

Social Networks and the Market for News

Rachel Kranton and David McAdams*

This draft prepared for: Penn State Theory Seminar

March 15, 2019

Abstract

This paper builds a simple model of the market for decision-relevant information (“news”) and examines how expanding social networks impact producers’ incentive to publish true news and consumers’ ability to discern what news is true. Having a highly-connected social network reduces producers’ incentive to invest in story quality if producers are paid when consumers encounter their stories (our baseline model), but increases their incentive to invest if producers are paid to influence consumer behavior. In either case, consumers’ ability to discern the truth is limited in equilibrium, even in the limit as each consumer follows infinitely many others. Introducing fake-news providers (who can only produce false stories) has an ambiguous effect on other news producers, perhaps triggering a downward spiral in news quality or perhaps motivating other producers to invest more in ensuring that their own stories are true. Contemporary news-media markets provide our main motivation, but the model applies more generally to a variety of settings—such as product marketing, charity fundraising, and political campaigns—in which decision-relevant information is shared socially.

*Kranton: Economics Department, Duke University. McAdams: Fuqua School of Business and Economics Department, Duke University. We thank seminar participants at Duke, EUI, GSE Barcelona, the 6th European Conference on Networks, and the 17th IO Theory Conference for helpful comments.

1 Introduction

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. As we use the term, “fake news” refers to published information that a news producer knows to be false. Many governments have used false propaganda to influence public opinion, at home and abroad.¹ Tabloid newspapers have long published questionable stories about the lives of celebrities. The current media environment is different, however, in that producers can more easily enter the market using online news-distribution channels and consumers can spread news more easily through social media. This paper builds a simple model of this contemporary market for news and examines how expanding social networks impact the likelihood that published news is true, what we refer to as “news veracity,” and consumers’ ability to discern what is true.

The model specifies the supply and demand for news and applies to a variety of settings in which decision-relevant information is shared socially. The model indicates dual roles of social networks. First, networks spread news. As more people become linked, more people can see a news story even if they did not encounter the story directly from the producer, what we refer to as the “broadcast”. Second, networks filter the news. Consumers receive informative private signals about each story they encounter and only share stories over the social network that they believe are sufficiently likely to be true. Consumers’ sharing behavior therefore can be informative about the likelihood of news truth. As more people become linked, more people can benefit from this evaluative process.

For any fixed news veracity, expanding the social network increases consumer welfare as consumers benefit from increased spread and filtering of the news. As the social network becomes more connected, however, news veracity is not fixed because producers’ incentive to invest in true stories depends on the social network. And as news veracity changes, consumers’ incentive to share stories also changes, which in turn impacts

¹In 1940, Britain deployed three thousand operatives to the United States to spread (sometimes false) propaganda under the guise of bona fide news reports (Boyd (2006), Cull (1995)), as a way to drum up popular support for entering the war effort against Nazi Germany.

consumers’ ability to learn from others’ sharing behavior. Our key innovation is to endogenize the quality of information that diffuses through a social network, and consumers’ ability to discern what is true, as a function of network structure.

We examine a non-cooperative model of news production in which each news producer decides whether to invest in news quality, with the interpretation that “high-quality” stories are always true while “low-quality” stories are false with some probability. We characterize all symmetric Nash equilibria of the resulting game, first in the baseline case when producers are paid per consumer who views their story (Sections 3-4), then in an extension allowing for producers also to be paid per consumer who chooses to “adopt” based on their story (Section 5). Then, in a second extension, we introduce “fake-news providers” who by assumption are only capable of producing false stories, and endogenize the prevalence of fake news and its equilibrium impact on consumers’ trust in published news. Introducing fake-news providers never increases overall news veracity, as one would expect, but has an ambiguous effect on other news producers’ investment decisions, in some cases triggering a downward spiral in news quality and in other cases motivating bona fide producers to invest more in ensuring that their own stories are true.

We consider a spectrum of news-producer revenue models, reflecting the variety of motivations among news providers operating in the contemporary media market. First, traditional brick and mortar newspapers, such as *The New York Times*, earn revenues from advertising that accompanies their articles both on paper and online. While consumers might base decisions on the news they see, these outlets do not typically earn revenue from those decisions. Second, fictitious-news websites, such as *denverguardian.com*, which famously published a story about a made-up FBI agent’s death being linked to Hillary Clinton,² earn revenues from advertising that accompanies the articles published online. Third, news outlets such as Fox News, MSNBC, and Breitbart support particular political positions and earn revenues from advertising but can also be supported by owners who care about advancing their own political views (see, e.g., Kroll (2017)³ on

²“This is a real news story about fake news stories” by Callum Borchers, *Washington Post*, November 7, 2016.

³“Ready for Trump TV? Inside Sinclair Broadcasting’s Plot to Take Over Your Local News” by Andy Kroll, *Mother Jones*, Nov/Dec 2017.

Sinclair Media’s owner David Smith). Finally, government-sponsored media, such as the British propagandists of 1940 or the Russian troll factories of today, publish news with the virtually sole aim of inducing consumers to take some desired action.⁴

Our results show how consumers’ social network shapes news veracity, depending on the revenue motives of the producer. Consider a network where each consumer follows many other consumers, what we refer to as a “crowd”. When producers are mostly paid based on views, they have little incentive to invest in story quality since even false stories will spread widely enough to be seen by most consumers. Equilibrium news veracity in this case can be just high enough for consumers to sometimes be willing to share, but not higher. By contrast, when producers are mostly paid based on adoption, they have a strong incentive to invest in story quality because consumers can make inferences based on their neighbors’ sharing behavior, to distinguish true stories from false ones and avoid adopting based on false stories. Even in this case, however, consumers’ ability to discern the truth is limited and some false stories are still produced in equilibrium. Indeed, even in the limit as consumers follow infinitely many others, consumers are unable in equilibrium to discern for sure which stories are true, i.e., there is no “wisdom of the crowd”.

The paper contributes to two distinct literatures: the literature on media markets and the literature on social learning, information transmission, and networks. Much of the previous work on media markets studies media bias. In Gentzkow & Shapiro (2006), media bias arises because consumers hold initial beliefs that make them less likely to find contrary news to be credible. News producers who earn revenues based on their reputations for accuracy then have an incentive to slant their news reports. In Besley & Pratt (2006) and Gentzkow, Glaeser & Goldin (2006), news producers face a tradeoff between earning advertising revenue or producing news biased in a sponsor’s favor. In that context, advertising revenue is shown to reduce political bias. In Ellman and Germano (2009), however, newspapers bias their news towards their advertisers.

⁴A well-known example of this category is the Heart of Texas which was revealed to be a fictitious advocacy group created in St. Petersburg that promoted Texas succession from the United States and other provocative positions. When the fake Facebook group called for a rally in against “the adoption of the Sharia law” in Texas, real people showed up to rally and counter rally.

The present paper considers news outlets that earn an exogenous (reduced form) mix of revenue from advertisements and from content sponsors. We focus on quality per se, with strategic consumers who desire true news, who receive private signals that are informative about news truth, and can share the news over a social network. A key insight that emerges from our analysis is that switching from an advertising-supported revenue model to a sponsor-supported model can increase the equilibrium truthfulness of the news. Several other recent papers also study features of contemporary media markets, such as proliferation of producers and competition for consumers’ limited attention (Chen and Suen (2018)), media bias when consumers have heterogeneous preferences and pass on news to like-minded individuals (Redilicki (2018)), and competition to break a story that leads to lower quality news (Andreottola and DeMoragas (2018)).

The demand side of our media market involves a simple model of consumer information transmission and social learning that is both similar to and different from other models in the literature. Consumers in our model are Bayesian, receive private signals, and update their beliefs about each news item based on others’ observed choices; so, consumers may in some cases ignore their own private information when deciding whether to “adopt”. However, unlike in the cascades literature (e.g., Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992)), there is only one round of social learning and consumers can observe multiple neighbors’ independent sharing decisions. As in Bloch, Demange and Kranton (2018) and Chatterjee and Dutta (2016), but unlike much of the literature on information diffusion in networks (e.g., Acemoglu, Ozdaglar, and ParandehGheibi (2010) and Banerjee, Chandrasekhar, Duflo, and Jackson (2013)), consumers in our model choose whether or not to pass on information to their neighbors. These decisions account for the network role in filtering the news. Finally, as in Galeotti and Goyal (2009) and Chatterjee and Dutta (2016), there are third parties outside the network, the producing firms, whose optimal decisions are affected by the network structure. These papers study how firms will target advertising, while the present paper endogenizes product quality. To the best of our knowledge, this paper is the first study of product diffusion on a network that considers how firms’ production incentives—and hence the product itself—are shaped by the network.

2 Model: The Market for News

The “market for news” consists of a single⁵ news producer and M news consumers and unfolds in three stages: producer investment at time $t = 0$, consumer sharing over a social network at time $t = 1$, and consumer adoption at time $t = 2$.

Consumers encounter decision-relevant information, what we refer to as “news”. Consumers’ ex ante belief that news items are true, which we refer to as “news veracity,” informs their choices whether to share news over a social network and whether to take an action based on each news item, e.g., whether to buy a product promoted in a marketing message or vote for a candidate featured in a political message. We refer to this action as “adoption”.

The producer decides whether to make a costly investment that increases the likelihood of news truth. In our baseline analysis, we assume that the producer is only paid for views, getting one unit of revenue per consumer who views a published story.⁶ In Section 5, we allow the producer to get revenue per consumer who adopts or a combination of revenue from both views and adoptions. The producer lacks commitment power.

At time $t = 0$, the producer has the opportunity to publish a news item, or “story,” and decides whether to invest in story quality. Low-quality stories are costless to produce and false with probability $f \in (0, 1]$, while high-quality stories cost $c_R > 0$ to produce and are always true. The cost of reporting c_R is an atomless random variable and i.i.d. across stories. Let $H(\cdot)$ denote the c.d.f. of c_R/M , the per-consumer cost of reporting, i.e., $Pr(c_R < x) = H(x/M)$. We assume that $H(1) = 1$, meaning that c_R is never greater than M , the maximal possible revenue that a story can earn.

At time $t = 1$, the story is broadcast to consumers, who then decide whether to share it over a social network. Each consumer sees the broadcast with independent probability

⁵Our single-producer analysis applies equally to a news market with finitely-many or a continuum of producers, as long as the stories that producers publish are unrelated to one another. See Section 5.3 for an extension in which there are multiple dissimilar producers.

⁶Consumers may encounter the same news item multiple times but, by assumption, the producer is only paid once per consumer who sees the news. For instance, social-media users may see many friends sharing the same *New York Times* story but only read it once—and hence only encounter advertisements once on the *New York Times* website.

$b \in (0, 1)$, what we refer to as “broadcast reach”. By assumption, consumers cannot directly observe story quality. However, when first encountering a story, each consumer receives a conditionally i.i.d. private signal $s_i \in \{T, F\}$, where $\Pr(s_i = T|\text{true}) = \rho_T$ and $\Pr(s_i = F|\text{false}) = \rho_F$. For simplicity, we assume that $\rho_T = \rho_F = \rho > \frac{1}{2}$.

Those who have seen the broadcast decide whether to share the news. Sharing occurs over a directed social network G , where $g_{ji} \in \{0, 1\}$ and $g_{ji} = 1$ indicates a communication link from consumer j to i , i.e., i “follows” j . N_i denotes the set of other consumers from whom i can receive information, with $d_i = \#(N_i)$. For simplicity, we focus on networks in which $d_i = d$ for all i . We refer to $d \geq 0$ as “social connectedness” and refer to networks with higher d as being “more connected”. By assumption, consumers would like to share true stories but not false stories. In particular, consumer i earns positive “sharing payoff” $\pi_{iT}^S > 0$ when sharing a true story, negative sharing payoff $-\pi_{iF}^S < 0$ when sharing a false story, or zero sharing payoff when not sharing. For simplicity, we assume that consumers have symmetric and equal sharing payoffs, i.e., $\pi_{iT}^S = \pi_{iF}^S = \pi^S > 0$ for all i . Consumers therefore prefer to share whenever they believe that a story’s likelihood of being true exceeds “sharing threshold” $p^S = \frac{1}{2}$.

At time $t = 2$, each consumer i who has seen the news decides whether to take an action referred to as “adopting”. Consumer i earns positive “adoption payoff” $\pi_{iT}^A > 0$ when adopting based on a true story, negative adoption payoff $-\pi_{iF}^A < 0$ when adopting based on a false story, and zero adoption payoff when not adopting. (After period 2, the story is revealed to be true or false. Consumers’ sharing and adoption payoffs are then realized.) For simplicity, we assume that consumers have symmetric adoption payoffs, i.e., $\pi_{iT}^A = \pi_T^A > 0$ and $\pi_{iF}^A = \pi_F^A > 0$ for all i . Consumers therefore prefer to adopt whenever they believe that a story’s likelihood of being true exceeds “adoption threshold” $p^A = \frac{\pi_F^A}{\pi_T^A + \pi_F^A}$. For concreteness, we assume that $p^A > \frac{1}{2}$, meaning that consumers require a greater degree of confidence in a story’s truth in order to adopt than they do in order to share.

Discussion. (a) We focus on a context in which the thing being produced is a *factual claim* that may or may not be true, but 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).

(b) The analysis can be easily adapted to a setting in which consumers are already aware of an opportunity but not adopting is the optimal action absent any new decision-relevant information. For example, suppose that “consumers” are leading scholars in a field who may be asked to write letters for a tenure case. Suppose further that, absent new information, the letter-writers will all advise against tenure. The tenure candidate (“producer”) has just finished a new working paper. Scholars in the field who read the paper form an impression of its quality and then decide whether to mention it to others at conferences. Those asked to write letters may not have seen or heard about the working paper but, if not, will account for that fact when evaluating its quality—exactly as consumers in our model do when seeing the initial broadcast but then not seeing anyone sharing the news.

(c) The producer in our model can be interpreted as representing a continuum of undifferentiated producers, each of whom makes their own independent choices when the opportunity to publish a story presents itself. Under this interpretation, our assumption that the producer lacks commitment power is without loss of generality, since no individual producer can significantly change consumers’ beliefs about the likelihood that a published story is true.

(d) The analysis can be extended in a straightforward way to allow for non-linear producer revenue. For example, suppose that a political candidate who is trailing in the polls has an option to spread a nasty rumor about her opponent just before the election. Our analysis can be applied to this sort of context, with “consumers” being voters who prefer the other candidate, “adoption” being to stay home from the polls (or switch their vote), and “revenue” being non-linear: zero if the candidate loses or one if she wins. Whether the candidate chooses to spread a rumor, and how much effort she invests in identifying the most damning possible rumor to spread, will therefore depend on how likely different-quality rumors will be to sway a sufficient number of voters and on the

costs associated with conducting opposition research.

3 Equilibrium Outcomes in the Market for News

This section analyzes the market for news in the baseline case when producer revenue depends only on consumer views, namely, that the producer earns one unit of revenue per consumer who views a story. In an extension in Section 5, we allow producer revenue to depend (also or instead) on consumers' adoption decisions.

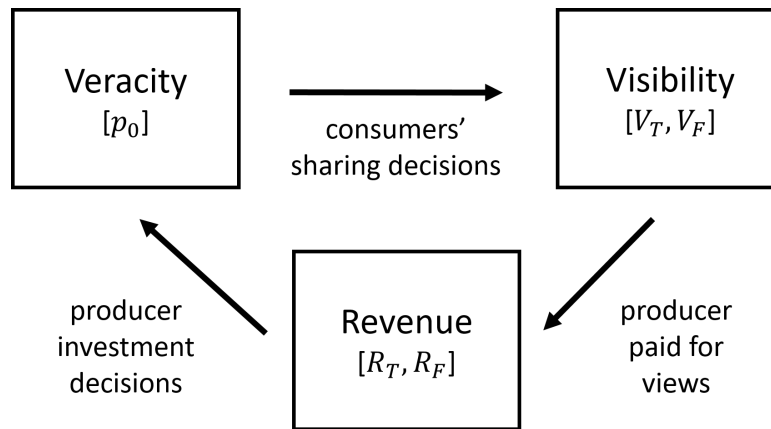


Figure 1: Schematic of the market for news when producers are paid for views.

Figure 1 illustrates the interaction between consumer and producer decisions in the market for news. Consumers' sharing decisions impact the producer's incentive to invest in story quality, by changing the visibility of true and false stories. Producer investment in turn impacts each consumer's incentive to share, by changing the likelihood that published stories are true. Finally, each consumer's incentive to adopt is impacted by both producer investment and others' sharing decisions. However, since the reporter is only paid for views, consumer adoption has no impact on producer investment.

3.1 Consumer best response

First, we consider the consumer's incentive to share news stories, depending on their own private signals and the ex ante likelihood of news truth, and the resulting extent to which consumers view true and false stories.

Time-1 belief updating. Let p_0 be the ex ante likelihood that published stories are true, what we refer to as “news veracity”. Suppose that consumer i has seen a story's time-0 broadcast and gets private signal $s_i \in \{T, F\}$. Consumer i updates her belief about the likelihood of news truth to $p_1(s_i; p_0)$. By Bayes' Rule, $p_1(T; p_0) = \frac{p_0 \rho}{p_0 \rho + (1-p_0)(1-\rho)}$ and $p_1(F; p_0) = \frac{p_0(1-\rho)}{p_0(1-\rho) + (1-p_0)\rho}$ or, equivalently,

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




Figure 2: News-veracity regions with different patterns of consumer sharing.

Consumer sharing. Since consumer i gets π^S (or $-\pi^S$) when sharing a true (or false) story, her expected sharing payoff is $p_1(s_i; p_0)\pi^S - (1 - p_1(s_i; p_0))\pi^S$ when sharing or zero when not sharing. Thus, consumer i strictly prefers to share (or not share) if $p_1(s_i; p_0) > \frac{1}{2}$ (or $p_1(s_i; p_0) < \frac{1}{2}$). Let z_{iT} and z_{iF} denote consumer i 's likelihood of sharing after seeing the broadcast and getting, respectively, private signal $s_i = T$ or $s_i = F$. We refer to the vector $\mathbf{z}_i = (z_{iT}, z_{iF})$ as consumer i 's “sharing rule”. To keep the analysis as simple as possible, we restrict attention to equilibria in which all consumers use the same sharing rule, what we call “symmetric sharing”. We say that “sharing is informative” if consumers use a sharing rule $\mathbf{z} = (z_T, z_F)$ such that $z_T > z_F$ since, in that case, consumers are more

likely to share true stories. On the other hand, “sharing is uninformative” if consumers use a sharing rule $\mathbf{z} = (z_T, z_F)$ such that $z_T = z_F$.

Consumers’ optimal sharing rule depends on news veracity.

- *Always-share region* $(\rho, 1]$: If news veracity $p_0 > \rho$, sharing is uninformative as consumers always find it optimal to share.
- *Never-share region* $[0, 1 - \rho)$: If news veracity $p_0 < 1 - \rho$, sharing is uninformative as consumers always find it optimal not to share.
- *Filtering region* $(1 - \rho, \rho)$: If news veracity $p_0 \in (1 - \rho, \rho)$, sharing is informative as consumers find it optimal to “filter the news,” sharing after a good signal but not after a bad signal. Let $\tilde{\mathbf{z}} \equiv (1, 0)$ denote the “filtering sharing rule”.

At the boundaries of the filtering region, consumers are sometimes indifferent whether to share. If news veracity $p_0 = \rho$, what we refer to as the “always-share threshold,” consumers are indifferent whether to share after getting a bad signal and hence find it optimal to use any sharing rule of the form $\mathbf{z} = (1, z_F)$. Similarly, if news veracity $p_0 = 1 - \rho$, the “never-share threshold,” consumers are indifferent after getting a good signal and have optimal sharing rules of the form $\mathbf{z} = (z_T, 0)$. Let $Z(p_0)$ denote consumers’ best-response correspondence, the set of optimal sharing rules for any given news veracity p_0 .

Proposition 1 summarizes our findings on best-response consumer sharing.

Proposition 1. *(i) In the never-share region, $Z(p_0) = (0, 0)$. (ii) In the always-share region, $Z(p_0) = (1, 1)$. (iii) In the filtering region, $Z(p_0) = \tilde{\mathbf{z}} \equiv (1, 0)$. (iv) At the never-share threshold, $Z(1 - \rho) = \{(z_T, 0) : z_T \in [0, 1]\}$. (v) At the always-share threshold, $Z(\rho) = \{(1, z_F) : z_F \in [0, 1]\}$.*

News visibility. Each consumer’s likelihood of viewing a story, what we refer to as the story’s “visibility” to that consumer, depends on broadcast reach (b), the number of others she follows in the social network (d), and the sharing rule \mathbf{z} used by her neighbors in the social network. (We use the word “neighbors” to refer to consumers who are

linked, with the context indicating the link’s direction. Here, a “neighbor” is anyone that consumer i follows.) Each neighbor j shares iff she sees the broadcast (probability b) and gets a private signal that induces sharing (probability $\rho z_T + (1 - \rho)z_F$ if the news is true or $(1 - \rho)z_T + \rho z_F$ if it is false). Each consumer’s ex ante likelihood of viewing a true story, by either seeing the broadcast directly and/or having at least one sharing neighbor, is therefore

$$V_T(\mathbf{z}) = 1 - (1 - b)(1 - b(\rho z_T + (1 - \rho)z_F))^d \quad (2)$$

Similarly, consumer i ’s likelihood of viewing a false story given sharing rule \mathbf{z} is

$$V_F(\mathbf{z}) = 1 - (1 - b)(1 - b((1 - \rho)z_T + \rho z_F))^d \quad (3)$$

Numerical example: news visibility. Suppose that broadcast reach $b = \frac{1}{12}$ and private-signal precision $\rho = \frac{2}{3}$. Given any news veracity in the never-share region ($p_0 < \frac{1}{3}$), consumers never share; so, consumers only become aware of a story if they see its initial broadcast and visibility $V_T(0, 0) = V_F(0, 0) = b = \frac{1}{12}$ for true and false stories. On the other hand, given any news veracity in the always-share region ($p_0 > \frac{2}{3}$), consumers who see the broadcast always share; so, each consumer becomes aware of a story whenever she or any of her d neighbors sees the broadcast and visibility is $V_T(1, 1) = V_F(1, 1) = \bar{V} = 1 - (1 - b)^{d+1} = 1 - \left(\frac{11}{12}\right)^{d+1}$. Finally, given any news veracity in the filtering region ($p_0 \in \left(\frac{1}{3}, \frac{2}{3}\right)$), consumers who see the broadcast share whenever they get a positive signal; so, each consumer shares true and false stories with ex ante probability $b\rho_T = \frac{1}{18}$ and $b(1 - \rho_F) = \frac{1}{36}$, respectively, yielding the following visibilities for true and false stories:

$$V_T(\tilde{\mathbf{z}}) = 1 - \frac{11}{12} \left(\frac{17}{18}\right)^d \quad \text{and} \quad V_F(\tilde{\mathbf{z}}) = 1 - \frac{11}{12} \left(\frac{35}{36}\right)^d$$

3.2 Producer best response

Next, we turn to the producer’s incentive to invest in story quality. Let $R_T(\mathbf{z})$ and $R_F(\mathbf{z})$ be the expected revenue of true and false stories, respectively, when consumers use sharing rule $\mathbf{z} = (z_T, z_F)$. Since the producer is paid for views and there are M 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 $\Delta V(\mathbf{z}) = V_T(\mathbf{z}) - V_F(\mathbf{z})$ is the incremental visibility of true stories over false ones for each consumer. Since low-quality stories are false with probability $f > 0$ and high-quality stories are always true, investing in story quality increases expected revenue by $Mf\Delta V(\mathbf{z})$ at cost c_R . The producer therefore maximizes expected profit by investing iff her realized reporting cost $c_R < Mf\Delta V(\mathbf{z})$, inducing news veracity

$$p_0(\mathbf{z}) = (1 - f) + fH(f\Delta V(\mathbf{z})) \quad (4)$$

where $H(f\Delta V(\mathbf{z}))$ is the ex ante probability that the producer invests and $H(1) = 1$ because of our assumption that reporting cost c_R is bounded above by M . Note that the reporter’s profit-maximizing strategy is uniquely determined at all reporting cost levels except $c_R = Mf\Delta V(\mathbf{z})$, which occurs with probability zero since c_R is drawn from an atomless distribution. Thus, the producer has an essentially-unique best response to any sharing rule. We refer to the resulting news veracity $p_0(\mathbf{z})$ as the “best-response news veracity”.

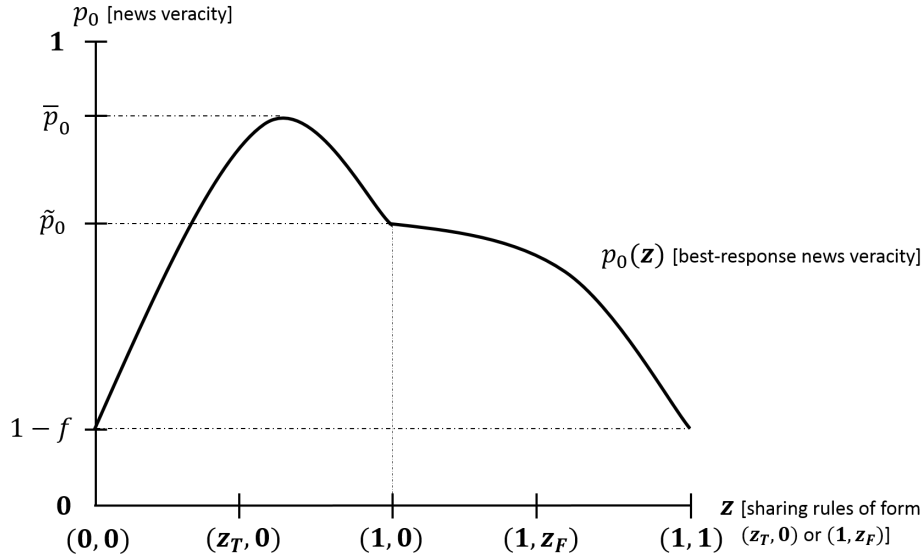


Figure 3: Illustration of the best-response news veracity function $p_0(\mathbf{z})$, the maximal best-response news veracity $\bar{p}_0 \equiv \max_{\mathbf{z}} p_0(\mathbf{z})$, and the filtering news veracity $\tilde{p}_0 \equiv p_0(1, 0)$.

Figure 3 illustrates key qualitative features of the best-response news veracity.⁷ First,

⁷To avoid confusion, note that the x-axis of this figure consists of all sharing rules that could poten-

the producer never invests if consumers always share or never share, i.e., $p_0(1, 1) = p_0(0, 0) = 1 - f$. Second, the producer sometimes invests so long as consumer sharing is informative, i.e., $p_0(\mathbf{z}) > 1 - f$ whenever $z_T > z_F$. Third, $p_0(z_T, 0)$ is continuous and single-peaked in z_T . Finally, $p_0(1, z_F)$ is continuous and strictly decreasing in z_F . Figure 3 also illustrates two news-veracity levels that will play an important role in the analysis: the “maximal best-response news veracity” $\bar{p}_0 \equiv \max_{\mathbf{z}} p_0(\mathbf{z})$ and the “filtering news veracity” $\tilde{p}_0 \equiv p_0(1, 0)$.

Lemma 1 gathers together useful facts about best-response news veracity.

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

Equilibrium concept. In the main text, we focus on Nash equilibria (NE) of the game played by the producer (deciding whether to invest) and consumers (deciding whether to share) in which all consumers use the same sharing rule; such NE are referred to as “symmetric Nash equilibria” (SNE). Moreover, in most of the analysis, we focus on the maximal news veracity that can be supported in any SNE, using notation p_0^* to denote the “maximal equilibrium news veracity”. In an effort to avoid acronyms, we use the term “equilibrium” in the main text when referring to SNE. See the Appendix for a full characterization of all NE.

3.3 Maximal equilibrium investment

This section characterizes the maximal news veracity p_0^* that can be supported in any equilibrium, which turns out to depend only on (i) the likelihood $1 - f$ that low-quality stories are true, (ii) private-signal precision ρ , (iii) the maximum best-response news veracity \bar{p}_0 , and (iv) the filtering news veracity \tilde{p}_0 .

tially be a best response for consumers, i.e., those of the form $(z_T, 0)$ or $(1, z_F)$.

Theorem 1. *If either $1 - f \geq \rho$ or $\bar{p}_0 < 1 - \rho$, then there is a unique equilibrium with zero investment and $p_0^* = 1 - f$. Otherwise, an equilibrium exists with positive investment and $p_0^* = \max\{1 - \rho, \min\{\tilde{p}_0, \rho\}\}$.*

We prove Theorem 1 through a series of five main steps.

Step 1: *In any equilibrium with investment, news veracity cannot be in the always-share region.* Suppose for a moment that an equilibrium exists in which the producer sometimes invests, inducing news veracity $\hat{p}_0 > 1 - f$. We begin by observing that \hat{p}_0 cannot exceed the always-share threshold ρ . To see why, suppose for the sake of contradiction that $\hat{p}_0 > \rho$. Consumers' unique best response would then be to always share, given which the producer's best response would be to never invest, contradicting the assumption that the producer sometimes invests.

An immediate implication is that, if low-quality stories are sufficiently likely to be true that $1 - f \geq \rho$, the producer never invests in any equilibrium. On the other hand, a “zero-investment equilibrium” always exists in this case in which (i) the producer never invests, anticipating that consumers will always share, and (ii) consumers always share, anticipating that all stories are low quality. Thus, whenever $1 - f \geq \rho$, the maximal equilibrium news veracity $p_0^* = 1 - f$.

Step 2: *In any equilibrium with investment, news veracity cannot be in the never-share region.* Suppose for the sake of contradiction that an equilibrium exists in which the producer sometimes invests, but invests sufficiently rarely that news veracity $\hat{p}_0 < 1 - \rho$. Consumers' unique best response would then be to never share, given which the producer's best response would be to never invest, contradicting the assumption that the producer sometimes invests.

An immediate implication is that, if the maximal best-response news veracity $\bar{p}_0 < 1 - \rho$, then all equilibria must have zero investment. To see why, suppose for the sake of contradiction that $\bar{p}_0 < 1 - \rho$ and there exists an equilibrium with investment. Let $\hat{\mathbf{z}}$ denote consumers' sharing rule in this equilibrium. The resulting news veracity $\hat{p}_0 = p_0(\hat{\mathbf{z}}) \leq \bar{p}_0$ by definition of \bar{p}_0 ; so, $\hat{p}_0 < 1 - \rho$, contradicting Step 2. On the other hand, a zero-investment equilibrium always exists in this case in which (i) the producer never

invests, anticipating that consumers will never share, and (ii) consumers never share, anticipating that all stories are low quality.⁸ Thus, whenever $\bar{p}_0 < 1 - \rho$, the maximal equilibrium news veracity $p_0^* = 1 - f$.

So far, we have seen that all equilibria must exhibit zero investment whenever *either* low-quality stories are sufficiently likely to be true that $1 - f \geq \rho$ (Step 1) *or* the maximal best-response news veracity $\bar{p}_0 < 1 - \rho$, which can only happen if low-quality stories are sufficiently unlikely to be true that $1 - f < 1 - \rho$ (Step 2). This completes the proof of the first part of Theorem 1. It remains to show that, when neither of these conditions holds, an equilibrium with positive investment exists and the maximal news veracity that can be supported in equilibrium equals $p_0^* = \max\{1 - \rho, \min\{\tilde{p}_0, \rho\}\}$. Fortunately, this complicated-looking expression can be written more simply in terms of three separate cases, depending on the filtering news veracity \tilde{p}_0 : (i) if $\tilde{p}_0 \in [1 - \rho, \rho]$, then $p_0^* = \tilde{p}_0$ (Step 3); (ii) if $\tilde{p}_0 > \rho$, then $p_0^* = \rho$ (Step 4); and (iii) if $\tilde{p}_0 < 1 - \rho$, then $p_0^* = 1 - \rho$ (Step 5).

Step 3: If the filtering news veracity $\tilde{p}_0 \in [1 - \rho, \rho]$, then the maximal equilibrium news veracity $p_0^ = \tilde{p}_0$.* If consumers filter the news, using sharing rule $\tilde{\mathbf{z}} = (1, 0)$, the producer's best response induces news veracity $p_0(1, 0) = \tilde{p}_0$. If $\tilde{p}_0 \in [1 - \rho, \rho]$,⁹ then consumers find it to be a best response to filter the news and an equilibrium exists with news veracity \tilde{p}_0 . Moreover, no equilibrium can exist with higher news veracity. Why? First, by Step 1, no equilibrium can have news veracity greater than ρ . Second, no equilibrium can have news veracity in the interval (\tilde{p}_0, ρ) since, if there were such an equilibrium, consumers would all filter the news and the producer would only invest enough to support news veracity \tilde{p}_0 . Finally and least obviously, no equilibrium can have news veracity equal to ρ when $\tilde{p}_0 < \rho$. In any such equilibrium, consumers must use a sharing rule of the form $\mathbf{z} = (1, z_F)$ for some $z_F \in [0, 1]$. The producer's best response then induces news veracity $p_0(1, z_F)$. However, by Lemma 1(ii), $p_0(1, z_F)$ is strictly decreasing in z_F , implying $p_0(1, z_F) \leq p_0(1, 0) = \tilde{p}_0 < \rho$, a contradiction. Overall, we conclude that the

⁸Since $\bar{p}_0 \geq 1 - f$, $\bar{p}_0 < 1 - \rho$ is only possible if $1 - f < 1 - \rho$. Thus, whenever $\bar{p}_0 < 1 - \rho$, consumers never share a story that they believe to be definitely low quality.

⁹Since $\tilde{p}_0 > 1 - f$ and $\tilde{p}_0 \leq \bar{p}_0$, $\tilde{p}_0 \in [1 - \rho, \rho]$ is only possible if $1 - f < \rho$ and $\bar{p}_0 \geq 1 - \rho$; so, Steps 1-2 do not apply.

maximal equilibrium news veracity $p_0^* = \tilde{p}_0$.

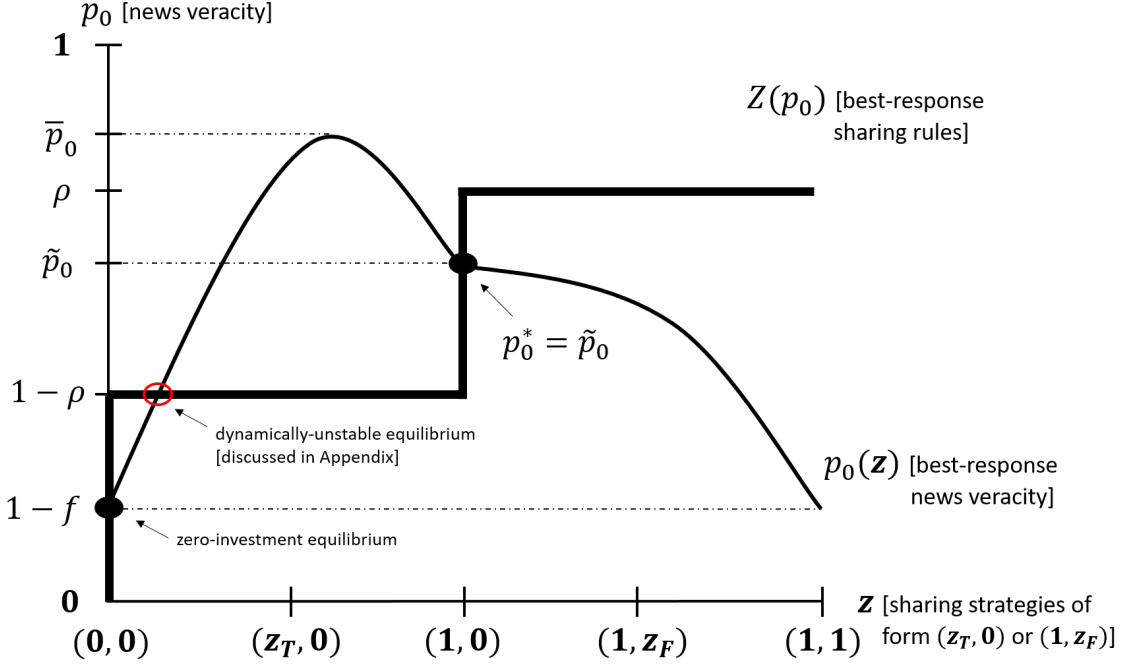


Figure 4: Illustration of the set of equilibria when filtering news veracity is in the filtering region ($\tilde{p}_0 \in [1 - \rho, \rho]$) and $1 - f < 1 - \rho$.

Figure 4 illustrates the set of equilibria in the most interesting subcase when low-quality stories are sufficiently unlikely to be true that $1 - f < 1 - \rho$: the x-axis depicts the set of potentially-optimal sharing rules; the y-axis is news veracity; the thick line depicts consumers' best-response correspondence; and the thin line depicts best-response news veracity (as in Figure 3). Each equilibrium corresponds to a crossing point of these curves, with the maximal crossing-point in this case being at the filtering rule $\tilde{\mathbf{z}} = (1, 0)$ and filtering news veracity \tilde{p}_0 . (In the Appendix, we show that two other equilibria also exist with news veracities equal to $1 - f$ and $1 - \rho$, and that the equilibrium with news veracity $1 - \rho$ is dynamically unstable with respect to simple adaptation dynamics.)

Step 4: If the filtering news veracity $\tilde{p}_0 > \rho$ and $1 - f < \rho$, then the maximal equilibrium news veracity $p_0^* = \rho$. Suppose next that $\tilde{p}_0 > \rho$. This is always true if $1 - f \geq \rho$, in which case Step 1 implies that all equilibria have zero investment. We therefore focus on

the more interesting case in which $\tilde{p}_0 > \rho$ and $1 - f < \rho$. No equilibrium exists with news veracity in the filtering region since, if it did, consumers' unique best response would be to filter the news, given which the producer's best response would induce news veracity $\tilde{p}_0 > \rho$, a contradiction. However, an equilibrium does exist with news veracity equal to ρ . Since $p_0(1, 0) = \tilde{p}_0 > \rho$ and $p_0(1, z_F)$ is continuous and strictly decreasing in z_F (Lemma 1(ii,v)), there exists a unique $z_F^* \in (0, 1)$ such that $p_0(1, z_F^*) = \rho$. An equilibrium therefore exists in which consumers use sharing rule $\mathbf{z}^* = (1, z_F^*)$ and the producer invests just often enough to induce news veracity equal to ρ . Since no equilibrium can have news veracity greater than ρ by Step 1, we conclude that the maximal equilibrium news veracity $p_0^* = \rho$.




Figure 5: Illustration of the equilibrium set when $\tilde{p}_0 < 1 - \rho < \bar{p}_0$.

Step 5: If the filtering news veracity $\tilde{p}_0 < 1 - \rho$ and $\bar{p}_0 \geq 1 - \rho$, then the maximal equilibrium news veracity $p_0^* = 1 - \rho$. Suppose finally that $\tilde{p}_0 < 1 - \rho$. This is always true if $\bar{p}_0 < 1 - \rho$, in which case Step 2 implies that all equilibria have zero investment. We therefore focus on the more interesting case in which $\tilde{p}_0 < 1 - \rho$ and $\bar{p}_0 \geq 1 - \rho$. No

equilibrium exists with news veracity in the filtering region since, if it did, consumers' unique best response would be to filter the news, given which the producer's best response would induce news veracity $\tilde{p}_0 < 1 - \rho$, a contradiction. In the same way, because $p_0(1, z_F) \leq \tilde{p}_0$ for all $z_F \in [0, 1]$ (Lemma 1(ii)), no equilibrium exists with news veracity equal to ρ . However, an equilibrium does exist with news veracity equal to $1 - \rho$. Since $\bar{p}_0 = p_0(\bar{z}_T, 0) \geq 1 - \rho$ for some $\bar{z}_T \in (0, 1)$ (Lemma 1(iv)), $p_0(1, 0) = \tilde{p}_0 < 1 - \rho$, and $p_0(z_T, 0)$ is continuous in z_T (Lemma 1(v)), there exists $z_T^* \in (\bar{z}_T, 1)$ such that $p_0(z_T^*, 0) = 1 - \rho$. An equilibrium therefore exists in which consumers use sharing rule $\mathbf{z}^* = (z_T^*, 0)$ and the producer invests just often enough to induce news veracity equal to $1 - \rho$. We conclude that the maximal equilibrium news veracity $p_0^* = 1 - \rho$.

Figure 5 illustrates the set of equilibria in the case when $\tilde{p}_0 < 1 - \rho < \bar{p}_0$. (In the Appendix, we show that two other equilibria also exist with news veracities equal to $1 - f$ and $1 - \rho$, and that the other equilibrium with news veracity $1 - \rho$ is dynamically unstable.)

3.4 Equilibrium Comparative Statics

In this section, we consider how the amount of investment that can be supported in equilibrium, as captured by the maximal equilibrium news veracity, varies with model parameters, focusing especially on social connectedness (d), private-signal precision (ρ), and broadcast reach (b).¹⁰ For clarity, we add notation in each subsection to indicate how the parameter of interest impacts endogenous variables, e.g., the maximal equilibrium news veracity is denoted as $p_0^*(d)$ in Section 3.4.1, $p_0^*(b)$ in Section 3.4.2, and $p_0^*(\rho)$ in Section 3.4.3.

¹⁰Other equilibrium comparative statics include that, holding all else fixed, (i) p_0^* is non-decreasing in market size M (due to economies of scale in reporting) and (ii) p_0^* is non-increasing in the reporting-cost distribution $H(\cdot)$.

3.4.1 Social connectedness

The effect of increasing social connectedness d on equilibrium news veracity $p_0^*(d)$ depends on how well consumers are already connected.¹¹

Proposition 2. *In the case when $1 - f \in (1 - \rho, \rho)$, $p_0^*(d)$ is single-peaked in d , i.e. there exists \bar{d} such that $p_0^*(d) \leq p_0^*(d')$ for all $d < d' \leq \bar{d}$ and all $d > d' \geq \bar{d}$.*

Proof. Since $1 - f \in (1 - \rho, \rho)$, Theorem 1 implies that $p_0^*(d) = \min\{\tilde{p}_0(d), \rho\}$, where $\tilde{p}_0(d) = p_0(\tilde{\mathbf{z}}; d)$ is the filtering news veracity. Note that, for any d', d , $p_0^*(d') > p_0^*(d)$ implies $\tilde{p}_0(d') > \tilde{p}_0(d)$ and $\tilde{p}_0(d') > \tilde{p}_0(d)$ implies $p_0^*(d') \geq p_0^*(d)$. To show that $p_0^*(d)$ is single-peaked in d , it therefore suffices to show that $\tilde{p}_0(d)$ is single-peaked in d .

By definition, $\tilde{p}_0(d)$ is the news veracity that results if all consumers use the filtering sharing rule $\tilde{\mathbf{z}} = (1, 0)$ and the producer invests optimally given such sharing behavior. By equation (4), $\tilde{p}_0(d) = 1 - f + fH(f\Delta V(\tilde{\mathbf{z}}; d))$, where $H(\cdot)$ is the c.d.f. of the producer's per-consumer reporting cost and $\Delta V(\tilde{\mathbf{z}}; d)$ is the extra visibility of true stories when consumers filter the news. 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, to show that $\tilde{p}_0(d)$ is single-peaked in d , it suffices to show that $\Delta V(\tilde{\mathbf{z}}; d)$ is single-peaked in d .

By equations (2,3), $\Delta V(\tilde{\mathbf{z}}; d) = (1-b) \left((1 - b(1 - \rho))^d - (1 - b\rho)^d \right)$ implying $\frac{d\Delta V(\tilde{\mathbf{z}}; d)}{dd} = (1 - b) \left(\ln(1 - b(1 - \rho))(1 - b(1 - \rho))^d - \ln(1 - b\rho)(1 - b\rho)^d \right)$. Re-arranging terms, we conclude that $\frac{d\Delta V(\tilde{\mathbf{z}}; d)}{dd} \geq 0$ iff

$$\left(\frac{1 - b\rho}{1 - b(1 - \rho)} \right)^d \geq \frac{\ln(1 - b(1 - \rho))}{\ln(1 - b\rho)} \in (0, 1) \quad (5)$$

Since $\rho > \frac{1}{2}$, $\frac{1-b\rho}{1-b(1-\rho)}$ and the left-hand-side of (5) is exponentially decreasing in d , while the right-hand-side of (5) does not depend on d . We conclude that $\Delta V(\tilde{\mathbf{z}}; d)$ is strictly increasing in d up to some critical level \bar{d} and strictly decreasing after \bar{d} , i.e., $\Delta V(\tilde{\mathbf{z}}; d)$ is single-peaked in d . \square

¹¹For ease of exposition, we focus here and in the next subsection on the case in which $1 - f \in (1 - \rho, \rho)$, so that consumers would prefer to filter the news if they believed that all stories were low quality. This case is especially convenient since, by Theorem 1, there is a unique symmetric Nash equilibrium and news veracity in this equilibrium takes the relatively simple form $p_0^*(d) = \min\{\tilde{p}_0(d), \rho\}$.

Discussion: When consumers are relatively poorly connected, adding more links increases the producer’s incentive to invest. Intuitively, this is due to the “viral” amplifying effect that social sharing has for true stories, since true stories are more likely to be shared. For instance, suppose that broadcast reach $b \approx 0$, private-signal precision $\rho = \frac{2}{3}$, consumers filter the news, and each consumer is followed by d others. Each consumer who sees the broadcast will expose approximately $\frac{2d}{3}$ others to the story when it is true but only about $\frac{d}{3}$ when it is false. On the other hand, when consumers are already well-connected, increasing social connectedness decreases the producer’s incentive to invest. For instance, suppose that $b = \frac{1}{2}$, $\rho = \frac{9}{10}$, and consumers filter the news. Increasing social connectedness from $d = 10$ to $d = 20$ increases the likelihood that consumers see true stories from 99.8% to 99.9997% and false stories from 70% to 82%. The *extra* likelihood that consumers see true stories, which determines the producer’s incentive to invest, therefore falls from about 30% to about 18%.

3.4.2 Broadcast reach

The effect of increasing broadcast reach on equilibrium news veracity depends on how widely the broadcast is already seen. As in the last subsection, we focus for convenience on the case in which $1 - f \in (1 - \rho, \rho)$.

Proposition 3. *In the case when $1 - f \in (1 - \rho, \rho)$, $p_0^*(b)$ is non-monotone with $\frac{dp_0^*(0)}{db} > 0$ and $\frac{dp_0^*(1)}{db} < 0$.*

Proof. The maximal equilibrium news veracity $p_0^*(b) = \min\{\tilde{p}_0(b), \rho\}$, where $\tilde{p}_0(b) = 1 - f + fH(f\Delta V(\tilde{\mathbf{z}}; b))$ and

$$\Delta V(\tilde{\mathbf{z}}; b) = (1 - b) \left((1 - b(1 - \rho))^d - (1 - b\rho)^d \right) \quad (6)$$

is the extra visibility of true news when consumers filter the news. Note that $\Delta V(\tilde{\mathbf{z}}; 0) = \Delta V(\tilde{\mathbf{z}}; 1) = 0$; so, $\tilde{p}_0(0) = \tilde{p}_0(1) = 1 - f$ which in turn implies $p_0^*(1) = p_0^*(0) = 1 - f$ since $1 - f < \rho$. In order to show that $\frac{dp_0^*(0)}{db} > 0$ and $\frac{dp_0^*(1)}{db} < 0$, it therefore suffices to show that $\frac{d\tilde{p}_0(0)}{db} > 0$ and $\frac{d\tilde{p}_0(1)}{db} < 0$. By equation (6) and by our assumption that the c.d.f.

$H(\cdot)$ is smooth, it suffices further to show that $\frac{d\Delta V(\tilde{\mathbf{z}};0)}{db} > 0$ and $\frac{d\Delta V(\tilde{\mathbf{z}};1)}{db} < 0$, where

$$\begin{aligned} \frac{d\Delta V(\tilde{\mathbf{z}}; b)}{db} = & - \left((1 - b(1 - \rho))^d - (1 - b\rho)^d \right) \\ & + (1 - b)d \left(\rho(1 - b\rho)^{d-1} - (1 - \rho)(1 - b(1 - \rho))^{d-1} \right) \end{aligned} \quad (7)$$

By (7), $\frac{d\Delta V(\tilde{\mathbf{z}};0)}{db} = d(2\rho - 1) > 0$ and $\frac{d\Delta V(\tilde{\mathbf{z}};1)}{db} = -(\rho^d - (1 - \rho)^d) < 0$. \square

Discussion: Proposition 3 considers the relatively simple case in which low-quality stories have an intermediate likelihood of being true ($1 - f \in (1 - \rho, \rho)$), but qualitatively similar results hold in other more complex cases. For example, consider the case in which low-quality stories are always false ($f = 1$) but with the simplifying assumptions that each consumer follows one other person ($d = 1$) and the producer's per-consumer reporting cost is uniformly distributed on $[0, 1]$ ($H(x) = x$). For any given sharing rule \mathbf{z} , the extra visibility of true news takes the form $\Delta V(\mathbf{z}; b) = (1 - b)b(2\rho - 1)(z_T - z_F)$. This is maximized when consumers filter the news, using sharing rule $\tilde{\mathbf{z}} = (1, 0)$; so, the filtering news veracity $\tilde{p}_0(b)$ equals the maximal best-response news veracity $\bar{p}_0(b)$ and, after simplifying equation (4), takes the form

$$\tilde{p}_0(b) = \bar{p}_0(b) = (1 - b)b(2\rho - 1)$$

By Theorem 1: if $\tilde{p}_0(b) < 1 - \rho$, then no investment is possible in any equilibrium, i.e., $p_0^*(b) = 0$; if $\tilde{p}_0(b) \in (1 - \rho, \rho)$, then $p_0^*(b) = \tilde{p}_0(b)$; and if $\tilde{p}_0(b) > \rho$, then $p_0^*(b) = \rho$. For instance, suppose further that $\rho = \frac{7}{8}$. If $b < \frac{1 - \sqrt{1/3}}{2} \approx 0.21$ or $b > \frac{1 + \sqrt{1/3}}{2} \approx 0.79$, then $\tilde{p}_0(b) = \frac{3(1-b)b}{4} < 1 - \rho = \frac{1}{8}$ and the unique equilibrium has zero investment and hence news veracity $p_0^*(b) = 0$. On other hand, if $b \in (0.21, 0.79)$, then some investment can be supported in equilibrium and $p_0^*(b) = \tilde{p}_0(b) \in (1/8, 3/16)$.

3.4.3 Private-signal precision

The effect of increasing private-signal precision ρ on equilibrium news veracity $p_0^*(\rho)$ depends on how precise consumers' private signals already are and on whether low-quality stories are more likely to be true or false.

Proposition 4. *Signal-precision thresholds $\bar{\rho} \geq \underline{\rho} \geq \frac{1}{2}$ exist such that (i) $p_0^*(\rho) = 1 - f$ for all $\rho \in (\frac{1}{2}, \underline{\rho})$, (ii) $p_0^*(\rho) = 1 - \rho$ and $p_0^*(\rho)$ is strictly decreasing in ρ for all $\rho \in (\underline{\rho}, \bar{\rho})$, and (iii) $p_0^*(\rho) > 1 - \rho$ and $p_0^*(\rho)$ is strictly increasing in ρ for all $\rho \in (\bar{\rho}, 1)$.*

Proof. We break the proof down into two basic cases, depending on whether low-quality stories are more likely to be true ($1 - f \geq \frac{1}{2}$) or false ($1 - f < \frac{1}{2}$). First, we establish a preliminary result that the filtering news veracity $\tilde{p}_0(\rho)$ and the maximal best-response news veracity $\bar{p}_0(\rho)$ are each strictly increasing in ρ . Consider any fixed sharing rule $\mathbf{z} = (z_T, z_F)$ that favors true-story sharing, i.e., such that $z_T > z_F$. Increasing ρ increases the visibility of true stories and decreases the visibility of false stories; see equations (2-3). This increases the producer's incentive to invest and the resulting best-response news veracity $p_0(\mathbf{z}; \rho)$; see equation (4). Consequently, both $\tilde{p}_0(\rho) = p_0(\tilde{\mathbf{z}}; \rho)$ and $\bar{p}_0(\rho) = \max_{\mathbf{z}} p_0(\mathbf{z}; \rho)$ are increasing in ρ .

Case 1: $1 - f \geq \frac{1}{2}$. Define $\bar{\rho} = \underline{\rho} = \frac{1}{2}$. Prop 4(i-ii) hold vacuously with respect to these thresholds and, because $1 - f \geq \frac{1}{2} > 1 - \rho$, $p_0^*(\rho) > 1 - \rho$ holds automatically; so, it suffices to show that $p_0^*(\rho)$ is strictly increasing in ρ . But this follows immediately from the fact that $p_0^*(\rho) = \min\{\tilde{p}_0(\rho), \rho\}$ (by Theorem 1) and that $\tilde{p}_0(\rho)$ and ρ are each strictly increasing in ρ .

Case 2: $1 - f < \frac{1}{2}$. Define $\bar{\rho}$ and $\underline{\rho}$ implicitly by the conditions

$$\bar{p}_0(\underline{\rho}) = 1 - \underline{\rho} \tag{8}$$

$$\tilde{p}_0(\bar{\rho}) = 1 - \bar{\rho} \tag{9}$$

$\bar{\rho}$ and $\underline{\rho}$ are uniquely defined by (8-9), with $\bar{\rho} \geq \underline{\rho} > \frac{1}{2}$. To see why, observe first that, if private signals were completely uninformative ($\rho = \frac{1}{2}$), the producer would have zero incentive to invest no matter what sharing rule consumers used; so, $\bar{p}_0(\frac{1}{2}) = \tilde{p}_0(\frac{1}{2}) = 1 - f < \frac{1}{2}$. On the other hand, since $\bar{p}_0(\rho) \geq \tilde{p}_0(\rho) > 1 - f$ for all $\rho > \frac{1}{2}$, $\bar{p}_0(\rho) > 1 - \rho$ and $\tilde{p}_0(\rho) > 1 - \rho$ for all $\rho > f$. Since $\bar{p}_0(\rho)$ and $\tilde{p}_0(\rho)$ are each strictly increasing and continuous in ρ (by Step 1; continuity is obvious), we conclude that equations (8-9) each have a unique solution, greater than $\frac{1}{2}$ and less than 1. Finally, $\bar{\rho} \geq \underline{\rho}$ follows from the fact that $\bar{p}_0(\rho) \geq \tilde{p}_0(\rho)$ for all ρ .

By Theorem 1, either $\bar{p}_0(\rho) < 1 - \rho$ and the maximal equilibrium news veracity $p_0^*(\rho) = 1 - f$ or $\bar{p}_0(\rho) \geq 1 - \rho$ and $p_0^*(\rho) = \max\{1 - \rho, \min\{\tilde{p}_0(\rho), \rho\}\}$. Suppose first that $\rho < \underline{\rho}$. By definition of $\underline{\rho}$, $\bar{p}_0(\rho) < 1 - \rho$; so, $p_0^*(\rho) = 1 - f$. This completes the proof of Prop 4(i). Suppose next that $\rho \in (\underline{\rho}, \bar{\rho})$. By definition of $\underline{\rho}$ and $\bar{\rho}$, $\tilde{p}_0(\rho) < 1 - \rho < \bar{p}_0(\rho)$; so, $p_0^*(\rho) = 1 - \rho$, which is obviously decreasing in ρ . This completes the proof of Prop 4(ii). Finally, suppose that $\rho > \bar{\rho}$. In this case, $\tilde{p}_0(\rho) > 1 - \rho$ and hence $p_0^*(\rho) = \min\{\tilde{p}_0(\rho), \rho\} > 1 - \rho$, which is increasing in ρ (shown in Case 1). This completes the proof of Prop 4(iii). \square




Figure 6: Increasing private-signal precision ρ increases maximal equilibrium news veracity $p_0^*(\rho)$ if filtering news veracity $\tilde{p}_0(\rho) > 1 - \rho$ (Scenario 1, upper panel) but decreases $p_0^*(\rho)$ if $\tilde{p}_0(\rho) < 1 - \rho < \bar{p}_0(\rho)$ (Scenario 2, lower panel).

Discussion: Increasing private-signal precision ρ has two main direct effects. First, the fil-

tering region expands as the always-share threshold increases and the never-share threshold decreases (see Figure 6). This reflects the fact that consumers make use of more-informative private signals given a wider range of prior beliefs. Second, consumers are more likely to share true stories and less likely to share false stories when signal precision is higher. Although higher signal precision induces more investment for any fixed sharing rule, increasing ρ also changes consumers' *equilibrium* sharing behavior in ways that can reduce the producer's incentive to invest. For this reason, the overall effect of increasing ρ on equilibrium news veracity $p_0^*(\rho)$ is ambiguous. If the filtering news veracity is initially high enough that $\tilde{p}_0(\rho) > 1 - \rho$ (see the upper panel of Figure 6), increasing private-signal precision to $\hat{\rho} > \rho$ increases equilibrium news veracity from $\min\{\tilde{p}_0(\rho), \rho\}$ to $\min\{\tilde{p}_0(\hat{\rho}), \hat{\rho}\}$. On the other hand, if $\tilde{p}_0(\rho) < 1 - \rho < \bar{p}_0(\rho)$ (see the lower panel of Figure 6) and $\hat{\rho}$ is sufficiently close to ρ that $\tilde{p}_0(\hat{\rho}) < 1 - \rho$, then increasing private-signal precision to $\hat{\rho} = \rho$ decreases equilibrium news veracity from $1 - \rho$ to $\max\{\tilde{p}_0(\hat{\rho}), 1 - \hat{\rho}\}$.¹²

4 The Curse (and Limited Wisdom) of the Crowd

In “Vox Populi” (Galton (1907)), Francis Galton examined hundreds of entries in a contest at the “West of England Fat Stock and Poultry Fair,” in which people guessed the dressed weight of an ox. Although individual guesses varied widely, the average guess (1997 lbs) was very close to the ox's true weight (1998 lbs). In the same way, in our model, consumers in a large news market could almost-certainly discern which stories are true if their private signals could somehow be aggregated.

Consumer sharing over a highly-connected social network, what we refer to as a “crowd,” provides one means by which consumers' information could potentially be aggregated in a decentralized way.¹³ If consumers all follow many others and all filter the news, a “wisdom of the crowd” will emerge and consumers will almost always act on true stories and almost never act on false stories. However, increasing the number of

¹²In the last main possibility, when $\bar{p}_0(\rho) < 1 - \rho$, slightly increasing ρ has no effect on the unique equilibrium, which exhibits zero investment, zero sharing, and news veracity $p_0^* = 1 - f < 1 - \rho$.

¹³Golub and Jackson (2010) and Mueller-Frank (2014) show how dispersed information can be successfully aggregated through repeated communication over a social network.

consumers’ *social* sources of information can also have indirect equilibrium effects, as consumers may change how often they share stories and the producer may invest more or less frequently in story quality.

In this section, we examine equilibrium outcomes in the market for news when consumers are in a crowd, focusing on the limit case of a complete network as network size goes to infinity and the cost of reporting scales with network size,¹⁴ what we refer to as “the crowd limit.” More precisely, we consider a sequence of news markets in which (i) market size M goes to infinity, (ii) social connectedness $d = M - 1$, meaning that everyone follows everyone, and (iii) the distribution of per-consumer reporting cost c_R/M is held fixed.

This section’s analysis addresses two main questions. First, how does equilibrium news veracity change as consumers grow increasingly connected? We find that, in the crowd limit, equilibrium investment *either* vanishes entirely *or* remains positive but falls to the lowest level consistent with consumer sharing. Either way, being in a crowd exposes consumers to more false news than if they had been in a less-well-connected social network—a finding that we refer to as “the curse of the crowd”.

Second, to what extent are consumers in a crowd able to learn from others’ sharing behavior? We find that, for some model parameters, consumers in a crowd filter the news, allowing consumers to discern which stories are true and avoid acting on false information. However, in such cases when a wisdom of the crowd emerges, equilibrium investment always vanishes in the crowd limit. Consider in particular the case in which low-quality stories are always false ($f = 1$). In the crowd limit, equilibrium investment does not vanish but individuals’ likelihood of sharing any given story goes to zero at such a rate that consumers typically have a small number of sharing neighbors and cannot determine with high confidence which stories are true and which are false—a finding that we refer to as “the limited wisdom of the crowd”.

Some additional terminology is helpful for stating our main findings: “investment

¹⁴If the reporting-cost distribution were held fixed, the per-consumer cost of reporting would go to zero and arbitrarily-small differences in true- and false-news visibility would induce maximal investment in the crowd limit. Scaling the cost of reporting with market size ensures that our findings here are not being driven by economies of scale.

vanishes” if the producer’s likelihood of investing in the maximal-veracity equilibrium converges to zero in the crowd limit; “individual sharing vanishes” if each consumer’s ex ante likelihood of sharing any given story in the maximal-veracity equilibrium converges to zero; “the crowd is wise” if, in the crowd limit, consumers are able to perfectly discern which stories are true from others’ sharing behavior in the maximal-veracity equilibrium; and “the crowd is ignorant” if consumers are able to infer nothing from others’ sharing behavior.

Shorthand notation is also useful: For each market size M , let $\tilde{p}_0(M)$ denote the filtering news veracity, $\bar{p}_0(M)$ the maximal best-response news veracity, and $p_0^*(M)$ the maximal equilibrium news veracity when the social network is a complete graph, i.e., $d = M - 1$, and reporting cost scales with market size, i.e., c_R/M is drawn from the same c.d.f. $H(\cdot)$ for all M . Let $\tilde{p}_0^\infty = \lim_{M \rightarrow \infty} \tilde{p}_0(M)$, $\bar{p}_0^\infty = \lim_{M \rightarrow \infty} \bar{p}_0(M)$, and $p_0^{*\infty} = \lim_{M \rightarrow \infty} p_0^*(M)$ denote the crowd limit of these variables.

Theorem 2. *In the crowd limit: (i) If $1 - f > \rho$, then investment is zero, consumers always share, and the crowd is ignorant. (ii) If $1 - f \in (1 - \rho, \rho)$, then investment vanishes, consumers filter the news, and the crowd is wise. (iii) If $1 - f < 1 - \rho$ and $\bar{p}_0^\infty < 1 - \rho$, then investment is zero, consumers never share, and the crowd is ignorant. (iv) If $1 - f < 1 - \rho < \bar{p}_0^\infty$, then investment does not vanish, individual sharing vanishes, and the crowd is neither wise nor ignorant.*

Proof of Theorem 2(i): Suppose that $1 - f > \rho$. By Theorem 1, investment is zero and consumers always share given any market size M ; so, $p_0^*(M) = 1 - f$ and sharing is uninformative for all M . It follows immediately that $p_0^{*\infty} = 1 - f$ and the crowd is ignorant.

Proof of Theorem 2(ii): Suppose that $1 - f \in (1 - \rho, \rho)$. By Theorem 1, investment is positive, inducing equilibrium news veracity $p_0^*(M) = \min\{\tilde{p}_0(M), \rho\}$, given any market size $M \geq 1$. Since $\bar{p}_0^\infty = 1 - f$ [[by the argument in ...; make a Lemma earlier about Delta V??]] and $1 - f < \rho$, we conclude that $p_0^{*\infty} = 1 - f \in (1 - \rho, \rho)$; so, investment vanishes in the crowd limit. Since $p_0^{*\infty} \in (1 - \rho, \rho)$, consumers find it optimal to filter the news in the crowd limit; by the Law of Large Numbers, consumers are able to discern

perfectly which stories are true in this limit, i.e., the crowd is wise.

Proof of Theorem 2(iii): Suppose that $1 - f < 1 - \rho$ and $\bar{p}_0^\infty < 1 - \rho$. By Theorem 1, investment is zero and consumers never share given any sufficiently large market size M ; so, $p_0^*(M) = 1 - f$ and sharing is uninformative for all sufficiently large M . It follows immediately that $p_0^{*\infty} = 1 - f$ and the crowd is ignorant.

Proof of Theorem 2(iv): Suppose that $1 - f < 1 - \rho$ and $\bar{p}_0^\infty > 1 - \rho$. For any sufficiently large market size M , $\bar{p}_0(M) > 1 - \rho$ and hence, by Theorem 1, $p_0^*(M) = 1 - \rho$ and consumer sharing in the maximal-veracity equilibrium takes the form $\mathbf{z}^*(M) = (z_T^*(M), 0)$. An immediate implication is that $p_0^{*\infty} = 1 - \rho > 1 - f$; so, investment does not vanish in the crowd limit.

Next, we show that individual sharing must vanish in the crowd limit. Suppose for the sake of contradiction that $\lim_{M \rightarrow \infty} z_T^*(M) = X > 0$. In the crowd limit, true stories would then earn extra expected revenue

$$\lim_{M \rightarrow \infty} M f \Delta V(\mathbf{z}^*(M); M) = f(1 - b) \lim_{M \rightarrow \infty} M \left((1 - b(1 - \rho)X)^{M-1} - (1 - b\rho X)^{M-1} \right),$$

which equals zero since $M(1 - b(1 - \rho)X)^{M-1}$ and $M(1 - b\rho X)^{M-1}$ each converge to zero. But then investment must vanish in the crowd limit, a contradiction.

Finally, we establish a Goldilocks-type result that consumers in a crowd cannot learn too little (i.e., the crowd cannot be ignorant) but also cannot learn too much (i.e., the crowd cannot be wise) in the maximal-veracity equilibrium.¹⁵ We have shown that $p_0^*(M) = 1 - \rho$ for all sufficiently large M ; so, the producer must invest with ex ante probability $\frac{f - \rho}{f}$ in any sufficiently large market. Since investing generates extra expected per-consumer revenue $f \Delta V(z_T^*(M), 0; M)$, it must be that $H(f \Delta V(z_T^*(M), 0; M)) = \frac{f - \rho}{f}$, which in turn requires that the extra visibility of true stories

$$\Delta V(z_T^*(M), 0; M) = H^{-1} \left(\frac{f - \rho}{f} \right) \in (0, 1) \quad (10)$$

for all sufficiently large M . In particular, $\lim_{M \rightarrow \infty} \Delta V(z_T^*(M), 0; M) > 0$, i.e., true stories

¹⁵When $1 - f < 1 - \rho$, another equilibrium always exists in which the producer never invests, consumers never share, and the crowd is ignorant.

must be viewed more often than false stories in the crowd limit. By equation (6),

$$\Delta V(z_T^*(M), 0; M) = (1 - b) \left((1 - b(1 - \rho)z_T^*(M))^{M-1} - (1 - b\rho z_T^*(M))^{M-1} \right) \quad (11)$$

Using the basic mathematical fact that $\lim_{M \rightarrow \infty} (1 - X/M)^M = e^{-X}$, equation (11) implies

$$\lim_{M \rightarrow \infty} \Delta V(z_T^*(M), 0; M) = (1 - b) \left(e^{-b(1-\rho)\lim_{M \rightarrow \infty} Mz_T^*(M)} - e^{-b\rho\lim_{M \rightarrow \infty} Mz_T^*(M)} \right) \quad (12)$$

Since equation (10) must hold for all sufficiently large M , equation (12) implies that $\lim_{M \rightarrow \infty} Mz_T^*(M) = C$, where $C > 0$ solves $e^{-b(1-\rho)C} - e^{-b\rho C} = H^{-1} \left(\frac{f-\rho}{f} \right) / (1 - b)$.¹⁶

Having characterized the rate at which consumer sharing vanishes in the crowd limit, we can now characterize the crowd-limit visibility of true and false stories:

$$V_T^{*\infty} \equiv \lim_{M \rightarrow \infty} V_T(z_T^*(M), 0; M) = 1 - (1 - b)e^{-b\rho C} \quad (13)$$

$$V_F^{*\infty} \equiv \lim_{M \rightarrow \infty} V_F(z_T^*(M), 0; M) = 1 - (1 - b)e^{-b(1-\rho)C} \quad (14)$$

The fact that the crowd cannot be ignorant and also cannot be wise follows immediately from (13-14). Why? Each consumer is more likely to have zero sharing neighbors when a story is false than when it is true ($e^{-b(1-\rho)C}$ versus $e^{-b\rho C}$). Seeing that no one has shared (or that someone has shared) is therefore informative about the likelihood of news truth; hence, the crowd is not ignorant. On the other hand, since true stories also sometimes fail to be shared, consumers seeing no one share cannot discern for certain whether the story is true or false; hence, the crowd also is not wise. \square

5 Advertising- vs Sponsor-Supported Content

Note to readers of this draft: The rest of paper, starting with this section, is rough and incomplete.

In many situations, those producing decision-relevant information do so because they want to influence others' decisions. For instance, a company launching a social-media

¹⁶Such a solution always exists when $\bar{p}_0^\infty > 1 - \rho$. In fact, there are two solutions: the larger solution corresponds to the limit of dynamically-stable equilibria and the smaller solution corresponds to the limit of dynamically-unstable equilibria. See Figure 5 for a visualization and the Appendix for details.

campaign promoting a new product ultimately wants consumers to buy their product (i.e., “adopt”), not just to see and share the message. Here we extend the baseline model so that the producer is paid $\alpha^V \in [0, 1]$ per view and $\alpha^A \in [0, 1]$ per adoption, with $\alpha^V + \alpha^A = 1$. This approach allows us to model the news producer as caring about views only (the baseline case considered earlier), adoptions only, or a combination of both views and adoptions, and to compare equilibrium outcomes under different revenue models. Later in Section 5.3, we will further extend the model to allow for multiple producers with different motivations.

For ease of exposition and analytical simplicity, we focus here in the main text on equilibrium outcomes in the crowd limit, when each consumer follows many others, and on the case when low-quality stories are always false ($f = 1$) and consumers have a higher belief threshold for adopting than for sharing ($p^A > p^S = 1/2$). See the Appendix for a more detailed analysis, in which we implicitly characterize all Nash equilibria given any finite social network, any minimal news veracity, and any adoption threshold.

5.1 Special case: revenue only from adoptions

Figure 7 illustrates the market for news in the special case when the producer is paid only for adoptions, what we refer to as “sponsor-supported content”. Consumers’ sharing decisions still effect the producer’s incentive to invest in story quality, but now only indirectly by impacting how widely stories are seen and how much consumers can infer about a story’s truthfulness from others’ sharing behavior. For clarity of the exposition, we will focus on this case first, and then extend the analysis later to allow for revenue from both views and adoptions.

Consumer incentive to adopt. Suppose that consumer i has the opportunity to adopt, having become aware of a story by seeing the time-0 broadcast and/or having at least one sharing neighbor. Given private signal $s_i \in \{T, F\}$ and $\sigma \geq 0$ sharing neighbors, consumer i will update her belief about the likelihood of news truth to $p_2(s_i, \sigma; p_0, \mathbf{z})$. By

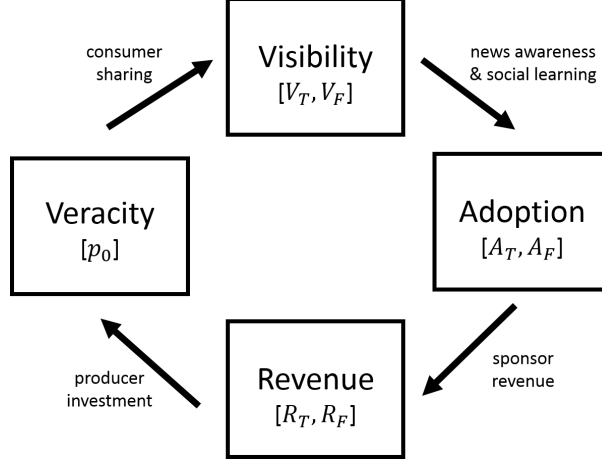


Figure 7: Schematic of the market for news when producers are paid for adoptions.

Bayes' Rule,

$$\frac{p_2(s_i, \sigma; p_0, \mathbf{z})}{1 - p_2(s_i, \sigma; p_0, \mathbf{z})} = \frac{p_1(s_i; p_0)}{1 - p_1(s_i; p_0)} \left(\frac{\Pr(\text{share}|\text{true}, \mathbf{z})}{\Pr(\text{share}|\text{false}, \mathbf{z})} \right)^\sigma \left(\frac{1 - \Pr(\text{share}|\text{true}, \mathbf{z})}{1 - \Pr(\text{share}|\text{false}, \mathbf{z})} \right)^{d-\sigma} \quad (15)$$

where $p_1(s_i; p_0)$ is consumer i 's time-1 belief based on her private signal alone (equation 1) and $\Pr(\text{share}|\text{true}, \mathbf{z}) = b(\rho z_T + (1 - \rho)z_F)$ and $\Pr(\text{share}|\text{false}, \mathbf{z}) = b((1 - \rho)z_T + \rho z_F)$ are each consumer's ex ante likelihood of sharing when the news is true or false, respectively, and consumers use sharing rule $\mathbf{z} = (z_T, z_F)$. Upon observing signal s_i and σ sharing neighbors, consumer i prefers to adopt when $p_2(s_i, \sigma; p_0, \mathbf{z}) > p^A$ and prefers not to share when $p_2(s_i, \sigma; p_0, \mathbf{z}) < p^A$. Let $A_T(p_0; \mathbf{z})$ and $A_F(p_0; \mathbf{z})$ denote the ex ante likelihood that any given consumer adopts when the story is true or false, respectively, and all consumers use the sharing rule \mathbf{z} .

Producer incentive to invest. True stories enjoy extra expected revenue $\Delta R(p_0; \mathbf{z}) = M\Delta A(p_0; \mathbf{z})$, where $\Delta A(p_0; \mathbf{z}) = A_T(p_0; \mathbf{z}) - A_F(p_0; \mathbf{z})$ is the extra likelihood that each consumer adopts based on a true story rather than a false one. Since by assumption low-quality stories are always false and high-quality stories are always true, the producer finds it optimal to invest in story quality so long as her realized per-consumer reporting cost c_R/M is less than $\Delta A(p_0; \mathbf{z})$. In particular, if consumers believe that published news is true with probability p_0 and use sharing rule \mathbf{z} , then optimal producer investment

induces news veracity $p_0(\mathbf{z}; p_0) = H(\Delta A(p_0; \mathbf{z}))$, where $H(\cdot)$ is the c.d.f. of c_R/M .

Note that an equilibrium exists with news veracity p_0 and symmetric sharing rule \mathbf{z} if and only if (i) \mathbf{z}^* is optimal for consumers given news veracity p_0^* , i.e., $\mathbf{z}^* \in Z(p_0^*)$, and (ii) p_0^* is the induced news veracity when the producer invests optimally in anticipation that consumers will use sharing rule \mathbf{z}^* and believe that news veracity equals p_0^* , i.e., $p_0^* = H(\Delta A(p_0^*; \mathbf{z}^*))$.

Equilibria in which consumers never share. The easiest case to consider is when consumers never share, i.e., when they use sharing rule $\mathbf{z} = (0, 0)$. Since consumers learn nothing from others' failure to share, a consumer who sees the broadcast and gets private signal s_i will prefer to adopt or not adopt, respectively, if her updated belief based on signal s_i is greater than or less than the adoption threshold p^A .

Can an equilibrium exist in which consumers never share and sometimes adopt? Recall that never sharing is a best response for consumers only when news veracity is less than or equal to the never-share threshold, i.e., $p_0 \leq 1 - \rho$, so that consumers' updated beliefs never exceed the sharing threshold $p^S = 1/2$. Because of our assumption here¹⁷ that the adoption threshold exceeds the sharing threshold, we conclude that, if news veracity is low enough that consumers never share, then consumers must also never adopt. Indeed, such an equilibrium always exists: the producer never invests, anticipating that consumers will never adopt, and consumers never share and never adopt, anticipating that all published stories are false.

Equilibria in which consumers always share. Suppose next that consumers always share, i.e., use sharing rule $\mathbf{z} = (1, 1)$. Because true and false stories are equally likely to be shared, consumers again infer nothing from others' sharing behavior. Whether a consumer finds it optimal to adopt depends on her private signal and whether news veracity is in a high, intermediate, or low region:

- *Always-adopt region:* If news veracity $p_0 > \bar{p}^A \equiv \frac{\rho p^A}{\rho p^A + (1-\rho)(1-p^A)}$, consumers find it

¹⁷The Appendix provides a more comprehensive analysis that allows for any adoption threshold $p^A \in (0, 1)$, as well as any minimal news quality $1 - f \in [0, 1)$.

optimal to always adopt, no matter what their private signal.¹⁸

- *Never-adopt region:* If news veracity $p_0 < \underline{p}^A \equiv \frac{(1-\rho)p^A}{(1-\rho)p^A + \rho(1-p^A)}$, consumers find it optimal to never adopt, no matter what their private signal.
- *Private-evaluation region:* If news veracity $p_0 \in (\underline{p}^A, \bar{p}^A)$, consumers find it optimal to adopt after a good private signal but not after a bad private signal.

Note that true stories are more likely to induce consumer adoption when news veracity is in the private-evaluation region, as each consumer's likelihood of getting a positive private signal is higher when the news is true, but that all stories generate the same amount of adoption when news veracity is in the always-adopt or never-adopt regions.

Can an equilibrium exist in which consumers always share and sometimes adopt? Yes, this is possible. In fact, so long as reporting costs are sufficiently low, such an equilibrium always exists with news veracity equal to \bar{p}^A . To see why, consider for a moment the related question of whether an equilibrium can exist with news veracity $p_0 \in (\rho, \bar{p}^A)$. In any such equilibrium, consumers always share (because $p_0 > \rho$) but only adopt after a good private signal (because $p_0 < \bar{p}^A$). Thus, consumers are equally aware of true and false stories, viewing every story with probability $\bar{V} = 1 - (1 - b)^{d+1}$, but more likely to adopt based on true stories. In particular, since consumers get a good signal with probability ρ or $1 - \rho$ when a story is true or false, respectively, each consumer's ex ante likelihood of adopting based on a true or false story equals $\rho\bar{V}$ or $(1 - \rho)\bar{V}$. Anticipating this, the producer finds it optimal to invest in story quality whenever per-consumer reporting cost $c_R/M < (2\rho - 1)\bar{V}$, which occurs with likelihood $H((2\rho - 1)\bar{V})$.

We conclude that an equilibrium with news veracity $p_0 \in (\rho, \bar{p}^A)$ does in fact exist, so long as the reporting-cost distribution $H(\cdot)$ is neither too high nor too low in the sense that $H^{-1}((2\rho - 1)\bar{V}) \in (\rho, \bar{p}^A)$. But what if the reporting-cost distribution is sufficiently low that $H^{-1}((2\rho - 1)\bar{V}) > \bar{p}^A$? No equilibrium can exist with news veracity in the interval (ρ, \bar{p}^A) because, if it did, the producer would have *too much* incentive to invest. However, in this case, an equilibrium always exists in which news veracity equals the always-adopt

¹⁸After getting a negative signal, consumer i 's updated belief after a bad private signal, $p_1(F; p_0) = \frac{(1-\rho)p_0}{(1-\rho)p_0 + \rho(1-p_0)}$, exceeds p^A iff $p_0 > \bar{p}^A$. Similarly, $p_1(T; p_0) < p^A$ iff $p_0 < \underline{p}^A$.

threshold \bar{p}^A . To see why, note that consumers are indifferent whether to adopt after getting a bad private signal and so find it optimal to adopt with probability one after a good signal and any probability $a_F \in [0, 1]$ after a bad signal. The extra likelihood that true stories induce adoption is then $(1 - a_F)(2\rho - 1)\bar{V}$, which varies from zero to $(2\rho - 1)\bar{V}$ as a_F varies from zero to one. Since $H^{-1}(0) = 0 < \bar{p}^A$ and $H^{-1}((2\rho - 1)\bar{V}) > \bar{p}^A$, there is a unique $a_F^* \in (0, 1)$ such that $H^{-1}((1 - a_F^*)(2\rho - 1)\bar{V}) = \bar{p}^A$. We conclude that, should consumers adopt after a bad signal with probability a_F^* , the producer will invest just enough to induce news veracity equal to \bar{p}^A , given which consumers are willing to randomize whether to adopt after a bad private signal.

Equilibria in which consumers filter the news. When consumers filter the news, true stories are more likely to induce consumer adoption than false stories, for a combination of reasons. First, since true stories are more likely to be shared, each consumer is more likely to become aware of the news and hence have the opportunity to adopt when a story is true. Second, because others' sharing behavior is informative, each consumer will be more likely to adopt when encountering a true story than when encountering a false story. Equilibria can indeed exist in which consumers filter the news (details in the Appendix) but, interestingly, such equilibria cannot exist if consumers have very many neighbors and the cost of investment is not especially high. To see why, note that there will be a wisdom of the crowd if consumers all filter the news, allowing consumers in the crowd limit to perfectly discern whether each story is true or false and only adopt when a story is true. Such perfect discernment by consumers creates the maximal possible incentive to invest, as the extra likelihood of true-story adoption equals 100% in the crowd limit. So long as the cost of investing in story quality is not so high that $H^{-1}(1) \leq \rho$, the producer's best response will then be to invest so often that the induced news veracity will be in the always-share region, a contradiction. On the other hand, if the reporting-cost distribution is sufficiently high that $H^{-1}(1) \in (1 - \rho, \rho)$, then an equilibrium does exist in the crowd limit in which consumers filter the news and the crowd is wise.¹⁹

¹⁹The last case when $H^{-1}(1) < 1 - \rho$ is trivial and uninteresting, with reporting costs so high that no investment can ever be supported in any equilibrium in the crowd limit.

5.2 Impact of the producer’s revenue model

Note to readers: This section is not yet written up, but here is the basic idea. In this section, we consider how varying the weight that the producer puts on views versus adoptions impacts equilibrium outcomes, focusing for simplicity and clarity on the crowd limit. When the producer only cares about consumer views and the cost of investment is low, equilibria in the crowd limit exhibit some investment (enough to support news veracity equal to $1 - \rho$) but each consumer’s likelihood of sharing vanishes in the limit fast enough that consumers fail to see some stories, even as their number of neighbors goes to infinity (Theorem 2). By contrast, when the producer only cares about consumer adoptions and the cost of investment is low, equilibria in the crowd limit exhibit more investment (enough to support news veracity equal to $\bar{p}^A > \rho$) and each consumer always shares.

Published stories are more likely to be true in the crowd-limit equilibrium when the reporter’s revenue is based on adoptions but, because consumers always share when revenue is based on adoptions, others’ sharing behavior is more informative when the reporter’s revenue is based on views. Consequently, the consumer-welfare impact of switching from a purely advertising-supported revenue model to a purely sponsor-supported model is ambiguous. As it turns out, however, best of all for consumers is when the producer has a mixture of motives, putting some weight on views and some on adoptions. Let $W^\infty(\alpha^A)$ denote consumers’ ex ante expected total welfare in the maximal-veracity equilibrium in the crowd limit, what we will refer to as “crowd welfare,” viewed as a function of the revenue weight α^A put on adoptions. We show that crowd welfare is not maximized at either $\alpha^A = 0$ or at $\alpha^A = 1$ but at some intermediate level, corresponding to a revenue model in which the reporter is paid for both views and adoptions.

5.3 Extension: multiple producers

Note to readers: This section is not yet written up, but here is the basic idea. Suppose that there is a population of news producers (“reporters”)—some paid for views and others paid for adoption—and consumers cannot tell which stories were written by which

type of reporter.²⁰ News veracity in this context is derived by averaging across the various reporters’ individual likelihoods of producing true news. In a large market in which widespread news filtering would allow consumers to almost-perfectly distinguish true from false stories, reporters motivated to drive adoption will have a strong incentive to invest in story quality while those motivated to get views will have almost no incentive to invest. Equilibrium news veracity is then determined by the fraction of reporters who are motivated by adoption. When the market is dominated by advertising-supported reporters, adding more sponsor-supported reporters leads news veracity to increase within the filtering region—making consumers and both types of reporters strictly better off than when only advertising-supported reporters were in the market.

6 Fake-news providers

Note to readers: This section is not yet written up, but here is the basic idea. Suppose that there is an unlimited supply of “fake-news providers” who can pay a cost c_F to produce a story that is definitely false (and who are incapable of producing true stories). The potential for economically-motivated²¹ fake news with free entry puts a new constraint on equilibrium sharing that, when binding at a sufficiently low level, can completely undermine equilibrium provision of high-quality news.

²⁰Of course, in many contexts, consumers can fairly easily infer the reporter’s motives for writing. For instance, someone writing an article for Huffington Post is likely motivated by views, while a political candidate making a public statement is motivated by actions / votes. Our single-producer model applies in such contexts, so long as the news producer lacks commitment power, but with different equilibrium outcomes for each producer depending on producer characteristics that impact model parameters. For instance, *New York Times* reporters enjoy a wider initial readership than most bloggers but also are likely subjected to more serious personal sanctions if their reporting is shown to be false. All else equal, we would therefore expect broadcast reach to be higher and reporting-cost distribution to be lower for a *New York Times* reporter than for a blogger.

²¹An “economically-motivated” fake-news provider is one who seeks to maximize expected profit and earns revenue from views and/or from adoption of their false stories. However, some fake-news providers may have more disruptive motives, for instance, to reduce consumers’ trust in news more generally. Such “enemies” could have much more damaging equilibrium impacts on the market for news.

A bit more detail: Note that the profitability of fake news depends only how often it is shared, which is increasing in consumers' sharing rule $\mathbf{z} = (z_T, z_F)$. Fake-news providers earn revenue from views, with total revenue $MV_F(\mathbf{z})$ since their stories are always false. If $c_F < M\underline{V}$, then fake-news providers strictly prefer to publish even if consumers never share. On the other hand, if $c_F > M\overline{V}$, then fake-news providers strictly prefer not to publish even if consumers always share. Focus now on the least trivial case when $M\Delta \leq c_F \leq M\overline{V}$. Among the sharing rules that could potentially be a best response for consumers (of the form $\mathbf{z} = (z_T, 0)$ and $\mathbf{z} = (1, z_F)$), there is a threshold sharing rule $\mathbf{z}^{fake} > (0, 0)$ given which fake-news providers are indifferent whether to produce. There are three main possibilities:

1. Consumers share less than \mathbf{z}^{fake} and zero fake news is produced.

This case arises whenever equilibrium sharing in the baseline model without fake-news providers is less than \mathbf{z}^{fake} . For instance, suppose that the maximal equilibrium news veracity (without fake news) $p^* < \rho$ but $\mathbf{z}^{fake} = (1, z_F^{fake})$ for some $z_F^{fake} > 0$. Since $p^* < \rho$, consumers never share after getting a negative signal and equilibrium sharing in the baseline model is less than \mathbf{z}^{fake} . In this case, the presence of fake-news providers has no effect on equilibrium outcomes since they choose not to produce fake news.

2. Consumers share more than \mathbf{z}^{fake} and an infinite supply of fake news is produced.

This possibility can never arise in equilibrium. Because of our assumption that there is an unlimited supply of fake-news providers (but a finite number of bona fide reporters), the news market would be overwhelmed by fake news and consumers would assess a zero likelihood that published stories are true—and hence never share, a contradiction since $\mathbf{z}^{fake} > (0, 0)$.

3. Consumers use sharing rule \mathbf{z}^{fake} and a finite amount of fake news is produced.

Equilibria with a finite amount of fake news can arise when equilibrium sharing in the baseline model without fake-news providers exceeds \mathbf{z}^{fake} . For instance, suppose that $p^* \in (1 - \rho, \rho)$ but $\mathbf{z}^{fake} = (z_T^{fake}, 0)$ for some $z_T^{fake} < 1$. Since $p^* \in (1 - \rho, \rho)$, consumers

in the baseline model use the filtering rule $\mathbf{z}^F = (1, 0) > \mathbf{z}^{fake}$. In this case, once fake-news providers are introduced, there cannot be an equilibrium with news veracity greater than $1 - \rho$ since, if there were, an infinite supply of fake news would be supplied and news veracity would be zero, a contradiction.

But can an equilibrium exist with at least some fake news? The answer depends on the likelihood $p_0(\mathbf{z}^{fake})$ that stories produced by bona fide reporters are true, derived in equation 4. If $p_0(\mathbf{z}^{fake}) \geq 1 - \rho$, then an equilibrium exists in which bona fide producers invest just often enough so that their stories are true with probability $p_0(\mathbf{z}^{fake})$ and just enough fake news is produced to drive the overall likelihood of news truth down to $1 - \rho$, so that consumers in turn find the sharing rule $\mathbf{z}^{fake} = (z_T^{fake}, 0)$ to be a best response. On the other hand, if $p_0(\mathbf{z}^{fake}) < 1 - \rho$, then overall news veracity in any equilibrium must be less than $1 - \rho$, implying that consumers never share which in turn implies that no fake news is produced.

Note that, in this last case, the presence of fake-news providers undermines the news market by eliminating any equilibrium in which bona-fide producers sometimes invest in story quality—even though, in the end, no fake news is produced. This finding can be interpreted in dynamic terms, whereby a news market in a positive-investment equilibrium is disturbed by the introduction of fake-news providers, who find it profitable to produce until consumers sufficiently lose faith in the news to stop sharing stories, at which point bona-fide producers no longer invest and fake-news providers no longer produce.

7 Concluding remarks

This paper is, to the best of our knowledge, the first to consider how the structure of consumers' social network impacts the quality of the information that spreads through the network. A key insight that emerges from our analysis, in the context of media markets, is that increased social connection may strengthen or weaken the quality of news, depending on the density of the network and on how news producers are paid. In large networks in which each consumer follows many others, the equilibrium quality of news is relatively poor in markets where news producers are paid when consumers

see their stories and relatively good in markets where news producers are paid when consumers *believe* the story enough to take an action. Even in that case, however, there is a limit to the overall news quality that can emerge in equilibrium and a limit to the amount of information aggregation that even an arbitrarily large network of consumers can achieve—the ideal of “the wisdom of crowds” is beyond reach.

References

- Acemoglu, Daron, Asuman Ozdaglar and Ali ParandehGheibi. 2010. “Spread of (Mis)information in Social Networks,” *Games and Economics Behavior*, 70 (2), pp. 194-227.
- Andreottola, Giovanni Andreottola and Antoni-Italo de Moragas. 2018. “Scandals, Media Competition and Political Accountability,” working paper.
- Banerjee, Abhijit. 1992. “A Simple Model of Herd Behavior,” *The Quarterly Journal of Economics*, 107(3), pp. 797-817.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. “The Diffusion of Microfinance,” *Science* 341(6144), pp. 363-370.
- Besley, Timothy and Andrea Prat. 2006. “Handcuffs for the Grabbing Hand? Media Capture and Government Accountability,” *American Economic Review*, 96(3), pp. 720-736.
- Bikhchandani, Sushil, David Hirshleifer and Ivo Welch. 1992. “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100(5), pp. 992-1026.
- Bloch, Francis, Gabrielle Demange, and Rachel Kranton. 2018. “Rumors and Social Networks,” *International Economic Review*, 59(2), pp. 421-448.
- Boyd, William. 2006. “The Secret Persuaders”, *The Guardian*, 19 August.
- Chatterjee, Kalyan and Bhaskar Dutta 2016. “Credibility and Strategic Learning in Networks,” *International Economic Review*, 57(3), pp. 759-786.
- Chen, Heng and Wing Suen. 2018. “Competition for Attention in the News Media Market,” working paper.

Cull, Nicholas. 1995. *Selling War: The British Propaganda Campaign Against American "Neutrality" in World War II*. Oxford: Oxford University Press.

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

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

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

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

Gentzkow, Matthew, Edward Glaeser, and Claudia Goldin. 2006. "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*. Chicago: University of Chicago Press.

Golub, Benjamin and Matthew O. Jackson. 2010. "Naive Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2(1), pp. 112–149.

Mueller-Frank, Manuel. 2014. "Does one Bayesian make a Difference," *Journal of Economic Theory*, 154, pp. 423–452.

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

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

A Omitted proofs and supplementary analysis

A.1 Proof of Lemma 1

Parts (i,ii,v). By equation (4), $p_0(z_T, z_F) = 1 - f + fH(f\Delta V(z_T, z_F))$, where $\Delta V(z_T, z_F) = (1-b) \left((1-b((1-\rho)z_T + \rho z_F))^d - (1-b(\rho z_T + (1-\rho)z_F))^d \right)$ by equations (2,3). Recall that, by assumption, the producer's reporting cost has a continuous c.d.f. and support on the interval $[0, M]$, implying that $H(0) = 0$ and that $H(c_R)$ is continuous and strictly increasing in c_R over the relevant range.²² Parts (i,ii) of Lemma 1 follow 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 (iii). Define $x(z_T) = \frac{\Delta V(z_T, 0)}{1-b}$. To prove part (iii), it suffices to show that $x(z_T)$ is strictly increasing in z_T over the interval $[0, \bar{z}_T]$ and strictly decreasing in z_T over $[\bar{z}_T, 1]$ for some $\bar{z}_T \in (0, 1]$. Note that $x'(z_T) = db \left(\rho(1 - b\rho z_T)^{d-1} - (1-\rho)(1 - b(1-\rho)z_T)^{d-1} \right)$. 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 $\bar{z}_T = 1$. Suppose next that $d \geq 2$. $x'(z_T) > 0$ iff $\frac{\rho}{1-\rho} > \left(\frac{1-b(1-\rho)z_T}{1-b\rho z_T} \right)^{d-1}$ 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 $\bar{z}_T \equiv \min\{\hat{z}_T, 1\}$.

Part (iv). Define $y(z_T, z_F) = \frac{\Delta V(z_T, z_F)}{1-b}$. To prove part (iv), we need to show that $y(\bar{z}_T, 0) \geq y(z_T, z_F)$ 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 $y(\bar{z}_T, 0) > 0$. Restricting attention to sharing rules with $z_T > z_F$, note that

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

since $1 - \rho < \rho$ and $\rho z_T + (1-\rho)z_F > (1-\rho)z_T + \rho z_F$. Finally, $y(\bar{z}_T, 0) \geq y(z_T, 0)$ for all $z_T \in [0, 1]$ by definition of \bar{z}_T . We conclude $y(\bar{z}_T, 0) \geq y(z_T, z_F)$ for all z_T, z_F . \square

²²Because $f \leq 1$, $b > 0$, and the producer is paid one unit of revenue per consumer who views the story, $0 \leq f\Delta V(z_T, z_F) < M$ no matter how consumers share the news.

A.2 Supplementary analysis related to Theorem 1

In Theorem 1 and the surrounding text, we focused on the question of whether an equilibrium exists with positive investment and on characterizing the maximal news veracity that can be supported in any symmetric Nash equilibrium (SNE). Here we augment that analysis by characterizing *all* Nash equilibria in the baseline case when the producer is paid only for views.²³ Recall by Steps 1-2 of the proof of Theorem 1 that (i) no Nash equilibrium exists with positive investment and news veracity in the always-share region and (ii) no Nash equilibrium exists with positive investment and news veracity in the never-share region. In what follows, we catalogue all equilibria in the remaining possibilities: Nash equilibria with no investment and hence news veracity $p_0 = 1 - f$; and Nash equilibria with positive investment and news veracity $p_0 \in (1 - \rho, \rho)$, $p_0 = \rho$, or $p_0 = 1 - \rho$.

Nash equilibria with minimal news veracity. Whether a Nash equilibrium exists with no investment depends on the likelihood $1 - f$ that low-quality stories are true, what we refer to as the “minimal news veracity.” There are three main cases.

First, suppose that $1 - f \geq \rho$. In this case, a Nash equilibrium exists in which consumers always share and the producer never invests; moreover, this is the unique Nash equilibrium. Uniqueness is obvious in the subcase when $1 - f > \rho$, since then consumers strictly prefer to always share no matter how what the producer does, giving the producer zero incentive to invest. Consider now the subcase when $1 - f = \rho$. The producer cannot invest with positive probability in any Nash equilibrium since, if she did, news veracity would exceed ρ , consumers would always invest, and the producer would have zero incentive to invest, a contradiction. Any Nash equilibrium must therefore have news veracity equal to ρ , making consumers indifferent whether to share after getting a bad private signal, i.e., each consumer i must use a sharing rule of the form $\mathbf{z}_i = (1, z_{iF})$ for some $z_{iF} \in [0, 1]$. Moreover, consumers must always share since, if any consumer(s) mixed after a bad signal with probability $z_{iF} < 1$, true stories would enjoy greater

²³Characterizing all Nash equilibria in the more general case when the reporter may also be paid for adoptions is substantially more complex and omitted.

visibility and the producer would sometimes prefer to invest, a contradiction.

Second, suppose that $1 - f \in (1 - \rho, \rho)$. In this case, all Nash equilibria exhibit positive investment. To see why, suppose that a Nash equilibrium exists in which the producer never invests. News veracity in this equilibrium equals $1 - f$, giving consumers an incentive to filter the news. This causes true stories to enjoy greater visibility, giving the producer an incentive to sometimes invest, a contradiction.

Finally, suppose that $1 - f \leq 1 - \rho$. In this case, a Nash equilibrium exists in which consumers never share and the producer never invests.

Best-response news veracity in a general directed graph. Let $\vec{z} = (z_i : i = 1, \dots, M)$ denote consumers' sharing-rule profile and let N_i be the set of others that consumer i follows. Each consumer i 's likelihood of viewing true and false stories is

$$\begin{aligned} V_{iT}(\vec{z}) &= 1 - (1 - b)\Pi_{j \in N_i}(1 - b(\rho z_{jT} + (1 - \rho)z_{jF})) \\ V_{iF}(\vec{z}) &= 1 - (1 - b)\Pi_{j \in N_i}(1 - b((1 - \rho)z_{jT} + \rho z_{jF})) \end{aligned}$$

so that the extra visibility of true stories to consumer i is

$$\Delta V_i(\vec{z}) = (1 - b) (\Pi_{j \in N_i}(1 - b((1 - \rho)z_{jT} + \rho z_{jF})) - \Pi_{j \in N_i}(1 - b(\rho z_{jT} + (1 - \rho)z_{jF}))).$$

The extra expected revenue associated with investing in story quality is now $f \sum_i \Delta V_i(\vec{z})$, inducing best-response news veracity $p_0(\vec{z}) = 1 - f + fH\left(f \frac{\sum_i \Delta V_i(\vec{z})}{M}\right)$.

Nash equilibria with news veracity in the filtering region. Suppose for a moment that a Nash equilibrium exists with positive investment and news veracity $p_0 \in (1 - \rho, \rho)$. Consumers strictly prefer to filter the news, using the sharing rule $\tilde{z} = (1, 0)$ and resulting in “filtering news veracity” $\tilde{p}_0 = p_0(\tilde{z}, \dots, \tilde{z})$. Thus, such a Nash equilibrium can only exist if $\tilde{p}_0 \in (1 - \rho, \rho)$; moreover, if that is the case, there is a unique Nash equilibrium with news veracity in the filtering region, in which all consumers filter the news and news veracity equals \tilde{p}_0 .

Nash equilibria with news veracity at the always-share threshold. Suppose for a moment that a Nash equilibrium exists with positive investment and news veracity

$p_0 = \rho$. Consumers' optimal sharing rules take the form $\mathbf{z}_i = (1, z_{iF})$ where $z_{iF} \in [0, 1]$ for all i . Since $z_{iT} = 1$ for all i and $\rho > 1/2$,

$$\frac{\partial \Delta V_i(\vec{z})}{\partial z_{jF}} = (1-b) \left((1-\rho) \prod_{k \neq j \in N_i} (1 - b(\rho + (1-\rho)z_{kF})) - \rho \prod_{k \neq j \in N_i} (1 - b(1-\rho + \rho z_{kF})) \right) < 0$$

for all i and all neighbors $j \in N_i$. Thus, $\frac{\partial p_0(\vec{z})}{\partial z_{jF}} < 0$ for all $j \in N_i$. We conclude that news veracity in any such Nash equilibrium cannot exceed the filtering news veracity \tilde{p}_0 and, in particular, that no such Nash equilibrium exists if $\tilde{p}_0 < \rho$. Similarly, if $\tilde{p}_0 = \rho$, there is a unique equilibrium with news veracity ρ , in which all consumers use the filtering sharing rule $\tilde{z} = (1, 0)$. Finally, if $\tilde{p}_0 > \rho$, the fact that $\frac{\partial p_0(\vec{z})}{\partial z_{jF}} < 0$ and $p_0((1, 1), \dots, (1, 1)) = 1 - f < \rho$ implies that there is a $(M - 1)$ -dimensional set of Nash equilibria with news veracity equal to ρ , in each of which consumers use a profile of sharing rules that creates the same overall incentive to invest in story quality.

This $(M - 1)$ -dimensional set of Nash equilibria includes a unique Nash equilibrium in which all consumers use the same sharing rule. In the main text, we characterized this unique symmetric equilibrium in the special case with all consumers follow d others.

Nash equilibria with news veracity at the never-share threshold. Suppose for a moment that a Nash equilibrium exists with positive investment and news veracity $p_0 = 1 - \rho$. Consumers' optimal sharing rules take the form $\mathbf{z}_i = (z_{iT}, 0)$ where $z_{iT} \in [0, 1]$ for all i , resulting in extra visibility

$$\Delta V_i(\vec{z}) = (1 - b) \left(\prod_{j \in N_i} (1 - b(1 - \rho)z_{jT}) - \prod_{j \in N_i} (1 - b\rho z_{jT}) \right)$$

for true stories. This results in extra revenue $\Delta R(\vec{z}) = f \sum_i \Delta V_i(\vec{z})$ for true stories and hence best-response news veracity $p_0(\vec{z}) = 1 - f + fH(\Delta R(\vec{z}))$. Building on the notational shorthand used in the main text, let $\bar{p}_0 = \max_{\vec{z}: z_{iF}=0 \text{ for all } i} p_0(\vec{z})$ denote the maximal equilibrium news veracity that can be supported when consumers use sharing rules of the form $\mathbf{z}_i = (z_{iT}, 0)$. As can be easily checked, $p_0(\vec{z})$ is continuous in z_{iT} for all i and $p_0((0, 0), \dots, (0, 0)) = 1 - f < 1 - \rho$. Thus, a Nash equilibrium with news veracity $1 - \rho$ exists iff $\bar{p}_0 \geq 1 - \rho$ and such an equilibrium exists with sharing-rule profile \vec{z} iff $p_0(\vec{z}) = 1 - \rho$.

The set of solutions to $p_0((z_{iT}, 0) : i = 1, \dots, M) = 1 - \rho$ is not easily characterized in the general case, because $p_0((z_{iT}, 0) : i = 1, \dots, M)$ is non-monotone in $(z_{iT} : i = 1, \dots, M)$. However, consider for a moment the special case examined in the main text: each consumer follows the same number d of others and uses the same sharing rule $(z_T, 0)$. Viewed now as a function of z_T , the extra visibility of true stories takes the form $\Delta V(z_T) = (1 - b) ((1 - b(1 - \rho)z_T)^d - (1 - b\rho z_T)^d)$, with derivative $\Delta V'(z_T) \geq 0$ iff $z_T \leq \bar{z}_T$, where \bar{z}_T was defined in the proof of Lemma 1(iv). Consequently, there may be up to two symmetric Nash equilibria with news veracity equal to $1 - \rho$, one in which consumers share with likelihood greater than \bar{z}_T after a good signal and another in which they share less than \bar{z}_T after a good signal. (As we discuss later in Appendix [[xx]], of these two symmetric Nash equilibria, the one with more sharing is dynamically stable while the one with less sharing is dynamically unstable.)

B Dynamic stability

Thm 1 characterizes all SNE, but not all of these equilibria are dynamically stable with respect to small changes in consumers' sharing rule. Let $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}}) = p_0(\mathbf{z}(1 - \epsilon) + \hat{\mathbf{z}}\epsilon)$ be the best-response news veracity that would result if consumers randomized between using sharing rule \mathbf{z} and sharing rule $\hat{\mathbf{z}}$, with weight $\epsilon \in (0, 1)$ on $\hat{\mathbf{z}}$.

Definition 1 (Dynamic stability). *A symmetric 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 $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}})$.²⁴ 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 2 (Dynamic stability). *(i) Any SNE with news veracity $p_0^* \notin \{1 - \rho, \rho\}$ is dynamically stable. (ii) Any SNE with news veracity $p_0^* = \rho$ is dynamically stable. (iii)*

²⁴Implicit 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 (1982)); straightforward details omitted to save space.

Any SNE with news veracity $p_0^* = 1 - \rho$ in which consumers use sharing rule $(z_T^*, 0)$ for some $z_T^* > 0$ is dynamically unstable if $z_T^* \in [0, \bar{p}_0]$ and dynamically stable if $z_T^* \in (\bar{p}_0, 1]$.

Proof: Given a SNE with sharing rule \mathbf{z}^* , say that sharing rule $\hat{\mathbf{z}}$ can “successfully invade” if $\hat{\mathbf{z}}$ is a better reply than \mathbf{z}^* given “perturbed news veracity” $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}}) = p_0(\mathbf{z}(1 - \epsilon) + \hat{\mathbf{z}}\epsilon)$ for all $\epsilon \approx 0$. By definition, a SNE with sharing rule \mathbf{z}^* is dynamically stable iff no other sharing rule $\hat{\mathbf{z}}$ can successfully invade.

Part One: $p_0^* \notin \{\rho, \rho\}$. Consumers have a unique best response $\mathbf{z}(p_0^*)$ equals $(1, 1)$ 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 $\epsilon \approx 0$, $p_0^\epsilon(\mathbf{z}; \hat{\mathbf{z}}) \approx p_0^*$. (Note by equation (4) that $p_0(\mathbf{z})$ is continuous in \mathbf{z} , a fact we used 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).

Part Two: $p_0^* = \rho$. Consumers strictly prefer to share given signal $s_i = T$ and are indifferent whether to share given signal $s_i = F$; so, the equilibrium sharing rule must take the form $\mathbf{z}^* = (1, z_F^*)$ for some $z_F^* \in [0, 1]$. For any $\hat{\mathbf{z}} \neq (1, z_F^*)$, 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 invade has three steps. First, consider any $\hat{\mathbf{z}}$ with $\hat{z}_T < 1$. After getting 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 getting 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 and hence $\hat{\mathbf{z}}$ cannot successfully invade. 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}^* and so cannot successfully invade. 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 $\hat{z}_F > z_F^*$. We again conclude that $\hat{\mathbf{z}}$ cannot successfully invade. This completes the proof of (ii).

Part Three: $p_0^* = 1 - \rho$. Consumers are indifferent whether to share after getting a positive signal $s_i = T$ and strictly prefer not to share after a negative signal $s_i = F$; so, the equilibrium sharing rule must take the form $\mathbf{z}^* = (z_T^*, 0)$ for some $z_T^* \in [0, 1]$. As in Part Two, one can easily show that any $\hat{\mathbf{z}}$ with $\hat{z}_F > 0$ cannot successfully invade; so, we will only consider sharing rules of the form $\hat{\mathbf{z}} = (\hat{z}_T, 0)$. Suppose first that $z_T^* \leq \bar{z}_T$ 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 $[0, \bar{z}_T]$; 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, sharing rule $\hat{\mathbf{z}} = (0, 0)$ *can* successfully invade; so, the SNE in question is dynamically unstable. Suppose next that $z_T^* > \bar{z}_T$ 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 $(\bar{z}_T, 1]$; 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 SNE in question is dynamically stable. This completes the proof of (iii).