Social Connectedness and the Market for Information

Rachel Kranton and David McAdams*

February 21, 2022

Abstract: This paper introduces a simple model of contemporary information markets: Consumers prefer high-quality information, judiciously sharing stories and posts. High-quality stories are costly to produce, and overall quality is endogenous. When producers' payoffs derive from how many consumers see their stories, quality is highest when social connectedness is neither too high nor too low. In highly-connected markets, low-quality stories are widely seen and dominate. Third-party misinformation can increase high-quality output, since consumers share more judiciously. When producers' payoffs depend on consumer actions (e.g, votes or purchases) and consumers are highly connected, consumers perfectly infer quality and misinformation has no impact.

Keywords: social-media networks, news veracity, social learning, misinformation, wisdom of the crowd

JEL Classification: C72, D62, D83

*Kranton: Economics Department, Duke University (email: rachel.kranton@duke.edu). McAdams: Fuqua School of Business and Economics Department, Duke University (email: david.mcadams@duke.edu), Durham, NC USA 27708. We thank seminar participants at Cornell, Dartmouth, Duke, EUI, GSE Barcelona, Johns Hopkins, Oxford, Northwestern, Penn State, UAB IAE (Barcelona),the 6th European Conference on Networks, the 17th IO Theory Conference, the 2019 Cowles Conference in Economic Theory, and the 6th Conference on Network Science and Economics for helpful comments. We also thank Guglhupf Bakery, Cafe, & Biergarten in Durham, NC for its hospitality while this research was conducted.

In September 2020, Facebook announced that users of its Messenger service would be limited to forwarding messages to only five people or groups at a time. As Facebook explained, "limiting forwarding is an effective way to slow the spread of viral misinformation and harmful content that has the potential to cause real world harm" (Facebook (2020)). But limiting sharing also reduces the spread of useful information and could therefore affect the incentives of bona fide information providers. Could limiting message sharing backfire, with unintended negative consequences on the producers and consumers of news and social media posts?

In this paper, we introduce a simple model of the contemporary information market to study the impact of such policies. On one side of the market are consumers, who choose whether to share decision-relevant information, called "stories," and who take actions based on stories that they encounter. On the other side are information producers, who decide whether to produce low-quality stories at zero cost or pay costs to produce high-quality stories by, for example, verifying sources in the case of a news story or testing a product in multiple ways in the case of a social influencer endorsing a product. High-quality stories are more valuable to consumers, with more accurate information about a policy issue, a political candidate, or a product. We assume that consumers cannot observe story quality directly but can (imperfectly) evaluate quality, and that consumers prefer only to share and take action on stories they believe are sufficiently likely to be high quality. Consumers rationally infer story quality from their own evaluation of the story and from the number of their social contacts who have shared the story with them. Producers receive benefits *either* based on the number of consumers who view their stories *or* based on the number of consumers who take action prompted by their stories.

We use this model to assess the equilibrium impact of people's "social connectedness," namely, how many others with whom they can share a story. We also study how misinformation injected into the market from outside sources affects the equilibrium quality of information produced. Our results emphasize how social connectedness and the nature of producer motivations impact equilibrium news quality.

Consider first a market where producers benefit from consumer views, as when newspapers earn revenue from advertising accompanying their stories or when social influencers gain stature when their posts are viewed more widely. If consumers have few sharing links, we show that adding connections increases equilibrium quality. Producers correctly anticipate that high-quality stories will be shared and hence viewed more frequently. But as consumers become very connected, producers have little incentive to invest in story quality since low-quality stories also spread and are viewed widely. Thus, story quality is highest when social connectedness is in an intermediate range.

Consider instead information markets where producers benefit when consumers take actions based on their stories, e.g., by supporting a particular candidate or purchasing a promoted product. If consumers are highly connected, we show that consumers can accurately infer in equilibrium whether a story is high quality based on their neighbors' sharing decisions, which in turn gives producers a strong equilibrium incentive to invest in story quality. In that context, policies like Facebook's which limit sharing can undermine consumers' ability to infer story quality and thereby induce information providers to produce fewer high-quality stories.

What about misinformation? Adam Mosseri, the head of product development for Facebook's news feed, wrote in 2017 that "a lot of fake news is [produced by] spammers ... masquerading as legitimate news publishers" (Facebook (2017)). If consumers are unable to easily distinguish between legitimate and fake news stories, they become part of the same "information market." While false stories injected into an information market always reduce consumers' (ex ante) confidence in the stories, we find that more misinformation can in some cases induce bona fide producers to invest more in quality and enable consumers to better infer which stories are high quality. The reason is that, as consumers expect that more stories are false, they become more judicious when deciding which stories to share, which in turn increases the informativeness of sharing decisions and enables their social contacts to better infer story quality.

The paper contributes to at least three bodies of literature:

Social learning, information transmission, and social connections. The demand side of our market specifies information transmission and social learning that is both similar to and different from other models. Consumers here receive private signals and rationally update beliefs about the quality of each story based on others' sharing decisions. Unlike in the cascades literature (e.g., Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)), consumers in the present paper observe multiple neighbors' independent sharing decisions in one round of social learning. Unlike much of the network literature on information diffusion (e.g., Acemoglu, Ozdaglar, and ParandehGheibi (2010) and Banerjee et al. (2013)), consumers in our model choose whether or not to pass on information to their neighbors (as in Bloch, Demange, and Kranton (2018) and Chatterjee and Dutta (2016)).

On the supply side, consumers' sharing decisions ultimately determine producers' incentives to invest in story quality. To the best of our knowledge, the present paper is the first to endogenize the "product" itself that spreads socially.[1] Previous research studies the effect of social-network structure on other producer decisions for a given product, such as relying on traditional versus word-of-mouth advertising (Galeotti and Goyal (2009)) or targeting consumers when launching a new product (Chatterjee and Dutta (2016), Bimpikis, Ozdaglar, and Yildiz (2016)).

[1] Many papers in diverse fields have examined how network structure impacts the decisions of a third party who cares about outcomes, e.g., a health authority deciding how best to control an epidemic (Peng et al. (2013)) or a supply-chain manager deciding how best to operate its warehouses (Beamon and Fernandes (2004)). The idea of endogenizing what passes through the network is rarely explored in these literatures, but there are exceptions, e.g., Read et al. (2015) on endogenous pathogen virulence and Bimpikis, Fearing, and Tahbaz-Salehi (2018) on upstream sourcing in a supply chain.

Our simple model highlights how suppliers' production incentives depend on social connectedness (the number of neighbors people have), abstracting from other aspects of network structure. The model captures the basic features of consumer sharing behavior in social networks; consumers *filter* the information that they receive and consumers *spread* information to their social contacts. These two forces, the filtering and spreading of information, form the bases of producer incentives. The effects of consumer sharing strategies on producer incentives and producer investment strategies on consumer sharing gives the equilibrium quality of information.

Media markets. Much previous work on news markets studies media bias. In Gentzkow and Shapiro (2006), news producers benefit from having a reputation for accuracy and thus have an incentive to slant their news towards consumers' initial beliefs. In Besley and Prat (2006) and Gentzkow, Glaeser, and Goldin (2006), earning revenue from advertising, rather than a sponsor, reduces bias. In Ellman and Germano (2009), however, newspapers bias their news towards their advertisers. The present paper does not consider political slants or opposing views. Consumers care only about story quality, where high-quality stories contain information that can prompt consumer actions such as purchasing or voting; we refer to expected quality as "veracity." A key finding is that, when every consumer "follows" many others, equilibrium veracity in an information market is lower when producers' payoffs depend on advertising (based on consumer views of their stories) than when producers are motivated to drive consumer action. Intuitively, the reason is that when consumers follow many others, all stories will be widely viewed regardless of quality (resulting in a low incentive to invest for producers "paid" for views) but consumers will ultimately be able to infer story quality from others' sharing behavior (resulting in a high incentive to invest for producers paid for actions). Our analysis thus serves as a jumping-off point for future research that examines the impact of media-market institutions and industrial organization when news stories travel broadly via social connections.[2]

Misinformation. In 1923, the Soviet Union launched the first modern black-propaganda office, with the aim of "manipulating a nation's intelligence system through the injection of credible but misleading data" (Safire (1989)), a tactic Joseph Stalin dubbed "dezinformatsiya (disinformation)" (Manning and Romerstein (2004)). State-sponsored disinformation efforts now abound[3] and are often online.[4] Consumers encounter false news from other sources as well, including individuals and social bots who spread conspiracy theories on social media and in memes.[5] The problem is so severe that the World Economic Forum has listed digital misinformation in online social media as one of the main threats to our society (Howell (2013a,b)). A large and varied literature examines implications of misinformation: how falsehoods and conspiracy theories spread differently than fact-based information on the Internet (del Vicario et al. (2016) and Vosoughi, Roy, and Aral (2018)); how exposure to misinformation can shape memory (Loftus (2005) and Zhu et al. (2010)); and how to identify misinformation and reduce its harmful impact (Qazvinian et al. (2011) and Shao et al. (2016)).

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

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

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

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

We evaluate the impact of misinformation, interpreted as the third-party broadcast of low-quality (inaccurate) stories, on equilibrium outcomes in the information market. A sufficiently large amount of misinformation ruins the market, in the sense that the only equilibrium that remains is a dysfunctional one with no consumer sharing and no high quality output from bona fide producers. However, smaller amounts of misinformation can ironically increase the equilibrium volume of high-quality stories from bona fide producers, since consumers rationally filter what information they share more stringently.

Overall, a key insight of our analysis is that *what consumers are able to infer* from others about story quality and *how much producers invest* in story quality are jointly determined in equilibrium and depend on the nature of bona fide producers' motivations. If producers benefit when consumers take actions based on their stories, then consumers with (infinitely) many connections are able to determine in equilibrium which stories are high quality; there is a "wisdom of the crowd." On the other hand, if producers benefit when consumers view their stories (regardless of what consumers then do) and each consumer follows many others, almost all stories are low quality and consumers learn essentially nothing from others' sharing behavior. The analysis thus uncovers basic forces at work in socially-connected information markets; the Conclusion discusses how the industrial organization of suppliers and other institutional structures overlay these forces.

The paper proceeds as follows. Section 1 presents the basic information market model. Section 2 characterizes equilibrium outcomes when producers benefit from consumer views of their stories, focusing on the equilibrium impact of social connectedness and misinformation injected into the market from third parties. Section 3 then examines an information market in which producers benefit from consumer actions based on their story, focusing especially on whether a "wisdom of the crowd" emerges when each consumer can observe many others' sharing behavior. The Conclusion outlines directions for further research on the organization of information markets.

1 Model: The Market for Information

Here we present a stylized model of socially-connected markets for information. The model specifies (i) suppliers producing information that could be high or low quality, (ii) consumers filtering and spreading information by evaluating information and sharing what they believe is sufficiently likely to be high quality, and (iii) consumer inferences about the information quality from neighbors' sharing decisions.

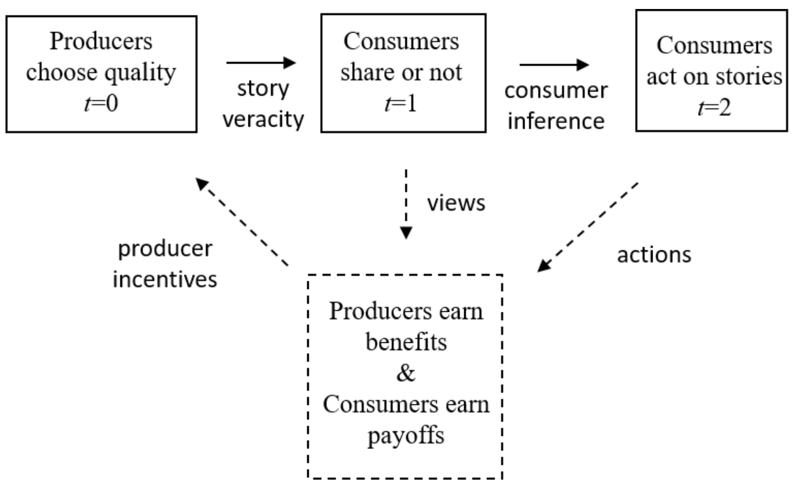

Figure 1: Illustration of the Information-Market Game

The information-market game proceeds in three phases, as shown in Figure 1. At time $t = 0$, each producer decides whether to produce a high- or low-quality piece of decision-relevant information. For shorthand, we will refer to a piece of information as a *story*. A high quality story contains factual claims which are sufficiently likely to be true that consumers want to share the story with their neighbors and to take actions based on it. (This simple assumption about consumers' preferences is sufficient for the study, placing a black box over underlying preferences such as altruism towards friends and families, desire to influence friends' opinions, or desire to be seen as correctly evaluating the quality of stories.) At time $t = 1$, each consumer who has seen a story's broadcast evaluates the story (modeled as a private signal, details below) and decides whether to share it with their neighbors. At time $t = 2$, each consumer who has encountered a story makes an inference about the story's quality and decides whether to take a related action, such as voting in an election or purchasing a product. At the end of the game, producers receive a benefit (e.g., revenue) based on how many consumers have viewed their story or how many consumers have acted on the basis of the story. Consumers also learn the quality of the story and realize their payoffs from having shared and/or acted on the story.

Time $t = 0$: information production. There is a unit-mass continuum of information producers.[6] Each producer has the opportunity to produce a "story," and decides whether to invest in story quality. The cost of producing a low-quality story is zero, while the cost of a high-quality story is c, where $c > 0$ is a privately-observed i.i.d. atomless random variable with c.d.f. $H(\cdot)$ and continuous p.d.f. $h(\cdot)$. We refer to c as the "investigation cost." In the case of news media, c is the cost of investigating and verifying information, say, by interviewing multiple sources for the story; in the case of a product endorsement by an influencer, c is the cost of testing the product and verifying its features.[7] We assume that $h(c) > 0$ for all $c \geq 0$; so, producers' investigation costs

[6] The analysis applies equally to a setting with finitely-many producers or even a single identifiable producer, as long as producers lack commitment power. The model thus applies to individual producers, each interacting with consumers in its own "information market" or "channel." A reputational cost of publishing a low-quality story can be incorporated into the model by adjusting the support of the reporting cost distribution, specified below.

[7] While we do not explicitly model how newspapers, media platforms, and others monitor the stories that are published, some policies can be viewed through the lens of our analysis as changing the parameters of the model. For instance, at a newspaper that rigorously fact-checks stories and fires employees who do sloppy work, reporters will find it costly *not* to invest in story quality. Editorial oversight there-

may be very small or very large. For ease of exposition, we focus on the case in which costs are always positive, i.e., $H(0) = 0$.[8] Let p_0 be the ex ante probability that each producer invests in high quality. For shorthand throughout the paper, we refer to p_0 as the "veracity" of a story.

There are countably-infinitely many consumers, of whom $M \geq 1$ can generate benefits for producers.[9] We consider two consumer sources of producers' benefits. First in Section 2, we suppose that producers benefit when consumers *view* their stories. For example, news-media outlets earn revenue by displaying advertisements next to news stories, and social-media content creators with more views enjoy better search- and recommendation-engine placement for their future content. In Section 3, we consider the case in which producers benefit when consumers *take action* based on their stories. In the news media, for instance, partisan news outlets desire to spur citizens to vote for a political candidate or to protest a policy.[10] Similarly, social-media influencers can earn revenue based on how many people purchase a product that they have recommended (Bradley (2021)).

A producer's expected benefit from publishing a story depends on the story's quality. Let B_H and B_L be the expected benefit that producers earn from high- and low-quality stories, respectively, and let $\Delta B = B_H - B_L$. These expected benefits depend on consumer behavior, which is modeled next.

fore lowers the distribution of investigation cost in the news market consisting of stories published in that newspaper.

[8]Extending the analysis to allow for negative investigation costs is straightforward. The main difference is that average story quality is bounded below by $H(0)$, the likelihood that a producer such as a news reporter or influencer is "intrinsically motivated" to publish high-quality stories.

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

[10]In future work, it would be interesting to extend the model to allow partisan producers to select the subject matter of their stories. In our model, producers only choose whether to invest in story quality.

Time $t = 1$: sharing and viewing stories. Each consumer sees the time-0 broadcast of each story with independent probability $b \in (0, 1)$.[11] Each consumer who sees the broadcast then decides whether to "share" the story after receiving a private signal about its quality (details below).

Consumers are socially linked to others, with a link from consumer i to consumer j indicating that i sees whatever stories j shares, i.e., i "follows" j. We use the word "neighbors" to describe consumers who are linked, with the context indicating the link's direction. For simplicity, we focus on the case in which each consumer follows $d \geq 0$ others and refer to d as "social connectedness."

Consumers are assumed to prefer to share high-quality stories[12] but not low-quality stories.[13] Ultimately, a consumer earns "sharing payoff" $\pi_H^S > 0$ from having shared a high-quality story, $-\pi_L^S < 0$ from having shared a low-quality story, and zero from not having shared a story. Consumers therefore find it optimal to share whenever they believe that a story's likelihood of being high quality (i.e., the story's veracity p_0) exceeds the "sharing threshold" $p^S = \frac{\pi_L^S}{\pi_H^S + \pi_L^S} \in (0, 1)$. For notational simplicity, we normalize $\pi_H^S = \pi_L^S$; so, $p^S = \frac{1}{2}$ and consumers share whenever they believe stories are more likely to be high than low quality.

By assumption, consumers cannot directly observe story quality. However, consumers can have story-specific expertise, personal experience, or access to other information with which to evaluate a story, modeled as an imperfectly informative private signal about the

[11] The case in which all consumers see the broadcast ($b = 1$) is relatively trivial and omitted. An earlier working-paper version (Kranton and McAdams (2020)) extends the analysis to allow consumers to have different likelihoods of seeing the broadcast, b_i, and different numbers of neighbors, d_i, among other asymmetries.

[12] As discussed in the Conclusion, a consumer's overall incentive to share a story could depend on observable characteristics such as novelty, as well as on unobservable characteristics such as quality. Our analysis focuses on the strategic issues created by the presence of an unobservable characteristic, holding observable characteristics fixed.

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

story's quality. In particular, each consumer i receives private signal $s_i \in \{H, L\}$ that matches the true state (High or Low quality) with probability $\rho_i \in (1/2, 1)$, where ρ_i, what we call "signal precision," is itself a random variable with c.d.f. $G(\cdot)$ and continuous density $g(\cdot)$.[14] We refer to signal H as "favorable" and signal L as "unfavorable."

Another interpretation of this private signal is that consumers imperfectly observe a producer's investment. In the case of the news media, for example, a consumer who reads a story does not know for sure whether a reporter has carefully or cursorily investigated sources. But a story in which the sources have been carefully checked is more likely to generate a favorable signal. Similarly, for social-media influencers, a consumer imperfectly discerns, by viewing a post, the number of times an influencer actually tested a recommended product. But a post promoting a product that has been tested more often is more likely to generate a favorable signal.

Let $p_1(s_i, \rho_i; p_0)$ be consumer i's interim belief given private signal s_i, precision ρ_i, and ex ante belief p_0. By Bayes' Rule,

$$\frac{p_1(s_i, \rho_i; p_0)}{1 - p_1(s_i, \rho_i; p_0)} = \frac{p_0}{1 - p_0} \times \frac{\rho_i}{1 - \rho_i} \text{ if } s_i = H \quad (1)$$

$$= \frac{p_0}{1 - p_0} \times \frac{1 - \rho_i}{\rho_i} \text{ if } s_i = L \quad (2)$$

We assume that (a) ρ_i is consumer i's private information, (b) (s_i, ρ_i) is conditionally i.i.d. across consumers and across stories, and (c) $g(\rho) > 0$ for all $\rho \in [1/2, 1]$. Assumption (a) simplifies the presentation but is not essential. The analysis is easily adapted to an alternative setting in which each consumer's expertise is observable to others. Assumption (b) guarantees that signals and sharing behavior about one story

[14] If all private signals have the same precision $\rho \in (1/2, 1)$, then the resulting equilibrium set possesses various "knife-edge" properties that obscure some of the key insights that emerge from our analysis. For example, in such a model, no equilibrium can ever have very high story veracity (greater than ρ) or very low but positive story veracity (between 0 and $1 - \rho$).

are uninformative about other stories, allowing us to consider each story in isolation. Assumption (c) ensures that, whenever some stories are high quality so that veracity is positive, each consumer has a chance of receiving a sufficiently favorable private signal that will lead them to want to share the story.

Time $t = 2$: learning and taking action based on a story. Consumers who have viewed a story decide at time $t = 2$ whether to take a related action, based on their own private signal and the sharing behavior of their neighbors. A consumer ultimately earns "action payoff" $\pi_H^A > 0$ from having acted on a high-quality story, $-\pi_L^A < 0$ from having acted on a low-quality story, and zero from not having acted. Consumers find it optimal to act on a story when its likelihood of being high quality exceeds "action threshold" $p^A = \frac{\pi_L^A}{\pi_H^A + \pi_L^A} \in (0, 1)$. For simplicity, we assume that $\pi_H^A = \pi_L^A$ so that $p^A = \frac{1}{2}$.[15]

The quality of the story is revealed at the end of $t = 2$, at which point producers' benefits and consumers' payoffs are realized.

Equilibrium in the Information Market Our solution concept is Bayesian Nash equilibrium. In any equilibrium, producer investment at $t = 0$ is optimal given consumers' subsequent behavior; consumer sharing at $t = 1$ is optimal given producer investment; and consumer actions at $t = 2$ are optimal given producer investment and consumer sharing. Beliefs are consistent with strategies; equilibrium veracity p_0 is then both the likelihood that producers invest in high quality and consumers' ex ante belief that any given story is high quality.

[15] The analysis easily extends to settings in which consumers have a higher or lower threshold for action than for sharing, but relatively little additional insight emerges.

2 Producers Benefit from Consumer Views

This section considers equilibria in the market for information when producers benefit from consumers viewing their stories. We have three main findings. First, an equilibrium with mostly high-quality stories can only exist when social connectedness is neither too high nor too low (Proposition 1). Second, when social connectedness is very high, almost all stories are low quality in all equilibria (Proposition 2). Finally, whenever most stories would otherwise be high quality, adding a small amount of "misinformation" (defined below) into the market leads to *more* high-quality stories being produced (Proposition 3). However, a sufficiently large amount of misinformation discourages bona fide producers from investing in information quality, leading to a "dysfunctional" market with only low-quality stories (Proposition 4).

We begin by characterizing the conditions under which an equilibrium exists with any given veracity p_0. The argument is organized as follows. First, we derive optimal consumer sharing, depending on consumers' beliefs about story veracity and their own private information. Then we derive the resulting expected benefit from high- and low-quality stories, and determine producers' best reply to consumers' strategies. An equilibrium exists with veracity p_0 if, when consumers believe that fraction p_0 of all stories are high quality, optimal consumer sharing induces producers to optimally invest fraction p_0 of the time.[16]

Optimal consumer sharing. A consumer who sees a story broadcast and receives private signal s_i with precision ρ_i finds it strictly optimal to share the story when their updated belief $p_1(s_i, \rho_i; p_0) > 1/2$ and strictly optimal not to share when $p_1(s_i, \rho_i; p_0) < 1/2$.[17]

[16]Because producers here only benefit from consumer views of stories, we do not need to consider what consumers learn from others' sharing behavior or how they act on that socially-acquired information. Such considerations play a central role in Section 3, when we consider producer benefits generated from consumer actions.

[17]If $p_1(s_i, \rho_i; p_0) = 1/2$, then consumer i is indifferent whether to share. However, because ρ_i is an

Let $\beta_H(p_0)$ and $\beta_L(p_0)$ be the ex ante likelihood that each consumer shares high- and low-quality stories, respectively, given veracity p_0.

If most stories are high-quality ($p_0 \geq 1/2$), consumers find it optimal to share unless they receive an unfavorable private signal ($s_i = L$) that is sufficiently precise ($\rho_i > p_0$) so that their updated belief is below 1/2. Such signals occur with probability $\int_{p_0}^{1}(1-\rho_i)g(\rho_i)\mathrm{d}\rho_i$ for high-quality stories and $\int_{p_0}^{1}\rho_i g(\rho_i)\mathrm{d}\rho_i$ for low-quality stories. Since consumers see the broadcast with probability b, we have that

$$\beta_H(p_0) = b\left(1 - \int_{p_0}^{1}(1-\rho_i)g(\rho_i)\mathrm{d}\rho_i\right) \text{ for all } p_0 \geq 1/2 \tag{3}$$

$$\beta_L(p_0) = b\left(1 - \int_{p_0}^{1}\rho_i g(\rho_i)\mathrm{d}\rho_i\right) \text{ for all } p_0 \geq 1/2. \tag{4}$$

If most stories are low-quality ($p_0 \leq 1/2$), consumers find it optimal to share only if they receive a favorable private signal ($s_i = H$) that is sufficiently precise ($\rho_i > 1 - p_0$) to raise their updated belief above 1/2. This happens with probability $\int_{1-p_0}^{1}\rho_i g(\rho_i)\mathrm{d}\rho_i$ for high-quality stories and $\int_{1-p_0}^{1}(1-\rho_i)g(\rho_i)\mathrm{d}\rho_i$ for low-quality stories; so,

$$\beta_H(p_0) = b\int_{1-p_0}^{1}\rho_i g(\rho_i)\mathrm{d}\rho_i \text{ for all } p_0 \leq 1/2 \tag{5}$$

$$\beta_L(p_0) = b\int_{1-p_0}^{1}(1-\rho_i)g(\rho_i)\mathrm{d}\rho_i \text{ for all } p_0 \leq 1/2. \tag{6}$$

For future reference, note that: $\beta_H(p_0) > \beta_L(p_0)$ for all $p_0 \in (0,1)$; $\beta_H(0) = \beta_L(0) = 0$ and $\beta_H(1) = \beta_L(1) = b$; $\beta'_H(p_0) > \beta'_L(p_0) > 0$ for all $p_0 < 1/2$; and $\beta'_L(p_0) > \beta'_H(p_0) > 0$ for all $p_0 > 1/2$.

Story visibility. Let $V_H(p_0; d)$ and $V_L(p_0; d)$ be the ex ante likelihood that each consumer views high- and low-quality stories, what we refer to as "story visibility," given story

atomless random variable, this occurs with zero probability.

veracity p_0 and social connectedness d. A consumer with d neighbors will view a given story *unless* they miss the broadcast and none of their neighbors share; so,

$$V_H(p_0; d) = 1 - (1 - b)(1 - \beta_H(p_0))^d \tag{7}$$

$$V_L(p_0; d) = 1 - (1 - b)(1 - \beta_L(p_0))^d. \tag{8}$$

Since there are M benefit-generating consumers and the producer is paid one unit of benefit (e.g., revenue) for each of them who views the story, high- and low-quality stories generate expected benefit $B_H(p_0; d) = MV_H(p_0; d)$ and $B_L(p_0; d) = MV_L(p_0; d)$, respectively. The extra expected benefit of high-quality stories is then

$$\Delta B(p_0; d) = M(1 - b)\left((1 - \beta_L(p_0))^d - (1 - \beta_H(p_0))^d\right). \tag{9}$$

Optimal producer investment. Given the sharing behavior that results from consumer belief p_0 and social connectedness d, producers find it optimal to invest when $c \leq \Delta B(p_0; d)$. The supply of high-quality stories (i.e., the probability of investment) is therefore $H(\Delta B(p_0; d))$.

Figure 2 illustrates key qualitative features of high-quality supply, which is the solid curve on the graph with p_0 on the x-axis and $H(\Delta B(p_0; d))$ on the y-axis. Note in particular that, when $p_0 > 1/2$, $H(\Delta B(p_0; d))$ is strictly decreasing and continuous in p_0. In words, as consumers grow even more confident about story quality, producers optimally respond by investing *less* frequently, since the incremental probability that a low-quality story is viewed is higher than the incremental probability that a high-quality story is viewed. Why is that? Differentiating (9) with respect to p_0 yields

$$\frac{\partial \Delta B(p_0; d)}{\partial p_0} = M(1 - b)d\left(\beta'_H(p_0)(1 - \beta_H(p_0))^{d-1} - \beta'_L(p_0)(1 - \beta_L(p_0))^{d-1}\right). \tag{10}$$

Because $\beta_H(p_0) > \beta_L(p_0)$ and $\beta'_H(p_0) < \beta'_L(p_0)$ for all $p_0 > 1/2$, we have $\frac{\partial \Delta B(p_0;d)}{\partial p_0} < 0$. Note also that $\Delta B(1;d) = 0$ for all d since $\beta_H(1) = \beta_L(1) = b$; so, $H(\Delta B(1;d)) = H(0) = 0$.

Equilibrium condition. Since producers optimally invest with probability $H(\Delta B(p_0;d))$, an equilibrium exists with story veracity p_0 if and only if

$$H(\Delta B(p_0;d)) = p_0. \tag{EQM}$$

Graphically, as illustrated in Figure 2, an equilibrium is a crossing-point of $H(\Delta B(p_0;d))$ and the 45-degree line. In Figure 2, two equilibria are shown: one in which most stories are high-quality, denoted by a filled circle, and another in which all stories are low quality ($p_0 = 0$), denoted by an empty circle. The latter possibility, which we refer to as a "dysfunctional informational market," is an equilibrium because producers have no incentive to invest in high quality if consumers never share, and consumers have no incentive to share if all stories are low quality.

In our analysis, we focus on the equilibrium with maximal investment and hence maximal equilibrium veracity. Let p_0^* denote the maximum veracity that can be supported in equilibrium. In what follows, we examine how p_0^* varies with various parameters of the model, focusing especially on the impact of social connectedness d and of misinformation volume m (defined below, with $m = 0$ in the baseline case).

Impact of social connectedness. Increasing social connectedness has a non-monotone effect on the maximal story veracity that can be supported in equilibrium, denoted now as $p_0^*(d)$. In particular, an equilibrium in which most stories are high quality can only exist when social connectedness d lies in an intermediate range.

Proposition 1. *Suppose that producers benefit from consumer views. Social-connectedness*

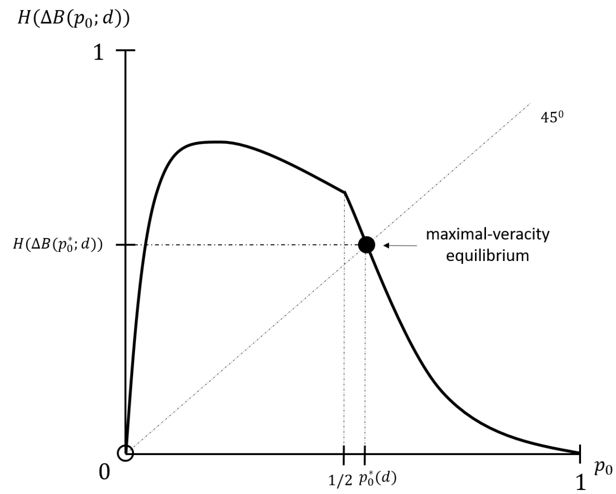

Figure 2: High-Quality Supply and Maximal Equilibrium Veracity $p_0^*(d)$.

thresholds $0 < \underline{d} \leq \overline{d} < \infty$ *exist such that* $p_0^*(d) > 1/2$ *if and only if* $\underline{d} < d < \overline{d}$.

Proof. $p_0^*(d) > 1/2$ if and only if $H(\Delta B(1/2; d)) > 1/2$. To see why, suppose first that $H(\Delta B(1/2; d)) > 1/2$. As shown earlier, $\Delta B(p_0; d)$ is strictly decreasing and continuous in p_0 when $p_0 > 1/2$. Since $H(\Delta B)$ is strictly increasing and continuous in ΔB, this means that $H(\Delta B(p_0; d))$ is strictly decreasing and continuous in p_0 when $p_0 > 1/2$. Moreover, $H(\Delta B(1; d)) = H(0) = 0$ for all d. Thus, there is a unique $p_0 \in (1/2, 1)$ such that $H(\Delta B(p_0; d)) = p_0$ and hence $p_0^*(d) > 1/2$. On the other hand, if $H(\Delta B(1/2; d)) \leq 1/2$, then $H(\Delta B(p_0; d)) < p_0$ for all $p_0 > 1/2$; so, no equilibrium exists with $p_0 > 1/2$ and hence $p_0^*(d) \leq 1/2$.

Next, we show that $\Delta B(p_0; d)$ is a single-peaked function of d for all $0 < p_0 < 1$. To see why, note that differentiating equation (9) with respect to d yields

$$\frac{\partial \Delta B(p_0; d)}{\partial d} = M(1-b)\left(\ln(1-\beta_L(p_0))(1-\beta_L(p_0))^d - \ln(1-\beta_H(p_0))(1-\beta_H(p_0))^d\right)$$

and hence $\frac{\partial \Delta B(p_0;d)}{\partial d} > 0$ if and only if $\frac{\ln(1-\beta_L(p_0))}{\ln(1-\beta_H(p_0))} > \left(\frac{1-\beta_H(p_0)}{1-\beta_L(p_0)}\right)^d$. Since $1 > \beta_H(p_0) > \beta_L(p_0) > 0$ for all $0 < p_0 < 1$, both sides of this inequality are between zero and one, but the left-hand side is constant while the right-hand side decreases exponentially with d. Thus, the inequality holds if and only if d is less than some threshold. Moreover, by inspection of equation (9), $\Delta B(p_0; 0) = 0$ and $\lim_{d\to\infty} \Delta B(p_0; d) = 0$. We conclude that $H\left(\Delta B(1/2; d)\right) > 1/2$ and hence $p_0^*(d) > 1/2$ if and only d lies in a (potentially empty) finite interval not including zero, as desired. □

Our next finding is that, as social connectedness d increases to infinity, the maximal equilibrium veracity $p_0^*(d)$ goes to zero. Thus, nearly all stories are low quality in every equilibrium when d is very large.

Proposition 2. *Suppose that producers benefit from consumer views.* $\lim_{d\to\infty} p_0^*(d) = 0$.

Proof. The proof follows very simply from the fact that $\lim_{d\to\infty} \Delta B(p_0; d) = 0$ for all p_0. Consider any fixed $p_0 > 0$, and note from equation (9) that the equilibrium condition $H(\Delta B(p_0; d)) = p_0$ is equivalent to

$$(1 - \beta_L(p_0))^d - (1 - \beta_H(p_0))^d = \frac{H^{-1}(p_0)}{M(1-b)} \tag{11}$$

As shorthand, let $X(p_0) \equiv \frac{H^{-1}(p_0)}{M(1-b)}$ and define $\hat{d}(p_0)$ implicitly by $(1-\beta_L(p_0))^{\hat{d}(p_0)} = X(p_0)$. For all $d > \hat{d}(p_0)$, the left-hand side of (11) is strictly less than $X(p_0)$; thus, no equilibrium exists with story veracity equal to p_0. Moreover, because consumers are more likely to share stories when veracity is higher, $(1 - \beta_L(p_0'))^d < (1 - \beta_L(p_0))^d$ for all $p_0' > p_0$. Since $X(p_0)$ is increasing in p_0, we conclude that $(1 - \beta_L(p_0'))^d - (1 - \beta_H(p_0'))^d < X(p_0')$ for all $d > \hat{d}(p_0)$ and all $p_0' > p_0$. Thus, $p_0^*(d) < p_0$ for all $d > \hat{d}(p_0)$. Since this argument applies to all $p_0 > 0$, we conclude that $\lim_{d\to\infty} p_0^*(d) = 0$, as desired. □

Impact of misinformation. We now consider how the presence of misinformation impacts the information market (abstracting from the motives of misinformation producers, which we discuss in the Conclusion). Define "misinformation" as low-quality stories injected into the information market from outside sources, which are distinct from the "bona fide producers" modeled above. Let $m \geq 0$ be the volume of misinformation. For any given level of investment by bona fide producers, misinformation reduces story veracity. At this reduced veracity level, consumers share stories less frequently and more judiciously. This change in consumer sharing in turn impacts producers' incentive to invest in quality.

Our previous equilibrium analysis is easily adapted to allow for misinformation. If consumers believe that fraction p_0 of stories are high quality, then their sharing behavior will induce bona fide producers to produce *volume* $H(\Delta B(p_0; d))$ of high-quality stories, as before. Since overall story volume is now $1+m$, the resulting *fraction* of stories that are high-quality is $\frac{H(\Delta B(p_0;d))}{1+m}$. The equilibrium condition therefore becomes $\frac{H(\Delta B(p_0;d))}{1+m} = p_0$ or, equivalently

$$H(\Delta B(p_0; d)) = p_0(1 + m). \tag{EQM-misinfo}$$

Adapting the proof of Proposition 1 in a straightforward way, an equilibrium with veracity $p_0 > 1/2$ exists if and only if $\Delta B(1/2; d)$ is large enough and m is small enough that $\frac{H(\Delta B(1/2;d))}{1+m} > 1/2$. We begin by showing that, when this condition holds, adding slightly more misinformation unambiguously increases the volume of high-quality stories produced.

Dropping "d notation" to simplify equations, let $p_0^*(m)$ denote the maximal equilibrium veracity, viewed now as a function of misinformation volume m. Note that the volume of high-quality stories in the maximal-veracity equilibrium is $(1 + m)p_0^*(m)$.

Proposition 3. *Suppose that producers benefit from consumer views. If $p_0^*(m) > 1/2$, then $\frac{\mathrm{d}\left((1+m)p_0^*(m)\right)}{\mathrm{d}m} > 0$.*

Proof. We show first that $p_0^*(m)$ is strictly decreasing in m. Because $p_0^*(m)$ is the highest story veracity level can be supported in equilibrium, $H(\Delta B(p_0^*(m))) = p_0^*(m)(1+m)$ and $H(\Delta B(p_0)) < p_0(1+m)$ for all $p_0 > p_0^*(m)$. But then, for any $m' > m$, we have $H(\Delta B(p_0)) < p_0(1+m')$ for all $p_0 \geq p_0^*(m)$; so, $p_0^*(m') < p_0^*(m)$.

Note that $\Delta B(p_0)$ is strictly decreasing in p_0 when $p_0 > 1/2$ (shown in the proof of Proposition 1) and that $H(\Delta B)$ is strictly increasing in ΔB. The fact that $p_0^*(m)$ is strictly decreasing in m therefore implies that the volume of high-quality stories, $H(\Delta B(p_0^*(m))) = (1+m)(p_0^*(m)$, is strictly increasing in m when $p_0^*(m) > 1/2$.

Finally, note that $H(\Delta B)$ is differentiable in ΔB (by the assumption that investigation cost c has a continuous p.d.f.) and $\Delta B(p_0)$ is differentiable in p_0 when $p_0 > 1/2$ (shown earlier). We conclude by the Implicit Function Theorem that $p_0^*(m)$ is differentiable in m when $p_0^*(m) > 1/2$, as desired. □

Proposition 3 shows that slightly increasing the amount of misinformation injected into the market can in some cases increase bona fide producer investment. In particular, if social connectedness is in the intermediate range specified in Proposition 1 so that most stories are high quality absent misinformation, a small amount of misinformation spurs bona fide producers to invest more in the maximal-veracity equilibrium.

But what if the volume of misinformation m is large? We have two findings. First, once m exceeds a threshold, the unique equilibrium is the dysfunctional one with zero producer investment in quality (Proposition 4(i)). On the other hand, for any given volume of misinformation, we also show that equilibria with positive investment exist whenever consumers' social connectedness is sufficiently high (Proposition 4(ii)).

Proposition 4. *Suppose that producers benefit from consumer views. (i) For any given social connectedness d, there is a threshold $\overline{m}(d) < \infty$ such that $p_0^*(m; d) = 0$ for all $m > \overline{m}(d)$. (ii) $\lim_{d\to\infty} \overline{m}(d) = \infty$.*

Proof. Define $\overline{m}(d) \equiv \sup_{p_0>0} \frac{H(\Delta B(p_0;d))}{p_0} - 1$. This is well-defined and finite because

(a) $H(\Delta B)$ is continuously differentiable in ΔB, (b) $\Delta B(p_0; d)$ is continuously differentiable in p_0, and (c) $\lim_{p_0 \to 0} \frac{H(\Delta B(p_0;d))}{p_0}$ is finite. To verify this last point, note that $\lim_{p_0 \to 0} \frac{H(\Delta B(p_0;d))}{p_0} = h'(0) \frac{\partial \Delta B(0;d)}{\partial p_0}$; $\frac{\partial \Delta B(0;d)}{\partial p_0} = M(1-b)d\left(\beta'_H(0) - \beta'_L(0)\right)$ by equation (10); and $\beta'_H(0) = bg(1)$ while $\beta'_L(0) = 0$ by equations (3,4). All together, then,

$$\lim_{p_0 \to 0} \frac{H(\Delta B(p_0;d))}{p_0} = dh'(0)Mb(1-b)g(1) \in (0, \infty).$$

By definition of $\overline{m}(d)$, $H(\Delta B(p_0;d)) < p_0(1+m)$ for all $m > \overline{m}(d)$; so, whenever misinformation volume exceeds $\overline{m}(d)$, the only equilibrium is the dysfunctional one with zero producer investment. This completes the proof of (i).

For any given m, define $\overline{d}(m) \equiv \frac{1+m}{h'(0)Mb(1-b)g(1)}$. By definition of $\overline{d}(m)$, we have $\lim_{p_0 \to 0} \frac{H(\Delta B(p_0;d))}{p_0} - 1 > m$ for all $d > \overline{d}(m)$. This implies that $H(\Delta B(\epsilon; d)) > \epsilon(1+m)$ for all sufficiently small $\epsilon > 0$. Since $H(\Delta B(p_0;d))$ is continuous in p_0 and $H(\Delta B(1;d)) = 0 < 1+m$, we conclude that some $p_0 > 0$ exists such that $H(\Delta B(p_0;d)) = p_0(1+m)$; so, $p_0^*(m;d) > 0$ and hence $\overline{m}(d) > m$ for all $d > \overline{d}(m)$. Because this argument applies to all $m > 0$, we conclude that $\lim_{d \to \infty} \underline{m}(d) = \infty$, as desired. This completes the proof of (ii). □

3 Producers Benefit from Consumer Actions

In this section, we study markets where producers benefit when consumers take a specific action based on their story. For instance, a social-media influencer earns money when people click on a link within their post and purchase the endorsed product (Bradley (2021)); a political lobbyist benefits when citizens believe in a message strongly enough to email their elected officials (Bergan (2009)); and partisan news media enjoys the support of political sponsors when stories sway votes (see Kroll (2017) on Sinclair Media and the political motivations of its CEO David Smith).

Our key finding is that a "wisdom of the crowd" emerges in equilibrium as consumers see the sharing decisions of many other consumers. In the $d \to \infty$ limit, consumers can perfectly discern which stories are high quality. Producers' equilibrium incentive to invest in story quality is also as high as possible. These findings contrast sharply with a key finding of Section 2; when producers benefit from views, there is a vanishingly small amount of high-quality investment in the $d \to \infty$ limit.

We begin by characterizing the conditions under which an equilibrium exists with any given story veracity p_0. The argument is organized as follows. First, given optimal consumer sharing (derived in Section 2), we describe how consumers learn from others' sharing behavior and when they choose to act on stories. We then derive the expected benefits of high- and low-quality stories and producers' incentives to invest in quality. An equilibrium exists with veracity p_0 if, when consumers believe that fraction p_0 of all stories are high quality, their sharing behavior leads to consumer inferences and resulting actions that in turn induces producers to invest fraction p_0 of the time.

Consumer learning and story impact. If consumers believe that fraction p_0 of stories are high quality, each consumer will share high- and low-quality stories, respectively, with ex ante likelihood $\beta_H(p_0)$ and $\beta_L(p_0)$, as derived in equations (3-6). Whether any given neighbor shares is therefore an informative (conditionally i.i.d.) binary "social signal." Observing that d_i out of d neighbors have shared, consumer i will update their interim belief $p_1(s_i, \rho_i; p_0)$, which accounts for their private information (s_i, ρ_i) and is given by equation (1), to their ex post belief $p_2(s_i, \rho_i, d_i; p_0, d)$, which also incorporates their socially-acquired information. By Bayes' Rule,

$$\frac{p_2(o_i; p_0, d)}{1 - p_2(o_i; p_0, d)} = \frac{p_1(s_i, \rho_i; p_0)}{1 - p_1(s_i, \rho_i; p_0)} \times \left(\frac{\beta_H(p_0)}{\beta_L(p_0)}\right)^{d_i} \left(\frac{1 - \beta_H(p_0)}{1 - \beta_L(p_0)}\right)^{d-d_i} \tag{12}$$

where here we introduce shorthand $o_i \equiv (s_i, \rho_i, d_i)$ for consumer i's overall "observation."

Let $\mathcal{A}(p_0, d) \equiv \{o_i : p_2(o_i; p_0, d) > 1/2\}$ be the subset of observations that induce consumers to act on the story when they follow d others and believe that fraction p_0 of stories are high quality.

Let $A_H(p_0; d) \equiv \sum_{o_i \in \mathcal{A}(p_0,d)} \Pr(o_i|\omega = high)$ and $A_L(p_0; d) \equiv \sum_{o_i \in \mathcal{A}(p_0,d)} \Pr(o_i|\omega = low)$ be the ex ante likelihood that each consumer acts on a high- and low-quality story, respectively, which we refer to as "story impact." $\Delta A(p_0; d) \equiv A_T(p_0; d) - A_F(p_0; d)$ is the extra impact of high-quality stories.

Producer investment and equilibrium condition. High- and low-quality stories generate expected benefit $MA_H(p_0; d)$ and $MA_L(p_0; d)$ for producers, respectively. (Recall that M is the number of benefit-generating consumers.) A producer with investigation cost c prefers to invest in high quality if $M\Delta A(p_0; d) \geq c$ but not if $M\Delta A(p_0; d) < c$, resulting in ex ante likelihood of investment $p_0 = H(M\Delta A(p_0; d))$. An equilibrium exists with veracity p_0 if and only if

$$H(M\Delta A(p_0; d)) = p_0. \tag{EQM-action}$$

Let $p_0^{*A}(d)$ be the maximum veracity that can be supported in equilibrium given social connectedness d, with superscript "A" denoting that producers now benefit from actions.

Producer investment and the "wisdom of the crowd." Because each consumer's likelihood of acting on a story is between zero and one, a producer's payoff from any story is between zero and M. Consequently, the extra expected benefit of high-quality stories is $M\Delta A(p_0; d) \leq M$. Moreover, this upper bound can never be achieved. To see why, note that $\Delta A(p_0; d) = 1$ is only possible if consumers always act on high-quality stories, i.e., $A_H(p_0; d) = 1$. But then consumers would be acting on stories regardless of their private and socially-acquired information, in which case $A_L(p_0; d) = 1$ as well.

Hence $\Delta A(p_0; d) = 0$, a contradiction.

We conclude that veracity must always be strictly less than $H(M)$ in any equilibrium given any finite social connectedness d. Nonetheless, as we show next, such *maximal-possible investment* can in fact be achieved in the $d \to \infty$ limit.

Proposition 5. *Suppose that producers benefit from consumer actions.* $\lim_{d\to\infty} p_0^{*A}(d) = H(M)$.

Proof. Consider any veracity level $p_0 \in (0, H(M))$, and suppose that consumers believe that fraction p_0 of stories are high quality. Each of consumer i's d neighbors shares high- and low-quality stories with ex ante likelihood $\beta_H(p_0)$ and $\beta_L(p_0)$, respectively, where $\beta_H(p_0) > \beta_L(p_0)$. Observing d neighbors' sharing behavior is therefore equivalent to observing d conditionally i.i.d. zero-one random variables X_j, $j = 1, ..., d$, with $\Pr(X_j = 1|high) = \beta_H(p_0)$ and $\Pr(X_j = 1|low) = \beta_L(p_0)$. By the Central Limit Theorem, each consumer can infer the true state of the world $\omega \in \{low, high\}$ in the $d \to \infty$ limit. Consumers will therefore always act on high-quality stories and never act on low-quality stories in the $d \to \infty$ limit. In particular, $\lim_{d\to\infty} A_H(p_0; d) = 1$ and $\lim_{d\to\infty} A_L(p_0; d) = 0$, and hence $\lim_{d\to\infty} \Delta A(p_0; d) = 1$

For any $p_0 < H(M)$, there exists $d(p_0)$ such that $H(M\Delta A(p_0; d)) > p_0$ for all $d > d(p_0)$. On the other hand, $H(M\Delta A(p_0'; d))$ is less than p_0' for all $p_0' \geq H(M)$ and all d. Since $\Delta A(p_0; d)$ is continuous in p_0 for any given d (straightforward details omitted), we conclude that an equilibrium exists with veracity *between* p_0 and $H(M)$—and hence that $p_0^{*A}(d) \in (p_0, H(M))$ whenever $d > d(p_0)$. Since this argument applies to all $p_0 < H(M)$, we have shown $\lim_{d\to\infty} p_0^{*A}(d) = H(M)$, as desired. $\square$

Proposition 5 and its proof have several important implications. Since consumers are able to discern stories' true quality in the $d \to \infty$ limit, the fact that low-quality stories are present in the information market does relatively little harm so long as consumers are highly connected. There is a "wisdom of the crowd," which enables consumers to avoid acting on low-quality stories, including misinformation.

4 Conclusion

This paper analyzes the market for decision-relevant information. Our model captures key features of contemporary media markets in which (i) consumers share stories over social networks, (ii) information producers cannot commit to publish high-quality stories, and (iii) information producers benefit (e.g., receive revenue) when consumers view their stories or when consumers act on the basis of their stories. Whether producers are motivated by consumer views or actions, consumers' social connectedness impacts equilibrium outcomes.

In an information market where producers benefit from consumer views, the fraction of stories that are high quality in equilibrium is non-monotone in social connectedness. If consumers are initially unconnected, adding social links induces producers to invest more in quality, since high-quality stories are more likely to be shared (and hence viewed) than low-quality stories. However, once consumers follow sufficiently many others, adding still more links reduces producers' incentives to invest, since low-quality stories which pass through consumers' filters are spread and thus viewed more widely. In particular, low-quality stories outnumber high-quality stories in all equilibria when social connectedness is either sufficiently low or sufficiently high. Moreover, high-quality stories completely disappear in the limit as social connectedness d goes to infinity, in the sense that the maximal equilibrium veracity converges to zero.

The equilibrium impact of misinformation on a views-supported information market depends on the quantity of misinformation in question. A sufficiently large volume of low-quality stories injected into the market by third parties can destroy the information market, in the sense that no bona fide providers produce high-quality stories and consumers never share the stories they see. However, smaller volumes of misinformation can have the ironic effect of promoting more high-quality production by bona fide information producers. This effect arises because consumers rationally respond to the presence of misinformation by sharing news more cautiously, which decreases the visibility of low-quality stories more than it decreases the visibility of high-quality stories.

In an information market where producers benefit from consumers' actions based on their stories, high levels of social connectedness are unambiguously beneficial. In the limit as consumers follow infinitely-many others, (i) high-quality story production is as large as it can possibly be, and (ii) consumers are able to perfectly infer story quality, enabling them to act on all high-quality stories and avoid acting on all low-quality stories, what we refer to as a "wisdom of the crowd."

This paper serves as a basis for the study of socially-networked information markets. Several directions for future work could build on our analysis.

A natural next step would be to endogenize social connectedness, allowing consumers to decide how many people to follow, possibly incurring a cost for each social link. Such endogenous link formation would have implications for the efficiency of social-media platforms. In actions-supported information markets, we find that high-quality stories arise in equilibrium if consumers are highly connected. But if stories are very likely to be high quality, consumers have little incentive to invest in social connections for inference purposes, potentially resulting in a sparse network that cannot support high-quality stories. In this context, platforms such as Facebook and Twitter that make it easier for consumers to follow one another (reducing link costs) might indirectly promote higher-quality story production by news journalists and by social influencers.

On the supply side, natural next steps are to consider the industrial organization of information providers and additional business models for information providers. For example, some news-media organizations have recently shifted back toward a subscription-based revenue model (New York Times (2015)). Whoever controls a subscription channel has an incentive to maximize its overall value to consumers, to increase subscribers' willingness to pay for channel access. However, subscriber engagement also drives advertiser and sponsor payoff,[18] and these different revenue sources generate potentially competing incentives, in ways that deserve further study. For instance, a channel that earns revenues only from subscribers might have an incentive to block readers from sharing content outside of its walled garden, while a channel that also earns advertising and/or content-sponsor revenue might prefer to enable stories to be more widely shared by subscribers.

Newspapers and social-media platforms can serve as intermediaries, screening the stories that consumers encounter and/or lending credibility to producers. For instance, a politician with an opinion might post it directly on Twitter or some other online channel such as Medium if the goal is just to grab attention,[19] but submit it to a newspaper such as the *Washington Post* which conducts editorial review if the goal is to change minds.

News-distribution platforms can also create their own individual information markets, by distinctively identifying the stories that consumers discover through their channels. For instance, at Facebook, a curation team consisting of journalists from partner news organizations decides which stories to highlight under the banner of "Today's Stories," creating a distinct information market with material re-published from original sources. Such curated channels could benefit consumers, by highlighting high-quality

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

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

stories by high-quality producers.[20] However, those who view these more-trustworthy stories might also share less judiciously, limiting how much others can learn from their sharing choices. In addition, a *dominant* curated channel might have anti-competitive and/or anti-democratic effects, if those curating the news seek to enhance the market power of existing producers and/or promote an ideological or partisan agenda.

Regarding misinformation, a natural next step would be to consider strategic actors who inject misinformation into the news market. The quantity of misinformation would then be endogenous and depend on the motives of these "misinformation producers," the incentives of bona fide producers, and the social connectedness of consumers. The pattern of equilibrium misinformation production naturally depends on whether misinformation producers benefit from views (like the fake-news site denverguardian.com that earns money from ads shown alongside its false content[21]), benefit when consumers act on their stories (like the Russian-sponsored Heart of Texas website[22]), or in the case of disinformation warfare when consumers are unable to act confidently on *any* story. When consumers are highly connected, our analysis shows that (i) consumers are likely to share but unlikely to act on false stories if bona fide producers are paid for actions, and (ii) consumers are unlikely to share or act on any story if bona fide producers are paid for views. If misinformation producers benefit from views and consumers' social connectedness is high, we would therefore expect more misinformation to be produced when bona fide producers are paid for actions than when they are paid for views, but

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

[21] Most famously, denverguardian.com published a false story linking Hillary Clinton to the death of an FBI agent (Borchers (2016)).

[22] The Russia-based Internet Research Agency created *Heart of Texas*, a fictitious advocacy group that promoted Texas secession from the United States and other provocative positions. When its Facebook page called for a protest against "the Islamification of Texas" in 2017, real people showed up to protest and counter-protest (Allbright (2017)).

that consumers will rarely act on these false stories. On the other hand, if misinformation producers benefit from actions, the fact that consumers are able to discern story quality from others' sharing behavior will deter misinformation producers from participating in the market.

Finally, future work could extend our analysis to allow for multidimensional story quality. Producers in the present paper invest in a single unobservable characteristic (called "quality") but, of course, producers also invest heavily in observable characteristics. Such investments directly affect consumers' incentives to share (e.g., consumers may want to share funny or shocking content, even if they suspect the story is low quality so the information contained is unlikely to be true) and act, while also indirectly affecting those decisions by shaping consumers' beliefs about unobservable characteristics. For instance, suppose that an information producer can invest in the unobservable quality and/or observable appeal of its stories, where "appeal" increases consumers' payoff when sharing a story but has no effect on their action payoff. If the cost of increasing a story's appeal is small relative to the cost of verifying the information contained in the story, producers paid for views may find it optimal to invest only in appeal, leading to equilibrium outcomes in which all stories have little information content but are appealing—widely shared because of their appeal but ineffective at driving action because no one believes the content. By contrast, producers paid for actions may find it optimal to *disinvest* in story appeal as a way of increasing the return to their investments in story quality, as doing so can cause consumers to share their stories more judiciously and thereby activate a "wisdom of the crowd"that drives widespread action on their high-quality content.

References

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

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

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

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

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

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

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

Bergan, Daniel E, "Does grassroots lobbying work? A field experiment measuring the effects of an e-mail lobbying campaign on legislative behavior," *American politics research*, 2009, *37* (2), 327–352.

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

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

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

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

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

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

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

Bradley, Sydney, "How much money Instagram influencers make," Business Insider July 30 2021. available at https://www.businessinsider.com/how-much-money-instagram-influencers-earn-examples-2021-6.

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

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

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

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

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

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

Facebook, "Working to Stop Misinformation and False News," April 6 2017.

_ , "Introducing a Forwarding Limit on Messenger," September 3 2020.

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

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

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

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

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

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

Kranton, Rachel and David McAdams, "Social Networks and the Market for News," Duke U. working paper 2020.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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