilde{\sigma}_F(p_0)} \Delta L(F, \sigma; \mathbf{z}) \tag{A20}$$

$$\max \Delta A(p_0) = \sum_{\sigma \geq \tilde{\sigma}_T(p_0)} \Delta L(T, \sigma; \mathbf{z}) + \sum_{\sigma \geq \tilde{\sigma}_F(p_0)} \Delta L(F, \sigma; \mathbf{z}) \tag{A21}$$

where $\tilde{\sigma}_T(p_0)$ and $\tilde{\sigma}_F(p_0)$ are continuous and strictly decreasing. Moreover, over this domain, $\Delta A(p_0)$ is single peaked and maximized at $p_0 = \frac{1}{2}$, i.e., $\max \Delta A(p_0) \leq \min \Delta A(p_0')$ for all $1 - \rho < p_0 < p_0' \leq \frac{1}{2}$ and for all $\frac{1}{2} \leq p_0' < p_0 < \rho$.

Proof. Part (i) is proven in the paragraphs preceding the lemma, so we focus here on part (ii). In any equilibrium with news veracity $p_0 \in (1 - \rho, \rho)$, consumers must use sharing rule $\tilde{\mathbf{z}} = (1, 0)$. The ex ante likelihood that each consumer shares is $b\rho$ when a story is true or $b(1 - \rho)$ when it is false. Thus, the relative likelihood of observation $o_i = (s_i, \sigma_i)$

when a story is true versus false takes the form

$$\frac{L_T(s_i, \sigma_i; \tilde{\mathbf{z}})}{L_F(s_i, \sigma_i; \tilde{\mathbf{z}})} = \frac{\Pr(s_i|\text{true})}{\Pr(s_i|\text{false})} \times \left(\frac{\rho}{1-\rho}\right)^{\sigma_i} \times \left(\frac{1-b\rho}{1-b(1-\rho)}\right)^{d-\sigma_i}$$

causing consumers to update their time-2 belief to $p_2(s_i, \sigma_i; \tilde{\mathbf{z}}, p_0)$ where, by Bayes' Rule, $\frac{p_2(s_i,\sigma_i;\tilde{\mathbf{z}},p_0)}{1-p_2(T,\sigma_i;\tilde{\mathbf{z}},p_0)} = \frac{p_0}{1-p_0} \times \frac{L_T(s_i,\sigma_i;\tilde{\mathbf{z}})}{L_F(s_i,\sigma_i;\tilde{\mathbf{z}})}$. In particular:

$$\frac{p_2(T, \sigma_i; \tilde{\mathbf{z}}, p_0)}{1 - p_2(T, \sigma_i; \tilde{\mathbf{z}}, p_0)} = \frac{p_0}{1-p_0} \times \left(\frac{\rho}{1-\rho}\right)^{\sigma_i+1} \times \left(\frac{1-b\rho}{1-b(1-\rho)}\right)^{d-\sigma_i} \tag{A22}$$

$$\frac{p_2(F, \sigma_i; \tilde{\mathbf{z}}, p_0)}{1 - p_2(F, \sigma_i; \tilde{\mathbf{z}}, p_0)} = \frac{p_0}{1-p_0} \times \left(\frac{\rho}{1-\rho}\right)^{\sigma_i-1} \times \left(\frac{1-b\rho}{1-b(1-\rho)}\right)^{d-\sigma_i} \tag{A23}$$

Since the action threshold $p^A = \frac{1}{2}$, consumer i strictly prefers to act whenever $\frac{p_2(s_i,\sigma_i;\tilde{\mathbf{z}},p_0)}{1-p_2(T,\sigma_i;\tilde{\mathbf{z}},p_0)} > 1$. Let $\tilde{\sigma}_T(p_0)$ and $\tilde{\sigma}_F(p_0)$ be the (possibly negative) levels of σ_i given which, respectively, the right-hand sides of (A22) and (A23) are equal to one. Since these expressions are exponentially increasing in σ_i, increasing in p_0, and continuous in (p_0, σ), $\tilde{\sigma}_T(p_0)$ and $\tilde{\sigma}_F(p_0)$ are well-defined, continuous, and strictly decreasing in p_0.[38] We conclude that consumer i strictly prefers to act after any observation $o_i \in \widetilde{\mathcal{O}}^{>}(p_0) \equiv \{(T, \sigma_i) : \sigma_i > \tilde{\sigma}_T(p_0)\} \cup \{(F, \sigma_i) : \sigma_i > \tilde{\sigma}_F(p_0)\}$ and is indifferent whether to act after any observation $o_i \in \widetilde{\mathcal{O}}^{=}(p_0) \equiv \{(T, \tilde{\sigma}_T(p_0)), (F, \tilde{\sigma}_F(p_0))\}$. Because $\sigma_T(p_0))$ and $\sigma_F(p_0)$ are each continuous and strictly decreasing in p_0, $\widetilde{\mathcal{O}}^{=}(p_0) = \emptyset$ for all but finitely-many p_0-levels.

Let $\tilde{\mathbf{a}}(p_0)$ denote an action rule consistent with optimal action, i.e., $a_{o_i}(p_0) = 1$ for all $o_i \in \widetilde{\mathcal{O}}^{>}(p_0)$ and $a_{o_i}(p_0) = 0$ for all $o_i \notin \widetilde{\mathcal{O}}^{>}(p_0) \cup \widetilde{\mathcal{O}}^{=}(p_0)$. Note that, except for finitely

[38] One can also show that (i) $\tilde{\sigma}_F(p_0) > \tilde{\sigma}_T(p_0) + 1$, (ii) $\tilde{\sigma}_T(p_0) < d$, and (iii) $\tilde{\sigma}_F(p_0) > 1$ (details omitted), i.e., (i) consumers need to see more neighbors share after a negative private signal, (ii) they act after a positive signal and seeing all of their neighbors share, and (iii) they do not act after a negative signal and seeing no one share.

many $p_0 \in (1-\rho, \rho)$, there is a unique optimal action rule, inducing action likelihoods

$$A_T(p_0) = \sum_{o_i \in \widetilde{\mathcal{O}}(p_0)} L_T(o_i; \tilde{\mathbf{z}}) \qquad \text{and} \qquad A_F(p_0) = \sum_{o_i \in \widetilde{\mathcal{O}}(p_0)} L_F(o_i; \tilde{\mathbf{z}})$$

Since $\widetilde{\mathcal{O}}^{>}(p_0) = \{(s_i, \sigma_i) : s_i \in \{T, F\}, \sigma_i > \sigma_{s_i}(p_0)\}$, this verifies equations (A20-A21) for these news-veracity levels, with $\max \Delta A(p_0) = \min \Delta A(p_0) = \sum_{o_i \in \widetilde{\mathcal{O}}^{>}(p_0)} \Delta L(o_i; \tilde{\mathbf{z}})$.

Let $X \subset (1-\rho, \rho)$ denote the finite set of news-veracity levels given which consumers are sometimes indifferent whether to act, i.e., $p_0 \in X$ if and only if $\widetilde{\mathcal{O}}^{=}(p_0) \neq \emptyset$. Without loss, label the elements of X in order: $X = \{x_1, x_2, ..., x_K\}$ with $1-\rho < x_1 \leq x_2 \leq x_K < \rho$. Note that, given any news-veracity level $p_0 \in (x_k, x_{k+1})$, consumers will act after the same set of observations; thus, $A_T(p_0)$ and $A_F(p_0)$ are constant over each of these subintervals. In particular, $A_T(p_0)$, $A_F(p_0)$, and hence $\Delta A(p_0)$ are each *step functions* over $(1-\rho, \rho) \backslash X$, with discontinuities at the news-veracity levels in X.

Now, consider any news-veracity level $p_0 \in X$. The extra likelihood of true-story action takes the form

$$\Delta A(\mathbf{a}, \tilde{\mathbf{z}}, p_0) = \sum_{o_i \in \widetilde{\mathcal{O}}(p_0)} \Delta L_i(o_i; \tilde{\mathbf{z}}) + \sum_{o_i \in \widetilde{\mathcal{O}}^{=}(p_0)} a_{o_i} \Delta L(o_i; \tilde{\mathbf{z}}), \tag{A24}$$

a special case of equation (A19). By definition, each observation $o_i \in \widetilde{\mathcal{O}}^{=}(p_0)$ must leave consumers indifferent whether to act, meaning that consumers must update their belief about the likelihood of news truth from p_0 to exactly $p^A = \frac{1}{2}$. For $p_0 \in X$ less than $\frac{1}{2}$, observation o_i must therefore be *more* likely to occur when the news is true than when it is false; so, $L_T(o_i; \tilde{\mathbf{z}}) > L_F(o_i; \tilde{\mathbf{z}})$ and hence $\Delta L_i(o_i; \tilde{\mathbf{z}}) > 0$ for all $o_i \in \widetilde{\mathcal{O}}^{=}(p_0)$. Consequently, as the action-mixing probabilities $(a_{o_i} : o_i \in \widetilde{\mathcal{O}}^{=}(p_0))$ range over the unit cube $[0,1]^{\#(\widetilde{\mathcal{O}}^{=}(p_0))}$, $\Delta A(\mathbf{a}, \tilde{\mathbf{z}}, p_0)$ varies from a minimum equal to $\lim_{\epsilon \to 0} \Delta A(p_0 - \epsilon) = \sum_{o_i \in \widetilde{\mathcal{O}}(p_0)} \Delta L(o_i; \tilde{\mathbf{z}})$ to a maximum equal to $\lim_{\epsilon \to 0} \Delta A(p_0 + \epsilon) =$

$\sum_{o_i \in \widetilde{\mathcal{O}}(p_0) \cup \widetilde{\mathcal{O}}^{=}(p_0)} \Delta L(o_i; \tilde{\mathbf{z}})$. This proves that (i) for all $p_0 < \frac{1}{2}$, $\Delta A(p_0)$ is interval-valued with minimal and maximal values given by equations (A20-A21), and (ii) for all $p_0 \in X$ less than $\frac{1}{2}$, the discontinuity in $\Delta A(p_0)$ at p_0 is an *upward* discontinuity.

On the other hand, for $p_0 \in X$ greater than $\frac{1}{2}$, observation o_i must be *less* likely to occur when the news is true; so, $L_T(o_i; \tilde{\mathbf{z}}) < L_F(o_i; \tilde{\mathbf{z}})$ and hence $\Delta L_i(o_i; \tilde{\mathbf{z}}) < 0$ for all $o_i \in \widetilde{\mathcal{O}}^{=}(p_0)$. But then, by the same argument as before, it must be that $\Delta A(\mathbf{a}, \tilde{\mathbf{z}}, p_0)$ varies from a minimum equal to $\lim_{\epsilon \to 0} \Delta A(p_0 + \epsilon) = \sum_{o_i \in \widetilde{\mathcal{O}}^{>}(p_0) \cup \widetilde{\mathcal{O}}^{=}(p_0)} \Delta L(o_i; \tilde{\mathbf{z}})$ to a maximum equal to $\lim_{\epsilon \to 0} \Delta A(p_0 - \epsilon) = \sum_{o_i \in \widetilde{\mathcal{O}}^{>}(p_0)} \Delta L(o_i; \tilde{\mathbf{z}})$. This again proves that $\Delta A(p_0)$ is interval-valued with minimal and maximal values given by equations (A20-A21), but now the discontinuities in $\Delta A(p_0)$ at $p_0 \in X$ are *downward* discontinuity.

Combining these observations, we conclude that $\Delta A(p_0)$ is single-peaked in p_0 over the filtering region and maximized at the action threshold $p_0 = \frac{1}{2}$, as desired. $\square$

Because producer cost c_R is drawn from an atomless distribution having support on $(0, \infty)$, induced news veracity $p_0 = H(M\Delta A)$ is a continuous, strictly increasing function of ΔA. We conclude that $\mathcal{P}(p_0)$, like $\Delta A(p_0)$ is a continuous, interval-valued, single-peaked function over the filtering region, maximized over this domain at $p_0 = 1/2$. Some immediate consequences: (i) an equilibrium exists with news veracity $p_0 > 1/2$ if and only if $H(M \max \Delta A(1/2)) > 1/2$; and (ii) an equilibrium exists with news veracity $p_0 \in (1 - \rho, \rho)$ if $\lim_{\epsilon \to 0} H(M\Delta A(1 - \rho + \epsilon)) > 1 - \rho$ and $\lim_{\epsilon \to 0} H(M\Delta A(\rho - \epsilon)) < \rho$.

Always-share threshold ($p_0 = \rho$) and never-share threshold ($p_0 = 1-\rho$). If news veracity $p_0 = \rho$ or $p_0 = 1 - \rho$, consumers must use a sharing rule of the form $(1, z_F)$ or $(z_T, 0)$, respectively. Each such sharing rule generates different observation likelihoods, affecting which observations are sufficiently "good news" to prompt consumers to act and thereby inducing a different extra likelihood of true-story action (ΔA). Computing $\Delta A(\rho)$ and $\Delta(1 - \rho)$, the range of values that ΔA can potentially take, is complex but ultimately straightforward given the analytical machinery introduced in Appendix C.1. Lemma A3 below collects key facts about $\Delta A(\rho)$ and $\Delta(1-\rho)$, useful later in other proofs.

Lemma A3. *(i)* $\Delta A(\rho) = \left[0, \overline{\Delta A}(\rho)\right]$, *where* $\overline{\Delta A}(\rho) \geq \lim_{\epsilon\to 0} \Delta A(\rho - \epsilon)$. *(ii)* $\Delta A(1-\rho) = \left[0, \overline{\Delta A}(1-\rho)\right]$, *where* $\overline{\Delta A}(1-\rho) \geq \lim_{\epsilon\to 0} \Delta A(1-\rho+\epsilon)$.

Proof. First, we show that $\min \Delta A(\rho) = 0$. Suppose that consumers use the always-share rule $\mathbf{z} = (1,1)$, which is optimal given news veracity $p_0 = \rho$. Consumers then learn nothing from others' sharing behavior and hence find it optimal to use the always-share rule. (Consumers are indifferent whether to act after receiving a bad private signal; so, all sharing rules of the form $(1, z_F)$ are optimal, including the always-share rule.) Under such optimal sharing and optimal action behavior, true and false stories are equally acted upon, resulting in $\Delta A = 0$; so, $0 \in \Delta A(\rho)$. Next, observe that optimizing consumers never act on false stories more often than true stories, since to do so would result in a negative expected action payoff, worse than never acting at all; so, $\min \Delta A(\rho) \geq 0$.

Next, we show that $\max \Delta A(\rho) \geq \widetilde{\Delta A} \equiv \lim_{\epsilon\to 0} \Delta A(\rho - \epsilon)$. Suppose that consumers use the filtering rule $\mathbf{z} = \tilde{\mathbf{z}} = (1,0)$. Consumers' updated beliefs after any given observation, and resulting action behavior, are characterized in the proof of Lemma A2. Let $\widetilde{\Delta A}(\rho)$ denote the resulting extra likelihood of true-story action—potentially an interval, if $p_0 = \rho$ is one of the finitely-many news-veracity levels at which consumers are indifferent whether to act after some observation. Since the filtering rule is optimal given news veracity $p_0 = \rho$, $\widetilde{\Delta A}(\rho) \subset \Delta A(\rho)$. And by the continuity argument of Lemma A2, it must be that $\widetilde{\Delta A} \in \widetilde{\Delta A}(\rho)$. All together, we conclude that $\widetilde{\Delta A} \leq \max \Delta A(\rho)$.

Finally, we show that $\Delta A(\rho)$ is an interval. For each optimal sharing rule $(1, z_F)$, let $\mathcal{A}(\rho; z_F)$ denote the set of optimal action rules. An optimal action rule must specify action probability $a_{o_i} = 1$ (or $= 0$) if consumers strictly prefer to act (or not to act) after observation o_i, or anything from the interval $[0,1]$ if consumers are indifferent. The set $\mathcal{A}(\rho; z_F)$ is therefore equivalent to a k-dimensional unit square, where k is the number of observations after which consumers are indifferent, or a singleton if consumers are never indifferent.

By definition, $\Delta A(\rho) = \cup_{z_F \in [0,1]} \cup_{\mathbf{a} \in \mathcal{A}(\rho; z_F)} \Delta A(\mathbf{a}, \mathbf{z})$. The extra likelihood that true stories are acted upon, $\Delta A(\mathbf{a}, \mathbf{z})$, is obviously continuous in $\mathbf{a}$ and, because the likelihood of any given observation is continuous in $\mathbf{z}$, must also be continuous in $\mathbf{z}$. Since $\mathcal{A}(\rho; z_F)$ is a product of intervals, we conclude that $\cup_{\mathbf{a} \in \mathcal{A}(\rho; z_F)} \Delta A(\mathbf{a}, \mathbf{z})$ is an interval for each z_F, and that the union of these intervals is also an interval.

This completes the proof of part (i). The proof of part (ii) is essentially identical and omitted to save space. $\square$

C.3 Proof of Thm 2

Preliminaries. For any finite social connectedness d, let $\mathcal{P}(p_0; d) = H(M \Delta A(p_0; d))$ be the continuous correspondence mapping any consumer belief p_0 to the interval of news veracities that can potentially arise given that belief when consumers share and act optimally and producers invest optimally. (The fact that $\mathcal{P}(p_0; d)$ is a continuous and interval-valued follows from the fact that $\Delta A(p_0; d)$ is a continuous and interval-valued (Lemmas A2-A3) and that $H(\cdot)$ is continuous, due to producer cost c_R being atomless.)

Given any news veracity $p_0 \in (1-\rho, \rho)$, consumers optimally filter the news, allowing consumers in the $d \to \infty$ limit to discern perfectly which stories are true; so, true stories are always acted upon while false stories are never acted upon in the limit, making true stories' extra likelihood of action as large as possible: $\lim_{d \to \infty} \Delta A(p_0; d) = 1$. Of course, for any finite d, any observation that induces consumers to act will sometimes occur for false stories as well as true ones; so, $\Delta A(p_0; d) < 1$ for any finite d.

We conclude that, for all $p_0 \in (1-\rho, \rho)$, $\max \mathcal{P}(p_0; d) < H(M)$ for all finite d but that $\lim_{d\to\infty} \min \mathcal{P}(p_0; d) = \lim_{d\to\infty} \max \mathcal{P}(p_0; d) = H(M)$ (or, more simply, "$\lim_{d\to\infty} \mathcal{P}(p_0; d) = H(M)$". By Lemma A3, we conclude further that $\lim_{d\to\infty} \mathcal{P}(\rho; d) = \lim_{d\to\infty} \mathcal{P}(1-\rho; d) = [0, H(M)]$. And by Lemma A2, $\mathcal{P}(p_0; d) = 0$ for all $p_0 < 1-\rho$ and all $p_0 > \rho$.

Proof of part (i). Suppose that $H(M) \leq 1-\rho$ and consider any finite d. For any $p_0 \geq 1-\rho$, $\max \mathcal{P}(p_0; d) < H(M) \leq p_0$; and for any $0 < p_0 < 1-\rho$, $\mathcal{P}(p_0; d) = 0 < p_0$. Thus, no equilibrium exists with positive news veracity and the dysfunctional equilibrium is the unique equilibrium, i.e., $p_0^{*A}(d) = 0$ for all d. Clearly, then, $p_0^{*A\infty} = 0$.

Proof of part (ii). Suppose that $1-\rho < H(M) \leq \rho$ and consider any large finite d. Because news veracity cannot possibly exceed $H(M)$ when producers invest optimally, equilibrium news veracity must be less than ρ; so, $p_0^{*A\infty} < \rho$. Because $\mathcal{P}(p_0; d) < H(M)$ but $\lim_{d\to\infty} \mathcal{P}(p_0; d) = H(M)$ for all $p_0 \in (1-\rho, \rho)$, $\min \mathcal{P}(1-\rho+) > 1-\rho$ and $\max \mathcal{P}(\rho-) < \rho$ for all sufficiently large d. By continuity, there exists some $p_0(d) \in (1-\rho, \rho)$ such that $p_0(d) \in \mathcal{P}(p_0; d)$. Again because $\lim_{d\to\infty} \mathcal{P}(p_0; d) = H(M)$, the maximal such crossing-point converges to $H(M) \in (1-\rho, \rho)$ as $d \to \infty$; so, $p_0^{*A\infty} = H(M)$. (As discussed earlier, the fact that all consumers filter the news in equilibrium allows consumers to perfectly discern which stories are true, based on others' sharing behavior.)

Proof of part (iii). Suppose that $H(M) > \rho$ and consider any large finite d. Because $\lim_{d\to\infty} \mathcal{P}(\rho; d) = [0, H(M)]$, $\rho \in \mathcal{P}(\rho; d)$ for all sufficiently large d; so, an equilibrium exists with news veracity equal to ρ. And because $\mathcal{P}(p_0; d) = 0$ for all $p_0 > \rho$, no equilibrium exists with news veracity higher than ρ.[39] We conclude that the maximal equilibrium news veracity $p_0^{*A}(d) = \rho$ for all sufficiently large d, and hence that $p_0^{*A\infty} = \rho$.

It remains to show that, in the limit-market equilibrium for this case, i.e., the $d \to \infty$

[39] Because $\mathcal{P}(p_0; d)$ is single-peaked over the filtering region (Lemma A2) and $\lim_{\epsilon\to 0} \mathcal{P}(\rho - \epsilon; d) = \lim_{\epsilon\to 0} \mathcal{P}(1-\rho+\epsilon; d) = H(M)$, no equilibrium exists with news veracity in the filtering region. Thus, when $H(M) > \rho$ and d is large, all equilibria have news veracity equal to ρ, equal to $1-\rho$, or equal to 0.

limit of these equilibria with news veracity equal to ρ, consumers cannot perfectly discern which stories are true. But this follows immediately from the fact that, if consumers could perfectly discern the truth in the limit, then they would act on all true stories and no false stories and true stories would generate M units of additional revenue. Optimal investment would then induce news veracity $p_0 = H(M) > \rho$, a contradiction. $\square$

C.4 Proof of Proposition 5

By Proposition 4, either $p_0^{*V\infty} = 0$ or $p_0^{*V\infty} = 1 - \rho$. We need to show that, if $p_0^{*V\infty} = 1 - \rho$, then $p_0^{*A\infty} > 1 - \rho$. Because all stories are seen with probability at least b, true stories' extra visibility is at most $1 - b$. Consequently, when producers are paid for views, true stories generate extra revenue of at most $M(1 - b)$ and optimal producer investment can never result in news veracity greater than $H(M(1 - b))$. We conclude that, if $p_0^{*V\infty} = 1 - \rho$, then it must be that $H(M(1 - b)) \geq 1 - \rho$. But then $H(M) > 1 - \rho$ and $p_0^{*A\infty} = \min\{H(M), \rho\}$ by Thm 2(ii-iii). $\square$