

Article

## From Dark to Light: The Many Shades of Sharing Misinformation Online

Miriam J. Metzger<sup>1,\*</sup>, Andrew J. Flanagin<sup>1</sup>, Paul Mena<sup>2</sup>, Shan Jiang<sup>3</sup> and Christo Wilson<sup>3</sup>

<sup>1</sup> Department of Communication, University of California, Santa Barbara, CA 93106, USA;  
E-Mails: metzger@ucsb.edu (M.J.M.), flanagin@ucsb.edu (A.J.F.)

<sup>2</sup> Writing Program, University of California, Santa Barbara, CA 93106, USA; E-Mail: pmena@ucsb.edu

<sup>3</sup> Khoury College of Computer Sciences, Northeastern University, Boston, MA 02115, USA;  
E-Mails: sjiang@ccs.neu.edu (S.J.), cbw@ccs.neu.edu (C.W.)

\* Corresponding author

Submitted: 30 June 2020 | Accepted: 7 August 2020 | Published: 3 February 2021

### Abstract

Research typically presumes that people believe misinformation and propagate it through their social networks. Yet, a wide range of motivations for sharing misinformation might impact its spread, as well as people's belief of it. By examining research on motivations for sharing news information generally, and misinformation specifically, we derive a range of motivations that broaden current understandings of the sharing of misinformation to include factors that may to some extent mitigate the presumed dangers of misinformation for society. To illustrate the utility of our viewpoint we report data from a preliminary study of people's dis/belief reactions to misinformation shared on social media using natural language processing. Analyses of over 2,5 million comments demonstrate that misinformation on social media is often disbelieved. These insights are leveraged to propose directions for future research that incorporate a more inclusive understanding of the various motivations and strategies for sharing misinformation socially in large-scale online networks.

### Keywords

credibility; fake news; misinformation; news sharing

### Issue

This article is part of the issue "Dark Participation in Online Communication: The World of the Wicked Web" edited by Thorsten Quandt (University of Münster, Germany).

© 2021 by the authors; licensee Cogitatio (Lisbon, Portugal). This article is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

### 1. Introduction

Few would dispute that the web has grown more wicked in recent years with the rise of 'fake news.' Indeed, the creation and spread of misinformation online for political or financial gain presents one of the greatest challenges to democratic society in modern history. Fake news and misinformation are said to increase political polarization, alter voters' perceptions of candidates and issues, and erode trust in critical democratic institutions (Allcott & Gentzkow, 2017; Ciampaglia, Mantzarlis, Maus, & Menczer, 2018; Hochschild & Einstein, 2015). Moreover, propelled largely through social media, fake news has been shown to spread faster and further than 'real' news (Lazer et al., 2018; Vosoughi, Roy, & Aral, 2018).

The danger of misinformation is that people will find such information to be credible. Fake news is thought to be dangerous because people are likely to believe the false information, and because it is spread through people's trusted contacts via social media, it will be similarly trusted by others. Hence, the presumed threat of fake news is that people will erroneously believe false information they encounter and that they will in turn propagate misinformation to others. Yet, there are likely many motivations for sharing information—including fake news—within one's social network that might impact both its spread and people's belief of it. For example, people might share misinformation for entertainment purposes, sarcastic reasons, or to illustrate a point counter to the message promoted in a false

news story. Under such circumstances, the danger of fake news may be less than feared or, perhaps, even mitigated or reversed.

This article interrogates common presumptions about sharing misinformation by analyzing people's motives for sharing fake news stories on social media. Drawing on literatures on news sharing, rumors, gossip, and urban legends suggests that a wide range of motivations exists for sharing misinformation, and that various motivations imply both positive and negative outcomes for democratic processes. Preliminary data from a research program that aims to understand dis/belief in fake news and misinformation in social media are then presented as proof-of-concept for some of the ideas about motives for sharing misinformation that are advanced in this article.

## 2. Motivations for Sharing News Information

Research on people's motivations to share news in general, and specifically via social media, is extensive. Motives for sharing news include: acting as an opinion leader, advocating for one's own beliefs, socializing, gaining social status, sharing experiences with others, and informing others. Although the news media decide which stories to publish, users make judgments about the value of what is published, which in turn affects news consumption. Through such processes of 'secondary gatekeeping,' whereby users share news stories with others (Shoemaker & Vos, 2009; Singer, 2014), information consumers propose what others should read, establishing "not only what is valuable to them as individuals but also what they believe will be important, interesting, entertaining, or useful to others" (Singer, 2014, p. 58). In this way, users filter, evaluate, and finally share news information.

This secondary gatekeeping capability has been enhanced by social media, where individuals are able to actively participate in the dissemination of news to large audiences. Indeed, users may experience a sense of agency when sharing news on social media by feeling that they have some control over what is posted on a social network (Oeldorf-Hirsch & Sundar, 2015). In this sense, the desire to act as a gatekeeper of information or an opinion leader may be a fundamental motivation for news sharing. Indeed, research has found that self-perceptions of opinion leadership influence users' news sharing intention (Ma, Sian Lee, & Hoe-Lian Goh, 2014) and that political expression enhances people's motivations to self-present as politically active on social media (Lane et al., 2019). Moreover, users may feel motivated to share news that supports their own views and contradicts dissenting beliefs since advocating for one's own beliefs can be a motivator of news sharing (Dafonte-Gómez, 2018), consistent with evidence of a relationship between news sharing and ideology spreading (Lottridge & Bentley, 2018).

In addition, studies have found that social gratifications and social status are motivators of news sharing

on social media (Bright, 2016; Choi, 2016; Lee & Ma, 2012). A sense of connection with the online community is developed through news sharing (Lee & Ma, 2012) and sharing information with others contributes to satisfying the need for social interaction, which "helps people clarify their opinions, and gives them an opportunity to work out their personal positions relative to media content" (Weeks & Holbert, 2013, p. 215). Additionally, social standing can be achieved by sharing news deemed useful to those receiving it, which can "make the person passing it on appear well informed and intelligent" (Bright, 2016, p. 346). News sharing may thus help people gain status among peers (Bruns, 2018; Choi, 2016; Lee & Ma, 2012), particularly if they feel that they were informed earlier than others (Kubey & Peluso, 1990).

In some cases, sharing news may be a by-product of people's psychological need to share experiences with others (Harber & Cohen, 2005; Kubey & Peluso, 1990), consistent with the notion that people who have experienced disturbing events are emotionally compelled to share their experiences: "An unintended but often real consequence of their efforts to unburden themselves, we believe, is to inform members of their social networks of valuable news" (Harber & Cohen, 2005, p. 383). In this way, people may share news to relieve their own feelings (Kubey & Peluso, 1990). This emotional component of news sharing has been examined by researchers who found that emotional arousal can explain transmission of news (Berger & Milkman, 2012; Dafonte-Gómez, 2018), such that news evoking high-arousal positive or negative emotions is shared more than news evoking low-arousal emotions.

Finally, researchers have also identified altruistic motives for news sharing (Bruns, 2018; Chadwick & Vaccari, 2019; Dafonte-Gómez, 2018; Kümpel, Karnowski, & Keyling, 2015). For example, surveys among U.K. social media users found that "to inform others" and "to express my feelings" were the most important motivations for news sharing (Chadwick & Vaccari, 2019). Users feel motivated to serve others' information needs and share meaningful news with their online community, to such a degree that social media users are even prone to adapt news sharing to their audience's interests (Rudat, Buder, & Hesse, 2014).

## 3. Motivations for Sharing Misinformation

While plenty of research exists on why people share news online, only a handful of studies have examined people's motivations for sharing 'fake' news stories (e.g., Chadwick & Vaccari, 2019; Chen & Sin, 2013; Chen, Sin, Theng, & Lee, 2015; Duffy, Tandoc, & Ling, 2019; Talwar, Dhir, Kaur, Zafar, & Alrasheedy, 2019). Moreover, most studies assume that people do not realize the information they share is false—in other words, that the sharing of misinformation, including fake news, is unintentional.

The picture is more complicated with misinformation, however, because people may or may not know

the information is false when they share it, and thus may spread the misinformation either intentionally or unintentionally (Lawrie, 2019). Unintentional fake news sharing may be motivated by self-expression and socialization (Chen et al., 2015). For instance, Chen and Sin (2013) found that people share misinformation mainly to obtain other's opinions, to express their own opinions, and to interact with others. Similarly, researchers have suggested that the motivation to build relationships can lead people to unintentionally share fake news (Duffy et al., 2019). In fact, the factors that drive people to share real news, including emotional impact and relevance, are likely to be the same factors that make fake news highly shared as well (Duffy et al., 2019).

However, when the sharer knows the information is false, then other motives likely come into play. Researchers suggest that people knowingly share fake news because it conforms to their prior views and in order to maintain positive social relations (Duffy et al., 2019). Intentional fake news sharing has been positively associated with social media use behaviors such as self-disclosure, online trust, fear of missing out, and social media fatigue (Talwar et al., 2019). More specifically, people who tend to disclose more information in social media generally, those who receive a fake news item from a trusted source, people who seek popularity and a sense of belonging (i.e., fear social exclusion), and who experience information overload are more likely to pass fake news on to others in their social networks.

In addition, research shows that people who seek to entertain, troll, or debate with others are more likely to engage in intentional misinformation sharing, while those who seek to persuade or inform others are less likely to do so (Chadwick, Vaccari, & O'Loughlin, 2018). This suggests that some people share fake news as a way to disrupt political dialogue, while others with more civic incentives share fake news as part of legitimate political debate:

Those motivated to debate may see sharing problematic news as a cultural norm; a practice that is simply part of 'what it takes' to engage politically on social media in order to attract attention and nudge others to take positions. (Chadwick et al., 2018, p. 4269)

Among participants in the study, 8.9% admitted to sharing news that they thought was made up when they shared it and 17.1% deliberately shared news that was exaggerated. Another study showed that 17.3% of British news sharers on social media knowingly shared news that they thought was made up (Chadwick & Vaccari, 2019). Similarly, a Pew Research Center survey found 14% of U.S. adults said they had shared a news story they knew was fake (Barthel, Mitchell, & Holcomb, 2016). And sharing fake news may be underreported due to social desirability biases that likely suppress self-reported data on misinformation sharing.

Finally, fake news sharing has some similarities to rumor, gossip, and urban legends. Research by Guerin

and Miyazaki (2006), for instance, suggests "the primary function of telling rumors, gossip, and urban legends is not to impart information to the listener or alleviate listener anxiety about the topic, but to entertain or keep the listener's attention, thereby enhancing social relationships" (p. 23). Others have argued that both rumor sharing and fake news sharing fulfill three motivations: "To cope with uncertainty, build relationships, and for self-enhancement" (Duffy et al., 2019, p. 3). In this fashion, people engage in rumor as a way to collectively reduce uncertainty, enhance interpersonal relationships, and feel positive about themselves (Bordia & DiFonzo, 2005), as may be the case with sharing news misinformation as well.

#### **4. Broadening the Range of Motivations for Sharing Misinformation**

There is still much to be learned about why people share misinformation. Theoretical explanations are nascent, and only a few empirical studies have been conducted to date (Talwar et al., 2019). There is certainly a host of reasons in addition to those discussed above why people share fake news and misinformation. For example, intentionally sharing misinformation may be done to engage in 'collective fact-checking,' which likely derives from the more general motivation for uncertainty reduction in information sharing as mentioned earlier. Fake news stories can be sensationalistic and often contain other elements such as extreme partisan information that could arouse suspicion about their veracity (Mourão & Robertson, 2019). Consequently, recipients of such stories may share the suspected misinformation with their friends and followers on social media in an effort to crowdsource its truthfulness, and thus reduce uncertainty as to the information's credibility.

Another potential consequence of the sensationalism of many fake news stories is that their absurdity can provide fodder for ridicule. Sarcasm or mockery of such stories may therefore be another reason to share news that one knows is false. This reason could be conceived as a specialized subset of sharing for socializing and entertainment purposes. In such cases, the misinformation is shared for the purpose of highlighting its falsity to others. This is different type of social entertainment motive than passing a story to a friend you think would find it interesting or spreading a titillating rumor for its shock value, which have been discussed previously in the news sharing literature (Lee & Ma; 2012; see also Kümpel et al., 2015).

A motivation that ultimately shares the same purpose as sarcasm is sharing misinformation for educational purposes. Although 'informing others' is a well-known reason for sharing news information (e.g., Chadwick & Vaccari, 2019), our concept of education differs from prior notions of imparting information to others because it is not about informing recipients of the *content* of the story, but rather about informing

them about the *credibility* of the story itself. Such sharing thus functions to debunk misinformation, and serves as a warning to others not to believe the information. Alerting recipients to potential misinformation may stem from simple altruistic motives, or may be self-serving if it is done to demonstrate one’s moral or intellectual superiority. If the latter, it may be a new variant of the ‘status seeking’ category of news sharing motives, as it casts the sharer as a well-informed opinion leader.

Thus, while social media users share fake news that they do not know or suspect to be false for the same reasons that they share any other type of news, there exist motivations that are unique to sharing misinformation that one knows to be false. An important issue that has not been explored is that a person’s motivation for sharing information is likely to have a strong impact on the credibility beliefs of the recipient. If a sharer appears to believe the information they share, the recipient may also. Yet, if a sharer appears not to believe the information (for instance, if the information is shared for sarcastic or educational purposes), the recipient may not either. But how can the recipient know what a sharer believes?

Although another person’s psychological motivation for sharing a piece of information and his or her credibility evaluation might be difficult for recipients to discern in many cases, the communication surrounding the shared information can offer useful clues. For example, a prominent fake news story shared on social media in 2016 said “Pope Francis Shocks World, Endorses Donald Trump for President.” On social media, this post could be accompanied by various captions or comments that may signal different degrees of belief in the misinformation. For example, if the sharer were to write: “See, this proves I was right all along, I knew he would support Trump!” or even “Oh no, this is awful!” this might signal that the sharer finds the information to be credible. On the other hand, if the commentary instead read: “Oh yeah right, sure he did,” or “Get a load of this, it’s the funniest thing I’ve seen in a long time!” or “What kind of dope would believe this bull\*\*\*\*???” it would indicate skepticism or disbelief on the part of the sharer (Figure 1).

Recent research supports the assertion that recipients of information notice cues regarding credibility from

sharers. A study by Colliander (2019) found that exposure to others’ comments that were critical versus supportive of a fake news story resulted in less favorable attitudes toward the news article, prompted participants to write more negative comments themselves about the article, and lowered their intentions to share the article in their social networks. Colliander concluded that “social media users seem to use the comments of other people as a guide for how to respond to disinformation online” (p. 208) based on human psychological tendencies toward conformity and a desire to maintain a positive self-concept. Winter, Bruckner, and Krämer (2015) observed similar social influence effects of negatively—versus positively-valenced mainstream news comments.

The foregoing discussion raises the possibility that the danger to democracy and to society due to misinformation and fake news circulating in social media may be mitigated at least partially to the extent that people spread fake news stories as a tool to engage with others in deliberative communication about its truthfulness or to signal its falsity. As U.S. Supreme Court Justice Louis Brandeis famously wrote, the remedy to ‘bad’ speech is not censorship, it is more speech, because through discussion and debate false information will be revealed as false. Sharing misinformation for purposes such as collective fact-checking, sarcasm, or education calls out the information as false, and thus may be good for deliberative democracy. In this way, and borrowing a term from van Heekeren (2020), social media commentary surrounding misinformation may function as a ‘curative’ to the problem of fake news, particularly over time as people reap the benefits of the collective sharing and debunking of misinformation.

To illustrate this potential, a preliminary study is next described in which the content of the communication surrounding misinformation is analyzed to gain insight into the range of reasons that people propagate such information, and how it may over time alter information interpretation. More specifically this research analyzed social media users’ commentary on known fake news stories as a means to better understand public perceptions of dis/belief in, and credibility assessments of, misinformation shared online.

**Comments Signaling Belief**

- OMG have you seen this?
- I knew he would support Trump!
- Oh no, this is awful!



**Comments Signaling Disbelief**

- Oh yeah right, *sure* he did. Fake news!
- Get a load of this, it’s the funniest thing I’ve seen in a long time!
- What kind of dope would believe this bull\*\*\*\*???

**Figure 1.** Example of a fake news post with examples of hypothetical commentary that signal a range of beliefs by the sharer. Source: “Pope Francis shocks world” (2016).

## 5. Preliminary Study of Dis/Belief in Misinformation Shared in Social Media

This study aimed to quantify the prevalence of dis/belief in misinformation at scale using a machine learning approach. The research was conducted in four steps. Step 1 involved identifying social media posts that contain misinformation and collecting user comments in response to those posts to analyze. Step 2 required reading the comments to determine if they reflect the commenter's belief or disbelief in the misinformation. Each comment was manually labeled by independent coders, and these data were used in Step 3 to examine language differences in comments expressing belief and disbelief, and to test if such differences could be used to build a classifier to identify dis/belief with reasonable accuracy. Finally, the classifier was leveraged in Step 4 to measure dis/belief at scale to answer the following RQs:

RQ1: To what extent does the public believe misinformation shared via social media?

If misinformation is not believed to a large extent, that would discount much of its presumed danger to society. And if evidence of disbelief is found in comments shared by users, there is potential for those comments to influence others' beliefs in a cascading manner:

RQ2: Is there is a time effect for expressed disbelief in misinformation, where the public gradually realizes the truth after a false claim is made, and therefore belief in false claims decreases over time?

This question is raised in light of possibilities of crowd-sourcing for misinformation detection, including via formal reporting mechanisms (i.e., letting users flag fake news; see Tschisatschek, Singla, Gomez Rodriguez, Merchant, & Krause, 2018) or by the social influence processes proposed herein, whereby recipients are influenced by sharers' critical comments on fake news stories. While the technical details of the research are presented in full elsewhere (citation blinded for review), the methods employed and research results at each step are summarized in the next section.

### 5.1. Step 1

A sample of social media posts containing at least some misinformation was collected from a census of PolitiFact's fact-checked articles between January 1 to June 1, 2019. Posts were manually identified as containing misinformation that originated from Twitter, which was used because its API allowed the collection of user comments written in response to posts containing misinformation at the time of data collection, unlike other social media platforms. Next, comments (i.e., tweets) in response to the posts were collected. Using the fact-checked posts as seeds, a 1% sample of the tweet stream

was queried to capture all comments to the seed posts. To reduce noise, only posts that had > 50 comments were retained, which resulted in 6,809 comments to analyze in Step 2.

### 5.2. Step 2

Once the comments written in response to posts containing misinformation were collected, they were manually annotated with belief and disbelief labels. Two independent, trained coders read each comment and provided a binary label: either disbelief (i.e., the person who wrote the comment does not appear to believe the misinformation) or belief (i.e., the person who wrote the comment does appear to believe the misinformation). These two labels are mutually exclusive but not necessarily complementary, that is, although a tweet is not expected to show both belief and disbelief, it can show neither.

The inter-annotator percent agreement (i.e., the number of agreed labels over the total count) was used to evaluate intercoder reliability. 66.7% of the coder pairs were above 80% agreement, 88.9% were above 70% agreement, and only two were below 60% agreement, suggesting an acceptable level of agreement among annotators, especially for a relatively subjective task. To obtain a final label for each comment, a third independent coder read through all cases where the two original coders disagreed and provided a final tie-breaking judgement. Out of 6,809 tweets, 2,399 (35.2%) were labeled as expressing disbelief, 1,282 (18.8%) were labeled as expressing belief, 3,128 (45.9%) were labeled as neither, and none (0%) were labeled as both.

### 5.3. Step 3

Using the labeled dataset, a lexicon-based exploratory analysis of the language used across the comments expressing belief and disbelief was conducted. Two lexicons were employed: (a) LIWC (Tausczik & Pennebaker, 2010), the most widely-used lexicon for understanding psychometric properties of language, and (b) ComLex (Jiang & Wilson, 2018), a more contextual lexicon built from social media comments to misinformation, containing additional domain-specific categories, e.g., 'fake,' 'fact,' 'hate speech.' This technique provided a frequency for each category in the lexicon, which allowed a comparison of the distributions of language frequencies between comments expressing belief, disbelief, and neither by performing t-tests with Bonferroni corrections.

Results showed that comments expressing disbelief contained significantly more falsehood awareness language, including referrals to falsehood such as 'lie, propaganda' and 'fake, false'; referrals to the truth 'fact, research'; and negative character portraits such as 'liar, crook' and 'stupid, dumb.' Comments expressing disbelief also contained more negative emotions and negation (e.g., 'no, not') and less positive emotions and discrepancy words (e.g., 'should, would'). By contrast,

comments expressing belief contained fewer falsehood awareness signals (e.g., ‘fake, false’) and negative character portrait (‘stupid, dumb’). Comments expressing belief also contained more exclamation (e.g., exclamation marks, ‘!, yay’); discrepancy words; and fewer negative reactions such as swearing (e.g., ‘damn, fuck’) and anger (e.g., ‘hate, kill’). All t-tests were significant at least  $p < .01$ .

Next, natural language processing models were applied to build a classifier to automatically identify and label comments that express dis/belief in misinformation. Different types of classifiers were experimented with and the neural transfer-learning based classifier RoBERTa was found to be the best in terms of accuracy of classification, as determined by F1 scores (binary F1 scores were around 0.86 for disbelief and 0.80 for belief).

#### 5.4. Step 4

The final step was to use the classifier on a larger dataset to label dis/belief expressed in responses to misinformation at scale. The dataset, collected previously by Jiang and Wilson (2018), contained 1,672,687 comments collected from Facebook, 113,687 from Twitter, and 828,000 from YouTube written in response to 5,303 fact-checked claims. These claims are drawn from the entire archive of Snopes and PolitiFact’s articles between their founding and January 9, 2018. The dataset was fed to the RoBERTa-based classifier from Step 3 to answer the research questions.


RQ1 asked about the prevalence of dis/belief in misinformation. To analyze this, all of the original claims were parsed into three mis/information types: (a) true, if the claims were rated as ‘true’ by Snopes or PolitiFact—these claims contain no misinformation; (b) mixed, if the claims were rated as ‘mostly true,’ ‘half true,’ or ‘mixed’—these claims contain some misinformation but also some truth; and (c) false, if the claims were rated as ‘mostly false,’ ‘false,’ or ‘pants on fire!’—these claims contain mostly falsehood. Comments in the dataset that were posted after its corresponding factcheck article was published were then filtered out. Finally, the remaining comments were grouped by the mis/information type, and their dis/belief labels were averaged.

Results showed that as the veracity of claims decreased, disbelief increased while belief decreased (Figure 2). Specifically, 12%, 14%, and 15% of comments expressed disbelief in response to true, mixed, and false claims, respectively; whereas 26%, 21%, and 20% of comments expressed belief in response to true, mixed, and false claims, respectively.

Notably, the majority of comments indicated no sure indication of belief or disbelief and therefore yield no information regarding people’s reactions to the claims spread via social media. Among comments indicating dis/belief, however, several implications emerge. First, an overall skepticism is evident from the data: not only do many people not believe fake information, but many people similarly disbelieve (12%) true information (and only 26% report believing true information). Thus, people appear generally distrustful of claims shared via social media in this study, suggesting a certain level of cautiousness that might guard against the deleterious effects of misinformation, consistent with the underlying rationale that not all misinformation is uniformly harmful. One potential explanation is that the partisan environment drives people to suspect any claims from the opposite ideological group regardless of veracity (Hochschild & Einstein, 2015). Another, though less likely, explanation is that people tend to be skeptical of all claims in social media, even when the claim is consistent with existing facts. Both explanations warrant further investigation.

Second, there is some indication that a portion of people commenting on misinformation are doing so to express their own disbelief in the content. This again supports the perspective that not all misinformation sharing implies tacit endorsement of the claims made therein, and provides evidence that at least some users are leveraging the power of social media in a fashion that can serve as a partial corrective to misinformation. If users are indicating that they do not believe the misinformation they see in social media, it suggests alternative, and perhaps useful, responses regarding content that is in turn shared with others.

That said, the difference in the prevalence of dis/belief across the mis/information types is small (e.g., 3% more people disbelieve versus believe false



**Figure 2.** Prevalence of dis/belief for true, mixed, and false claims on social media. Notes: For disbelief (left graph), as the veracity of the claims decreases, the prevalence of expressed disbelief increases. For belief (right graph), as the veracity of the claims decreases, the prevalence of expressed belief also decreases.

statements). Yet, at scale, even small differences represent a substantial volume of occurrences where people are more likely to disbelieve versus believe false claims. Moreover, taken alone, the fact that in this study 15% of users report not believing misinformation is notable.

The final analysis sought to answer RQ2, which asked whether belief in false claims decreases over time. The data revealed a very small-time effect, where disbelief increases 0.001% and belief decreases 0.002% per day after the initial false claim appeared. Notably, at scale this small-time effect would yield a large number of changes in dis/belief. This effect could be a result of seeing prior comments that are critical of the claim (i.e., a sharer's expressed disbelief in the misinformation), but further studies that move beyond machine learning techniques are needed to determine if this explanation is correct.

For example, social scientific methodologies such as surveys would be useful to elicit the variety of reasons why people share misinformation that they know to be false in social media, how often they accompany it with commentary signaling its veracity, and also the extent that social media users report receiving such information from their network contacts. Experiments similar to Colliander (2019) could shed light on the power of negative comments to alter recipients' credibility beliefs, especially if conducted unobtrusively where actual behavioral data (e.g., misinformation sharing) could be observed.

In sum, the results of our preliminary study show evidence that when people share misinformation via social media, they do not always take it at face value. For example, over a third of the comments analyzed in Step 2 reflected disbelief in tweets containing misinformation, and the wider analysis in Step 4 revealed that about one in six people disbelieved false information they encountered in various social media platforms. Moreover, our analyses found that belief in false claims decreases slightly over time, which might at least partly be a result of exposure to social comments critical of the misinformation, although future research will need to confirm this potential explanation. These data should prompt scholars to expand their thinking on why people share misinformation beyond unintentional sharing of (believed) misinformation, to suggest a healthy process of intentional social debunking of fake news that is rarely examined in the literature.

## 6. Conclusion

Although the creation and spread of misinformation online represent a serious challenge to democratic society, the nature and extent of the problem might to some degree be overstated. If, as our preliminary evidence suggests, misinformation on social media is often disbelieved, and to the extent that those sharing it are doing so for reasons that expose and help to stem the spread of misinformation, then shared misinformation

is in fact not universally harmful and its propagation is not always and necessarily detrimental. By examining research on motivations for sharing news information generally, and misinformation specifically, we derive a range of motivations—including entertainment, sarcasm, and education—that broaden our understanding of the sharing of misinformation to include factors that may help to mitigate its danger.

Yet, because research to date has almost exclusively characterized information sharers as passive diffusers of misinformation, and information consumers as universally receptive to it, little is known about the prevalence and dynamics of information consumers' potentially defensive or constructive reactions to misinformation. To address this deficiency research must consider a host of factors that consider the sharing and consumption of misinformation as a more active, critical, and strategic process. For instance, future research is required to establish baseline data on the wide range of motivations invoked by those sharing misinformation, the extent to which they are aware of its in/accuracy, and the degree of intentionality in its sharing with others.

Relatedly, research must assess the extent to which misinformation that is shared with the intention of debunking it is in fact interpreted appropriately, and the degree to which social, psychological, and contextual factors affect such outcomes. Moreover, the complex network dynamics of social media information sharing (e.g., which network nodes are most influential, in what domains), coupled with the relational attributes that guide the mis/interpretation of shared misinformation (e.g., the relations among information sources and recipients), need to be examined to understand social influence processes in this context. Finally, paralinguistic cues (e.g., emojis, likes, etc.) should be examined as markers of endorsement, signals of alternative motivations (e.g., sarcasm), and cues demonstrating the extent of shared group identity, all of which may be interrelated in complex ways to affect misinformation belief.

In sum, insights into the wide range of motivations for sharing misinformation via social media suggest that the spread of misinformation should be viewed as a dynamic process among active users driven by a wide range of factors. Viewed this way, the negative aspects of misinformation may to some degree be mitigated and can only be fully understood as future research incorporates a more inclusive understanding of the various motivations and strategies for sharing misinformation socially in large-scale online networks.

## Acknowledgments

This research was supported in part by NSF grant IIS-1553088. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

## Conflict of Interests

The authors declare no conflict of interests.

## References

- Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2), 211–236.
- Barthel, M., Mitchell, A., & Holcomb, J. (2016, December 15). Many Americans believe fake news is sowing confusion. *Pew Research Center*. Retrieved from <http://www.journalism.org/2016/12/15/many-americans-believe-fake-news-is-sowing-confusion>
- Berger, J., & Milkman, K. L. (2012). What makes online content viral? *Journal of Marketing Research*, 49(2), 192–205.
- Bordia, P., & DiFonzo, N. (2005). Psychological motivations in rumor spread. In G. A. Fine, C. Heath, & V. Campion-Vincent (Eds.), *Rumor mills: The social impact of rumor and legend* (pp. 87–101). New York, NY: Aldine Press.
- Bright, J. (2016). The social news gap: How news reading and news sharing diverge. *Journal of Communication*, 66, 343–365.
- Bruns, A. (2018). *Gatewatching and news curation*. New York, NY: Peter Lang.
- Chadwick, A., & Vaccari, C. (2019). *News sharing on UK social media: Misinformation, disinformation, and correction* (Report No. O3C 1). London: Loughborough University. Retrieved from [https://repository.lboro.ac.uk/articles/News\\_sharing\\_on\\_UK\\_social\\_media\\_misinformation\\_disinformation\\_and\\_correction/9471269](https://repository.lboro.ac.uk/articles/News_sharing_on_UK_social_media_misinformation_disinformation_and_correction/9471269)
- Chadwick, A., Vaccari, C., & O’Loughlin, B. (2018). Do tabloids poison the well of social media? Explaining democratically dysfunctional news sharing. *New Media & Society*, 20(11), 4255–4274.
- Chen, X., & Sin, S. J. (2013). Misinformation? What of it? Motivations and individual differences in misinformation sharing on social media. *Proceedings of the American Society for Information Science and Technology*, 50(1), 1–4. Retrieved from [www.asis.org/asist2013/proceedings/submissions/posters/23poster.pdf](http://www.asis.org/asist2013/proceedings/submissions/posters/23poster.pdf)
- Chen, X., Sin, S. J., Theng, Y., & Lee, C. S. (2015). Why students share misinformation on social media: Motivation, gender, and study-level differences. *The Journal of Academic Librarianship*, 41, 583–592.
- Choi, J. (2016). Why do people use news differently on SNSs? An investigation of the role of motivations, media repertoires, and technology cluster on citizens’ news-related activities. *Computers in Human Behavior*, 54, 249–256.
- Ciampaglia, G. L., Mantzarlis, A., Maus, G., & Menczer, F. (2018). Research challenges of digital misinformation: Toward a trustworthy web. *AI Magazine*, 39(1), 65–74.
- Colliander, J. (2019). “This is fake news”: Investigating the role of conformity to other users’ views when commenting on and spreading disinformation in social media. *Computers in Human Behavior*, 97, 202–215.
- Dafonte-Gómez, A. (2018). Audience as medium: Motivations and emotions in news sharing. *International Journal of Communication*, 12, 2133–2152.
- Duffy, A., Tandoc, E., & Ling, R. (2019). Too good to be true, too good not to share: The social utility of fake news. *Information, Communication & Society*, 23(13), 1–15.
- Guerin, B., & Miyazaki, Y. (2006). Analyzing rumors, gossip and urban legends through their conversational properties. *Psychological Record*, 56, 23–34.
- Harber, K., & Cohen, D. (2005). The emotional broadcaster theory of social sharing. *Journal of Language and Social Psychology*, 24, 382–400.
- Hochschild, J. L., & Einstein, K. L. (2015). *Do facts matter? Information and misinformation in American politics*. Norman, OK: University of Oklahoma Press.
- Jiang, S., & Wilson, C. (2018). Linguistic signals under misinformation and fact-checking: Evidence from user comments on social media. *Proceedings of the ACM on Human-Computer Interaction*, 2, 1–23.
- Kubey, R. W., & Peluso, T. (1990). Emotional response as a cause of interpersonal news diffusion: The case of the space shuttle tragedy. *Journal of Broadcasting and Electronic Media*, 34, 69–76.
- Kümpel, A. S., Karnowski, V., & Keyling, T. (2015). News sharing in social media: A review of current research on news sharing users, content, and networks. *Social Media + Society*. 1(2). <https://doi.org/10.1177/2056305115610141>
- Lane, D. S., Lee, S. L., Liang, F., Kim, D. H., Shen, L., Weeks, B. E., & Kwak, N. (2019). Social media expression and the political self. *Journal of Communication*, 69(1), 49–72.
- Lawrie, L. (2019). *The fake news crisis of 2016: The influence of political ideologies and news trust on news consumer “innocent sharing”* (Unpublished doctoral dissertation). Wichita State University, Wichita, KA.
- Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., . . . Rothschild, D. (2018). The science of fake news. *Science*, 359(6380), 1094–1096.
- Lee, C. S., & Ma, L. (2012). News sharing in social media: The effect of gratifications and prior experience. *Computers in Human Behavior*, 28(2), 331–339.
- Lottridge, D., & Bentley, F. R. (2018, April 21). Let’s hate together: How people share news in messaging, social, and public networks. In R. Mandryk & M. Hancock (Eds.), *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). New York, NY: Association for Computing Machinery.
- Ma, L., Sian Lee, C., & Hoe-Lian Goh, D. (2014). Understanding news sharing in social media: An explanation from the diffusion of innovations theory. *Online Information Review*, 38(5), 598–615.




- Mourão, R. R., & Robertson, C. T. (2019). Fake news as discursive integration: An analysis of sites that publish false, misleading, hyperpartisan and sensational information. *Journalism Studies*, 20(14), 2077–2095.
- Oeldorf-Hirsch, A., & Sundar, S. S. (2015). Posting, commenting, and tagging: Effects of sharing news stories on Facebook. *Computers in Human Behavior*, 44, 240–249.
- Pope Francis shocks world, endorses Donald Trump for president, releases statement. (2016, July). *WTOE 5 News*. <https://wtoe5news.com>
- Rudat, A., Buder, J., & Hesse, F. W. (2014). Audience design in Twitter: Retweeting behavior between informational value and followers' interests. *Computers in Human Behavior*, 35, 132–139.
- Shoemaker, P. J., & Vos, T. P. (2009). *Gatekeeping theory*. New York, NY: Routledge.
- Singer, J. B. (2014). User-generated visibility: Secondary gatekeeping in a shared media space. *New Media & Society*, 16, 55–73.
- Talwar, S., Dhir, A., Kaur, P., Zafar, N., & Alrasheedy, M. (2019). Why do people share fake news? Associations between the dark side of social media use and fake news sharing behavior. *Journal of Retailing and Consumer Services*, 51, 72–82.
- Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, 29, 24–54.
- Tschiatschek, S., Singla, A., Gomez Rodriguez, M., Merchant, A., & Krause, A. (2018). Fake news detection in social networks via crowd signals. In P.-A. Champin, F. Gandon, & L. Médini (Eds.), *Companion proceedings of The Web Conference* (pp. 517–524). New York, NY: Association for Computing Machinery.
- van Heekeren, M. (2020). The curative effect of social media on fake news: A historical re-evaluation. *Journalism Studies*, 21(3), 306–318.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.
- Weeks, B. E., & Holbert, R. L. (2013). Predicting dissemination of news content in social media: A focus on reception, friending, and partisanship. *Journalism & Mass Communication Quarterly*, 90, 212–232.
- Winter, S., Bruckner, C., & Krämer, N. C. (2015). They came, they liked, they commented: Social influence on Facebook news channels. *Cyberpsychology, Behavior, and Social Networking*, 18(8), 431–436.

#### About the Authors



**Miriam J. Metzger** is a Professor of Communication at the University of California, Santa Barbara. Her research lies at the intersection of media, information technology, and trust, centering on how information technology alters humans' understandings of information credibility, privacy, and how networked communication technologies affect behavior.



**Andrew J. Flanagin** is a Professor in the Department of Communication at the University of California, Santa Barbara, where he is a former Director of the Center for Information Technology and Society. His work broadly considers processes of social influence in digitally-mediated environments, with emphases on the use of social media for information sharing and assessment; people's perceptions of the credibility of information gathered and presented online; and processes of collective organizing as influenced by the use of contemporary technologies.



**Paul Mena** teaches journalism and news writing at the University of California, Santa Barbara. He holds a PhD in Mass Communication from the University of Florida. His research interests include misinformation, credibility, fact-checking, journalism, and social media.



**Shan Jiang** is a PhD Student and Research Assistant in the Khoury College of Computer Sciences at Northeastern University. His research interests revolve around computational social science and natural language processing. His current work investigates how storytellers, platforms, and audiences interact in the online misinformation ecosystem. He has published award-winning papers and served on the program committees at WWW, CSCW, ICWSM, and AAAI.



**Christo Wilson** is an Associate Professor in the Khoury College of Computer Sciences at Northeastern University. He is a Founding Member of the Cybersecurity and Privacy Institute at Northeastern, and serves as Director of the Bachelors in Cybersecurity program. Professor Wilson’s research focuses on online security and privacy, with a specific interest in using ‘algorithm audits’ to increase transparency and accountability of black-box online platforms.