Social Connectedness and the Market for News

Rachel Kranton and David McAdams[1]

this draft for 6th Annual Conference on Network Science and Economics only.

March 26, 2021

Abstract: This paper introduces a simple model of the contemporary news market where consumers share stories in social networks. In this model, consumers want to share true news and producers incur costs to produce true news. News veracity is endogenous, shaped by consumer behavior which filters and spreads news stories, based on more or less accurate private signals as to a story's truth. When producer revenues derive from the total number of consumers who view a story (e.g., revenue from accompanying advertising), veracity is high in networks that are not too dense. In highly dense networks, however, even false news spreads widely, so the incentive for high quality stories is low. Adding third-party misinformation can increase equilibrium true-news production, as consumers respond by being more judicious when sharing stories. When producer revenues come from consumers' actions based on stories (e.g, voting), veracity is higher in dense networks, and consumers make better inferences about news truth.

Keywords: social networks, news veracity, misinformation

[1]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, Oxford, 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. Producers cannot commit to a production strategy or prevent consumers from seeing their news stories. The model thus provides a benchmark for how the news market operates absent institutions and interventions, such as subscription services and curation by social medial platforms discussed in the Conclusion. We characterize outcomes first in the baseline case when producers are paid per consumer who views their story, as in advertising that accompanies a news story. We consider the impact of network connections on news veracity as well as the impact 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,[2] earn revenues from advertising accompanying their articles. While consumers might base decisions on their stories, these outlets do 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.[3] 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.[4]

Our results emphasize how consumer sharing and the producers' sources of revenues determine 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 false also stories spreads widely to consumers. Thus, news veracity is high only for an intermediate range of network connectedness.

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-news production. Knowing that false stories are being injected into the market from

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

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

[4]Britain 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 counter-protest (Allbright, 2017).

outside sources, there is a direct effect on consumers who become more judicious when deciding which stories to share. This behavior generates an indirect of misinformation on bona fide producers who have more incentive to invest in publishing high-quality news.

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.

Thus, by endogenizing the product itself, the analysis reveals that producers' incentives are key to how much consumers learn from others (in equilibrium) in a social network. When producers' revenues derive from consumers' actions, in infinitely dense networks consumers are able to determine in equilibrium which stories are true and which are false; there is a "wisdom of the crowd". On the other hand, when producers revenues derive from consumers' views, consumers essentially learn nothing from their neighbors' decisions in equilibrium, since all stories in the limit are false.

The paper contributes to three distinct literatures:

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. Consumers thereby both filter news stories and and spread stories to others. These decisions ultimately determine producers' incentives. To the best of our knowledge, this is the first paper to endogenize the product which spreads in a network setting.[5] 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.

Media markets. Much previous work on news markets studies media 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 only about the veracity of news. A key finding 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. The analysis thus serves as jumping off point for institutions that correct the problems that arise when producers earns revenues from consumers viewing

[5]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).

their stories and news travels broadly through social networks. [6]

Misinformation. 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[7] and are often online.[8] Consumers encounter false news from other sources as well, including individuals and social bots who spread conspiracy theories on social media and in memes.[9] 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).[10]

We evaluate the impact of misinformation, interpreted as the third-party broadcast

[6] 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).

[7] 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).

[8] 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/).

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

[10] 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).

of false stories on bona fide news producers' equilibrium investment in news truth. While large misinformation campaigns unambiguously undermine the news market, smaller efforts can lead to increased true-news production. Consumers pass on fewer stories, so bona fide producers' incentive to broadcast true stories increases since 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 2.2 considers how increasing social connections affects news veracity. In Section 3, we study large markets in which producers' revenues derive from consumers taking actions based on their stories. The Conclusion outlines directions for future research.

1 Model: The Market for News

The market for decision-relevant information, which we refer to as "news," consists of a large number N of consumers, of whom M generate revenue for producers.[11] 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.[12] 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.[13] (Our analysis extends easily to the case when low-quality

[11] The distinction between revenue-generating and non-revenue-generating consumers allows us to increase the number of links in the social network while holding producers' revenue base fixed; see Section 2.2.

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

[13] We focus on a context in which the object produced is a *factual claim*, but the analysis applies 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 political candidate (with "consumers" being political donors).

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 continuously differentiable 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 sees whatever news j shares, 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 revenue-generating consumer follows $d \geq 0$ others and refer to d as "social connectedness;" networks with higher d are "more connected."

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. Let p_0 denote the proportion of stories that are true, referred to as "news veracity." All produced stories are "broadcast," seen by each consumer with independent probability $b \in (0, 1]$.[14]

At $t = 1$, each consumer who see a story's broadcast decides whether to share the story with her neighbors. By assumption, no consumer can directly observe story quality but each consumer might have story-specific expertise, personal experience, or access to other information to help evaluate a story, modeled as follows: With probability $\iota \in (0, 1]$, each consumer receives an informative private signal, T or F, and with probability $1 - \iota$ receives an uninformative private signal. Informative signals match the true state with probability $\rho_i \in (1/2, 1)$, where ρ_i ("signal precision") is itself a random variable with continuous differentiable distribution $G(\cdot)$ and full support on the interval $(1/2, 1)$. Formally, each consumer i receives a signal $s_i \in \{\emptyset, T, F\}$, where $\Pr(s_i = \emptyset|\text{true}) = \Pr(s_i = \emptyset|\text{false}) = 1 - \iota$ and $\Pr(s_i = T|\text{true}) = \Pr(s_i = F|\text{false}) = \iota\rho_i$. Let $\hat{\rho} =$

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

$E[\rho_i]$ denote the ex ante likelihood that informative signals match the true state. We assume that (i) the realization of (ρ_i, s_i) is consumer i's private information, (ii) (ρ_i, s_i) is conditionally i.i.d. across consumers and across stories, and (iii) the distribution $G(\cdot)$ is common knowledge.[15]

Consumers prefer to share true stories[16] but not false stories.[17] 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, $p^S = \frac{1}{2}$ and consumers prefer to share a story if they believe it is more likely to be true than not.

At $t = 2$, consumers view the stories shared by their neighbors. Each consumer who has seen a story then 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 exceeds "action threshold" $p^A = \frac{\pi_F^A}{\pi_T^A + \pi_F^A} \in (0, 1)$. Also for simplicity, we assume that $\pi_T^A = \pi_F^A$ so that $p^A = \frac{1}{2}$. (Appendix C extends the analysis to allow consumers to have a higher or a lower standard for action than for sharing.)

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 considered in Section 2, producers earn a unit of revenue for each consumer that has viewed its story. In Section 3, we study markets where producers earn a unit of revenue for each consumer that has

[15] Assumption (i) simplifies the presentation but is not essential. The analysis is easily adapted to an alternative setting in which each consumer's expertise is observable to others; see Appendix B. Assumption (ii) guarantees that signals and sharing behavior about one story are uninformative about other stories, allowing us to consider each story in isolation.

[16] 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.

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

acted on its story. (In Appendix C, we analyze an extension in which producers earn revenue from both views and actions.)

We solve for the perfect Bayesian Nash equilibria of these games and focus on equilibria that are dynamically stable, with details of the stability concept below.

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.[18]

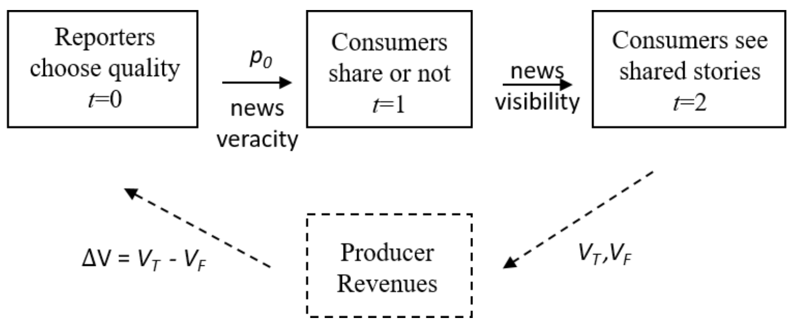

Figure 1: News Market: Revenues from Views

Figure 1 illustrates this news-market game. Producers decide whether to produce high- or low-quality news, which determines p_0. News veracity impacts consumers' sharing decisions, which determine the ex ante likelihood that any given consumer sees a 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 solve for equilibria of this game by working backward, first considering consumers' incentives to share given beliefs over news veracity, and then considering producers'

[18]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 producer revenue to be non-linear in the number of consumers who view the story.

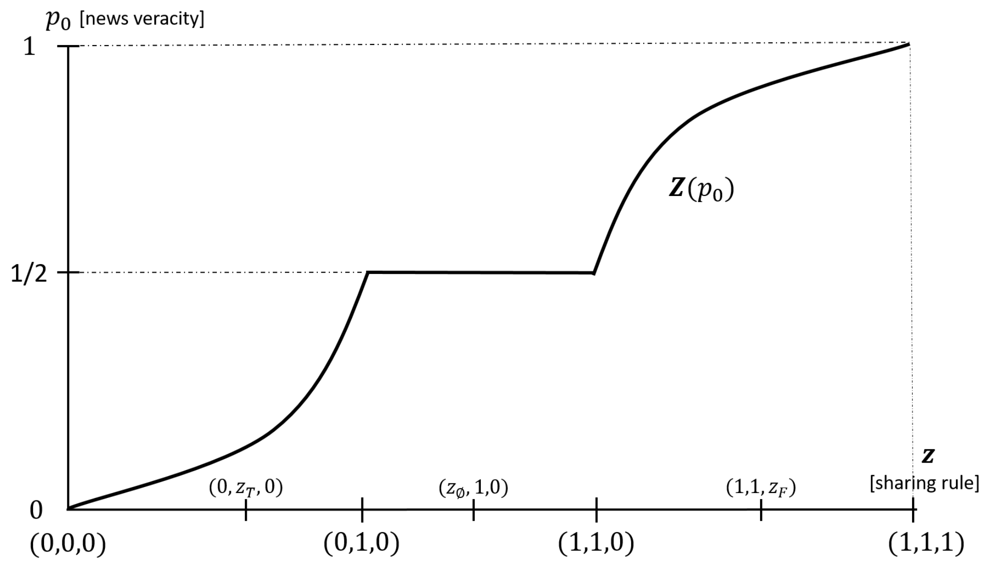

Figure 2: Illustration of consumers' best-response correspondence $\mathbf{Z}(p_0)$.

incentive to invest. Because producer revenue only depends on consumer views, we can ignore consumers' time-2 action decisions when deriving equilibria.

Consumer sharing. Consumers who see a broadcast story decide whether or not to share it. Those who receive an uninformative private signal ($s_i = \emptyset$) strictly prefer to share if $p_0 > 1/2$ and are indifferent whether to share if $p_0 = 1/2$. Those who receive an informative private signal ($s_i \in \{T, F\}$) update their beliefs about the likelihood of news truth to $p_1(s_i, \rho_i; p_0)$ using Bayes' Rule and strictly prefer to share whenever $p_1(s_i, \rho_i; p_0) > 1/2$. For each private signal $s_i \in \{\emptyset, T, F\}$, let $z_{is_i} = \Pr(i \text{ shares}|s_i)$ be the ex ante conditional likelihood that consumer i shares after receiving signal s_i. Each consumer i's sharing strategy can be summarized with the vector $\mathbf{z}_i = (z_{i\emptyset}, z_{iT}, z_{iF})$, referred to as consumer i's "sharing rule."

Let $\mathbf{Z}(p_0)$ denote the set of sharing rules that are optimal for consumers given news veracity p_0, as illustrated in Figure 2. When $p_0 = 1/2$, those who receive signal T prefer to share, those who receive signal F prefer not to share, and those who receive signal $\emptyset$ are indifferent whether to share; thus, $\mathbf{Z}(p_0) = \{(z_\emptyset, 1, 0) : z_\emptyset \in [0,1]\}$. For all $p_0 \neq 1/2$, there is a unique optimal sharing rule. When $p_0 < 1/2$, consumers prefer not to share unless they have received signal T with precision greater than $1 - p_0$; thus, $\mathbf{Z}(p_0) = (0, 1 - G(1-p_0), 0)$ for all $p_0 < 1/2$. Similarly, when $p_0 > 1/2$, consumers prefer to share unless they have received signal F with precision greater than p_0; thus, $\mathbf{Z}(p_0) = (1, 1, G(p_0))$ for all $p_0 > 1/2$.

Let $\mathcal{Z} = \cup_{p_0 \in [0,1]} \mathbf{Z}(p_0)$ denote the set of sharing rules that can potentially be optimal for consumers, those of the form $(z_\emptyset, 1, 0)$, $(0, z_T, 0)$ or $(1, 1, z_F)$ for some $z_\emptyset, z_T, z_F \in [0,1]$.

News visibility. Let $\sigma_T(\mathbf{z})$ and $\sigma_F(\mathbf{z})$ be the ex ante likelihood that each consumer shares a true and false story, respectively, when following sharing rule $\mathbf{z}$. For each story, each consumer sees the broadcast with probability b and receives an informative signal with probability ι, where infomative signals' ex ante likelihood of being T equals $\hat{\rho}$ if the story is true or $1 - \hat{\rho}$ if the story is false. Given sharing rule $\mathbf{z} = (z_\emptyset, z_T, z_F)$, each consumer's ex ante likelihood of sharing a true story is then $\sigma_T(\mathbf{z}) = b((1-\iota)z_\emptyset + \iota(\hat{\rho} z_T + (1-\hat{\rho})z_F))$. Similarly, each consumer's ex ante likelihood of sharing a false story $\sigma_F(\mathbf{z}) = b((1-\iota)z_\emptyset + \iota((1-\hat{\rho})z_T + \hat{\rho} z_F))$. True and false stories' ex ante likelihood of being seen ("visibility") by a consumer is therefore, respectively,

$$V_T(\mathbf{z}) = 1 - (1-b)(1-\sigma_T(\mathbf{z}))^d \quad (1)$$

$$V_F(\mathbf{z}) = 1 - (1-b)(1-\sigma_F(\mathbf{z}))^d. \quad (2)$$

(Each consumer sees a story unless they miss the broadcast and each of their d neighbors does not share.) The extra visibility of a true story is

$$\Delta V(\mathbf{z}) = (1-b)\left((1-\sigma_F(\mathbf{z}))^d - (1-\sigma_T(\mathbf{z}))^d\right). \quad (3)$$

Producer investment. Next, we turn to each producer's decision whether to incur reporting cost c_R ("invest") 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" $M\Delta V(\mathbf{z})$.

Producers maximize expected profit by investing whenever $c_R < M\Delta V(\mathbf{z})$, which results in fraction $P(\mathbf{z}) \equiv H(M\Delta V(\mathbf{z}))$ of all stories being true. $P(\mathbf{z})$ is the news-veracity level that arises when producers invest optimally in response to consumer sharing rule $\mathbf{z}$, what we call the "producers' best-reply."

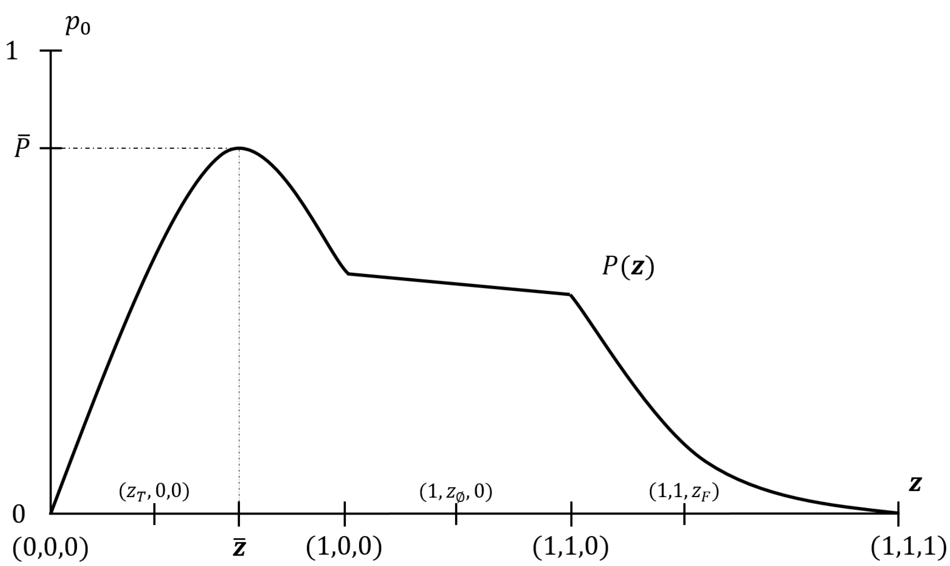

Figure 3: Illustration of producers' best-reply $P(\mathbf{z})$.

Figure 3 illustrates how producers' best reply varies with consumer sharing, over the range of potentially-optimal sharing rules, i.e., those of the form $(z_\emptyset, 1, 0)$, $(0, z_T, 0)$, or $(1, 1, z_F)$ for some $z_\emptyset, z_T, z_F \in [0, 1]$. Some observations:

First, $P(0, 0, 0) = P(1, 1, 1) = 0$. If true and false stories are equally shared, producers have no incentive to invest in news quality.

Second, starting at the left-hand-side of the figure, $P(0, z_T, 0)$ is increasing in z_T up to a threshold $\overline{z}_T \leq 1$, at sharing rule $\overline{\mathbf{z}} = (0, \overline{z}_T, 0)$, after which it is decreasing in z_T. (If $\overline{z}_T = 1$, then $P(0, z_T, 0)$ is increasing in z_T over the entire interval $[0, 1]$.) Because a signal T is more likely when a story is true, consumer sharing after a signal T magnifies the visibility of true stories more than false stories, but only when sharing is sufficiently rare at the lower levels of z_T (and/or social connectedness d is sufficiently low).[19] Let $\overline{P}$ denote the *maximal best-response news veracity*, the highest news veracity that can arise from optimal producer investment, achieved at sharing rule $\overline{\mathbf{z}}$:

$$\overline{P} \equiv \max_{\mathbf{z}} P(\mathbf{z}) = P(\overline{\mathbf{z}}) = H\left(M \Delta V(\overline{\mathbf{z}})\right).$$

Third, moving towards the right, $P(z_\emptyset, 1, 0)$ is decreasing in $z_\emptyset$ and $P(1, 1, z_F)$ is decreasing in z_F. To see why $P(1, 1, z_F)$ is decreasing in z_F, recall that signal F is more likely when a story is false. Sharing after signal F therefore increases the spread of false stories more than it increases the spread of true stories, reducing producers' incentive to invest in news truth. The reason why $P(z_\emptyset, 1, 0)$ is decreasing in $z_\emptyset$ is similar, although less obviously so, since signal $\emptyset$ is equally likely for true and false stories. However, because true stories have a "head start" in being shared, the incremental impact on visibility from increased uninformed sharing is greater for false stories than for true stories. See the proof of Lemma 1 for details.

Lemma 1 gathers together useful facts about the producers' best-reply.

Lemma 1. *(i)* $P(0, 0, 0) = 0$ *and* $P(0, z_T, 0)$ *is strictly increasing in* z_T *over the interval*

[19] We show in the proof of Lemma 1 that the critical value $\overline{z}_T$ is strictly less then one whenever d is sufficiently large. However, the magnitudes involved are not big. For example, if $b = \frac{1}{2}$ and ρ_i has a degenerate distribution at $\rho = \frac{2}{3}$, then $\overline{z}_T \approx 0.823$ if $d = 5$ and $\overline{z}_T \approx 0.53$ if $d = 9$. To see why, suppose that consumers use a sharing rule of the form $(0, z_T, 0)$ for some $z_T \approx 1$. Each consumer's ex ante likelihood of *not* viewing a true or false story is, respectively, $1 - V_T(0, z_T, 0) \approx (1 - b)(1 - b\hat{\rho})^d$ and $1 - V_F(0, z_T, 0) \approx (1 - b)(1 - b(1 - \hat{\rho}))^d$. The relative likelihood that false stories are not seen, $\frac{1 - V_F(0, z_T, 0)}{1 - V_T(0, z_T, 0)}$, grows exponentially with d. When d is large, increased sharing (even if only after signal T) therefore increases the visibility of false stories more than true ones, reducing producers' incentive to invest.

$[0, \overline{z}_T]$ *and strictly decreasing in* z_T *over the interval* $[\overline{z}_T, 1]$ *for some* $\overline{z}_T \in (0, 1]$. *(ii)* $P(z_\emptyset, 1, 0)$ *is non-increasing in* $z_\emptyset$*, and strictly decreasing in* $z_\emptyset$ *if* $\iota < 1$. *(iii)* $P(1, 1, z_F)$ *is strictly decreasing in* z_F*, with* $P(1, 1, 1) = 0$. *(iv)* $P(\mathbf{z})$ *is maximized at* $\mathbf{z} = (0, \overline{z}_T, 0)$.

Proof. See the Appendix. □

2.1 Equilibrium characterization

This section characterizes and examines the equilibria of the news-market game, focusing especially on the equilibria with the highest and lowest news veracity. Without loss of generality, we focus on equilibria in which all consumers use the same sharing rule.[20]

An equilibrium exists with news veracity p_0 and sharing rule $\mathbf{z}$ if and only if $p_0 = P(\mathbf{z})$ and $\mathbf{z} \in \mathbf{Z}(p_0)$. That is, given consumer sharing rule $\mathbf{z}$, optimal producer investment results in veracity $p_0 = P(\mathbf{z})$. And given news veracity p_0, consumers use a sharing rule in the best-response set $\mathbf{Z}(p_0)$. Figure 4 illustrates consumers' and producers' best-reply curves, and the resulting equilibrium set, in two scenarios. In each panel, the x-axis depicts the set $\mathcal{Z}$ of potentially-optimal sharing rules, while the y-axis depicts the range of feasible news-veracity levels. The black line depicts $P(\mathbf{z})$, the producers' best-reply curve. The blue line depicts $\mathbf{Z}(p_0)$, the consumer-sharing best-reply correspondence.

Equilibria may or may not be "stable" in the sense of evolutionary stability (Samuelson and Zhang (1992)). An equilibrium $(p_0, \mathbf{z})$ is stable[21] if $P(Z(p_0 + \epsilon)) < p_0 + \epsilon$ and $P(Z(p_0 - \epsilon)) > p_0 - \epsilon$ for all sufficiently small $\epsilon > 0$, and unstable if $P(Z(p_0 + \epsilon)) > p_0 + \epsilon$ or $P(Z(p_0 - \epsilon)) < p_0 - \epsilon$ for all sufficiently small $\epsilon > 0$. Graphically, an equilibrium is stable (or unstable) if the producer best-reply curve is more (or less) downward-sloping than the consumer best-reply curve at the equilibrium point, i.e., the producer best-reply

[20] Equilibria with asymmetric sharing can exist but, as shown in the proof of Thm 1, the news veracity in such equilibria is never higher than in the maximal-veracity equilibrium with symmetric sharing.

[21] We relegate the formal definition of stability to the Appendix for ease of exposition.

curve crosses from above (or from below).

For instance, in Figure 4 panel (a), the equilibrium with zero news veracity is unstable. Why? Consider introducing a small mass of producers who always produce true stories, increasing news veracity from 0 to $\epsilon \approx 0$. Consumers will then find it optimal to share whenever they receive a sufficiently precise T signal, causing true stories to spread more widely than false ones. This in turn induces producers with sufficiently low costs to invest, increasing true-news production and causing still more consumer sharing. Because $P(\mathbf{z})$ is steeper than $\mathbf{Z}^{-1}(\mathbf{z})$ at the origin, this mutual-adaptation process leads investment and sharing to spiral upward and to the right, away from the original equilibrium. As we elaborate in the next section, this instability depends on the social connectedness of consumers' social network, which affects the spread of stories and hence the slope of the producer best reply.

Equilibria in Figure 4 are denoted by circles, with filled circles for stable equilibria and unfilled circle for unstable ones. Let p_0^* denote the highest news veracity in any equilibrium, what we call the "maximal equilibrium news veracity." Let $\mathbf{z}^*$ be consumers' sharing rule in the equilibrium that achieves news veracity p_0^*.

We establish several results concerning the equilibrium set, summarized below in Theorem 1.

First, a "dysfunctional equilibrium" with zero news veracity always exists. This equilibrium is simple to construct: If all stories are false, so that $p_0 = 0$, consumers find it optimal to never share. And if consumers never share, all stories have equal visibility and producers have no incentive to invest in news truth. This equilibrium is pictured at the origin in each of the panels of Figure 4.

Second, some false stories are produced in any equilibrium. To see why, suppose that an equilibrium existed in which all stories were true. Consumers would then share any story they see, regardless of their private signal, causing true and false stories to be equally shared and hence equally viewed. Producers would then have no incentive to invest in news truth, causing all stories to be false, a contradiction.

Third, there is at most one equilibrium in which news veracity exceeds 1/2. If $p_0 > 1/2$, then consumers find it optimal to share whenever *either* they have received the signal T *or* they are uninformed *or* they received signal F with precision $\rho_i < p_0$. In particular, consumers' equilibrium sharing rule must take the form $\mathbf{z} = (1, 1, z_F)$, with $z_F = G(p_0) \in (0, 1)$. Over this range of sharing rules, producers' best-reply $P(\mathbf{z})$ is strictly decreasing in z_F (Lemma 1) while consumers' inverse best-reply $\mathbf{Z}^{-1}(\mathbf{z})$ is strictly increasing in z_F. Thus, there is at most one crossing point of the best replies with veracity greater than 1/2. Moreover, because $\mathbf{Z}^{-1}(\mathbf{z})$ and $P(\mathbf{z})$ are continuous and $\mathbf{Z}^{-1}(1, 1, 1) = 1 > 0 = P(1, 1, 1)$, such an equilibrium exists if and only if $P(1, 1, 0) > \mathbf{Z}^{-1}(1, 1, 0) = 1/2$.

Fourth, $p_0^* = 1/2$ if and only if $P(0, 1, 0) \geq 1/2 \geq P(1, 1, 0)$. If $P(1, 1, 0) > 1/2$, then an equilibrium exists with news veracity greater than 1/2 (shown in the previous point). On the other hand, if $P(0, 1, 0) < 1/2$, the fact that $P(z_\emptyset, 1, 0)$ is non-increasing in $z_\emptyset$ and $P(1, 1, z_F)$ is strictly decreasing in z_F (Lemma 1) implies that all equilibria must have news veracity strictly less than 1/2. The only remaining possibility is when $P(0, 1, 0) \geq 1/2 \geq P(1, 1, 0)$. In this case, the fact that $P(z_\emptyset, 1, 0)$ is non-increasing and continuous in $z_\emptyset$ implies that there exists $\hat{z}_\emptyset \in [0, 1]$ such that $P(\hat{z}_\emptyset, 1, 0) = 1/2$. Since consumers find all sharing rules of the form $(\hat{z}_\emptyset, 1, 0)$ to be optimal when $p_0 = 1/2$, we conclude that an equilibrium exists with news veracity equal to 1/2 and hence that 1/2 is the maximal equilibrium news veracity, i.e., $p_0^* = 1/2$.

Theorem 1 collects these main findings concerning the set of equilibria.

Theorem 1. *(i)* $p_0^* < 1$. *(ii)* $p_0^* > 1/2$ *if and only if* $P(1, 1, 0) > 1/2$ *and, if so, there is a unique equilibrium with news veracity* $p_0 > 1/2$. *(iii)* $p_0^* = 1/2$ *if and only if*

$P(0,1,0) \geq 1/2 \geq P(1,1,0)$. *(iv) A "dysfunctional equilibrium" always exists in which all stories are false and consumers never share.*

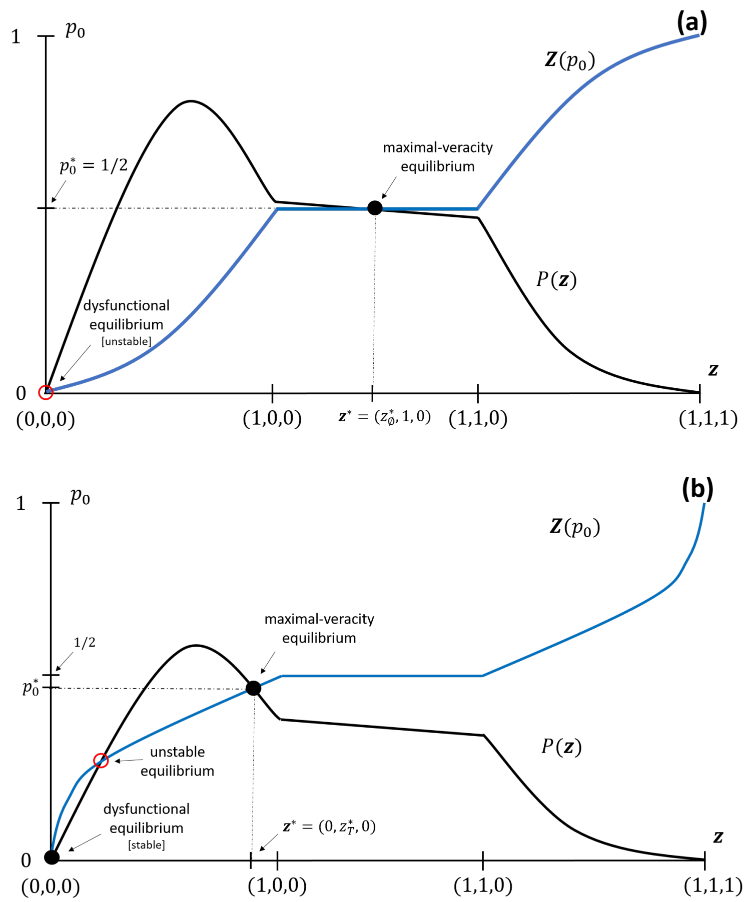

Figure 4: Illustration of the equilibrium set in two scenarios. In panel (a), the dysfunctional equilibrium is unstable and the maximal equilibrium news veracity $p_0^* = 1/2$. In panel (b), the dysfunctional equilibrium is stable and $p_0^* < 1/2$.

2.2 Social Connectedness and News Veracity

This section examines how social connectedness, d, impacts the equilibrium set and, in particular, how much true news can be produced in equilibrium. Our main finding is that the maximal equilibrium news veracity, denoted here as $p_0^*(d)$, is non-monotone in social connectedness, with zero (or nearly zero) true stories in equilibrium whenever d is sufficiently small or sufficiently large.

First, if consumers are unconnected ($d = 0$), then no true news is ever produced in equilibrium. If $d = 0$, then true and false stories are equally visible and producers have zero incentive to invest. The only equilibrium in this case is the dysfunctional equilibrium, with zero producer investment and zero consumer sharing; that is, $p_0^*(0) = 0$.

Second, an equilibrium with positive news veracity exists whenever consumers' social network is sufficiently connected, in particular, whenever $d > \hat{d}$ where

$$\hat{d} \equiv \frac{1}{g(1)h(0)Mb(1-b)\iota(2\hat{\rho}-1)}.$$

To understand this threshold, consider how social connectedness affects the slope of the producers' best-reply curve versus the consumers' best-reply curve at the origin. Because $P(\mathbf{z}) = H(M\Delta V(\mathbf{z}))$ and $\Delta V(0, z_T, 0) = (1-b)\left((1-b\iota(1-\hat{\rho})z_T)^d - (1-b\iota\hat{\rho}z_T)^d\right)$, we have

$$\frac{\mathrm{d}P(0,0,0)}{\mathrm{d}z_T} = h(0)M \times db(1-b)\iota(2\hat{\rho}-1).$$

Because $\mathbf{Z}(p_0) = (0, 1-G(1-p_0), 0)$ for $p_0 < 1/2$, we have $\frac{\mathrm{d}\mathbf{Z}^{-1}(0,0,0)}{\mathrm{d}z_T} = \frac{1}{g(1)}$. We conclude that $\frac{\mathrm{d}P(0,0,0)}{\mathrm{d}z_T} > \frac{\mathrm{d}\mathbf{Z}^{-1}(0,0,0)}{\mathrm{d}z_T}$ if and only if $d > \hat{d}$. This implies that, whenever $d > \hat{d}$, (i) the dysfunctional equilibrium is unstable and (ii) $P(0, z_T, 0) > \mathbf{Z}^{-1}(0, z_T, 0)$ for all small z_T. But because $P(1,1,1) = 0 < 1 = \mathbf{Z}^{-1}(1,1,1)$, condition (ii) in turn implies there must be a sharing rule $\mathbf{z} > 0$ satisfying $P(\mathbf{z}) = \mathbf{Z}^{-1}(\mathbf{z})$ and hence supported in equilibrium. So, $p_0^*(d) > 0$ for all $d > \hat{d}$.

Third, an equilibrium in which the majority of stories are true ($p_0 > 1/2$) can only exist, if at all, when social connectedness is neither too high nor too low. By Theorem 1, such an equilibrium exists if and only if $P(0,1,0) = H(M\Delta V(0,1,0)) > 1/2$. But

$\Delta V(0,1,0) = (1-b)\left((1-b\iota(1-\hat{\rho}))^d - (1-b\iota\hat{\rho})^d\right)$ is single-peaked[22] in d and goes to zero as d goes to infinity. Thus, *either* all equilibrium have news veracity less than 1/2 regardless of social connectedness, i.e., $p_0^*(d) < 1/2$ for all d, *or* $p_0^*(d) \geq 1/2$ iff social connectedness lies in an intermediate interval.

Finally, the maximal equilibrium news veracity goes to zero as d goes to infinity, implying that each consumer's ex ante likelihood of sharing converges to zero as well. Details are in the Appendix but, to gain intuition, suppose for the sake of contradiction that equilibrium news veracity greater than $\hat{p}_0 > 0$ could be supported in the $d \to \infty$ limit. In such an equilibrium, consumers must use a sharing rule greater than $\mathbf{z}(\hat{p}_0) > 0$, causing the ex ante likelihood that any given consumer shares to be at least $\sigma_T(\mathbf{z}(\hat{p}_0)) > 0$ and $\sigma_F(\mathbf{z}(\hat{p}_0)) > 0$ for true and false stories, respectively. But then the visibility of true and false stories must both converge to one in the $d \to \infty$ limit (see equations (1,2)), implying that the extra visibility of true stories must converge to zero. Producers must therefore have zero incentive to invest in the limit, contradicting the presumption that news veracity does not vanish in the limit.

Theorem 2. *When producers are paid for views, the maximal news veracity is non-monotonic in consumer connectedness: (i)* $p_0^*(0) = 0$*; (ii)* $p_0^*(d) > 0$ *for finite* $d > \hat{d}$*; (iii)* $p_0^*(d) > 1/2$ *when* d *is neither too large nor too small, or potentially not at all for any* d*. (iv)* $\lim_{d\to\infty} p_0^*(d) = 0$*;*

Proof. Parts (i-iii) were proven in the text. Part (iv) is proven in the Appendix. □

2.3 Misinformation and Equilibrium News Truth

This section examines the impact of misinformation on the news market.

[22] $\frac{\mathrm{d}\left((1-b\iota(1-\hat{\rho}))^d-(1-b\iota\hat{\rho})^d\right)}{\mathrm{d}d} > 0$ iff $\frac{\log(1-b\iota\hat{\rho})}{\log(1-b\iota(1-\hat{\rho}))} > \left(\frac{1-b\iota(1-\hat{\rho})}{1-b\iota\hat{\rho}}\right)^d$. Since $\frac{\log(1-b\iota\hat{\rho})}{\log(1-b\iota(1-\hat{\rho}))}$ and $\left(\frac{1-b\iota(1-\hat{\rho})}{1-b\iota\hat{\rho}}\right)^d$ each exceed one, this condition holds only up to some critical threshold $\overline{d}$. We conclude that $\Delta V(0,1,0)$ and hence $P(0,1,0)$ is increasing in d up to $\overline{d}$ and then decreasing in d.

Suppose that there is a mass $m \geq 0$ of "misinformation agents" who only produce false stories, in addition to a unit mass of "bona fide producers" with reporting costs drawn from distribution $H(\cdot)$. Total quantity $1+m$ of stories is produced, fraction $\frac{1}{1+m}$ by bona fide producers.[23] An equilibrium exists with news veracity p_0 and sharing rule $\mathbf{z}$ if and only if $p_0 = \frac{P(\mathbf{z})}{1+m}$ and $\mathbf{z} \in \mathbf{Z}(p_0)$.[24]

Let $p_0^*(m)$ denote the maximal equilibrium news veracity, viewed now as a function of misinformation volume m, and let $\mathbf{z}^*(m)$ be the corresponding equilibrium sharing rule. Increasing the volume of misinformation can have three sorts of equilibrium effects, depending on how much misinformation is already in the market and how much true news can be supported in equilibrium absent misinformation.

Possibility #1: Dysfunction. Define $\overline{m} \equiv \sup_{\mathbf{z}>(0,0)\in\mathcal{Z}} \frac{P(\mathbf{z})}{\mathbf{Z}^{-1}(\mathbf{z})} - 1$, or $\overline{m} \equiv 0$ if $P(\mathbf{z}) < \mathbf{Z}^{-1}(\mathbf{z})$ for all non-zero sharing rules $\mathbf{z}$ in the potentially-optimal set $\mathcal{Z}$. If $m > \overline{m}$, then $P(\mathbf{z}; m) \equiv \frac{P(\mathbf{z})}{1+m}$ is everywhere below consumers' best-reply curve except at the origin, i.e., $P(\mathbf{z}) < \mathbf{Z}^{-1}(\mathbf{z})$ for all non-zero $\mathbf{z} \in \mathcal{Z}$, as in Figure 5(b). Thus, the dysfunctional equilibrium is the unique equilibrium; in this case, $p_0^*(m) = 0$ and $\mathbf{z}^*(m) = \mathbf{0} \equiv (0,0,0)$. On the other hand, if $m < \overline{m}$, then $\hat{\mathbf{z}} > \mathbf{0}$ exists such that $P(\hat{\mathbf{z}}) > \mathbf{Z}^{-1}(\hat{\mathbf{z}})$, implying that an equilibrium exists with sharing rule greater than $\hat{\mathbf{z}}$; in this case, $p_0^*(m) > \mathbf{Z}^{-1}(\hat{\mathbf{z}}) > 0$.

Possibility #2: More misinformation leads to more true news if $\mathbf{z}^*(m) > \overline{\mathbf{z}}$. Define $\underline{m} \equiv \max\{m \geq 0 : \mathbf{z}^*(m) \geq \overline{\mathbf{z}}\}$, or $\underline{m} \equiv 0$ if $\mathbf{z}^*(0) < \overline{\mathbf{z}}$. If $m < \underline{m}$, then producers' best reply $P(\mathbf{z})$ is decreasing in $\mathbf{z}$ at the equilibrium sharing rule $\mathbf{z}^*(m)$. Marginally

[23] As discussed in footnote 12, 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).

[24] Our previous analysis without misinformation can be re-interpreted to characterize all equilibria for any given $m > 0$. In particular, suppose that producers' costs are drawn from c.d.f. $\widehat{H}(c_R) = \frac{H(c_R)}{1+m}$, i.e. there is an atom of mass $\frac{m}{1+m}$ at $c_R = \infty$. The producer best-reply curve is then $\widehat{P}(\mathbf{z}) = \frac{P(\mathbf{z})}{1+m}$, generating the same set of equilibria as in our model here with mass $m > 0$ of misinformation agents.

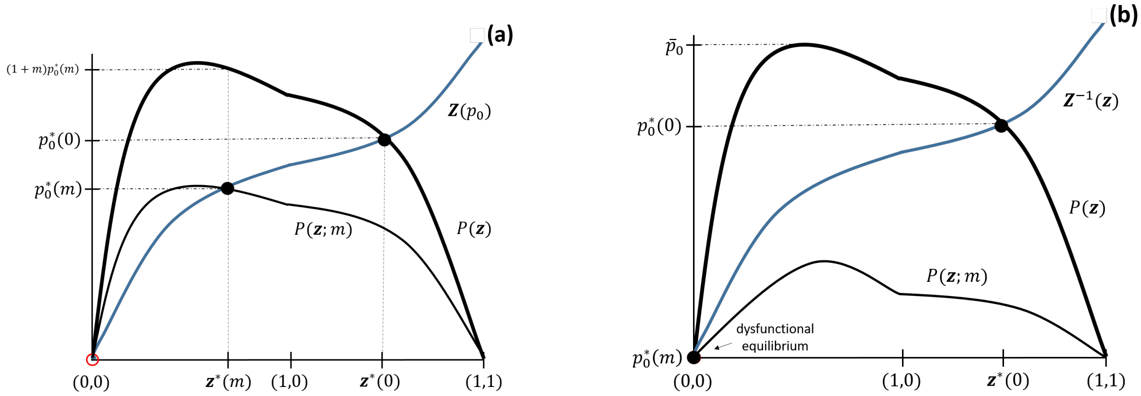

Figure 5: The impact of misinformation. In panel (a), news veracity falls from $p_0^*(0)$ to $p_0^*(m)$ but the volume of true news produced in the maximal-veracity equilibrium increases from $p_0^*(0)$ to $(1+m)p_0^*(m)$. In panel (b), the quantity of misinformation is sufficiently large that only the dysfunctional equilibrium remains. EDIT FIGURE.

increasing the quantity of misinformation therefore leads to less consumer sharing but *more* true-news production by bona-fide producers in the maximal-veracity equilibrium. See Figure 5(a), where increasing misinformation volume from 0 to m causes equilibrium sharing to fall from $\mathbf{z}^*(0)$ to $\mathbf{z}^*(m)$ and the volume of true-news production to increase from $p_0^*(0)$ to $(1+m)p_0^*(m)$.

Possibility #3: More misinformation leads to less true news if $\mathbf{0} < \mathbf{z}^*(m) < \bar{\mathbf{z}}$*.* Finally, if $\underline{m} < m < \overline{m}$, then an equilibrium exists with positive news veracity but producers' best reply $P(\mathbf{z})$ is increasing in $\mathbf{z}$ at the equilibrium sharing rule $\mathbf{z}^*(m)$. Marginally increasing the quantity of misinformation now leads to less consumer sharing and *less* true-news production by bona-fide producers in the maximal-veracity equilibrium.

Proposition 1 summarizes these findings:

Proposition 1. *(i) If $m < \underline{m}$, then the quantity of true news in the maximal-veracity equilibrium is increasing in m. (ii) If $\underline{m} < m < \overline{m}$, then the quantity of true news in the maximal-veracity equilibrium is decreasing in m. (iii) If $m > \overline{m}$, then the dysfunctional equilibrium is the unique equilibrium.*

3 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 or buying a product. Our key finding in this setting is that, in the $d \to \infty$ limit as consumers follow many others, (i) an equilibrium "wisdom of the crowd" emerges, meaning that consumers can perfectly discern in the limit which stories are true, and (ii) producers' incentive to invest in news truth is as high as it can possibly be. By contrast, when producers are paid for views, our previous analysis showed that all stories are false in the $d \to \infty$ limit (Theorem 2).

Figure 6 illustrates the news-market game when producers are paid for actions. Producers decide whether to invest in news truth (at time $t = 0$) and consumers who see a story's broadcast decide whether to share (at time $t = 1$). Consumers then make inferences about the story's truth based on their neighbors' sharing behavior and decide whether or not to act on the story (at time $t = 2$).

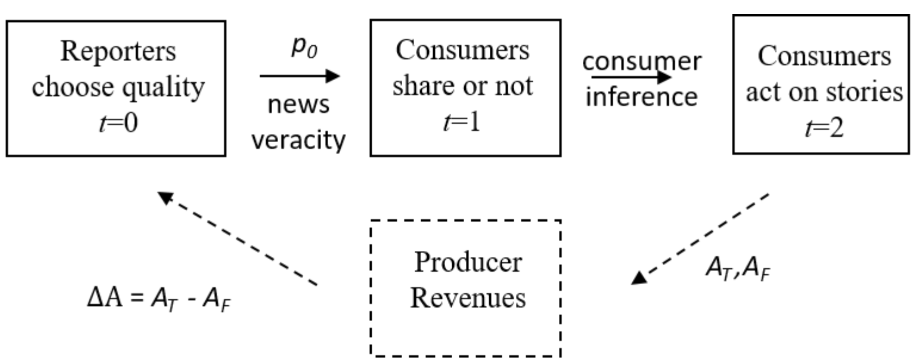

Figure 6: Schematic of an actions-supported news market.

Consumer learning and action. What one learns from other consumers' sharing behavior depends on the sharing rule that they are using and how many others one is following. Let $o_i = (s_i, \rho_i, d_i)$ be shorthand for consumer i's "observation" prior to deciding whether or not to act, consisting of their private signal $s_i \in \{T, F\}$ and its precision ρ_i

as well as the number d_i out of d neighbors who share a story. Let $p_2(o_i; p_0, \mathbf{z}, d)$ denote consumer i's Bayes-updated belief at time $t = 2$, after observation o_i, when news veracity is p_0, others are using sharing rule $\mathbf{z}$, and social connectedness equals s. (The formula for $p_2(o_i; p_0, \mathbf{z}, d)$ is provided in the Appendix.)

Let $A_T(p_0; \mathbf{z}, d) = \Pr(p_2(o_i; p_0, \mathbf{z}, d) > 1/2 | true)$ denote the ex ante likelihood that consumer i's time-2 belief will be high enough to act when the story is true. Similarly, let $A_F(p_0; \mathbf{z}, d) = \Pr(p_2(o_i; p_0, \mathbf{z}, d) > 1/2 | false)$ denote each consumer's likelihood of acting on false stories. $\Delta A(p_0; \mathbf{z}, d) = A_T(p_0; \mathbf{z}, d) - A_F(p_0; \mathbf{z}, d)$ is the extra likelihood that consumers act on true stories. $M\Delta A(p_0; \mathbf{z}, d)$ is the extra expected revenue earned by true stories.

Producer investment. Producers find it optimal to invest and produce a true story whenever their cost c_R is less than $M\Delta A(p_0; \mathbf{z}, d)$, which occurs with ex ante probability $H(M\Delta A(p_0; \mathbf{z}, d))$. Because $A_T(p_0; \mathbf{z}, d) \leq 1$ and $A_F(p_0; \mathbf{z}, d) \geq 0$, the extra expected revenue earned by true stories can never exceed M. News veracity is therefore bounded above by $p^{max} \equiv H(M)$. We will refer to p^{max} as "the maximal-possible news veracity."

Let $P_A(\mathbf{z}; d) = H(M\Delta A(\mathbf{Z}^{-1}(\mathbf{z}); \mathbf{z}, d))$. $P_A(\mathbf{z}; d)$ is the ex ante likelihood that producers optimally invest, given that consumers use sharing rule $\mathbf{z}$, believe that news veracity equals $\mathbf{Z}^{-1}(\mathbf{z})$ (making sharing rule $\mathbf{z}$ optimal), and optimally choose whether or not to act. We refer to $P_A(\mathbf{z}; d)$ as "the producer best-reply curve." Lemma 2 establishes some important basic properties of the producer best-reply.

Lemma 2. *(i)* $P_A(\mathbf{0}; d) = 0$ *and* $P_A(\mathbf{1}; d) = 0$ *for all* d. *(ii) Consider any* $\mathbf{z} \in \mathcal{Z} \setminus \{\mathbf{0}, \mathbf{1}\}$. $P_A(\mathbf{z}; d) > 0$ *for all* d *and* $\lim_{d\to\infty} P_A(\mathbf{z}; d) = p^{max}$.

Proof. In the Appendix. □

The second part of Lemma 2 is especially important. To gain intuition, suppose that consumers believe that news veracity equals $\hat{p}_0 \in (0,1)$ and so use the sharing rule $\hat{\mathbf{z}} = \mathbf{Z}^{-1}(\mathbf{z}) \in \mathcal{Z} \setminus \{\mathbf{0}, \mathbf{1}\}$. Because $\hat{z}_T > \hat{z}_F$, true stories are more likely to be shared than false ones, making each consumer's sharing behavior an informative "social signal" about news truth. Having many neighbors allows a consumer to observe many (conditionally i.i.d.) social signals and hence, by the Law of Large Numbers, make arbitrarily accurate inferences about which stories are true. Anticipating that such a "wisdom of the crowd" will emerge, producers then have maximal incentive to invest in news truth in $d \to \infty$ limit, since only true stories will be acted upon, generating news veracity p^{max} in the limit. Consequently, an equilibrium cannot exist with news veracity $\hat{p}_0 \in (0, p^{max})$ whenever d is sufficiently large.

Equilibrium. An equilibrium with news veracity p_0 and sharing rule $\mathbf{z}$ exists if and only if (i) $\mathbf{z}$ is an optimal sharing rule given news veracity p_0, i.e., $\mathbf{z} \in \mathbf{Z}(p_0)$, and (ii) consumers' extra likelihood of acting on true stories given news veracity p_0 and sharing rule $\mathbf{z}$ induces producers to invest just enough to support news veracity p_0, i.e., $H(M\Delta A(p_0; \mathbf{z})) = p_0$. As previously, let p_0^* denote the maximal equilibrium news veracity.

Consider any potentially-optimal sharing rule $\mathbf{z} \in \mathcal{Z}$. An equilibrium exists with this sharing rule if and only if two conditions hold. First, news veracity must equal $\mathbf{Z}^{-1}(\mathbf{z})$, since only then do consumers find it optimal to use that sharing rule. Second, sharing rule $\mathbf{z}$ and news veracity $\mathbf{Z}^{-1}(\mathbf{z})$ must together induce consumers to act in a way that producers invest just enough to support news veracity $\mathbf{Z}^{-1}(\mathbf{z})$. In other words, the set of equilibria corresponds to the set of crossing-points of the consumer best-reply curve (same as in Section 2) and the producer best-reply curve (defined above).

3.1 Social Connectedness and the Wisdom of the Crowd

From consumers' point of view, the ideal outcome in a news market would be (i) for producers to invest maximally in news truth, resulting in news veracity p^{max} ("maximal producer investment"), and (ii) for consumers to be able to discern which stories are true and which are false ("wisdom of the crowd"). Here we show that, if producers are paid for actions, this combination is achieved in the limit as social connectedness d goes to infinity. In particular, once d is sufficiently large, all stable equilibria of an actions-supported news market are "nearly ideal" in the sense that news veracity is approximately equal to p^{max} and consumers are nearly certain to act on any true story but nearly certain not to act on any false story (Theorem 3). By contrast, views-supported news markets devolve into dysfunction as d goes to infinity, with nearly all published stories being false in all equilibria once d is large (Theorem 2).

Theorem 3. *When producers are paid for actions and d is large, the maximal possible news veracity can be nearly achieved in equilibrium:* $\lim_{d\to\infty} p_0^*(d) = p^{max}$.

Proof. Consider any news-veracity level $\hat{p}_0 \equiv p^{max} - \epsilon$ for some small $\epsilon > 0$, and let $\hat{\mathbf{z}} \equiv \mathbf{z}(\hat{p}_0)$ denote consumers' best-reply sharing rule. Because $\hat{p}_0 \in (0,1)$, we have $\lim_{d\to\infty} P_A(\hat{\mathbf{z}}; d) = p^{max}$ (Lemma 2(ii)), and hence that $P_A(\hat{\mathbf{z}}; d) > \hat{p}_0$ for all sufficiently large d. Since $P_A(\mathbf{z}; d)$ is continuous in $\mathbf{z}$ and $P_A(\mathbf{1}; d) = 0$ for all d (Lemma 2(i)), we conclude that an equilibrium exists with news veracity greater than any given $\hat{p}_0 \approx p^{max}$ whenever d is sufficiently large. In particular, the maximal equilibrium news veracity $p_0^*(d)$ converges to p^{max} as d goes to infinity. By definition of p^{max}, producers only have the incentive to invest enough to support news veracity p^{max} if true stories are certain to be acted on and false stories are certain not to be acted on, i.e., if there is a "wisdom of the crowd." The fact that $\lim_{d\to\infty} p_0^*(d) = p^{max}$ therefore implies that consumers also enjoy a wisdom of the crowd in the maximal-veracity equilibrium in the $d \to \infty$ limit. □

Theorem 3 establishes desirable properties of the maximal-veracity equilibrium when consumers are highly connected, but what about other equilibria? As it turns out, p^{max} is the *only* news-veracity level that can be arise in stable equilibrium as $d \to \infty$.

Equilibria do not exist with intermediate news veracity when d is large. Fixing any news veracity $\tilde{p}_0 \in (0, p^{max})$, suppose for the sake of contradiction that there were a sequence of equilibria with news veracity $p_0(d)$ converging to $\tilde{p}_0$. In such equilibria, each consumer's decision whether or not to share (the "social signal" their neighbors receive) would have non-vanishing equilibrium informativeness. Each consumer must then be able to perfectly discern in the $d \to \infty$ limit which stories are true and which are false. But then producer investment must converge to the maximal possible level, so that equilibrium news veracity converges to p^{max} in the limit, a contradiction.

The dysfunctional equilibrium is unstable when d is large. Much as in the revenue-from-views model of Section 2, the dysfunctional equilibrium becomes unstable once consumers are sufficiently connected. The difference now is that the social signals spread over consumers' social network spread not just *awareness* but also *confidence* in true stories more than false stories.

To see the point, suppose that consumers were to believe that news veracity equals $\epsilon \approx 0$. Each consumer will share a story if and only if they have gotten an "extremely favorable signal," i.e., signal T with precision $\rho_i > 1 - \epsilon$. This rarely occurs, but is an order of magnitude rarer when a story is false. In particular, a consumer who sees the broadcast will get an extremely favorable signal with probability $(1-G(1-\epsilon))\Pr(T|\rho_i > 1-\epsilon, true) > b(1-G(1-\epsilon))(1-\epsilon) \approx \epsilon b g(1)$ if the story is true, or with probability $b(1-G(1-\epsilon))\Pr(F|\rho_i > 1-\epsilon, false) < b(1-G(1-\epsilon)\epsilon \approx \epsilon^2 b g(1)$ if the story is false. Anyone who sees a neighbor sharing will infer that they must have gotten an extremely favorable signal, and update their own beliefs accordingly. In particular, because an extremely favorable signal is (by definition) enough to make one believe that the story is more likely true than false, anyone who sees someone sharing a story will have enough confidence in that story to act as long as their own private signal is not too negative. In this way, sharing over the social network creates a "call-to-action effect" that differentially benefits true stories, and that grows larger as social connectedness grows.

This intuition is formalized in the Appendix, in the proof of Lemma 3.

Lemma 3. *The dysfunctional equilibrium is unstable whenever social connectedness* $d > 2\hat{d}$*, where* $\hat{d}$ *is the threshold in Theorem 2.*

Proof. In the Appendix. □

4 Conclusion

This paper analyzes the market 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. In both types of markets, consumer connectedness profoundly impacts the qualitative features of equilibrium outcomes.

In a views-supported news market, the amount of true news produced in equilibrium is non-monotone in social connectedness. If consumers are initially unconnected, adding social links induces producers to invest more in news truth, since true stories are more likely to be shared (and hence viewed) than false ones. However, once consumers follows sufficiently many others, adding still more links reduces producers' equilibrium incentive to invest, as even false stories spread over the network and are widely viewed. In particular, false stories always outnumber true stories in all equilibria social connectedness is either sufficiently high or sufficiently low. Moreover, true stories completely disappear in the limit as social connectedness d goes to infinity, in the sense that the maximal equilibrium news veracity converges to zero.

Misinformation agents intent on undermining a news market can do so by injecting a sufficiently large volume of additional false stories into the market. However, small volumes of misinformation can have the ironic effect of promoting more true-news production by bona-fide news producers. This is because consumers respond to the presence of misinformation by sharing news more cautiously, which in turn reduces the false-story revenue more than it reduces true-story revenue.

In an actions-supported news market, increasing social connectedness is more unambiguously beneficial. Indeed, in the limit as consumers follow infinitely-many others, consumers enjoy the "ideal outcome" in equilibrium whereby (i) true-news production is as large as it can possibly be and (ii) consumers are able to perfectly discern which stories are true, enabling them to act on all true stories and avoid acting on all false stories, what we refer to as a "wisdom of the crowd." In such large actions-supported market, misinformation has no appreciable effect on outcomes, since no one acts on misinformation due to the equilibrium wisdom of the crowd.

This paper serves as a basis for the study of 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. Such endogenous link-formation would have implications for the efficiency of media platforms. In an actions-supported news market, we showed that high-veracity news can arise in equilibrium if consumers are densely connected. 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 have shifted toward subscription-based revenue (New York Times 2015). Whoever controls a subscription channel 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[25] 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, screening 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,[26] but submit it to a reputable paper such as the *Washington Post* for editorial

[25] 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).

[26] 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.

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.[27] 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 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

[27] 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).

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.

AppendixS

Appendix A provides details on dynamic stability of equilibria. Appendix B contains proofs of results not established in the text.

Conference readers: The Appendix is incomplete. Several proofs need to be edited, and some proofs are not yet provided here.

A Dynamic stability

We consider equilibria that are stable outcomes, where our stability concept corresponds to "evolutionary stability" in asymmetric games, with producers and consumers as co-adapting populations; see Samuelson and Zhang (1992). As we illustrate below, this stability restriction imposes familiar conditions on how the best reply functions cross each other from above and below.

Definition 1 (Equilibrium stability)**.** *An equilibrium with sharing rule* $\mathbf{z}$ *is "stable" if, for all* $\hat{\mathbf{z}} \neq \mathbf{z}$ *and all* $\epsilon \approx 0$*,* $\mathbf{z}$ *is a strictly better response for consumers than* $\hat{\mathbf{z}}$ *given perturbed news veracity* $P(\mathbf{z}(1-\epsilon)+\hat{\mathbf{z}}\epsilon)$*. Similarly,* $\mathbf{z}$ *is "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* $P(\mathbf{z}(1-\epsilon)+\hat{\mathbf{z}}\epsilon)$*.*

An equilibrium is stable if $P(\mathbf{Z}(p_0+\epsilon)) < p_0+\epsilon$ and $P(\mathbf{Z}(p_0-\epsilon)) > p_0-\epsilon$ for all $\epsilon \approx 0$. In particular, an equilibrium with news veracity p_0 and sharing rule $\mathbf{z}$ is stable if $\frac{1}{\mathbf{Z}'(p_0)} > P'(\mathbf{z})$, so that consumers' inverse best-reply curve $\mathbf{Z}^{-1}(\mathbf{z})$ crosses producers' best-reply curve $P(\mathbf{z})$ from below. Similarly, an equilibrium is unstable if $\frac{1}{\mathbf{Z}'(p_0)} > P'(\mathbf{z})$, so that $\mathbf{Z}^{-1}(\mathbf{z})$ crosses $P(\mathbf{z})$ from above.

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 2 (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 3 (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* $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}})$*.*[28] *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 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 1(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 (**??**), $p_0(\mathbf{z})$ 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

[28] 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.

$s_i = F$; 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 1(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 1(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 1(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). □

B Proofs

B.1 Proof of Lemma 1

Part (i). Define $x(z_T) = \frac{\Delta V(0, z_T, 0)}{1-b}$. Since $p_0(\mathbf{z}) = H(M\Delta V(\mathbf{z}))$, 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 $\overline{z}_T \in (0, 1]$. For sharing rules of the form $\mathbf{z} = (0, z_T, 0)$, each consumer's likelihood of sharing a true story $\sigma_T(\mathbf{z}) = b\iota\hat{\rho})z_T$ and likelihood of sharing a false story $\sigma_F(\mathbf{z}) = b\iota(1-\hat{\rho}))z_T$. By equations (1-2),

$$x(z_T) = (1 - b\iota(1-\hat{\rho})z_T)^d - (1 - b\iota\hat{\rho}z_T)^d$$
$$x'(z_T) = db\iota\left(\hat{\rho}(1 - b\iota\hat{\rho}z_T)^{d-1} - (1-\hat{\rho})(1 - b\iota(1-\hat{\rho})z_T)^{d-1}\right)$$

Suppose first that $d = 1$. Since $x'(z_T) = b\iota(2\rho - 1) > 0$, $x(z_T)$ 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 $\frac{\hat{\rho}}{1-\hat{\rho}} > \left(\frac{1-b\iota(1-\hat{\rho})z_T}{1-b\iota\hat{\rho}z_T}\right)^{d-1}$ which, after re-arranging, can be written as $z_T < \hat{z}_T \equiv \frac{\left(\frac{\hat{\rho}}{1-\hat{\rho}}\right)^{\frac{1}{d-1}} - 1}{b\iota\left(\hat{\rho}\left(\frac{\hat{\rho}}{1-\hat{\rho}}\right)^{\frac{1}{d-1}} - (1-\hat{\rho})\right)}$. 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 (ii). Define $x(z_F) = \frac{\Delta V(1, 1, z_F)}{1-b}$. We need to show that $x(z_F)$ is strictly decreasing in

z_F. For sharing rules of the form $\mathbf{z} = (1, 1, z_F)$, we have $\sigma_T(\mathbf{z}) = b(1 - \iota(1 - \hat{\rho})(1 - z_F)$ and $\sigma_F(\mathbf{z}) = b(1 - \iota\hat{\rho}(1 - z_F))$. So,

$$x(z_F) = (1 - b(1 - \iota\hat{\rho}(1 - z_F))^d - (1 - b(1 - \iota(1 - \hat{\rho})(1 - z_F))^d$$
$$x'(z_F) = db\iota\left((1 - \hat{\rho})(1 - \sigma_T(1, 1, z_F))^{d-1} - \hat{\rho}(1 - \sigma_F(1, 1, z_F))^{d-1}\right)$$

$x'(z_F) < 0$ follows from $\sigma_T(1, 1, z_F) > \sigma_F(1, 1, z_F)$ and $\hat{\rho} > 1/2$.

Part (iii). Define $x(z_\emptyset) = \frac{\Delta V(z_\emptyset, 1, 0)}{1-b}$. For sharing rules of the form $\mathbf{z} = (z_\emptyset, 1, 0)$, we have $\sigma_T(\mathbf{z}) = b((1 - \iota)z_\emptyset + \iota\hat{\rho})$ and $\sigma_F(\mathbf{z}) = b((1 - \iota)z_\emptyset + \iota(1 - \hat{\rho}))$. So,

$$x(z_\emptyset) = (1 - b((1 - \iota)z_\emptyset + \iota(1 - \hat{\rho})))^d - (1 - b((1 - \iota)z_\emptyset + \iota\hat{\rho}))^d$$
$$x'(z_\emptyset) = db(1 - \iota)\left((1 - \sigma_T(z_\emptyset, 1, 0))^{d-1} - (1 - \sigma_F(z_\emptyset, 1, 0))^{d-1}\right)$$

If $\iota = 1$, then $x(z_\emptyset)$ does not depend on $z_\emptyset$. The fact that $x'(z_\emptyset) < 0$ when $\iota < 1$ follows from $\sigma_T(1, 1, z_F) > \sigma_F(1, 1, z_F)$.

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

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

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

B.2 Proof of Theorem 2(iv)

Proof. EDIT PROOF. Differentiating (3) yields

$$\frac{\mathrm{d}\Delta V(\mathbf{z};d)}{\mathrm{d}d} = (1-b)\left(\log(1-\sigma_F(\mathbf{z})(1-\sigma_F(\mathbf{z}))^d - \log(1-\sigma_T(\mathbf{z})(1-\sigma_T(\mathbf{z}))^d\right)$$
$$\gtrless 0 \text{ iff } \frac{\log(1-\sigma_F(\mathbf{z}))}{\log(1-\sigma_T(\mathbf{z}))} \gtrless \left(\frac{1-\sigma_T(\mathbf{z})}{1-\sigma_F(\mathbf{z})}\right)^d \tag{4}$$

Since $0 < \sigma_F(\mathbf{z}) < \sigma_T(\mathbf{z}) < 1$, the left-hand side of (4) is between zero and one, while the right-hand side is exponentially decreasing in d. Thus, there exists a threshold $\overline{d}(\mathbf{z}) \geq 1$ such that (i) $\Delta V(\mathbf{z}; d) > \Delta V(\mathbf{z}; d-1)$ for all $d \leq \overline{d}(\mathbf{z})$ and (ii) $\Delta V(\mathbf{z}; d+1) < \Delta V(\mathbf{z}; d)$ for all $d \geq \overline{d}(\mathbf{z})$. We conclude that $P(\mathbf{z}; d)$ is single-peaked in d and maximized at $\overline{d}(\mathbf{z})$. Part (iv) then follows since for each sharing rule $\mathbf{z}$, $\Delta V(\mathbf{z}; d)$ is strictly decreasing in d for all $d \geq \lceil \overline{d}(\mathbf{z}) \rceil$. As $d \to \infty$, $\Delta V(\mathbf{z}; d) \to 0$ for all $\mathbf{z} \in \mathcal{Z}$. Hence, $M\Delta V(\mathbf{z}; d)$ converges to 0 as $d \to \infty$, and, hence, the maximal equilibrium news veracity $p_0^*(d)$ must also converge to zero. □

B.3 Proof of Lemma 2.

INSERT PROOF.

B.4 Proof of Lemma 3.

INSERT PROOF.