

**Misinformation, Persuasion, and News Media
on Social Networks**

by

Chin-Chia Hsu

Submitted to the Institute for Data, Systems, and Society
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Social and Engineering Systems and Statistics
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2022

© Massachusetts Institute of Technology 2022. All rights reserved.

Author
Institute for Data, Systems, and Society
June 30, 2022

Certified by
Ali Jadbabaie
JR East Professor of Civil and Environmental Engineering, MIT
Thesis Supervisor

Certified by
Muhamet Yildiz
Professor of Economics, MIT
Thesis Committee Member

Certified by
Dirk Bergemann
Douglass and Marion Campbell Professor of Economics, Yale
Thesis Committee Member

Certified by
Amir Ajorlou
Research Scientist, MIT
Thesis Committee Member

Accepted by
Fotini Christia
Chair, Social and Engineering Systems Program

Misinformation, Persuasion, and News Media on Social Networks

by

Chin-Chia Hsu

Submitted to the Institute for Data, Systems, and Society
on June 30, 2022, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Social and Engineering Systems and Statistics

Abstract

Social media platforms have become a popular source of news and information for a large segment of society. Users can receive information, share digital content, or attend to online publishers for the latest news. However, the recent proliferation of misinformation has affected people’s perception of the veracity of online information and, in turn, their social behavior. In this environment of real and false information, this dissertation studies two aspects of user behavior through the lens of *persuasion*: (1) sharing online news, and (2) consumer choice of news media.

The first part focuses on the dissemination of online news on social media platforms such as Twitter. I propose two frameworks: in the first I focus on non-strategic agents and in the second one I proceed with a game-theoretic setting.

In the first model, agents choose to share news based on whether it can move their followers’ beliefs closer to their own in the aggregate, and the current size of news spread, without considering news spreading in the future. I describe the dynamics of news spread arising from individual decisions and uncover the mechanisms that lead to a sharing cascade. I elucidate an association between the news precision levels that maximize the probability of a cascade and the wisdom of the crowd.

The second model concerns a binary vote and rational agents who share news to make their followers cast the same vote as they do while strategically speculating on others’ sharing decisions and news spread at the steady state. I characterize the underlying news spread as an endogenous Susceptible-Infected (SI) epidemic process and derive agents’ sharing decisions and the size of the sharing cascade at the equilibrium of the game. I show that lower credibility news can result in a larger cascade than fully credible news provided that the network connectivity surpasses a *connectivity limit*. I further delineate the relationship between cascade size, network connectivity, and news credibility in terms of polarization and diversity in prior beliefs.

The second part of this dissertation investigates how subscribers with diverse prior

beliefs choose between two ideologically opposing news media (intermediaries) that are motivated to influence the public opinion, through their roles of news verification and selective disclosure. The news media may access some news about the state of the world, which may or may not be informative and they can choose whether to verify it. The news media then decide whether to disclose the news, aiming to persuade their subscribers to take the optimal action about the state based on their own beliefs. I show that centrists choose to subscribe to the intermediary with the opposing view, thereby exhibiting anti-homophily. By contrast, extremists exhibit homophily and prefer the intermediary with ideology that aligns with theirs.

This dissertation contributes to the growing literature on people's behavior of news consumption by offering game-theoretic frameworks built on a persuasion motive.

Thesis Supervisor: Ali Jadbabaie

Title: JR East Professor of Civil and Environmental Engineering, MIT

Thesis Committee Member: Muhamet Yildiz

Title: Professor of Economics, MIT

Thesis Committee Member: Dirk Bergemann

Title: Douglass and Marion Campbell Professor of Economics, Yale

Thesis Committee Member: Amir Ajorlou

Title: Research Scientist, MIT

Acknowledgments

This thesis has been tremendously benefited from many people's wisdom and generous support.

First and foremost, I am deeply grateful to my advisor Professor Ali Jadbabaie for his dedicated mentorship and huge support during the vicissitudes of my Ph.D. study. Ali has been passionately offering me his constructive advice and perspectives about my research. Especially, Ali enlightened and guided me through the ups and downs of my academic career. I am fortunate to be nurtured by Ali to become an independent researcher.¹

I would like to thank Dr. Amir Ajorlou, my thesis committee member and my co-author, for his valuable inputs and advice in the past five years. Amir is a guru in mathematics, always attentive to the accuracy and robustness of our approaches. Amir also provided me with many helpful suggestions for my research career. I learned greatly from Amir and I am indebted to him for his continuous support.

I am also grateful to my thesis committee members, Professor Muhamet Yildiz (my co-author) and Professor Dirk Bergemann. I first met Muhamet in his course on game theory. I was fortunate to receive helpful comments from Muhamet for my first project and later collaborate with him on an exciting and challenging topic. Muhamet is adept at illustrating his ideas and thinking thoroughly without leaving any ambiguity that may induce confusion. Writing a paper with Muhamet is always an enlightening moment to me. Having Dirk on my committee is a serendipity, which happened during his visit to MIT in my last year of Ph.D. study. Dirk is sharp; he quickly understood the gist of my research while I was explaining it on a blackboard within an hour. Dirk offered many constructive suggestions from his perspective as an economist, which are very beneficial to me.

I appreciate the support from Institute for Data, Systems, and Society (IDSS). I

¹This thesis was supported by ARO MURI W911NF-18-S0003, W911NF-19-1-0217 and a Vannevar Bush Fellowship from the Office of Secretary of Defense.

would like to thank Professors Munther Dahleh, John Tsitsiklis, Ali Jadbabaie, and Fotini Christia for their great effort to create and foster the interdisciplinary Ph.D. program that allow us to explore unknowns not restricted to traditional disciplinary boundaries in academia. I also thank Professors Devavrat Shah and Sasha Rakhlin to lead the statistic center at IDSS. I am deeply indebted to many IDSS staff. I want to pay special tribute to Ms. Elizabeth Miles for her warmheartedness, passion and generous support to IDSS students, which make us, especially international students like me, feel at home. I thank Ms. Laura Dorson and Ms. Kim Strampel for their administrative assistance to make IDSS thrive.

I want to show my thankfulness to all of my friends with whom I shared many valuable moments and memories in the past six years: my SES and lab family, Qi & Jinglong, Yan, Yi & Lei, Yuan, Hanwei, Minghao, Ian, Eaman, Amir, Paolo; my cherished friends in Boston, Yifu, Sifan, Shengdi, Chuliang & Haoran, Jingfan & Joy, Junang, Intae, Iksung, Fangji, Kevin, Lucas, Landon, Patrick; my college cohort of NTU EE 15'; specially, my funny Taiwanese friends, Janet, Pin-Yi & James, Liang-Hsun & Ting-An, Tzyy-Shyang & Victoria, Wei, Po-Han, Jia-Ying, Ang-Yu & Ya-Lien, Wei-Tung, Hung-Hsun, Chung-Yueh, Alex, Mason, Szu-Yu & Pei-Yu; and many others. Thank you all for adding so many flavors to my Ph.D. life.

Finally, this thesis is dedicated to my beloved Mom and Dad. Without their unconditional love I would not have been healthy, happy, and able to pursue my Ph.D. at MIT.

Contents

1	Introduction	15
2	News Sharing of Myopic Individuals on Social Networks	21
2.1	Background	21
2.2	Model	28
2.3	Agent's Decision of Sharing the News	34
2.3.1	Decision Rule	34
2.3.2	Broadcast Size	38
2.4	Prevalence of News over the Network	39
2.4.1	Spread Dynamics	41
2.4.2	Effect of the Network Structure	45
2.5	News Precision Levels Maximizing Cascade Probability	48
2.5.1	In a Population with Highly-Diverse Perspectives	51
2.5.2	When the Truth is Surprising to a Population with Homogeneous or Less Diverse Perspectives	51
2.5.3	In a Wise Population with Homogeneous Perspectives	52
2.6	Discussion	53
3	A Game-theoretic Model of News Sharing on Social Networks	55
3.1	Background	55
3.2	Model	61

3.3	Equilibrium Analysis	67
3.3.1	News spread emerging from news-sharing decisions	67
3.3.2	News-sharing decisions in light of news spread	68
3.3.3	Equilibrium structure	72
3.3.4	Effects of news magnitude and connectivity on news spread size	75
3.4	News Cascade, Scale of Spread, and Credibility	77
3.5	News Cascade in Populations with Polarized Perspectives	81
3.6	Discussion	85
4	Subscription Networks, Verification, and Media Bias	87
4.1	Background	87
4.2	Model	93
4.3	Equilibrium Structure	100
4.3.1	Belief updating and subscribers' action strategies	101
4.3.2	Decision making for disclosure and verification	105
4.4	Equilibria without Verification Option	111
4.4.1	Variability in players' reactions to news with uncertain veracity	111
4.4.2	Subscription choices of centrists and extremists	113
4.5	Equilibria with Verification Option	117
4.6	Comparative Statics on Subscription Choices	122
4.6.1	Intermediaries' perspectives and the set of extremists	122
4.6.2	Strict homophily and indifference among the extremist subscribers	123
4.7	Extensions	127
4.7.1	Disclosure strategies for verified news	127
4.7.2	Verification result is perfectly observable to subscribers	129
4.7.3	Utility function and action space	135
4.8	Discussion	135
5	Conclusion and Future Work	137

A	Supplementary Material for Chapter 2	141
A.1	Proofs of Section 2.3	142
A.2	Proofs of Section 2.4	143
A.3	Proof of Section 2.5	150
B	Supplementary Material for Chapter 3	157
B.1	Proofs of Section 3.3	157
B.2	Proofs of Section 3.4	161
B.3	Proofs of Section 3.5	163
C	Supplementary Material for Chapter 4	165
C.1	Proofs of Section 4.3	165
C.2	Proofs of Section 4.4	173
C.3	Proofs of Section 4.5	175
C.4	Proofs of Section 4.7	182
C.5	Additional Results and Discussion	185

List of Figures

2-1	Timeline of our model.	31
2-2	Illustration of the network generation process by configuration model from a sequence of (in-degree, out-degree) pairs.	41
2-3	The mean degree of the line graphs for Zipf and Poisson degree distributions with common expected degrees $\mu_Z = \mu_P = 1$ and correlation coefficient $\rho_Z = \rho_P \in (0, 0.5)$	47
2-4	Illustration of the three cases in Theorem 2.2 for a sample case with $\bar{\mu} = 0, \sigma_\theta = 2, (\theta - \bar{\mu})^2 = 2, C = 4$, and $\mu^{\text{Line}}(\mathcal{G}) = \{2.2, 6350.3\}$. (Left) Cascade likelihood versus the credibility β and the diversity in perspectives σ_μ^2 . (Right) The set of optimal credibility Ω^* as the correspondence of the diversity in perspectives σ_μ^2	50
2-5	Optimal credibility levels $\Omega^*(\sigma_\mu^2)$ for $\sigma_\mu^2 \in [0, 4]$ for a sample case with $C = 4, \sigma_\theta = 2, \mu^{\text{Line}}(\mathcal{G}) = 2.2$, and three levels of collective wisdom $(\theta - \bar{\mu})^2 \in \{0, 2, 6\}$ illustrating cases (ii)-(iii) in Theorem 2.2.	53
3-1	The network connectivity limit $\lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G)$ as a function of polarization $\bar{\mu}$ and in-group diversity σ_μ when $C = 2, x = 3$, and G is the CDF for standard normal distribution. The colormap indicates the value of connectivity limit.	84
4-1	The subscription choices of the centrist (green) and extremist (red) subscribers.	91

4-2 Timeline of the subscription-verification-disclosure game described in Section 4.2 95

4-3 The posterior mean as a function of news x when news veracity is uncertain and the signal structures are Gaussian variables: $\sigma_\theta = \sigma_\epsilon = 1.5$ ($\sigma_T^2 = 4.5$) and $q = 0.2, \mu_i = 2$ 112

4-4 Verification and disclosure strategy of leader R and the subscription choices in the true-news-disclosing equilibria as described in Proposition 4.3. 119

4-5 Illustration of intermediary R 's decision-making of verifying left-leaning news that satisfies $\mathbb{E}_R[\theta|x, T] < 0$ in Example 4.2. 125

4-6 The range of strict homophily and indifference for the right-leaning extremists' subscription choices varies with the verification cost in Example 4.2. 127

List of Tables

2.1	Network statistics for the follower graph of Twitter (Data source: Kwak et al. 2010).	49
-----	---	----

Chapter 1

Introduction

Social networks have become a major source of news and information for a large segment of the population across the political spectrum. In particular, social media giants such as Facebook and Twitter have amassed a large number of users who primarily consume their news via these platforms. Furthermore, publishers increasingly use social media to inform the public of unknown facts or the latest news (Shearer and Matsa 2018). Politicians also launch their political propaganda on social media more often than via conventional channels to promulgate their ideology and mobilize voters (Williams and Gulati 2013). A myriad of various information is circulating on online social networks.

This thesis focuses on three elements underlying users' interaction on social networks in this informational environment: misinformation, persuasion, and news media.

Misinformation

Social media have also become a conduit for spread of inaccurate news and misinformation. Many popular and widespread posts or memes on social media continue to be inaccurate, misleading, or false (Gillin 2017, Silverman et al. 2017, Silverman and Pham 2018). This is due to the fact that any user on social media can readily publicize

digital content and share others' posts, often without verification, fact-checking, or significant third party filtering.

A flurry of recent empirical findings (Takayasu et al. 2015, Vosoughi et al. 2018) suggest that false news and misinformation are more likely to disseminate on social networks, spreading faster, deeper, and more widely than true information. More recent studies show that difference between false-news and true news dissemination may be substantially explained by the difference in peer-to-peer transmission rate of news (Juul and Ugander 2021). Even though recent literature across multiple disciplines has outlined the impact of prevalent misinformation on public opinions, especially on elections (Bessi and Ferrara 2016, Allcott and Gentzkow 2017), little progress has been made on why and how inaccurate information can trigger a large sharing cascade among the users on social media.

Persuasion and strategic communication

At the core of this thesis lies the motivation of individuals behind news sharing and information disclosure. A potential reason is that one wants to *persuade* her peers, moving their beliefs closer to her own belief (Hovland et al. 1953, Che and Kartik 2009, Berger 2014). In particular, in collective decision-making (e.g. presidential election or referendum) where each individual's utility is a function of others' actions, the persuasion motive arises since one's belief may differ from her peers' beliefs and the optimal actions can be subjectively dissimilar. When receiving some new information, one can influence her peers' beliefs and in turn their actions by sharing the information and altering their information sets. Likewise, one may choose not to share the information if it otherwise persuades her peers to take undesired actions from her view (e.g. vote for the opposing candidate). The persuasion motive therefore leads to strategic communication among individuals, related to disclosure games, (Milgrom 1981, Grossman 1981), cheap talk (Crawford and Sobel 1982), and persuasion games (Milgrom and

Roberts 1986). This thesis will focus on the scenario in which agents can not distort the received news nor send other messages; they only choose whether or not to share the news they hear (e.g., retweeting on Twitter).

News media and their roles

Unlike creative digital contents such as memes and posts, certain types of information or news is not easily accessible by the general public, and decision makers often rely on news intermediaries, such as news media or print publishers, are relied on for such information. In contrast to a limited number of trusted news agencies available a few decades ago, the advance of digital technology and popularity of social media have boosted online news business, fostering a wide range of information sources for news consumers (Newman et al. 2019). Despite numerous intermediaries, news consumers selectively attend to a few of them owing to their limited quantity of mental effort during a short time span (Kahneman 1973).

News media play the vital role of fact-checking the information before any news release, based on their journalists' expertise and web of sources. The importance of fact-checking information can not be emphasized enough: Nowadays digital social platforms have become a hotbed for digital misinformation to spread wide and fast (Del Vicario et al. 2016, Vosoughi et al. 2018). Without fact-checking or significant third-party filtering beforehand, the inundation of misinformation or false news poses threats to our society, inciting political polarization and social unrest (Howell 2013).

However, news intermediaries often have their own ideological biases and are motivated to influence the public opinion (Gentzkow et al. 2015, Puglisi and Snyder Jr 2015). For example, they can choose to promote news that can persuade their subscribers towards their own political agenda or conceal the unfavorable information that otherwise moves their readers' beliefs away. With diverging information, in a fragmented society the ideological divide among the population is being entrenched.

Thesis outline

My thesis contributes to the theoretical literature on individual behavior of sharing online news and choosing news media for information in an environment with real and false news. I study the two types of behavior through the lens of persuasion and strategic communication.

In Chapter 2, I propose a framework to study agents' news-sharing decisions. Specifically, the agents receiving some news choose whether to share it based on the level to which the news can move their followers' beliefs closer to their own in aggregate. In particular, the agents myopically make their decision in the sense that they do not foresee the news spreading in the future as a result of others' decisions. I characterize the dynamics of the news spread arising from the agents' decisions and uncover the mechanisms that lead to a sharing cascade. I further identify a connection between news the precision levels that maximize the probability of a cascade and the prior wisdom of the crowd. This chapter is based on Hsu et al. (2020a).

In Chapter 3, on the other hand, I present a game-theoretic analysis of individual decision making of sharing news on social networks, based on Hsu et al. (2022). There is a binary vote and agents share news so as to make their followers cast the same vote as they do while strategically pondering on others' sharing decisions and news spread in the future. The highlight result shows that lower credibility news can trigger a larger cascade than credible news when the social network is highly connected. Moreover, I demonstrate that sharp polarization in prior beliefs in the population prompts more sharing of lower credibility news, leading to larger cascade size.

In Chapter 4, I provide a theory of how people select news intermediaries as information sources, based on Hsu et al. (2020b). The news media, caring for the welfare of all subscribers, have a motive to persuade the subscribers to take the optimal action about the state based on their own beliefs. When there is some news, a leader can decide whether to fact-check the veracity of the news and then choose whether to disclose the news to her own subscribers or not. Our main results give insights into how

individuals would seek information when information is private or costly to obtain, while considering the persuasion motive of strategic news providers who are partisan.

Finally, in Chapter 5 I conclude and propose a few future directions that I would like to explore in the interaction of platforms and artificial intelligence.

Chapter 2

News Sharing of Myopic Individuals on Social Networks

2.1 Background

Popularity, ease of access, as well as wide and fast dissemination of digital content are some of the major reasons why many users share news and information on social media. News publishers use social media to instantly inform their online readers of the latest news (Shearer and Matsu 2018). Candidates running for political office launch their political campaigns on social media to advertise their ideology and mobilize voters (Williams and Gulati 2013). Social media platforms such as Facebook and Twitter are now a popular instrument for political activists to organize the public by sharing information, drawing awareness, and coordinating their activities (Steinert-Threlkeld et al. 2015, Enikolopov et al. 2015, Hendel et al. 2017).

Social media have also been weaponized as an effective tool in information warfare: online users often share contents without verification; a vulnerability that is increasingly being exploited by malicious actors to manipulate public opinion by spreading false news and misinformation on social media (Bessi and Ferrara 2016, Allcott and Gentzkow 2017, Weedon et al. 2017, Shao et al. 2018). Many fast-spreading social

media posts or memes turn out to be false or inaccurate (Gillin 2017, Silverman et al. 2017, Silverman and Pham 2018).

The observed prevalence of many false posts or tweets seem to suggest that false news and misinformation are more likely to circulate on social networks, spreading faster, deeper, and more widely than accurate information. Recently, Vosoughi et al. (2018) empirically corroborated this idea using data from Twitter: they find that false and inaccurate information travels wider, faster, and deeper, and hypothesize novelty as a potential cause for the prevalence of misinformation.

In this chapter, we aim to complement such empirical findings by providing a theory that rationalizes the individuals' decision making process that leads to news sharing and by characterizing conditions under which these decisions can potentially trigger a wide spread of false or inaccurate information.

Why do people tend to spread the news that they find to be more novel? One potential rationale lies in our *desire for like-minded others* (Bahns et al. 2017, Jost et al. 2018): we would like to persuade others and assimilate their opinions to our own. After all, new information can change opinions; if an agent finds a piece of news so novel and surprising that significantly affects her opinion, the persuasion motive creates an incentive to modify her peers' information sets by sharing the news with them, and as a result, change their opinions as well. Sharing news can also be a result of an affirmation motive: When having different/opposing opinions from friends, one may share with them the news that affirms his/her viewpoint. Indeed, Del Vicario et al. (2016) find empirical evidence that Facebook users involved in sharing cascades are highly homogeneous in their polarization. As we will show in this work, our model is capable of rationalizing both.

To this end, we study the news-sharing decision process of individuals on a social network. News is a noisy observation of an underlying state of the world, initially shared with a limited number of (randomly chosen) agents. Agents are endowed with independent heterogeneous priors on the state. Whenever an agent receives the news,

she updates her belief on the state via Bayes' rule and then chooses whether to broadcast a copy of the news to her followers (e.g., retweet the news): She makes her decision so as to maximize her expected utility, deliberating the effect of her action on her followers' beliefs. Each agent's utility is a function of the average distance between her belief and the beliefs of her followers, capturing her desire for like-minded others. We assume that there is a cost incurred with broadcasting, which can represent people's inherent tendency to refrain from taking an action (Jeuland 1979, Su 2009).

Moreover, we consider non-strategic agents in the sense that they *myopically* make their sharing decisions depending on information at the time of receiving the news, without projecting the news spread in the future: While assessing the marginal gain in sharing the news with their followers, the agents only weigh in the possibility that a follower may have already received the news from other agents in the past, and particularly, do not speculate on the likelihood that an uninformed follower may hear the news from others at some point in the future. In Chapter 3 we will otherwise focus on strategic agents who decide whether to share news while accounting for the effect of others' news-sharing decisions on news spread in the future.

We emphasize that, instead of a dichotomy between true and false news, this model concerns news that conveys noisy information about an underlying state of the world. This corresponds to the scenarios in which the state has not been uncovered nor identified at the time of publishing articles. The extent to which agents take into account the news vary with the news agencies' investment in estimating the state, or relatedly, their credibility in news reporting. We focus on news with credibility that is known to the population, corresponding to the old news media who possess long history and well-established reputation in the news market.

The network that depicts social interactions within the population is exogenously constructed using a simple random graph model that is similar to the frameworks adopted by Galeotti and Goyal (2009), Fainmesser and Galeotti (2016, 2020): Each agent only knows her own set of followers and followees but has no information about

the interactions among other agents in the network, except for the joint degree distribution.¹

This notion that agents’ knowledge of the network are limited to their neighbors is also empirically corroborated by Breza et al. (2018). We use a random graph model to construct the followers’ network from a sequence of pairs of in-degrees and out-degrees that are generated according to a joint degree distribution. Indeed, random networks can serve as a tractable framework for modeling the structure of social interactions (e.g., the following relations) within a population. They are extensively used as a convenient modeling abstraction to facilitate the modeling and analysis of the diffusion processes on networks (Rapoport 1957, Erdős and Rényi 1959, Newman 2003, Jackson 2008).

As our first contribution, we show how *surprise* in the news and *affirming* one’s own perspective naturally arise from the utility-maximizing decision of agents on broadcasting. In particular, we show that the utility resulting from a broadcast action, exclusive of the cost, can be decomposed into three parts: (i) A part favoring a broadcast when the news *affirms* the view of the agent against the average prior perspective in the population; (ii) A quadratic gain incentivizing broadcast for *surprising news* (with surprise defined as deviation of news from the average prior perspectives, discounted by news credibility); (iii) A part independent from the news and agents’ perspectives, solely resulting from reducing the diversity in followers’ perspectives.

As our second result, we then characterize the dynamics of the news spread that is endogenously generated by micro-level news-sharing decisions, and establish necessary and sufficient conditions for emergence of a cascade. The cascade condition simply suggests that a news item goes viral if and only if, each follower who receive the news early on in the process will pass the news to more than one new agent on average. The condition depends on both the network structure and the likelihood of sharing

¹We abstract the network formation using random network modeling, unlike the set of models (Sethi and Yildiz 2016, Hsu et al. 2020b) that endogenized the social network as a result of individual strategic behavior of following or connecting to others in order to garner relevant information.

for a newly-informed follower. We show that the cascade likelihood over a network is increasing with the average degree of its corresponding *line graph*, a dual graph where the nodes are the edges of the original graph and two nodes form an edge if the two corresponding edges in the original graph are incident. This enables us to compare the effect of different network structures on the spread of news by examining expected degrees of their corresponding line graphs. For network structures generated from a joint Poisson degree distribution and a Zipf (power-law) degree distribution with a common expected degree and correlation coefficient, we show that cascades are more likely to emerge on networks generated from the Zipf distribution.

Finally, we offer theoretical insights into how the fundamental motives underlying individual news-sharing decisions would determine the news precision levels that maximize the ex-ante likelihood of a cascade. We elucidate the connections between optimal news precision levels and *the diversity of perspectives* in the population and *the wisdom of the crowd* (Galton 1907, Surowiecki 2005), particularly for well-connected social networks whose line graphs have a mean degree of at least two. This assumption ensures all the results to still hold for a general class of utility functions elaborated on in Section 2.2. Indeed, it is not a restrictive assumption given the apparent well-connectedness of the existing social media platforms. For the follower network on Twitter (as the motivating application of our work), we analyze a dataset which was collected and kindly shared by Kwak et al. (2010): the dataset comprises over 41.7 million nodes and 1.47 billion links, with a value of 6350.3 for the average degree of its corresponding line graph.

We show that in a well-connected social network, the diversity of perspectives facilitates the spread of news at all precision levels. In a highly-diverse population, a wide range of precision levels result in viral news, including but not limited to the truth; the range of such precision levels shrinks as the population becomes less diverse. On the other hand, in a population with moderately-diverse or homogeneous perspectives, the truth has to be surprising enough to the public so as to trigger a cascade; in a

population with individual perspectives concentrated around the truth, however, the truth has no chance of going viral and the probability of a cascade is maximized at some nonzero precision level.

Our work joins the recent line of theoretical research on spread of misinformation. Nguyen et al. (2012), Budak et al. (2011), Törnberg (2018) computationally investigated misinformation diffusion using mechanic models with exogenous rules of relaying information to others. This work is also related to the theories on information diffusion (Bala and Goyal 1998, Watts 2002, Jackson 2008, Acemoglu et al. 2010, Sadler 2020).

Our work also connects to the literature on controlling the spread of information on networks via word-of-mouth. Ajorlou et al. (2018) characterize the optimal dynamic pricing policy in a setting where the information about a new product spreads only via word of mouth among buyers, and hence the seller should use the price both to control the profit and to spread the information in the network. The engagement of buyers in the spread, however, is determined by their purchase decisions based on their valuations of the product, independent from the network. Campbell et al. (2017) suggest self-promotion as an expert as the incentive to engage in word-of-mouth spread of information and study how a firm could maneuver this process by strategically limiting the release of information and advertising only to possibly high types. Although this model can explain the diminishing desire of agents to engage in the spread, it is unable to explain behaviors related to surprise or accuracy of information since no uncertainty, or belief, is present in their model.

Our work is further related to the recent paper by Candogan and Drakopoulos (2020), where the authors focus on optimal signaling and mechanism design for an online platform, aiming to balance the engagement and misinformation on social networks. Unlike our model in which agents engage in the news spread by broadcasting the news, in Candogan and Drakopoulos (2020) the platform directly provides a personalized recommendation to each individual as to whether engage with the content or not, without the element of news sharing.

Individuals in our model acknowledge that the news is generated among the noisy process of identifying the state and incorporate news value into their beliefs using the known news credibility on a continuous scale. In this regard, we share the similar approach to investigating the circulating information on platforms with Allon et al. (2021); they studied the evolution of beliefs of consumers who selectively consume only few posts each time when the platforms present a menu of posts that are informative about the true state.

Our choice of the persuasion motive for an agent to influence her followers' opinions via sharing the news is related to the literature on disclosure games (Milgrom 1981, Grossman 1981); cheap talk (Crawford and Sobel 1982); and persuasion games as coined by Milgrom and Roberts (1986). In our setting, agents can only choose whether or not to disclose the news they receive (e.g., retweeting on Twitter) without distorting or coarsening the news. In particular, our model shares the similar utility formulation with the model of Che and Kartik (2009), which concerns one leader and one follower and the leader would like to move this followers belief close to hers via strategic disclosure of information.

The effect of news precision on agents' broadcasting decisions and consequently emergence of a news-sharing cascade also has connections to information design literature (Rayo and Segal 2010) and Bayesian persuasion (as coined in Kamenica and Gentzkow 2011) literature, where the focus is on the problem of persuading a rational agent to take a desired action by controlling her informational environment in a symmetric information setting. The key distinction between Bayesian persuasion and the classical literature on signaling games (e.g., Crawford and Sobel 1982) is that once the signal is realized it cannot be manipulated or distorted as is the case in this paper. In the Bayesian persuasion literature, however, the main focus is on controlling the informativeness of the public information and usually no element of network processes or dynamic information disclosure, as involved in our model, is present.

The rest of the chapter is organized as follows. In Section 2.2, we describe our

model and discuss its assumptions. In Section 2.3 we characterize individual news-sharing decision rules, where we show how affirmation and surprise motives arise from utility-maximizing broadcast actions. The endogenous dynamics of news spread in networks as well as necessary and sufficient conditions for a news cascade to emerge are derived and discussed in Section 2.4. In Section 2.5, we formulate the likelihood of a news cascade as a function of the precision level and demonstrate the connections of the optimal noise levels to the diversity in individual perspectives and the collective wisdom in a population. Section 2.6 discusses our findings. The formal proofs of all the lemmas and theorems are collected in the Appendix A.

2.2 Model

The main object of study in this paper is a social network consisting of a large population of agents \mathcal{V} with size $|\mathcal{V}| = n$ (our particular case of interest is when $n \rightarrow +\infty$). Agents are indexed by $i \in \{1, 2, \dots, n\}$ and form a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. We say that agent j *follows* agent i if and only if $(j, i) \in \mathcal{E}$. We denote the set of agent i 's followers by $N^{\text{in}}(i) \triangleq \{j \in \mathcal{V} | (j, i) \in \mathcal{E}\}$. We assume that the set of followers and followees of agent i is her private information. For an edge $e = (j, i)$, we may alternatively refer to j as its tail and i as its head, which we denote with $\text{Head}(e)$ and $\text{Tail}(e)$, respectively. We will work with an abstract model of a network characterized by its joint distribution of in-degrees and out-degrees. The joint distribution of in-degrees and out-degrees is common knowledge while the graph \mathcal{G} itself is not. We postpone the details of the network to Section 2.4.

There is an unobservable state of the world $\theta \in \mathbb{R}$. Following standard Bayesian models, agents agree to disagree (Aumann 1976); they have heterogeneous prior beliefs on the distribution from which the state is drawn. We assume that agent i has a normal prior belief with mean μ_i and variance σ_θ^2 :

$$\theta \sim_i \mathcal{N}(\mu_i, \sigma_\theta^2).$$

Following Sethi and Yildiz (2016), we refer to the mean of agent i 's belief as her *perspective* and to its variance σ_θ^2 as the *uncertainty* level of her belief.² The prior perspective μ_i is agent i 's private information; individual prior perspectives, however, are assumed to be independently and identically distributed according to

$$\mu_i \sim \mathcal{N}(\bar{\mu}, \sigma_\mu^2),$$

where $\bar{\mu}$ represents the *average or aggregate prior perspective* and σ_μ^2 quantifies the variance of the individual perspectives on the state in the population. We refer to σ_μ^2 as the diversity of perspectives, and to $\mu_i - \bar{\mu}$ as the *ex-ante bias* of agent i . Furthermore, we measure the collective wisdom, or *wisdom of the crowd* (Galton 1907, Surowiecki 2005), by the squared difference between the state and the average perspective in the population $(\theta - \bar{\mu})^2$; the smaller this deviation, the wiser the population.

At the onset, the unobservable state θ is realized. A sender has access to a noisy observation of θ to which we refer as news x . The noise is assumed to be additive, independent from the state, and normally distributed. We can hence view news x as a realization of a random variable X , where

$$\begin{aligned} X &= \theta + \epsilon, \\ \epsilon &\sim \mathcal{N}(0, \sigma_\epsilon^2), \end{aligned}$$

where σ_ϵ is the noise level in the news.³ The sender then releases the news x to a (randomly and uniformly chosen) subset of the agents.

An agent who received the news updates her belief on the state via Bayes' rule.⁴

²We assume the same precision for prior beliefs, reflecting that agents are *ex-ante* equally informed about the state.

³The precision can reflect the credibility of the source, or can be directly inferred from the news (e.g., from the confidence level when news is a report on the result of a polling survey). The distribution of the noise and its precision are common knowledge among agents who received the news.

⁴Only a single piece of news x is transmitted. Since the source is included with the news (e.g., a link to a certain post from a news agency), agents can distinguish copies of the same news, thus updating their beliefs only the first time they see the news.

In particular, the posterior belief of agent i who has received a copy of the news x is

$$\theta|x \sim_i \mathcal{N}\left((1 - \beta)\mu_i + \beta x, (1 - \beta)\sigma_\theta^2\right), \quad (2.1)$$

where $\beta \triangleq \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2} \in [0, 1]$. We refer to β as the *credibility* of the news, as it determines how much weight agents place on the news when updating their perspectives on the state. Specifically, $\beta = 1$ corresponds to news with infinite precision, i.e., the case in which the news reveals the true state; by contrast, $\beta = 0$ represents the case in which the news is completely uninformative. After receiving the news and updating her belief, agent i has to immediately decide⁵ whether to broadcast a copy of the news to her followers, referred to as her broadcast action $a_i^B \in \{0, 1\}$ (this can correspond to tweeting).

Given the desire to have like-minded followers, an agent receiving the news broadcasts it to her followers only if it sufficiently reduces the difference between her belief and her followers' beliefs. To quantify the difference between two normally distributed beliefs, $P_1 = \mathcal{N}(\mu_1, \sigma_1^2)$ and $P_2 = \mathcal{N}(\mu_2, \sigma_2^2)$, let us define a distance measure between P_1 and P_2 as⁶

$$W_2^2(P_1; P_2) = (\mu_1 - \mu_2)^2. \quad (2.3)$$

We then consider the following utility function for agent i taking broadcast action

⁵Agents typically do not spend much time and attention deliberating whether to share news now or postpone their decision to a future time.

⁶All the results presented in the paper still hold for a more general distance of the form

$$W_2^2(P_1; P_2) = (\mu_1 - \mu_2)^2 + \gamma(\sigma_1 - \sigma_2)^2, \quad (2.2)$$

measuring the imbalance of information besides disagreement in perspectives, where $\gamma \geq 0$ controls the relative importance of the two terms. Since a normally distributed belief is fully characterized by its mean and standard deviation, this will fully capture distance between agents' beliefs. For the sake of clarity, however, we present the results for the case $\gamma = 0$, and prove the validity of the results for the general case in the Appendix.

times $t \in \{0\} \cup \mathbb{N}$ by those agents who just receive a copy of the news. Our objective is to characterize conditions under which a news cascade can emerge as the result of micro-level utility-maximizing broadcast decisions. The formal definition of a cascade will be stated in Section 2.4.

Discussion of the Model

Here we comment on some of the assumptions made in our model and the rationale beneath those assumptions.

Belief-based utilities

This choice of utility is motivated by the observation that in many social settings, agents may base their information-sharing decisions on incentivizing others to take similar actions in an upcoming decision, rather than merely maximizing a single utility function. An example would be whether sharing a piece of news would persuade others to vote for one's party in a number of upcoming regional and national elections. Given the belief-driven nature of these actions, one may attain the similarity in her peers' actions by making a decision of sharing information that persuades peers to have similar beliefs.

Cost of broadcasting. As can be seen from (2.4), there is a cost associated with broadcasting the news in our model. This cost could represent people's inherent tendency to refrain from taking an action as many empirical studies have suggested that people show "status quo bias" while making decisions (Samuelson and Zeckhauser 1988). This is also known as consumers' inertia in the marketing literature studying consumers' brand-choice (Jeuland 1979, Su 2009). In our context of browsing posts, news, or tweets on social media, the default action of status quo is "not sharing". Faced with extensive amount of information on such platforms, agents usually scroll through page fast sifting the digital content. A post or tweet must gain enough interest from an agent to make her pause browsing and share the post.

The cost also partly accounts for agents' discomfort about public discussions under a post on social media. A post on Facebook or a tweet on Twitter may sometimes serve as a venue for intense discussions or even toxic debates. As reported in Newman et al. (2018, 2019), agents on social media are now more aware of the negative consequences that might ensue from their shared contents. It is important to note that, despite such concerns, retweets and content-sharing are still quite common on social platforms. This suggests that the cost is not prohibitive and that agents are willing to share news when they find sufficient value in doing so.

Being unaware of news. One implicit assumption in our model is that agents do not even know that there is a piece of news about the state unless they hear it. Many models also concern the players who are unaware of products, opportunities, or rumors until they observe their neighbors adopted or communicated.⁸ Sadler (2020) studied a diffusion game in which strategic players who are informed of the opportunity via their neighbors' adoption decide whether to adopt it as well considering network externality. Merlino et al. (2020) investigated the diffusion of rumors among strategic agents who can fact-check the rumor once they receive the rumor and know its existence.

In a setting of full rationality, however, an agent who knows that some news may exist and yet has not received it strategically interprets not hearing news as possibly a result of her followee deciding not to share it. This level of strategic speculation primarily concerns niche applications in which the access to information may be costly or restricted to some players, such as strategic communication at private meetings or organizations. As an example, in Chapter 4 we will investigate news subscription choices of fully rational subscribers who actively seek information, accounting for media bias in news disclosure. In the current context of news sharing on Twitter-like platforms where many tweets on a wide range of topics are circulating and users are passively fed digital contexts from others, such level of strategic speculation for the existence

⁸One can see Bass (1969), Campbell (2013) for diffusion based on non-strategic information transmission.

of specific news is not realistic. Moreover, if an agent speculates that there is news circulating in the society, she can readily look for it via digital search engines.

2.3 Agent's Decision of Sharing the News

2.3.1 Decision Rule

An agent i receiving a copy of the news x broadcasts it to her followers if and only if the marginal reduction in the expected average distance between her belief and her followers' surpasses the associated cost. Based on the utility function in (2.4), agent i decides to broadcast the news x if and only if

$$\mathbb{E}_i \left[W_2^2(P^i; P^k) | k \in N^{\text{in}}(i) \wedge a_i^B = 1 \right] + C < \mathbb{E}_i \left[W_2^2(P^i; P^k) | k \in N^{\text{in}}(i) \wedge a_i^B = 0 \right] \quad (2.5)$$

where the expectation is taken over the agent i 's belief on her followers' beliefs conditioned on her broadcast action $a_i^B \in \{0, 1\}$. To evaluate the expected payoff of action $a_i^B \in \{0, 1\}$, therefore, agent i needs to speculate on the effect of her action on her followers' beliefs: Choosing to broadcast the news ($a_i^B = 1$) ensures that all the followers will hear the news. If agent i decides not to broadcast it ($a_i^B = 0$), she still needs to weigh in the possibility that some of her followers may have already heard the news from others, in order to evaluate her expected utility.

Denote with q the probability that a randomly chosen follower of agent i has already heard the news from other agents at the time she makes her decision, and let $\bar{q} = 1 - q$. The independence of q from agent i is a consequence of the independent matching process of the in-stubs and out-stubs in the *configuration model*, as we will discuss in Section 2.4. The value of q is indeed a function of the dynamics of the news spread. Later in Section 2.4, after we explicitly develop the dynamics of the news spread, we will elaborate on how the probability q can be endogenously inferred from the network metrics known to the public and the current scale of the news spread in the population.

For the time being, however, let us assume the value of q as given. The decision rule (2.5) can then be written as

$$\begin{aligned} \mathbb{E}_i \left[W_2^2(P_1^i; P_1^k) | k \in N^{\text{in}}(i) \right] + C < q \times \mathbb{E}_i \left[W_2^2(P_1^i; P_1^k) | k \in N^{\text{in}}(i) \right] \\ + (1 - q) \times \mathbb{E}_i \left[W_2^2(P_1^i; P_0^k) | k \in N^{\text{in}}(i) \right], \end{aligned} \quad (2.6)$$

where P_1^k and P_0^k denote the beliefs conditioned on having received the news or not, respectively. P_0^k is the same as the prior while P_1^k is the posterior belief updated according to (2.1):

$$\begin{aligned} P_0^k &= \mathcal{N}(\mu_k, \sigma_\theta^2), \\ P_1^k &= \mathcal{N}((1 - \beta)\mu_k + \beta x, (1 - \beta)\sigma_\theta^2). \end{aligned}$$

Agents' prior perspectives are their private information and are independent from their positions in the network—in particular their in- and out-degrees. As a result, although agent i 's belief on the joint degree distribution of her followers differs from the commonly known joint distribution of in- and out-degrees in the population,⁹ her belief on their prior perspectives is the same as the common prior on the perspectives. This means that, for agent i and any of her followers $k \in N^{\text{in}}(i)$, $\mu_k \sim_i \mathcal{N}(\bar{\mu}, \sigma_\mu^2)$. Plugging in the expectation terms in (2.6), we can find the following threshold rules for agents' broadcast decisions:

Lemma 2.1 *Given the probability q and the received news x , agent i broadcasts the*

⁹This is related to *the friendship paradox*, as we explain in Section 2.4.

news if and only if the following holds:¹⁰

$$2 \underbrace{(1 - \beta)(\mu_i - \bar{\mu})}_{\text{interim bias of agent } i} \underbrace{\beta(x - \bar{\mu})}_{\text{interim surprise}} + \underbrace{\beta^2(x - \bar{\mu})^2}_{\text{magnitude of surprise}} + \underbrace{(1 - (1 - \beta)^2)\sigma_\mu^2}_{\text{reduction in the variance of perspectives}} > \frac{C}{1 - q}. \quad (2.7)$$

We use the decision rule in (2.7) to explain how the *surprise* and *affirmation* motives for sharing emerge naturally from the utility-maximizing broadcast decisions of agents. Let us define $S_\beta(x) \triangleq \beta(x - \bar{\mu})$ as the *interim surprise* in the news. Also define $B_\beta(\mu_i) \triangleq (1 - \beta)(\mu_i - \bar{\mu})$ as the *interim bias* of agent i , which is the difference between agent i 's posterior perspective and average posterior perspective *conditioned on hearing any realized news*.

The interim surprise measures the effective level of surprise in the news. The value $(x - \bar{\mu})$ alone represents the deviation of the news x from the average perspective $\bar{\mu}$ and its polarity. However, this value per se does not capture the significance of news credibility when Bayesian agents process the news. To see why, consider the example of receiving a report of discovering extraterrestrial life on Mars. This would be surprising if it is announced by NASA; by contrast, the report conveys little surprise if it is circulated on an online forum from an anonymous source. We therefore denote $S_\beta(x)$ as the effective surprise perceived by a Bayesian agent hearing news x . The interim bias is the difference between agent i 's posterior perspective and the average perspective of a random follower, hypothetically assumed to observe the news.¹¹ The direction of ex-ante and interim biases are the same but the magnitude of interim bias is smaller since agents put positive weight on the new information; the more credible the news,

¹⁰The following inequality is derived before (2.7):

$$\left(((1 - \beta)\mu_i + \beta x) - \bar{\mu} \right)^2 - \left(((1 - \beta)\mu_i + \beta x) - ((1 - \beta)\bar{\mu} + \beta x) \right)^2 > \frac{C}{q} - (1 - (1 - \beta)^2)\sigma_\mu^2.$$

From this we can see the gain in utility by decreasing the differences in the deviation of the mean of an agent's posterior from the aggregate mean of the public before and after they see the news.

¹¹This is similar to the setting used by Che and Kartik (2009). We also adopt their terminology.

the smaller the interim bias.

We say that the news x *affirms* the bias of agent i if the interim bias and the interim surprise have the same sign, i.e., $\text{sign}(S_\beta(x)) = \text{sign}(B_\beta(\mu_i))$; otherwise, we say that the news is *opposing* her bias. This affirmation effect appears as the first term in (2.7). From this, we can observe that an affirmative news improves the utility of a broadcast, while an opposing news forces against it. Moreover, the magnitude of affirmation/opposition effect is stronger for agents who are more biased away from the average perspective.

The second part in (2.7) represents the magnitude of the surprise, which improves the utility of broadcast regardless of its polarity. For affirmative news, the surprise makes the broadcast action more favorable. When news is opposing an agent's bias, its negative effect on the utility (first term in (2.7)) can be compensated by a sufficiently large surprise. We can rearrange the first two terms in (2.7) as

$$\left(2B_\beta(\mu_i) + S_\beta(x)\right) \times S_\beta(x). \quad (2.8)$$

Fix $B_\beta(\mu_i)$ and consider an opposing news item x such that $\text{sign}(S_\beta(x)) = -\text{sign}(B_\beta(\mu_i))$. When $|S_\beta(x)|$ is small, the opposition effect dominates the surprise and (2.8) is negative, forcing a no-broadcast action. However, as $|S_\beta(x)|$ increases, the surprise in the news eventually dominates, flipping agent i 's polarity and making (2.8) positive, favoring a broadcast action.

The third term in (2.7) is independent of both the news content and individual perspectives. This part corresponds to the utility that a broadcast action yields by concentrating the followers' perspectives further around their mean. Finally, note that the right-hand side of (2.7) is increasing as a function of the probability q , which captures the diminishing return of a broadcast as the news spreads wider and the followers are more likely to hear the news from other agents.

2.3.2 Broadcast Size

In this subsection, we use the micro-level decision rule of broadcasting characterized in (2.7) to identify the fraction of the newly-informed agents who will share the news with their followers. This enables us to analyze the dynamics of the news spread without the need to know each individual's perspective involved in the spread, thereby bridging the individual behavior at the micro-level with the collective behavior at the macro-level.

When the news is positively biased against the aggregate perspective, or $x - \bar{\mu} > 0$, the broadcast condition of (2.7) can be written as

$$\mu_i - \bar{\mu} > \frac{C}{2\bar{q}\beta(1-\beta)(x-\bar{\mu})} - \frac{(2-\beta)\sigma_\mu^2}{2(1-\beta)(x-\bar{\mu})} - \frac{\beta(x-\bar{\mu})}{2(1-\beta)}.$$

Given \bar{q} (that is, the probability that a follower has not seen the news yet), we can find the probability of broadcasting x for an agent who has just received the news:¹²

$$\begin{aligned} & \mathbb{P}\left(a_i^B = 1 \mid \text{agent } i \text{ received the news}\right) \\ &= \mathbb{P}\left(\mu_i - \bar{\mu} > \frac{C}{2\bar{q}\beta(1-\beta)(x-\bar{\mu})} - \frac{(2-\beta)\sigma_\mu^2}{2(1-\beta)(x-\bar{\mu})} - \frac{\beta(x-\bar{\mu})}{2(1-\beta)}\right) \\ &= \Phi\left(\frac{(2-\beta)\sigma_\mu}{2(1-\beta)(x-\bar{\mu})} + \frac{\beta(x-\bar{\mu})}{2(1-\beta)\sigma_\mu} - \frac{C}{2\bar{q}\beta(1-\beta)(x-\bar{\mu})\sigma_\mu}\right), \end{aligned}$$

where $\Phi(\cdot)$ is the standard normal CDF. Let us define

$$\begin{aligned} K_\beta(x) &= \frac{C}{2\beta(1-\beta)|x-\bar{\mu}|\sigma_\mu}, \\ \eta_\beta(x) &= \frac{\sigma_\mu}{2(1-\beta)|x-\bar{\mu}|}(2-\beta) + \frac{\beta|x-\bar{\mu}|}{2(1-\beta)\sigma_\mu}. \end{aligned}$$

¹²Since we assume independence of individual perspectives and agents broadcast the news instead of forwarding the news to a subset of her followers they select, receiving the news is independent from the agent's perspective. The case would be different if we consider homophily in perspectives.

It is easy to verify that for a general x ,

$$\mathbb{P}\left(a_i^B = 1 \mid \text{agent } i \text{ received the news } x\right) = \Phi\left(\eta_\beta(x) - \frac{K_\beta(x)}{\bar{q}}\right). \quad (2.9)$$

The expression in (2.9) is conditional on the value of \bar{q} at the time when an agent decides about her broadcast action. As will be discussed in detail in the next section, \bar{q} is a temporal variable which can be endogenously derived at each time instant, given the dynamics of news spread. It is also worth noting that the broadcast size (2.9) is the result of aggregating the broadcasting decisions of agents. Furthermore, for given \bar{q} and β , enlarging the difference of the news from the average perspective, i.e., $|x - \bar{\mu}|$, can increase the broadcast size $\Phi\left(\eta_\beta(x) - \frac{K_\beta(x)}{\bar{q}}\right)$.

2.4 Prevalence of News over the Network

Based on the micro-level decision rule derived in Section 2.3 for news sharing, we now study the conditions under which a cascade will emerge in steady state. We start by deriving the spread dynamics and then characterize the prevalence of the news in the network at steady state.

We say that a news cascade happens if and only if the set of the agents who have received the news in steady state reaches a size of $\Theta(n)$, where n is the number of agents in the population. Our analysis concerns the case in which $n \rightarrow \infty$, i.e., a large population of asymptotically infinite size. We identify conditions under which news initially shared only with a randomly chosen subset of size $o(n)$ of the population eventually spreads to a non-zero constant fraction of the population. The initial seeding of the news is assumed to be large enough to ensure that at least some agents in the giant component of the network \mathcal{G} , if any, would receive the news at $t = 0$; otherwise, a news cascade is never possible, even if everybody hearing the news broadcasts it to her neighbors. Furthermore, it is guaranteed that the fraction of agents who have received the news converges to a steady-state value as time goes to infinity since the fraction of

such informed agents in the network is non-decreasing in time and trivially bounded from above by 1.

We characterize the network by the joint distribution of its in-degrees and out-degrees. Following the standard notation, we denote with $P(\ell, d)$ the joint probability distribution of out-degrees and in-degrees, where d and ℓ stand for out-degree and in-degree respectively. Moreover, let $P^{\text{out}}(\cdot)$ and $P^{\text{in}}(\cdot)$ represent the marginal distributions of out-degrees and in-degrees respectively, and $P^{\text{out}}(\cdot|\ell)$ and $P^{\text{in}}(\cdot|d)$ be the conditional degree distributions.

The network is generated from a realization of in-degrees and out-degrees based on the *configuration model* (see Newman (2003) or Jackson (2008) for more details on the configuration model). As illustrated in Figure 2-2, the basic idea of generating a network from a sequence of in-degrees and out-degrees in the configuration model is to consider a “in-stub” (“out-stub”) for each in-link (out-link). The edges are then constructed by each time pairing unassigned in-stubs and out-stubs uniformly at random until all the stubs are assigned.¹³ Note that $\mathbb{E}[\ell] = \mathbb{E}[d]$ since the total in-stubs should match the total out-stubs (both count the total number of links) and we therefore refer to the expected in-degree/out-degree of the network as the expected degree for simplicity.

Regarding the agents’ knowledge of the underlying network, we follow the frameworks of Sundararajan (2008), Galeotti et al. (2010), and Fainmesser and Galeotti (2016): each agent only knows her own set of followers and followees but has no information about the interactions among other agents in the network, except for the joint degree distribution.

¹³It is possible to have self-links, or multiple links between two nodes under this process. In large networks, however, the chances of any node having duplicate or self-link tends to zero under some mild assumptions on the generated degree sequence. These issues are discussed in details in Newman (2003) and Jackson (2008).

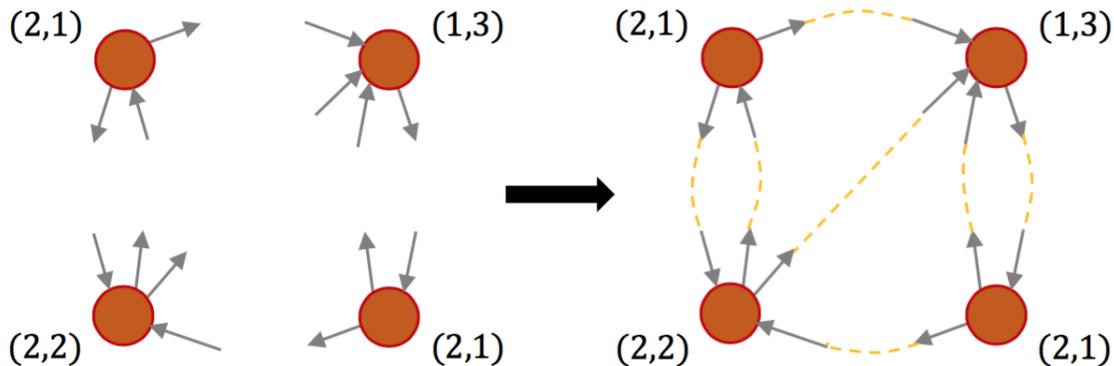


Figure 2-2: Illustration of the network generation process by configuration model from a sequence of (in-degree, out-degree) pairs.

2.4.1 Spread Dynamics

The following notations and definitions prove useful in characterizing the dynamics of the spread. For any subset \mathcal{A} of the agents, we define $|\mathcal{A}|$ as the probability of a randomly chosen agent belonging to \mathcal{A} . In other words, $|\mathcal{A}|$ represents the size of \mathcal{A} normalized by the population size. We also use \mathcal{A}^{in} to denote the set of edges with their head in \mathcal{A} , and $|\mathcal{A}^{\text{in}}|$ as the probability that a randomly chosen edge $e \in \mathcal{E}$ lies in \mathcal{A}^{in} . \mathcal{A}^{out} and $|\mathcal{A}^{\text{out}}|$ are defined similarly. Finally, the set of those agents not in \mathcal{A} is denoted by $\bar{\mathcal{A}}$.

Let $t = 0$ be the time when the news is initially released to a subset of (randomly and uniformly chosen) agents in the population. Denote with $r_\beta(x, t)$ the set of agents who receive the news x with credibility β at time t and let $b_\beta(x, t) \subseteq r_\beta(x, t)$ be the set of those who decide to share the news with their followers at time t . The set of all agents who have received the news up to time t is denoted by $\mathcal{R}_\beta(x, t) = \cup_{\tau=0}^t r_\beta(x, \tau)$. Similarly, $\mathcal{B}_\beta(x, t) = \cup_{\tau=0}^t b_\beta(x, \tau)$. The fraction of agents who have received the news in the steady state is $\lim_{t \rightarrow \infty} |\mathcal{R}_\beta(x, t)|$ which we denote with $|\mathcal{R}_\beta(x, \infty)|$.

The broadcast decision rule in Lemma 2.1 requires an agent (who just received a copy of the news and is deciding whether to broadcast) to estimate the probability

q that a randomly chosen follower of hers has already heard the news. Recall from Section 2.2 that, whether a follower has heard the news or not is her private information and hence unknown to the agent making the broadcast decision. Nonetheless, we assume that the decision-making agent can observe the current scale of the news spread in the entire population (that is, $|\mathcal{R}_\beta(x, t)|$), from which she can estimate q . This makes q a temporal variable, endogenously evolving with the news spread, and hence we denote it hereafter with $q_\beta(x, t)$. An important property here is that agents observing $|\mathcal{R}_\beta(x, t)|$ at time t will have identical estimates for $q_\beta(x, t)$, which is in particular independent of their in-degrees (out-degrees); this is due to the independence matching process of the in-stubs and out-stubs (followers and followees) in the configuration model. Indeed, it is easy to see that

$$q_\beta(x, t) = \mathbb{P}\left(\text{Tail}(e) \in \mathcal{R}_\beta(x, t) \mid \text{Head}(e) \in r_\beta(x, t)\right). \quad (2.10)$$

In the following lemma, we describe how to find $q_\beta(x, t)$ from the publicly observable cascade size $|\mathcal{R}_\beta(x, t)|$ and the common knowledge of the joint degree distribution.

Lemma 2.2 *An agent receiving the news at time t and observing the current size of news spread $|\mathcal{R}_\beta(x, t)|$ can find $q_\beta(x, t)$, the probability that a randomly chosen follower of hers has seen the news, using the following equations:*

$$\begin{aligned} 1 - q_\beta(x, t) &= \sum_{d=1}^{\infty} \frac{d}{\mathbb{E}[d]} (1 - |\mathcal{B}_\beta^{\text{in}}(x, t-1)|)^{d-1} P^{\text{out}}(d), \\ 1 - |\mathcal{R}_\beta(x, t)| &= \sum_{d=0}^{\infty} (1 - |\mathcal{B}_\beta^{\text{in}}(x, t-1)|)^d P^{\text{out}}(d). \end{aligned} \quad (2.11)$$

Based on Lemma 2.2, the agent making the broadcast decision can derive the probability $q_\beta(x, t)$ needed in Lemma 2.1 from $|\mathcal{B}_\beta^{\text{in}}(x, t-1)|$, that is, the fraction of the edges through which the news has traversed (i.e., triggered by the broadcast) by time $t-1$; this fraction $|\mathcal{B}_\beta^{\text{in}}(x, t-1)|$ can be itself estimated from the size of the spread at time t (i.e., $|\mathcal{R}_\beta(x, t)|$) from (2.11). It should be noted that the probability $q_\beta(x, t)$

of a randomly chosen follower having seen the news is generically different from the probability that a randomly chosen individual from the population has seen the news (that is, $|\mathcal{R}_\beta(x, t)|$). This result is rooted in the difference between the out-degree distributions of a follower and a random member of the population, reflecting the fact that it is more likely to be followed by an agent with a higher number of followees (out-degrees)—this phenomenon echoes *the friendship paradox* (Feld 1991).

In order to characterize the evolution of the spread it suffices to find the update rule for $|\mathcal{B}_\beta^{\text{in}}(x, t)|$. Noting that $|\mathcal{B}_\beta^{\text{in}}(x, t + 1)| = |\mathcal{B}_\beta^{\text{in}}(x, t)| + |b_\beta^{\text{in}}(x, t + 1)|$, and that $|b_\beta^{\text{in}}(x, t + 1)| = \Phi(\eta_\beta(x) - \frac{K_\beta(x)}{\bar{q}_\beta(x, t+1)})|r_\beta^{\text{in}}(x, t + 1)|$ from (2.9), we can come up with the following update rule for $|\mathcal{B}_\beta^{\text{in}}(x, t)|$ (see the proof of Theorem 2.1 for the details):

$$\begin{aligned} & |\mathcal{B}_\beta^{\text{in}}(x, t + 1)| - |\mathcal{B}_\beta^{\text{in}}(x, t)| \\ &= \frac{\Phi(\eta_\beta(x) - \frac{K_\beta(x)}{\bar{q}_\beta(x, t+1)})}{\mathbb{E}[d]} \sum_{d=0}^{\infty} \mathbb{E}[\ell|d] ((1 - |\mathcal{B}_\beta^{\text{in}}(x, t - 1)|)^d - (1 - |\mathcal{B}_\beta^{\text{in}}(x, t)|)^d) P^{\text{out}}(d). \end{aligned}$$

We can use the continuous-time mean-field approximation of the above second-order nonlinear dynamics to approximate the asymptotic behavior of the spread, as described in the following Theorem 2.1.

Theorem 2.1 *Denote by $|\mathcal{B}_\beta^{\text{in}}(x, \infty)|$ the fraction of the links through which the news has passed in the steady state. Then, $|\mathcal{B}_\beta^{\text{in}}(x, \infty)|$ is the largest solution of $G(|\mathcal{B}_\beta^{\text{in}}(x, \infty)|) = 0$, where*

$$G(|\mathcal{B}_\beta^{\text{in}}(x, \infty)|) = \int_0^{|\mathcal{B}_\beta^{\text{in}}(x, \infty)|} \left(-1 + \frac{\Phi(\eta_\beta(x) - \frac{K_\beta(x)}{\bar{q}(y)})}{\mathbb{E}[d]} \sum_{d=0}^{\infty} d \mathbb{E}[\ell|d] (1 - y)^{d-1} P^{\text{out}}(d) \right) dy,$$

and

$$\bar{q}(y) = \frac{1}{\mathbb{E}[d]} \sum_{d=1}^{\infty} d (1 - y)^{d-1} P^{\text{out}}(d),$$

if $G(1) < 0$; otherwise, $|\mathcal{B}_\beta^{\text{in}}(x, \infty)| = 1$. The steady state size of the spread $|\mathcal{R}_\beta(x, \infty)|$ is given by

$$|\mathcal{R}_\beta(x, \infty)| = 1 - \sum_{d=0}^{\infty} (1 - |\mathcal{B}_\beta^{\text{in}}(x, \infty)|)^d P^{\text{out}}(d).$$

We can use Theorem 2.1 to derive necessary and sufficient conditions for news x to cascade, i.e., $|\mathcal{R}_\beta(x, \infty)| > 0$, as given in Proposition 2.1.

Proposition 2.1 *Consider the model of news sharing established in Section 2.2. For a news realization x with credibility β a cascade happens almost surely if and only if¹⁴*

$$\frac{\mathbb{E}[\ell d]}{\mathbb{E}[d]} \Phi(\eta_\beta(x) - K_\beta(x)) > 1. \quad (2.12)$$

The cascade condition (2.12) depends on the network structure through the quantity $\frac{\mathbb{E}[\ell d]}{\mathbb{E}[d]}$. This is indeed the average in-degree/out-degree¹⁵ of the *line graph* $L(\mathcal{G})$ of the original network \mathcal{G} . Recall that the vertices in the line graph $L(\mathcal{G})$ are edges of the original graph \mathcal{G} , and (x, y) is an edge in $L(\mathcal{G})$ if and only if $x = (r, s)$ and $y = (s, t)$ are two edges in \mathcal{G} for some $r, s, t \in \mathcal{V}(\mathcal{G})$ (Aigner 1967). The average degree in the line graph quantifies the expected number of edges that can be *triggered* by a randomly chosen edge in the original graph. We denote with $\mu^{\text{Line}}(\mathcal{G})$ the average (or mean) degree of the line graph $L(\mathcal{G})$. We can rewrite the cascade condition (2.12) as

$$\mu^{\text{Line}}(\mathcal{G}) \Phi(\eta_\beta(x) - K_\beta(x)) > 1.$$

Recalling the broadcast probability from (2.9), the cascade condition simply suggests that a cascade emerges if and only if, after receiving the news, each follower in expectation will pass the news to more than one new agent (where the expectation is taken over

¹⁴In fact, not all the random realizations of the network generated from configuration model have a giant component, even if the condition of the proposition is met. It happens, however, in almost sure sense as $n \rightarrow +\infty$. The same asymptotic behavior applies to our model.

¹⁵Of course these two quantities are the same, and we henceforth refer to them simply as the average degree.

both the broadcast action and the joint degree distribution), assuming a cascade has not happened yet. For a hypothetical scenario where every agent receiving the news broadcasts it to her followers (i.e. the broadcast size is always 1), (2.12) is reduced to $\mu^{\text{Line}}(\mathcal{G}) > 1$. That is, the necessary and sufficient conditions for the existence of a nonzero-measured giant component in a given directed network generated from the configuration model.¹⁶

2.4.2 Effect of the Network Structure

We note that in the cascade condition (2.12), given the credibility β , the set of news that meets the cascade condition is increasing¹⁷ in $\mu^{\text{Line}}(\mathcal{G})$, manifesting the structural fact that cascades are induced when there are on average more links through which a follower can further share the news. The monotonicity enables us to compare the susceptibility of different network structures to the spread of imprecise/false news by examining the expected degrees of their corresponding line graphs.

To illustrate the effect of the joint degree distribution on the average degree of the line graph, we compare the average degree of the line graph for network structures generated from two different joint degree distributions: joint Poisson and Zipf (power-law) degree distributions.

We first introduce a useful identity for the average degree of the line graph for a network \mathcal{G} with joint degree distribution $P(\ell, d)$:

$$\mu^{\text{Line}}(\mathcal{G}) = \frac{\mathbb{E}[\ell d]}{\mathbb{E}[d]} = \frac{\text{cov}(\ell, d) + (\mathbb{E}[d])^2}{\mathbb{E}[d]} = \frac{\rho\sigma_\ell\sigma_d}{\mathbb{E}[d]} + \mathbb{E}[d], \quad (2.13)$$

where σ_ℓ (resp. σ_d) represents the standard deviation of in-degree (resp. out-degree) distribution and ρ is the correlation coefficient between in- and out-degree distributions.

¹⁶This condition in terms of the mean degree of the line graph accommodates the version for the emergence of a giant component in undirected networks. See the details in, e.g., Newman (2003) and Jackson (2008).

¹⁷Denote a collection of set $\{\mathcal{S}_\alpha\}_{\alpha \in \mathcal{I}}$ indexed by an ordered index set \mathcal{I} . We say that the sequence $\{\mathcal{S}_\alpha\}_{\alpha \in \mathcal{I}}$ is increasing if for any $a, b \in \mathcal{I}$ such that $a \preceq b$, then $\mathcal{S}_a \subseteq \mathcal{S}_b$.

The equation (2.13) enables us to represent $\mu^{\text{Line}}(\mathcal{G})$ in terms of the statistics of the joint degree distribution and to investigate the effects of each component, *ceteris paribus*.

In what follows, we use (2.13) to juxtapose the mean degree of the line graph for networks generated from joint Poisson and Zipf (power-law) degree distributions—with similar expected in-degrees/out-degrees and a common correlation coefficient.

Definition 2.1 *A joint Poisson degree distribution is represented by the following PMF:*

$$P^{\text{Pois}}(d, \ell) = \sum_{m=0}^{\min\{d, \ell\}} \frac{\lambda^{d-m} \lambda^{\ell-m} \lambda_c^m}{(d-m)! (\ell-m)! m!} e^{-(2\lambda + \lambda_c)}, \quad \lambda \geq 0, \lambda_c \geq 0. \quad (2.14)$$

A joint Zipf degree distribution is represented by the following CCDF:

$$P^{\text{Zipf}}(d \geq m_1, \ell \geq m_2) = \left(1 + \frac{m_1 + m_2}{h}\right)^{-\alpha}, \quad (2.15)$$

where $h > 0$, $\alpha > 0$, and m_1, m_2 are non-negative integers.¹⁸

For Poisson and Zipf degree distributions, we can express the average degrees of the corresponding line graphs in terms of their expected degrees and correlation coefficients, as Lemma 2.3 shows.

Lemma 2.3

- (i) *The line graph of a joint Poisson degree distribution with expected degree $\mu_P > 0$ and correlation coefficient $\rho_P \in [0, 1]$, has a mean degree of $\mu_P^{\text{Line}} = \rho_P + \mu_P$.*
- (ii) *The line graph of a joint Zipf degree distribution with expected degree $\mu_Z > 0$ and correlation coefficient $\rho_Z \in (0, \frac{1}{2})$, has a mean degree of $\mu_Z^{\text{Line}} = \frac{\rho_Z(\mu_Z + 1)}{1 - 2\rho_Z} + \mu_Z$.*

Proposition 2.2 compares the average degrees of line graphs for these two degree distributions with the common expected degree and correlation coefficient.

¹⁸We adapt the definitions of bivariate Poisson (Kawamura 1973) and Zipf distributions (Yeh 2002) to joint degree distributions by requiring $\mathbb{E}[d] = \mathbb{E}[\ell]$. See the proof of Lemma 2.3 for the details.

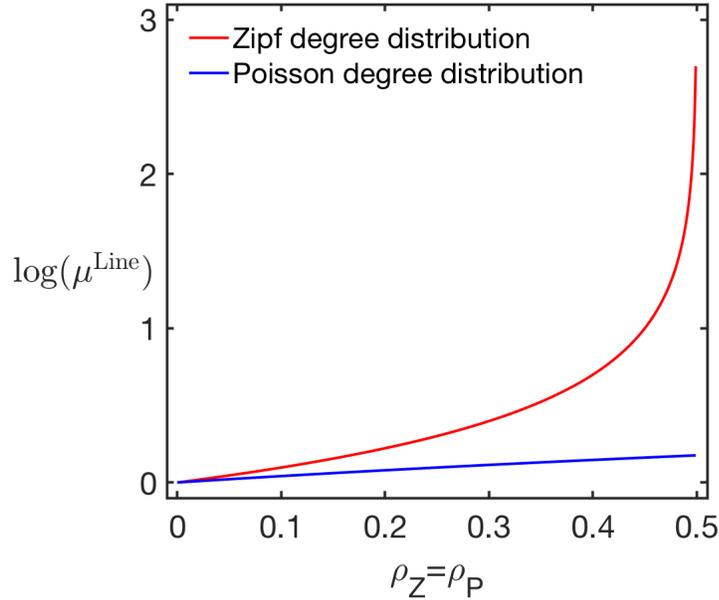


Figure 2-3: The mean degree of the line graphs for Zipf and Poisson degree distributions with common expected degrees $\mu_Z = \mu_P = 1$ and correlation coefficient $\rho_Z = \rho_P \in (0, 0.5)$.

Proposition 2.2 Consider a joint Poisson degree distribution and a joint Zipf degree distribution with a common expected degree $\mu_P = \mu_Z > 0$ and a common correlation coefficient $\rho_P = \rho_Z \in (0, \frac{1}{2})$. Then,

- (i) $\mu_Z^{\text{Line}} > \mu_P^{\text{Line}}$.
- (ii) A network with the Zipf degree distribution is more permeable to news spread than a network with the Poisson degree distribution.

Figure 2-3 shows the mean degree of the line graphs for Zipf and Poisson degree distributions with common expected degrees on a log scale, varying their common correlation coefficient. The rapid growth of μ_Z^{Line} as the correlation coefficient increases suggests that for Zipf degree distribution, even networks with small expected degrees could allow a wide range of news to cascade, provided sufficiently high correlation coefficient between the in-degrees and out-degrees.

2.5 News Precision Levels Maximizing Cascade Probability

Proposition 2.1 establishes the necessary and sufficient conditions for news items to go viral as a function of the credibility of its source, network statistics, and information structure of the agents in the network. Given the inherent noise in the news that not only controls the distribution of the news but also determines its credibility, a natural question is then to identify the noise level that maximizes the *ex-ante* likelihood of news going viral *before the news is realized*.

Given the variance of the priors σ_θ^2 , changing the noise level σ_ϵ^2 is the same as varying the credibility $\beta \in [0, 1]$. For a piece of news x , a cascade emerges almost surely if and only if the condition (2.12) is satisfied. Therefore, the ex-ante likelihood of a news cascade is the mass (probability) of all news x for which (2.12) holds. The problem of finding the credibility level that maximizes the ex-ante likelihood of a cascade can therefore be formulated as

$$\begin{aligned} & \underset{\beta \in [0, 1]}{\text{maximize}} \mathbb{P}_X \left(\mu^{\text{Line}}(\mathcal{G}) \Phi(\eta_\beta(X) - K_\beta(X)) > 1 \right), & (2.16) \\ & \text{subject to:} \\ & X \sim \mathcal{N}(\theta, \frac{1 - \beta}{\beta} \sigma_\theta^2), \\ & K_\beta(X) = \frac{C}{2\beta(1 - \beta)\sigma_\mu |X - \bar{\mu}|}, \\ & \eta_\beta(X) = \frac{\sigma_\mu}{2(1 - \beta)|X - \bar{\mu}|} (2 - \beta) + \frac{\beta}{2(1 - \beta)\sigma_\mu} |X - \bar{\mu}|. \end{aligned}$$

The rest of this section is dedicated to analyzing the solutions to this optimization problem. We limit our analysis to the networks with $\mu^{\text{Line}}(\mathcal{G}) > 2$, to which we refer as *well-connected networks*. This ensures all the results to still hold for the general distance of the form (2.2), as we show in the Appendix. Indeed, there is not much interest, if any, in the case $\mu^{\text{Line}}(\mathcal{G}) \leq 2$ given the well-connectedness of the social

$\mathbb{E}[d]$	σ_ℓ	σ_d	$\mathbb{E}[\ell d]$	ρ	$\mu^{\text{Line}}(\mathcal{G})$
35.3	2419.7	354.7	223867.1	0.259	6350.3

Table 2.1: Network statistics for the follower graph of Twitter (Data source: Kwak et al. 2010).

platforms: For the follower network of Twitter, which is the motivating application of our work, we have $\mu^{\text{Line}}(\mathcal{G}) \approx 6350.3$; Table 2.1 lists some of the basic network statistics.¹⁹

The next theorem relates the optimal credibility levels to the diversity of perspectives and deviation of the truth from the average of perspectives in a population.

Theorem 2.2 *Given the variance of the perspectives σ_μ^2 , let $\Omega^*(\sigma_\mu^2)$ be the set of values β that maximize the likelihood of a cascade as formulated in (2.16). Assume that $\mu^{\text{Line}}(\mathcal{G}) > 2$. Then,*

(i) *If $\sigma_\mu^2 > C$, $\Omega^*(\sigma_\mu^2) = [\bar{\beta}^*(\sigma_\mu^2), 1]$ and a cascade emerges almost surely, where*

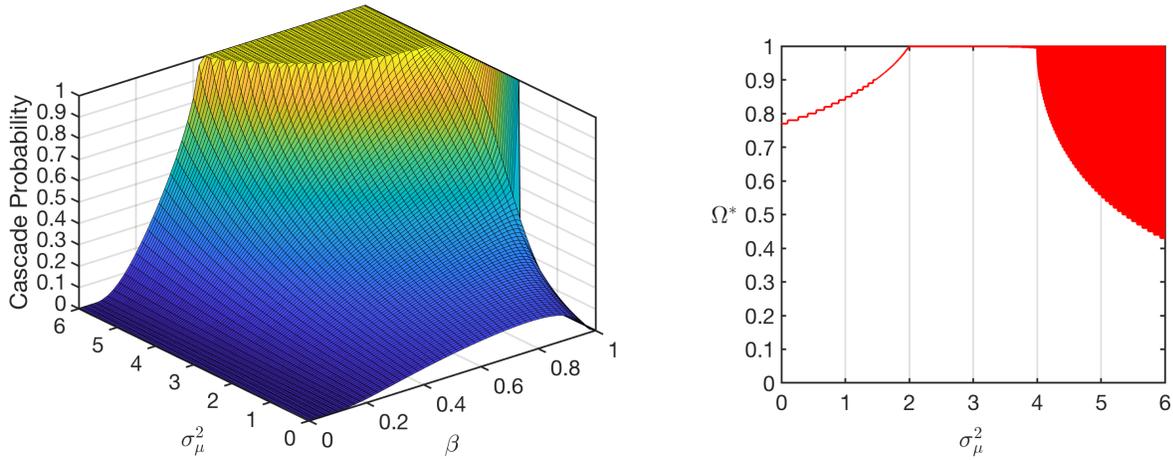
$$\bar{\beta}^*(\sigma_\mu^2) = 1 - \sqrt{1 - \frac{C}{\sigma_\mu^2}}.$$

Furthermore, $\bar{\beta}^(\sigma_\mu^2)$ is strictly decreasing in σ_μ^2 .*

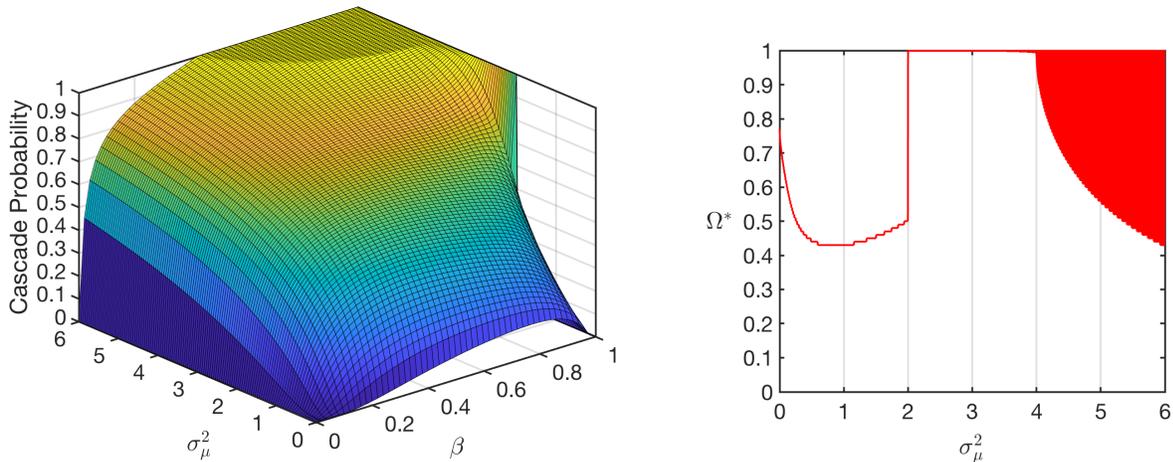
(ii) *If $\sigma_\mu^2 \leq C < \sigma_\mu^2 + (\theta - \bar{\mu})^2$, then $\Omega^*(\sigma_\mu^2) = \{1\}$ and a cascade emerges almost surely.*

(iii) *If $\sigma_\mu^2 + (\theta - \bar{\mu})^2 \leq C$, then $1 \notin \Omega^*(\sigma_\mu^2)$ and particularly, the truth never cascades.*

¹⁹We used the open-source dataset collected and kindly shared by Kwak et al. (2010) (available at <http://an.kaist.ac.kr/traces/WWW2010.html>) of over 41.7 million nodes and 1.47 billion links to obtain the statistics reported in Table 2.1. In fact, the well-connectedness we define in this paper can be observed in many empirical network analyses on the follower graphs of Twitter or other social media (Java et al. 2007, Mislove et al. 2007, Zhou et al. 2010, Alipourfard et al. 2020): their results show that the expected degrees are greater than 2 and the correlation between in- and out-degrees is positive, indicating that the mean degrees of the line graphs are larger than 2 according to the identity (2.13).



(a) $\mu^{\text{Line}}(\mathcal{G}) = 2.2$



(b) $\mu^{\text{Line}}(\mathcal{G}) = 6350.3$

Figure 2-4: Illustration of the three cases in Theorem 2.2 for a sample case with $\bar{\mu} = 0, \sigma_{\theta} = 2, (\theta - \bar{\mu})^2 = 2, C = 4$, and $\mu^{\text{Line}}(\mathcal{G}) = \{2.2, 6350.3\}$. (Left) Cascade likelihood versus the credibility β and the diversity in perspectives σ_{μ}^2 . (Right) The set of optimal credibility Ω^* as the correspondence of the diversity in perspectives σ_{μ}^2 .

The three cases in Theorem 2.2 are illustrated in Figure 2-4 via an example. In what follows, we explain the underpinnings of the results summarized in the above theorem, while providing insights into the relation between the cascade likelihood and the diversity and the wisdom of a population.

2.5.1 In a Population with Highly-Diverse Perspectives

Recall from Lemma 2.1 in Section 2.3.1 that, when the size of the spread is almost zero (i.e., news has not yet gone viral), agent i broadcasts news x if and only if

$$2 \underbrace{(1 - \beta)(\mu_i - \bar{\mu})}_{\text{interim bias of agent } i} \underbrace{\beta(x - \bar{\mu})}_{\text{interim surprise}} + \underbrace{\beta^2(x - \bar{\mu})^2}_{\text{magnitude of surprise}} + \underbrace{(1 - (1 - \beta)^2)\sigma_\mu^2}_{\text{reduction in the variance of perspectives}} > C. \quad (2.17)$$

Examining $\beta = 1$, we see that in a population with wide diversity in perspectives ($\sigma_\mu^2 > C$), the truth can incentivize an agent with any perspective to share the news and always trigger a cascade. In such a population, however, news may also go viral almost surely for a wide range of other credibility levels β . To see this, consider sufficiently precise news (e.g., β close to 1) whereby the gain resulting from reducing the variance of perspectives in (2.17) surpasses the cost of broadcasting. For news having such a level of credibility, at least half of those followers who receive the news will in turn share it with their own followers: if the news is not affirming the position of agent i with perspective $\mu_i = \bar{\mu} + \delta$, then it is otherwise affirming an agent with perspective $\bar{\mu} - \delta$ and therefore the broadcast condition (2.17) holds for at least one of them. As a result, such a credibility level will always result in a cascade no matter what the content of the news is, in a well-connected network where $\mu^{\text{Line}}(\mathcal{G}) > 2$.

2.5.2 When the Truth is Surprising to a Population with Homogeneous or Less Diverse Perspectives

When the diversity of perspectives is not large, i.e., $\sigma_\mu^2 \leq C$, the gain from reducing the variance of followers' perspectives will not make up for the cost, even when this gain is fully extracted by making the followers unanimous on the truth, i.e., the maximum of σ_μ^2 . The surprise effect then becomes prominent in compensating the cost of broadcasting. Specifically, the truth cascades if and only if $(\theta - \bar{\mu})^2 > C - \sigma_\mu^2$, that

is, the truth needs to carry sufficient surprise ($|S_1(\theta)| = |\theta - \bar{\mu}|$) to the public. It is also easy to see that in this case $\beta = 1$ is the only credibility level for which a cascade will always emerge; any $\beta < 1$ can result in news realizations that are not surprising enough (e.g., those realizations close to $\bar{\mu}$) to satisfy the broadcast condition in (2.17) to trigger a cascade.

2.5.3 In a Wise Population with Homogeneous Perspectives

If a population is wise and their perspectives are homogeneous in the sense that $(\theta - \bar{\mu})^2 + \sigma_\mu^2 \leq C$, then the truth has no chance of going viral. Highly imprecise news is also unlikely to trigger a cascade: note that for $\beta \ll 1$ interim surprise $\beta(x - \bar{\mu})$ is concentrated around zero ($\beta(x - \bar{\mu}) \sim \mathcal{N}(\beta(\theta - \bar{\mu}), \beta(1 - \beta)\sigma_\theta^2)$). Therefore, cascade probability is maximized at some moderate level of credibility $\beta \in (0, 1)$.

Given the diversity, the difference between the cascade probability for extreme values of the news precision (i.e., β close to 0 or 1) and the optimal cascade probability widens further in more connected networks with larger $\mu^{\text{Line}}(\mathcal{G})$. As already discussed in Section 2.4.2, the cascade probability is increasing in $\mu^{\text{Line}}(\mathcal{G})$ at all precision levels. Therefore, for networks with large $\mu^{\text{Line}}(\mathcal{G})$ such as the follower graph of Twitter (with $\mu^{\text{Line}}(\mathcal{G}) \approx 6350.3$), the optimal cascade probability is pushed further away from 0, resulting in a sharper drop from the peak cascade probability to nearly zero for highly precise and highly imprecise news. This can be observed from Figures 2-4(a) and 2-4(b), by comparing the cascade probability for extreme values of β (i.e., close to 0 and 1) with its optimal value.

Figure 2-5 depicts the optimal credibility for three different levels of population wisdom, illustrating case (ii) and (iii) in Theorem 2.2 when the diversity is not wide ($\sigma_\mu^2 \leq C$). We see that the minimum level of diversity necessary for the truth to cascade is higher for wiser populations, as expected by the condition $\sigma_\mu^2 > C - (\theta - \bar{\mu})^2$ from Theorem 2.2. We can also observe that, fixing the level of diversity, the truth almost never goes viral in a sufficiently wise population (when $(\theta - \bar{\mu})^2 \leq C - \sigma_\mu^2$).

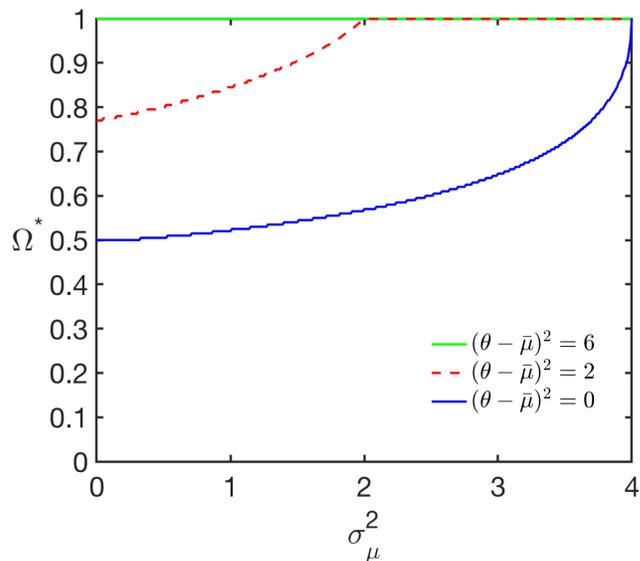


Figure 2-5: Optimal credibility levels $\Omega^*(\sigma_\mu^2)$ for $\sigma_\mu^2 \in [0, 4]$ for a sample case with $C = 4$, $\sigma_\theta = 2$, $\mu^{\text{Line}}(\mathcal{G}) = 2.2$, and three levels of collective wisdom $(\theta - \bar{\mu})^2 \in \{0, 2, 6\}$ illustrating cases (ii)-(iii) in Theorem 2.2.

As a final remark, we would like to generalize an observation that can be made from Theorem 2.2 on the monotonicity of the set of credibility levels that would almost surely trigger a cascade: the more diverse the perspectives of a population, the wider the range of these credibility levels. We can indeed prove that in a well-connected population (with $\mu^{\text{Line}}(\mathcal{G}) > 2$) the diversity of perspectives facilitates the spread of the news for all credibility levels, as the following proposition shows.

Proposition 2.3 *In a well-connected social network where $\mu^{\text{Line}}(\mathcal{G}) > 2$, the cascade probability is increasing in the variance of perspectives σ_μ^2 for all credibility levels.*

2.6 Discussion

We provided a theory for news-sharing behavior on social networks. Agents make sharing decisions aiming to reduce the average distance between their beliefs and the beliefs of their followers, which captures their desire for like-minded others. We showed that both the *surprise* and *affirmation* motives naturally emerge from the utility-maximizing behavior of agents.

We further characterized the endogenous dynamics of the news spread resulting from the agents' sharing decisions, subsequently deriving necessary and sufficient conditions for a news cascade to emerge. We used this condition to show the monotonicity of the cascade likelihood in the average degree of the line graph of the network. To demonstrate the effect of joint degree distributions on the susceptibility of networks to news cascades, we then compared two network structures generated from joint Zipf and Poisson degree distributions with common expected degrees and correlation coefficients, showing that news is always more likely to cascade on the network with the joint Zipf degree distribution.

Lastly, we revealed a connection between the precision levels maximizing ex-ante likelihood of a cascade and the diversity of perspectives, as well as the wisdom of the population. In particular, we showed that in a well-connected network, the cascade likelihood is increasing in the diversity of perspectives for all precision levels. In a very diverse population, a wide range of precision levels may result in viral news, including the truth. The range of such precision levels shrinks as the population becomes less diverse. By contrast, in a population with moderately-diverse or homogeneous perspectives, the truth has to be sufficiently surprising to the public so as to cascade; when the perspectives are so concentrated around the truth, however, the truth has no chance of going viral and the cascade probability is maximized at some nonzero precision level.

In this chapter, we studied the decision making of non-strategic agents who decide whether to share news contingent on the current size of spread without speculating on the future news spreading. However, size of news spread is (weakly) increasing in time and an agent will tend to be more reluctant to share news if she can foresee that the myopic individual news-sharing decisions leads to a (weakly) wider spread than at the current time. In the next chapter, we turn to focus on strategic agents under a game-theoretic setting and provide new insights into news-sharing behavior.

Chapter 3

A Game-theoretic Model of News Sharing on Social Networks

3.1 Background

In Chapter 2, we considered myopic agents in the sense that they assess the marginal gain in sharing news with their followers given the current spread and do not internalize or assume about the likelihood that an uninformed follower may hear the news from others at some point in the future. Different from this approach, in this chapter, we propose a simple, game-theoretic setting that is focused on news-sharing decisions of individuals on social networks. Using the proposed game-theoretic model, we revisit several fundamental questions regarding news dissemination on social networks: (i) Under what conditions does a piece of news that initially goes to only an infinitesimal subset of the population spread to a significant fraction of the population? (ii) Is it possible that low credibility news spreads wider than highly credible news? (iii) What is the association between population characteristics, news credibility and size of news spread?

Focusing on the persuasion motive for sharing, in the model we study the decision process of news sharing for a continuum of agents who are connected on a social

network. The agents have heterogeneous subjective prior beliefs on an unobservable real-valued state. They each will later take a binary action, like voting, that matches the sign of the state based on their own beliefs. Nature generates a single piece of real-valued news about the state and the news credibility that represents the sensitivity of agents' beliefs in reaction to the news; this is consistent with the approach to news generation in Chapter 2. Initially, a small and random subset of the population are informed of the news and its credibility.

When an agent observes the news and its credibility, she updates her belief, fusing her prior belief and the news. She then chooses whether to share the news with all of her followers (e.g., retweet the news), deliberating how effectively the news can persuade her followers to cast the vote in line with her belief. As in Chapter 2, we still assume that news sharing is not costless.

Due to the nominal cost associated with sharing, agents' sharing decisions feature strategic substitutability. Specifically, if an agent speculates that her followers are likely to hear the news through other agents, she may choose not to share to avoid the cost. As a result, each agent has to strategically evaluate the impact of *other agents' news-sharing decisions* on her followers' information sets while making her sharing decision.

Our approach to modeling the network is the same simple random graph model as the one in Chapter 2. However, we will in particular consider Poisson distribution for the number of one's followees and followers.

As our first result, we characterize the news-sharing strategies in the game. The sharing rule is a function of an agent's belief mean about the state and the effective persuasiveness power of the news: Agents with more extreme beliefs have stronger incentives (or stronger disincentives, depending on whether the news is in alignment with their beliefs) to share news. We establish the system of equations that govern the interdependency between individual sharing strategies and the scale of news spread that endogenously emerges. Using the equations we quantify the cascade size at the

equilibrium and show the necessary and sufficient conditions for occurrence of a news cascade. Given news credibility, we show that more extreme news increases the spread size and that news spreads wider on more connected networks.

As our second contribution, we identify the necessary and sufficient condition under which news with lower credibility can trigger a larger cascade on social networks. We further discover that when news is quite neutral to both actions, lower credibility can make news diffusion larger than full credibility. The beliefs of informed agents about the state are concentrated around the small news value when observing high level of credibility. Consequently they find insufficient benefit of making their followers' actions aligned with the sign of the state, having weak incentives to persuade their followers. By contrast, though lower credibility undermines the persuasiveness of news, the extremist agents still hold relatively extreme beliefs about the state and prefer to share, leading to a larger cascade than credible news.

The condition also indicates a limit on network connectivity (the average number of followers) beyond which fully credible news no longer generates a cascade with larger size than lower credibility news. This result is rooted at the strategic substitutability of sharing decisions: The ease of diffusion caused by greater connectivity suggests that informed agents share the news only if they hold more extreme beliefs and, in turn, stronger sharing incentives. For news with lower credibility, since it does not move the beliefs of extremists towards the news value as much as fully credible news, more fraction of informed agents hold extreme beliefs and participate in sharing, resulting in a larger sharing cascade.

Finally, we study news sharing and diffusion in populations exhibiting belief polarization. We offer theoretical insights into how the intensity of polarization and the level of perspective diversity within each ideology group (in-group diversity) may facilitate spillovers of less credible news and affect the limit on network connectivity.

We find that increased polarization *decreases* the limit on network connectivity level above which lower credibility news becomes more viral than fully credible news. This

is due to the fact that a large fraction of agents who hold extremist beliefs will still share the news after hearing the news with lower credibility. Despite the conventional wisdom that diversifying opinions across the ideological spectrum may mitigate the spread of less credible information in highly-polarized populations, our analysis shows that the impact of increased belief diversity is not necessarily monotone and can even result in wider spread of less credible news: When there is low in-group diversity in perspectives, increased diversity heavily reduces the fraction of the extremists who are in line with the ideology of their own group and impedes sharing of less credible news, increasing the network connectivity limit. However, since increased diversity also increases the number of the extremists whose perspectives have the opposite ideological leaning from their group, the aggregate fraction of extremists is instead rising up when the perspectives are too diverse, propelling the sharing of news with lower credibility and reducing the connectivity limit.

Our work contributes to the recent line of theoretical research on news sharing and spread of misinformation. Several papers have studied the prevalence of misinformation arising among strategic individuals whose motive is to share and act on the truth. Kranton and McAdams (2020) focused on news markets on which news consumers can not fact-check news and need to speculate how likely a news producer may report low-quality news given that producing high-quality news is costly. Moreover, in Kranton and McAdams (2020) news-sharing is limited to only a single round, and happens over a d -regular network. On the other hand, Merlino et al. (2020) enabled each individual hearing a rumor to imperfectly inspect the truth with probability of success increasing in the effort level the individual chooses, and individuals communicate to others the opinion in line with their beliefs as the inspection result is not decisive. Papanastasiou (2020) considered perfect inspection on an article and studied decisions of costly inspection and sharing for a stream of strategic agents who have heterogeneous beliefs and can not observe any inspection actions (nor the results) but only sharing decisions by previous agents. In contrast, in our model an individual shares news or

not to persuade her social contacts to take the action in alignment with her own belief.

This chapter significantly generalizes our preliminary results on news sharing, reported in Hsu et al. (2021). First, unlike the full focus of Hsu et al. (2021) on Gaussian distributions and news that is not fully credible, the present work allows for non-parametric distributions over agents' ex-ante expectations about the unobservable state and characterizes the equilibrium for any news credibility level. This generalized model enables us to quantify the impact of different population characteristics such as polarization and in-group diversity on the size of sharing cascades for different levels of credibility, and in turn on the levels of connectivity required for less credible news to trigger a larger cascade. We deliver new insights into the mechanisms through which news, network connectivity, and beliefs in a population affect individual news-sharing decisions and spread size of news. Specifically, the main contribution of Hsu et al. (2021) was to determine the set of news values that can trigger a cascade, i.e., whether or not the news spreads to a positive fraction of the population (irrespective of the cascade size), and to identify the sufficient and necessary condition under which the spread size is increasing in credibility. By contrast, this paper shows the comparative statics on spread size and characterizes the sufficient and necessary condition under which less credible news can spread to a larger fraction of the population. Moreover, in this paper we uncover the relationship between polarization, in-group diversity, and network connectivity limits, shedding light on spread of less credible news in polarized societies. As a particular case of interest, whereas the specific setting in Hsu et al. (2021) (with no polarization) would suggest a lower connectivity limit if in-group diversity is increased, our new result demonstrates that the effect of in-group diversity can in fact be non monotone, or even reversed, in a polarized population.

Another related work to our paper is by Acemoglu et al. (2021), where the authors investigate the fact-checking behavior of agents and sharing news on social networks, assuming that agents are craving for future shares on the news they shared while having a reputation concern of being caught sharing misinformation. Our work differs in

three fundamental aspects about individual motives and actions. First, in our model agents are motivated to persuade their followers via news sharing; an informed agent may choose not to share high credibility news if the news suggests lack of value about matching the unobservable state. By contrast, in Acemoglu et al. (2021) agents avoiding sharing misinformation strictly prefer to spread the news that is more likely to be truthful from their own perspectives. Second, our work excludes the option of costly fact-check by agents and the reputation concern. Consequently, the force is not present that larger polarization increases the pressure of being caught by the other group’s fact-checking and inhibits the spread of misinformation. Instead, in our model strong polarization means that there are more extremists who have strong incentives to share news, potentially boosting spillovers of less credible news. Lastly, while the sharing decisions of our game are strategic substitutes, the consideration of one’s reputation perceived by other peers makes their model feature strategic complementarity.

Our focus on individuals’ strategic news sharing on social networks is also related to the recent literature on strategic interactions in large random graphs (Parise and Ozdaglar 2021, Board and Meyer-ter Vehn 2021).

The rest of this chapter proceeds as follows. We describe our model and discuss the assumptions in Section 3.2. We analyze the equilibrium of our model in Section 3.3, elucidating the interaction between sharing decisions and scale of news spread. In Section 3.4 we identify conditions for lower credibility news to trigger a larger cascade than fully credible news, using which we characterize the limit on network connectivity. We then provide the implications of our model for news spread in populations with polarized perspectives. We discuss in Section 3.6. The technical proofs of all lemmas and theorems are all relegated to the Appendix B.

3.2 Model

We consider a unit-mass continuum of agents, each indexed by a real number $i \in \mathcal{I} \triangleq [0, 1]$. Social interactions are characterized by a directed network $\mathcal{G} = (\mathcal{I}, \mathcal{E})$. The population is concerned with an unobservable state $\theta \in \mathbb{R}$, for which there will be a voting and each agent casts her binary vote. Ex ante the agents have heterogeneous subjective prior beliefs about the state. Prior to the voting, each agent interacts with her social contacts and, if observing some piece of news, may share the news with her followers, influencing their beliefs about the state.

The game proceeds in three stages 0, 1, 2. In stage 0, Nature makes three moves: (1) Each agent is independently assigned a prior belief about the state, drawn from a common distribution; (2) the social network \mathcal{G} is generated based on a simple random network model; (3) news is generated and given to some random agents with infinitesimal size of the population. All agents are initially unaware of the existence of the news except for the agents who received the news from Nature. In stage 1 agents interact: Upon seeing some news, an agent has to decide immediately whether to broadcast the news or not. We also call stage 1 as interaction stage. Stage 2 is the voting stage where each agent simultaneously casts a binary vote. The utilities resulting from sharing decisions and votes are realized after the voting is completed. We provide the specifics as below.

Network of social interactions. The directed network $\mathcal{G} = (\mathcal{I}, \mathcal{E})$ is exogenously constructed, representing follower–followee relationships among the agents. We say that agent j *follows* agent i if and only if $(j, i) \in \mathcal{E}$: In this relationship, agent j is agent i 's follower and agent i is agent j 's followee. For each agent i , we denote the set of her followers as $N_i^{\text{in}} \triangleq \{j \in \mathcal{I} | (j, i) \in \mathcal{E}\}$ and let $\ell_i \triangleq |N_i^{\text{in}}|$ be the cardinality of this set, to which we refer as agent i 's in-degree. Similarly, we write $N_i^{\text{out}} \triangleq \{j \in \mathcal{I} | (i, j) \in \mathcal{E}\}$ for the set of agent i 's followees and let $d_i \triangleq |N_i^{\text{out}}|$ denote agent i 's out-degree.

Network \mathcal{G} is constructed as follows. Each agent i is randomly assigned with two in-

tegers as her in-degree ℓ_i and out-degree d_i respectively. We assume that the in-degrees and out-degrees of agents are uncorrelated and follow a common discrete degree measure $P^{\text{deg}}(\cdot)$. Assuming that the draws for agents' in- and out-degree are independent of each other, we say that the agents with in-degree ℓ (out-degree d) account for the fraction $P^{\text{deg}}(\ell)$ ($P^{\text{deg}}(d)$) of the population by informal use of the law of large numbers. Hereafter, we denote as $P^{\text{deg}}(\cdot)$ the degree distribution of the network. Agent i with out-degree d_i is then randomly assigned with d_i others according to an atomless weighted-uniform distribution on the population $[0, 1]$, where the weights are proportional to the in-degrees of the agents.¹

Throughout this paper we make a particular choice for the degree distribution:

Assumption 1 *We consider Poisson degree distribution with parameter λ , i.e., $P^{\text{deg}}(d) = \frac{\lambda^d}{d!} e^{-\lambda}$, $d = 0, 1, 2, \dots$. The parameter λ is the expected degree of Poisson distribution to which we refer as the network connectivity. We focus on the range $\lambda > 1$ for which there exists at least one weakly-connected component in the network comprising a positive fraction of the population.*²

The abstract network model is common knowledge to all the agents while the realization \mathcal{G} is not: Agents have no knowledge about the interactions among others and only know their local positions in the network \mathcal{G} . Specifically, each agent only knows the sets of her followers N_i^{in} and followees N_i^{out} , which are her private information.

Prior beliefs and types of agents. The agents have heterogeneous subjective prior beliefs about the state θ . We assume that their prior beliefs only differ with their real-valued prior expectations (means) about the state. Individual prior means are their

¹Specifically, an agent with in-degree ℓ is ℓ times likely to be drawn as a followee than an agent with in-degree 1, i.e., the probability density that agent i is sampled is $\frac{\ell_i}{\mathbb{E}[\ell]}$ where $\mathbb{E}[\ell]$ is the average degree. See Galeotti and Goyal (2009), Fainmesser and Galeotti (2016, 2020) for this modeling approach.

²Many empirical studies indeed support large average degrees of social networks (Java et al. 2007, Mislove et al. 2007, Kwak et al. 2010, Zhou et al. 2010) and hence $\lambda \leq 1$ is not of practical interest. For example, the dataset which was collected and kindly shared by Kwak et al. (2010) for the follower network on Twitter (as the motivating application of the present work) comprises over 41.7 million nodes and 1.47 billion links, with a value of 35.2 for the average degree.

private information; we denote each agent i 's prior mean about the state as $\mu_i \in \mathbb{R}$, and call it agent i 's *perspective* (or type). However, it is commonly known that each agent's perspective is independent and identically distributed according to a common CDF F , i.e., for any agent i

$$\mu_i \sim F.$$

Assumption 2 *We assume that CDF F is continuously differentiable with corresponding density function f and full support on the real line. The mean and variance of function f are finite. For simplicity of exposition, we assume that $F(0) = \frac{1}{2}$. We write \mathcal{F} for the space of the CDFs.*

News, credibility, and belief updating. Nature generates a single piece of *news* $x \in \mathbb{R}$, informative about the state and includes the news source in the news so that agents can distinguish copies of the same news. We assume that all agents agree on the precision of the news source; the objective precision is measured as the *credibility* $\beta \in [0, 1]$, which quantifies the sensitivity of agents' beliefs in reaction to the news.

When an agent i receives the news (x, β) , she then knows that the news is circulating in the population and updates her belief about the state. We denote as $z_i \in \{\emptyset, (x, \beta)\}$ her private history of being informed or not: We set $z_i = \emptyset$ when agent i has not received any news while $z_i = (x, \beta)$ if she has received the news. We abstract the belief updating process by focusing only on each agent i 's belief mean, which we denote as $\mathbb{E}_i[\theta | \mu_i, z_i]$ given her type (μ_i) and whether she is informed or not (z_i) . When agent i has not received the news ($z_i = \emptyset$) and hence stays unaware of the news, her belief mean stays unchanged, i.e., $\mathbb{E}_i[\theta | \mu_i, \emptyset] = \mu_i$. For an agent who has observed the news, we particularly consider the following linear updating rule for the mean of her belief:

Assumption 3 *For each agent i , if she observes the news (x, β) , her belief mean is*

updated as

$$\mathbb{E}_i[\theta|\mu_i, (x, \beta)] = \beta x + (1 - \beta)\mu_i, \quad (3.1)$$

where the credibility β is the sensitivity of the updated mean to the observed news.

This linear updating rule is in the same form as the Bayesian updating rule (2.1) derived in Chapter 2, where agents hold Gaussian priors on the state and news is generated as the state with additive Gaussian noise. In fact, this particular updating rule for the mean of one's belief can be rationalized by considering a distribution in the exponential family for news generation and the associated conjugate prior as an agent's prior on the state (Diaconis and Ylvisaker 1979).

Sharing decision, voting, and utility function. The interaction stage takes place in discrete time $t = \{0\} \cup \mathbb{N}$. At $t = 0$, the news (x, β) is released to a random subset of the population with an infinitesimal mass $\delta > 0$. Any agent i who just received the news has to immediately make a decision $s_i \in \{0, 1\}$ as whether to share the news with her followers or not, which is described as her *sharing* or *broadcast* decision. Sharing news, however, incurs a fixed cost, which will be specified in the utility functions. We denote agent i 's sharing strategy as $s_i(\mu_i, z_i, F) : \mathbb{R} \times \{\emptyset, (x, \beta)\} \times \mathcal{F} \rightarrow \{0, 1\}$ where $s_i(\mu_i, \emptyset, F) = 0$ as she can share news only if she has received some news. The inclusion of F as an argument captures agent i 's consideration of her followers' perspectives on the state in her decision making. By convention, we write s as the profile of sharing strategies; we denote by $s((x, \beta), F)$ the sharing strategy profile for news (x, β) and F .

We assume that upon observing the news, agents are unaware of when the news spread started: They know the calendar of the game but do not know the time corresponding to $t = 0$. Moreover, we assume the interaction stage continues long enough for the news spread to reach its steady state.³ We define the spread size of news as

³Tracking online mainstreams and social media activities, many empirical studies (Leskovec et al. 2009, Schema et al. 2019) have found that the lifetime of online news circulating in a society is mostly under two weeks; this time span is relatively short compared to the duration until the voting that is almost always announced well in advance.

follows:⁴

Definition 3.1 *Given λ, F, δ , and sharing strategy profile s , the spread size for news x with credibility β is the fraction of the population who have received the news at the steady state and we denote it by $q(\lambda, F, \delta, s((x, \beta), F))$.*

In the voting stage, each agent i casts a binary vote $a_i \in \{-1, +1\}$ that matches the sign of her belief mean. Specifically, based on the belief updating (3.1), agent i votes using the following rule:⁵

$$a_i(\mu_i, z_i) = \text{sign}(\mathbb{E}_i[\theta|\mu_i, z_i]). \quad (3.2)$$

Let a denote the profile of the votes.

After the voting ends, agent i 's utility resulting from her sharing decision s_i and the agents' votes a is realized, and is given by

$$u_i(\theta, s_i, a) = \theta(w^I a_i + w^L \bar{a}_i^L + w^G \bar{a}) - C \times \mathbb{I}_{\{s_i=1\}}, \quad (3.3)$$

where

$$\bar{a}_i^L \triangleq \frac{1}{|N_i^{\text{in}}|} \sum_{k \in N_i^{\text{in}}} a_k; \quad \bar{a} \triangleq \int_{j \in [0,1]} a_j dj;$$

the weights $w^I, w^L, w^G \in (0, 1)$ are common for all the agents, satisfying $w^I + w^L + w^G = 1$, and $C > 0$ represents a universal cost associated with broadcasting news. The utility function (3.3) captures an agent's preference for alignment of collective actions with

⁴The present work concerns the size of news-sharing cascades arising from individuals' sharing decisions, with no reference to structural properties of news-sharing dynamics on online social networks (see Vosoughi et al. (2018), Zhao et al. (2020)). This choice enables tractable analysis with focus on the individuals' motivation for sharing news, and is also consistent with the findings by Juul and Ugander (2021) that the structural difference in news spread can be largely explained by difference in spread size, and moreover, by the transmission rate of news between peers.

⁵For simplicity we assume that agents take the vote +1 when their belief mean is zero. This simplification does not affect agents' strategies of sharing in the following analysis as we consider a continuum of agents and the measure of such agents is zero.

the state at individual, local, and global scales, comprising a convex combination of three components: (1) agent i 's individual vote a_i ; (2) the average vote of agents' followers; (3) the average vote of the population. Agent i derives positive utility when the weighted average of collective votes matches the sign of the state and the gain/loss is increasing in the magnitude of state. Since the agents have different subjective beliefs on the state, agent i has an incentive to influence, via sharing the news, her followers' beliefs and in turn their votes.

We solve for Bayesian Nash equilibria (BNE) of the game, characterizing the news-sharing strategies of the agents. Note that we exclude from each agent i 's strategy formulation the private information about her followees and followers $N_i^{\text{in}}, N_i^{\text{out}}$. Indeed, in our model the strategies are independent of the identities of one's followees and followers as well as her in-degree and out-degree, as we will demonstrate in Section 3.3.

Our particular interest is whether the news can spread to a non-zero fraction of the population when the size of the initial seeding is infinitesimal.

Definition 3.2 *Given λ, F , and strategy profile s , we say that a cascade of news (x, β) emerges almost surely if and only if $\lim_{\delta \rightarrow 0} q(\lambda, F, \delta, s((x, \beta), F)) > 0$; we refer to the limit as the corresponding cascade size (or spread size).*

We use the following terminology when discussing our results. The binary votes to match the sign of an unobservable state naturally splits the information space into positive and negative. We refer to positive side (+1) as *right-leaning* and negative side (-1) as *left-leaning*, using which we describe the sign of an agent's perspective and the sign of news. In particular we refer to $x = 0$ as *neutral* news. The absolute value of news x , i.e., $|x|$, is the *news magnitude*, which measures how extreme the news is away from the neutral point. Finally, we say that news is fully credible when it has the full level of credibility, i.e., $\beta = 1$.

3.3 Equilibrium Analysis

We first derive the mechanics of how news spread emerges from the agents' news-sharing decisions. We then elucidate on agents' decision making given their knowledge about the network structure and speculation on news spread, deriving their equilibrium news-sharing strategies.

3.3.1 News spread emerging from news-sharing decisions

Based on the random network model and the fact that individual perspectives are private information, we compute the size of news spread q at its steady state using susceptible-infection model (SI model, see Jackson 2008), the network structure (Poisson network with connectivity λ), and *the fraction of informed agents who share the news*.

Let $P_F(s((x, \beta), F))$ denote the fraction of informed agents who share the news (x, β) according to sharing strategy profile s , measured with respect to perspective CDF F :⁶

$$P_F(s((x, \beta), F)) \triangleq \mathbb{P}(s_i(\mu_i, (x, \beta), F) = 1 | i \in \mathcal{I}, \mu_i \sim F).$$

Since each following link is drawn uniformly at random from the population, the probability that at the steady state a random following link points to an agent who has shared the news equals the probability that it follows an agent who has received the news (q) *and* would like to forward the news (P_F). At the steady state, the fraction of the agents who has not received the news at the whole interaction stage is $1 - q$, which are the agents who did not receive it at $t = 0$ and has not received it from any of her followees. We can hence write the equation for the probability that a random

⁶Since we assume independence of individual perspectives and agents broadcast the news instead of forwarding the news to a subset of her followers they select, receiving the news is independent from the agent's perspective. Therefore, the chance of sharing news for an agent who receives the news equals the fraction of the population who will share the news upon seeing it.

agent does not hear the news at the steady state:

$$1 - q = (1 - \delta) \sum_{d=0}^{\infty} P^{\text{out}}(d) (1 - qP_F(s((x, \beta), F)))^d.$$

Using $P^{\text{out}}(\cdot) \sim \text{Poisson}(\lambda)$ and taking $\delta \rightarrow 0$, we obtain the following result:

Lemma 3.1 *Given λ, F , and sharing strategy profile s , the spread size $q(\lambda, F, s)$ for news x with credibility β satisfies*

$$q = 1 - e^{-q\lambda P_F(s((x, \beta), F))}. \quad (3.4)$$

Moreover, $q = 0$ if $\lambda P_F(s((x, \beta), F)) \leq 1$.⁷

Different from the conventional version of SI model in which the probability of information transmission is an exogenous constant, in our model the probability of news sharing is a function of sharing strategy profile s , which endogenously hinges on the size of news spread, as will be shown later.

3.3.2 News-sharing decisions in light of news spread

We adopt the conventional notation s_{-i} for the profile of sharing strategies other than i . Upon receiving the news and observing its credibility (x, β) , agent i 's sharing decision $s_i \in \{0, 1\}$ is her best response to agents' voting rule a and others' sharing strategies s_{-i} , if and only if there is no deviation strictly improving her expected utility:

$$\mathbb{E}_i \left[u_i(\theta, s_i, a) | s_{-i}, \mu_i, z_i = (x, \beta) \right] \geq \mathbb{E}_i \left[u_i(\theta, 1 - s_i, a) | s_{-i}, \mu_i, z_i = (x, \beta) \right], \quad (3.5)$$

where the expectation is taken over agent i 's belief about all the other agents' perspectives on the state and their private histories at the voting stage. To evaluate (3.5),

⁷It can be shown that $q = 0$ is the only solution to (3.4) if and only if $\lambda P_F(s((x, \beta))) \leq 1$. When $\lambda P_F(s((x, \beta))) > 1$, the spread size is the non-zero solution to (3.4) since $q = 0$ is not a stable solution due to the infinitesimal fraction of the population who initially received the news. See the proof of Theorem 3.1 in the Appendix and Jackson (2008) for more details.

agent i needs to speculate on the influence of her sharing decision s_i , as well as other agents' sharing decisions, on any agent's private history. However, it suffices to consider only whether her followers are informed or not. In a population with a continuum of agents, the sharing decision of a single agent with a finite number of followers has a negligible effect on the spread size of news. Therefore, given any s_{-i} , agent i 's sharing decision only influences the average vote of her followers \bar{a}_i^L , without affecting the average vote of the population \bar{a} ; equivalently, agent i evaluates whether persuading her followers is worthy of the cost of sharing.

Specifically, choosing to broadcast the news ($s_i = 1$) ensures that all of her followers will hear the news. On the other hand, if she decides not to broadcast ($s_i = 0$), she has to weigh in the chance that some of her followers may end up hearing the news from other agents at the steady state of the news spread. For Poisson networks, one can show that this chance equals the spread size of the news:

Lemma 3.2 *On a Poisson network with connectivity λ , given F and profile of sharing strategies s , for any follower of agent i , say $k \in N_i^{\text{in}}$, the probability that follower k receives the news (x, β) from agents other than agent i at the steady state of the news spread is equal to the spread size of the news $q(\lambda, F, s((x, \beta), F))$.⁸*

Therefore, for any agent i 's follower $k \in N_i^{\text{in}}$, the probability that follower k is informed at the steady state of the news spread conditional on agent i 's sharing decision s_i and other agents' sharing strategies s_{-i} is estimated as:⁹

$$\mathbb{P}(z_k = (x, \beta) | s_i, s_{-i}) = \begin{cases} 1, & \text{if } s_i = 1; \\ q(\lambda, F, s((x, \beta), F)), & \text{if } s_i = 0. \end{cases} \quad (3.6)$$

⁸For general out-degree distributions, the identity does not necessarily hold due to two factors: (i) the out-degree distribution of a follower should be adjusted by the well-known friendship paradox; (ii) the chance is evaluated conditioned on the event that a follower with out-degree d can receive the news only from any of her other $d - 1$ followees.

⁹Given any s_{-i} the value $q(\lambda, F, s_i = 0, s_{-i})$ in (3.6) is independent of s_i because the sharing decision of a single agent does not affect the spread size of news in .

The average vote of agent i 's followers can then be computed as

$$\begin{aligned} & \mathbb{E}_i \left[\frac{1}{|N_i^{\text{in}}|} \sum_{k \in N_i^{\text{in}}} a_k(\mu_k, z_k) \mid s_i, s_{-i}, z_i = (x, \beta) \right] \\ &= \begin{cases} \mathbb{E}_i [a_k \mid k \in N_i^{\text{in}}, z_k = (x, \beta)], & \text{when } s_i = 1; \\ q \mathbb{E}_i [a_k \mid k \in N_i^{\text{in}}, z_k = (x, \beta)] + (1 - q) \mathbb{E}_i [a_k \mid k \in N_i^{\text{in}}, z_k = \emptyset], & \text{when } s_i = 0. \end{cases} \end{aligned} \quad (3.7)$$

As individual perspectives are private information and each follower k uses the same voting rule (cf. (3.2)), we observe from (3.7) that all informed agents are faced with the same evaluation process, independent of the identity and number of their own followers, and consequently deploy the same sharing strategy in the game.¹⁰ We henceforth refer to s as the common sharing strategy.

The following lemma characterizes an agent's decision rule in news sharing:

Lemma 3.3 *Given x, β, λ, F , and the endogenous spread size $q(\lambda, F, s((x, \beta), F))$ arising from the strategy s , agent i strictly prefers to broadcast the news if the following inequality holds:*

$$\underbrace{\left[\beta x + (1 - \beta) \mu_i \right]}_{\text{agent } i\text{'s belief mean}} \quad \underbrace{\left[2(1 - q) F^{\text{flip}}(x, \beta) \right]}_{\text{effective persuasive power of the news}} \quad > \quad \frac{C}{wL}, \quad (3.8)$$

where

$$F^{\text{flip}}(x, \beta) \triangleq F(0) - F\left(\frac{-\beta x}{1 - \beta}\right),$$

measures the fraction of agents who will change their favored votes upon receiving the news (x, β) . Agent i does not broadcast the news if the reversed inequality of (3.8) holds.

Three elements are identified as indicated in (3.8): belief mean, effective persuasive

¹⁰As the set of followees provides no additional information and is not involved in the utility function, the sharing decision of each agent does not depend on N_i^{out} .

power of the news, and the adjusted cost of sharing.

The effective persuasive power of the news measures the probability that a random follower will flip her favored vote *directly due to* the sharing by the agent making a decision. Specifically, the sign of $F^{\text{flip}}(x, \beta)$ indicates the direction of vote flipping, which is determined by the sign of news: Right-leaning (resp. left-leaning) news can only change the favored vote from -1 to $+1$ (resp. $+1$ to -1). The magnitude, i.e., $|F(0) - F(\frac{-\beta x}{1-\beta})|$ is the fraction of agents who flip their favored votes when seeing news x . However, considering the probability that a follower may also receive the news from other agents, this difference has to be discounted by $1 - q$ to reflect the effective fraction that can be persuaded *only* by the decision maker, showing the strategic substitutability of news-sharing decisions. Finally, the multiplicity 2 accounts for the change in aggregating votes when one agent flips her favored vote.

We observe that the magnitude of persuasive power is increasing in both news magnitude and credibility: Farther the news is away from being neutral or higher the news credibility is, then a larger fraction of the population will switch their favored votes upon seeing the news. In particular, when $\beta = 0$, no agents share the news since the news has zero persuasive power, generating no benefits of sharing. It is a trivial scenario where news does not spread; for simplicity we hereafter exclude $\beta = 0$ from our discussion.

The adjusted cost reflects an agent's care about the average votes of her followers in her sharing decision: A larger weight w^L means a higher level of concern and increases her willingness to persuade her followers. For simplicity of exposition, we henceforth use notation C to represent the adjusted cost, normalized with the weight w^L .

In (3.8) the product of belief mean and the effective power of persuasiveness is the gain from sharing the news, exhibiting monotonicity in perspective given all the other parameters. Agents with more extreme perspectives place higher value on their votes (absolute value of their belief means), and hence have stronger incentives to share or not to share the news, depending on whether the sign of news is in line with their beliefs.

This observation suggests a structural form for sharing strategy that we characterize in the following definition:

Definition 3.3 *We say that agent i deploys a threshold sharing strategy if given F and β , there exists a set of thresholds $\mu^{\text{th}}((x, \beta), F)$ for each x such that*

$$s(\mu_i, (x, \beta), F) = \begin{cases} \mathbb{I}\{\mu_i > \mu^{\text{th}}((x, \beta), F)\}, & \text{when } x \geq 0; \\ \mathbb{I}\{\mu_i < \mu^{\text{th}}((x, \beta), F)\}, & \text{when } x < 0. \end{cases} \quad (3.9)$$

where $\mathbb{I}\{\cdot\}$ denotes the indicator function.

3.3.3 Equilibrium structure

Lemma 3.1, 3.2 and 3.3 together identify the interplay between the scale of news spread at macro-level and individual threshold strategies via the fraction of informed agents who share the news. Noting that the sharing decision rule in Lemma 3.3 becomes the same for the informed agents when $\beta = 1$, we proceed with characterizing the equilibrium in the two cases of $\beta < 1$ and $\beta = 1$.

Not fully credible news ($\beta < 1$)

Lemma 3.3 establishes a strategy following the form (3.9) whereby the threshold μ^{th} must satisfy the indifference equation¹¹

$$\left[\beta x + (1 - \beta)\mu^{\text{th}} \right] \left[2 \left(1 - q(\lambda, F, s((x, \beta), F)) \right) F^{\text{flip}}(x, \beta) \right] = C. \quad (3.10)$$

The following theorem identifies the equilibrium for $\beta < 1$.

Theorem 3.1 *Given any $\beta < 1$, for any x, λ, F , an equilibrium exists and it is unique, with the following properties:*

¹¹The agents whose perspectives equal the threshold feel indifferent between sharing or not; for simplicity we assume that they choose not to share the news. Again this simplification does not affect agents' strategies of sharing as the measure of critical agents is zero. When $x = 0$ the decision rule (3.8) suggests no sharing by any informed agents, which is reflected by the corresponding threshold of positive infinity in (3.10).

(i) *The agents deploy a threshold sharing strategy.*

(ii) *The common threshold $\mu^{\text{th}*}$, the steady-state spread size q^* and the sharing fraction P_F^* are the unique solution of the following system of equations:*

$$\begin{aligned} \mu^{\text{th}*} &= \frac{C}{2(1-q^*)F^{\text{flip}}(x,\beta)} - \beta x && \text{(indifference condition)} \\ P_F(s^*) &= \begin{cases} 1 - F(\mu^{\text{th}*}) & \text{when } x \geq 0; \\ F(\mu^{\text{th}*}), & \text{when } x < 0, \end{cases} && \text{(sharing fraction)} \\ q^* &= 1 - e^{-q^* \lambda P_F(s^*)}. && \text{(spread dynamics)} \end{aligned}$$

For news that is not fully credible, the beliefs of informed agents about the state are still heterogeneous: The informed agents with extreme beliefs on one ideological wing strictly prefer to share the news whereas those holding extreme beliefs with the opposite ideology strictly prefer not to share. The equilibrium sharing fraction is therefore non-zero.

Fully credible news ($\beta = 1$)

When $\beta = 1$, for the informed agents, the mean of their beliefs are identical and equal to the news value. Lemma 3.3 hence yields the identical criterion for all the informed agents for news sharing:

$$|x| \left(1 - q(\lambda, F, s((x, 1), F)) \right) \begin{matrix} \geq \\ \leq \end{matrix} C. \quad (3.11)$$

As a result, the informed agents may (a) all strictly prefer sharing (when “ $>$ ” holds in (3.11)); or (b) all strictly prefer not sharing (“ $<$ ” holds); or (c) all feel indifferent (“ $=$ ” holds). Case (a), in which each agent shares the news (i.e., $s \equiv 1$), is an

equilibrium if

$$|x| > \frac{C}{1 - q(\lambda, F, s \equiv 1)},$$

where $q(\lambda, F, s \equiv 1)$ is in fact the size of the giant component of network \mathcal{G} when every agent shares the news. Similar for case (b) where no agent shares the news (i.e., $s \equiv 0$), the spread size is 0, and not sharing is an equilibrium if

$$|x| < C.$$

For the last case (c), any strategy that assigns sharing to an arbitrary subset of perspectives sustains an equilibrium if the induced fraction of agents who share the news given by $P_F \in [0, 1]$, and the consequent spread size q satisfy

$$\begin{aligned} |x| &= \frac{C}{1 - q} \\ q &= 1 - e^{-q\lambda P_F}. \end{aligned}$$

The equilibria in all three cases are in fact in the class of threshold strategies: The cases (a) and (b) in fact correspond to a threshold strategy (3.9) with thresholds of infinite magnitude, while for case (c) the threshold is finite. The next theorem summarizes the above discussion deriving the closed-form solutions for spread size, the fraction of informed agents sharing news, and the sharing threshold at the equilibrium for $\beta = 1$.

Theorem 3.2 *Suppose $\beta = 1$. Given x, λ, F there exists an equilibrium possessing the following properties:*

- (i) *The agents deploy a threshold sharing strategy.*
- (ii) *The common threshold $\mu^{\text{th}*}$, the steady-state spread size q^* and the sharing frac-*

tion P_F^* are given by:

$$(q^*, P_F^*)(x) = \begin{cases} (0, 0), & \text{if } |x| \in [0, C]; \\ \left(1 - \frac{C}{|x|}, \frac{-\ln(\frac{C}{|x|})}{\lambda(1 - \frac{C}{|x|})}\right), & \text{if } |x| \in (C, \frac{C}{1-q^G(\lambda)}]; \\ (q^G(\lambda), 1), & \text{if } |x| > \frac{C}{1-q^G(\lambda)}, \end{cases} \quad (3.12)$$

$$\mu^{\text{th}*}(x) = \begin{cases} F^{-1}(1 - P_F^*(x)), & \text{when } x \geq 0; \\ F^{-1}(P_F^*(x)), & \text{when } x < 0, \end{cases}$$

where $q^G(\lambda)$ is the size of the giant component of \mathcal{G} , i.e., the largest solution to $q = 1 - e^{-q\lambda}$.

Based on Theorem 3.1 and 3.2, the equilibrium spread size and sharing fraction are continuous functions of credibility β in the range $[0, 1]$:

Corollary 3.1 *The steady-state spread size q^* , the sharing fraction P_F^* , and the common sharing threshold $\mu^{\text{th}*}$ are continuous functions of β . Specifically,*

$$\lim_{\beta \rightarrow 1} q^*(\beta) = q^*(1).$$

With the focus on threshold sharing strategies (3.9), the sharing threshold is also continuous at $\beta = 1$. This continuity property will later be used to identify the conditions under which fully credible news does not maximize the cascade size.

3.3.4 Effects of news magnitude and connectivity on news spread size

Theorem 3.2 shows that given $\beta = 1$ and the sign of news, the spread size is increasing in news magnitude (cf. (3.12)). This can be explained by inspecting the sharing decision making in (3.8). Recall that the persuasive power of news (i.e., $|F(0) - F(\frac{-\beta x}{1-\beta})|$)

reaches its full level $\frac{1}{2}$ at $\beta = 1$ regardless of news magnitude. As informed agents are fully convinced and hold their belief means unanimously at the news value, given any q , more extreme news suggests higher expectation about the state, creating a stronger incentive for the informed agents to share the news to persuade their followers. As a result, at the equilibrium the scale of news spread has to be commensurate with the news magnitude. Moreover, for news with moderate magnitude (i.e., $|x| \in (C, \frac{C}{1-q^G(\lambda)}]$), the spread size q^* is on the exact level that makes all of informed agents feel indifferent between sharing or not, and the threshold sharing strategy translates to the fraction of sharing P_F^* that sustains the spread size. The above argument for the monotonicity of spread size in news magnitude in fact holds for any level of credibility.

Lemma 3.4 *For any β, λ and F , fixing the sign of news, the equilibrium spread size q^* and the sharing fraction P_F^* are both increasing in news magnitude $|x|$. Moreover, the equilibrium threshold $\mu^{th*}((x, \beta), F)$ for right-leaning (left-leaning) news is decreasing (increasing) in news magnitude.*

Besides the monotone effect of news magnitude, we also demonstrate that higher network connectivity enlarges the scale of diffusion for any news and any credibility.

Lemma 3.5 *Given x, β and F , the equilibrium spread size q^* is increasing in network connectivity λ .*

More social connections in the population make the shared news reach to more followers on average, facilitating news diffusion and increasing the scale of spread. However, higher network connectivity in fact yields a smaller fraction of sharing at the equilibrium. It is because, in reaction to the ease of news spread, the informed agents are less motivated to share the news due to the strategic substitutability arising from the cost of sharing.

Though the spread size is increasing both in news magnitude and in network connectivity, we observe that the mechanisms are fundamentally different. For the former,

informed agents are incentivized to share extreme news since they have large expectations on the state and the more extreme news can persuade more followers to flip their votes, resulting in a larger fraction of sharing. In contrast, high network connectivity facilitates news spread via more social interactions while discouraging informed agents from sharing. In the next section we will see how such distinct mechanisms affect the relative cascade size of news spread for different levels of credibility.

3.4 News Cascade, Scale of Spread, and Credibility

A more intriguing question is how the credibility level may affect the size of news diffusion in the population. By Theorem 3.2 and the continuity of equilibrium spread size at $\beta = 1$, we can identify whether fully credible news ($\beta = 1$) maximizes the size of a news cascade or lower credibility can otherwise induce a larger news cascade. The following proposition addresses the key question:

Proposition 3.1 *Fix F and λ , and denote the size of the giant component by $q^G(\lambda)$. Then,*

- (i) when $|x| \leq C$, no informed agents share fully credible news x and the cascade size is minimized at $\beta = 1$ with value 0;*
- (ii) when $|x| > \frac{C}{1-q^G}$, each informed agent shares fully credible news x and the cascade size is maximized at $\beta = 1$ with value q^G ;*
- (iii) when $|x| \in (C, \frac{C}{1-q^G}]$, for which agents at equilibrium feel indifferent between sharing the fully credible news or not at the equilibrium, less credible news results in a larger cascade if and only if the equilibrium threshold for sharing $\mu^{\text{th}*}$ for*

$\beta = 1$ satisfies¹²

$$\begin{aligned}\mu^{\text{th}*}((x, 1), F) &> x, \text{ when } x > 0, \\ \mu^{\text{th}*}((x, 1), F) &< x, \text{ when } x < 0.\end{aligned}\tag{3.13}$$

where

$$\mu^{\text{th}*}((x, 1), F) = \begin{cases} F^{-1}\left(1 - \frac{-\ln(\frac{C}{|x|})}{\lambda(1-\frac{C}{|x|})}\right), & \text{when } x > 0; \\ F^{-1}\left(\frac{-\ln(\frac{C}{|x|})}{\lambda(1-\frac{C}{|x|})}\right), & \text{when } x < 0.\end{cases}$$

For the ease of exposition, we discuss the implications of Proposition 3.1 and the ensuing results for right-leaning news $x > 0$; the discussion for left-leaning news $x < 0$ is similar. We focus on the range of news with moderate magnitude (cf. case (iii)). Indeed, the high connectivity that is usually observed on empirical social networks (cf. footnote 2) suggests that the giant component size $q^G(\lambda)$ is nearly 1.¹³ This implies that even when the cost of sharing is small, $\frac{C}{1-q^G(\lambda)}$ can still be significantly large and the interval $(C, \frac{C}{1-q^G(\lambda)}]$ can cover a wide range of news magnitude of interest.

The rationale for the condition (3.13) in case (iii) is as follows. Fix the spread size $q^* = 1 - \frac{C}{|x|}$ that corresponds to the case $\beta = 1$ while reducing credibility β from 1. According to the decision rule of sharing (3.8), the effective persuasive power of news remains nearly the same as $\frac{1}{2}$, suggesting that informed agents with perspectives greater than x strictly prefer to share (β less than 1). When the condition $\mu^{\text{th}*}((x, 1), F) > x$ holds, more informed agents are incentivized to share the less credible news, meaning that spread size q^* can not be sustained and the scale of news increases.

According to Lemma 3.4, the equilibrium threshold $\mu^{\text{th}*}((x, \beta), F)$ is decreasing

¹²Hsu et al. (2021) assumed F to be a normal distribution with zero mean. Under this specific setting the corresponding version of condition (3.13) is also necessary for $\beta = 1$ not to be the global maximizer of the cascade size. However, the authors did not further quantify the effects of x, λ and σ_μ on this sufficient and necessary condition.

¹³Using the empirical average degree $\lambda = 35.2$ for Poisson network, the giant component is computationally 1.

in news magnitude of right-leaning news. This, together with case (i), (ii) and the condition (3.13), implies that fully credible news can not spread wider than lower credibility news if the news is not sufficiently extreme:

Proposition 3.2 *Given λ , F , and the sign of news, the spread of lower credibility news is not smaller than fully credible news for news with sufficiently small magnitude. Specifically, there exist two unique thresholds $\underline{x}(\lambda, F) < 0 < \bar{x}(\lambda, F)$ such that the spread of lower credibility news is larger than fully credible news if and only if $x \in (\underline{x}, \bar{x})$.*

When news is small or almost neutral, high credibility will concentrate the beliefs of informed agents to be around small positive values. As agents find low expected values for the state, only those agents with quite extreme perspectives (who hold relatively extreme updated belief mean after seeing the news) share the news, yielding a large threshold $\mu^{\text{th}*}((x, \beta), F)$. On the other hand, when the credibility is reduced, informed agents are less convinced and their updated expectations about the state are more dispersed compared to observing fully credible news. Despite the compromised persuasiveness of news due to lower credibility that discourages agents from sharing, there are more informed agents who still perceive high expected values for the state and have strong incentives to share, propelling the spread of news with lower credibility. This finding echoes the result of Acemoglu et al. (2021), which suggested that reduced extremism of news increases the virality of misinformation.

From Theorem 3.2 we observe that given news $x > C$ ($x < -C$) with $\beta = 1$, the equilibrium sharing fraction P_F^* is decreasing in λ and the equilibrium threshold $\mu^{\text{th}*}$ is increasing (decreasing) in λ . Specifically, for small connectivity λ , we have $\frac{C}{1-q^G(\lambda)} \leq |x|$ and the equilibrium sharing fraction P_F^* is 1; the fraction starts to decrease when the connectivity becomes sufficiently large such that $\frac{C}{1-q^G(\lambda)} > |x|$. This suggests that there exists some threshold $\lambda^{\text{th}}(x, F)$ such that the condition (3.13) holds if and only if $\lambda > \lambda^{\text{th}}(x, F)$. By measuring the fraction of agents who share news based on condition (3.13) for $\mu^{\text{th}*}((x, 1), F)$ and x , we can derive an indifference equation for the threshold

as

$$\underbrace{\frac{-\ln \frac{C}{|x|}}{\lambda^{\text{th}}(1 - \frac{C}{|x|})}}_{=P_F^*(x)} = \begin{cases} 1 - F(x), & \text{when } x > 0; \\ F(x), & \text{when } x < 0. \end{cases} \quad (3.14)$$

The next corollary formalizes this observation:

Corollary 3.2 *Given F , for news x with magnitude $|x| > C$, lower credibility news can trigger a larger cascade than fully credible news when the network connectivity is sufficiently high. Specifically there exists a unique value $\lambda^{\text{th}}(x, F)$ such that lower credibility news results in a cascade with larger size than fully credible news if and only if $\lambda > \lambda^{\text{th}}(x, F)$ where*

$$\lambda^{\text{th}}(x, F) = \begin{cases} \frac{-\ln \frac{C}{|x|}}{(1 - \frac{C}{|x|})(1 - F(x))}, & \text{when } x > 0; \\ \frac{-\ln \frac{C}{|x|}}{(1 - \frac{C}{|x|})F(x)}, & \text{when } x < 0. \end{cases}$$

Moreover, given the sign of news, $\lambda^{\text{th}}(x, F)$ is increasing in news magnitude $|x|$. We refer to the threshold λ^{th} as network connectivity limit.

Acknowledging the aiding effect of network connectivity on news diffusion (cf. Lemma 3.5), informed agents are discouraged from sharing the news due to the strategic substitutability; their updated expected values about the state have to be more extreme so that they will share, which suggests an increased threshold for the equilibrium sharing strategy.

However, the extent of such increment in the threshold varies significantly for different levels of news credibility: The indifference equation for the sharing threshold (cf. Theorem 3.1) shows that a small increment in spread size q (thanks to larger connectivity) will be translated to a significant increase in μ^{th} when credibility level β is high. Indeed, in this case the updated means barely allow for subjective perspec-

tives, the threshold on perspective has to be much more extreme to result in large enough updated belief means that motivate sharing. In contrast, when observing news with lower credibility than full level, agents with large right-leaning perspectives still hold large updated belief means and would like to share the news. Consequently, on a densely-connected network, lower credibility news triggers a cascade that outsizes the spread of fully credible news.

Corollary 3.2 suggests that the connectivity limit $\lambda^{\text{th}}(x, F)$ serves as a natural index for the level of susceptibility of a population to wider spread of lower credibility news than credible news: The network connectivity can not surpass the threshold λ^{th} or lower credibility news can generate a larger cascade than fully credible news. In the next section, we evaluate the effects of perspective CDF F on the connectivity limit, studying how the distribution of perspectives may facilitate spread of news with lower credibility.

3.5 News Cascade in Populations with Polarized Perspectives

In this section, we study the societies comprising two groups of agents who exhibit polarization in their perspectives on the state. We investigate how the extent of polarization and the diversity of perspectives within each group may affect the spread of news for different levels of credibility.

Formally, we assume that each agent i is equally likely to orient her perspective toward either left or right side of perspective spectrum, which is her private information and denoted as $e_i \in \{-1, +1\}$. We refer to the group of agents with orientation -1 ($+1$) as left-wing (right-wing). Agent i 's perspective is then independently drawn from

a distribution that conditional on her orientation e_i :

$$\mu_i|e_i \sim \begin{cases} \frac{1}{\sigma_\mu} g\left(\frac{\mu+\bar{\mu}}{\sigma_\mu}\right), & \text{when } e_i = -1; \\ \frac{1}{\sigma_\mu} g\left(\frac{\mu-\bar{\mu}}{\sigma_\mu}\right), & \text{when } e_i = +1, \end{cases} \quad (3.15)$$

where g is some continuous density function, $\bar{\mu} \geq 0$ measures the *polarization* between the two groups, and $\sigma_\mu > 0$ represents the diversity of perspectives within each group. We will also refer to σ_μ as *in-group diversity*. Similar to Assumption 2, we impose some regularity conditions on g :

Assumption 4 *The density function g is continuously-differentiable, log-concave, and symmetric around zero, with full support on the real line, with mean and variance. We denote as G the associated CDF for g .*

The CDF F of perspectives given polarization $\bar{\mu}$ and in-group diversity σ_μ can be expressed as

$$F(\mu) = \frac{1}{2}G\left(\frac{\mu + \bar{\mu}}{\sigma_\mu}\right) + \frac{1}{2}G\left(\frac{\mu - \bar{\mu}}{\sigma_\mu}\right). \quad (3.16)$$

Note that $F(0)$ is a constant $\frac{1}{2}$ and also $F(-\mu) = 1 - F(\mu)$ for any μ . The distributions (3.15) and (3.16) are common knowledge.

From Proposition 3.1, we can see that lower credibility ($\beta < 1$) can trigger a larger cascade than fully credible news ($\beta = 1$) only if the news has sufficiently small magnitude ($|x| < C$) or the network is sufficiently connected ($\lambda > \lambda^{\text{th}}(x, F)$). Below we investigate the effects of the parameters $\bar{\mu}, \sigma_\mu, G$ on the connectivity limit, using the more explicit notation $\lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G)$. We rewrite the connectivity limit given in (3.14), for any news x , as

$$\lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G) = \frac{-\ln \frac{C}{|x|}}{\left(1 - \frac{C}{|x|}\right)F(-|x|)} = \frac{-\ln \frac{C}{|x|}}{\left(1 - \frac{C}{|x|}\right)\left(\frac{1}{2}G\left(\frac{-|x|+\bar{\mu}}{\sigma_\mu}\right) + \frac{1}{2}G\left(\frac{-|x|-\bar{\mu}}{\sigma_\mu}\right)\right)}. \quad (3.17)$$

The first result concerns the impact of polarization on the network connectivity limit:

Proposition 3.3 *For any density g , in-group diversity σ_μ , and any news x with magnitude $|x| > C$, increased polarization decreases network connectivity limit, i.e., λ^{th} is decreasing in $\bar{\mu}$. Moreover, $\lambda^{\text{th}}(x, \bar{\mu} = 0, \sigma_\mu, G) = \frac{-\ln \frac{C}{|x|}}{(1 - \frac{C}{|x|})G(-\frac{|x|}{\sigma_\mu})}$.*

Stronger polarization in a population leads to a larger fraction of extremists, who, as we have already argued, assume a higher value for the persuasiveness effect of the news aligned with their large belief mean when $\beta < 1$. Though stronger polarization reduces the fraction of the extremists who have the opposite ideological leaning from their group (e.g., right-leaning agents in left-wing group), such extremists only account for a small fraction (on the tail of the perspective distribution of left-wing group). As a result, the fraction of extremists overall increases, facilitating the sharing of lower credibility news. The news with lower credibility can hence trigger a larger cascade than fully credible news even on the networks with smaller connectivity, suggesting a decreased connectivity limit.

Figure 3-1 illustrates the effect of polarization $\bar{\mu}$ and in-group diversity σ_μ on network connectivity limit λ^{th} when function g is the standard normal distribution. Given any level of in-group diversity, we can see that the network connectivity limit decreases as the polarization is increased.

On the other hand, the effect of increased in-group diversity on the network connectivity limit is not necessarily monotone, as stated in the following proposition:

Proposition 3.4 *For any density g (and corresponding G) and any news x with magnitude $|x| > C$,*

- *when polarization is weak, increased in-group diversity decreases the network connectivity limit. Specifically, when $\bar{\mu} \leq |x|$, λ^{th} is decreasing in σ_μ and moreover, $\lim_{\sigma_\mu \rightarrow 0} \lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G) = +\infty$;*

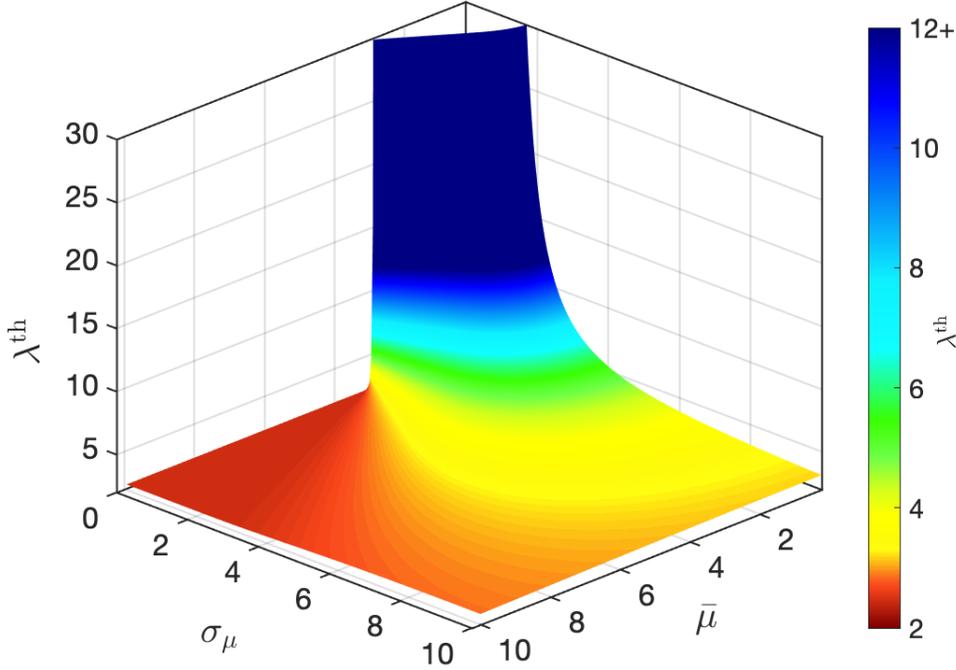


Figure 3-1: The network connectivity limit $\lambda^{\text{th}}(x, \bar{\mu}, \sigma_{\mu}, G)$ as a function of polarization $\bar{\mu}$ and in-group diversity σ_{μ} when $C = 2, x = 3$, and G is the CDF for standard normal distribution. The colormap indicates the value of connectivity limit.

- when polarization is strong, increased in-group diversity increases the network connectivity limit only when the in-group diversity is small. Formally, when $\bar{\mu} > |x|$, λ^{th} is increasing in σ_{μ} up to some $\hat{\sigma}(x, \bar{\mu}, G) > 0$ and then decreasing. Moreover, $\lim_{\sigma_{\mu} \rightarrow 0} \lambda^{\text{th}}(x, \bar{\mu}, \sigma_{\mu}, G) = \frac{-2 \ln \frac{C}{|x|}}{1 - \frac{C}{|x|}}$.

In a population with weak polarization and low in-group diversity, the agents' perspectives are highly concentrated around the neutrality with a small fraction of extremists. Increasing the in-group diversity from a low level instead results in more agents with extreme perspectives for both ideological leanings, lowering the network connectivity limit. Figure 3-1 exemplifies that given some polarization level $\bar{\mu} \leq x$ the network connectivity limit approaches the infinity when the diversity is nearly zero and that the connectivity limit is decreasing in diversity.

On the other hand, in a population with strong polarization, increased in-group diversity has a non-monotone effect on the connectivity limit. In this case, there is

a large fraction of extremists and enlarged diversity affects the fraction of extremists in two ways: (i) decreases the number of the extremists whose perspectives align with their group’s orientation, while (ii) *increases* the number of the extremists whose perspectives are opposing to their group’s orientation and are on the tail of the perspective distribution of the group.

When there is low in-group diversity in the population, the decrease from (i) out-sizes the increase from (ii), reducing the fraction of extremists and raising up the connectivity limit. For example, with fixed $\bar{\mu} = 5$ in Figure 3-1, λ^{th} increases from 2.5 to around 4 as σ_{μ} is increased from 0 to 5. However, due to the log-concavity of the perspective distribution (i.e., g), the relative magnitude of effect (i) to (ii) is decreasing in diversity and consequently the network connectivity limit starts to fall when the diversity becomes sufficiently large. We can see in Figure 3-1 that given $\bar{\mu} = 5$, λ^{th} is decreasing for $\sigma_{\mu} > 5$.

3.6 Discussion

In this chapter, we characterized the Bayesian Nash equilibrium of the news-sharing game, identifying individual decision rule of sharing and the resulting endogenous news cascade. We further delineated the sufficient and necessary conditions under which news with lower credibility can trigger a larger spread than fully credible news. As our main result, we showed that low credibility news can result in a larger cascade than fully credible news when the network is highly connected. Moreover, we investigated news spread in polarized populations. We found that strong polarization relaxes the connectivity limit necessary for lower credibility news to generate a larger cascade. We also illustrated that increased diversity of beliefs within each ideological group does not necessarily reduce the spread of news with lower credibility: When the polarization is weak or when both groups have quite diverse perspectives, raising up the level of in-group diversity can instead prompt more sharing of lower credibility news.

We delineate a few comparisons between this game-theoretic model and the non-strategic model in Chapter 2 as below.

In both models, each informed agent does not inspect news and updates her belief by placing real-valued weights on the news and her own prior mean according to the news generation process.

Though both models feature persuasion motives underlying agents' sharing decisions, there is a difference in agents' preferences. In the non-strategic model, each agent desires to make her followers beliefs close to hers according to some distance measure on the space of beliefs. By contrast, in this game-theoretic model each agent prefers aligned binary and the marginal change of her expected utility is increasing in the magnitude of her expectation about the state. This variation in utility formulation leads to a difference in the effect of diversity on news spread between the two models. In Chapter 2, we showed that $\beta = 1$ maximizes the cascade probability since it can significantly concentrate the beliefs of followers with diverse perspectives. However, in this chapter, increased diversity suggests a large fraction of extremists in a population who would share less credible news given their strong incentives to persuade others.

The strategic model makes tractable the analysis on spread size. In Theorem 2.1 we can see that the spread size at the steady state is an integral of the incremental spread size at each time. By contrast, strategic agents speculate on the cascade size at the steady state, of which the equation can be formulated without involving an integral over time, as shown in Theorem 3.1.

Also, in order for tractable analysis on spread size, in this chapter we particularly considered Poisson degree distribution for our network setting. On the other hand, the non-strategic model used the configuration model for the construction of random networks, considering a general joint degree distribution in the population. Nonetheless, its analysis is restricted to the cascade probability given the intractability of spread size in this model.

Chapter 4

Subscription Networks, Verification, and Media Bias

4.1 Background

One key assumption that was made in chapters 2 and 3 was that individuals who do not receive any news are unaware of its existence. While this is a reasonable assumption in certain scenarios such as information spread on Twitter, it certainly is not true in all cases. People sometimes do look for new information for their interests or understanding of the world. For example, when there is a developing story about a natural disaster or a political turmoil, people usually rely on news media or print publishers for the latest information that is not easily accessible by the general public. In this chapter, we turn to people's behavior of seeking information from news media.

In contrast to a limited number of trusted news agencies available a few decades ago, the advance of digital technology and popularity of social media have boosted online news business, fostering a wide range of information sources for news consumers (Newman et al. 2019). However, news consumers can only selectively attend to a few of them owing to their limited quantity of mental effort during a short time span (Kahneman 1973).

Different segments of the society then may consume quite distinct pieces of news from different news intermediaries, without a set of commonly accepted facts or truths (Mitchell et al. 2014, Jurkowitz et al. 2020). News intermediaries often have their own ideological biases and are motivated to influence the public opinion. For example, they can choose to promote news that can persuade their subscribers towards their own political agenda (Allcott and Gentzkow 2017) or conceal the unfavorable information that otherwise moves their readers' beliefs away. With diverging information, in a fragmented society the ideological divide among the population is being entrenched.

Despite their selective news disclosure to persuade the public, news intermediaries play a vital role of fact-checking the information before any news release, based on their journalists' expertise and web of sources. The importance of fact-checking information can not be emphasized enough: Nowadays digital social platforms have become a hotbed for digital misinformation to spread wide and fast (Del Vicario et al. 2016, Vosoughi et al. 2018). Without fact-checking or significant third-party filtering beforehand, the inundation of misinformation or false news poses threats to our society, inciting political polarization and social unrest (Howell 2013).

In this chapter, we investigate the subscription choices of subscribers with diverse perspectives between biased news intermediaries who are motivated to influence the public opinion, through the lens of news verification and selective disclosure in an environment with real and false news. We shed light on three questions: (i) What news does an intermediary choose to verify? (ii) How do intermediaries' disclosure decision depend on their knowledge of the veracity of the news? (iii) How do subscribers react to disclosed news while inferring its veracity and, especially, what do they infer when there is no disclosure?

Our model concerns two biased intermediaries (she) and a continuum of subscribers (he). There is an unobservable real-valued state that is subject to a vote, in which each subscriber will simultaneously take a binary action. The subscribers share the same preference, aiming to match the sign of the state, while the intermediaries are both

concerned with the aggregate utility of the subscribers. The players, either intermediaries or subscribers, have heterogeneous subjective priors on the unobservable state, and hence ex-ante their favored actions differ. Prior to the vote, only the intermediaries may access a piece of information about the state and can choose whether to disclose it to their own subscribers. In order to make an informed decision in the vote, each subscriber initially selects exactly one of the intermediaries, hoping to receive information about the state from the selected intermediary. We refer to the subscription choice of a subscriber as (weak) homophily when the subscriber (weakly) prefers the intermediary who ex-ante advocates for the same binary action; otherwise, we say it is anti-homophily.

The information a news intermediary may receive is not perfectly informative about the state: the news can be informative (true) or uninformative (false) about the state: Informative news is some noisy observation about the state with a noise distribution that is ex-ante known to all the players. On the other hand, uninformative news is independent from the state. Only the intermediaries can verify whether the news they observe is informative or not, if they spend a cost for verification. After receiving some news and before deciding whether to disclose it to her subscribers, a news intermediary may choose to verify it at a cost. Each subscriber only observes his news intermediary's disclosure decision, either some news or no disclosure, but he cannot observe his intermediary's verification decision and result. Based on his prior belief and his news intermediary's strategy, the subscriber takes the action that matches the sign of his posterior mean about the state.

As our first contribution, we propose a game-theoretic framework to study the strategic communication between news media and subscribers in an environment with true and false news, examining how the media's function of filtering and verification affect subscription choices. We identify the equilibrium structure, characterizing news intermediaries' disclosure and verification strategies and subscribers' belief updating. We show that each intermediary has no incentives to verify the news that is aligned

with her perspective but discloses such news for sure. This suggests that only the news that is against her perspective may be concealed by the intermediary, and her subscribers will update their own belief towards the ideologically opposite wing when there is no disclosure by the intermediary.

As our second result, we demonstrate how the intermediaries' option to selectively disclose news leads to *anti-homophily* among the centrists. We investigate the particular case in which the verification option is absent and the intermediaries only make their disclosure decisions. The subscribers care about the news that may be concealed by their intermediaries out of the persuasion motive. For the centrists, if they follow the intermediary with aligned ideology, the possibility of news being concealed makes them strategically choose the action opposite from their perspective when there is no disclosure from this intermediary, suffering a loss from not matching the sign of the state when there is actually no news. However, if the centrists subscribe to the opposite intermediary, the concealed news by the intermediary does not change their preferred actions and any news that may change their actions is disclosed, leading to higher expected payoffs for the centrists. On the other hand, the extreme subscribers with stronger ideological bias have no concern about news concealment by the intermediary with aligned ideology, since the news that can change their favored actions also changes the intermediary's favored one, making the intermediary disclose it to persuade her subscribers.¹

We further demonstrate that when the intermediaries can verify news, in the *true-news-disclosing* equilibria in which the intermediaries disclose the verified news if and only if it is true, weak homophily emerges among the extremists whereas the centrists still make anti-homophilic subscription choices. Each intermediary verifies the news only if the news can flip her favored action to the opposite when it turns out to be informative. In order to signal the verification result, which is unobservable to her

¹These results subsume our preliminary work in Hsu et al. (2020b), which fully focused on an environment with no false news and assumed that prior beliefs and noise in informative news follow Gaussian distributions.

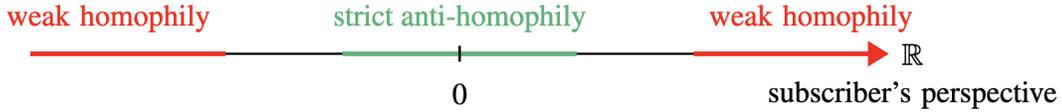


Figure 4-1: The subscription choices of the centrist (green) and extremist (red) subscribers.

subscribers, the intermediary chooses to disclose the news if and only if it is informative. Though the intermediary with aligned ideology verifies news that may affect the centrists' actions, the necessary signalling for the verification result makes the centrists unable to distinguish false news from concealed unverified news when observing no disclosure from the intermediary. The centrists then end up not taking their optimal actions when the news is false, deriving no benefits from the verification strategy. By contrast, for the extremists, some of the news that flips their favored actions when it is informative will be verified and disclosed by the intermediary with aligned ideology, who they therefore weakly prefer to follow.

We consider a variation of our model in which the verification result for any news is perfectly observable to subscribers when it is disclosed. This perfect observability creates an incentive for the intermediaries to verify the news in line with their perspectives so as to persuade their subscribers with the opposite ideology. The equilibrium subscription choices depend on the relative strength of the intermediaries' motives between learning the veracity of opposing news and verifying aligned news for persuasion purpose. We provide an example to show that when the players think the aligned news is highly likely to be true but the opposing news is likely to be false, the intermediaries are only incentivized to verify the aligned news for persuasion, leading to anti-homophily among the subscribers.

Our work joins the recent flourishing literature on the political economy of news market (Prat and Strömberg 2013, Anderson et al. 2016), studying both supply-side and demand-side forces behind media bias. The news intermediaries strategically decide what news to verify and to disclose in order to sway their subscribers' opinions: On the supply side they have their own political agendas while on the demand side the

subscribers possess heterogeneous subjective prior beliefs.²

Our work focuses on how decision makers select their information sources while accounting for the media bias, rationalizing how *homophily* and *anti-homophily* may arise among different segments of the subscribers. We share the result of homophily with the following papers. Calvert (1985) proposed a model with two information sources and two alternatives as in ours and suggested that a decision maker with a bias in preference on the alternatives would prefer the information source in favor of her bias rather than the unbiased one. Banerjee and Somanathan (2001) showed that the experts may conceal their information if the prior beliefs of the experts and the decision maker differ. Suen (2004) suggested that homophily emerges when there are binary actions and the leaders can only indicate what action to take rather than disclose the exact information they have. Similarly, Sethi and Yildiz (2019) obtained homophily when the information sources disclose their posteriors non-strategically, and the bias of the sources are unknown.

On the other hand, Che and Kartik (2009) considered a one-leader-one-follower model with a continuum of actions and the leader aims to persuade the follower to take her favored real-valued action by endogenously acquiring information and selectively disclosing the information. They demonstrated that when a follower can choose from a list of leaders with diverse opinions, a leader with some difference in opinion is preferred to a like-minded one. Jackson and Tan (2013) studied a similar model of strategic information transmission with binary signals. They do observe anti-homophily with the same intuition that an expert with an opposite bias in prior belief shares the information that changes one's favored action whereas an expert with an aligned bias does not. We obtain the subscription network formation in a slightly richer model (with continuum of signals). The network formation prior to the news intermediaries' reporting decisions

²A nonexhaustive list for papers on supply side includes Balan et al. (2004), Baron (2006), Germano (2008), Chan and Suen (2009), Ellman and Germano (2009), Anderson and McLaren (2012). For a list of work on demand side, see Strömberg (2004), Mullainathan and Shleifer (2005), Gentzkow and Shapiro (2006), Chan and Suen (2008), Duggan and Martinelli (2011). See Gentzkow et al. (2015) for more details on media and review of the related theoretical literature.

adds substantial complexity to the problem, since subscription choices are endogenously determined with the leaders' strategies in the equilibrium.

A key component distinguishes our model from the above work is the coexistence of true news and false news and the news intermediaries' option to learn the veracity of news. Many recent works have studied social media users' behavior in the environments with coexistence of real news and false news (Papanastasiou 2020, Merlino et al. 2020, Acemoglu et al. 2022, Hsu et al. 2022). In our model only the intermediaries are able to fact-check news and decide whether to disclose the news. In this regard, our paper is related to Kranton and McAdams (2020), who studied whether news producers' make efforts to create true news in order to earn revenue from views or sharing, resulting in the distribution of false and true news. In contrast, our work concerns news intermediaries who verify news to learn the veracity of news and persuade their subscribers to have aligned ideological stance.

The rest of the chapter proceeds as follows. We describe our model and discuss the assumptions in Section 4.2. Section 4.3 explains the belief updating process. In Section 4.4 and 4.5 we study subscription choices when the verification option is unavailable or available to the intermediaries respectively. Section 4.6 examines how subscription choices change with the intermediaries' ideological biases. Section 4.7 discusses a few extensions of our model and Section 4.8 summarizes our findings. The technical proofs of all lemmas and theorems are relegated to the Appendix C.

4.2 Model

The players are two news intermediaries $i \in \{L, R\}$ (she) and a unit-mass continuum of *anonymous* subscribers $j \in [0, 1]$ (he) with Lebesgue measure λ . Each subscriber j is to take a binary action $a_j \in \{-1, +1\}$ and his payoff depends on an unobservable

state $\theta \in \mathbb{R}$. Specifically, his payoff from taking action a_j is

$$u_j(\theta, a_j) = \theta a_j,$$

so that he aims to match the sign of the state. We write a for the action profile of all subscribers. Before taking an action, subscribers are able to learn about the state from a news intermediary. The payoffs of intermediaries are simply the average payoff of the subscribers:

$$u_i(\theta, a) = \int_{j \in [0,1]} u_j(\theta, a_j) d\lambda.$$

for all measurable action profiles a .

Prior beliefs and perspectives. The players have heterogeneous subjective prior beliefs about the state θ . In particular, the subjective prior belief of player i is

$$f_i(\theta) = f(\theta - \mu_i), \tag{4.1}$$

where f is a continuous density function associated with zero expected value and $\mu_i \in \mathbb{R}$ is the a priori expected value of θ according to player i . We refer to μ_i as the *perspective* of player i . The players' prior beliefs differ only at their expectations (perspectives) on the state as shown in (4.1). We assume that each player's prior belief and perspective are common knowledge.³ Observe that, although players would have identical preferences if they knew the state, they have differing subjective beliefs, and hence they may have differing preferences on optimal action when they are uncertain. In particular, ex ante, each player k prefers action $+1$ if $\mu_k > 0$ and prefers action -1

³Since the subscribers are anonymous, each subscriber in fact only needs to know his own μ_j and the perspectives of the intermediaries μ_L, μ_R .

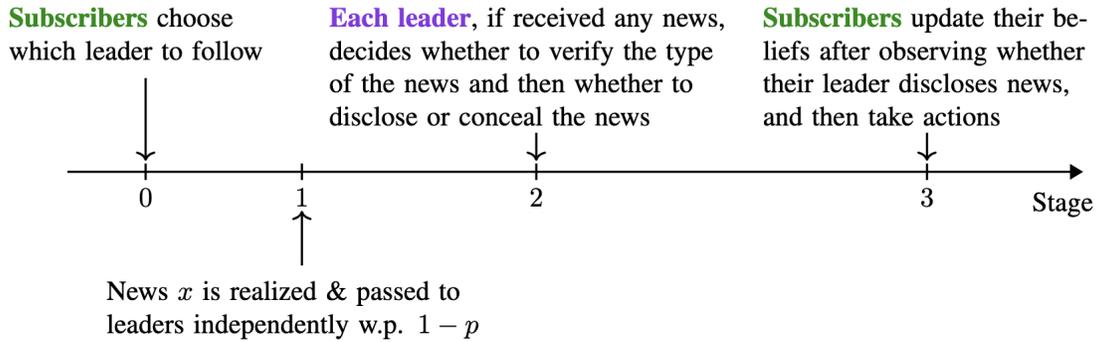


Figure 4-2: Timeline of the subscription-verification-disclosure game described in Section 4.2

if $\mu_k < 0$. The intermediaries' perspectives are

$$\mu_L < 0 \text{ and } \mu_R > 0,$$

so that they have opposing preferences ex-ante.

We denote the distribution of subscribers' perspectives induced from measure λ as a CDF $G_0 : \mathbb{R} \rightarrow [0, 1]$; we assume that G_0 is strictly increasing with range $(0, 1)$ so that the density of followers' perspectives is non-zero everywhere on the real line. We also write \mathcal{G} as the space of measure functions that map \mathbb{R} to $[0, 1]$ and are strictly increasing (note that not all elements of this space are CDFs).

Timeline. The game proceeds in four stages 0, 1, 2, 3 as depicted in Figure 4-2. Each stage will be described in detail below. Here we summarize. Stage 0 represents the subscription stage, at which each subscriber subscribes to an intermediary. At stage 1, Nature generates a piece of news, which may be informative about the state, and privately discloses it to some of the intermediaries (with some probability). At stage 2, intermediaries move. If an intermediary observes the news, then she decides whether to verify it and whether to disclose it to her own subscribers. At stage 3, each subscriber observes whether and what his intermediary discloses, updates his belief about the state, and takes an action. All of these are common knowledge.

Subscription choices. At stage 0, each subscriber simultaneously makes a subscription choice to follow exactly one of the two intermediaries at zero cost. For each subscriber j , we denote his subscription with $s_j \in \{L, R\}$, which is his private information. For any relationship $s_j = i$, we refer to i as subscriber j 's intermediary and to j as intermediary i 's subscriber. We write $s : [0, 1] \rightarrow \{L, R\}$ as the profile of the subscription choices. For simplicity of analysis, we assume that a $\delta > 0$ fraction of the subscribers are "noise subscribers", in that they are randomly assigned to one of the intermediaries (with equal probabilities); the distribution of the noise subscribers' perspectives is the same as the distribution of the perspective for overall population.

The subscription profile s and the noise subscribers induce a measure on the subscribers' perspectives for each intermediary $i \in \{L, R\}$. Specifically let $G_i^s(\mu)$ denote the fraction of all subscribers who subscribe to intermediary i and have perspectives no more than μ .⁴

$$G_i^s(\mu) = \frac{\delta}{2}G_0(\mu) + (1 - \delta)\lambda(\{j \in [0, 1] | s_j = i, \mu_j \leq \mu\}),$$

where the term $\frac{\delta}{2}G_0(\mu)$ accounts for the noise subscribers and the second term accounts for the subscribers who make their own subscription choices. We assume that the measure functions G_L^s and G_R^s are public information to all the players; we denote this common history at the end of subscription stage by $G^s = (G_L^s, G_R^s)$.

News and intermediary types. With probability $1 - p$, each intermediary i privately receives a piece of *news*, a real number $x \in \mathbb{R}$, generated as follows. There are three random variables: *true news* x_T that is informative about the state θ , *false news* x_F that is independent of the state θ , and a binary variable $\omega \in \{T, F\}$ that indicates the *veracity* of the news. The binary variable ω is drawn according to Bernoulli distribution

⁴As each subscriber chooses exactly one intermediary, the induced measures satisfy $G_R^s(\mu) + G_L^s(\mu) = G_0(\mu)$ for any $\mu \in \mathbb{R}$ and any subscription profile s .

$\mathbb{P}(\omega = F) = q$ with $q < 1$. The news is determined as:

$$x = \begin{cases} x_T, & \text{if } \omega = T, \\ x_F, & \text{if } \omega = F. \end{cases}$$

When $\omega = T$, the news $x = x_T$ is informative about the state; the news $x = x_F$ is uninformative when $\omega = F$. We will use the terms true and informative (resp., false and uninformative) interchangeably. When an intermediary i observes a piece of news x , she does not know whether it is true (i.e. $x = x_T$) or false (i.e. $x = x_F$). She would have updated her belief about the state if she knew that it is true, and she would not update it if she knew that it is false. We assume that the probability of receiving a piece of news does not depend on the realization of (x_T, x_F, ω) . We set an intermediary i 's type as $t_i = x$ if she observes x for some $x \in \mathbb{R}$; we set her type as $t_i = \emptyset$ if she does not observe any news. Later in the game, the intermediaries may also garner additional private information as we will see below.

Verification and disclosure. Upon observing her type $t_i = x$, an intermediary i decides whether to verify the news at a verification cost $C \geq 0$. Her verification decision is denoted by $v_i \in \{0, 1\}$ where $v_i = 1$ means verifying; v_i is a function of x and the subscription distributions G^s . If she verifies the news, she learns its veracity ω . Type $t_i = \emptyset$ can only choose $v_i = 0$. The resulting histories are denoted by $(G^s, t_i, 1, \omega)$ or $(G^s, t_i, 0)$, depending on whether she verifies the news or not, respectively.

Next, intermediary i sends a message $m_i \in \{t_i, \emptyset\}$ to all of her subscribers; we denote the message space by $\mathcal{M} \triangleq \emptyset \cup \mathbb{R}$. When $t_i = x$, she has the choice of disclosing the news by choosing $m_i = x$ or concealing it by choosing $m_i = \emptyset$. When $t_i = \emptyset$, she can only send a message $m_i = \emptyset$. Note that when $t_i = x$, m_i is a function of the histories above.

Subscribers' actions. Each subscriber j observes only G^s , his subscription s_j and

his intermediary's message m_{s_j} . If $m_{s_j} = x$, then subscriber j knows that intermediary s_j received the news x . If $m_{s_j} = \emptyset$, subscriber j observes that no news is disclosed and cannot distinguish whether s_j conceals information or has no information. Given his private history $h_j = (G^s, s_j, m_{s_j})$, subscriber j updates his belief about the state using Bayes' rule and then simultaneously takes a binary action a_j aiming to match the sign of the state. We write $\pi_j(h_j)$ for his posterior expectation of θ and $a_j(h_j) \in \{-1, +1\}$ for his action.

Beliefs and information structure. Informative news is a noisy observation of the state; specifically, we assume that $x_T = \theta + \epsilon$ where ϵ is an additive, zero-mean, and independent continuous random noise with continuous density function f_ϵ . From player i 's perspective, the ex-ante joint probability density of θ and informative news x_T is then given by $f(\theta - \mu_i)f_\epsilon(x_T - \theta)$. We then find the ex-ante marginal probability density of informative news from player i 's perspective:

$$f_T(x; \mu_i) \triangleq \int f(\theta - \mu_i)f_\epsilon(x - \theta)d\theta.$$

On the other hand, uninformative news x_F is generated according to a continuous density function f_F . Throughout this paper, we make the following regularity assumptions on the density functions f , f_ϵ , and f_F ; recall that f is the density function of θ for $\mu_i = 0$.

Assumption 5 *The density functions f , and f_ϵ, f_F are absolutely integrable, symmetric around zero, have full support on the real line, and continuously differentiable almost everywhere. The density f and f_ϵ are also assumed to be log-concave. We assume that for any perspective $\mu_i \geq 0$ (resp. informative news $x \geq 0$), there exists some informative news $x < 0$ (resp. some perspective $\mu_i < 0$) such that for any $\theta > 0$*

$$\frac{f(\theta - \mu_i)f_\epsilon(x - \theta)}{f(-\theta - \mu_i)f_\epsilon(x + \theta)} < 1. \tag{4.2}$$

The symmetry of f around zero and its log-concavity capture a player's perception that a priori the state is more unlikely to be farther away from her perspective. The log-concavity dictates that f and f_ϵ have light tails for large values, leading to some properties of a player's posterior mean when observing informative news: (a) the posterior mean is increasing in her perspective (f); (b) it is increasing in the news value (f_ϵ); (c) it lies between her perspective and the news value. Finally, condition (4.2) means that the likelihood ratio of θ to $-\theta$ is upper bounded by 1 for any $\theta > 0$; in turn, player i observing true news x reckons that the state is more likely to be negative and holds negative posterior mean. Given the symmetry of the density functions, the last assumption in fact requires that for any player i , there must exist some true news x that can change the sign of her expectation. Moreover, it also dictates that for any true news, there must exist some player who does not change the sign of her expectation.⁵ An example is the family of Gaussian distributions, which results in a linear updating rule of one's expectation fusing her perspective and news. We will elucidate the properties of belief updating in Section 4.3.

Solution concept and terminology

Our solution concept is pure-strategy perfect Bayesian equilibrium (PBE). By convention, we drop subscripts to denote profiles; for example, we write v for $\{v_i\}_{i \in \{L,R\}}$. An assessment $(\hat{s}, \hat{v}, \hat{m}, \hat{a}, \hat{\pi})$ is a PBE if (i) it is *sequentially rational* (i.e., each player plays a best response to other players' strategies at every information set given his beliefs) and (ii) the beliefs are derived from the strategies using the Bayes rule on every equilibrium path that is reached with positive probability. We assume that each intermediary i does not disclose if she feels indifferent between disclosure or not; similarly, each intermediary i does not verify news if verifying it does not strictly improve her expected utility.

⁵Since f and f_ϵ are log-concave, the likelihood ratio can be regarded as a product of two monotone likelihood ratios. One can show that when $\mu_i \geq 0$ then (4.2) holds only if $x < 0$. Otherwise, the ratio is a product of two strictly increasing functions and surpasses 1 for any $\theta > 0$.

We identify the equilibria and discuss the implication of intermediaries' and subscribers' strategies for the following scenarios:

- Section 4.4 considers the scenario when verification is not possible (v_i is restricted to 0).
- Section 4.5 discusses the equilibria when intermediaries can choose to verify news at a cost.
- Section 4.7 discusses a few variations of our model with illustrative examples.

We use the following terminology. The action space $\{-1, +1\}$ naturally splits the space of real numbers into three ideological leanings: *right-leaning* individuals with positive prior μ , who prefer action $+1$ ex-ante, *left-leaning* individuals with negative prior μ , who prefer action -1 ex-ante, and *neutral* individuals with $\mu = 0$, who are ex-ante indifferent; the measure of neutral subscribers is zero. We say that a subscriber j exhibits *homophily* if she subscribes to the intermediary who shares his ideological leaning, i.e., $\text{sign}(\mu_j) = \text{sign}(\mu_{s_j})$; we say that she exhibits *anti-homophily* if she subscribes to the one with the opposite ideological leaning. Moreover, we say that j exhibits weak (resp. strict) homophily if she weakly (resp. strictly) prefers the intermediary who shares his own ideological leaning ex-ante. Weak and strong anti-homophily are defined similarly.

4.3 Equilibrium Structure

In this section, we identify structural properties of the communication between intermediaries and subscribers after the subscription stage and the common history G^s is observed for any pure-strategy perfect Bayesian equilibrium.

4.3.1 Belief updating and subscribers' action strategies

We first study the Bayesian updating process of a player's belief about the state given her prior belief and information set, introducing a few notions that are useful for following analysis. In particular, we will use the function $\pi_j(G^s, s_j, m_{s_j})$ to represent subscriber j 's posterior mean about the state that is updated using Bayes' rule based on his intermediary's verification and disclosure strategies.

When observing news x , a player incorporates the news into her belief on the state via Bayes' rule and using her knowledge about the veracity of the news. Specifically, for each intermediary i , she can choose to verify the news and learn its veracity. She holds one of three possible beliefs about news veracity: (a) true news ($\omega = T$); (b) false news ($\omega = F$); (c) uncertain about the news veracity. On the other hand, a subscriber j who receives news x from his intermediary (i.e., $m_{s_j} = x$) updates his expectation $\pi_j(G^s, s_j, x)$ while inferring whether the news was verified and, if so, whether it is true or false. As we focus on pure-strategy equilibria, we can find that subscriber j 's belief on the veracity of the news is also one of the three cases (a)-(c).⁶ For each case of (a)-(c), player i 's posterior mean about the state is formed as below:

(a) true news ($\omega = T$): we denote player i 's posterior expectation on the state for informative news as $\mathbb{E}_i[\theta | x, T]$, which is given by:

$$\mathbb{E}_i[\theta | x, T] = \frac{\int \theta f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta}{f_T(x; \mu_i)}. \quad (4.3)$$

(b) false news ($\omega = F$): a player i 's posterior expectation is the same as her perspective μ_i (we will sometimes denote it as $\mathbb{E}_i[\theta | x, F]$ for purpose of exposition).

(c) uncertain about the news veracity: player i weighs in the likelihood that the news is informative based on her prior on the state and the generation process of news.

⁶For subscriber j , the verification and disclosure strategies of his intermediary s_j corresponding to each case are described as follows: (a) intermediary s_j verifies news x , and discloses it if and only if it is true; (b) intermediary s_j verifies news x , and discloses it if and only if it is false; (c) intermediary s_j discloses news x regardless of his verification decision and the result.

We can find the likelihood that news x is true based on player i 's prior belief:

$$\mathbb{P}_i(T|x) = \frac{(1-q)f_T(x; \mu_i)}{qf_F(x) + (1-q)f_T(x; \mu_i)}.$$

We write $\mathbb{E}_i[\theta | x]$ as her posterior mean for this case and it is computed as

$$\begin{aligned} \mathbb{E}_i[\theta | x] &= \mathbb{P}_i(T|x) \mathbb{E}_i[\theta | x, T] + (1 - \mathbb{P}_i(T|x))\mu_i, \\ &= \mu_i + \mathbb{P}_i(T|x)(\mathbb{E}_i[\theta | x, T] - \mu_i), \end{aligned} \tag{4.4}$$

where we observe that $\mathbb{E}_i[\theta | x]$ is a convex combination of $\mathbb{E}_i[\theta | x, T]$ and μ_i .

On the other hand, when there is no disclosure, subscriber j strategically updates his belief on the state, speculating on the possibilities that his intermediary may have accessed some news but decided to conceal it. He adjusts his expectation by accounting for the possible cases in which informative news ends up being concealed by his intermediary:⁷

$$\pi_j(G^s, s_j, \emptyset) = \mu_j + \mathbb{P}_j(\exists x_T \text{ that is concealed}) (\mathbb{E}_j[\theta | \exists x_T \text{ that is concealed}] - \mu_j). \tag{4.5}$$

Below we describe several properties of a belief function that are important for our analysis; for the purpose of exposition, we denote a belief function of perspective μ_i and news x by $g(\mu_i, x)$, which can represent either function $\mathbb{E}_i[\theta | x, T]$ or $\mathbb{E}_i[\theta | x]$.

Symmetry: A belief function $g(\mu_i, x)$ is said to be *symmetric* if for any μ_i and any x

$$g(\mu_i, x) = -g(-\mu_i, -x).$$

⁷Specifically, there are three cases that can lead to no disclosure: (i) his intermediary did not receive any news; (ii) his intermediary received the news and concealed it without verification (true or false news both ended up being concealed); (iii) his intermediary received the news, verified it, and decided to conceal it after learning the news veracity.

Monotonicity in news: A belief function $g(\mu_i, x)$ is said to be monotone in news if it is strictly increasing in x for any μ_i .

Unboundedness of news: A belief function $g(\mu_i, x)$ is said to exhibit unboundedness of news if for every prior belief, there exists some news x that switches the sign of the expectation, i.e., for any μ_i , $\lim_{x \rightarrow +\infty} g(\mu_i, x) > 0$ and $\lim_{x \rightarrow -\infty} g(\mu_i, x) < 0$.

Monotonicity in perspective: A belief function $g(\mu_i, x)$ is said to be monotone in perspective if it is strictly increasing in μ_i for any news x .

Unboundedness of prior beliefs: A belief function $g(\mu_i, x)$ is said to exhibit unboundedness of prior beliefs if for any news x , there exist a positive prior and a negative prior for both of which the sign of the expectation does not change, i.e.,

$$\lim_{\mu_j \rightarrow +\infty} g(\mu_i, x) > 0 \text{ and } \lim_{\mu_j \rightarrow -\infty} g(\mu_i, x) < 0.$$

Lemma 4.1 *The function $\mathbb{E}_i[\theta | x, T]$ is symmetric, monotone in news, monotone in perspective, and satisfies unboundedness of news and unboundedness of prior beliefs properties. The function $\mathbb{E}_i[\theta | x]$ is symmetric, monotone in perspective, and satisfies unboundedness of prior beliefs property.*

The next lemma shows that when there is no disclosure, the posterior mean of a subscriber also exhibits monotonicity in his perspective and the property of unboundedness of prior beliefs.

Lemma 4.2 *Given any measurable strategies v_i, m_i of intermediary i , for any of her subscriber j (with $s_j = i$), when there is no disclosure from intermediary i , his posterior mean is increasing in his perspective μ_j , i.e., $\pi_j(G^s, i, \emptyset)$ is increasing in μ_j . Moreover,*

$$\lim_{\mu_j \rightarrow +\infty} \pi_j(G^s, s_j, \emptyset) > 0 \text{ and } \lim_{\mu_j \rightarrow -\infty} \pi_j(G^s, s_j, \emptyset) < 0.$$

Thanks to the monotonicity in perspective and unboundedness of prior beliefs as characterized in Lemma 4.1 and 4.2, the sign of a subscriber's posterior mean, whether he observes disclosure or not, can be found by comparing his perspective with a unique threshold, which is a function of the message, his intermediary, and the common history according to the intermediaries' verification and disclosure strategies. Specifically, given

any G^s and any message m_i sent by intermediary i , there exists a unique threshold, denoted by $\widehat{\mu}_i(G^s, m_i) \in \mathbb{R}$, such that for any of her subscribers j ($s_j = i$)

$$\pi_j(G^s, i, m_i) > 0 \text{ if and only if } \mu_j > \widehat{\mu}_i(G^s, m_i),$$

where the threshold $\widehat{\mu}_i(G^s, m_i)$ is the unique perspective such that $\pi_j(G^s, i, m_i) = 0$. Since each subscriber j aims to match the sign of the state based on his belief, this observation suggests a structural form for subscribers' action strategies as characterized in the following definition:

Definition 4.1 *We say that subscriber j deploys a threshold action strategy if for any private history (G^s, s_j, m_{s_j}) , there exists a threshold $\widehat{\mu}_{s_j}(G^s, m_{s_j})$ such that*

$$a_j(G^s, s_j, m_{s_j}) = +1 \text{ if and only if } \mu_j > \widehat{\mu}_{s_j}(G^s, m_{s_j}).$$

Lemma 4.3 *In any equilibrium, each subscriber j deploys threshold action strategy. Moreover, for any private history $h_j = (G^s, s_j, m_{s_j})$, the threshold perspective $\widehat{\mu}_{s_j}^*(G^s, m_{s_j})$ is the unique perspective such that $\pi_j^*(h_j) = 0$ and the action yields expected utility $|\pi_j^*(h_j)|$.*

Considering that subscribers deploy threshold action strategy, given any common history and any news, each intermediary can evaluate the fractions of her subscribers choosing +1 or -1 for her disclosure choices by using the perspective measures G^s and the thresholds. Moreover, for any news, the decision making for disclosure always involves the threshold associated with no disclosure; we especially define critical perspective as follows:

Definition 4.2 *We refer to the threshold on perspective for no disclosure by intermediary i , i.e., $\widehat{\mu}_i(G^s, \emptyset)$, as the critical perspective for intermediary i . We denote with k_i a representative subscriber of intermediary i who holds the critical perspective and call*

him critical subscriber.⁸

For each intermediary there must exist a critical subscriber, since both of the perspective measures G_L^s, G_R^s have full support for any subscription profile s due to the noise subscribers. We will alternatively write μ_{k_i} as the critical perspective. We leverage the critical perspectives for analysis of intermediaries' decision making in the next section.

4.3.2 Decision making for disclosure and verification

An intermediary i receiving news x (i.e., $t_i = x$) decides whether to verify and whether to disclose the news. Since her message $m_i \in \mathcal{M}$ can influence only her subscribers, intermediary i equivalently aims to maximize *the aggregate utility of her subscribers* while considering their threshold action strategies:

$$\mathbb{E}_i \left[\int_{j \in [0,1] | s_j = i} \theta a_j^*(G^s, i, m_i) d\lambda \mid G^s, x \right] - C \mathbb{1}_{\{v_i=1\}} \quad (4.6)$$

Though intermediary i does not observe the identity of her subscribers, she can aggregate her subscribers' utilities using measure G_i^s and threshold perspectives according to Lemma 4.3. Abusing the notation $a_j^*(G^s, i, m_i, \mu_j)$ with $s_j = i$, we can rewrite the aggregate utility in (4.6) as

$$\begin{aligned} \mathbb{E}_i \left[\int_{\mathbb{R}} \theta a_j^*(G^s, i, m_i, \mu_j) dG_i^s \mid G^s, x \right] &= \mathbb{E}_i \left[\theta [G_i^s(+\infty) - 2G_i^s(\widehat{\mu}_i^*(G^s, m_i))] \mid G^s, x \right] \\ &= \mathbb{E}_i [\theta \mid G^s, x] [G_i^s(+\infty) - 2G_i^s(\widehat{\mu}_i^*(G^s, m_i))], \end{aligned} \quad (4.7)$$

where intermediary i 's expectation about the state is conditionally independent from the aggregate actions of her subscribers, and in fact its sign determines her favored action that she wants to persuade her subscribers to take. Her expectation on the state

⁸Among a continuum of subscribers, there may exist multiple ones who have the critical perspective.

may be changed depending on her subsequent verification decision v_i and verification result. On the other hand, the fraction of her subscribers taking +1 (or -1) is a function of her message m_i . (The function value depends on her disclosure and verification decisions which affect her subscribers' inference on the space of news and veracity.) Below we analyze intermediaries' decision making and identify the properties in any equilibrium using backward induction.

Disclosure decision

If intermediary i verified news x ($v_i = 1$) and learned its veracity ω , she has posterior mean $\mathbb{E}_i[\theta | x, \omega]$ and chooses whether to disclose it so as to maximize (4.7).

$$m_i^*(G^s, x, 1, \omega) = \emptyset \Leftrightarrow \emptyset \in \operatorname{argmax}_{m_i \in \{\emptyset, x\}} \mathbb{E}_i[\theta | x, \omega] [G_i^s(+\infty) - 2G_i^s(\widehat{\mu}_i^*(G^s, m_i))].$$

On the other hand, if she did not verify it ($v_i = 0$) and holds posterior mean $\mathbb{E}_i[\theta | x]$,

$$m_i^*(G^s, x, 0) = \emptyset \Leftrightarrow \emptyset \in \operatorname{argmax}_{m_i \in \{\emptyset, x\}} \mathbb{E}_i[\theta | x] [G_i^s(+\infty) - 2G_i^s(\widehat{\mu}_i^*(G^s, m_i))].$$

Since $G_i^s(\mu_j)$ is strictly increasing in μ_j , intermediary i receiving news x effectively makes her disclosure decision by comparing $\widehat{\mu}_i^*(G^s, x)$ and $\widehat{\mu}_i^*(G^s, \emptyset)$, choosing the lower one if and only if she prefers +1. In another way, intermediary i ponders on how her critical subscriber would react to the news if observing it: If the news would sway her critical subscriber to strictly prefer the action as she prefers, intermediary i chooses to disclose it or her critical subscriber would instead feel indifferent. The following lemma formalizes this observation and further characterizes each intermediary's disclosure strategy and the ideological leaning of her critical subscriber.

Lemma 4.4 (disclosure strategy) *In any equilibrium, for any common history G^s and news x ,*

- (i) *intermediary i discloses news x if and only if she and her critical subscriber have*

the same belief sign when her critical subscriber receives the news, i.e.,

$$m_i^*(G^s, x, 0) = x \text{ if and only if } \mathbb{E}_i[\theta | x] \pi_{k_i^*}^*(G^s, i, x) > 0,$$

and for any $\omega \in \{T, F\}$

$$m_i^*(G^s, x, 1, \omega) = x \text{ if and only if } \mathbb{E}_i[\theta | x, \omega] \pi_{k_i^*}^*(G^s, i, x) > 0.$$

- (ii) if news x will not change the sign of her expectation even when it is informative, intermediary i chooses the same disclosure decision whether she learns news veracity or not, i.e., for any x such that $\mu_i \mathbb{E}_i[\theta | x, T] > 0$, then $m_i^*(G^s, x, 1, T) = m_i^*(G^s, x, 1, F) = m_i^*(G^s, x, 0)$;
- (iii) intermediary i and her critical subscriber have the same ideology, i.e., $\mu_i \mu_{k_i^*} > 0$;
- (iv) intermediary i discloses news x if it is not against her ideology, i.e., for any x such that $\mu_i x \geq 0$, her equilibrium disclosure strategy is $m_i^*(G^s, x, 1, T) = m_i^*(G^s, x, 1, F) = m_i^*(G^s, x, 0) = x$.

Part (ii) directly follows the result of part (i): If an intermediary prefers the action in line with her perspective even if the news is true, it then suggests that she prefers the action regardless of her verification decision and result (recall that $\mathbb{E}_i[\theta | x]$ is a convex combination of $\mathbb{E}_i[\theta | x, T]$ and μ_i , cf. (4.4)), and that she chooses the same disclosure action for all of the three cases. Intuitively, because the verification result cannot be observed by her subscribers, the intermediary has the incentive to deviate to make the same disclosure decision. Specifically, suppose at an equilibrium intermediary i chooses different disclosure decisions, say $m_i^*(G^s, x, 1, T) \neq m_i^*(G^s, x, 1, F)$ for example. This suggests that $\hat{\mu}_i^*(G^s, x) \neq \hat{\mu}_i^*(G^s, \emptyset)$ (otherwise, she just conceals), but intermediary i would then deviate for the case of true news or false news.

Part (iii) states that the ideological leaning of an intermediary and her critical subscriber must be aligned. The intuition is as follows. Without loss of generality,

consider intermediary R and suppose her critical subscriber ex-ante (weakly) prefers the opposite wing. According to intermediary R 's disclosure strategy and the monotonicity of belief functions in perspective, intermediary R conceals news only when she weakly prefers $+1$ whereas her critical subscriber would weakly prefer -1 if seeing the news. This indicates that when there is no disclosure, the critical subscriber forms negative posterior expectation, which is contradictory by definition. Therefore, the critical perspective must have the same sign as the associated intermediary.

The ideological alignment between an intermediary and her critical perspective suggests that the intermediary discloses any news in line with her ideology (cf. part (iv)). It is because, regardless of the veracity of the aligned news, the intermediary and her critical subscriber will both strictly prefer the same action which aligns with their perspectives. Moreover, disclosing any aligned news in turn sustains this ideological alignment: Now that only the news with the opposite sign from the intermediary's perspective may be concealed, when there is no disclosure, her subscribers accounts for the possibility that some opposing news was concealed and adjust their expectation towards the opposite wing. Then the critical perspective must be in line with the intermediary's so as to satisfy the indifference condition (cf. Lemma 4.3).

Verification decision

Knowing that she will play disclosure strategy m_i^* , intermediary i decides whether to verify news x so as to maximize (4.6):

$$v_i^*(G^s, x) = 0 \Leftrightarrow 0 \in \operatorname{argmax}_{v_i \in \{0,1\}} \mathbb{E}_i \left[\theta [G_i^s(+\infty) - 2G_i^s(\hat{\mu}_i^*(G^s, m_i^*(h_i^{v_i})))] \mid G^s, x, v_i \right] - C \mathbb{1}_{\{v_i=1\}}, \quad (4.8)$$

where

$$h_i^{v_i} = \begin{cases} (G^s, x, 1, \omega), & \text{if } v_i = 1; \\ (G^s, x, 0), & \text{if } v_i = 0. \end{cases}$$

Precisely, if intermediary i chooses not to verify the news ($v_i = 0$), the expected value of (4.8) becomes

$$\mathbb{E}_i[\theta | x] [G_i^s(+\infty) - 2G_i^s(\hat{\mu}_i^*(G^s, m_i^*(G^s, x, 0)))] , \quad (4.9)$$

On the other hand, to evaluate whether verifying the news ($v_i = 1$) is more beneficial, intermediary i estimates the likelihood of the news being true (or false) and accounts for her subsequent optimal disclosure decisions while incurring the verification cost:

$$\begin{aligned} & \mathbb{P}_i(T|x) \mathbb{E}_i[\theta | x, T] [G_i^s(+\infty) - 2G_i^s(\hat{\mu}_i^*(G^s, m_i^*(G^s, x, 1, T)))] , \\ & + (1 - \mathbb{P}_i(T|x)) \mu_i [G_i^s(+\infty) - 2G_i^s(\hat{\mu}_i^*(G^s, m_i^*(G^s, x, 1, F)))] - C. \end{aligned} \quad (4.10)$$

According to Lemma 4.4, intermediary i chooses the same disclosure action for the news that does not change her favored action regardless of its veracity. In this case, we can observe that her expected utility from no verification (4.9) is larger than verifying the news (4.10) by cost C , and intermediary i decides not to verify it. It implies the necessary condition for intermediary i to verify news, as stated in the following lemma:

Lemma 4.5 (necessary condition for verification) *In any equilibrium, for any G^s and news x , intermediary i verifies news x only if the news, if it is informative, will flip the sign of her expectation to the opposite, i.e.,*

$$v_i^*(G^s, x) = 1 \text{ only if } \mu_i \mathbb{E}_i[\theta | x, T] < 0.$$

Moreover, if intermediary i verifies the news, she then chooses different disclosure decisions for true news and false news after learning its veracity, i.e., if $v_i^(G^s, x) = 1$, then $m_i^*(G^s, x, 1, T) \neq m_i^*(G^s, x, 1, F)$.*

When observing some news that can change an intermediary's favored action if it is informative, she has an incentive to verify it and learn its veracity, based on which she can tailor her disclosure decision and accordingly persuade her subscribers to take her

favored actions. Moreover, if an intermediary verifies some news at the equilibrium, she then has to choose different disclosure decisions for true and false news respectively, say disclosing the news if and only if it is true (resp. if and only if it is false). Otherwise, she can simply choose not to verify the news and make the same disclosure decision, deriving the same expected aggregate utility of her subscribers without spending the cost of verification.

The separating disclosure decisions asymmetrically signal the verification result to the subscribers: When the news is true (resp. false) and disclosed, the subscribers know that it is true (resp. false), whereas when the news is false (resp. true) and concealed, the subscribers observe no disclosure, unable to distinguish it from no news or other concealed news. We will see more details about the effect of this asymmetrical signaling on subscription choices in Section 4.5.

The next theorem summarizes the equilibrium structure identified in Lemma 4.3, 4.4, and 4.5.

Proposition 4.1 *Any perfect Bayesian equilibrium features the following properties:*

- (i) *Each subscriber plays a threshold action strategy.*
- (ii) *Each intermediary's ideology and her critical subscriber's are aligned.*
- (iii) *Each intermediary discloses the news whose sign is not opposite from her ideology.*
- (iv) *Each intermediary verifies news only if the news will switch her belief sign when it is true.*
- (v) *If an intermediary verifies news and learns its veracity, she chooses different disclosure decisions for true news and false news.*

4.4 Equilibria without Verification Option

In this section, in order to understand how media bias in selective disclosure affects subscription choices, we investigate the equilibria assuming that none of the intermediaries have the option of verifying news (corresponding to infinite cost of verification, i.e., $C = +\infty$). Any player i receiving news x updates her expectation about the state as $\mathbb{E}_i[\theta | x]$. We will study the characteristic of this belief function, illustrating the variability in players' reactions to extreme news with uncertain veracity. We in turn categorize the subscribers as *centrists* and *extremists*, and investigate their equilibrium subscription choices.

4.4.1 Variability in players' reactions to news with uncertain veracity

Below we exemplify the expectation of a player i for news with unknown veracity ($\mathbb{E}_i[\theta | x]$) by considering a Gaussian system for prior beliefs (f) and generation process of news (f_ϵ, f_F).

Example 4.1 Let ϕ be the PDF of the standard normal distribution. Consider $f(x) = \frac{1}{\sigma_\theta} \phi(\frac{x}{\sigma_\theta})$, $f_\epsilon(\epsilon) = \frac{1}{\sigma_\epsilon} \phi(\frac{\epsilon}{\sigma_\epsilon})$ and $f_F(x) = \frac{1}{\sigma_F} \phi(\frac{x}{\sigma_F})$ where standard deviations $\sigma_\theta, \sigma_\epsilon, \sigma_F$ are positive. Then

$$f_T(x; \mu_i) = \frac{1}{\sigma_T} \phi\left(\frac{x - \mu_i}{\sigma_T}\right)$$

where $\sigma_T \triangleq \sqrt{\sigma_\theta^2 + \sigma_\epsilon^2}$. The posterior mean of player i for informative news x is

$$\mathbb{E}_i[\theta | x, T] = \frac{\sigma_\epsilon^2}{\sigma_T^2} \mu_i + \frac{\sigma_\theta^2}{\sigma_T^2} x;$$

her posterior mean for unknown veracity is updated as:

$$\mathbb{E}_i[\theta | x] = (1 - \mathbb{P}_i(\omega = T|x)) \mu_i + \mathbb{P}_i(\omega = T|x) \left(\frac{\sigma_\epsilon^2}{\sigma_T^2} \mu_i + \frac{\sigma_\theta^2}{\sigma_T^2} x \right),$$

where $\mathbb{P}_i(T|x) = \frac{\frac{1-q}{\sigma_T} \phi(\frac{x-\mu_i}{\sigma_T})}{\frac{q}{\sigma_F} \phi(\frac{x}{\sigma_F}) + \frac{1-q}{\sigma_T} \phi(\frac{x-\mu_i}{\sigma_T})}$.

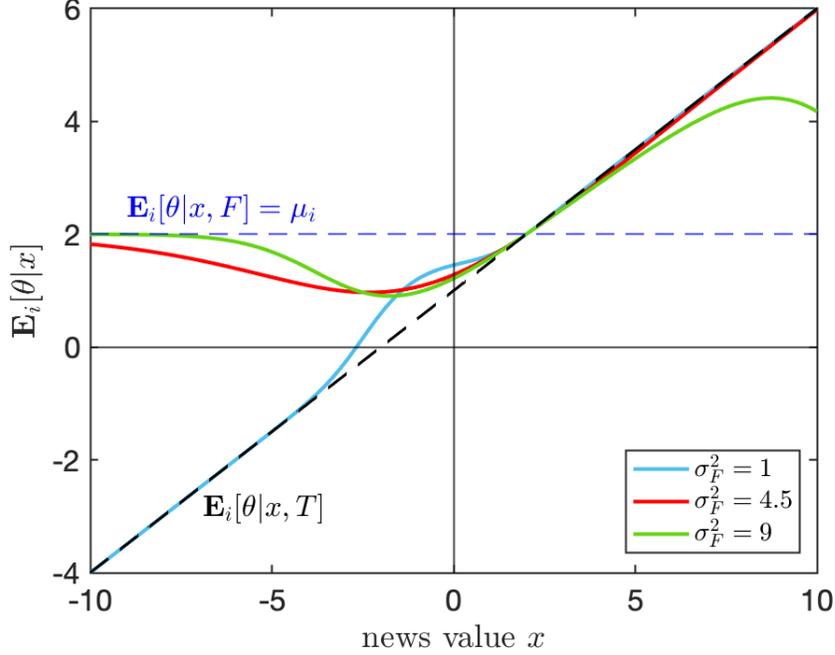


Figure 4-3: The posterior mean as a function of news x when news veracity is uncertain and the signal structures are Gaussian variables: $\sigma_\theta = \sigma_\epsilon = 1.5$ ($\sigma_T^2 = 4.5$) and $q = 0.2$, $\mu_i = 2$.

For Gaussian distributions, the posterior mean $\mathbb{E}_i[\theta|x]$, as a function of news, features a variety in the player’s reaction to the news that is away from her perspective, as illustrated in Figure 4-3 for three different cases. The asymptotic trend of the posterior mean varies with the relative density between the tails of the distributions of true and false news. When the variance in false news (σ_F) is smaller than that in true news (σ_T) (cf. blue curve), the player reckons that the news that is away from her perspective, regardless of which ideological extreme the news lies in, is more likely to be true and updates her belief mean closer to the posterior mean for true news. By contrast, if the variance in false news is greater (cf. green curve), she regards the extreme news as likely to be false and her belief mean is updated close to her perspective. Lastly, when the variances are equal (cf. red curve), the player thinks that the extreme news in line with her perspective is likely to be true whereas the

extreme news on the opposite ideology is probably false. We will leverage this finding for many of our examples in the following sections.

From Figure 4-3 we also observe that there may exist some extreme perspectives such that the sign of corresponding posterior mean is unchanged for any news with unknown veracity (cf. red and green curve). Based on the monotonicity of belief function in perspective, the range for such perspective values can be identified by the following bounds:

$$\begin{aligned}\nu_L &\triangleq \max\{\mu_i \in \mathbb{R} \mid \mathbb{E}_i[\theta \mid x] \leq 0 \forall x \in \mathbb{R}\}; \\ \nu_R &\triangleq \min\{\mu_i \in \mathbb{R} \mid \mathbb{E}_i[\theta \mid x] \geq 0 \forall x \in \mathbb{R}\}.\end{aligned}$$

Note that $\nu_L \in [-\infty, 0)$, and that $\nu_L = -\nu_R$ by the symmetry of f, f_e, f_F around zero. We now define *centrist* and *extremist* subscribers as follows:

Definition 4.3 *Given any belief system and generation process of news, we refer to any subscriber j with $\mu_j \leq \mu_L^{\text{ext}}$ or $\mu_j \geq \mu_R^{\text{ext}}$ as an extremist, where*

$$\mu_L^{\text{ext}} \triangleq \max\{\mu_L, \nu_L\} \text{ and } \mu_R^{\text{ext}} \triangleq \min\{\mu_R, \nu_R\};$$

otherwise, we call subscriber j as a centrist.

The definition of extremists is based on whether their perspectives are relatively more extreme than an intermediary or they can not be influenced to change their favored actions via any news with uncertain veracity.

4.4.2 Subscription choices of centrists and extremists

First, the disclosure strategy stated in Lemma 4.4 can be reduced to the following rule

$$m_i^*(G^s, x, 0) = x \text{ if and only if } \mathbb{E}_i[\theta \mid x] \mathbb{E}_{k_i^*}[\theta \mid x] > 0. \quad (4.11)$$

Based on (4.11), we identify the equilibria and the subscription choices in the following proposition:

Proposition 4.2 *When verification option is not possible, equilibria exist and all the equilibria feature the unique intermediaries' disclosure strategies and the two unique critical perspectives for each intermediary. The critical subscribers are both centrists, i.e., μ_k^* is between 0 and μ_i^{ext} . In any equilibrium,*

- **strict anti-homophily among all the centrists:** *for any subscriber j , if $\mu_j \in (\mu_L^{\text{ext}}, 0)$, then $s_j^* = R$; on the other hand, if $\mu_j \in (0, \mu_R^{\text{ext}})$, then $s_j^* = L$.*
- **indifference among all the extremists:** *for any subscriber j such that $\mu_j \notin (\mu_L^{\text{ext}}, \mu_R^{\text{ext}})$ or $\mu_j = 0$, subscribing to either intermediary results in the same ex-ante expected utility.*

The equilibrium critical subscribers must be centrists. If the critical subscribers are otherwise extremists (ideologically more extreme than their own intermediary or do not change their favored action for any news), the disclosure strategy (4.11) and the monotonicity of belief function in perspective together suggest that an intermediary conceals news only when it switches her favored action but will not change her critical subscriber's. Then the critical subscriber in turn strictly prefers the action aligned with his ideology when there is no disclosure, making the indifference condition unsatisfied.

The qualitative properties of equilibrium subscription choices are robust to the ex-ante probability of news being false and the variability in players' reaction to extreme news (cf. Figure 4-3). In fact, the subscription choices and disclosure strategies are rooted at the change in the intermediaries' favored actions in reaction to news and the monotonicity of belief mean in perspective, as we will explain below.

When making a subscription choice, each subscriber accounts for potential utility loss incurred in the event of no disclosure where he may take an action that is not optimal if he were to observe his intermediary's type (some news x or no news). Specifically, the nondisclosure event can result from multiple possibilities (either no news or some

concealed news that opposes his intermediary's perspective), but the subscriber cannot distinguish the underlying reason, only to take one action which may be sub-optimal for part of the possible causes. The subscriber chooses an intermediary to minimize this expected loss from the mismatch between his best-response action in the nondisclosure event and his optimal action if he could observe the information set of his intermediary.

If subscribing to the intermediary with opposite ideology, a subscriber effectively takes his optimal action for any news realization, achieving his ex-ante best expected payoff. According to Lemma 4.4, the news that is concealed by the opposite intermediary must be against her ideology, that is, must be in line with the subscriber's perspective. As the subscriber's optimal action for no news also aligns with his perspective, it suggests that when there is no disclosure the subscriber will take the action aligned with his ideology, which is optimal for concealed news and no news.

On the other hand, if making homophilic subscription choices, due to the identification problem in the non-disclosure event, centrists cannot effectively take their optimal actions for some news advocating the opposite ideology from theirs, incurring some loss in their ex-ante expected utility. Specifically, no disclosure can be attributed to either no news or to some concealed news that opposes the intermediary's ideology and can sway her critical subscriber, who is a centrist with the same ideology, to take the opposite action. It suggests that the centrist would maintain his ex-ante favored action if he knew there was no news, but for some news concealed by this intermediary with aligned ideology he would prefer the opposite action if he could observed the news. The centrist then reckons that ex-ante there must be some news realizations (or the case of no news) for which his corresponding optimal action is opposite from the single action he will take in the non-disclosure event. As a result, the centrist makes an anti-homophilic subscription choice.

In contrast, the extremists are still able to achieve their best expected payoff even if making homophilic subscription choices. When there is no disclosure, extremists choose the action in line with their own perspective and they know that it is optimal for

whatever news realizations that lead to no disclosure. First, for an extremist who never switches the sign of his expectation whatever news he observes (i.e., the magnitude of his perspective is no less than $|\nu_i|$), he always takes the action aligned with his ideology, which is his optimal action for any news realization. Indeed, his action strategy is independent of both intermediaries' disclosure decisions and hence he finds indifference. Second, for an extremist who is more extreme than the intermediary on the same wing, he knows that any opposing news that will flip his favored action is certainly disclosed by the aligned intermediary as well and he can then choose his optimal action. It is because such opposing news also flips the intermediary's favored action, incentivizing her to disclose it to persuade her subscribers. Consequently, the extremist knows that the action aligned with his perspective is optimal for any case that leads to no disclosure and can obtain his best ex-ante expected payoff.

Finally, as a special case, we can find the equilibria when ex-ante there is no false news (i.e., $q = 0$) and $\nu_R = -\nu_L = +\infty$:

Corollary 4.1 *When ex-ante there is no false news, i.e., $q = 0$, in any equilibrium,*

- **strict anti-homophily among all the centrists:** *For any subscriber j , if $\mu_j \in (\mu_L, 0)$, then $s_j^* = R$; on the other hand, if $\mu_j \in (0, \mu_R)$, then $s_j^* = L$.*
- **indifference among all the extremists:** *For any intermediary i and any subscriber j such that $\mu_j \notin (\mu_L, \mu_R)$ or $\mu_j = 0$, there exists an equilibrium in which $s_j^* = i$.*

Corollary 4.1 generalizes our preliminary result in Hsu et al. (2020b), in which players' beliefs and the additive noise in informative news are Gaussians. As news must be informative, for each subscriber there must exist some informative news that can change his favored action. The extremists find indifference between the intermediaries because both intermediaries disclose the informative news that flips their favored actions.

4.5 Equilibria with Verification Option

In this section we study the general model in which only the two intermediaries can verify the veracity of news. As Proposition 4.1 shows, each intermediary verifies the news only if it is opposing her ideology. Moreover, since verification results are completely unobservable to subscribers, upon receiving some news, a subscriber has to infer the news veracity and update his belief according to his intermediary strategies; the intermediaries need to signal verification results via their disclosure decisions. In the following discussion, we concentrate on the equilibria in which each intermediary discloses the news that is verified if and only if it is informative.

Definition 4.4 *We refer to the class of equilibria in which each intermediary chooses to disclose the news that is verified if and only if it is true as true-news-disclosing equilibria.*

In the class of true-news-disclosing equilibria, when an intermediary disclosed unfavorable news against her ideological perspective, her subscribers reckon that if the news was verified, it must be that this news is informative and has changed their intermediary's favored action.

We note that when ex-ante there is no false news (i.e. $q = 0$), any news is informative for sure and intermediaries do not verify, considering whether to disclose informative news. The corresponding equilibrium subscription choices are as characterized in Corollary 4.1. In the following discussion we therefore focus on the case $q > 0$ and we study the equilibria for $C = 0$ and $C > 0$.

Equilibrium with zero verification cost $C = 0$

Costless fact-checking is beneficial for the intermediaries when their own favored action may be changed by informative news; in fact, we show that the necessary condition in Lemma 4.5 is also sufficient:

Lemma 4.6 (sufficient condition for verification when the cost is zero) *When $q > 0$ and $C = 0$, in any equilibrium, for any common history G^s and news x , intermediary i verifies the news if it will change the sign of her expectation about the state when it is true, i.e., $v_i^*(G^s, x) = 1$ if $\mu_i \mathbb{E}_i[\theta | x, T] < 0$.*

The following proposition identifies subscription choices in the true-news-disclosing equilibria.

Proposition 4.3 *Given $q > 0$ and both intermediaries can verify news at cost $C = 0$, true-news-disclosing equilibria exist and feature the unique intermediaries' strategies and the unique set of critical perspectives. The critical subscribers are centrists, i.e., $\mu_{k_i^*}$ is between 0 and μ_i^{ext} . In any true-news-disclosing equilibrium,*

- **strict anti-homophily among the centrists:** *for any subscriber j with $\mu_j \in [\mu_{k_L^*}, 0)$, then $s_j^* = R$; on the other hand, if $\mu_j \in (0, \mu_{k_R^*}]$ then $s_j^* = L$;*
- **strict homophily among all the extremists:** *for any subscriber j , if $\mu_j \leq \mu_L^{\text{ext}}$, then $s_j^* = L$; on the other hand, if $\mu_j \geq \mu_R^{\text{ext}}$, then $s_j^* = R$;*
- *for any subscriber j with $\mu_j = 0$, she feels indifferent between the two intermediaries.*

Figure 4-4 illustrates the strategies of intermediary R and the subscription choices in the true-news-disclosing equilibria when $C = 0$. According to Lemma 4.5 and 4.6, each intermediary verifies news if and only if it will flip her favored action when it turns out to be informative, and then chooses different disclosure decisions to signal the verification result. Below we examine the subscription choices of extremists and centrists through the two functions of intermediaries: verifying news and selective disclosure. For simplicity we discuss the rationale for the right-leaning subscribers.

A right-leaning subscriber cares about the veracity of the left-leaning news that is extreme since his favored action will be switched to -1 if the news is informative. For the other news values, including any right-leaning news, he prefers action $+1$

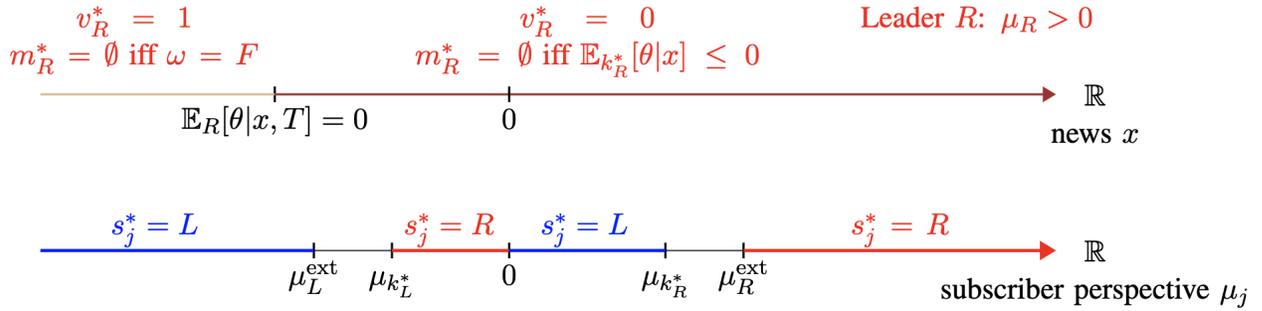


Figure 4-4: Verification and disclosure strategy of leader R and the subscription choices in the true-news-disclosing equilibria as described in Proposition 4.3.

regardless of its veracity. If subscribing to intermediary L , since only right-leaning news (unverified or verified as false news) or no news can lead to no disclosure by intermediary L , the right-leaning subscriber takes $+1$ in the non-disclosure event and it is his optimal action for both concealed news and no news. However, since intermediary L does not verify any left-leaning news and only discloses it (cf. Proposition 4.1), the subscriber can only take the same action for left-leaning news whether it is true or false, suggesting that he does not benefit from intermediary L 's function of verifying news.

By contrast, intermediary R verifies left-leaning extreme news, but concealing false news to signal verification results couples with no news and her concealment of some unverified left-leaning news in the non-disclosure event, interfering right-leaning subscribers' extraction of this verification benefit. Indeed, the extremists can derive the benefit from intermediary R 's verifying left-leaning news whereas the centrists who are not as extreme as critical subscriber of intermediary R (i.e., their perspectives are between 0 and $\mu_{k_R^*}$) cannot. This result leads to homophily among all the extremists but anti-homophily among the centrists; we explain it below.

For a right-leaning extremist, any left-leaning news that flips his favored action to -1 if it is true also flips intermediary R 's. Therefore, the extremist knows that intermediary R will verify the left-leaning news and disclose it when it is true; in turn he can take his optimal action -1 as long as intermediary R observes such news. For

all the other news and the case of no news, the extremist maintains the sign of his expectation about the state and prefers action $+1$. It then suggests that even when there is no disclosure from intermediary R , the extremist has no concern that there exists some news that switches his favored action (otherwise, it would be disclosed by intermediary R), and he takes his optimal action $+1$, regardless of the underlying cause of no disclosure. Consequently, by subscribing to intermediary R , the right-leaning extremist effectively takes his optimal action for any news value and any veracity, achieving his best ex-ante expected payoff. It is better than choosing intermediary L , who does not verify any left-leaning news whose veracity matters to the extremist's action.

On the other hand, for a right-leaning centrist who is not as extreme as critical subscriber of intermediary R , the identification problem he faces in the non-disclosure event makes him unable to derive the benefit from intermediary R 's verifying left-leaning extreme news as the extremists do. Specifically, no disclosure can also result from concealment of left-leaning news that is not verified by intermediary R (such news can be true), besides no news and false left-leaning news. With a perspective close to 0, the centrist thinks that informative news is around 0 with high probability and reckons that no disclosure is more likely attributed to concealment of unverified left-leaning news that is true (intermediary R verifies only extreme news), therefore choosing -1 . As a result, for any left-leaning news value that is verified by intermediary R , the centrist ends up taking action -1 for both true news (it is disclosed and switches his favored action to -1) and false news (it is concealed and he chooses -1), unable to take his optimal actions correspondingly.⁹ Moreover, choosing -1 in the non-disclosure event incurs loss if there is in fact no news or false news and his optimal action is $+1$.

⁹Given any news value, the expected utility resulting from taking action -1 is weakly worse than choosing an action to maximize expected utility based on the news value while speculating its veracity. Therefore, in terms of the left-leaning news value verified by intermediary R , the centrist derives weakly lower ex-ante expected utility from intermediary R than L : He effectively takes -1 for the news value if following intermediary R whereas intermediary L discloses the news so that he observes it and makes his best response.

Overall, considering that he cannot benefit from intermediary R 's verification and may take a sub-optimal action in the non-disclosure event, the centrist finds it strictly more beneficial to subscribe to intermediary L .

Equilibrium with non-zero verification cost $C > 0$

We consider the scenario in which the verification comes with positive cost, i.e., $C > 0$.¹⁰ Due to the cost, each intermediary does not verify all of the opposing news that, if it is informative, changes her favored action. To reckon whether verifying news is worth the cost, the intermediary has to account for not only the likelihood of the news being true or false and her belief mean for each case, but also the fraction of her subscribers who can be persuaded.

We find that the anti-homophily still arises among the centrists: The identification problem in the non-disclosure event still makes centrists unable to benefit from verification by the intermediary with aligned ideology and also leads to utility loss due to misaligned optimal actions for underlying causes of no disclosure, pushing the centrists to make anti-homophilic subscription choices. On the other hand, there is *weak* homophily exhibited in the extremists' subscription choices: They either strictly prefer the aligned intermediary or feel indifferent. As discussed in Proposition 4.3, an extremist strictly prefers the aligned intermediary because she verifies the news that potentially changes the sign of his expectation about the state and discloses it if it is informative. However, when news verification is costly, it can happen that any such news is not verified by the aligned intermediary; in this case, the extremist then finds indifferent between the two intermediaries. The following proposition summarizes the equilibrium subscription choices when $C > 0$.

Proposition 4.4 *Given $q > 0$ and both intermediaries can verify news at cost $C > 0$, in any true-news-disclosing equilibrium, the critical subscribers are centrists, i.e., $\mu_{k_i}^*$*

¹⁰The cost is implicitly assumed on a moderate level so that the intermediaries still have incentives to verify some news. Otherwise, the results would be equivalent to the case when verification is not possible (cf. Proposition 4.2).

is between 0 and μ_i^{ext} , and strict anti-homophily arises among the centrists whereas weak-homophily emerges among all the extremists.

4.6 Comparative Statics on Subscription Choices

In this section, we study how the subscription choices may change with the ideological biases of the intermediaries and the cost of verification. In particular, we investigate the fine grains of the weak homophily among the extremist subscribers, characterizing the subsets of the extremists who strictly prefer the aligned intermediary and who finds indifference. Our discussion will be focused on right-leaning players for simplicity.

4.6.1 Intermediaries' perspectives and the set of extremists

From Proposition 4.2, 4.3, and 4.4, we observe that the equilibrium subscription choices of right-leaning subscribers do not vary with intermediary L 's ideological bias. It is because right-leaning subscribers, whose favored actions may be switched only by left-leaning news, find that intermediary L deploys the same verification and disclosure strategy for any left-leaning news regardless of her ideological bias. In turn the right-leaning subscribers, whose favored actions may be switched only by left-leaning news, derive the identical ex-ante expected utility if subscribing to intermediary L .

By contrast, since intermediary R with a different perspective deploys different strategies for left-leaning news, the subscription choices of right-leaning subscribers change with intermediary R 's ideological bias. Indeed, the boundary μ_R^{ext} that distinguishes right-leaning extremists and centrists for their different subscription choices depends on intermediary R 's perspective μ_R as well as the bound ν_R . Specifically, the boundary μ_R^{ext} is the smaller value of μ_R and ν_R . We discuss how the boundary scales with intermediary R 's perspective in two cases: ν_R is infinite or finite.

When $\nu_R = +\infty (= -\nu_L)$, e.g., when ex-ante there is no false news ($q = 0$) or $\sigma_T > \sigma_F$ in Example 4.1, the boundary equals intermediary R 's perspective. When

intermediary R ' perspective moves to the extreme (i.e., μ_R is increased), there are more right-leaning centrists, who tend to subscribe to intermediary L due to their concern about left-leaning news close to the neutrality being concealed by intermediary R . If intermediary L also has an extreme perspective, i.e., the two intermediaries are sharply polarized, there is a wide range of centrists who make anti-homophilic subscription choices.

On the other hand, when ex-ante false news exists and the bound ν_R is finite (e.g., $\sigma_T \leq \sigma_F$ in Example 4.1), there are extremists whose favored actions can be influenced only by informative news: The right-leaning extremists with perspectives no less than ν_R always weakly prefer intermediary R regardless of her perspective, since only intermediary R may verify left-leaning news and disclose it if it is true. In other words, the range of the right-leaning centrists, who exhibit anti-homophilic subscription behavior, does not scale up with intermediary R 's perspective, limited by the bound ν_R .

4.6.2 Strict homophily and indifference among the extremist subscribers

Weak homophily arises among all the extremists no matter what level of verification cost is. In the following example, we further study how the subsets of the extremists who strictly prefer the aligned intermediary and who feel indifference can vary with the verification cost and the intermediaries' perspectives.

Example 4.2 (strict homophily and indifference among the extremists) Suppose that f, f_ϵ, f_F are Gaussian distributions as described in Example 4.1 and that $\sigma_T^2 = \sigma_F^2$. We assume that the intermediaries have biases large enough such that $|\mu_i| > |\nu_i|$. We also assume that there are more extremist subscribers than the centrists for both ide-

ological wings; particularly, for $i \in \{L, R\}$,

$$|G_0(\nu_i) - G_0(\mu_i)| > |G_0(0) - G_0(\nu_i)|$$

We analyze the verification strategy of intermediary R and right-leaning subscribers; the case for left-leaning players can be similarly derived by symmetry.

Intermediary R verifies news x only if it can flip her favored action when it is true, i.e., $\mathbb{E}_R[\theta|x, T] < 0$. Given equilibrium perspective measures (G_L^{s*}, G_R^{s*}) and the cost, one can further show that in a true-news-disclosing equilibrium, for any news x such that $\mathbb{E}_R[\theta|x, T] < 0$, intermediary R verifies the news only if

$$\mathbb{P}_R(T|x)\mathbb{E}_R[\theta|x, T] \quad \underbrace{2(G_R^{s*}(\widehat{\mu}_R^*(\emptyset)) - G_R^{s*}(\widehat{\mu}_R^*(x)))}_{\text{the fraction who are persuaded if } x \text{ is true and disclosed}} > C, \quad (4.12)$$

where $\widehat{\mu}_R^*(x)$ stands for the perspective threshold when the subscribers know that news x is true, i.e., $\mathbb{E}_j[\theta|x, T] = 0$ if $\mu_j = \widehat{\mu}_R^*(x)$. Recall that the equilibrium critical subscribers are centrists; we have the relation $\widehat{\mu}_R^*(\emptyset) < \nu_R < \mu_R < \widehat{\mu}_R^*(x)$. Intuitively, the condition (4.12) means that deviating to not verifying and concealing the news (intermediary R prefers +1 for any news with unknown veracity) is not beneficial.¹¹ Moreover, intermediary R can verify the news in an equilibrium if the condition (4.12) holds; below we study the equilibria in which she chooses to verify when (4.12) is satisfied.

Using two simple bounds on the fraction of her subscribers who are persuaded and the fact that $\mathbb{P}_R(T|x)\mathbb{E}_R[\theta|x, T]$ is unimodal in the range of news x for which $\mathbb{E}_R[\theta|x, T] < 0$, we identify the subsets of extreme left-leaning news that intermediary R verifies and does not verify in the following analysis, which is illustrated in Figure 4-5.

¹¹Specifically, when intermediary R verifies news x in an equilibrium, she discloses true news whereas conceals false news. If she deviates to not verifying it, then given the news with unknown veracity her posterior expectation about the state is positive ($\mu_R > \nu_R$) and she will conceal the news because $\widehat{\mu}_R^*(\emptyset) < \widehat{\mu}_R^*(x)$. The condition (4.12) dictates that intermediary R derives higher expected payoff from verifying the news so as she will not deviate.

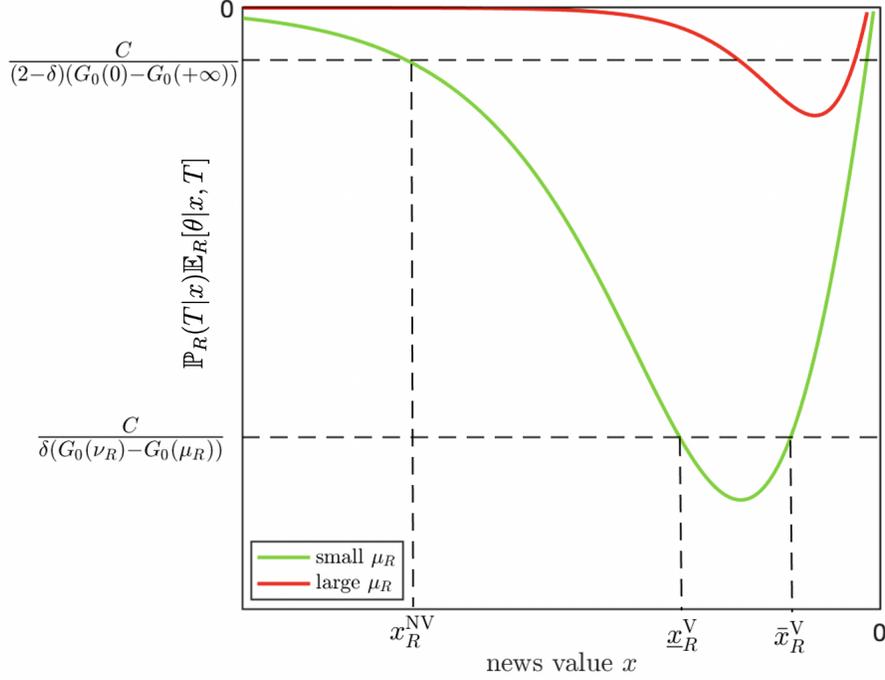


Figure 4-5: Illustration of intermediary R 's decision-making of verifying left-leaning news that satisfies $\mathbb{E}_R[\theta|x, T] < 0$ in Example 4.2.

- **News that is verified and the extremists who make homophilic choices:** The magnitude of the fraction is lower bounded by $|(G_0(\nu_R) - G_0(\mu_R))\frac{\delta}{2}|$, which results from the noise subscribers with perspectives in the range (ν_R, μ_R) . Intermediary R verifies the news x if

$$\mathbb{P}_R(T|x)\mathbb{E}_R[\theta|x, T](G_0(\nu_R) - G_0(\mu_R))\delta > C. \quad (4.13)$$

Intermediary R verifies news x if $x \in (\underline{x}_R^V, \bar{x}_R^V)$ where $\underline{x}_R^V \leq \bar{x}_R^V < 0$ both satisfy the indifference condition in (4.13). Consequently, any subscriber j with perspective $\mu_j \in [\mu_R^{\text{ext}}, \mu_R^{\text{SH}})$, where μ_R^{SH} is the unique perspective such that $\mathbb{E}_j[\theta | \underline{x}_R^V, T] = 0$ when $\mu_j = \mu_R^{\text{SH}}$, strictly prefers intermediary R because she verifies some news that will affect his action if it is true.

- **News that is not verified and the extremists who find indifference:** The magnitude of the fraction is at most $|(G_0(0) - G_0(+\infty))(1 - \frac{\delta}{2})|$, which comprises

all the right-leaning subscribers (except for some noise subscribers). Intermediary R does not verify news x if

$$\mathbb{P}_R(T|x)\mathbb{E}_R[\theta|x, T](G_0(0) - G_0(+\infty))(2 - \delta) < C. \quad (4.14)$$

Intermediary R does not verify news x if $x \leq x_R^{\text{NV}}$ for some threshold $x_R^{\text{NV}} < 0$ that satisfies the indifference condition in (4.14). For any subscriber j with perspective $\mu_j \geq \mu_R^{\text{ind}}$, where μ_R^{ind} is the unique perspective such that $\mathbb{E}_j[\theta | x_R^{\text{NV}}, T] = 0$ when $\mu_j = \mu_R^{\text{ind}}$, he feels indifferent between the intermediaries since no intermediaries verify the news of which he desires to learn the veracity.

Figure 4-6 depicts how the extremists' subscription behavior of strict homophily and indifference vary with the verification cost. When the cost is zero, intermediary R verifies any news x for which $\mathbb{E}_R[\theta|x, T] < 0$, and strict homophily arises among all the extremists (cf. Proposition 4.3). As the cost is increased, the two bounds on the fraction in Figure 4-5 move downwards, capturing intermediary R 's increased reluctance to verify news. As a result, more right-leaning subscribers on the very extreme feel indifferent between the intermediaries (μ_R^{ind} moves towards 0) and less right-leaning subscribers strictly prefer intermediary R (μ_R^{sH} moves towards 0). When the cost becomes sufficiently large and the intermediaries find verification not beneficial for any news, all the extremists feel indifferent, which corresponds to Proposition 4.2 when verification is not possible.

We also remark that the range of news x for which $\mathbb{E}_R[\theta|x, T] < 0$ is decreasing in her perspective, and so is the magnitude of the product $\mathbb{P}_R(T|x)\mathbb{E}_R[\theta|x, T]$: Intermediary R with a more extreme perspective thinks that left-leaning news is more unlikely to be true and even if it is true her posterior mean tends to be on the right wing. This monotonicity is reflected in Figure 4-5 by the relative position of the two curves for small and large perspectives. As a result, given a moderate verification cost, a more extreme intermediary does not verify a large set of news, leading to more right-leaning

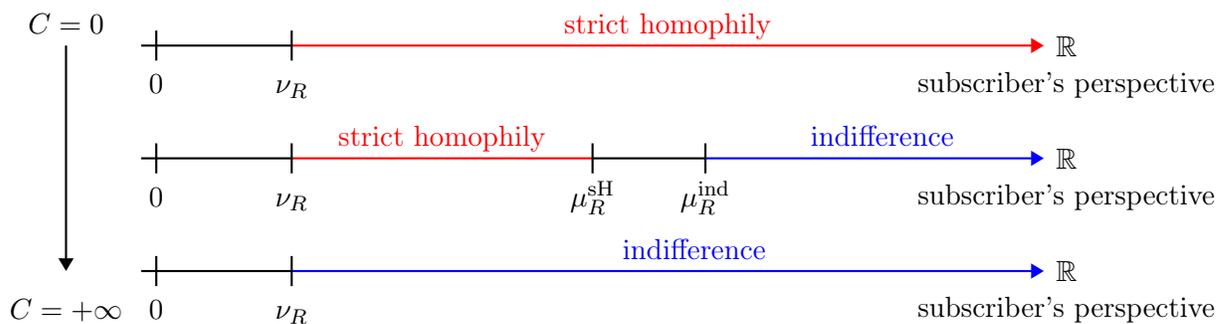


Figure 4-6: The range of strict homophily and indifference for the right-leaning extremists' subscription choices varies with the verification cost in Example 4.2.

extremists finding indifference between the two intermediaries.

4.7 Extensions

We study a few variations of our model. We point out that the variety of disclosure strategies for the set of verified news can arise in equilibria, other than the class of true-news-disclosing equilibria. We then investigate another setting under which verification result is perfectly observable to subscribers for any disclosed news. Finally, we discuss an alternative to our modeling on the utility function and action space.

4.7.1 Disclosure strategies for verified news

In Section 4.5 we concentrate on the true-news-disclosing equilibria. In fact, there can be infinitely many PBEs featuring different disclosure strategies for the set of verified news, since for each news value in the verification region an intermediary may choose one of the two separating disclosure decisions. For example, consider the case with zero verification cost and suppose after verifying news, an intermediary *discloses the news if and only if it is uninformative*. Believing in this disclosure strategy, her critical subscriber, when seeing the news that should be verified by the intermediary, maintains his posterior mean equal to his perspective and hence takes the action in line with the

intermediary's perspective. The intermediary in turn has no incentive to deviate to conceal the news, which will otherwise make her critical subscriber feel indifferent and consequently result in less fraction of subscribers taking her favored action. Neither will she deviate to disclose the news if it is true: Her critical subscribers would regard the disclosed news as false and in turn choose the unfavorable action in this case. Such an equilibrium exists.

The strategy that disclosing verified news if and only if it is false changes the equilibrium subscription choices of the centrists and the extremists from the result in Proposition 4.3. In contrast to the true-news disclosing equilibria, the centrists can derive the verification benefits from the intermediary with aligned ideology as they are able to take their optimal actions for true news and false news respectively when it is verified. If the verification gain outweighs the expected loss from inability to take their optimal actions when there is actually no news, the centrists make homophilic subscription choices.

The extremists, on the other hand, no longer derive the verification benefits from the intermediary with aligned ideology. The reason is similar to why the centrists can not derive such benefits in the true-news-disclosing equilibria: The extremists end up with the same action in line with their perspectives no matter whether the verified news is true or false. The extremists may instead turn to the opposite intermediary since this intermediary simply discloses such news (without verification) and the extremists can evaluate the likelihood of the news being true, choosing their optimal actions.

The diversity in equilibrium disclosure strategies is a common result in signaling games that an equilibrium is sensitive to receivers' beliefs about the implications of disclosure and no disclosure. When any information about verification results can not be transmitted to subscribers, the disclosure strategy couples with the beliefs of subscribers about news veracity when observing disclosed news. We relate such a variety of equilibria to the intricacies of readers' linguistic interpretation about the news coverage by intermediaries and only focus on the true-news-disclosing equilibria.

4.7.2 Verification result is perfectly observable to subscribers

We consider a variation of the model in which for any news that is disclosed by an intermediary, her subscribers can perfectly observe the verification result, i.e., whether it was verified and, if it was, whether the news is true or false. This perfect observability can model absolute trust the subscribers have in their intermediaries' report about veracity of news. The perfect observability creates a new incentive for the intermediaries: For the news that is aligned with their perspectives, the intermediaries are motivated to prove that the news is informative so as to sway the beliefs of their subscribers to full extent. As we will see, when this new motive for the intermediaries to show that the aligned news is true is stronger than the motive to learn the veracity of the opposing extreme news that can change their favored actions, they tend to verify aligned news rather than opposing news and anti-homophily can emerge among the extremist subscribers.

First, we identify some properties of intermediaries' disclosure strategy in the following lemma:

Lemma 4.7 (disclosure strategy when verification result is perfectly observable)

In any equilibrium, given any common history G^s ,

- (i) *intermediary i discloses news x if and only if she and her critical subscriber have the same belief sign when seeing the news and its verification result:*

$$m_i^*(G^s, x, 0) = x \text{ if and only if } \mathbb{E}_i[\theta | x] \mathbb{E}_{k_i^*}[\theta | x] > 0,$$

and for any $\omega \in \{T, F\}$

$$m_i^*(G^s, x, 1, \omega) = x \text{ if and only if } \mathbb{E}_i[\theta | x, \omega] \mathbb{E}_{k_i^*}[\theta | x, \omega] > 0.$$

- (ii) *intermediary i 's critical subscriber is ideologically aligned with her but has a less extreme perspective, i.e., $\mu_{k_i^*}$ is between 0 and μ_i .*

(iii) intermediary i discloses any news that is not against her ideology, i.e., for any x such that $\mu_i x \geq 0$, $m_i^*(G^s, x, 1, T) = m_i^*(G^s, x, 1, F) = m_i^*(G^s, x, 0) = x$.

(iv) moreover, for any news x that switches intermediary i 's favored action when it is true, she discloses the news if she verifies it, i.e., for x such that $\mu_i \mathbb{E}_i[\theta | x, T] < 0$, $m_i^*(G^s, x, 1, T) = m_i^*(G^s, x, 1, F) = x$.

By the same insights from Lemma 4.4 and 4.5, each intermediary discloses any news that is not against their perspectives and her critical subscribers have the same ideology as her: The intermediary discloses the aligned news so that her critical subscribers do not adjust their belief means towards the opposite wing if observing no disclosure.

Below we discuss the verification decisions and characterize the equilibria; we in particular study two cases when $C = 0$ and $C > 0$.

No verification cost $C = 0$

With zero verification cost, each intermediary would like to learn whether opposing extreme news is true or false so that she may switch her favored action and accordingly persuade her subscribers (cf. Lemma 4.6). Given the perfect observability on verification results, the intermediary maximizes her expected utility by disclosing both true news and false news, instead of deploying separating disclosure decisions to signal the veracity (cf. Lemma 4.5).¹² Consequently, her subscribers with the same ideology can derive the benefits from her verifying opposing extreme news and take their optimal actions according to its veracity.

Moreover, an intermediary may verify some news that aligns with her perspective because she can sway more subscribers to take her favored action (the one her perspective dictates) if she shows that this news is true. This means that the intermediary is also appealing to the subscribers on the opposite wing, as she may verify the news

¹²We can show that in any equilibrium the critical perspective must lie between the neutrality and the corresponding intermediary's perspective; disclosing both true and false news is the intermediary's best response.

that is against these subscribers and discloses it. However, there is a trade-off in this verification decision: When the news is false, the intermediary has to disclose it to avoid pushing her subscribers to the opposite wing due to their speculation about concealment of opposing news, and then her subscribers take the actions in line with their own perspectives. She in turn loses some fraction of opposing subscribers who would have taken her favored action if she simply disclosed the news without verification (the centrists on the opposite wing think that the unverified news can be true and tend to choose her favored action). The trade-off suggests that the intermediary does not necessarily verify all the aligned news.

We demonstrate that weak homophily arises among the extremists:

Proposition 4.5 *Consider $q > 0$ and both intermediaries can verify news at cost $C = 0$. When verification results are perfectly observable to subscribers for any disclosed news, in any equilibrium weak homophily arises among the subscribers who are more extreme than the intermediaries.*

An extremist derives his maximal ex-ante expected utility if subscribing to the intermediary with aligned ideology, who verifies and discloses any opposing extreme news. On the other hand, the extremist may also attain his best ex-ante payoff from the opposite intermediary and hence find indifference, if any such opposing news that possibly switches his favored action is also verified by the opposite intermediary, who wants to prove its veracity for persuasion purpose. Consequently, weak homophily arises among the extremists.

The equilibrium subscription choices of the centrists are not clear: They have to weigh the benefit from the intermediaries' verification strategies against the potential loss from the identification problem in the non-disclosure event. Their subscription choices are sensitive to the functional forms of the belief system and news generation.

Verification requires a cost $C > 0$

Costly verification discourages the intermediaries to verify news. However, the

extents of this impact on the two motives, to learn news veracity or to prove informativeness, may significantly differ, contingent on the belief system and news generation. We provide an example below to illustrate that the motive to learn news veracity can be overwhelmed by the cost whereas the motive to prove the informativeness stays stronger than the disincentive, leading to anti-homophily among all the subscribers.

Example 4.3 (Anti-homophily among all the subscribers) Suppose that f, f_ϵ, f_F are Gaussian distributions as described in Example 4.1 and that $\sigma_T^2 = \sigma_F^2$. We assume that the intermediaries have commensurate biases, i.e., $\mu_L = -\mu_R$. and their biases are large such that $|\mu_i| > |\nu_i|$.

Suppose that the subscribers are mainly concentrated on four perspective values in terms of their ideological leaning and magnitude of bias: Left or Right, Centrists or Extremists. In particular, the CDF of perspectives is assumed to be

$$G_0(\mu) = (1 - \alpha) \sum_{\mu' \in \{\mu_L - \eta, -\eta, \eta, \mu_R + \eta\}} \frac{1}{4} \mathbb{1}_{\{\mu \geq \mu'\}} + \alpha H(\mu),$$

where $\eta > 0, \alpha > 0$ are both infinitesimal real numbers and $H(\cdot)$ is a CDF with full support and $H(0) = \frac{1}{2}$. We assume verification cost $C > (1 + \alpha)\mu_R$.

There exists an equilibrium in which all the subscribers choose the intermediary with the opposite ideology from their own. We reason as below; suppose the subscription choices hold in the equilibrium,

- Each intermediary i *does not* verify any opposing news x ; intermediary i discloses it if and only if the centrists with the aligned ideology will take her favored action when seeing unverified news:

For any opposing news x that may change her favored action ($\mu_i \mathbb{E}_i[\theta | x, T] < 0$), intermediary i does not verify it because the utility gain from verification is upper

bounded by¹³

$$\underbrace{\mathbb{P}_i(T|x)|\mathbb{E}_i[\theta|x, T]| \frac{1}{2}(1+\alpha)}_{\omega = T \text{ and persuade the aligned extremists}} + \underbrace{(1 - \mathbb{P}_i(T|x))|\mu_i| \frac{1}{2}(1+\alpha)}_{\omega = F \text{ and persuade the aligned centrists}} < (1+\alpha)(1 - \mathbb{P}_i(T|x))|\mu_i| < C,$$

where $\mathbb{P}_i(T|x)|\mathbb{E}_i[\theta|x, T]| < (1 - \mathbb{P}_i(T|x))|\mu_i|$ since intermediary i favors the action in line with her perspective when the news veracity is unknown ($|\mu_i| > |\nu_i|$). On the other hand, for opposing news x such $\mu_i \mathbb{E}_i[\theta|x, T] \geq 0$, the only benefit from verification is to persuade the opposing Centrists if it is false. However, this gain is at most $\frac{1}{2}(1+\alpha)(1 - \mathbb{P}_i(T|x))|\mu_i| < C$ and she does not verify it. As a result, intermediary i does not verify any opposing news.

- Each intermediary i verifies news x that is not against her ideology if

$$\begin{aligned} \mu_i \mathbb{E}_{-i}[\theta|x] < 0 \text{ and} & \tag{4.15} \\ \mathbb{P}_i(T|x)|\mathbb{E}_i[\theta|x, T]| \frac{1}{2}(1-\alpha)(1-\frac{\delta}{2}) > (1 - \mathbb{P}_i(T|x))|\mu_i| \frac{1}{2}(1+\alpha)(1-\frac{\delta}{2}) + C. \end{aligned}$$

She then discloses the news for both true and false news. The first condition in (4.15) dictates that the news, if not verified, can not persuade the opposite intermediary, nor the opposing Extremists, to take the action in line with her perspective; intermediary i then has an incentive to verify and to show that the news is true. The second condition requires that, given the anti-homophily in the subscription choices, the gain from showing the true news and persuading the opposing Extremists with perspective magnitude $|\mu_i| + \eta$ outweighs the verification cost and the loss if the news turns out to be false and the opposing Centrists will otherwise take the opposite action. By the assumption that

¹³The upper bound is computed by considering that intermediary i can persuade all the aligned Extremists (aligned Centrists) when the news is true (false) as well as all the other subscribers on the same wing in the population.

$|\mu_i| > |\nu_i|$ and intermediary i thinks the aligned extreme news is likely to be true (i.e., $\mathbb{P}_i(T|x)|\mathbb{E}_i[\theta|x,T]| \rightarrow +\infty$ when the magnitude of aligned news $|x|$ increases), the set of news that satisfies the conditions (4.15) and is verified is non-empty.

- For each subscriber, the news of which he would like to learn the veracity for his action is only verified by the opposite intermediary. Besides, the opposite intermediary only conceals the news that aligns with the subscriber's ideology; it suggests that his action in response to no disclosure is optimal for the case of no news and any concealed news. Consequently, all the subscribers find it weakly better to subscribe to the opposite intermediary.

Example 4.3 reveals a coordination problem between the intermediaries' strategies and subscription choices when verification is costly and the intermediaries have an incentive to prove informativeness of aligned news (due to the perfect observability of verification results). Specifically, for an intermediary to verify aligned news, not only does she have to think that the news is likely to be informative, but also has to be followed by many subscribers on the opposite wing so that she is incentivized to persuade them and verify the news. This in turn justifies these opposing subscribers' anti-homophilic choices.

We remark that this anti-homophily among the extremists is not contradictory to our previous results that they make (weakly-)homophilic subscription choices. When a channel to perfectly convey the verification result is present, both intermediaries may verify the news of which the extremists would like to learn the veracity and the equilibrium subscription behavior relies on the belief system and the coordination between the intermediaries and the subscribers.

4.7.3 Utility function and action space

In this work, we consider a binary action space $\{-1, +1\}$ and the subscribers' utility functions to be the product of their own binary actions and the state of the world. An alternative for this modeling is the continuous action space and each subscriber aims to minimize the quadratic loss between his real-valued action and the state (see Che and Kartik 2009). This alternative modeling, however, leads to intractable analysis when an intermediary's welfare aggregation of the subscribers' utilities is evaluated. Specifically, the common approach to deal with the quadratic loss is to decompose it into two parts: the variance in her own belief and the aggregate quadratic loss between her optimal action and her subscribers'. As we have seen in Section 4.3 which studies players' belief updating, the functional form for an intermediary's optimal action can be quite different from her subscribers' for the case when the news is concealed; this adds much technical complexity to analysis of the model.¹⁴

4.8 Discussion

The goal of this chapter was to investigate how subscribers choose their information sources when the partisan news intermediaries are motivated to persuade the public opinion while selectively deciding what news to verify and to disclose. We proposed a framework for modeling the coexistence of true and false news and illustrated the variety in one's belief updating when observing news with unknown veracity. We showed that when verification option is not possible and the intermediaries can only decide whether to disclose the news, anti-homophily emerges among the centrist subscribers whereas the extremists feel indifferent between the intermediaries. We proved that, when verification option is available before any news disclosure, anti-homophily still arises among the centrists whereas the extremists make weakly-homophilic subscription choices. Furthermore, we demonstrated that the larger the intermediaries' ideological

¹⁴See Appendix C for our speculation on this variation of the model.

biases are, the wider the range of perspectives around the neutrality with which the subscribers exhibit anti-homophilic subscription behavior.

Chapter 5

Conclusion and Future Work

This Ph.D. thesis provided a few theories that rationalize individual behavior of communication on social network through a persuasion motive when circulating online information can be inaccurate or false. Specifically, this thesis focused on two aspects of people's behavior considering their persuasion motives underlying communication: (1) news sharing, (2) choice of news media.

For the first component, I proposed two theoretical frameworks to study people's decisions of sharing news. In Chapter 2 where I focused on non-strategic individuals, I studied the network effect on news cascade and identified the effects of perspective diversity and the wisdom of the crowd on the optimal levels of news precision that maximizes cascade probability. In Chapter 3 I presented a game-theoretic setting and showed that there is a natural network connectivity limit above which less credible information spreads faster. I further demonstrated how this network connectivity limit is affected by polarization and diversity of perspectives in a population.

In Chapter 4, I built a theory of strategic communication between news consumers and news media to understand the decision making underlying news consumers' selection of their information sources. I showed how centrists and extremists may make different choices in terms of ideological alignment while they ponder on news media's strategic moves of news verification and selective disclosure.

This thesis complemented the recent growing empirical literature on news consumption behavior of social media users. In particular, this thesis made a theoretical contribution to several disciplines: information diffusion on social networks, misinformation on social media, and political economy on news media. On the other hand, this thesis suggested several empirical predictions, such as the non-monotone effect of diversity of perspectives on connectivity limit as well as the difference in subscription choices between centrists and extremists, which are worth investigating empirically for future research.

Beyond this thesis, I describe a few interesting research directions below:

News spread on social networks with homophily in ideology

Chapters 2 and 3 both consider a homogeneous network structure in the sense that follower-followee relationships are independent from agents' individual beliefs (or ideology). However, much empirical evidence shows that people tend to build relationships with others who have similar opinions or personal characteristics. I am extending the settings in Chapter 3 to homophilic social networks and have derived some preliminary results.

Strategic fabrication of news

In this thesis, creation of false news (or inaccurate news) is an exogenous process and common knowledge to all individuals. This source of false news, however, can be strategic and endogenous. Specifically, a malicious principal aiming to influence the public opinion can strategically fabricate and disseminate misinformation or conspiracies on social media. Aware of this intention, the individuals on social media in turn strategically speculate on how much misinformation is circulating and adjust their beliefs about veracity of online news. This extension can give us some insights into what types of news may be fabricated (extreme or more neutral?) and the composition of real news and false news on social media.

Interaction between users and algorithms of platforms

The process of news spreading and belief manipulation is meddled by *news feed*

algorithms of social media platforms, which selectively determine what posts a user may see. In this regard, I plan to derive some principles about how a news feed algorithm should be designed to elude the interference of misinformation. Specifically, social media platforms have to choose whether they should detect suspiciously false news and even fact-check the information for their users before pushing the information to other users. If the platforms find some posts dubious, should they present the information to their users with warnings or conceal them for the users' sake? Given that verifying tons of dubious news requires lots of resources, is it possible to create a mechanism or use behavioral nudging to engage users in verifying news by themselves?

On the other hand, many platforms also select what online news media to recommend for their users (e.g. Apple News, Google News, or Facebook News). It is still unclear whether news recommendation algorithms deployed by these platforms lead to receiving information from diverse views or result in reinforcement of one's perspective and even polarization in a society. So far there has been little progress on managerial policies of the platforms and their effects on the public opinion.

Appendix A

Supplementary Material for Chapter 2

A more general distance between the beliefs: The distance (2.3) between beliefs P_1 and P_2 can be extended to include the square of differences in standard deviations as

$$W_2^2(P_1; P_2) = (\mu_1 - \mu_2)^2 + \gamma(\sigma_1 - \sigma_2)^2,$$

which also measures the information imbalance, where $\gamma \geq 0$ controls for the relative importance of two terms. Since a normally distributed belief is fully characterized by its mean and standard deviation, this will fully capture the distance between any two agents' beliefs.

While the discussions in the main body correspond to the case $\gamma = 0$ for simplicity, we present all the proofs for the generalized distance measure in (2.2). It is also worth noting that when $\gamma = 1$ the measure becomes what is known as the *Wasserstein distance* between two Gaussian distributions. The Wasserstein distance between two probability distributions μ, ν on real line is defined as

$$W_2^2(\mu, \nu) = \inf_{X \sim \mu, Y \sim \nu} \mathbb{E} \left[d(X, Y)^2 \right].$$

A.1 Proofs of Section 2.3

Proof of Lemma 2.1:

First, let us calculate the expected distance between the agent i and her followers if they all see the news.

$$\begin{aligned}\mathbb{E}_i \left[W_2^2(P_1^i; P_1^k) | k \in N^{\text{in}}(i) \right] &= \mathbb{E}_i \left[(1 - \beta)^2 (\mu_i - \mu_k)^2 | k \in N^{\text{in}}(i) \right], \\ &= (1 - \beta)^2 \mathbb{E}_i \left[(\mu_i - \bar{\mu} + \bar{\mu} - \mu_k)^2 | k \in N^{\text{in}}(i) \right], \\ &= (1 - \beta)^2 (\mu_i - \bar{\mu})^2 + (1 - \beta)^2 \sigma_\mu^2.\end{aligned}$$

Similarly, we can compute the expected distance if none of agent i 's followers receives the news.

$$\begin{aligned}\mathbb{E}_i \left[W_2^2(P_1^i; P_0^k) | k \in N^{\text{in}}(i) \right] &= \mathbb{E}_i \left[(\beta x + (1 - \beta) \mu_i - \mu_k)^2 + \gamma (\sqrt{1 - \beta} \sigma_\theta - \sigma_\theta)^2 | k \in N^{\text{in}}(i) \right], \\ &= \mathbb{E}_i \left[(\mu_i - \bar{\mu} + \bar{\mu} - \mu_k)^2 | k \in N^{\text{in}}(i) \right] + \gamma (1 - \sqrt{1 - \beta})^2 \sigma_\theta^2, \\ &= (\beta x + (1 - \beta) \mu_i - \bar{\mu})^2 + \sigma_\mu^2 + \gamma (1 - \sqrt{1 - \beta})^2 \sigma_\theta^2.\end{aligned}$$

By plugging in the corresponding values above and rearranging, we have

$$\begin{aligned}2 \underbrace{(1 - \beta)(\mu_i - \bar{\mu})}_{\text{interim bias of agent } i} &+ \underbrace{\beta(x - \bar{\mu})}_{\text{interim surprise}} + \underbrace{\beta^2(x - \bar{\mu})^2}_{\text{magnitude of surprise}} \\ &+ \underbrace{(1 - (1 - \beta)^2)\sigma_\mu^2}_{\text{reducing the variance of perspectives}} + \underbrace{\gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2}_{\text{reducing the difference in uncertainty}} > \frac{C}{1 - q}.\end{aligned}\tag{A.1}$$

The new term resulting from accounting for the information balance captures the broadcasting agent's incentive to inform her followers (as measured by the uncertainty of their beliefs) so that they are at the same informativeness level of the agent. Lemma 2.1 is

an immediate consequence of (A.1) when no consideration on information balance is involved. ■

A.2 Proofs of Section 2.4

Proof of Lemma 2.2 and Theorem 2.1:

For the generalized belief distance (2.2) that incorporates the variances of individual beliefs, we redefine the function $K_\beta(x)$ and $\eta_\beta(x)$ that appears in the broadcast size (2.9):

$$K_\beta(x) = \frac{C}{2\beta(1-\beta)|x-\bar{\mu}|\sigma_\mu},$$

$$\eta_\beta(x) = \frac{\sigma_\mu}{2(1-\beta)|x-\bar{\mu}|} \left(2 - \beta + \gamma \frac{1 - \sqrt{1 - \beta} \sigma_\theta^2}{1 + \sqrt{1 - \beta} \sigma_\mu^2} \right) + \frac{\beta|x-\bar{\mu}|}{2(1-\beta)\sigma_\mu}.$$

The following analysis is also applied to the case in our main context by making $\bar{\mu} = \gamma = 0$ in function $K_\beta(x)$ and $\eta_\beta(x)$. Since the credibility of news β and the realized news x are both given, we omit the argument of these terms and make the time variables as subscripts for simplicity. Additionally, given a set \mathcal{A}_t of agents, we use $\mathcal{A}_t(d) \subseteq \mathcal{A}_t$ as the subset of these agents with out-degree d . Given an edge $e \in \mathcal{E}$ whose tail has not yet seen the news by time t , let us define p_t as the probability of receiving it from $\text{Head}(e)$ at time $t+1$. That is,

$$p_t = \mathbb{P}\left(\text{Head}(e) \in b_t \mid \text{Tail}(e) \notin \mathcal{R}_t\right),$$

and write $\bar{p}_t = 1 - p_t$. An out-stub of an agent in $\bar{\mathcal{R}}_t$ can not be connected to any of the in-stubs of agents in $\bar{\mathcal{B}}_{t-1}$, but is equally likely to connect to in-stubs of agents in $\bar{\mathcal{B}}_{t-1}$. Given an edge $e \in \mathcal{E}$ with its tail in $\bar{\mathcal{R}}_t$, the probability of having its head in b_t

is thus given by the density of the in-stubs of b_t in $\bar{\mathcal{B}}_{t-1}^{\text{in}}$. Therefore,

$$p_t = \frac{|b_t^{\text{in}}|}{1 - |\mathcal{B}_{t-1}^{\text{in}}|},$$

using which and Bayes' rule update we can find

$$\begin{aligned} \bar{q}_t &= \frac{\mathbb{P}\left(\text{Tail}(e) \notin \mathcal{R}_t, \text{Head}(e) \in r_t\right)}{\mathbb{P}\left(\text{Head}(e) \in r_t\right)}, \\ &= \frac{\mathbb{P}\left(\text{Tail}(e) \notin \mathcal{R}_t\right)\mathbb{P}\left(\text{Head}(e) \in r_t \mid \text{Tail}(e) \notin \mathcal{R}_t\right)}{\mathbb{P}\left(\text{Head}(e) \in r_t, e \in \mathcal{E}\right)}, \\ &= \frac{|\bar{\mathcal{R}}_t^{\text{out}}| \times \frac{|r_t^{\text{in}}|}{1 - |\mathcal{B}_{t-1}^{\text{in}}|}}{|r_t^{\text{in}}|}, \\ &= \frac{|\bar{\mathcal{R}}_t^{\text{out}}|}{1 - |\mathcal{B}_{t-1}^{\text{in}}|}. \end{aligned} \tag{A.2}$$

From (2.10) it follows that

$$\mathbb{P}\left(\text{Head}(e) \notin b_t \mid \text{Tail}(e) \notin \mathcal{R}_t, e \in \mathcal{E}\right) = 1 - p_t = \frac{1 - |\mathcal{B}_t^{\text{in}}|}{1 - |\mathcal{B}_{t-1}^{\text{in}}|},$$

using which and by the independent following process in the configuration model we can show that

$$|\bar{\mathcal{R}}_{t+1}(d)| = \left(\frac{1 - |\mathcal{B}_t^{\text{in}}|}{1 - |\mathcal{B}_{t-1}^{\text{in}}|}\right)^d |\bar{\mathcal{R}}_t(d)|.$$

Applying this recursively along with $|\bar{\mathcal{R}}_0(d)| = P^{\text{out}}(d)$, we get

$$|\bar{\mathcal{R}}_{t+1}(d)| = (1 - |\mathcal{B}_t^{\text{in}}|)^d P^{\text{out}}(d).$$

As a result,

$$\begin{aligned}
|\bar{\mathcal{R}}_{t+1}| &= \sum_{d=0}^{\infty} (1 - |\mathcal{B}_t^{\text{in}}|)^d P^{\text{out}}(d), \\
|\bar{\mathcal{R}}_{t+1}^{\text{out}}| &= \frac{1}{\mathbb{E}[d]} \sum_{d=0}^{\infty} d (1 - |\mathcal{B}_t^{\text{in}}|)^d P^{\text{out}}(d), \\
|\bar{\mathcal{R}}_{t+1}^{\text{in}}| &= \frac{1}{\mathbb{E}[d]} \sum_{d=0}^{\infty} \mathbb{E}[\ell|d] (1 - |\mathcal{B}_t^{\text{in}}|)^d P^{\text{out}}(d).
\end{aligned}$$

We can simplify the expression for \bar{q}_t in (A.2) using the above to get

$$\begin{aligned}
\bar{q}_t &= \frac{\frac{1}{\mathbb{E}[d]} \sum_{d=1}^{\infty} d (1 - |\mathcal{B}_{t-1}^{\text{in}}|)^d P^{\text{out}}(d)}{1 - |\mathcal{B}_{t-1}^{\text{in}}|} \\
&= \frac{1}{\mathbb{E}[d]} \sum_{d=1}^{\infty} d (1 - |\mathcal{B}_{t-1}^{\text{in}}|)^{d-1} P^{\text{out}}(d).
\end{aligned}$$

This means that \bar{q}_t needed in the decision rule of Lemma 2.1, can be estimated from $|\mathcal{B}_{t-1}^{\text{in}}|$, which is the fraction of the edges through which the news has passed (i.e., covered by the broadcast) by time $t - 1$. This can be itself estimated from the size of the spread by time t from

$$|\bar{\mathcal{R}}_t| = \sum_{d=0}^{\infty} (1 - |\mathcal{B}_{t-1}^{\text{in}}|)^d P^{\text{out}}(d).$$

This proves Lemma 2.2.

The update rule for $|\mathcal{B}_t^{\text{in}}|$ can be obtained noting that $|\mathcal{B}_{t+1}^{\text{in}}| = |\mathcal{B}_t^{\text{in}}| + |b_{t+1}^{\text{in}}|$ and $|b_{t+1}^{\text{in}}| = \Phi(\eta - \frac{K}{\bar{q}_{t+1}}) |r_{t+1}^{\text{in}}|$, and that

$$\begin{aligned}
|r_{t+1}^{\text{in}}| &= |\bar{R}_t^{\text{in}}| - |\bar{R}_{t+1}^{\text{in}}| \\
&= \frac{1}{\mathbb{E}[d]} \sum_{d=0}^{\infty} \mathbb{E}[\ell|d] ((1 - |\mathcal{B}_{t-1}^{\text{in}}|)^d - (1 - |\mathcal{B}_t^{\text{in}}|)^d) P^{\text{out}}(d).
\end{aligned}$$

This leads to

$$|\mathcal{B}_{t+1}^{\text{in}}| - |\mathcal{B}_t^{\text{in}}| = \frac{\Phi(\eta - \frac{K}{\bar{q}_{t+1}})}{\mathbb{E}[d]} \sum_{d=0}^{\infty} \mathbb{E}[\ell|d] \left((1 - |\mathcal{B}_{t-1}^{\text{in}}|)^d - (1 - |\mathcal{B}_t^{\text{in}}|)^d \right) P^{\text{out}}(d).$$

To approximate the asymptotic behavior of the spread, we now use the continuous-time mean-field approximation of the above second-order nonlinear dynamics. More specifically, we approximate $|\mathcal{B}_t^{\text{in}}|$ with a continuous-time process y_t with $\dot{y}_t \approx |\mathcal{B}_t^{\text{in}}| - |\mathcal{B}_{t-1}^{\text{in}}|$ and $\ddot{y}_t \approx |\mathcal{B}_{t+1}^{\text{in}}| - 2|\mathcal{B}_t^{\text{in}}| + |\mathcal{B}_{t-1}^{\text{in}}|$. The dynamics of y_t is given by

$$\ddot{y}_t = \dot{y}_t \times \left(-1 + \frac{\Phi(\eta - \frac{K}{\bar{q}(y_t)})}{\mathbb{E}[d]} \sum_{d=0}^{\infty} d \mathbb{E}[\ell|d] (1 - y_t)^{d-1} P^{\text{out}}(d) \right), \quad (\text{A.3})$$

where

$$\bar{q}(y) = \frac{1}{\mathbb{E}[d]} \sum_{d=1}^{\infty} d (1 - y)^{d-1} P^{\text{out}}(d).$$

We can solve (A.3) using a change of variables $z_t = \dot{y}_t$:

$$\frac{dz_t}{dy_t} = \frac{\ddot{y}_t}{\dot{y}_t} = \left(-1 + \frac{\Phi(\eta - \frac{K}{\bar{q}(y_t)})}{\mathbb{E}[d]} \sum_{d=0}^{\infty} d \mathbb{E}[\ell|d] (1 - y_t)^{d-1} P^{\text{out}}(d) \right),$$

leading to

$$\dot{y}_t - \dot{y}(0) = \int_{y(0)}^{y_t} \left(-1 + \frac{\Phi(\eta - \frac{K}{\bar{q}(y)})}{\mathbb{E}[d]} \sum_{d=0}^{\infty} d \mathbb{E}[\ell|d] (1 - y)^{d-1} P^{\text{out}}(d) \right) dy,$$

where the initial conditions are $\dot{y}(0) = y(0) = 0$. Define

$$g(y) = -1 + \frac{\Phi(\eta - \frac{K}{\bar{q}(y)})}{\mathbb{E}[d]} \sum_{d=0}^{\infty} d \mathbb{E}[\ell|d] (1 - y)^{d-1} P^{\text{out}}(d).$$

Then, $y = 0$ is unstable if $g(0) > 0$, or equivalently

$$\frac{\mathbb{E}[\ell d]}{\mathbb{E}[d]} \Phi(\eta - K) > 1,$$

and is stable if $g(0) < 0$, or

$$\frac{\mathbb{E}[\ell d]}{\mathbb{E}[d]} \Phi(\eta - K) < 1.$$

For the case $g(0) = 0$, we need to look at $\frac{d}{dy}g(y)|_{y=0}$. It is easy to verify that $\frac{d}{dy}g(y)|_{y=0} < 0$, implying the stability of $y = 0$ in this case. The above analysis proves (2.12), which states that a news cascade happens if and only if

$$\frac{\mathbb{E}[\ell d]}{\mathbb{E}[d]} \Phi(\eta - K) > 1.$$

Moreover, when a news cascade emerges, $|\mathcal{B}_\infty^{\text{in}}|$ is the unique nonzero solution to $G(|\mathcal{B}_\infty^{\text{in}}|) = 0$, where

$$G(|\mathcal{B}_\infty^{\text{in}}|) = \int_0^{|\mathcal{B}_\infty^{\text{in}}|} \left(-1 + \frac{\Phi(\eta - \frac{K}{\bar{q}(y)})}{\mathbb{E}[d]} \sum_{d=0}^{\infty} d \mathbb{E}[\ell|d] (1-y)^{d-1} P^{\text{out}}(d) \right) dy,$$

if $G(1) < 0$; otherwise, $|\mathcal{B}_\infty^{\text{in}}| = 1$. The steady state size of the spread, denoted by $|\mathcal{R}_\infty|$, can be calculated from $|\mathcal{B}_\infty^{\text{in}}|$ according to

$$|\mathcal{R}_\infty| = 1 - \sum_{d=0}^{\infty} (1 - |\mathcal{B}_\infty^{\text{in}}|)^d P^{\text{out}}(d).$$

■

Proof of Proposition 2.1: This is already proved as part of Theorem 2.1. ■

Proof of Lemma 2.3:

- (i) Given non-negative parameters λ_{01} , λ_{10} and λ_{11} , the bivariate Poisson distribu-

tion (see Kawamura (1973)) is defined as

$$\mathbb{P}_{\text{Pois}}(\ell, d) = \sum_{m=0}^{\min\{\ell, d\}} \frac{\lambda_{10}^{d-m} \lambda_{01}^{\ell-m} \lambda_{11}^m}{(d-m)! (\ell-m)! m!} e^{-(\lambda_{10} + \lambda_{01} + \lambda_{11})},$$

with the following statistics

$$\begin{aligned} \mathbb{E}_{\text{Pois}}[\ell] &= \sigma_{\ell, \text{Pois}}^2 = \lambda_{10} + \lambda_{11}, \\ \mathbb{E}_{\text{Pois}}[d] &= \sigma_{d, \text{Pois}}^2 = \lambda_{01} + \lambda_{11}, \\ \mathbb{E}_{\text{Pois}}[\ell d] &= (\lambda_{10} + \lambda_{11})(\lambda_{01} + \lambda_{11}) + \lambda_{11}. \end{aligned}$$

Degree distributions should satisfy $\mathbb{E}_{\text{Pois}}[\ell] = \mathbb{E}_{\text{Pois}}[d]$, requiring that $\lambda_{10} = \lambda_{01} = \lambda$ and $\lambda_{11} = \lambda_c$. The joint Poisson degree distribution (2.14) then follows, with the statistics

$$\begin{aligned} \mu_P &:= \mathbb{E}_{\text{Pois}}[\ell] = \mathbb{E}_{\text{Pois}}[d] = \lambda + \lambda_c \\ \sigma_P^2 &:= \sigma_{\ell, \text{Pois}}^2 = \sigma_{d, \text{Pois}}^2 = \lambda + \lambda_c \\ \rho_P &:= \frac{\mathbb{E}_{\text{Pois}}[\ell d] - \mu_P^2}{\sigma_P^2} = \frac{\lambda_c}{\lambda + \lambda_c} \in [0, 1]. \end{aligned}$$

Using the relation (2.13), we obtain

$$\mu_P^{\text{Line}} = \frac{\mathbb{E}_{\text{Pois}}[\ell d]}{\mathbb{E}_{\text{Pois}}[d]} = \frac{\rho_P \sigma_P^2}{\mu_P} + \mu_P = \rho_P + \mu_P, \quad (\text{A.4})$$

where $\rho_P \in [0, 1]$ and $\mu_P > 0$.

(ii) Given positive parameters k_1, k_2 , and α , the bivariate Zipf distribution is defined in terms of its complement CDF (see (Yeh 2002)):

$$\mathbb{P}_{\text{Zipf}}(\ell \geq m_1, d \geq m_2) = \left(1 + \frac{m_1}{k_1} + \frac{m_2}{k_2}\right)^{-\alpha}, m_1, m_2 = 0, 1, \dots$$

with the following statistics

$$\begin{aligned}
\mathbb{E}_{\text{Zipf}}[\ell] &= \sum_{m_1 \geq 1} \left(1 + \frac{m_1}{k_1}\right)^{-\alpha}, \\
\sigma_{\ell, \text{Zipf}}^2 &= \sum_{m_1 \geq 1} (2m_1 - 1) \left(1 + \frac{m_1}{k_1}\right)^{-\alpha} - (\mathbb{E}_{\text{Zipf}}[\ell])^2, \\
\mathbb{E}_{\text{Zipf}}[d] &= \sum_{m_2 \geq 1} \left(1 + \frac{m_2}{k_2}\right)^{-\alpha}, \\
\sigma_{d, \text{Zipf}}^2 &= \sum_{m_2 \geq 1} (2m_2 - 1) \left(1 + \frac{m_2}{k_2}\right)^{-\alpha} - (\mathbb{E}_{\text{Zipf}}[d])^2, \\
\mathbb{E}_{\text{Zipf}}[\ell d] &= \sum_{m_1 \geq 1} \sum_{m_2 \geq 1} \left(1 + \frac{m_1}{k_1} + \frac{m_2}{k_2}\right)^{-\alpha}.
\end{aligned}$$

The requirement $\mathbb{E}_{\text{Zipf}}[\ell] = \mathbb{E}_{\text{Zipf}}[d]$ for degree distributions enforces $k_1 = k_2 = k$, resulting in the joint Zipf degree distribution (2.15). Similarly, we derive the equations for the statistics

$$\begin{aligned}
\mu_Z &:= \mathbb{E}_{\text{Zipf}}[\ell] = \mathbb{E}_{\text{Zipf}}[d] = \sum_{m \geq 1} \left(1 + \frac{m}{k}\right)^{-\alpha} \\
\sigma_Z^2 &= \sum_{m \geq 1} (2m - 1) \left(1 + \frac{m}{k}\right)^{-\alpha} - \mu_Z^2 \\
\mathbb{E}_{\text{Zipf}}[\ell d] &= \sum_{m_1 \geq 1} \sum_{m_2 \geq 1} \left(1 + \frac{m_1 + m_2}{k}\right)^{-\alpha} = \sum_{m \geq 1} (m - 1) \left(1 + \frac{m}{k}\right)^{-\alpha}.
\end{aligned}$$

Observe, from the above, that

$$\mathbb{E}_{\text{Zipf}}[\ell d] = \frac{\sigma_Z^2 + \mu_Z^2 - \mu_Z}{2}.$$

Therefore,

$$\rho_Z := \frac{\mathbb{E}_{\text{Zipf}}[\ell d] - \mu_Z^2}{\sigma_Z^2} = \frac{\sigma_Z^2 - \mu_Z^2 - \mu_Z}{2\sigma_Z^2}, \tag{A.5}$$

which is smaller than $\frac{1}{2}$ because $\mu_Z > 0$. Moreover, observe that

$$\begin{aligned}\sigma_Z^2 - \mu_Z^2 - \mu_Z &= \sum_{m \geq 1} (2m - 2) \left(1 + \frac{m}{k}\right)^{-\alpha} - 2\mu_Z^2, \\ &= 2 \left[\sum_{m \geq 1} (m - 1) \left(1 + \frac{m}{k}\right)^{-\alpha} - \mu_Z^2 \right], \\ &= 2\mathbb{E}_{\text{Zipf}}[\ell d] > 0,\end{aligned}$$

and therefore $\rho_Z > 0$. By (A.5), σ_Z^2 can be written in terms of μ_Z and ρ_Z ,

$$\sigma_Z^2 = \frac{\mu_Z^2 + \mu_Z}{1 - 2\rho_Z},$$

and by (2.13), we derive that for $\rho_Z \in (0, 1/2)$ and $\mu_Z > 0$,

$$\mu_Z^{\text{Line}} = \frac{\mathbb{E}_{\text{Zipf}}[\ell d]}{\mathbb{E}_{\text{Zipf}}[d]} = \frac{\rho_Z \sigma_Z^2}{\mu_Z} + \mu_Z = \frac{\rho_Z(\mu_Z + 1)}{1 - 2\rho_Z} + \mu_Z. \quad (\text{A.6})$$

This completes the proof. ■

Proof of Proposition 2.2:

Part (i) is an immediate result from (A.4) and (A.6) with $\mu_P = \mu_Z > 0$ and $\rho_P = \rho_Z \in (0, 1/2)$ and part (ii) follows according to Proposition 2.1. ■

A.3 Proof of Section 2.5

Proof of Theorem 2.2:

We state the general form of the theorem for the general measure of distance between individual beliefs defined in (2.2).

Theorem A.1 (*General version of Theorem 2.2*) *Given the variance of the perspectives σ_μ^2 , denote with $\Omega^*(\sigma_\mu^2)$ the set of values β that maximize the likelihood of a cascade as formulated in (2.16). Assume that $\mu^{\text{Line}}(\mathcal{G}) > 2$ and let $\Delta = \Phi^{-1}\left(\frac{1}{\mu^{\text{Line}}(\mathcal{G})}\right)$. Then,*

(i) If $\sigma_\mu^2 + \gamma\sigma_\theta^2 > C$, then $\Omega^*(\sigma_\mu) = [\bar{\beta}^*(\sigma_\mu^2), 1]$ and a cascade emerges almost surely, where $\bar{\beta}^*(\sigma_\mu^2)$ satisfies

$$((1 - \beta)^2 - 1)\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C = 0.$$

Furthermore, $\bar{\beta}^*(\sigma_\mu^2)$ is strictly decreasing in σ_μ^2 for any $\gamma \geq 0$.

(ii) If $\sigma_\mu^2 + \gamma\sigma_\theta^2 \leq C < \sigma_\mu^2 + \gamma\sigma_\theta^2 + (\theta - \bar{\mu})^2$, then $\Omega^*(\sigma_\mu^2) = \{1\}$ and a cascade emerges almost surely.

(iii) If $\sigma_\mu^2 + \gamma\sigma_\theta^2 + (\theta - \bar{\mu})^2 \leq C$, then $1 \notin \Omega^*(\sigma_\mu^2)$ and particularly, the truth never causes a cascade.

Below, we prove the theorem in its general form.

According to (2.16), a realization of the news x results in a cascade if and only if

$$-\frac{C}{2\beta(1 - \beta)|x - \bar{\mu}|\sigma_\mu} + \left(2 - \beta + \gamma \frac{1 - \sqrt{1 - \beta}\sigma_\theta^2}{1 + \sqrt{1 - \beta}\sigma_\mu^2}\right) \frac{\sigma_\mu}{2(1 - \beta)|x - \bar{\mu}|} + \frac{\beta}{2(1 - \beta)\sigma_\mu}|x - \bar{\mu}| > \Delta,$$

or equivalently $V(|x - \bar{\mu}|) > 0$, where we define

$$V(z) = \beta z^2 - 2(1 - \beta)\Delta\sigma_\mu z + \gamma \frac{1 - \sqrt{1 - \beta}\sigma_\theta^2}{1 + \sqrt{1 - \beta}\sigma_\mu^2} - \frac{C}{\beta} + (2 - \beta)\sigma_\mu^2. \quad (\text{A.7})$$

The discriminant of the above quadratic polynomial is

$$\begin{aligned} \Gamma(\beta, \sigma_\mu^2) &= (1 - \beta)^2\Delta^2\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C - \sigma_\mu^2 + (1 - \beta)^2\sigma_\mu^2 \quad (\text{A.8}) \\ &= ((1 - \beta)^2(\Delta^2 + 1) - 1)\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C. \end{aligned}$$

If the discriminant $\Gamma(\beta, \sigma_\mu^2) \leq 0$, then $V(|x - \bar{\mu}|) > 0$ holds for almost all x (where $V(z)$ is defined in (A.7)). This implies that a cascade happens almost surely when for news having credibility β . It is important, however, to note that this is not the only

possibility for a cascade to happen almost surely: For a value of β with $\Gamma(\beta, \sigma_\mu^2) > 0$, it still holds that $V(|x - \bar{\mu}|) > 0$ almost surely (or, equivalently $V(z) > 0$ for $z > 0$) if and only if the larger root of (A.7) is non-positive. This is translated to $(1 - \beta)\Delta\sigma_\mu + \sqrt{\Gamma(\beta, \sigma_\mu^2)} \leq 0$, which requires both $\Delta < 0$ and $\Gamma^{<0}(\beta, \sigma_\mu^2) \leq 0$ (and $\Gamma(\beta, \sigma_\mu^2) > 0$), where

$$\Gamma^{<0}(\beta, \sigma_\mu^2) = ((1 - \beta)^2 - 1)\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C. \quad (\text{A.9})$$

Observe that the discriminant functions (A.8) and (A.9) are both decreasing in β and that $\Gamma(\beta, \sigma_\mu^2) \geq \Gamma^{<0}(\beta, \sigma_\mu^2)$ for any β and any Δ .

First consider the case when $\Gamma(1, \sigma_\mu^2) < 0$, i.e., $\sigma_\mu^2 + \gamma\sigma_\theta^2 > C$ (cf. case (i)). Due to the monotonicity in β , we derive $\Gamma(\beta, \sigma_\mu^2) \leq 0$ for $\beta \in [\bar{\beta}_1^*, 1]$ where $\bar{\beta}_1^*$ is the unique solution to $\Gamma(\bar{\beta}_1^*, \sigma_\mu^2) = 0$ and a cascade emerges with probability 1 for credibility in this range. If $\Delta < 0$, the other possibility requires $\Gamma^{<0}(\beta, \sigma_\mu^2) \leq 0$ and $\Gamma(\beta, \sigma_\mu^2) > 0$, and there is an almost always cascade for $\beta \in [\bar{\beta}_2^*, \bar{\beta}_1^*]$ where $\bar{\beta}_2^*$ is the unique solution to $\Gamma^{<0}(\bar{\beta}_2^*, \sigma_\mu^2) = 0$. The relation $\bar{\beta}_2^* \leq \bar{\beta}_1^*$ is due to the inequality $\Gamma(\beta, \sigma_\mu^2) \geq \Gamma^{<0}(\beta, \sigma_\mu^2)$. As a result, for networks with $\mu^{\text{Line}}(\mathcal{G}) > 2$ ($\Delta < 0$), the values of β for which a cascade happens almost surely are in form of $\beta \in [\bar{\beta}^*, 1]$, where $\bar{\beta}^*$ is the unique solution to $\Gamma^{<0}(\bar{\beta}^*, \sigma_\mu^2) = 0$. For any $\beta \notin [\bar{\beta}^*, 1]$ defined as above, the cascade probability would be less than 1 and hence not optimal; the calculation for the corresponding cascade probability would follow the same derivation in the next case.

Moreover, for the unique solution $\bar{\beta}^*(\sigma_\mu^2)$ to $\Gamma^{<0}(\bar{\beta}^*, \sigma_\mu^2) = 0$ when $\Delta < 0$, it is easy to verify that for any $\gamma \geq 0$

$$\frac{d\bar{\beta}^*(\sigma_\mu^2)}{d\sigma_\mu} = \frac{2((1 - \bar{\beta}^*)^2 - 1)\sigma_\mu}{2(1 - \bar{\beta}^*)\sigma_\mu^2 + \gamma\sigma_\theta^2 \frac{1 - \sqrt{1 - \bar{\beta}^*}}{\sqrt{1 - \bar{\beta}^*}}} \leq 0,$$

and therefore the lower bound $\bar{\beta}^*(\sigma_\mu^2)$ is decreasing in σ_μ^2 .

Now we consider the case $\Gamma(1, \sigma_\mu^2) \geq 0$, that is, $\sigma_\mu^2 + \gamma\sigma_\theta^2 \leq C$. It then follows that

$\Gamma(\beta, \sigma_\mu^2) \geq \Gamma^{<0}(\beta, \sigma_\mu^2) \geq C - \sigma_\mu^2 - \gamma\sigma_\theta^2 \geq 0$ for all $\beta \in [0, 1]$, resulting in one non-negative root and the other non-positive root for (A.7). Therefore, the cascade-triggering news x satisfies

$$|x - \bar{\mu}| > \frac{(1 - \beta)\Delta\sigma_\mu + \sqrt{((1 - \beta)^2(\Delta^2 + 1) - 1)\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C}}{\beta} := h(\beta, \sigma_\mu^2), \quad (\text{A.10})$$

where $h(\beta, \sigma_\mu^2) \geq 0$, the non-negative root. The probability of a cascade emerging (2.16) can therefore be written as a function of β :

$$\begin{aligned} P^{\text{cascade}}(\beta) &\triangleq 1 - \Phi\left(\frac{(1 - \beta)\Delta\sigma_\mu + \sqrt{((1 - \beta)^2(\Delta^2 + 1) - 1)\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C}}{\beta} - (\theta - \bar{\mu})\right) \\ &\quad - \Phi\left(\frac{(1 - \beta)\Delta\sigma_\mu + \sqrt{((1 - \beta)^2(\Delta^2 + 1) - 1)\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C}}{\beta} - (\theta - \bar{\mu})\right) \\ &= 1 - \Phi\left(\frac{(1 - \beta)\Delta\sigma_\mu + \sqrt{((1 - \beta)^2(\Delta^2 + 1) - 1)\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C}}{\sqrt{\beta(1 - \beta)}} - (\theta - \bar{\mu})\sqrt{\frac{\beta}{1 - \beta}}\right) \\ &\quad + \Phi\left(\frac{(1 - \beta)\Delta\sigma_\mu + \sqrt{((1 - \beta)^2(\Delta^2 + 1) - 1)\sigma_\mu^2 - \gamma(1 - \sqrt{1 - \beta})^2\sigma_\theta^2 + C}}{\sqrt{\beta(1 - \beta)}} - (\theta - \bar{\mu})\sqrt{\frac{\beta}{1 - \beta}}\right), \end{aligned}$$

which is strictly smaller than 1 for $\beta \in (0, 1)$. From the above formula, we can observe that as $\beta \rightarrow 0$, the probability of cascade $P^{\text{cascade}}(\beta) \rightarrow 1 - \Phi(+\infty) + \Phi(-\infty) = 0$. As

$\beta \rightarrow 1$, $P^{\text{cascade}}(\beta)$ has the following asymptotic:

$$P^{\text{cascade}}(\beta) \approx 1 - \Phi\left(\frac{\sqrt{C - \sigma_\mu^2 - \gamma\sigma_\theta^2} - (\theta - \bar{\mu})}{\sqrt{1 - \beta\sigma_\theta}}\right) + \Phi\left(\frac{-\sqrt{C - \sigma_\mu^2 - \gamma\sigma_\theta^2} - (\theta - \bar{\mu})}{\sqrt{1 - \beta\sigma_\theta}}\right).$$

Using this equation, we can see that as $\beta \rightarrow 1$,

Case (ii) $P^{\text{cascade}}(\beta) \rightarrow 1$ if $|\theta - \bar{\mu}| > \sqrt{C - \sigma_\mu^2 - \gamma\sigma_\theta^2}$, and

Case (iii) $P^{\text{cascade}}(\beta) \rightarrow 0$ if $|\theta - \bar{\mu}| \leq \sqrt{C - \sigma_\mu^2 - \gamma\sigma_\theta^2}$.

This completes the proof. ■

Proof of Proposition 2.3:

Proposition 2.3 needs no modification for the general measure of distance between individual beliefs (2.2).

Recall that for networks with $\mu^{\text{Line}}(\mathcal{G}) > 2$, or $\Delta < 0$, Theorem 2.2 already shows that, when $\sigma_\mu^2 + \gamma\sigma_\theta^2 > C$, the cascade probability is 1 for $\beta \in [\bar{\beta}^*(\sigma_\mu^2), 1]$ where the endpoint $\bar{\beta}^*(\sigma_\mu^2)$ is the unique solution to $\Gamma^{<0}(\beta, \sigma_\mu^2) = 0$ (A.9) and is strictly decreasing in σ_μ^2 . It means that after σ_μ^2 is increased, the news of credibility $\beta \in [\bar{\beta}^*(\sigma_\mu^2), 1]$ still cascades almost surely. Now it remains to examine whether the cascade probability is increasing (a) for $\beta < \bar{\beta}^*(\sigma_\mu^2)$ and $\sigma_\mu^2 + \gamma\sigma_\theta^2 > C$ and (b) when $\sigma_\mu^2 + \gamma\sigma_\theta^2 \leq C$.

According to the proof of Theorem 2.2, both cases (a) and (b) will result in one non-negative root and the other non-positive root for (A.7) (and $\Gamma(\beta, \sigma_\mu^2) \geq 0$), meaning that the quantity $h(\beta, \sigma_\mu^2)$ would be non-negative in (A.10). That is, news x would cascade if and only if $|x - \bar{\mu}| > h(\beta, \sigma_\mu^2) \geq 0$. Now we would like to show that in case (a) and (b) where $h(\sigma_\mu^2) \geq 0$, the function $h(\sigma_\mu^2)$ decreases in σ_μ^2 , or equivalently decreases in σ_μ . Consequently, the set of news realizations that generate a cascade would be enlarged, resulting in higher cascade probability.

By simple calculations, we can find that

$$\sigma_\mu \frac{\partial h(\beta, \sigma_\mu^2)}{\partial \sigma_\mu} = h(\beta, \sigma_\mu^2) - \frac{C - \gamma(1 - \sqrt{1 - \beta})^2 \sigma_\theta^2}{\beta \sqrt{\Gamma(\beta, \sigma_\mu^2)}}.$$

We further rewrite $h(\beta, \sigma_\mu^2)$ as

$$h(\beta, \sigma_\mu^2) = \frac{C - \gamma(1 - \sqrt{1 - \beta})^2 \sigma_\theta^2 - \beta(2 - \beta)\sigma_\mu^2}{\beta(\sqrt{\Gamma} - (1 - \beta)\Delta\sigma_\mu)}.$$

For $\Delta < 0$ (i.e., $\mu^{\text{Line}}(\mathcal{G}) > 2$) and $h(\beta, \sigma_\mu^2) \geq 0$, we can see that

$$0 \leq \frac{C - \gamma(1 - \sqrt{1 - \beta})^2 \sigma_\theta^2 - \beta(2 - \beta)\sigma_\mu^2}{\beta(\sqrt{\Gamma} - (1 - \beta)\Delta\sigma_\mu)} \leq \frac{C - \gamma(1 - \sqrt{1 - \beta})^2 \sigma_\theta^2}{\beta(\sqrt{\Gamma} - (1 - \beta)\Delta\sigma_\mu)} \leq \frac{C - \gamma(1 - \sqrt{1 - \beta})^2 \sigma_\theta^2}{\beta\sqrt{\Gamma}}.$$

Thus, for $\Delta < 0$ and for any β , $\frac{\partial h(\beta, \sigma_\mu^2)}{\partial \sigma_\mu} \leq 0$ in case (a) and (b), and the cascade probability is increasing in σ_μ (also increasing in σ_μ^2). With the result for the case when $\beta \in [\bar{\beta}^*(\sigma_\mu^2), 1]$ and $\sigma_\mu^2 + \gamma\sigma_\theta^2 > C$, the proof is complete. \blacksquare

Appendix B

Supplementary Material for Chapter 3

B.1 Proofs of Section 3.3

Proof of Lemma 3.2.

Denote by Q the probability that a follower receives the news at the steady state from other agents according to the strategy profile s ; we reason similarly as how the equation for q is derived. However, we need to adjust two elements: (i) a follower with out-degree d has not received it from any of her “other” followees, conditioned on the event that the decision maker does not share the news; (ii) the out-degree distribution of a follower should be adjusted by the well-known *friendship paradox*. We have:

$$1 - Q = (1 - \delta) \sum_{d=1}^{\infty} \frac{dP^{\text{out}}(d)}{\mathbb{E}[d]} (1 - qP_F(s((x, \beta))))^{d-1}.$$

Using $P^{\text{out}}(\cdot) \sim \text{Poisson}(\lambda)$ and taking $\delta \rightarrow 0$, we obtain that at the steady state $Q = q$. ■

Proof of Lemma 3.3:

According to (3.2), we obtain

$$\begin{aligned}
& \mathbb{E}_i[a_k | k \in N_i^{\text{in}}, z_k = (x, \beta)] \\
&= \mathbb{P}[a_k(\mu_k, (x, \beta)) = +1 | \mu_k \sim F] - \mathbb{P}_{\mu_k \sim F}[a_k(\mu_k, (x, \beta)) = -1 | \mu_k \sim F] \\
&= 1 - 2F\left(\frac{-\beta x}{1 - \beta}\right),
\end{aligned}$$

and similarly,

$$\mathbb{E}_i[a_k | k \in N_i^{\text{in}}, z_k = \emptyset] = 1 - 2F(0).$$

Plug them into (3.7) and the sharing rule (3.8) is the immediate result. ■

Proof of Theorem 3.1:

WLOG, we discuss the case $x \geq 0$. The equilibrium when $\delta \rightarrow 0$ should satisfy the following system of equations:

$$\begin{aligned}
q^* &= 1 - e^{-q^* \lambda P_F(s^*)}, \\
P_F(s^*) &= 1 - F\left(\frac{\frac{C}{2(1-q^*)F^{\text{HIP}}(x, \beta)} - \beta x}{1 - \beta}\right)
\end{aligned}$$

To prove the existence and uniqueness of an equilibrium, we define a continuous function $h(q)$ for $q \in [0, 1]$:

$$\begin{aligned}
h(q) &\triangleq 1 - e^{-q \lambda P_F(s)} - q, \\
P_F(s) &= 1 - F\left(\frac{\frac{C}{2(1-q)F^{\text{HIP}}(x, \beta)} - \beta x}{1 - \beta}\right)
\end{aligned} \tag{B.1}$$

By the continuity and $h(0) = 0$ and $h(1) = -1$ (because $P_F(s) = 0$ for $q = 1$), there must exist at least one such solution to (B.1), proving the existence of an equilibrium.

Though $q = 0$ is always a solution, the zero spread size may not be an equilibrium

for infinitesimal seeding size δ when there are other positive solutions (see Jackson (2008) for details). We now identify the condition under which there are other positive solutions to $h(q)$, i.e., a news cascade emerges, and then show the uniqueness of the equilibrium.

Emergence of a news cascade:

Take derivative (B.1) of w.r.t. q ,

$$h'(q) = \lambda e^{-\lambda q P_F(s)} \left(P_F(s) + q \frac{dP_F(s)}{dq} \right) - 1.$$

Note that for $\beta < 1$, P_F is differentiable in the range $[0, 1]$ (for $x = 0$ or $\beta = 0$, $P_F(s)$ is the constant zero). As $h(0) = 0, h(1) = -1$, if $h'(0) = -1 + \lambda P_F(s)|_{q=0} > 0$, due to the continuity, some positive solutions must exist.

We now show that $h'(0) > 0$ is also necessary. Noting that the derivative $\frac{dP_F(s)}{dq}$ must be non-positive due to the strategic substitutability and $P_F(s)$ is maximized at $q = 0$. Consequently, we can find that when $h'(0) \leq 0$, then $h'(q) < 0$ for any $q > 0$, indicating that $h(q)$ is strictly decreasing and $q = 0$ is the only root.

Uniqueness of the equilibrium:

We have shown that if $\lambda P_F(s)|_{q=0} \leq 1$, $q^* = 0$ is the only solution and $P_F^*, \mu^{\text{th}*}$ are in turn determined, resulting in a unique equilibrium. If $\lambda P_F(s)|_{q=0} > 1$, there are other positive solutions and $q^* \neq 0$ (Jackson 2008); for this case, there is only one positive solution and the equilibrium is therefore unique. To see this, suppose there are two positive solutions $q^1 > q^2 > 0$. Then we have $P_F(s)|_{q=q^1} < P_F(s)|_{q=q^2}$ due to strategic substitutability. However, the equation $q = 1 - e^{-q\lambda P_F(s)}$, or $P_F(s) = \frac{-\ln(1-q)}{\lambda q}$, dictates that the solutions should satisfy $P_F(s)|_{q=q^1} > P_F(s)|_{q=q^2}$. Contradiction! As a result, only one positive solution for q exists, with $P_F^*, \mu^{\text{th}*}$ in turn determined, and the equilibrium is unique. ■

Proof of Theorem 3.2:

For convenience, define $\tilde{q}(x) \triangleq 1 - \frac{C}{|x|}$. Given λ, F , for $\beta = 1$, according to

Lemma 3.3, we can see that $P_F(s(x, 1)) = 1$ for any $q < \tilde{q}(x)$ while $P_F(s(x, 1)) = 0$ for any $q > \tilde{q}(x)$. We keep the argument news x and β only as a shorthand. Note that the cascade size is bounded above by the size of giant component $q^G(\lambda)$ when every agent shares news ($P_F \equiv 1$). We have three cases depending on the magnitude of x :

- 1) $|x| \in [0, C]$: $\tilde{q}(x, 1) \leq 0$ and then $P_F(x, 1) = 0$ for any $q \in (0, 1] \Rightarrow q^*(x, 1) = 0, P_F^*(x, 1) = 0$.
- 2) $|x| \in (C, \frac{C}{1-q^G(\lambda)})$: In this case $\tilde{q}(x) \in [0, q^G(\lambda)]$. Observe that $h(q) > 0$ for $q < \tilde{q}(x)$ and $h(q) < 0$ for $q > \tilde{q}(x)$. Let $P_F(x, 1)|_{q=\tilde{q}(x)} = \frac{-\ln(1-\tilde{q}(x))}{\lambda\tilde{q}(x)}$, and it can be shown that $P_F(x, 1) \in [0, 1]$ and $h(\tilde{q}(x)) = 0$. Therefore, the equilibrium is $q^*(x, 1) = \tilde{q}(x)$ and $P_F^*(x, 1) = \frac{-\ln(1-\tilde{q}(x))}{\lambda\tilde{q}(x)}$.
- 3) $|x| > \frac{C}{1-q^G(\lambda)}$: $\tilde{q}(x) > q^G(\lambda)$ and then $P_F(x) = 1$ for any $q \leq q^G(\lambda)$. It is obvious that $q^*(x, 1) = q^G(\lambda), P_F^*(x, 1) = 1$.

The equilibrium sharing threshold $\mu^{\text{th}*}(x, 1)$ is obtained by applying inverse of CDF F^{-1} to the sharing fraction $P_F^*(x, 1)$. The existence of an equilibrium with is then obvious. ■

Proof of Corollary 3.1:

When $\beta < 1$, it is easy to see that for any news x , $q^*(x, \beta)$, $P_F^*(x, \beta)$ and $\mu^{\text{th}*}(x, \beta)$ are continuous in $\beta \in [0, 1)$. We now show that they are continuous at $\beta = 1$, i.e., their limits when $\beta \rightarrow 1$ are equal to their corresponding values at $\beta = 1$.

- 1) $|x| \leq [0, C]$: If $\lim_{\beta \rightarrow 1} q^*(x, \beta) > 0$, then for $x \geq 0$

$$\lim_{\beta \rightarrow 1} \mu^{\text{th}*}(x, \beta) = \lim_{\beta \rightarrow 1} \frac{\frac{C}{2(1-q^*)(F(0)-F(\frac{-\beta x}{1-\beta}))} - \beta x}{1 - \beta} = \lim_{\beta \rightarrow 1} \frac{\frac{C}{1-q^*} - x}{1 - \beta} = +\infty,$$

and hence $\lim_{\beta \rightarrow 1} P_F^*(x, \beta) = 0$, resulting in $\lim_{\beta \rightarrow 1} q^*(x, \beta) = 0$. Then obviously $\lim_{\beta \rightarrow 1} q^*(x, \beta) = 0$ and $\lim_{\beta \rightarrow 1} P_F^*(x, \beta) = 0$. The reasoning for $x < 0$ is similar.

- 2) $|x| \in (C, \frac{C}{1-q^G(\lambda)}]$: If $\lim_{\beta \rightarrow 1} q^*(x, \beta) < \tilde{q}(x)$, then $\lim_{\beta \rightarrow 1} P_F^*(x, \beta) = 1$, resulting in $\lim_{\beta \rightarrow 1} q^*(x, \beta) = q^G(\lambda) \geq 1 - \frac{C}{|x|}$. Contradiction. Similarly $\lim_{\beta \rightarrow 1} q^*(x, \beta) > \tilde{q}(x)$ can not hold as well. Therefore $\lim_{\beta \rightarrow 1} q^*(x, \beta) = \tilde{q}(x)$. Moreover, $P_F^*(x, 1) = \frac{-\ln(1-\tilde{q}(x))}{\lambda\tilde{q}(x)}$.
- 3) $|x| > \frac{C}{1-q^G(\lambda)}$: Since q^* is bounded above by $q^G(\lambda)$, we hence obtain $\lim_{\beta \rightarrow 1} P_F^*(x, \beta) = 1$, which sustains that $\lim_{\beta \rightarrow 1} q^*(x, \beta) = q^G(\lambda)$.

Therefore, with the results from Theorem 3.2, $q^*(x, \beta)$ and $P_F^*(x, \beta)$ are both continuous at $\beta = 1$. As $\mu^{\text{th}*}(x, \beta)$ is a continuous function of $P_F^*(x, \beta)$ (inverse CDF F^{-1}), it is also continuous at $\beta = 1$. ■

Proof of Lemma 3.4:

It is self-evident for the case when $\beta = 1$ by Theorem 3.2. For $\beta < 1$, we observe that, given news sign, if for any $q \in [0, 1]$, the corresponding value of function $1 - e^{-q\lambda P_F(s((x, \beta), F))}$ is increased in $|x|$, the solution q^* (and corresponding fraction of sharing P_F^*) is increased as well. Therefore, it suffices to show that for any q , P_F is increasing in news magnitude $|x|$; this can be easily proved by noting the monotonicity of the sharing threshold $\mu^{\text{th}}((x, \beta), F) = \frac{\frac{C}{2(1-q)(F(0)-F(\frac{-\beta x}{1-\beta}))} - \beta x}{1-\beta}$ in news magnitude. ■

Proof of Lemma 3.5: According to Theorem 3.2, given $\beta = 1$, when $|x| \leq C$ the equilibrium spread size is zero whereas when $|x| > C$ it is $\min\{1 - \frac{C}{|x|}, q^G(\lambda)\}$; in either case the equilibrium spread size is increasing in λ . For $\beta < 1$, the argument is similar to the above proof: Since all else equal the function value of $1 - e^{-q\lambda P_F(s((x, \beta), F))}$ is increasing in λ , the result is then obvious. ■

B.2 Proofs of Section 3.4

Proof of Proposition 3.1:

Case (i) and (ii) are direct results of Theorem 3.2. We now prove the necessary and sufficient condition for $\beta = 1$ to be a local minimizer when $|x| \in (C, \frac{C}{1-q^G(\lambda)}]$. We compute the derivative of q^* w.r.t. β when $\beta \rightarrow 1$, i.e., $\lim_{\beta \rightarrow 1} \frac{\partial q^*}{\partial \beta}$. In the following we only compute the case for $x > 0$ and for $x < 0$ the approach is similar. To evaluate $\lim_{\beta \rightarrow 1} \frac{\partial q^*}{\partial \beta}$, we use the continuity of μ^{th^*} at $\beta = 1$.

$$\mu^{\text{th}^*}((x, 1), F) = \lim_{\beta \rightarrow 1} \mu^{\text{th}^*}((x, \beta), F) = \lim_{\beta \rightarrow 1} \frac{\frac{C}{2(1-q^*)F^{\text{flip}}(x, \beta)} - \beta x}{1 - \beta}.$$

Applying l'hospital rule to the last term and using the fact that $\lim_{|\mu| \rightarrow \infty} \mu^2 f(\mu) = 0$ (as the mean and variance are both finite), we derive

$$\begin{aligned} \lim_{\beta \rightarrow 1} \frac{\frac{C}{2(1-q^*)F^{\text{flip}}(x, \beta)} - \beta x}{1 - \beta} &= \lim_{\beta \rightarrow 1} \frac{-C \left(-\frac{dq^*}{d\beta} F^{\text{flip}}(x, \beta) + (1-q^*) f\left(\frac{-\beta x}{1-\beta}\right) \frac{x}{(1-\beta)^2} \right)}{2((1-q^*)F^{\text{flip}}(x, \beta))^2 - 1} - x, \\ &= \frac{-x^2 \lim_{\beta \rightarrow 1} \frac{dq^*}{d\beta}}{C} + x. \end{aligned}$$

As a result, when $x > 0$, $\lim_{\beta \rightarrow 1} \frac{dq^*}{d\beta} < 0$ iff $\mu^{\text{th}^*}((x, 1), F) > x$. Similarly, repeating the approach, one can show that when $x < 0$, $\lim_{\beta \rightarrow 1} \frac{dq^*}{d\beta} < 0$ iff $\mu^{\text{th}^*}((x, 1), F) < x$. The proof is completed. \blacksquare

Proof of Proposition 3.2:

WLOG, we prove for positive news; we aim to show the existence of $\bar{x}(\lambda, F)$.

According to Proposition 3.1, for $\beta = 1$, when $x \leq C$, no agents share the news whereas when $x > \frac{C}{1-q^G(\lambda)}$, every agent shares the news. It also means that $\mu^{\text{th}^*}(\{x = C, 1\}, F) = +\infty$ and $\mu^{\text{th}^*}(\{x = \frac{C}{1-q^G(\lambda)}, 1\}, F) = -\infty$. Therefore, in the range $(C, \frac{C}{1-q^G(\lambda)}]$, the equilibrium sharing threshold $\mu^{\text{th}^*}((x, 1), F)$ is decreasing in x , from $+\infty$ to $-\infty$ (cf. Lemma 3.4). We can then easily see that there exists unique $\bar{x}(\lambda, F) > C$ such that $\mu^{\text{th}^*}(\{x, 1\}, F) > x$ iff $x < \bar{x}(\lambda, F)$. \blacksquare

B.3 Proofs of Section 3.5

Proof of Proposition 3.3:

Observing that in (3.17), $\bar{\mu}$ affects λ^{th} through function $F(-|x|)$, we investigate their effects on $F(-|x|)$. When $\bar{\mu} = 0$, $\lambda^{\text{th}}(x, 0, \sigma_\mu, G) = \frac{-\ln \frac{c}{|x|}}{(1-\frac{c}{|x|})G(-\frac{|x|}{\sigma_\mu})}$. The derivative of $F(-|x|)$ w.r.t. $\bar{\mu}$ is

$$\frac{dF(-|x|)}{d\bar{\mu}} = \frac{1}{2\sigma_\mu} \left[g\left(\frac{-|x| + \bar{\mu}}{\sigma_\mu}\right) - g\left(\frac{-|x| - \bar{\mu}}{\sigma_\mu}\right) \right] = \frac{1}{2\sigma_\mu} \left[g\left(\frac{\bar{\mu} - |x|}{\sigma_\mu}\right) - g\left(\frac{\bar{\mu} + |x|}{\sigma_\mu}\right) \right] > 0,$$

since g is even and decreasing in $[0, +\infty)$ and $|\bar{\mu} + |x|| > |\bar{\mu} - |x||$. Therefore $F(-|x|)$ is strictly increasing in $\bar{\mu}$ and then $\lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G)$ is decreasing in $\bar{\mu}$. \blacksquare

Proof of Proposition 3.4:

We can similarly derive the effect of σ_μ on λ^{th} by examining its effect on $F(-|x|)$:

$$\begin{aligned} \frac{dF(-|x|)}{d\sigma_\mu} &= \frac{1}{2\sigma_\mu} \left[g\left(\frac{-|x| + \bar{\mu}}{\sigma_\mu}\right) \left(\frac{x - \bar{\mu}}{\sigma_\mu}\right) + g\left(\frac{-|x| - \bar{\mu}}{\sigma_\mu}\right) \left(\frac{x + \bar{\mu}}{\sigma_\mu}\right) \right], \\ &= \frac{1}{2\sigma_\mu} \left[g\left(\frac{\bar{\mu} - |x|}{\sigma_\mu}\right) \left(\frac{x - \bar{\mu}}{\sigma_\mu}\right) + g\left(\frac{\bar{\mu} + |x|}{\sigma_\mu}\right) \left(\frac{\bar{\mu} + x}{\sigma_\mu}\right) \right], \end{aligned}$$

which is positive iff

$$\frac{g\left(\frac{\bar{\mu} + |x|}{\sigma_\mu}\right)}{g\left(\frac{\bar{\mu} - |x|}{\sigma_\mu}\right)} > \frac{\bar{\mu} - |x|}{\bar{\mu} + |x|}. \quad (\text{B.2})$$

Note that for $\bar{\mu} \leq |x|$, (B.2) holds for any σ_μ , implying that $F(-|x|)$ is increasing in σ_μ . When $\bar{\mu} > |x|$, we first show that the LHS of (B.2) is increasing in σ_μ . Take the

derivative, we find that

$$\frac{d}{d\sigma_\mu} \frac{g\left(\frac{\bar{\mu}+|x|}{\sigma_\mu}\right)}{g\left(\frac{\bar{\mu}-|x|}{\sigma_\mu}\right)} = -\frac{g\left(\frac{\bar{\mu}+|x|}{\sigma_\mu}\right)}{g\left(\frac{\bar{\mu}-|x|}{\sigma_\mu}\right)} \underbrace{\left[\frac{g'\left(\frac{\bar{\mu}+|x|}{\sigma_\mu}\right)}{g\left(\frac{\bar{\mu}+|x|}{\sigma_\mu}\right)} \frac{(\bar{\mu}+|x|)}{\sigma_\mu^2} - \frac{g'\left(\frac{\bar{\mu}-|x|}{\sigma_\mu}\right)}{g\left(\frac{\bar{\mu}-|x|}{\sigma_\mu}\right)} \frac{(\bar{\mu}-|x|)}{\sigma_\mu^2} \right]}_{(*)}.$$

Since $\frac{g'(y)}{g(y)}$ is decreasing due to log-concavity of g and $\frac{g'(y)}{g(y)} < 0$ for $y > 0$, the part $(*)$ is negative and consequently LHS of (B.2) is increasing in σ_μ . Observing that LHS of (B.2) approaches to 0 when $\sigma_\mu \rightarrow 0$ while to 1 when $\sigma_\mu \rightarrow +\infty$, the monotonicity then suggests that there exists some value $\hat{\sigma}(x, \mu, G) > 0$ such that $F(-|x|)$ is decreasing in σ_μ iff $\sigma_\mu \leq \hat{\sigma}(x, \mu, G)$.

For the two cases discussed above:

- When $\bar{\mu} \leq |x|$, we have $\lim_{\sigma_\mu \rightarrow 0} F(-|x|) = 0$. Hence, $\lim_{\sigma_\mu \rightarrow 0} \lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G) = +\infty$ and $\lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G)$ is decreasing in σ_μ .
- When $\bar{\mu} > |x|$, we have $\lim_{\sigma_\mu \rightarrow 0} F(-|x|) = \frac{1}{2}$ and thus $\lim_{\sigma_\mu \rightarrow 0} \lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G) = \frac{-2 \ln \frac{C}{|x|}}{1 - \frac{C}{|x|}}$. As $F(-|x|)$ is decreasing in σ_μ iff $\sigma_\mu \leq \hat{\sigma}(x, \mu, G)$, then $\lambda^{\text{th}}(x, \bar{\mu}, \sigma_\mu, G)$ is increasing in σ_μ up to $\hat{\sigma}(x, \bar{\mu}, G)$ and then decreasing. ■

Appendix C

Supplementary Material for Chapter 4

C.1 Proofs of Section 4.3

Proof of Lemma 4.1.

Symmetry: by the symmetry in the information structure.

We claim one more important property here.

Non-expansion of beliefs: for any x , $|\mathbb{E}_i[\theta | x, T] - \mu_i| \leq |x - \mu_i|$. For a perspective μ_i , the difference $\mathbb{E}_i[\theta | x, T] - \mu_i$ equals

$$\frac{\int (\theta - \mu) f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta}{\int f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta} = \frac{\int t f(t) f_\epsilon((x - \mu_i) - t) dt}{\int f(t) f_\epsilon((x - \mu_i) - t) dt}. \quad (\text{C.1})$$

Chambers and Healy (2012, Proposition 3) showed that under Assumption 5 the above expression (C.1) is a convex combination of 0 and $x - \mu_i$ and hence $|\mathbb{E}_i[\theta | x, T] - \mu_i| \leq |x - \mu_i|$. We will use $\mathbb{E}_0[\theta | x, T]$ to refer to the posterior expectation corresponding to zero perspective and informative news x ; note that the expression (C.1) equals $\mathbb{E}_0[\theta | x - \mu_i, T]$.

Monotonicity in informative news: Given that x is the sum of two independent variables θ and ϵ and that f_i and f_ϵ are both log-concave, according to Efron's theorem (Efron 1965), the belief function $\mathbb{E}_i[\theta | x, T]$ is non-decreasing in x . Moreover, following

the proof provided by Saumard and Wellner (2014), we can find that the equality holds only if θ or ϵ is almost surely a constant. This suggests that $\mathbb{E}_i[\theta | x, T]$ is *strictly* increasing in x .

Unbounded information:

Now we show that for any μ_i , $\lim_{x \rightarrow +\infty} \mathbb{E}_i[\theta | x, T] > 0$ and $\lim_{x \rightarrow -\infty} \mathbb{E}_i[\theta | x, T] < 0$. It is obviously true for $\mu_i = 0$. For $\mu_i \neq 0$, according to (4.3), we only need to focus on the numerator:

$$\begin{aligned} \int \theta f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta &= \int_0^{+\infty} \theta f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta + \int_{-\infty}^0 \theta f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta, \\ &= \int_0^{+\infty} \theta f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta - \int_0^{+\infty} \theta f(-\theta - \mu_i) f_\epsilon(x + \theta) d\theta, \\ &= \int_0^{+\infty} \underbrace{\left(1 - \frac{f(-\theta - \mu_i) f_\epsilon(x + \theta)}{f(\theta - \mu_i) f_\epsilon(x - \theta)}\right)}_{(*)} \theta f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta. \end{aligned}$$

Alternatively, we can write the convolution sum as

$$\int \theta f(\theta - \mu_i) f_\epsilon(x - \theta) d\theta = - \int_0^{+\infty} \underbrace{\left(1 - \frac{f(\theta - \mu_i) f_\epsilon(x - \theta)}{f(-\theta - \mu_i) f_\epsilon(x + \theta)}\right)}_{(*)} \theta f(-\theta - \mu_i) f_\epsilon(x + \theta) d\theta.$$

The ratio $(*)$ is the relative likelihood between $-\theta$ and θ given μ_i and x_i .¹

¹In fact $(*)$ can be seen as the ratio of two ratios that possess monotone likelihood property in θ :

$$\frac{\frac{f(-\theta - \mu_i)}{f(\theta - \mu_i)}}{\frac{f_\epsilon(x - \theta)}{f_\epsilon(x + \theta)}}$$

The magnitude of numerator measures the likelihood ratio of $-\theta$ to θ based on perspective μ_i while the denominator means based on the news generation process the likelihood ratio between θ to $-\theta$ based

By Assumption 5 that for any μ_i there exists x such that the ratio (*) is decreasing in θ . Since (*) equals 1 for $\theta = 0$, we prove that for any μ_i , there exist $x^-(\mu_i)$, $x^+(\mu_i)$ such that $\mathbb{E}_i[\theta | x^-, T] < 0$ and $\mathbb{E}_i[\theta | x^+, T] > 0$. Using the monotonicity of $\mathbb{E}_i[\theta | x, T]$ in true news x , we show the result.

Monotonicity in perspective:

For $\mathbb{E}_i[\theta | x, T]$, the result follows Efron's theorem by noticing that given x we can regard $f_\epsilon(x - \theta)$ as a density function of θ and $f(\theta - \mu_i)$ as one of some ϵ' independent of θ such that $\theta + \epsilon' = \mu_i$.

To show the monotonicity of $\mathbb{E}_i[\theta | x]$ in μ_i , we use the property that it is a convex combination of $\mathbb{E}_i[\theta | x, T]$ and μ_i . Consider any $\mu_1 > \mu_2$. If $x \in [\mu_2, \mu_1]$, then $\mathbb{E}_1[\theta | x, T] \geq x \geq \mathbb{E}_2[\theta | x, T]$ by the convexity property; since the likelihood weights in the convex combination are nonzero, it is obvious that $\mathbb{E}_1[\theta | x] > \mathbb{E}_2[\theta | x]$ for this case. If $x < \mu_2 < \mu_1$, then $f_T(x; \mu_2) > f_T(x; \mu_1)$ and $\ell_2(x, T) > \ell_1(x, T)$. Because $\mathbb{E}_2[\theta | x, T] < \mu_2$ and $\mathbb{E}_2[\theta | x, T] < \mathbb{E}_1[\theta | x, T]$, we have

$$\begin{aligned} \mathbb{E}_2[\theta | x] &= \ell_2(x, T)\mathbb{E}_2[\theta | x, T] + (1 - \ell_2(x, T))\mu_2, \\ &< \ell_1(x, T)\mathbb{E}_2[\theta | x, T] + (1 - \ell_1(x, T))\mu_2, \\ &< \ell_1(x, T)\mathbb{E}_1[\theta | x, T] + (1 - \ell_1(x, T))\mu_1 = \mathbb{E}_1[\theta | x]. \end{aligned}$$

We can similarly analyze the case when $x > \mu_1 > \mu_2$.

Unboundedness of prior beliefs:

Now we show that for any x , $\lim_{\mu_i \rightarrow +\infty} \mathbb{E}_i[\theta | x, T] > 0$ and $\lim_{\mu_i \rightarrow -\infty} \mathbb{E}_i[\theta | x, T] < 0$. It is obviously true for $x = 0$. Similar to our approach in the proof for unboundedness of informative news, for any $x \neq 0$ we can find that there exist some $\mu_i^+(x)$ and $\mu_i^-(x)$ such that $\mathbb{E}_i[\theta | x, T] > 0$ for $\mu_i = \mu_i^+$ and $\mathbb{E}_i[\theta | x, T] < 0$ for $\mu_i = \mu_i^-$; by the monotonicity, we show the result.

on the noise distribution. If both ratios are decreasing in θ , it means that perspective μ_i suggests that the set of positive value for θ is more likely whereas x suggests the opposite side. In this case, if the ratio for perspective μ_i decreases faster in θ than its counterpart, the expected value will have the sign of the ideology that x indicates.

Moreover, for the case when news veracity is uncertain, given any x ,

$$\lim_{\mu_i \rightarrow +\infty} \mathbb{E}_i[\theta | x] = \lim_{\mu_i \rightarrow +\infty} \mathbb{P}_i(\omega = T|x) \mathbb{E}_i[\theta | x, T] + (1 - \mathbb{P}_i(\omega = T|x))\mu_i > 0.$$

Similarly we can derive for the case when $\mu_i \rightarrow -\infty$. The proof is complete. \blacksquare

Proof of Lemma 4.2.

We define some notations for the proof. Let TC_i denote the set of "true" news that will be concealed according to the strategies v_i, m_i and similarly define FC_i for the set of false news that is concealed; we rewrite the updating function when there is no disclosure (4.5) as

$$\pi_j(i, \emptyset) = \frac{p\mu_j + (1-p) \left[q \int_{x \in FC_i} \mu_j f_F(x) dx + (1-q) \int_{x \in TC_i} \mathbb{E}_j[\theta | x, T] f_T(x; \mu_j) dx \right]}{p + (1-p) \left[q \int_{x \in FC_i} f_F(x) dx + (1-q) \int_{x \in TC_i} f_T(x; \mu_j) dx \right]}, \quad (\text{C.2})$$

$$= \mu_j + \frac{(1-p)(1-q) \int_{x \in TC_i} \mathbb{E}_0[\theta | x - \mu_j, T] f_T(x; \mu_j) dx}{\underbrace{p + (1-p) \left[q \int_{x \in FC_i} f_F(x) dx + (1-q) \int_{x \in TC_i} f_T(x; \mu_j) dx \right]}_{P_j}}$$

where P_i represents the probability of the no disclosure by intermediary i based on subscriber j 's prior belief and strategy v_i, m_i .

Take the derivative of (C.2) w.r.t. μ_i , we obtain²

$$1 + \frac{(1-p)(1-q)}{P_j^2} \left(P_j \int_{x \in TC_i} -\frac{d}{dx} \mathbb{E}_0[\theta|x - \mu_j, T] f_T(x; \mu_j) dx + P_j \int_{x \in TC_i} \mathbb{E}_0[\theta|x - \mu_j, T] \frac{d}{d\mu_j} f_T(x; \mu_j) dx \right. \\ \left. - (1-p)(1-q) \int_{x \in TC_i} \frac{d}{d\mu_j} f_T(x; \mu_j) dx \int_{x \in TC_i} \mathbb{E}_0[\theta|x - \mu_j, T] f_T(x; \mu_j) dx \right)$$

If $TC_i = \emptyset$, it is trivial to see that the difference is positive. Suppose $TC_i \neq \emptyset$. We decompose the expression in within the large parentheses into three terms:

$$\underbrace{P_j \int_{x \in TC_i} -\frac{d}{dx} \mathbb{E}_0[\theta|x - \mu_j, T] f_T(x; \mu_j) dx}_{(a)} \\ + (p + (1-p)q \int_{x \in FC_i} f_F(x) dx) \underbrace{\int_{x \in TC_i} \mathbb{E}_0[\theta|x - \mu_j, T] \frac{d}{d\mu_j} f_T(x; \mu_j) dx}_{(b)} \\ + (1-p)(1-q) \left(\int_{x \in TC_i} f_T(x; \mu_j) dx \int_{x \in TC_i} \mathbb{E}_0[\theta|x - \mu_j, T] \frac{d}{d\mu_j} f_T(x; \mu_j) dx \right. \\ \left. - \int_{x \in TC_i} \frac{d}{d\mu_j} f_T(x; \mu_j) dx \int_{x \in TC_i} \mathbb{E}_0[\theta|x - \mu_j, T] f_T(x; \mu_j) dx \right) \underbrace{(c)}.$$

We will show that the derivative is positive by combining (a)-(c) with the constant 1.

(a): Since $0 < \frac{d}{d\mu_j} \mathbb{E}_j[\theta|x, T] = \frac{d}{d\mu_j} (\mu_j + \mathbb{E}_0[\theta|x - \mu_j, T]) = 1 - \frac{d}{dx} \mathbb{E}_0[\theta|x - \mu_j, T]$, combining (a) with the term 1, we have

$$1 + \frac{(1-p)(1-q)}{P_j^2} \left(P_j \int_{x \in TC_i} -\frac{d}{dx} \mathbb{E}_0[\theta|x - \mu_j, T] f_T(x; \mu_j) dx \right), \\ > 1 - \frac{1}{P_j} (1-p)(1-q) \int_{x \in TC_i} f_T(x; \mu_j) dx > 0.$$

²The interchange of integration and derivative is justified as the sets of concealed news are measurable and the density functions are log-concave.

(b): It is easy to see that $\mathbb{E}_0[\theta|x - \mu_j, T] > 0$ iff $x > \mu_j$. Note that given μ_j , the function $f_T(x; \mu_j)$ is symmetric around $x = \mu_j$ and log-concave in x . We can also show that $f_T(x; \mu_j) = f_T(x - \mu_j; 0)$ and then derive that $f_T(x; \mu_j)$ is log-concave in μ_j . Therefore

$$\frac{d}{d\mu_j} f_T(x; \mu_j) < 0, \text{ iff } x < \mu_j.$$

We then show that (b) is non-negative.

(c): Divide it by $(\int_{x \in TC_i} f_T(x; \mu_j) dx)^2$ and we obtain

$$\begin{aligned} & \int_{x \in TC_i} \mathbb{E}_0[\theta|x - \mu_j, T] \frac{\frac{d}{d\mu_j} f_T(x; \mu_j)}{f_T(x; \mu_j)} \frac{f_T(x; \mu_j)}{\int_{x \in TC_i} f_T(x; \mu_j)} dx \\ & - \int_{x \in TC_i} \frac{\frac{d}{d\mu_j} f_T(x; \mu_j)}{f_T(x; \mu_j)} \frac{f_T(x; \mu_j)}{\int_{x \in TC_i} f_T(x; \mu_j)} dx \int_{x \in TC_i} \mathbb{E}_0[\theta|x - \mu_j, T] \frac{f_T(x; \mu_j)}{\int_{x \in TC_i} f_T(x; \mu_j)} dx. \end{aligned}$$

Since $\frac{d}{d\mu_j} f_T(x; \mu_j) = -\frac{d}{dx} f_T(x - \mu_j; 0)$, we can find that $\frac{\frac{d}{d\mu_j} f_T(x; \mu_j)}{f_T(x; \mu_j)}$ is increasing in x by log-concavity. Note that $\mathbb{E}_0[\theta|x - \mu_j, T]$ is also increasing in x and the total sum is non-negative by Chebyshev sum inequality.

Now to prove that $\lim_{\mu_j \rightarrow +\infty} \pi_j(i, \emptyset) > 0$. Let's consider $\mu_j > 0$ and then increase μ_j to $+\infty$. We only need to consider the numerator of (C.2):

$$\begin{aligned} & \underbrace{\left(p + (1-p)q \int_{x \in FC_i} f_F(x) dx \right) \mu_j + (1-p)(1-q) \int_{\substack{x \in TC_i \\ \mathbb{E}_j[\theta|x, T] < 0}} \mathbb{E}_j[\theta|x, T] f_T(x; \mu_j) dx}_{(*)} \\ & + (1-p)(1-q) \underbrace{\int_{\substack{x \in TC_i \\ \mathbb{E}_j[\theta|x, T] \geq 0}} \mathbb{E}_j[\theta|x, T] f_T(x; \mu_j) dx}_{(**)}. \end{aligned}$$

The term (**) is non-negative. Equivalently we want to show that the term (*) is positive for some perspective that is sufficiently large. Let's simply focus on the range

Since $\mu_j > 0$, then only $x < 0$ can make $\mathbb{E}_j[\theta | x, T] < 0$. We can see that given any x such that $\mathbb{E}_j[\theta | x, T] < 0$, $f_T(x; \mu_j)$ is decreasing in the range $\mu_j > 0$. Together with the fact that $\mathbb{E}_j[\theta | x, T]$ is increasing in μ_j (and hence the set of news x such that $\mathbb{E}_j[\theta | x, T] < 0$ is decreasing in μ_j), we can derive that the red term is increasing in μ_j . Because the blue term is strictly increasing in μ_j and the derivative is bounded away from zero ($p + (1 - p)q \int_{x \in FC_i} f_F(x) dx > p$) for any μ_j , suggesting that for sufficiently large μ_j , the numerator is positive. ■

Proof of Lemma 4.4.

The result holds for any common history G^s and we dismiss the argument G^s for simplicity. WLOG, we consider intermediary R . Also, since $\mathbb{E}_i[\theta | x, T]$ is strictly in news value (cf. Lemma 4.1), we know that for any player i , $\mathbb{E}_i[\theta | x, T] > 0$ iff $x > x_T^{\text{th}}(\mu_i)$ where $x_T^{\text{th}}(\mu_i)$ is the unique solution to $\mathbb{E}_i[\theta | x, T] = 0$. We will simply refer to the set of x such that $\mu_R \mathbb{E}_R[\theta | x, T] > 0$ as $x > x_T^{\text{th}}(\mu_i)$.

(i): For this part, we illustrate the case when the news is not verified and the cases for true and false news follow the same logic.

(if part) If $\mathbb{E}_R[\theta | x] \pi_{k_R^*}(x) > 0$, it is obvious that she can persuade her critical subscribers to strictly prefer her favored action if disclosing news x instead of finding indifference.

(only if part) Suppose intermediary R discloses news x but $\mathbb{E}_R[\theta | x] \pi_{k_R^*}(x) \leq 0$. When it is $\mathbb{E}_R[\theta | x] = 0$ or $\pi_{k_R^*}(x) = 0$, intermediary R feels indifferent because for the former her utility is zero while for the latter she cannot change the fractions of her subscribers whether she discloses or not. If she and her critical subscribers strictly prefer the opposite actions when observing news x , intermediary R conceals it since she can make her critical subscribers feel indifferent, relatively leaning to her side. Contradiction.

(ii): For any $x > x_T^{\text{th}}(\mu_R)$, intermediary R favors +1 regardless of her verification decision and news veracity. Then based on part (i), it is easy to see that

$$m_R^*(G^s, x, 1, T) = m_R^*(G^s, x, 1, F) = m_R^*(G^s, x, 0). \quad \blacksquare$$

(iii): We check the numerator of posterior mean when there is no disclosure (C.2) and show that the critical perspective for intermediary R must be positive. As intermediaries' disclosure decisions can depend on news veracity, we write TC_R^* (FC_R^*) for the set of true (false) news that is concealed by intermediary R at the equilibrium. According to (4.5) and TC_R^* , FC_R^* , we write the equation for critical perspective as:

$$0 = p\mu_{k_R^*} + (1-p) \underbrace{\int_{\substack{x \in TC_R^* \cap FC_R^* \\ x > x_T^{\text{th}}(\mu_R)}} (qf_F(x) + (1-q)f_T(x; \mu_{k_R^*})) \mathbb{E}_{k_R^*}[\theta|x] dx}_{(a)} \quad (\text{C.3})$$

$$+ (1-p) \underbrace{\left[\int_{\substack{x \in FC_R^* \\ x \leq x_T^{\text{th}}(\mu_R)}} qf_F(x) \mu_{k_R^*} dx + \int_{\substack{x \in TC_R^* \\ x \leq x_T^{\text{th}}(\mu_R)}} (1-q)f_T(x; \mu_{k_R^*}) \mathbb{E}_{k_R^*}[\theta|x, T] dx \right]}_{(b)}.$$

The term (a) concerns the news value $x > x_T^{\text{th}}(\mu_R)$ for which intermediary R chooses the same disclosure action (cf. (ii)); consequently, if she conceals some news value $x > x_T^{\text{th}}(\mu_R)$, both true and false news are concealed and her subscribers integrate the likelihood of such news value and corresponding expectation. On the other hand, the term (b) represents the other news that intermediary R may choose different disclosure decisions depending on her verification result.

Suppose there exists an equilibrium with $\mu_{k_R^*} < 0$. By (i) and monotonicity in perspective, in term (a) it must be that $\mathbb{E}_{k_R^*}[\theta|x] \leq 0$ since intermediary R discloses any news x such that $\mathbb{E}_{k_R^*}[\theta|x] > 0$. Also, it is obvious to see that (b) is non-positive: For any $x \leq x_T^{\text{th}}(\mu_R) < 0$, the sign of the expectation based on perspective $\mu_{k_R^*}$ must be negative. Therefore the sum is negative (note that $p\mu_{k_R^*} < 0$); contradiction. It is not possible that $\mu_{k_R^*} = 0$ because when $\mu_{k_R^*} = 0$, then $\mathbb{E}_{k_R^*}[\theta|x] < 0$ for any $x \in (x_T^{\text{th}}(\mu_R), 0)$ and the sum is negative. Consequently $\mu_{k_R^*} > 0$.

(iv): For any $x \geq 0$, if intermediary R discloses the news, based on (ii), each of her subscriber j updates his expectation as $\mathbb{E}_j[\theta|x]$ as they cannot tell if it is true or false. Because for such news intermediary R strictly prefer +1 and her critical subscribers

have positive expectation if seeing such news, intermediary R discloses it (otherwise she loses some fraction of her subscribers taking $+1$ as critical subscribers feel indifferent when there is no disclosure). ■

Proof of Lemma 4.5.

As argued in the context, intermediary R does not verify any news $x > x_T^{\text{th}}(\mu_R)$. For news $x = x_T^{\text{th}}(\mu_R)$, which makes $\mathbb{E}_R[\theta|x, T] = 0$, we have $\mathbb{E}_R[\theta|x_T^{\text{th}}(\mu_R)] = (1 - \ell_R(x_T^{\text{th}}(\mu_R), T))\mu_R$, suggesting that she feels indifferent between disclosing or not if the news is true. It is then easy to see that she chooses $m_R^*(x, 1, F) = m_R^*(x, 0)$ and derives the same expected payoff for either verifying it or not; she does not verify it.

If $v_R^*(x) = 1$, she has to deploy a separating disclosure strategy; otherwise, she can obtain the same expected payoff by choosing not to verify it and making the same disclosure decision. ■

C.2 Proofs of Section 4.4

Proof of Proposition 4.2.

Corollary 4.1 is a special case of the proof.

Existence of equilibrium and unique critical perspectives and disclosure strategies:

In (4.11) the disclosure strategy m_R^* is formulated as a function of the critical perspective $\mu_{k_R^*}$. We can then write the fixed-point equation for $\mu_{k_R^*}$ according to (4.5):

$$p\mu_{k_R^*} + (1 - p) \int_{x|m_R^*(x)=\emptyset} [qf_F(x) + (1 - q)f_T(x; \mu_{k_R^*})] \mathbb{E}_{k_R^*}[\theta|x] dx = 0, \quad (\text{C.4})$$

where the set of concealed news (i.e., $\{x|m_R^*(x) = \emptyset\}$) is determined using (4.11). Note that the solutions m_R^* and $\mu_{k_R^*}$ to (C.4) (if any) are not dependent on G^s , meaning that intermediary R plays m_R^* with corresponding $\mu_{k_R^*}$ for any common history G^s . In turn, any solution to (C.4) suggests the existence of an equilibrium and the equilib-

rium subscription choices are made by simply comparing the intermediaries' disclosure strategies m_L^* , m_R^* resulting from (C.4).

There exists a solution in range $(0, \mu_R^{\text{ext}})$: The function on the left-hand side of (C.4) is continuous in μ_{k_R} and, based on our previous proofs, the function value is negative when $\mu_{k_R} = 0$ whereas positive when $\mu_{k_R} = \mu_R^{\text{ext}}$. Consequently, a solution exists.

We now show that the function is increasing in $\mu_{k_R} \in [0, \mu_R^{\text{ext}}]$, implying the uniqueness of the solution. First, since $\mu_{k_R} < \mu_R$, $\mathbb{E}_{\mu_{k_R}}[\theta|x] \leq 0$ for any concealed news $\{x \in \mathbb{R} | m_R^*(x) = \emptyset\}$. Second, $\mathbb{E}_{\mu_{k_R}}[\theta|x]$ is increasing in μ_{k_R} for any x ; it also means that the set of concealed news is decreasing in μ . Finally, f_F is not a function of μ_{k_R} while $f_T(x; \mu)$ is *decreasing* in $\mu_{k_R} \geq 0$ given any $x < 0$ (recall that intermediary R only conceals left-leaning news by Lemma 4.4). By calculus the function in (C.4) is increasing in μ_{k_R} .

Subscription choices:

WLOG, consider subscriber j with $\mu_j \geq 0$; his ex-ante expected utility from subscribing to intermediary i , who plays disclosure strategy m_i^* , is

$$\mathbb{E}_j \left[a_j^* \theta \mid s_j = i \right] = \mathbb{E}_j \left[\mathbb{E}_j [a_j^* \theta \mid s_j, m_{s_j}^*] \mid s_j = i \right] = \mathbb{E}_j \left[\left| \mathbb{E}_j [\theta \mid s_j, m_{s_j}^*] \right| \mid s_j = i \right].$$

where the expectation is taken w.r.t. subscriber j 's prior belief about the distributions of true and false news. Expanding the term further, we obtain

$$\underbrace{\left| p\mu_j + (1-p) \mathbb{E}_j \left[\mathbb{E}_i [\theta \mid x] \mathbb{1}_{\{m_i^*(x)=\emptyset\}} \right] \right|}_{\text{match the sign of his posterior mean there is no disclosure}} + \underbrace{(1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j [\theta \mid x] \right| \mathbb{1}_{\{m_i^*(x)=x\}} \right]}_{\text{match the sign of his posterior mean when hearing some news}},$$

$$\begin{aligned} &\leq p|\mu_j| + (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j [\theta \mid x] \right| \mathbb{1}_{\{m_i^*(x)=\emptyset\}} \right] + (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j [\theta \mid x] \right| \mathbb{1}_{\{m_i^*(x)=x\}} \right], \\ &= p|\mu_j| + (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j [\theta \mid x] \right| \right]. \end{aligned}$$

If he follows intermediary L , i.e., $s_j = L$, he can obtain the maximum; it is because only right-leaning news may be concealed, for which $\mu_j \mathbb{E}_j[\theta \mid x] \geq 0$, and the equality

holds. On the other hand, if $s_j = R$,

- (i) For $\mu_j \in (0, \mu_{k_R^*})$: Since $\mathbb{E}_{\mu_{k_R^*}}[\theta | x] \leq 0$ for any concealed news x by intermediary R and $\mathbb{E}_j[\theta | x]$ is strictly increasing in μ_j , we obtain that $\mathbb{E}_j[\theta | x] < 0$ for any x concealed by intermediary R . It is the opposite sign against μ_j and hence his ex-ante utility is strictly smaller than that if following intermediary L ;
- (ii) For $\mu_j \in [\mu_{k_R^*}, \mu_R^{\text{ext}})$: There must exist *some nonzero-measured subset* of the concealed news for which $\mathbb{E}_j[\theta | x] < 0$, and as a result, he derives strictly higher payoff from subscribing to intermediary L . To see this claim, suppose not and $\mathbb{E}_j[\theta | x] \geq 0$ for any concealed news. Given that $\mu_j < \nu_R$, there exists some nonzero measure set of news for which $\mathbb{E}_j[\theta | x] < 0$.³ For convenience, denote the set as $A(\mu_j) \triangleq \{x \in \mathbb{R} | \mathbb{E}_j[\theta | x] < 0\}$. Since $\mu_j < \mu_R$, by continuity and the monotonicity in perspective, we have $\mathbb{E}_R[\theta | x] \geq 0$ for some subset of $A(\mu_j)$ and intermediary R would conceal this subset of news because $\mathbb{E}_{k_R^*}[\theta | x] < \mathbb{E}_j[\theta | x] < 0$ for any $x \in A(\mu_j)$. Contradiction.
- (iii) For $\mu_j = 0$ or $\mu_j \geq \mu_R^{\text{ext}}$, it is easy to see that $\mu_j \mathbb{E}_j[\theta | x] \geq 0$ holds for any concealed news by intermediary R and the equality is satisfied, suggesting indifference between the two intermediaries. ■

C.3 Proofs of Section 4.5

Proof of Lemma 4.6.

WLOG, we discuss for intermediary R . When $C = 0$, we show that for any $x < x_T^{\text{th}}(\mu_R)$, it is always strictly beneficial for intermediary R to verify.

For any $x < x_T^{\text{th}}(\mu_R)$ the decision $v_R^*(x) = 0$ is not sustainable. Suppose in an equilibrium intermediary R does not verify x for some news $x < x_T^{\text{th}}(\mu_R)$. If intermediary R subsequently chooses $m_R^*(x, \emptyset) = x$, i.e., to reveal the unverified news, then she can

³The nonzero measure is implied by the regularity conditions in Assumption 5.

deviate to verify and then choose to conceal either true or false news. First note that after seeing x her subscribers' posterior mean are updated as $\mathbb{E}_j[\theta | x]$ as they believe it is unverified; the corresponding perspective threshold $\mu_R^{\text{th}*}(x)$ is positive (because $x < 0$). If $\mu_R^{\text{th}*}(x) > \widehat{\mu}_R^*(\emptyset)$, then intermediary R is incentivized to verify the news and will conceal the false news as she can make more subscribers vote $+1$. Similarly, she deviates when $\widehat{\mu}_R^*(x) < \widehat{\mu}_R^*(\emptyset)$. If $\widehat{\mu}_R^*(x) = \widehat{\mu}_R^*(\emptyset)$, then intermediary R would just conceal the news in the beginning; contradiction.

If intermediary R subsequently chooses to conceal ($m_R^*(x, 0) = \emptyset$), then she can deviate to verify it and choose $m_R^*(x, 1, F) = x$ if the news is false to derive a higher expected payoff: When her subscribers see the off-equilibrium disclosure of news x , they know that intermediary R verified the news (it could not be disclosing without verification; otherwise, intermediary R would not play $m_R^*(x, 0) = \emptyset$ in the beginning.). They acknowledge her incentive when the news is false ($\omega = F$) for which intermediary R prefers $+1$: The fraction of her subscribers taking $+1$ is $G_R^{s*}(+\infty) - G_R^{s*}(0)$, more than $G_R^{s*}(+\infty) - G_R^{s*}(\widehat{\mu}_R^*(\emptyset))$ ($\widehat{\mu}_R^*(\emptyset) > 0$ cf. Lemma 4.4) when the news is concealed. This justifies intermediary R 's deviation.

Now we prove that intermediary R will not deviate to choose no verification when she verifies some news $x < x_T^{\text{th}}(\mu_R)$. We simply consider disclosure strategy $m_R^*(x, 1, T) = \emptyset, m_R^*(x, 1, F) = x$, i.e., disclosing the news iff it is false. As her subscribers believe this disclosure strategy, they preserve their perspective as their belief means when seeing news x . If intermediary R decides not to verify the news and then simply discloses it, her subscribers would think the news is false and take their actions accordingly. In this case, if the news is in fact true, for which intermediary R prefers -1 , then she loses some expected payoff since she could have made more of subscribers to take -1 by concealing the news ($\mu_j \in [0, \widehat{\mu}_R^*(\emptyset))$). Similarly, if she simply conceals the news without verification, she could have gained more expected utility if the news is false, by making those subscribers with $\mu_j \in [0, \widehat{\mu}_R^*(\emptyset))$ take $+1$. Consequently, when $C = 0$, in any equilibrium intermediary R verifies news x if $x < x_T^{\text{th}}(\mu_R)$. ■

Proof of Proposition 4.3.

We study the class of *true-news-disclosing* equilibria. We brief the results from Lemma 4.4 to 4.6: (1) $v_R^*(x) = 1$ iff $x < x_T^{\text{th}}(\mu_R)$; (2) for any $x \geq x_T^{\text{th}}(\mu_R)$, $m_R^*(x, 0) = \emptyset$ iff $\mathbb{E}_R[\theta|x]\mathbb{E}_{k_R^*}[\theta|x] \leq 0$; (3) for any $x < x_T^{\text{th}}(\mu_R)$, then $m_R^*(x, 1, T) = x$ and $m_R^*(x, 1, F) = \emptyset$.

The critical subscribers are centrists:

We prove that if there exists any true-news-disclosing equilibrium, then $\mu_{k_R^*} < \mu_R^{\text{ext}}$. Given that intermediary R conceals news if it is verified and false, we rewrite (C.3) as

$$0 = p\mu_{k_R^*} + (1-p) \underbrace{\int_{x \in \mathbb{R} | v_R^*(x)=0, m_R^*(x,0)=\emptyset} \left(qf_F(x) + (1-q)f_T(x; \mu_{k_R^*}) \right) \mathbb{E}_{k_R^*}[\theta|x] dx}_{(a)} \tag{C.5}$$

$$+ (1-p) \underbrace{\int_{x \in \mathbb{R} | v_R^*(x)=1} qf_F(x) \mu_{k_R^*} dx}_{(b)}.$$

Suppose $\mu_{k_R^*} \geq \mu_R^{\text{ext}}$. In (a), $\mathbb{E}_{k_R^*}[\theta|x] \geq 0$ for any unverified news (if for some x such that $\mathbb{E}_{k_R^*}[\theta|x] < 0$, then it must be the case that $\nu_R > \mu_{k_R^*} \geq \mu_R$. However, it leads to $\mathbb{E}_R[\theta|x] < 0$ and x should be disclosed.). We can see that the sum is positive; contradiction.

No deviations for true-news-disclosing strategy:

For any news that is verified, it must be that $x < x_T^{\text{th}}(\mu_R)$ and we have the relation that $\widehat{\mu}_R^*(x) > \mu_R > \widehat{\mu}_R^*(\emptyset) (= \mu_{k_R^*})$ because her subscribers believe that if news x is disclosed, then it is true and flips intermediary R 's favored action. It is not hard to find that intermediary R has no incentives to deviate from the true-news-disclosing strategy.⁴

Existence of equilibrium and unique critical perspectives and disclosure

⁴For any news $x < x_T^{\text{th}}(\mu_R)$, the off-equilibrium disclosure choice (if intermediary R does not verify it) is to conceal the news iff $\mathbb{E}_R[\theta|x] \geq 0$, i.e., her favored action (+1) is different from the action her critical subscribers will take when seeing the news (-1).

WLOG, consider subscriber j with $\mu_j \geq 0$. If he follows intermediary L , i.e., $s_j = L$, subscriber j votes +1 both when he sees verified and informative news as well as when there is no disclosure from intermediary L (because $\mu_j \geq 0 > \mu_{k_L^*}$):

$$\begin{aligned}
& \mathbb{E}_j [a_j^* \theta \mid s_j = L], \tag{C.7} \\
& = (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j [\theta \mid x] \right| \mathbb{1}_{\{v_L^*(x)=0, m_L^*(x,0)=x\}} \right] + (1-p) \underbrace{\mathbb{E}_j \left[\mathbb{E}_j [\theta \mid x, T] \mathbb{1}_{\{\omega=T, v_L^*(x)=1, m_L^*(x,T)=x\}} \right]}_{\text{he votes +1}} \\
& \quad + \underbrace{\left(p\mu_j + (1-p) \mathbb{E}_j \left[\mu_j \mathbb{1}_{\{\omega=F, v_L^*(x)=1, m_L^*(x,F)=\emptyset\}} \right] + (1-p) \mathbb{E}_j \left[\mathbb{E}_j [\theta \mid x] \mathbb{1}_{\{v_L^*(x)=0, m_L^*(x,0)=\emptyset\}} \right] \right)}_{\text{when not receiving any news}}, \\
& = (1-p) \left(\mathbb{E}_j \left[\left| \mathbb{E}_j [\theta \mid x] \right| \mathbb{1}_{\{v_L^*(x)=0, m_L^*(x,0)=x\}} \right] + \mathbb{E}_j \left[\mathbb{E}_j [\theta \mid x] \mathbb{1}_{\{v_L^*(x)=1 \text{ or } v_L^*(x)=0, m_L^*(x,0)=\emptyset\}} \right] \right) + p\mu_j, \\
& = (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j [\theta \mid x] \right| \right] + p\mu_j,
\end{aligned}$$

where the last equality comes from the two facts: (1) Any news x such that $v_L^*(x) = 1$ or $(v_L^*(x) = 0, m_L^*(x, 0) = \emptyset)$ must be positive, making $\mathbb{E}_j[\theta \mid x]$ positive as well; (2) the two subsets of news are a partition of real line.

On the other hand, if $s_j = R$:

- (i) For $\mu_j \in [0, \mu_{k_R^*}]$, subscriber j votes -1 when observing verified and true news as well as votes -1 when there is no disclosure from intermediary R .⁵ We rewrite

⁵The critical subscriber with $\mu_{k_R^*}$ feels indifferent between the binary actions when there is no disclosure. For convenience of analysis, we make him vote -1 here.

(C.6) as

$$\begin{aligned}
& \mathbb{E}_j \left[a_j^* \theta \mid s_j = R \right], \\
& = (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x] \right| \mathbb{I}_{\{v_R^*(x)=0, m_R^*(x,0)=x\}} \right] + (1-p) \mathbb{E}_j \left[\underbrace{-\mathbb{E}_j[\theta \mid x, T]}_{\text{he votes } -1} \mathbb{I}_{\{v_R^*(x)=1, \omega=T, m_R^*(x,T)=x\}} \right] \\
& \quad - \underbrace{\left(p\mu_j + (1-p) \mathbb{E}_j \left[\mu_j \mathbb{I}_{\{v_R^*(x)=1, \omega=F, m_R^*(x,F)=\emptyset\}} \right] + (1-p) \mathbb{E}_j \left[\mathbb{E}_j[\theta \mid x] \mathbb{I}_{\{v_R^*(x)=0, m_R^*(x,0)=\emptyset\}} \right] \right)}_{\text{when not receiving any news; he votes } -1}, \\
& = (1-p) \left(\mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x] \right| \mathbb{I}_{\{v_R^*(x)=0, m_R^*(x,0)=x\}} \right] - \mathbb{E}_j \left[\mathbb{E}_j[\theta \mid x] \mathbb{I}_{\{v_R^*(x)=1 \text{ or } v_R^*(x)=0, m_R^*(x,0)=\emptyset\}} \right] \right) - p\mu_j.
\end{aligned}$$

Hence, the difference in ex-ante expected utility between choosing intermediary L over R is

$$\begin{aligned}
& \mathbb{E}_j \left[a_j^* \theta \mid s_j = L \right] - \mathbb{E}_j \left[a_j^* \theta \mid s_j = R \right], \\
& = (1-p) \mathbb{E}_j \left[\left(\left| \mathbb{E}_j[\theta \mid x] \right| + \mathbb{E}_j[\theta \mid x] \right) \mathbb{I}_{\{v_R^*(x)=1 \text{ or } v_R^*(x)=0, m_R^*(x,0)=\emptyset\}} \right] + 2p\mu_j.
\end{aligned}$$

For $\mu_j \in (0, \mu_{k_R^*}]$, the difference is greater than zero and subscriber j strictly prefers intermediary L . When $\mu_j = 0$, the difference is zero ($\mathbb{E}_j[\theta \mid x] \leq 0$ for any x such that $v_R^*(x) = 1$ or $v_R^*(x) = 0, m_R^*(x,0) = \emptyset$) and he feels indifferent between the intermediaries.

- (ii) For $\mu_j \geq \mu_R^{\text{ext}}$, subscriber j can take his optimal vote when observing verified and true news (it may be +1 or -1) while votes +1 when there is no disclosure from

intermediary R :

$$\begin{aligned}
& \mathbb{E}_j \left[a_j^* \theta \mid s_j = R, v_R^*, m_R^* \right], \\
&= (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x] \right| \mathbb{1}_{\{v_R^*(x)=0, m_R^*(x,0)=x\}} \right] + (1-p) \mathbb{E}_j \left[\underbrace{\left| \mathbb{E}_j[\theta \mid x, T] \right|}_{\text{optimal vote}} \mathbb{1}_{\{v_R^*(x)=1, \omega=T, m_R^*(x,T)=x\}} \right] \\
& \quad + \underbrace{\left(p\mu_j + (1-p) \mathbb{E}_j \left[\mu_j \mathbb{1}_{\{v_R^*(x)=1, \omega=F, m_R^*(x,F)=\emptyset\}} \right] + (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x] \right| \mathbb{1}_{\{v_R^*(x)=0, m_R^*(x,0)=\emptyset\}} \right] \right)}_{\text{when not receiving any news, he votes } +1^6}, \\
&= (1-p) \left(\mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x, T] \right| \mathbb{1}_{\{v_R^*(x)=1, \omega=T, m_R^*(x,T)=x\}} \right] + \mathbb{E}_j \left[\mu_j \mathbb{1}_{\{v_R^*(x)=1, \omega=F, m_R^*(x,F)=\emptyset\}} \right] \right) \\
& \quad + (1-p) \mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x] \right| \mathbb{1}_{\{v_R^*(x)=0\}} \right] + p\mu_j,
\end{aligned}$$

where in the last equality we use the fact that $\mathbb{E}_j[\theta \mid x] \geq 0$ for any x such that $v_R^*(x) = 0$ (i.e., $x \geq x_T^{\text{th}}(\mu_R)$). We then can find the difference in ex-ante expected utility between choosing intermediary R over L as

$$\begin{aligned}
& \mathbb{E}_j \left[a_j^* \theta \mid s_j = R \right] - \mathbb{E}_j \left[a_j^* \theta \mid s_j = L \right], \tag{C.8} \\
&= (1-p) \left(\mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x, T] \right| \mathbb{1}_{\{v_R^*(x)=1, \omega=T, m_R^*(x,T)=x\}} \right] + \mathbb{E}_j \left[\mu_j \mathbb{1}_{\{v_R^*(x)=1, \omega=F, m_R^*(x,F)=\emptyset\}} \right] \right. \\
& \quad \left. - \mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x] \right| \mathbb{1}_{\{v_R^*(x)=1\}} \right] \right)
\end{aligned}$$

which is positive due to the triangular inequality and $\mathbb{E}_j[\theta \mid x, T] < 0$ for a nonzero-measured subset of the verified news ($\{x \mid v_R^*(x) = 1\}$); intermediary R is thus strictly preferred. \blacksquare

Proof of Proposition 4.4.

⁶In fact any subscriber j with $\mu_j > \mu_{k_R^*}$ votes +1.

We study the strategies of right-leaning players in the class of true-news-disclosing equilibria.

Verification and disclosure strategy:

By Lemma 4.5, intermediary R does not verify any news $x \geq x_T^{\text{th}}(\mu_R)$. Due to the positive verification cost, intermediary R does not necessarily verify any news $x < x_T^{\text{th}}(\mu_R)$. For unverified news, she conceals the news if the critical subscribers would weakly prefer the opposite action when seeing the news. For verified news, she discloses it iff it is true.

The critical subscribers are centrists: The reasoning follows the proof of Proposition 4.3.

Subscription choices:

The ex-ante expected utility of right-leaning subscribers who subscribe to intermediary L is the same as (C.7) since they take +1 for any news value that is verified by intermediary L .

On the other hand, if $s_j = R$:

- (i) For $\mu_j \in [0, \mu_{k_R}^*]$, the analysis follows the proof of Proposition 4.3.
- (ii) For $\mu_j \geq \mu_R^{\text{ext}}$, one will obtain the difference between following intermediary R and L as described in (C.8). The difference is non-negative due to the triangular inequality; the zero holds iff $\mathbb{E}_j[\theta | x, T] \geq 0$ for any x such that $v_R^*(x) = 1$. ■

C.4 Proofs of Section 4.7

Proof of Lemma 4.7.

(i): Since verification result is perfectly observable to her subscribers when news is disclosed, it is easy to see that after observing the verification result intermediary R conceals the news if and only if her critical subscribers, observing the verification result, would weakly prefer the action that is not favored by her.

(ii): We rewrite the equation for the critical perspective (C.2):

$$\begin{aligned}
0 &= p\mu_{k_R^*} + (1-p) \int_{x \in \mathbb{R} | v_R^*(x)=0, m_R^*(x,0)=\emptyset} \left(qf_F(x) + (1-q)f_T(x; \mu_{k_R^*}) \right) \mathbb{E}_{k_R^*}[\theta|x] dx \\
&\quad + (1-p) \int_{x \in \mathbb{R} | v_R^*(x)=1, m_R^*(x,1,F)=\emptyset} qf_F(x) \mu_{k_R^*} dx \\
&\quad + (1-p) \int_{x \in \mathbb{R} | v_R^*(x)=1, m_R^*(x,1,T)=\emptyset} (1-q)f_T(x; \mu_{k_R^*}) \mathbb{E}_{k_R^*}[\theta|x, T] dx.
\end{aligned}$$

If $\mu_{k_R^*} < 0$, she would conceal any right-leaning news (for which she prefers +1) that will make her critical subscribers take action -1 (either it is verified or not). Since her critical subscribers strictly prefer -1 for any for any $x \leq 0$, they would have negative posterior mean when there is no disclosure. Contradiction.

We now show that $\mu_{k_R^*} < \mu_R$. Suppose not, then intermediary R conceals the news (with its verification result) only when she weakly prefers -1 and her critical subscribers would weakly prefer +1 if seeing the news; as a result her critical subscribers' posterior mean would be positive when there is no disclosure. Contradiction.

(iii): It is immediate from (i) and (ii).

(iv): For any x such that $\mu_R \mathbb{E}_R[\theta|x, T] < 0$, if intermediary R verifies it and learns that it is true, she prefers -1 and knows that her critical subscriber, whose perspective is more left-leaning than hers, will take -1 when observing the news and its verification result. She hence discloses the news. Similarly, if intermediary R verifies it and the news is false, she prefers +1 and her critical subscriber will take +1 when seeing the false news; consequently she discloses it. \blacksquare

Proof of Proposition 4.5.

We still focus on the right-leaning players. We use the insights in previous proofs. **Subscription choices of the subscribers who are more extreme than intermediaries:**

We first show that $v_R^*(x) = 1$ if $x < x_T^{\text{th}}(\mu_R)$. By Lemma 4.7, if intermediary R

verifies some $x < x_T^{\text{th}}(\mu_R)$ and learns its veracity, she will then disclose the news either it is true or false. Given the news value, in both cases she attains the maximal fraction of her subscribers taking the action she favors. Concealing or disclosing unverified news only result in sub-optimal fractions of her subscribers who take her favored action. Given that $C = 0$, it is strictly beneficial for her to verify the news (and then discloses it).

For an extremist subscriber j with $\mu_j \geq \mu_R$, whichever intermediary he chooses, he takes +1 when there is no disclosure. Moreover, he takes +1 for any news value $x \geq x_T^{\text{th}}(\mu_R)$ if it is disclosed. Therefore, one can derive that the difference in his ex-ante expected utility between choosing intermediary R and L comes from their strategies for news that can change his favored action, i.e., for news $x < x_T^{\text{th}}(\mu_j)$. Note that $x_T^{\text{th}}(\mu_j) \leq x_T^{\text{th}}(\mu_R)$ and observe that for any news $x < x_T^{\text{th}}(\mu_R) < 0$, intermediary L may not verify it but will disclose it (cf. Lemma 4.7). Subscriber j then weakly prefers intermediary R who verifies any news that may change his favored action if it is true and enables him to take his optimal action. Specifically,

$$\begin{aligned} & \mathbb{E}_j \left[a_j^* \theta \mid s_j = R \right] - \mathbb{E}_j \left[a_j^* \theta \mid s_j = L \right], \\ &= (1 - p) \left(\mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x, T] \right| \mathbb{1}_{\{\omega=T, v_L^*(x)=0, x < x_T^{\text{th}}(\mu_j)\}} \right] + \mathbb{E}_j \left[\mu_j \mathbb{1}_{\{\omega=F, v_L^*(x)=0, x < x_T^{\text{th}}(\mu_j)\}} \right] \right. \\ & \quad \left. - \mathbb{E}_j \left[\left| \mathbb{E}_j[\theta \mid x] \right| \mathbb{1}_{\{v_L^*(x)=0, x < x_T^{\text{th}}(\mu_j)\}} \right] \right) \end{aligned}$$

which is non-negative and the zero holds iff intermediary L also verifies any news $x < x_T^{\text{th}}(\mu_j)$. ■

C.5 Additional Results and Discussion

The change of the critical perspective with the intermediary's perspective

We discuss the comparative statics for the case when verification is available and costless. The following lemma identifies how the magnitude of an intermediary's perspective affects her critical perspective at the equilibria.

Lemma C.1 *When verification comes with no cost, i.e., $C = 0$, if intermediary i has a more extreme perspective, the unique critical perspective for her is also more extreme. Specifically, for intermediary i , the magnitude of the equilibrium critical perspective $|\mu_{k_i}^*|$ is increasing in $|\mu_i|$.*

Proof:

For the equilibrium in Proposition 4.2, the critical perspective $\mu_{k_R}^* (= \widehat{\mu}_R^*(\emptyset))$ for intermediary R satisfies

$$p\mu_{k_R}^*(\emptyset) + (1-p) \left[\int_{x \in C_R^*} \left(qf_F(x) + (1-q)f_T(x; \mu_{k_R}^*) \right) \mathbb{E}_{k_R^*}[\theta|x] dx \right] = 0.$$

Fix $\mu_{k_R}^*$ and increase μ_R to $\mu_{\tilde{R}} = \mu_R + \tilde{\eta}$ for some infinitesimal $\tilde{\eta} > 0$. Then the set of concealed news C_R^* , for which $\mathbb{E}_{k_R^*}[\theta|x] \leq 0$, will be enlarged because $\mathbb{E}_{\tilde{R}}[\theta|x] > \mathbb{E}_R[\theta|x]$ for every x and the intermediary with $\mu_{\tilde{R}}$ would like to conceal more news. Consequently, the sum will be negative, and $\mu_{k_R}^*$ has to be increased, so that $\mathbb{E}_{k_R^*}[\theta|x]$ is increased for every x and then the set of the concealed news is decreased, to satisfy the equation.

For the equilibrium in Proposition 4.3, the critical perspective for intermediary R

satisfies

$$0 = p\mu_{k_R^*} + (1-p) \left[\int_{\substack{x \in FC_R^* \\ x < x_T^{\text{th}}(\mu_R)}} qf_F(x)\mu_{k_R^*} dx \right] \\ + (1-p) \left[\int_{\substack{x \in TC_R^* \cap FC_R^* \\ x \in [x_T^{\text{th}}(\mu_R), 0)}} \left(qf_F(x) + (1-q)f_T(x; \mu_{k_R^*}) \right) \mathbb{E}_{k_R^*}[\theta|x] dx \right].$$

Again fix $\mu_{k_R^*}$ and consider $\mu_{\bar{R}} = \mu_R + \tilde{\eta}$ for some infinitesimal $\tilde{\eta} > 0$. The set of verified news is decreased (note that $x_T^{\text{th}}(\mu_{\bar{R}}) < x_T^{\text{th}}(\mu_R)$); the set of the news that is concealed without verification can be increased as the intermediary may want to conceal some news in the range $[x_T^{\text{th}}(\mu_{\bar{R}}), x_T^{\text{th}}(\mu_R))$ if it makes $\mathbb{E}_{k_R^*}[\theta|x] \leq 0$. Consequently, the sum will be negative, and $\mu_{k_R^*}$ has to be increased to satisfy the equation. ■

Intuitively, when a intermediary has a more extreme perspective, she only verifies the opposing news that is very extreme, leaving a wider range of opposing news unverified, for which she retains her favored action whereas her critical subscribers prefer the opposite action. Therefore, the intermediary would conceal more news that is unfavorable. Expecting this strategic disclosure, the subscribers will adjust their belief mean more towards to the opposite wing when observing no disclosure.

Verification decisions for news aligned with an intermediary's ideology

As argued in the main context, for any news aligned with intermediary i 's ideology, intermediary i has an incentive to prove its informativeness while she may lose some subscribers taking her favored action if the news turns out to be false. We provide the details of verification decision making as below.

According to Lemma 4.7, intermediary i discloses any news aligned with her ideology. Since subscribers can perfectly observe the verification result for any disclosed news, we can find the threshold perspectives for observing news as a function of news value and verification result, with no dependence on intermediaries' strategies. Specif-

ically, for any x , we denote by $\hat{\mu}(x, T)$ the unique perspective such that the corresponding posterior mean when seeing true news x is zero, i.e., $\mathbb{E}_i[\theta | x, T] = 0$ for $\mu_i = \hat{\mu}(x, T)$. Similarly, for the case when there is no verification, we denote by $\hat{\mu}(x)$ the unique perspective such that $\mathbb{E}_i[\theta | x] = 0$ for $\mu_i = \hat{\mu}(x)$. For any false news, the threshold perspective is zero. By the convexity and monotonicity in perspective, $\hat{\mu}(x)$ is between $\hat{\mu}(x, T)$ and 0.

Given common history G^s , if intermediary i chooses to verify aligned news x , her expected utility is

$$\begin{aligned} & \ell_i(x, T) \mathbb{E}_i[\theta | x, T] \left[G_i^s(+\infty) - 2G_i^s(\hat{\mu}(x, T)) \right] \\ & + (1 - \ell_i(x, T)) \mu_i \left[G_i^s(+\infty) - 2G_i^s(0) \right] - C. \end{aligned}$$

If she does not verify it, she derives expected utility

$$\mathbb{E}_i[\theta | x] \left[G_i^s(+\infty) - 2G_i^s(\hat{\mu}(x)) \right].$$

The difference between verifying and not is then given by

$$\begin{aligned} & \ell_i(x, T) \mathbb{E}_i[\theta | x, T] \underbrace{2 \left[G_i^s(\hat{\mu}(x)) - G_i^s(\hat{\mu}(x, T)) \right]}_{\text{persuade more opposing subscribers}} \\ & - (1 - \ell_i(x, T)) \mu_i \underbrace{2 \left[G_i^s(0) - G_i^s(\hat{\mu}(x)) \right]}_{\text{lose opposing subscribers with perspectives close to neutrality}} - C. \end{aligned}$$

The alternative modeling of continuous action space and utility function of quadratic loss

Despite the challenge in analysis, we do expect some different incentives under this alternative setting. When the verification results are unobservable to the subscribers, the intermediaries may now have an incentive to verify the aligned news and disclose the news if and only if it is true. The reason is that the exact adjustments in the

subscribers' belief mean matter: If the aligned news is true, then the intermediary discloses it and her subscribers move their belief mean closer to the news value, as well as to the intermediary's belief. On the other hand, if the news is false, the intermediary may prefer that her subscribers keep their perspectives rather than update their beliefs speculating the news is likely true. The intermediary conceals the news when it is false and has no incentives to deviate to disclose it, which otherwise makes the subscribers think the news is true and consequently overreact. Similarly, the reasoning is applied to the case of extreme opposing news as the intermediary significantly updates her belief towards the opposite wing when the news is true and wants to persuade her subscribers as well. For the news that around the neutrality, whether it is true or false may not cause a significant shift in the subscribers' belief means and the intermediary may only decide whether to disclose the news.

As the intermediaries disclose verified news if and only if it is true, the extremists can derive the verification benefits whereas the centrists may not due to the indistinguishable false news signalling from the concealed news. However, it is uncertain which intermediary the extremists would choose considering that both intermediaries verify extreme news on both wings.

Bibliography

- D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi. Spread of (mis) information in social networks. *Games and Economic Behavior*, 70(2):194–227, 2010.
- D. Acemoglu, A. Ozdaglar, and J. Siderius. Misinformation: Strategic sharing, homophily, and endogenous echo chambers. Technical report, National Bureau of Economic Research, 2021.
- D. Acemoglu, A. Ozdaglar, and J. Siderius. A model of online misinformation. 2022.
- M. Aigner. On the linegraph of a directed graph. *Mathematische Zeitschrift*, 102:56–61, 1967.
- A. Ajorlou, A. Jadbabaie, and A. Kakhbod. Dynamic pricing in social networks: The word-of-mouth effect. *Management Science*, 64:971–979, 2018.
- N. Alipourfard, B. Nettasinghe, A. Abeliuk, V. Krishnamurthy, and K. Lerman. Friendship paradox biases perceptions in directed networks. *Nature communications*, 11(1):1–9, 2020.
- H. Allcott and M. Gentzkow. Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2):211–36, 2017.
- G. Allon, K. Drakopoulos, and V. Manshadi. Information inundation on platforms and implications. *Operations Research*, 69(6):1784–1804, 2021.
- S. P. Anderson and J. McLaren. Media mergers and media bias with rational consumers. *Journal of the European Economic Association*, 10(4):831–859, 2012.
- S. P. Anderson, J. Waldfogel, and D. Stromberg. *Handbook of Media Economics, vol 1A*. Elsevier, 2016.
- R. J. Aumann. Agreeing to disagree. *The annals of statistics*, pages 1236–1239, 1976.
- A. J. Bahns, C. S. Crandall, O. Gillath, and K. J. Preacher. Similarity in Relationships as Niche Construction: Choice, Stability, and Influence Within Dyads in a Free Choice Environment. *Journal of Personality and Social Psychology*, 112(2):329–355, 2017.
- V. Bala and S. Goyal. Learning from neighbours. *The review of economic studies*, 65(3): 595–621, 1998.
- D. J. Balan, P. DeGraba, and A. L. Wickelgren. Ideological persuasion in the media. *Available at SSRN 637304*, 2004.
- A. Banerjee and R. Somanathan. A simple model of voice. *The Quarterly Journal of Economics*, 116(1):189–227, 2001.
- D. P. Baron. Persistent media bias. *Journal of Public Economics*, 90(1-2):1–36, 2006.
- F. M. Bass. A new product growth for model consumer durables. *Management science*, 15(5):215–227, 1969.

- J. Berger. Word of mouth and interpersonal communication: A review and directions for future research. *Journal of consumer psychology*, 24(4):586–607, 2014.
- A. Bessi and E. Ferrara. Social bots distort the 2016 us presidential election online discussion. *SSRN*, 2016.
- S. Board and M. Meyer-ter Vehn. Learning dynamics in social networks. *Econometrica*, 89(6):2601–2635, 2021.
- E. Breza, A. G. Chandrasekhar, and A. Tahbaz-Salehi. Seeing the forest for the trees? an investigation of network knowledge, 2018.
- C. Budak, D. Agrawal, and A. El Abbadi. Limiting the spread of misinformation in social networks. In *Proceedings of the 20th International Conference on World Wide Web, WWW '11*, page 665–674, New York, NY, USA, 2011. Association for Computing Machinery.
- R. L. Calvert. The value of biased information: A rational choice model of political advice. *The Journal of Politics*, 47(2):530–555, 1985.
- A. Campbell. Word-of-mouth communication and percolation in social networks. *The American Economic Review*, 103(6):2466–2498, 2013.
- A. Campbell, D. Mayzlin, and J. Shin. Managing buzz. *RAND Journal of Economics*, 48(1):203–229, 2017.
- O. Candogan and K. Drakopoulos. Optimal signaling of content accuracy: Engagement vs. misinformation. *Operations Research*, 68(2):497–515, 2020.
- C. P. Chambers and P. J. Healy. Updating toward the signal. *Economic Theory*, 50(3):765–786, 2012.
- J. Chan and W. Suen. A spatial theory of news consumption and electoral competition. *The Review of Economic Studies*, 75(3):699–728, 2008.
- J. Chan and W. Suen. Media as watchdogs: The role of news media in electoral competition. *European Economic Review*, 53(7):799–814, 2009.
- Y.-K. Che and N. Kartik. Opinions as incentives. *Journal of Political Economy*, 117(5):815–860, 2009.
- V. P. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 50:431–1451, 1982.
- M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi. The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3):554–559, 2016.
- P. Diaconis and D. Ylvisaker. Conjugate priors for exponential families. *The Annals of statistics*, pages 269–281, 1979.
- J. Duggan and C. Martinelli. A spatial theory of media slant and voter choice. *The Review of Economic Studies*, 78(2):640–666, 2011.
- B. Efron. Increasing properties of polya frequency function. *The Annals of Mathematical Statistics*, pages 272–279, 1965.
- M. Ellman and F. Germano. What do the papers sell? a model of advertising and media bias. *The Economic Journal*, 119(537):680–704, 2009.
- R. Enikolopov, A. Makarin, and M. Petrova. Social media and protest participation: Evidence from russia. *SSRN*, pages 1–99, 2015.

- P. Erdős and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.
- I. P. Fainmesser and A. Galeotti. Pricing Network Effects. *Review of Economic Studies*, 83(1):165–198, 2016.
- I. P. Fainmesser and A. Galeotti. Pricing network effects: Competition. *American Economic Journal: Microeconomics*, 12(3):1–32, 2020.
- S. L. Feld. Why your friends have more friends than you do. *American Journal of Sociology*, 96(6):1464–1477, 1991.
- A. Galeotti and S. Goyal. Influencing the influencers: a theory of strategic diffusion. *The RAND Journal of Economics*, 40(3):509–532, 2009.
- A. Galeotti, S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv. Network games. *The review of economic studies*, 77(1):218–244, 2010.
- F. Galton. Vox populi. *Nature*, 1907.
- M. Gentzkow and J. M. Shapiro. Media bias and reputation. *Journal of Political Economy*, 114(2):280–316, 2006.
- M. Gentzkow, J. M. Shapiro, and D. F. Stone. Media bias in the marketplace: Theory. In *Handbook of media economics*, volume 1, pages 623–645. Elsevier, 2015.
- F. Germano. On commercial media bias. *Available at SSRN 1374241*, 2008.
- J. Gillin. Politifact’s guide to fake news websites, Apr 2017. Accessed May 10, 2020.
- S. J. Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483, 1981.
- I. Hendel, S. Lach, and Y. Spiegel. Consumers’ Activism: the Cottage Cheese Boycott. *RAND Journal of Economics*, 48(4):972–1003, 2017. ISSN 17562171. doi: 10.1111/1756-2171.12212.
- C. I. Hovland, I. L. Janis, and H. H. Kelley. *Communication and persuasion*. Yale University Press, 1953.
- L. Howell. *Global risks 2013: Eighth edition*. 2013.
- C.-C. Hsu, A. Ajorlou, and A. Jadbabaie. News sharing, persuasion, and spread of misinformation on social networks. *Available at SSRN 3391585*, 2020a.
- C.-C. Hsu, A. Ajorlou, M. Yildiz, and A. Jadbabaie. Information disclosure and network formation in news subscription services. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 5586–5591, 2020b.
- C.-C. Hsu, A. Ajorlou, and A. Jadbabaie. Persuasion, news sharing, and cascades on social networks. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021.
- C.-C. Hsu, A. Ajorlou, and A. Jadbabaie. Persuasion, news sharing, and cascades on social networks. *Available at SSRN 3934010*, 2022.
- M. O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA, 2008.
- M. O. Jackson and X. Tan. Deliberation, disclosure of information, and voting. *Journal of Economic Theory*, 148(1):2–30, 2013.
- A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter: An analysis of a microblogging community. In *International Workshop on Social Network Mining and Analysis*, pages 118–138. Springer, 2007.

- A. P. Jeuland. Brand choice inertia as one aspect of the notion of brand loyalty. *Management Science*, 25(7):671–682, 1979.
- J. T. Jost, S. van der Linden, C. Panagopoulos, and C. D. Hardin. Ideological asymmetries in conformity, desire for shared reality, and the spread of misinformation. *Current Opinion in Psychology*, 23:77–83, 2018.
- M. Jurkowitz, A. Mitchell, E. Shearer, and M. Walker. *U.S. Media Polarization and the 2020 Election: A Nation Divided*. Pew Research Center, 2020.
- J. L. Juul and J. Ugander. Comparing information diffusion mechanisms by matching on cascade size. *Proceedings of the National Academy of Sciences*, 118(46), 2021. ISSN 0027-8424. doi: 10.1073/pnas.2100786118.
- D. Kahneman. *Attention and effort*, volume 1063. Citeseer, 1973.
- E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6): 2590–2615, 2011.
- K. Kawamura. The structure of bivariate poisson distribution. In *Kodai Mathematical Seminar Reports*, volume 25, pages 246–256. Department of Mathematics, Tokyo Institute of Technology, 1973.
- R. Kranton and D. McAdams. Social networks and the market for news. *Working Paper*, 2020.
- H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *WWW '10: Proceedings of the 19th international conference on World wide web*, pages 591–600. ACM, 2010.
- J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '09*, page 497–506. Association for Computing Machinery, 2009.
- L. P. Merlino, P. Pin, and N. Tabasso. Debunking rumors in networks. *arXiv preprint arXiv:2010.01018*, 2020.
- P. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.
- P. Milgrom and J. Roberts. Relying on the information of interested parties. *The RAND Journal of Economics*, pages 18–32, 1986.
- A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 29–42, 2007.
- A. Mitchell, J. Gottfried, J. Kiley, and K. E. Matsa. *Political Polarization & Media Habits*. Pew Research Center, 2014.
- S. Mullainathan and A. Shleifer. The market for news. *American economic review*, 95(4): 1031–1053, 2005.
- M. E. Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.
- N. Newman, R. Fletcher, A. Kalogeropoulos, D. Levy, and R. Nielsen. Reuters institute digital news report. *Reuters Institute for the Study of Journalism*, 2018.

- N. Newman, R. Fletcher, A. Kalogeropoulos, and R. Nielsen. *Reuters institute digital news report*. Reuters Institute for the Study of Journalism, 2019.
- N. P. Nguyen, G. Yan, M. T. Thai, and S. Eidenbenz. Containment of misinformation spread in online social networks. In *Proceedings of the 4th Annual ACM Web Science Conference*, WebSci '12, page 213–222. Association for Computing Machinery, 2012.
- Y. Papanastasiou. Fake news propagation and detection: A sequential model. *Management Science*, 66(5):1826–1846, 2020.
- F. Parise and A. Ozdaglar. Analysis and interventions in large network games. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:455–486, 2021.
- A. Prat and D. Strömberg. The political economy of mass media. *Advances in economics and econometrics*, 2:135, 2013.
- R. Puglisi and J. M. Snyder Jr. Empirical studies of media bias. In *Handbook of media economics*, volume 1, pages 647–667. Elsevier, 2015.
- A. Rapoport. A contribution to the theory of random and biased nets. *Bulletin of Mathematical Biophysics*, 19(4):257–277, 1957.
- L. Rayo and I. Segal. Optimal Information Disclosure. *The Journal of Political Economy*, 118(5):949–987, 2010.
- E. Sadler. Diffusion games. *American Economic Review*, 110(1):225–70, 2020.
- W. Samuelson and R. Zeckhauser. Status quo bias in decision making. *Journal of risk and uncertainty*, 1(1):7–59, 1988.
- A. Saumard and J. A. Wellner. Log-concavity and strong log-concavity: a review. *Statistics surveys*, 8:45, 2014.
- Schema, G. N. Initiative, and Axios. The lifespan of news stories, 2019. <https://newslifespan.com/>. Accessed October 1, 2021.
- R. Sethi and M. Yildiz. Communication with unknown perspectives. *Econometrica*, 84(6):2029–2069, 2016.
- R. Sethi and M. Yildiz. Culture and communication. *Available at SSRN 3263560*, 2019.
- C. Shao, G. L. Ciampaglia, O. Varol, K.-C. Yang, A. Flammini, and F. Menczer. The spread of low-credibility content by social bots. *Nature communications*, 9(1):4787, 2018.
- E. Shearer and K. E. Matsa. News use across social media platforms 2018. *Pew Research Center*, 2018.
- C. Silverman and S. Pham. These are 50 of the biggest fake news hits on facebook in 2018, Dec 2018. Accessed May 10, 2020.
- C. Silverman, J. Lytvynenko, and S. Pham. These are 50 of the biggest fake news hits on facebook in 2017, Dec 2017. Accessed May 10, 2020.
- Z. C. Steinert-Threlkeld, D. Mocanu, A. Vespignani, and J. Fowler. Online social networks and offline protest. *EPJ Data Science*, 4(1):1–9, 2015.
- D. Strömberg. Mass media competition, political competition, and public policy. *The Review of Economic Studies*, 71(1):265–284, 2004.
- X. Su. A model of consumer inertia with applications to dynamic pricing. *Production and Operations Management*, 18(4):365–380, 2009.

- W. Suen. The self-perpetuation of biased beliefs. *The Economic Journal*, 114(495):377–396, 2004.
- A. Sundararajan. Local network effects and complex network structure. *The BE Journal of Theoretical Economics*, 7(1), 2008.
- J. Surowiecki. The wisdom of crowds. *Anchor*, 2005.
- M. Takayasu, K. Sato, Y. Sano, K. Yamada, W. Miura, and H. Takayasu. Rumor diffusion and convergence during the 3.11 earthquake: a Twitter case study. *PLoS one*, 10(4): e0121443, 2015.
- P. Törnberg. Echo chambers and viral misinformation: Modeling fake news as complex contagion. *PLoS one*, 13(9):e0203958, 2018.
- S. Vosoughi, D. Roy, and S. Aral. The spread of true and false news online. *Science*, 359 (March):1146–1151, 2018.
- D. J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771, 2002.
- J. Weedon, W. Nuland, and A. Stamos. Information operations and facebook. Technical report, 2017.
- C. B. Williams and G. J. J. Gulati. Social networks in political campaigns: Facebook and the congressional elections of 2006 and 2008. *New Media & Society*, 15(1):52–71, 2013.
- H.-C. Yeh. Six multivariate zipf distributions and their related properties. *Statistics & probability letters*, 56(2):131–141, 2002.
- Z. Zhao, J. Zhao, Y. Sano, O. Levy, H. Takayasu, M. Takayasu, D. Li, J. Wu, and S. Havlin. Fake news propagates differently from real news even at early stages of spreading. *EPJ Data Science*, 9(1):1–14, 2020.
- Z. Zhou, R. Bandari, J. Kong, H. Qian, and V. Roychowdhury. Information resonance on Twitter: watching Iran. In *Proceedings of the first workshop on social media analytics*, pages 123–131, 2010.