Understanding the Sorting Algorithm: Emotion Contagion and Comment Ranking on a Politically Polarizing News Article

Jinzhi Zhou Susan C. Herring Indiana University, Bloomington

Abstract

Online news platforms tend to sort and rank comments in different ways, other than chronologically displaying them. Previous research has found that bias in ranking algorithms can promote political bias and contribute to ideological polarization. We compared chronologically sequenced versus algorithmically ranked comments on a controversial *Foxnews.com* article using Linguistic Inquiry and Word Count software and content analysis. Findings reveal that the ranking algorithm promotes comments with a positive emotional tone and discourages negative comments, suggesting that the algorithm is partially neutralizing the ideological bias of the *Foxnews.com* platform. Ranking was also affected by comment length and upvotes and downvotes.

Introduction

Virtual interaction during the COVID-19 pandemic became particularly common since people were stuck at home. Pandemic-related news especially attracted attention, and online news sites saw an increase in reader-posted comments (Eisele et al., 2022). Conflicting political ideologies were often on display in these comments, with readers on the Left critical of politicians and their followers on the Right who downplayed the seriousness of the disease, while readers on the Right derided those on the Left as weak "snowflakes" whose advocacy of safety measures such as sheltering in place masked an excuse not to work. A vivid example of this conflict played out on January 25, 2021, at the height of the pandemic, when Foxnews.com, a right-leaning news outlet in the United States, posted a news article that reported on a controversial social and political issue related to the transmission of the coronavirus: tensions between school districts and Teachers Unions as regards when schools should resume inperson classes. News stories that feature controversial or negative events have been found to attract more attention and to motivate repeated commenting (Weber, 2014). The Foxnews.com article, with the headline, "Thousands of Chicago teachers not heading back to classrooms following union vote, will remain remote," received almost 3,000 comments in a single day. News articles involving controversial topics not only generate more comments, but also "more hostility in those discussions" (Ksiazek, 2018, p. 666). Many of the comments on the Foxnews.com news article expressed strong negative emotions, especially hostility directed towards teachers and the teachers' union, which are associated with the political Left in the United States.

Editors and platforms are naturally concerned about the detrimental effects of negative comments, especially those that are abusive or hateful. Some online news websites, such as NPR, Reuters, and CNN, have closed their comment sections, with some redirecting their audiences to engage through the organizations' social media accounts (Finley, 2015; Reimer et al., 2023). In contrast, Foxnews.com seems to encourage users to leave comments and interact with other users on its platform. The news comment section offers features like "upvotes" and "downvotes" and indicates the number of replies a comment receives. Further, like many platforms, Foxnews.com allows users to sort comments either chronologically or by "best" comments. Online news platforms commonly rank comments as a strategy to increase user engagement in online discussions (Park et al., 2020). However, ranking mechanisms necessarily use selective criteria, which can result in bias. Such bias can affect the tone of subsequent comments, potentially exacerbating political polarization (Shmargad & Klar, 2020). It is not clear according to what criteria Foxnews.com identifies and ranks "best" comments, since the algorithm is proprietary, and the platform does not explain the criteria it uses.¹ Does the sorting algorithm take emotional expression into consideration? For example, does it promote positive comments in order to encourage a more positive, civil tone in the discourse (cf. Goldenberg & Gross, 2020)? Or does the algorithm preferentially promote negative comments in order to attract more readers and commenters (cf. Weber, 2014)?

Researchers who have analyzed the emotional valence of comments and its effect on subsequent discourse have reported mixed findings. Goldenberg and Gross (2020) claim that the majority of social media comments are positive, consistent with a human preference to feel and express positive emotions; moreover, positive comments are more often liked and shared than negative comments on some platforms. However, comments on other platforms, including Foxnews.com, tend to skew negative (Masulo Chen et al., 2019). Other research has found that negatively valenced comments promote more user engagement, and that negative emotions such as sadness and anger spread faster than positive emotions on social media (Fan et al., 2017; Kwon & Gruzd, 2017). A negativity bias has also been observed in reactions to news stories, based on the human tendency to react more strongly to negative than positive information (Soroka et al., 2019). However, although Goldenberg and Gross (2020) note that digital platforms encourage emotion contagion in user comments, for example by selectively presenting certain kinds of news stories or by having a 'like' button, none of these studies has considered how algorithmic ranking affects and is affected by the emotionality of online comment threads. This is especially important to understand when the comments are responding to news stories that are politically or ideologically polarizing (cf. Shmargad & Klar, 2020), given the potential for polarizing emotions to spread.

In this article, we contribute to addressing this gap by analyzing the comment thread on the aforementioned *Foxnews.com* article as a case study. We manually extracted and coded a large number of comments and analyzed them in two ways: in chronological sequence, looking for evidence of emotion contagion (Goldenberg & Gross, 2020), and ranked by "best" comments, in an attempt to identify the algorithm's ranking criteria. Specifically, we used Linguistic Inquiry and Word Count (LIWC) software, supplemented by manual content analysis, to

analyze how the emotions of later comments compared with those of earlier comments and how those of top-ranked comments compared with those of lower-ranked comments, as well as to analyze factors such as comment length and upvotes and downvotes, which previous research has suggested play a role in comment ranking (Diakopoulos, 2015b; Shmargad & Klar, 2020).

We found that although the unranked comments showed no significant change in emotionality over time – most were negative – the comparison of chronologically sequenced versus algorithmically ranked comments suggests factors that determine comment ranking on *Foxnews.com*. Although the news story attracted ideologically polarized comments, especially attacks on teachers, teachers' unions, and Democrats from readers on the political Right, the ranking algorithm promoted comments with a positive emotional tone and discouraged negative comments, suggesting that the algorithm is partially neutralizing, rather than exacerbating, the ideological bias of the *Foxnews.com* platform. Longer comments were also ranked more highly than shorter comments, consistent with the goal of fostering more considered argumentation. The study demonstrates the utility of using linguistic analysis methods to understand what kinds of comments the sorting algorithm privileges and to what extent emotionally valenced comments are contagious on the news site, as well as shedding light on the relationship between these factors and ideological polarization in comments on *Foxnews.com*.

Relevant Literature

Emotionality

Emotionality is a central concept in the present study, because emotionality plays an important role in news engagement and online news dynamics (Choi et al., 2021; Eisele et al., 2022). Following Eisele et al. (2022), we define emotionality as the expression of emotions that are "the demonstration of a feeling [...] discursively manifested in emotional expressions in the comments, considering both negative and positive emotions" (p. 4). Positive emotions, signaling optimal well-being, include joy, interest, contentment, and love; negative emotions include anxiety, sadness, anger, fear and the like (Fredrickson, 2004).

Studies have investigated how news content affects the emotionality of readers' comments. Bösch et al. (2018) found that the emotions from news articles positively affect emotions in readers' comments, based on analysis of a large number of online news articles and readers' comments on them. However, in an experimental study, Petit et al. (2021) found that users were prone to demonstrate negative emotions (such as flaming) when the opinion expressed from the controversial topic in the news article was opposed to their own standpoint. Eisele et al. (2022) investigated the dynamics of emotionality of user comments in response to news coverage in the context of COVID-19 in two Austrian newspapers over the first half year of 2020. The researchers found increased positive and negative emotions in both the articles and

the comments during the lockdown period. They also found that news content involving political decision-makers and their images provoked emotionality in comments.

Other research has analyzed the influence of emotionality in news articles and news comments on when and why online users post comments. In their study of four major elite newspapers' Facebook pages, Choi et al. (2021) found that users were less likely to comment on news articles that conveyed positive emotions, while the negative emotions of "sadness" triggered more reader engagement. Relatedly, Ziegele et al. (2018) found that news articles that included controversy and damage caused by the news events on individuals or institutions increased the willingness of online participants to write comments, and also led participants to leave more uncivil comments. Overall, increased exposure to emotions, whether positive or negative, has been found to lead to increased engagement in online media platforms (Goldenberg & Gross, 2020). However, there are limits to this effect: The incivility of some news comment sections discourages readers from commenting (e.g., Engelke, 2019).

The study by Ziegele et al. (2018) also found that uncivil and off-topic comments led participants to leave more uncivil comments. In other words, "uncivil, aggressive, and off-topic comments" had a tendency to "heat up the subsequent debate" (p. 14). This is an example of emotion contagion, whereby people's emotional expression becomes more similar to the emotional expression of others (Goldenberg & Gross, 2020). Specifically, it is an example of an emotional cascade, where "exposure to emotions elicits similar emotions in perceivers, who then express their emotion by either replying or further sharing the content" (Goldenberg & Gross, 2020, p. 323). While there is considerable evidence that emotional contagion occurs in comment threads, there is no consensus on what type of emotion is most contagious. Goldenberg and Gross (2020) found that positive tweets got more likes and retweets than negative tweets and less emotional tweets on Twitter. Conversely, in an experimental study, Masullo Chen and Lu (2017) found that both civil and uncivil disagreement caused negative emotion, and uncivil disagreement led people to respond back uncivilly. Kwon and Gruzd (2017) also found offensive language to be contagious in comments on Donald Trump's campaign videos on YouTube.

In summary, previous research has focused on the relationship between the emotionality of news articles and that of users' comments, as well as what type of emotions attract more user engagement and are more likely to result in emotional contagion (for more extensive reviews, see Goldenberg & Gross, 2020 and Reimer et al., 2023). Somewhat paradoxically, the *Foxnews.com* article at the center of this study attracted very negative emotion expression and very active reader engagement, despite being neutral on its face. Moreover, the intensity of the negative emotions expressed in the comments makes it a good candidate for emotional contagion, although it is unclear as yet to what extent this takes place.

Interestingly, Reimer et al. (2023) found that "irony, sarcasm, cynicism" were the most frequently studied forms of (presumably negative) emotion in news comments. Such comments have high entertainment value for readers, although they can adversely affect the

credibility of the commenter and the news platform (Ziegele & Jost, 2020). These non-bona fide forms of expression are also difficult to detect using automated methods (e.g., Muresan et al., 2016; Thelwall et al., 2012), as we discuss further below.

Comment Ranking

Online news platforms tend to sort and rank comments in different ways, other than chronologically displaying them. They may sort by "best comments," "most relevant comments," number of likes/upvotes, or number of replies. To accomplish this ranking, either editors manually select recommended comments or platforms automatically sort using ranking algorithms which are based on certain criteria.

Through reviewing the literature, Diakopoulos (2015b) identified 12 editorial criteria (e.g., argument quality, emotionality, personal experience, and brevity) for identifying high quality user comments. He analyzed how these criteria were demonstrated in the *New York Times* "Picks" comments that were manually selected by editors, and the study confirmed that some criteria (e.g., argument quality, personal experience, readability) did manifest in those comments. The study found weaker evidence in support of emotionality as a criterion in the selected comments. Interestingly, in contrast to previous studies (e.g., Wahl-Jorgensen, 2002), Diakopoulos (2015b) found that brevity was not a criterion for editors; instead, longer comments were preferred by "Picks" editors.

In contrast to manual ranking, algorithms are designed to be "automatic," without any regular human intervention (Gillespie, 2014). While editorial criteria for manually selected comments tend to be explicit, it is unclear what criteria are used for ranking algorithms on news platforms; their underlying criteria are hidden (Gillespie, 2014). This is a problem, because "these criteria embed a set of choices and value propositions, which may be political or otherwise biased, that determine what gets pushed to the top" (Diakopoulos, 2015a, p. 41). Yet, it is unclear "how these criteria are measured, how they are weighed against one another, what other criteria have also been incorporated, and when, if ever, these criteria will be overridden" (Gillespie, 2014, p. 176). Burrell (2016, p. 1) cites "intentional corporate or state secrecy" and technical reasons as contributing to the lack of transparency in a particular classification decision. In part due to the lack of transparency in how they work, ranking algorithms are understudied (Shmargad & Klar, 2020).

Moreover, algorithms vary significantly across social media platforms, which generally do not make their algorithms public, either. An exception is Reddit, which made their algorithms for ranking user comments explicit (Shmargad & Klar, 2020). Reddit determines the ranking of news based on popularity, i.e., number of likes, shares, and comments (Shmargad & Klar, 2020). Research has shown that sorting and ranking comments by number of likes and upvotes drives users to make more comments and become more engaged. Organizing and sorting comments based on emotionality can also be helpful for structuring the commenting experience and meeting readers' expectations (Diakopoulos & Naaman, 2011). Digital media

platforms may attempt to maximize users' emotions through algorithms that particularly promote comments with positive emotions (Goldenberg & Gross, 2020). Goldenberg and Gross (2020) found that tweets with emotional expression, particularly positive emotion, predict that users receive more likes and retweets on Twitter. However, likes and upvotes could also lend themselves to malicious manipulation (Risch & Krestel, 2020; Park et al., 2020). Gillespie (2014) points out that Reddit must "constantly seek out and correct instances of organized downvoting, and these tactics cannot be made public" (p. 176). Moreover, ranking news articles by their popularity has been found to impact people's attitudes towards politics, potentially contributing to ideological polarization (Shmargad & Klar, 2020).

Taken together, these findings suggest that there could be multiple potential criteria deciding ranking algorithms, including number of words, number of upvotes and likes, and emotionality, and that these criteria could have differing effects on the discourse of a platform.

Research Questions and Hypotheses

In this study, we address three main research questions. We articulate and justify our questions and hypotheses below.

- **RQ 1:** Is there a relationship between the emotional quality and the chronological order of comments on the *Foxnews.com* article?
 - RQ 1a: Does the positivity of comments on the article change significantly over time?
 - RQ 1b: Does the negativity of comments on the article change significantly over time?

Based on the literature cited in the previous section, we hypothesize the following:

- H1a: The positivity of comments on the article will decrease significantly over time.
- H1b: The negativity of comments on the article will increase significantly over time.

Our assumption is that if emotion contagion occurs, there will be increasing emotionality in the comments over time, and that this is more likely to occur with negative emotion, as found, for example, by Ziegele et al. (2018). That is, negative comments, when and if they occur early in the comment thread, should cause subsequent comments to be more negative.

RQ 2: Is there a relationship between the emotional quality and the ranking of best comments on the *Foxnews.com* article?

- RQ 2a: Is there a relationship between positivity and the ranking of best comments?
- RQ 2b: Is there a relationship between negativity and the ranking of best comments?

Given editors' and platforms' concerns about the detrimental effects of negative comments, we hypothesize that the *Foxnews.com* ranking algorithm promotes positive comments over

negative comments, extrapolating from the findings of Goldenberg and Gross (2020) for social media platforms such as Twitter. Thus:

H2a: Higher-ranked comments will be more positive.

H2b: Higher-ranked comments will be less negative.

RQ 3: What other factors (comment length, upvotes/downvotes) impact the ranking of best comments on the *Foxnews.com* article?

Diakopoulos (2015b) posits that an editorial shift towards a preference for longer comments is taking place, due in part to the existence of fewer production constraints online than in traditional media. Since the comments we are studying are online, we hypothesize that the *Foxnews.com* ranking algorithm will also favor longer comments. Thus:

H3a: Higher-ranked comments will be longer than lower-ranked comments.

Previous studies have found that upvoted comments tend to be highly ranked (e.g., Park et al., 2020; Shmargad & Klar, 2020). Conversely, downvoted comments should be ranked lower. Thus, we hypothesize:

H3b: Higher-ranked comments will have more upvotes than lower-ranked comments.

H3c: Lower-ranked comments will have more downvotes than higher-ranked comments.

Methods

Theoretical and Analytical Framework

To address the above research questions, this study draws on computer-mediated discourse analysis (CMDA), a "methodological toolkit and a set of theoretical lenses through which to make observations and interpret the results of empirical analysis" of online language (Herring, 2004, p. 4). CMDA aims to identify patterns in language use that may not be evident to the casual observer or to the discourse participants themselves – in this case, commenters on the *Foxnews.com* article. Such patterns manifest within individual messages as well as across multiple messages, including in sequences of messages posted over time, making CMDA well-suited for analyzing patterns in online news comment threads.

CMDA methods can be applied to four levels of language: structure, meaning, interaction, and social behavior (Herring, 2004). Our analysis of online news comments involves both structure and meaning, and it employs two kinds of CMDA methods. One is "language-focused content analysis" (Herring, 2004, p. 4), which has been used to study meaning and social phenomena such as politeness (Kim & Herring, 2018; Wardoyo, 2019). Relevant to the present study, Wardoyo (2019) found that violations of positive politeness, especially sarcasm, were contagious in comments on three popular YouTube videos, in the sense of encouraging more

replies. In contrast, Kim and Herring (2018) did not find any effect of positive politeness violations, including sarcasm, on reply frequency in comments on a Korean news site.

The CMDA toolkit also includes automated and semi-automated methods of analysis, which are especially helpful for analyzing patterns at the level of structure, such as message length and word frequencies. We use LIWC for this purpose. Structure and meaning are conflated in LIWC: Emotion expression, for example, involves meaning, but the LIWC dictionary program categorizes emotional terms into structurally identifiable units that can be counted automatically.

Use of LIWC to Detect Emotion

LIWC (Pennebaker et al., 2001) is a dictionary-based text analysis software program that counts the frequency of words "that reflect different emotions, thinking styles, social concerns, and even parts of speech" (liwc.wpengine.com), hence capturing people's social and psychological states. The LIWC2015 version includes around 90 linguistic categories that refer to collections of words, e.g., articles, prepositions, and pronouns, as well as more subjective categories, such as positive and negative emotion words, which were selected and evaluated by human judges (Tausczik & Pennebaker, 2010). Positive emotion words like "love," "nice," and "sweet," and negative emotion words like "hurt," "ugly," and "nasty" provide psychological cues to people's emotional states and intentions (Tausczik & Pennebaker, 2010). In addition, LIWC2015 calculates a summary variable called emotional tone. The LIWC scores for positive and negative emotions represent percentages out of the total number of words in the sample. In contrast, the score for emotional tone is calculated on a 100-point scale; a high score reveals a more positive form of discourse, whereas a low score indicates "greater anxiety, sadness, or hostility" (LIWC2015 Operator's Manual).

Numerous studies have used LIWC to investigate people's behavior and internal states, particularly emotions, through measuring their use of different language categories. For example, Diakopoulos (2015b) utilized a set of LIWC categories to calculate the Personal Experience score for comments in the *New York Times* (NYT). Comments that were picked by NYT editors had a higher average score than those not picked, implying that personal experience was one of the criteria NYT editors used when selecting good comments. Bösch et al. (2018) used the German version of LIWC to measure sentiment scores for online newspaper articles and readers' comments. Zheng et al. (2022) used LIWC to assess the effects of the emotional valence of tweets on information sharing related to COVID topics on Twitter. Kahn et al. (2007) conducted three experiments to assess whether LIWC emotion counts were sensitive to verbal expression of amusement (positive emotion) and sadness (negative emotion); their results corroborated the validity of LIWC for measuring emotion expression.

LIWC has also been used in CMDA studies, as noted above. Kapidzic and Herring (2011) applied it to the analysis of gender differences in comments in teen chatrooms. Kleanthous and Otterbacher (2019) used LIWC to analyze reactions to TED talks about robotics, focusing on

comment length, emotional tone, authenticity, analytical thinking, and clout. Zhu and Kadirova (2022) used the software to analyze social and cognitive presence in students' comments on YouTube videos. In the present study, we use LIWC2015 to examine linguistic features of comments, including word count and three categories related to emotionality: positive emotion, negative emotion, and emotional tone. We supplement automated LIWC analysis with manual content analysis to analyze the emotional valence and the presence of sarcasm in selected comments on the *Foxnews.com* article over time.

Data

Data Source

Our data for this case study are comments that responded directly to the Foxnews.com article. The news story reported that the Chicago Teachers Union voted against in-person instruction, although the district of Chicago Public Schools wanted K-8 teachers and staff to return to school. This was not a simple situation. On the one hand, the district expected teachers to return to school because they wanted to provide the same option to students as those in private and parochial schools, where students had been learning safely in classrooms. The district also expressed concerns about the drop in grades, attendance, and enrollment, especially among Black and Latinx students. On the other hand, the Union said that while teachers wanted to return to in-person instruction, they were worried about the spread of the coronavirus, since the district had not prepared for a return, and hoped that Union members could get vaccines and other safeguards before returning. Overall, the article itself adopted a fairly even-handed tone when reporting on the tensions between the district and the Union. However, its publication on Foxnews.com, a politically right-leaning media platform, predisposed commenters to adopt critical positions towards the Teachers Union and the teachers who were concerned about the coronavirus, both of which were associated with the political Left in the United States.

The news story attracted a total of 2806 public comments, all in English, virtually all of which were posted on January 25, 2021 (except for three posted on January 27). The comments addressed a variety of topics, including calling for firing the teachers, criticism of the Teachers Union, dissatisfaction with current pedagogy in public schools and the ineffectiveness of remote teaching, and transmission of the coronavirus.

Data Collection

Comments on the *Foxnews.com* website are threaded. Under each comment, the numbers of upvotes and downvotes, as well as replies, are indicated. Only 10 comments are displayed on the first page, and users must unfold sub-level comments and load "more comments" to continue reading. There are three ways of sorting comments, by oldest, newest, and best. The default setting is sorting by best.

CMDA data sampling techniques include sampling by time (chronological sequence of posting) and according to a judgment criterion (e.g., "top posts") (Herring, 2004). Out of the

2806 publicly accessible comments, in April 2021 we collected two datasets: the first 1000 first-level comments sorted chronologically, starting with the "oldest," and the first 1000 first-level comments sorted by rank, starting with the "best." We limited our data collection to first-level comments to focus on direct reactions to the news story, since replies to comments are more likely to digress from the topic of the story (Herring & Chae, 2021). We manually downloaded² all comments in PDF format and sorted by "oldest" and "best," resulting in more than 400 pages for each dataset. We then limited our collection to 1000 comments for each set due to the time-consuming nature of this process. For each dataset, we organized the comments and their associated numbers of replies, upvotes, and downvotes in spreadsheets.

Analytical Methods

Our analysis comprised an automated (quantitative) and a manual (qualitative) component. The quantitative analysis consisted of five steps. 1) Given the brevity of many online news comments,³ in order to create units large enough for meaningful LIWC analysis, we first divided the 1000 oldest comments evenly into 10 segments according to chronological order, such that each segment consists of 100 comments.⁴ Similarly, we divided the 1000 best comments evenly into 10 segments according to their rank. Thus, the segments for the oldest comments are time segments, and the segments for the best comments are ranking segments. For both datasets, our unit of analysis was the 100-comment segment. 2) Next, we ran the text of every segment through the LIWC2015 academic version, generating average results for word count, positive emotions, negative emotions, and emotional tone. 3) We compared the average results from our data with related genres that are offered in LIWC, including "news article," "social media," and "scientific writing." 4) We generated charts comparing our LIWC results and the results for upvotes and downvotes for the oldest and the best comments. We also examined the position of oldest comments in best comments and how many oldest comments were selected as best comments in order to determine the extent to which the two measures (oldest and best) are independent of one another. 5) We conducted Pearson correlation analysis among the variables of emotionality, word count, ranking order, upvotes, and downvotes. The first author collected the data and ran the LIWC analyses, and both authors interpreted the results.

We then conducted a content analysis of selected oldest and best comments that showed especially high or low scores in positive emotion, since the positive emotion results showed greater variance than the negative emotion results. Specifically, we manually examined the 200 oldest comments in the 5th and 7th segments and the 200 best comments in the 5th and 6th segments to explore what happened in those comments that might account for the unusually high or low positive emotion scores assigned by LIWC2015 to these segments. In this process, we coded for positive, negative, neutral, and sarcastic comments based on the categories used by Bourlai and Herring (2014) in analyzing multimodal Tumbler content. Emotional valence was coded as positive, negative, or neutral. Sarcasm (non-bona fide communication) was coded for presence or absence. Sarcasm was coded because of its frequent presence in the comments, and because it could be misinterpreted by LIWC2015 as expressing one emotion when the opposite emotion was actually intended (Muresan et al., 2016; Tausczik & Pennebaker, 2010).⁵

Both authors jointly coded the 400 selected comments, and cases of potential disagreements were discussed until consensus was reached.



Figure 1 summarizes the steps followed in analyzing the data.

Figure 1. Data analysis procedure

Results

LIWC Results

Comparison across Different Discourse Genres

Table 1 shows the results for word counts and emotional expression in our two datasets, the *Foxnews.com* article, and three comparison genres available through LIWC: news articles (in general), scientific articles, and social media. Several observations can be made based on these results. First, the score for negative emotions in the *Foxnews.com* story (1.13) is lower than that for the other genres, as is the score for positive emotions (1.51). This suggests that the language of this news story is not very emotional, in support of our impression that the article adopts a relatively neutral stance toward the issues it reports.⁶ According to these measures, the article is even less emotional than scientific writing, which is known for its impersonal, objective style.

Table 1 also shows that there are differences in emotional expression between the *Foxnews.com* article and comments on the article. The comments are more emotional than the article; they are both more negative than the article (and news articles in general) and more positive than the *Foxnews.com* article.

Finally, Table 1 shows overall differences between our two comment samples. The first 1000 oldest comments are shorter than the top 1000 best comments, and they are more negative (2.24 vs. 1.99), including having a lower emotional tone (28.99 vs. 33.48). Indeed, the oldest comments are more negative according to these measures than any of the other genres in Table

Discourse genre	Word count	Emotional tone	Positive emotions	Negative emotions
Comments - oldest	29997	28.99	2.43	2.24
Comments - best	36270	33.48	2.43	1.99
Fox News story	794	N/A^7	1.51	1.13
News articles	N/A	53.5	3.1	1.6
Social media	N/A	63.35	4.57	2.1
Scientific writing	N/A ⁸	43.61	2.32	1.45

1. However, there is no overall difference between the two comment samples in the expression of positive emotions.

Table 1. LIWC results for word count and emotionality in different discourse genres

Trends over Time for Oldest Comments

Figure 2a depicts the variations in positive emotion words over time in the 1000 oldest comments. The figure suggests that positive emotion words fluctuate and that there is no clear trend of positivity over time. Similarly, Figure 2b shows that the scores for negative emotion fluctuate in the oldest comments, and that there is no clear trend for negativity over time.





Trends over Time for Best Comments

As we move from the highest-ranked to the lowest-ranked comments in the best comments, we see a downward trend for positive emotion, with some variations, as shown in Figure 3a. That is, the highest-ranked comments are the most positive. Figure 3b shows a corresponding upward trend for negative emotion in the best comments dataset, indicating a strong, linear relationship between negative emotion words and message ranking. Moreover, there is less variation in the trendline for negative emotion than in the trendline for positive emotion.



Figure 3a & 3b: Positive emotion (left) and negative emotion (right) in best comments

Emotional Tone

Figure 4a & 4b display the scores for emotional tone for the oldest and the best comments. Emotional tone fluctuates over time for the oldest comments, as shown in Figure 4a, with no overall clear trend. In contrast, Figure 4b shows a clear downward trend for the emotional tone of the best comments, indicating that the highest-ranked comments have a higher emotional tone. These results mirror those for positive emotion in Figure 3a and demonstrate an even stronger trend.





Other Factors

Our analyses considered three other possible explanatory factors underlying ranking of best comments: comment length, upvotes, and downvotes.

Comment length. Comment length, operationalized as number of words per segment, changed significantly both over time and as rank decreased, but in opposite directions. Figure 5a shows that earlier comments were shorter than later ones, whereas Figure 5b shows that the highest-ranked comments were longer than lower-ranked comments.



Figure 5a & 5b. Length of oldest comments (left) and best comments (right)

Upvotes and downvotes. Figures 6a through 6d show the trends of upvotes and downvotes of oldest and best comments. Average frequencies of upvotes and downvotes were calculated and plotted for each time segment. The significant downward linear trend of upvotes for the oldest comments indicates that the older the comment, the more upvotes it received (Figure 6a). In contrast, Figure 6b shows a long-tailed distribution of upvotes of best comments, indicating that the first 100 best comments got the most upvotes.

The trendline of downvotes shows similar patterns. The older comments received significantly more downvotes, and this relationship is also linear (Figure 6c). The best 100 comments received the most downvotes, which is shown in the long-tailed distribution in Figure 6d. These results show that comment age and comment rank overlap to some extent for upvotes and downvotes.



Figure 6a & 6b. Upvotes and downvotes of oldest comment (left) and best comments (right)



Figure 6c & 6d. Downvotes of oldest comment (left) and of best comments (right)

Relationship of age to rank. Figure 7a and 7b show that there is no significant overall relationship between the age and the rank of the comments. The best comments are the oldest, but only for the first 100 top-ranked comments (7b). Conversely, the oldest comments are not the best; rather, the comments in the 7th and 8th time segments were ranked the best (7a). This is evidence that the age and the rank of a comment beyond the 100 top-ranked comments are independent of one another. Statistical analyses supporting this independence are summarized in Table 2.





Summary of Factors Influencing Comment Ranking

Table 2 summarizes the factors that were found to influence comment ranking. Low values for best comments indicate higher ranking. Best comments are strongly positively correlated (.794**) with negative emotion, indicating that comments with more negative emotions were ranked lower. Best comments are negatively correlated with tone (-.755*) and word count (-.783**), indicating that comments with higher emotional tone and longer comments were ranked higher. Older comments are highly negatively correlated with (receive more) upvotes (-.915**) and downvotes (-.929**).

	Best comments	Positive emotion	Negative emotion	Tone	Word count	Upvotes	Downvotes	Position in oldest comments
Best	1	485	.794**	755*	783**	572	475	.349
Positive	485	1	481	.870**	.381	.272	.245	180
Negative	.794**	481	1	848**	688*	586	503	.540
Tone	755*	.870**	848**	1	.622	.503	.440	409
Word count	783**	.381	688*	.622	1	.774**	.731*	619
Upvotes	572	.272	586	.503	.774**	1	.992**	915**
Downvotes	475	.245	503	.440	.731*	.992**	1	929**
Position in oldest comments	.349	180	.540	409	619	915**	929**	1

Table 2. Correlations for best comments and oldest comments with multiple variables

* Correlation is significant at the 0.05 level (2-tailed)

** Correlation is significant at the 0.01 level (2-tailed)

Content Analysis

Figures 2 and 3 above show segment-by-segment variation between high and low scores, especially in positive emotion. In the oldest comments dataset, comments in the 6th segment had the highest (rising) score for positive emotion and also a relatively high score for negative emotion, whereas the 8th segment had the lowest (falling) score for positive emotion. Similarly, in the best comments dataset, the 6th segment had the highest (rising) score for positive emotion and a relatively high score for negative emotion. To explore what might have caused these fluctuations, we manually examined the 400 comments from these four segments by coding them for emotion – positive, negative, or neutral – and sarcasm. Sarcasm was considered independently of the other emotion codes because it could, in principle, co-occur with any of them, although it was most often negative in connotation. The content analysis results are presented in Table 3.

Segment from Figures 2a & 3a	Positive	Negative	Neutral	Total	Sarcasm
Oldest 6th (rising)	1	79	20	100	14
Oldest 8th (falling)	3	76	21	100	7
Best 6th (rising)	2	67	31	100	10
Best 7th (falling)	0	84	16	100	11

Table 3. Results of content analysis of comment emotion

Table 3 shows that the comments were overwhelmingly coded as negative, followed by neutral. There are very few positive comments in the four-segment sample. Table 3 also shows different patterns for the oldest and the best comments. The pattern for the oldest comments is that rising and falling segments are differentiated by sarcasm; rising comments included more sarcasm, which LIWC might have mistakenly classified as positive. A different pattern is evident for the best comments: Rising and falling segments are differentiated by neutral and negative comments. The rising segment has fewer negative comments and more neutral comments compared to the falling segment.

Examples of each category are provided below. The source segment is indicated in parentheses.

First, comments expressing *positive* emotion were rare in the comment thread, and their positivity is relatively weak. This is evident in examples 1-3, which lack strongly positive words.

- (1) "Add teachers to the priority of vaccination." (Oldest 6th)
- (2) "Maybe 'The Census Cowboy' can help out here." (Oldest 8th)
- (3) "In my kids' rural district, one principal died of Covid and many teachers have gotten sick, including both of one of [my] kid's classroom teachers. We all want inclass teaching, but we have to do our part to protect the teachers." (Best 6th)

In a thread that is highly critical of teachers overall, comment 1 is supportive of the teachers' goal to obtain the then-new Covid-19 vaccination. Comment 2 makes a constructive suggestion that acknowledges the expertise of one of the previous commenters. In addition to being supportive of teachers, comment 3 seeks consensus by acknowledging the position of those taking the opposite side in the argument.

Comments expressing *negative* emotion make up the majority of our data. Negativity is most often directed toward the teachers who are reluctant to return to the classroom during the pandemic. These comments include strong, unambiguously negative words such as *useless* (example 4), *ruined* (example 5), and *WORST, uncaring, selfish, greedy, cesspool,* and *goons* (example 6).

- (4) "What's the use? Teachers are useless." (Oldest 8th)
- (5) "Fire every one of these liberal educators. They've ruined the mind's [sic] of our youth and the fabric of society." (Oldest 8th)
- (6) "Chicago teachers are the absolute WORST examples of 'educators' I have ever had the misfortune to see. Most uncaring, selfish and greedy lot...No wonder the place is such a cesspool- their children are brought up by these goons." (Best 7th)

Examples 7-9 illustrate comments that we coded as emotionally *neutral*. Although the ideological position of the commenters can be inferred from their content, the comments are neutrally worded and leave room for rational debate, unlike examples 4-6.

- (7) "Been teaching in the classroom since September...some in class some at home. The virus spread in schools is very low...get the kids back, we are losing them." (Best 6th)
- (8) "There's a reason children in India, Japan hell, even China are smarter than our kids." (Oldest 8th)
- (9) "Parents should be able to spend their tax dollars on whatever school they choose." (Best 7th)

Note that LIWC does not have a category for emotionally neutral words, and therefore the higher number of neutral comments in Table 3 cannot by itself explain the higher ranking of the 6th segment of the best comments by LIWC. The lower frequency of negative words in that segment could be a factor, though.

Finally, examples 10-12 illustrate comments that were coded as *sarcastic*. Their surface form is positive, including words like *LOVE* (example 10), *good*, *well-being* (example 11), and *thank you* (example 12). However, their intended meanings are negative – the commenters do not love Democrats, do not think that the teachers have the well-being of the children in mind, or feel appreciation toward the teachers' union; quite the contrary.

- (10) "Democrats! Don't ya just LOVE em!" (Oldest 6th)
- (11) "It's so good the teachers are thinking about the well-being of the children." (Oldest 6th)
- (12) "Thank you teachers union, now it's time for you to allow the schools to fire these teachers. It's clear they don't want to do their jobs. Remember Libs, we must follow the science, the science says that the schools are one of the safest places for no getting the virus!" (Best 6th)

The last sentence in example 12 is also sarcastic, but instead of conveying the opposite of what it says on the surface, it mocks liberal commenters by pretending to "remind" them of what is presumed to be a liberal trope.

Sarcastic comments were considerably more frequent than positive comments in our datasets, as the content analysis results in Table 3 suggest. It seems highly likely that some words classified by LIWC as positive were in fact intended sarcastically, and thus that the frequency of positive terms was artificially inflated. Table 3 suggests that this may especially be the case in the oldest comments, whereas frequency of sarcasm in the best comments did not appear to affect fluctuations in the LIWC results in the two segments analyzed in Table 3. Why this

should be so, and whether sarcasm is a factor that the ranking algorithm takes into account, is unclear at this time and should be explored with a larger data sample.

Discussion

Our first research question asked: *Is there a relationship between the emotional quality and the chronological order of comments on the Foxnews.com article*? We found no significant relationship over time between comment order and positive or negative emotion, nor was there a relationship between comment order and LIWC's emotional tone variable. These results are contrary to our hypotheses that positive emotion would decrease and negative emotion would increase over time – that is, that negative emotion would be contagious (e.g., Ziegele et al., 2018). The comments are consistently negative over time, as well as more negative overall than any of the other genres included in Table 1.

A possible explanation for this lack of variation is that the commenters are similarly minded in their stance toward the article; they might be responding to the article's content, rather than influenced by other commenters' emotional expressions (Goldenberg & Gross, 2020). Another possible explanation for the lack of emotional contagion in comments is the *Foxnews.com* platform itself, which displays a limited number of comments per page⁹ and requires that readers click on comments to unfold reply threads. Moreover, the platform displays best comments, rather than comments in chronological sequence, by default. As a consequence, readers might not see or read many previous messages before commenting.

Our second research question asked: *Is there a relationship between the emotional quality and the ranking of best comments*? Consistent with our hypothesis, negativity was significantly disfavored in best comments. Moreover, we identified a trend associating positive emotion with best comments, especially in the first five segments, although it did not reach overall significance. Thus, the emotionality of comments impacts comment ranking. The results for emotional tone further corroborate these findings, in that comments that have a higher emotional tone were ranked higher. These findings constitute evidence that some news platforms, like some social media platforms, algorithmically promote positivity and discourage negativity (Goldenberg & Gross, 2020).

Our last research question asked: *What other factors (comment length, upvotes/downvotes) impact the ranking of best comments?* As regards comment length, longer comments were ranked higher than shorter comments, consistent with our hypothesis and the findings of Diakopoulos (2015b) for *New York Times* articles but inconsistent with other previous studies (e.g., Wahl-Jorgensen, 2002). This finding could be due in part to the existence of fewer production constraints (including word limits) in online publications compared with offline publications, as Diakopoulos (2015b) suggests. Longer comments are also significantly less negative, as the correlation analysis in Table 2 shows.

As regards upvotes and downvotes, the picture is less clear. Upvotes and downvotes are strongly correlated: Comments that get more upvotes also get more downvotes. However, neither are correlated with best comments overall, although the top 100 best comments received the most upvotes and downvotes, in a long-tailed distribution. Only chronological comment order was strongly correlated with upvotes and downvotes, with the oldest comments receiving the most of each type of feedback. This is understandable, in that the longer a comment is publicly available on a platform, the more opportunity readers have to upvote or downvote it. This same reasoning could explain why the top 100 best comments received the most upvotes and downvotes, since the best comments are displayed by default, and only a few comments are displayed on a page. It may be that those best comments attract more upvotes and downvotes because they get more exposure through the platform's interface, rather than that they are ranked higher because they got more upvotes and downvotes. This would be an example of interface bias (Friedman & Nissenbaum, 1996), rather than algorithmic bias. Thus, our findings do not clearly support previous findings that upvoted comments are highly ranked (e.g., Park et al., 2020; Shmargad & Klar, 2020) and that popularity determines comment ranking.

Taken together, these findings shed considerable light on the workings of the *Foxnews.com* ranking algorithm. We identified significant linear correlations between top-ranked comments and (avoidance of) negativity and longer comments; this suggests that the algorithm takes these factors into account. At the same time, the ranking algorithm only weakly favored positivity, perhaps because the incidence of positive comments in our data was so low. The emotional tone results are stronger, showing that comments with higher (positive) emotional tone are ranked higher. Longer comments may be encouraged because they allow for more thoughtful expression of views and experiences and may therefore help promote more civil discussions. Finally, the algorithm does not appear to take upvotes or downvotes into account systematically, although this may be a good thing. Algorithms that consider factors beyond popularity may help counteract the efforts of malicious users who attempt to manipulate the ranking algorithm by strategically upvoting comments that align with their own opinions (Risch & Krestel, 2020).

Previous research has found that bias in ranking algorithms can promote political bias and contribute to ideological polarization (Shmargad & Klar, 2020). The *Foxnews.com* ranking algorithm promotes comments with a positive emotional tone and discourages negative comments. The more positive and neutral comments in our dataset are supportive of teachers and the teachers' union, unlike the dominant ideology of Fox News viewers and readers at that time. This suggests that, in this case at least, the ranking algorithm is partially neutralizing the ideological bias of the *Foxnews.com* platform. This may seem surprising, in view of the strongly ideological orientation of the Fox News media outlet (Jones, 2012). However, given the current amount of toxicity in online comments in general (Salminen et al., 2020), and the negativity of *Foxnews.com* comments in particular (Masulo Chen et al., 2019), mitigation of negativity may be considered desirable by the platform to avoid discouraging people from reading and commenting (Engelke, 2019).

Still, a number of characteristics of the algorithm and how it operates remain unclear. These include how different factors are weighted, when the initial ranking takes place, and how often it is revised. Moreover, how the algorithm treats sarcasm and comments that are emotionally neutral is unknown.

Conclusion

Algorithms increasingly influence our daily life through search engines, online news websites, and social media (Diakopoulos, 2015a; Gran et al., 2021). Understanding how ranking algorithms work is necessary to shed light on the biases they incorporate (Diakopoulos, 2015a) and how they influence the tone of online discourse. This study contributes novel insights into the mechanisms of comment ranking on the popular *Foxnews.com* platform. The ranking criteria support previous findings as regards emotionality and message length, while calling into question the degree to which measures of popularity such as upvotes and downvotes are considered by the algorithm. Further, the study contributes to the current literature on emotion contagion of comments on online news articles by analyzing a case where the emotionality of the commenters' ideological commitments. We suggested that the lack of emotion contagion in the chronologically sequence comments could be due to the behaviors of similarly minded commenters together with bias in the platform's interface.

The study's findings have implications for the design of ranking algorithms. Accuracy is important, especially during health crises such as the Covid-19 pandemic, when misinformation could be legitimized if it is highly ranked. Moreover, ranking algorithms need to be able to detect sarcasm. This is especially critical on news sites and political sites, where sarcastic comments proliferate and could cause algorithms to classify them as positive and thereby promote them. In our content analysis, sarcasm did not appear to affect comment ranking, but the sample size for the analysis was small. With a larger sample, content analysis could help improve supervised machine learning by providing manually labelled data to train algorithms to identify sarcasm and to evaluate their performance (Reimer et al., 2023). Human moderators and users who report misinformation could also work together with artificial intelligence to accomplish these goals, which are currently challenging for artificial intelligence alone (but cf. Muresan et al. [2016], who use lexical and pragmatic factors to recognize sarcasm from positive and negative emotions in Twitter posts).

An obvious limitation of any case study is its sample. Although we analyzed a large number of comments, they all came from one news article on *Foxnews.com*. The findings of our study may therefore not generalize to other articles or news websites. Another limitation is the inability of LIWC to recognize sarcasm from context, which may have influenced the absolute values for positivity and emotional tone, although this should not affect their relative values over time and in the top-ranked comments. Moreover, the effect of positivity in the ranking algorithm did not reach statistical significance; this may be due to the small number of positive

comments in the data overall. Future work could benefit from using computational methods to scrape and analyze comments on news articles with a larger number of genuinely positive comments on *Foxnews.com* and other news sites.

Finally, while ranking algorithms can impact online discourse, there are limits to their impactfulness. Although the *Foxnews.com* algorithm favors positive emotional expression, its rankings do not lead to increased positivity in the subsequent comment thread. Rather, the emotional tone of comments on *Foxnews.com* reflects the political polarization in US society at large. Thus, comment ranking alone is not sufficient to eliminate online polarization; people will still disagree with one another. Masullo Chen and Lu (2017) found that disagreement caused negative emotions and aggressive intentions. At the same time, Gil de Zúñiga et al. (2018) found that civil and reasoned disagreements can lead people to reconsider their political beliefs. The practice of displaying best comments as the default is a potentially useful step toward mitigating negative tendencies, especially when those comments are longer and more reasoned. However, the *Foxnews.com* case that we analyzed suggests that further measures are needed if the goal is to promote civil and productive discussion of controversial news content.

Notes

- 1. The first author reached out to the *Foxnews.com* digital team and asked what criteria the ranking algorithm uses to sort best comments. However, she received only a general response about the site's ability to turn comments on or off depending on the nature of the article.
- 2. Extracting data from news sites can be time- and labor-consuming, in part due to the nested structure of comments on such platforms, which precludes the use of automated methods of data collection.
- 3. Although some comments are short, all of them were deemed to be substantive; thus, all were included in the analysis. However, we omitted the longest comment. It was an outlier at 629 words; the next-longest comments had around 260 words.
- 4. A more natural grouping might have been produced by determining breakpoints based on when the comment was posted; however, this was not possible, since at the time we collected the comments, their timestamps indicated only the date of posting, and all but the last three comments were posted on the same day.
- 5. Tausczik and Pennebaker (2010, p. 30) acknowledge that LIWC "ignore[s] context, irony, sarcasm, and idioms ... like any computerized text analysis program."
- 6. On the surface, at least. The content of the article was likely to trigger *Foxnews.com* readers, and both the author of the article and *Foxnews.com* were no doubt aware of that.
- 7. By the time we conducted this part of the analysis, the version of LIWC available on the LIWC website had updated to LIWC-22, which no longer includes the Emotional tone variable or the category of Scientific Writing.

- 8. Neither LIWC2015 nor LIWC-22 provides word counts for the comparison genres.
- 9. The first page displays 10 comments; the second and subsequent pages display 25 comments each.

References

- Bourlai, E., & Herring, S. C. (2014). Multimodal communication on Tumblr: "I have so many feels!" *Proceedings of WebSci'14*, June 23–26, Bloomington, IN.
- Bösch, K., Müller, O., & Schneider, J. (2018). Emotional contagion through online newspapers. In the 26th European Conference on Information Systems (ECIS), Portsmouth, UK, 11-28.
- Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, *3*(1), 2053951715622512.
- Choi, J., Lee, S. Y., & Ji, S. W. (2021). Engagement in emotional news on social media: Intensity and type of emotions. *Journalism & Mass Communication Quarterly*, 98(4), 1017-1040.
- Diakopoulos, N. (2015a). Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism*, *3*(3), 398-415.
- Diakopoulos, N. (2015b). Picking the NYT picks: Editorial criteria and automation in the curation of online news comments. *International Symposium on Online Journalism*, 6(1), 147-166.
- Diakopoulos, N., & Naaman, M. (2011, March). Towards quality discourse in online news comments. In Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (pp. 133-142).
- Eisele, O., Litvyak, O., Brändle, V. K., Balluff, P., Fischeneder, A., Sotirakou, C., ... & Boomgaarden, H. G. (2022). An emotional rally: exploring commenters' responses to online news coverage of the COVID-19 crisis in Austria. *Digital Journalism*, 10(6), 952-975.
- Engelke, K. M. (2019). Enriching the conversation: Audience perspectives on the deliberative nature and potential of user comments for news media. *Digital Journalism*, *8*, 447-466.
- Fan, R., Xu, K., & Zhao, J. (2016). Higher contagion and weaker ties mean anger spreads faster than joy in social media. arXiv. Published online August 12, 2017. http://arxiv.org/abs/1608.03656
- Finley, K. (2015, October 8). A brief history of the end of the comments. *Wired*. https://www.wired.com/2015/10/brief-history-of-the-demise-of-the-comments-timeline/

- Fredrickson, B. L. (2004). The broaden–and–build theory of positive emotions. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 359(1449), 1367-1377.
- Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. ACM Transactions on Information Systems (TOIS), 14(3), 330-347.
- Gil de Zúñiga, H., Barnidge, M., & Diehl, T. (2018). Political persuasion on social media: A moderated moderation model of political discussion disagreement and civil reasoning. *The Information Society*, 34:5, 302-315.
 DOI: <u>10.1080/01972243.2018.1497743</u>
- Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 167-195). MIT Press.
- Goldenberg, A., & Gross, J. J. (2020). Digital emotion contagion. *Trends in Cognitive Sciences*, 24(4), 316-328.
- Herring, S. C. (2004). Computer-mediated discourse analysis: An approach to researching online behavior. In S. A. Barab, R. Kling, & J. H. Gray (Eds.), *Designing for virtual communities in the service of learning* (pp. 338-376). Cambridge University Press.
- Herring, S. C., & Chae, S. (2021). Prompt-rich CMC on YouTube: To what or to whom do comments respond? In *Proceedings of the Fifty-fourth Hawaii International Conference on System Sciences* (*HICSS-54*). <u>https://homes.luddy.indiana.edu/herring/HICSS.2021.herring.chae.pdf</u>
- Jones, J. P. (2012). Fox News and the performance of ideology. *Cinema Journal*, *51*(4), 178-185.
- Kahn, J. H., Tobin, R. M., Massey, A. E., & Anderson, J. A. (2007). Measuring emotional expression with the Linguistic Inquiry and Word Count. *The American Journal of Psychology*, 120(2), 263-286.
- Kapidzic, S., & Herring, S. C. (2011). Gender, communication, and self-presentation in teen chatrooms revisited: Have patterns changed? *Journal of Computer-Mediated Communication*, 17(1), 39-59.
- Kim, Y., & Herring, S. C. (2018, January). Is politeness catalytic and contagious? Effects on participation in online news discussions. In *Proceedings of the 51st Hawaii International Conference on System Sciences*. IEEE.
- Kleanthous, S., & Otterbacher, J. (2019, June). Shaping the reaction: Community characteristics and emotional tone of citizen responses to robotics videos at TED versus YouTube. In Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization (pp. 325-330).
- Ksiazek, T. B. (2018). Commenting on the news: Explaining the degree and quality of user comments on news websites. *Journalism Studies*, *19*(5), 650-673.

- Kwon, K. H., & Gruzd, A. (2017). Is offensive commenting contagious online? Examining public vs interpersonal swearing in response to Donald Trump's YouTube campaign videos. *Internet Research*, 27(4), 991–1010. https://doi.org/10.1108/IntR-02-2017-0072
- Linguistic Inquiry and Word Count (LIWC) 2015. http://liwc.wpengine.com.
- Masullo Chen, G., & Lu, S. (2017). Online political discourse: Exploring differences in effects of civil and uncivil disagreement in news website comments. *Journal of Broadcasting & Electronic Media*, 61(1), 108-125.
- Masullo Chen, G., Riedl, M. J., Shermak, J. L., Brown, J., & Tenenboim, O. (2019). Breakdown of democratic norms? Understanding the 2016 US presidential election through online comments. *Social Media*+ *Society*, 5(2), 2056305119843637.
- Muresan, S., Gonzalez-Ibanez, R., Ghosh, D., & Wacholder, N. (2016). Identification of nonliteral language in social media: A case study on sarcasm. *Journal of the Association for Information Science and Technology*, 67(11), 2725-2737.
- Norman, G. (2021, January 25). *Thousands of Chicago teachers not heading back to classrooms following union vote, will remain remote.* Foxnews.com. https://www.foxnews.com/us/thousands-of-chicago-teachers-are-not-heading-back-toclassrooms-today-following-union-vote
- Park, I., Shim, H., Kim, J. H., Lee, C., & Lee, D. (2020). The effects of popularity metrics in news comments on the formation of public opinion: Evidence from an internet portal site. *The Social Science Journal*, June, 1-16. DOI: 10.1080/03623319.2020.1768485
- Pennebaker, J. W., Francis, M. E., & Booth, R. J. (2001). *Linguistic Inquiry and Word Count* (*LIWC*): *LIWC 2001*. Erlbaum.
- Petit, J., Li, C., & Ali, K. (2021). Fewer people, more flames: How pre-existing beliefs and volume of negative comments impact online news readers' verbal aggression. *Telematics and Informatics*, 56 (January), 101471.
- Reimer, J., Häring, M., Loosen, W., Maalej, W., & Merten, L. (2023). Content analyses of user comments in journalism: A systematic literature review spanning communication studies and computer science. *Digital Journalism*, 11(7), 1328-1352.
- Risch, J., & Krestel, R. (2020, May). Top comment or flop comment? Predicting and explaining user engagement in online news discussions. In *Proceedings of the International AAAI Conference on Web and Social Media* (Vol. 14, pp. 579-589).
- Salminen, J., Sengün, S., Corporan, J., Jung, S. G., & Jansen, B. J. (2020). Topic-driven toxicity: Exploring the relationship between online toxicity and news topics. *PloS One*, 15(2), e0228723.
- Shmargad, Y., & Klar, S. (2020). Sorting the news: How ranking by popularity polarizes our politics. *Political Communication*, *37*(3), 423-446.

- Soroka, S., Fournier, P., & Nir, L. (2019). Cross-national evidence of a negativity bias in psychophysiological reactions to news. *Proceedings of the National Academy of Sciences*, 116(38), 18888-18892.
- 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(1), 24-54.
- Thelwall, M., Buckley, K., & Paltoglou, G. (2012). Sentiment strength detection for the social web. *Journal of the American Society for Information Science and Technology*, 63(1), 163-173.
- Wahl-Jorgensen, K. (2002). Understanding the conditions for public discourse: Four rules for selecting letters to the editor. *Journalism Studies*, *3*(1), 69-81.
- Wardoyo, C. (2019). Contagiousness of politeness on YouTube. Paradigm, 2(2), 139-148.
- Weber, P. (2014). Discussions in the comments section: Factors influencing participation and interactivity in online newspapers' reader comments. *New Media & Society*, 16(6), 941-957.
- Zheng, H., Goh, D. H. L., Lee, E. W. J., Lee, C. S., & Theng, Y. L. (2022). Understanding the effects of message cues on COVID-19 information sharing on Twitter. *Journal of the Association for Information Science and Technology*, 73(6), 847-862.
- Zhu, M., & Kadirova, D. (2022). Self-directed learners' perceptions and experiences of learning computer science through MIT open courseware. *Open Learning: The Journal* of Open, Distance and e-Learning, 37(4), 370-385.
- Ziegele, M., & Jost, P. B. (2020). Not funny? The effects of factual versus sarcastic journalistic responses to uncivil user comments. *Communication Research*, 47(6), 891-920.
- Ziegele, M., Weber, M., Quiring, O., & Breiner, T. (2018). The dynamics of online news discussions: Effects of news articles and reader comments on users' involvement, willingness to participate, and the civility of their contributions. *Information, Communication & Society*, 21(10), 1419-1435.

Biographical Notes

Jinzhi Zhou [zhoujinz@iu.edu] is a Ph.D. candidate in Learning Sciences at Indiana University, Bloomington. Her current research interests include computer-supported collaborative learning and computer-mediated communication. She has expertise in quantitative analysis, conversation analysis, and computer-mediated discourse analysis.

Susan C. Herring [herring@indiana.edu] is Professor of Information Science and Linguistics at Indiana University, Bloomington, where she also directs the Center for Computer-Mediated Communication. She specializes in computer-mediated discourse analysis and multimodal CMC. She is the current Editor-in-Chief of *Language@Internet*.