

Rumors in Retweet: Ideological Asymmetry in the Failure to Correct Misinformation

Matthew R. DeVerna¹ , Andrew M. Guess²,
Adam J. Berinsky³, Joshua A. Tucker⁴,
and John T. Jost⁴ 

Personality and Social
Psychology Bulletin
1–15
© 2022 by the Society for Personality
and Social Psychology, Inc
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/01461672221114222
journals.sagepub.com/home/pspb


Abstract

We used supervised machine-learning techniques to examine ideological asymmetries in online rumor transmission. Although liberals were more likely than conservatives to communicate in general about the 2013 Boston Marathon bombings (Study 1, $N = 26,422$) and 2020 death of the sex trafficker Jeffrey Epstein (Study 2, $N = 141,670$), conservatives were more likely to share rumors. Rumor-spreading decreased among liberals following official correction, but it *increased* among conservatives. Marathon rumors were spread twice as often by conservatives pre-correction, and nearly 10 times more often post-correction. Epstein rumors were spread twice as often by conservatives pre-correction, and nearly, eight times more often post-correction. With respect to ideologically congenial rumors, conservatives circulated the rumor that the Clinton family was involved in Epstein's death 18.6 times more often than liberals circulated the rumor that the Trump family was involved. More than 96% of all fake news domains were shared by conservative Twitter users.

Keywords

rumor, misinformation, correction, social media, political psychology

Received August 25, 2021; revision accepted June 30 2022

In a homogenous social medium, rumor is set in motion and continues to travel by its appeal to the strong personal interests of the individuals involved in the transmission. The powerful influence of these interests harnesses the rumor largely as a rationalizing agent, requiring it not only to express but also to explain, justify and provide meaning for the emotional interest at work. (Allport & Postman, 1946)

In 2013, Abdulrahman Alharbi—a Saudi man—was questioned by the Federal Bureau of Investigation (FBI) about his involvement in the Boston Marathon bombing. Although he was cleared within 24 hours, many social media users continued to claim that he was responsible. This illustrates an important concern; correcting misinformation is not easy. Psychologists have long understood that falsehoods often persist, even after a valid correction has been provided—a phenomenon dubbed the “continued influence effect” (Lewandowsky et al., 2012). Research on political communication yields similar conclusions (Ecker & Ang, 2019; Nyhan & Reifler, 2015; Swire et al., 2017; Thorson, 2016). In extreme cases, people may exhibit a “backfire effect,” redoubling their commitment to a belief that has seemingly been debunked (Jerit & Zhao, 2020).

The peculiarities of modern social media networks and the dynamics of ideological conflict and polarization may exacerbate the continued influence effect. One worry is that because of “echo chambers” belief-inconsistent corrective information may not reach the appropriate target (Barberá et al., 2015; Del Vicario et al., 2016; Sasahara et al., 2021). Even if people are exposed to corrective information, they may engage in cognitive dissonance reduction (Vraga, 2015) and fail to update their beliefs (Sinclair et al., 2020). In the present research program, we tracked the diffusion of misinformation on Twitter, exploring the role of ideology in the processing of corrective information over time.

We focus on two case studies that highlight the complex dynamics of misinformation diffusion in the real world. The

¹Indiana University, Bloomington, USA

²Princeton University, NJ, USA

³Massachusetts Institute of Technology, Cambridge, USA

⁴New York University, New York City, USA

Corresponding Author:

Matthew R. DeVerna, Luddy School of Informatics, Computing, and Engineering, Indiana University, 700 N. Woodlawn Avenue, Bloomington, IN 47408, USA.

Email: mdeverna@iu.edu

first pertained to rumors surrounding the Boston Marathon bombings in 2013. The second concerned the death of accused sex trafficker Jeffrey Epstein while awaiting trial in 2019. By harvesting large corpuses of Twitter messages, we tracked the spread of misinformation before and after attempts at correction took place.

We also investigated the possibility that there would be an ideological asymmetry—in terms of initial rumor-spreading and the subsequent likelihood of correcting misinformation. Political conservatives are more likely than liberals to possess an “intuitive” (vs. “analytic”) thinking style and to exhibit intolerance of uncertainty, as well as less integrative complexity, cognitive reflection, and verbal reasoning ability (Hodson & Busseri, 2012; Jost, 2021; Zmigrod et al., 2021). Some evidence suggests that conservatives are especially avoidant of dissonant information (Vraga, 2015) and reluctant to update their beliefs in light of contradictory evidence (Sinclair et al., 2020).

In the United States, political conservatives—including supporters of the QAnon Movement—are more likely to engage in conspiratorial thinking in general and to endorse a wide range of specific conspiracy theories (Van der Linden et al., 2021). Compared to liberals, conservatives are more receptive to pseudo-profound but meaningless statements (Evans et al., 2020; Pfattheicher & Schindler, 2016; Sterling et al., 2016). They are more susceptible to false information concerning potentially dangerous events (Calvillo et al., 2020; Fessler et al., 2017; Losee et al., 2020) and more likely to spread low quality information from unreliable sources (Mosleh et al., 2021). Consistent with research in psychology, studies of political communication show that rumors, “fake news,” and conspiracy theories spread more quickly and extensively in the social networks of conservatives (Benkler et al., 2017; Guess et al., 2019, 2020; Marwick & Lewis, 2017; Miller et al., 2016; Vosoughi et al., 2018). Furthermore, conservatives are generally more distrusting than liberals of politicians, government officials, journalists, scientists, academics, and other representatives of “officialdom” (e.g., Azevedo & Jost, 2021; Kraft et al., 2015; Miller et al., 2016; Pew Research Center, 2021). For these reasons, we hypothesized that conservatives might be more susceptible to rumor-mongering and less responsive to corrective attempts, in comparison with liberals.

On the other hand, the fact that conservatives exhibit more conscientiousness, conventional thinking, and respect for authorities might mean that they would be less susceptible to rumor-mongering and more responsive to official announcements that debunk misinformation (Carney et al., 2008; Graham et al., 2009; Lambert & Chasteen, 1997). Moreover, it has been claimed that liberals and conservatives are equally prone to motivated reasoning and biased information processing (Crawford, 2012; Ditto et al., 2019; Frimer et al., 2017), in which case we would expect no significant ideological differences in the spread of false rumors or the failure to update beliefs in a rational manner. Previous

research is inconclusive. According to a literature review by Jerit and Zhao (2020), “Evidence from several previous studies suggests that conservatives are more prone than liberals to backfire effects . . . Yet, other studies show that motivated cognition occurs across the whole ideological spectrum” (p. 85).

Ideological asymmetries in either direction should be of interest to psychologists, given that the psychological study of rumor transmission has—for at least 70 years—assumed that there is something fundamental, if not universal, about human nature that makes us vulnerable to political and other forms of misinformation under circumstances of ambiguity and emotional involvement (Di Fonzo & Bordia, 2007; Festinger et al., 1948; Rosnow, 1991). For instance, Allport and Postman (1946) wrote that: “As old as human society itself, rumor has flourished in wars and depressions, in peace and prosperity.” This is because “Our minds protest against chaos: from childhood we are asking *why, why?*” It has long been recognized that rumors may be spread for ideological reasons—to discredit political adversaries and share propagandistic versions of current events (Rosnow & Fine, 1976, pp. 100–106). To our knowledge, however, no previous study has directly addressed the question of whether liberals and conservatives are equally likely to believe or spread rumors—and whether they are equally adept at updating their beliefs and behaviors once the misinformation has been officially corrected.

The Psychology of Rumor

This work focuses on the spread of rumors to address a growing interest in, and need to understand, the psychology of misinformation (Pennycook & Rand, 2021; Van Der Linden et al., 2020; Lazer et al., 2017). We believe the study of rumor spreading is highly relevant as it seems possible that the psychological mechanisms driving people to spread rumors may underlie all of these phenomena.

Rumors belong to an especially interesting class of information-sharing, because of their high emotional and motivational content. As Allport and Postman (1946) pointed out, people do not spread rumors about people or events that they care little about, and many rumors satisfy “a primary emotional urge” by, for example, “permitting one to slap at the thing one hates” (p. 503). Rosnow and Fine (1976) likewise note that “some emotional arousal in the form of anxiety also seems essential” to the act of rumor-spreading (p. 30). At the same time, rumors address cognitive—or perhaps epistemic—functions, what Allport and Postman referred to as “effort after meaning,” that is, the continuous human desire to “extract meaning from our environment” and the “pursuit of a ‘good closure,’” a subjectively satisfying way of “explaining, justifying, and providing meaning for the emotional interest at work” (pp. 503–504). Thus, putting these various elements together, rumor is a distinctive phenomenon that involves the spread of emotionally salient

“information, neither substantiated nor refuted,” that is, “most often fueled by a desire for meaning, a quest for clarification and closure” (Rosnow & Fine, 1976, p. 4).

Consistent with this formulation, Allport and Postman (1946) proposed a theory of rumor intensity as a multiplicative function of two variables, namely emotional intensity and informational ambiguity. They argued that people do not spread rumors when emotional intensity is zero (i.e., when they do not care) or when informational ambiguity is zero (i.e., when they have direct, incontrovertible knowledge about the phenomenon of interest). Conversely, people should be most likely to spread rumors when intensity and ambiguity are both very high, as in the two cases we examine here, namely, the Boston Marathon bombing in 2013 and the suspicious death of Jeffrey Epstein in 2019.

A book-length treatment of the psychology of rumor by Rosnow and Fine (1976) found plenty of support for Allport and Postman’s (1946) general conception. The former pair of authors concluded, for instance, that, “Where the need is for information or clarity, rumor fills the void of ambiguity and uncertainty . . . Thus rumors arise when there is any exciting or mysterious event that has not been fully explained” (p. 75). There are two other aspects of Rosnow and Fine’s analysis that are especially useful for the present purposes of our investigation: (a) the authors observe that the domain of politics is an especially fertile ground for the spread of rumors (especially in the United States, where libel laws are less strict than in other countries, such as the United Kingdom) and (b) they acknowledge that there are important individual differences that affect susceptibility to rumor: “not every individual in society necessarily hears the rumor, and not everyone who has heard it believes it” (p. 75).

Although many of the political rumors described by Rosnow and Fine in the middle of the 20th century appear to have had a right-wing (and frequently anti-Communist) bent, these authors did not consider the possibility that there might be an ideological asymmetry in rumor transmission. However, as detailed in the previous section, whether there may be an ideological asymmetry with respect to rumor spreading remains an open question. Although it is extremely likely that leftists and rightists would both be capable of experiencing high levels of emotional intensity (sometimes about the same issues and sometimes about different issues), there is good reason to think that the other variable in Allport and Postman’s (1946) equation, namely, the psychological response to informational ambiguity and the pursuit of cognitive closure, is indeed subject to ideological asymmetry.

Study 1

Method

In Study 1, we investigated the spread of two rumors pertaining to the Boston Marathon bombings (Starbird et al., 2014). One focused on Abdulrahman Alharbi, the “Saudi man” who

was questioned by the FBI about his involvement. Three days after the bombings, on April 18, 2013, Alharbi was exonerated, and investigators released images of two alternative suspects. In the early morning hours of April 19, Dzhokhar and Tamerlan Tsarnaev were confirmed as primary suspects. The second case involves the rumor that the bombings represented a “false flag” attack (i.e., a covert operation by the U.S. government). In relatively short order, mainstream news sources reported that Alharbi was no longer a suspect, switched their focus to the Tsarnaev brothers, and refuted the notion that the bombing was a false flag attack (Gross, 2013). In this study, we traced the spread of two logically inconsistent rumors—that Alharbi was involved in the Boston bombings, and the bombings were the result of a false flag operation by the U.S. government—both before and after they were debunked.

To track real-time reactions to misinformation and corrective information regarding the attacks, we used tools created by the Social Media and Political Participation (SMaPP) Lab at New York University to collect Twitter messages in real time using the streaming application programming interface (API) based on keywords listed in the Supplement (see Table S1). Our analysis is based on messages harvested from April 15 to April 25, 2013 that made reference to the Boston Marathon bombings. (Because of technical problems, we were unable to save tweets posted on April 20.) All tweet IDs are available on the Harvard Dataverse: <https://doi.org/10.7910/DVN/TYCTGN>.

User ideology was estimated in 2016 with the use of a validated method based on individual users’ patterns of “following” political elites—such as elected representatives and media figures (see Barberá et al., 2015, for a description of the method and evidence for its validity). We dichotomized the sample into two groups: liberal (less than 0, the midpoint) and conservative (greater than 0). (No user had an estimated score of exactly 0.) Of the total number of ideologically classifiable users who tweeted about the Boston Marathon ($N = 487,156$), 67.4% were liberal and 32.6% were conservative. Thus, liberals were more likely to communicate about the Boston Marathon bombing in general, but this is probably because Democrats are significantly more active on Twitter than Republicans (Pew Research Center, 2021). To estimate ideological base rates for Twitter users at the time, we used Barberá et al.’s (2015) method to classify over 12 million users in January of 2016 (a few years after the period of data collection, but the closest period for which we have reliable data) who followed at least three political accounts. Of these politically interested users, we found that 64.5% were classified as liberal ($n = 7,976,620$) and 35.5% as conservative ($n = 4,384,807$).

The “Saudi” rumor. For the rumor about Alharbi’s involvement, we focused on 77,147 tweets (sent by 43,027 unique users) that referenced the rumor about the bomber being “Saudi” or a “Saudi national” (see keywords and sample tweets for both

Table 1. Rumors, Keywords, and Sample Tweets Pertaining to the Boston Marathon Bombing (Data Collection for Study 1).

Rumors	Keywords	Sample tweets
Saudi bomber	saudi (case insensitive)	Breaking: authorities ID a Saudi national as a suspect in Boston Marathon bombings http://t.co/4hfqLc3n0A Did the Boston terrorist come in on Obama's \"Saudis go to the front of the line\" flying policy? I find it interesting that 9/11 and potentially today's attack involved Saudi nationals. You know . . . Our ally in the Mideast
False flag operation	false flag, falseflag	False Flag Attack Iran? http://t.co/VwsIwrWb8z I Boston Bomb Blast to bring in Martial Law in 4 USA No Doubt a False Flag E . . . \"Anyone with a backpack needs to come to a Boston police station to cross off the list\"—said no police statement ever #falseflag

rumors in Table 1). We were able to classify the ideological positions of 22,749 (or 53%) of the users who tweeted about this rumor. Of these users, 15,326 were classified as conservative (67.4% of the users who could be classified), and 7,423 as liberal (32.6%). In terms of rumor-specific tweets, 42,189 were sent by conservatives (82.8% of those that could be classified ideologically, 54.7% of the total), and 8,754 were sent by liberals (or 17.2% of those that could be classified; 11.3% of total). Overall, conservative users in our sample sent 2.75 tweets about the Saudi rumor on average, compared to 1.18 sent by liberal users.

The “False Flag” rumor. According to the “false flag” rumor, the Boston Marathon bombing was staged by the government to justify violations of civil liberties. Using the search terms of “false flag” and “falseflag,” we were able to identify 19,756 tweets sent by 11,804 users. We were able to classify the ideological positions of 4,968 (or 42.1%) of the users who tweeted a total of 9,139 times about the false flag rumor (accounting for 46% of the total number of tweets about this rumor). Of these users, 2,830 were classified as conservative (57.0% of the users who could be classified), and 2,138 as liberal (43.0%). In terms of rumor-specific tweets, 5,776 (63.2% that could be classified, 29.2% of the total) were sent by conservatives, and 3,363 (36.8% that could be classified, 17.0% of the total) were sent by liberals. Overall, conservative users in our sample sent 2.04 tweets about the false flag rumor on average, compared to 1.57 sent by liberal users.

Supervised learning technique to identify rumor-spreading tweets. A challenge with Twitter data collected via keyword-based search methods is that posts intended to debunk the rumor—by questioning, expressing skepticism, or referring ironically to that rumor—may be incorrectly grouped with posts that were intended to spread the rumor. To address this issue, we used supervised learning techniques to classify tweets as attempting to spread the rumor (vs. not).

Supervised learning relies on a training set that consists of labeled data—in this case, a random set of hand-coded tweets. Treating this representative subset as the ground truth, we then apply automated text classification to code

the remainder of the data. The specific model we use is L2 penalized (“ridge”) logistic regression. Penalization helps the model perform feature selection when there is a large number of features. In the case of this tweet data, we have as many as 75,000 features comprising 1, 2, and 3 grams extracted directly from text. Ridge regression performs particularly well when predictors are highly correlated, as may be the case when we have a distinct set of words and phrases that often go together during a highly publicized incident. We use this classification procedure to separate “real,” credulous mentions of a rumor from attempts to debunk or contextualize it.

Two coders performed the labeling task on tweets related to the “Saudi national” and “false flag” rumors. In the first case, to generate the set of tweets we randomly sampled 200 before and 200 after the alleged bombers’ capture to ensure adequate representation of words and phrases used both during the initial period of uncertainty and after the official “correction” had been issued. In both cases, the coders labeled each tweet as genuinely “spreading” the rumor—which includes retweets—or not, which included unrelated tweets and attempts to debunk, question, or even mock the rumor. (All coding was done blind to ideology.) We found that 90% of the “Saudi” tweets and 84% of the “false flag” tweets were classified with the same rating, and in both cases, we retained only the codes for which there was unanimous agreement. The vast majority of tweets were classified as genuine attempts to spread the rumor, leaving the results essentially unchanged from an earlier analysis based exclusively on keyword selection. For ease of exposition, we present the results from the latter, more refined analysis.

Results

The daily proportions of tweets about the Boston Marathon bombing that mention the Saudi rumor are shown in the left panel of Figure 1. Overall, the proportion of tweets in our collection containing the rumor about the suspect being a Saudi national is low. At its peak on April 17, 2013, there were slightly more than 18,000 tweets about the rumor. The figure shows a clear increase in the share of

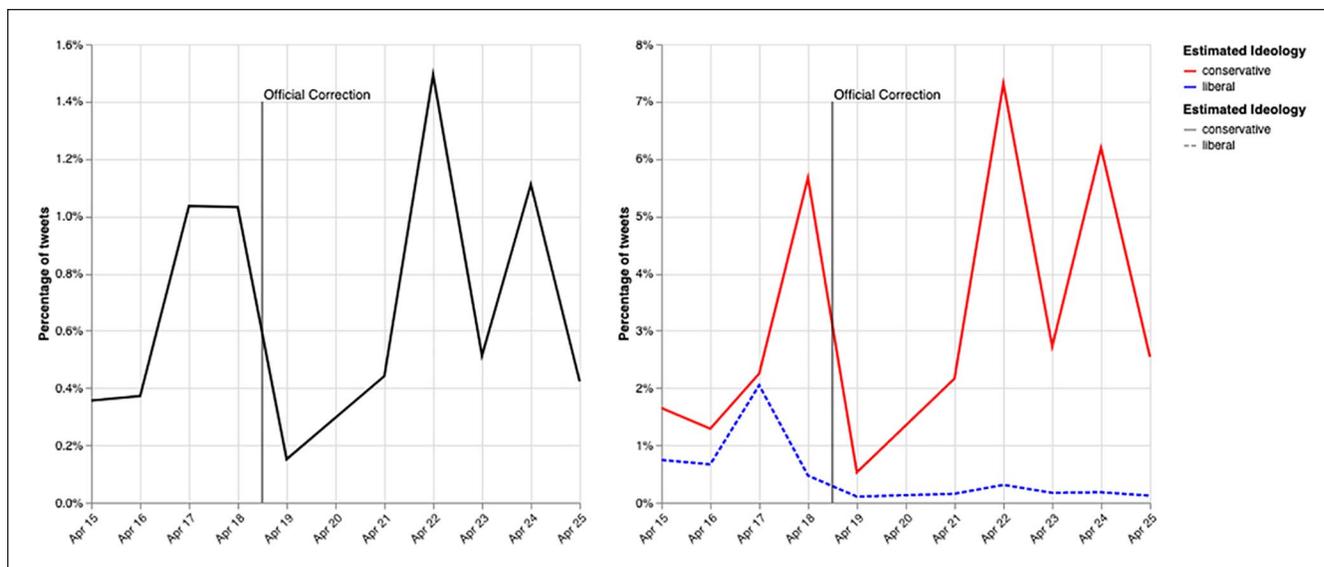


Figure 1. Percentage of Saudi rumor tweets per day for the total sample (left) and for liberal and conservative social media users separately (right).

Note: The panel on the left shows the percentage of tweets about the Saudi rumor, relative to the total number of tweets in the Boston Marathon collection sent on the same day. The panel on the right shows the percentage of the tweets about the Saudi rumor sent by users who were classified as liberal or conservative, relative to the total number of tweets in the Boston Marathon collection sent by liberal and conservative users, respectively, on the same day. The vertical line indicates when official corrections were provided to the public, that is, the announcement of information about alternative suspects (and the official debunking of the original rumor).

tweets mentioning the rumor leading up to the emergence of the news about the Tsarnaev brothers early the morning of April 19. The share plummets following this revelation. However, within a few days, there is a rebound in tweets about the Saudi rumor.

Importantly, this pattern differs greatly according to the ideology of the users circulating the rumor. The right panel of Figure 1 shows the proportion of tweets (and retweets) by liberals and conservatives mentioning the Saudi rumor on each day of data collection. The proportion of messages by liberal Twitter users changes very little over the entire period; it generally hovers below 0.5%. By contrast, the proportion for conservatives is higher than that for liberals on every day. Furthermore, the share of rumor-related tweets varies over time. The first spike in conservatives' spreading of the rumor occurs on April 18, 1 day before the correction and 3 days after the bombing. A marked decrease occurs the next day, following the correction. However, there is a significant rebound in the proportion of tweets by conservatives, but not liberals, after the suspects were identified as non-Saudi. In fact, the largest spikes in rumor-spreading occur during this post-correction period.

A similar ideological asymmetry in message patterns emerges when we analyze the rumor that the attacks were a "false flag" operation by the U.S. government (see Figure 2). Conservatives circulated the rumor more than liberals both prior to and after the correction. In contrast to the first case, similar proportions of liberals and conservatives revisited the rumor in the days after it was corrected, but conservatives

were more likely to tweet about it over time. After April 23, a statistically significant gap reemerges between liberals' and conservatives' tweeting about the rumor (in an ordinary least-squares [OLS] model, an interaction between ideology and a linear time trend beginning on April 23 confirms this at $p < .001$). To ensure that these results were not sensitive to our coding of ideology or to specific cut-off points between liberals and conservatives, we re-estimated the analyses after excluding observations that were close to the estimated midpoint of the latent ideological continuum. The results were effectively identical (see Supplement, Figures S1 and S2).

In Table 2, we compare both types of rumor-spreading for liberal and conservative Twitter users separately (and for the sample as a whole) before and after the official correction. Results show that before the correction was made, liberals in this sample were responsible for sharing the rumor about the Saudi bomber 7,740 times and the false flag rumor 1,868 times. During the same period conservatives spread the Saudi rumor 20,326 times and the false flag rumor 2,868 times. Thus, before the correction was made, conservatives were responsible for spreading the rumors roughly 2.4 times more than liberals.

After the correction was made, liberals' rumor-spreading behavior decreased. During this period, liberals spread a total of 2,509 rumors (1,014 about the Saudi bomber and 1,495 about the false flag), whereas conservatives spread 24,821 rumors (21,913 about the Saudi bomber and 2,908 about the false flag), which was slightly a larger number than before the correction was made. After the correction was made,

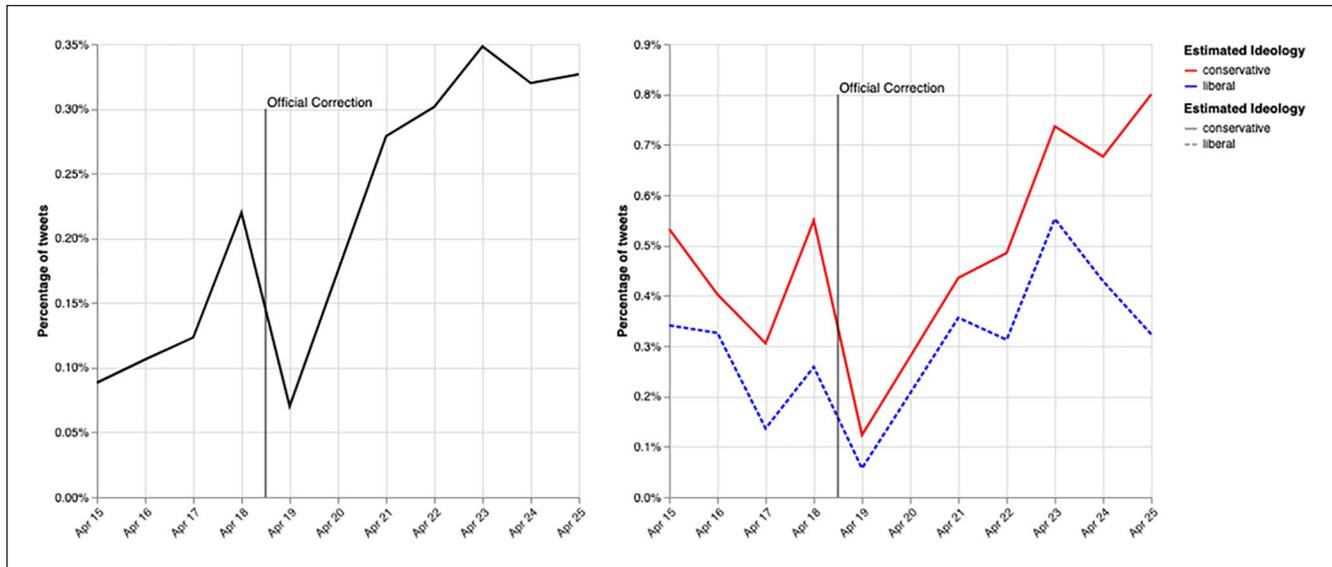


Figure 2. Percentage of false flag rumor tweets per day for the total sample (left) and for liberal and conservative social media users separately (right).

Note: The panel on the left shows the percentage of tweets about the false flag rumor, relative to the total number of tweets in the Boston Marathon collection. The panel on the right shows the percentage of the tweets about the false flag rumor sent by users who were classified as liberal or conservative, relative to the total number of tweets in the Boston Marathon collection sent by liberal and conservative users, respectively, on the same day. The vertical line indicates when official corrections were provided to the public (i.e., official debunking of the notion that the bombings represented a “false flag” attack).

Table 2. Rumor-Specific Tweets Sent by Liberal, Conservative, and Total Users Before and After the Official Correction Took Place (Study 1).

	Pre-correction period	Post-correction period
Liberal Twitter users ($n = 7,423$)		
Tweets spreading Saudi rumor	7,740	1,014
No. of users who spread Saudi rumor	6,746	853
Tweets per user	1.15	1.19
Liberal Twitter users ($n = 2,138$)		
Tweets spreading false flag rumor	1,868	1,495
No. of users who spread false flag rumor	1,350	935
Tweets per user	1.38	1.60
Conservative Twitter users ($n = 15,326$)		
Tweets spreading Saudi rumor	20,326	21,913
No. of users who spread Saudi rumor	10,130	8,639
Tweets per user	2.01	2.54
Conservative Twitter users ($n = 2,830$)		
Tweets spreading false flag rumor	2,868	2,908
No. of users who spread false flag rumor	1,715	1,498
Tweets per user	1.67	1.94
Total users ($n = 43,027$)		
Tweets spreading Saudi rumor	49,092	28,177
No. of users who spread Saudi rumor	34,391	13,098
Tweets per user	1.43	2.15
Total users ($n = 11,804$)		
Tweets spreading false flag rumor	9,564	10,192
No. of users who spread false flag rumor	6,716	6,055
Tweets per user	1.42	1.68

Note: “Total Users” includes those who could not be ideologically classified.

conservatives spread these two rumors approximately 9.9 times more than liberals did.

Discussion

With respect to two major rumors pertaining to the 2013 Boston Marathon bombings, we observed that conservatives were more likely than liberals to spread misinformation initially and return to it later, even after it was officially debunked. This was the case despite the fact that—by a ratio of nearly 2 to 1—liberals were more likely to communicate about the bombings in general. No doubt, perceptions of source credibility and trustworthiness played a key role. The claim that Alharbi was involved in the Boston Marathon bombing was pushed by Glenn Beck, a right-wing commentator, and the “false flag” rumor was heavily promoted by right-wing conspiracy theorist Alex Jones. As a result, both rumors circulated prominently in the right-wing blogosphere.

The fact that there were no false rumors about the Boston Marathon bombing regularly promulgated by left-leaning actors is consistent with the possibility that conspiratorial thinking is more common on the right—at least during the present historical period in the United States (Van der Linden et al., 2021). Because rumors about the bombing, which took place during the Obama Presidency, were more popular among conservative and right-wing media personalities and platforms, it was not possible in this study to distinguish unambiguously between the effects of (a) ideological congruence between unreliable media sources on the right and their audience members and (b) a more general ideological asymmetry in susceptibility to rumor, misinformation, and motivated reasoning. We conducted a second study focusing on a scandal that took place during the Trump Presidency, namely the death of Jeffrey Epstein.

Study 2

Method

On July 6, 2019 Jeffrey Epstein was indicted on charges of sex trafficking and conspiracy (Brown vs. Maxwell; Dershowitz vs. Giuffre, 2018). According to publicly released court documents, it was alleged that “numerous prominent American politicians” (and other well-known individuals) were involved in sexually abusing teenage girls along with Epstein (Brown vs. Maxwell; Dershowitz vs. Giuffre, 2018). These shocking allegations, combined with the fact that Epstein had been convicted as a sex offender in Florida in 2008, generated massive media attention. On August 10, 2019, Epstein was found dead in his prison cell from what was subsequently ruled a suicide. The news of Epstein’s death shocked the public and spread quickly, fueled by the fact that he had been placed on suicide watch after a failed attempt only a few weeks earlier (Schapiro & Dienst,

2019). Epstein’s political connections fueled rumors and conspiracy theories, especially on social media platforms. We analyzed the dynamics pertaining to three of the most popular rumors, which were that (a) Epstein did not in fact kill himself (i.e., that he was either murdered or still alive), (b) the Clinton family was involved in having Epstein murdered, and (c) the Trump family was involved in having him murdered.

To follow social media users’ reactions to various rumors and attempts at correction, we initially harvested over 33 million Twitter messages making reference to Jeffrey Epstein from July 9 to December 31, 2019. We utilized Twitter’s V1 API filter endpoint to stream tweets in real-time, beginning 1 month prior to his death, based on a list of keywords and hashtags related to Epstein’s sex trafficking case at the time (see Supplement, Table S2). Following Epstein’s death, we added several more keywords and hashtags on August 10 and 15 that were in wide circulation at the time. Any message containing one or more of these keywords was captured and stored within the collection. As was the case with respect to the Boston Marathon bombing, liberals were more likely than conservatives to tweet about the Epstein case in general. For the total sample of approximately 33 million tweets sent by users whom we could classify ideologically, 59.6% were sent by liberals and 40.4% by conservatives.

In the aftermath of Epstein’s death, several rumors emerged. For the purposes of exploring ideological symmetries and asymmetries we focused on three of the most prominent rumors: (a) a general rumor that was not explicitly political, namely that Epstein did not actually commit suicide (Epstein didn’t kill himself or EDKH); (b) a rumor that was expected to be especially congenial to conservatives, that the family of Bill and Hillary Clinton was involved in having Epstein murdered to cover up their alleged involvement in Epstein’s criminal activity (Clinton body count or CBC); and (c) a rumor that was expected to be especially congenial to liberals, namely that the family of Donald Trump was involved in having Epstein murdered to cover up his alleged involvement in Epstein’s criminal activity (Trump body count or TBC).¹

To compare rumor-spreading behavior before and after a credible correction, we used August 16—the date that Epstein’s highly publicized autopsy report was released to the public—as the cut-off point. The pre-correction period, then, began on August 10 (the date Epstein died), and the post-correction period ran from August 16 to December 31, 2019 (when we stopped data collection).

Supervised learning technique to identify rumor-spreading tweets. To identify tweets focusing on these three rumors, we utilized a keyword matching procedure based on hashtags listed in Table 3. This procedure produced 526,882 tweets that mentioned one or more of the three rumors. To remove messages that were sent for reasons other than rumor-spreading (e.g., to mock or question a rumor), three human coders first

Table 3. Rumors, Keywords, and Sample Tweets Pertaining to the Death of Jeffrey Epstein (Data Collection for Study 2).

Rumors	Keywords	Sample tweets
Epstein didn't kill himself (EDKH)	epsteindidntkillhimself Epstein didn't kill himself	we know that epstein didn't kill himself. but, do we know the body examined by the coroner was, in fact, epstein? obama is not an american citizen, brennan is a communist, the climate change crisis is a nwo attempt to control the world's population, hillary is guilty as hell, there's only 2 genders, mccarthy was right, tds is real, the un is useless & epstein didn't kill himself. l is for the way you look at me. o is for the only one i see. v is very, very extraordinary. e is for the fact that epstein didn't kill himself.
Clinton body count (CBC)	clintonbodycount clinton body count	opps, there's another one added to the clinton body count. that didn't take long. the fbi is on the case #clintonbodycount #clintoncrimefamily #epstien #epsteinsuicide #epsteinmurder #billclinton #hillaryclinton epstein's prison guards charged with falsifying jail records. two federal prison officers were arrested and charged on tuesday with conspiracy and falsifying jail records following the "clinton suicide" of billionaire pedophile epstein.
Trump body count (TBC)	trumpbodycount trump body count	so. we're expected to believe, shortly after finding jeffrey epstein injured in federal prison, a pedophile with dirty secrets of the rich & famous was not on suicide watch and killed himself? nope. not buying it. #trumpbodycount #epsteinmurder jeffrey epstein was under watch of the donald trump administration and bill barr in the justice department. and there was a previous attempt at suicide we are told. no excuses! #trumpbodycount prince andrew #epsteinfiles #epsteinmurder #amjoy let's all keep in mind that barr's dad was the first to hire epstein. prepare for evidence to disappear. #trumpbodycount

EDKH = Epstein Didn't Kill Himself; CBC = Clinton body count; TBC = Trump body count.

Note: For the collection of rumor-specific tweets we used machine-learning methods to match tweets based on textual patterns (rather than hashtags) because this approach is more comprehensive and uses regular tweet text as well as hashtags and keywords. If a user were to include an apostrophe in a hashtag (e.g., #epsteindidntkillhimself), Twitter would exclude the text trailing the apostrophe (so the hashtag would not work). Therefore, we omitted concatenated text forms (that included apostrophes in hashtags). For sample tweets, all text was converted to lowercase format and links that were automatically shortened (and have since been deleted on Twitter) were removed from the text.

manually labeled a random subset of 4,550 tweets each (13,650 tweets in total) as either "rumor spreading" or "not rumor spreading."² We then used the hand-coded data set to train three separate supervised machine-learning models; one model for each rumor. All tweet IDs are available on the Harvard Dataverse.³

We used the "autotune" feature of the *fasttext* library (created by Facebook for efficient text classification) to identify the best hyperparameters for our data set. The models performed very well, as indicated by high levels of precision, recall, F1, and F2 scores (see Supplement, Table S3). Using these classification models, we then labeled all rumor-specific tweets according to their respective model (or rumor type). When labeling tweets, *fasttext* estimates the continuous probability that the returned prediction label ("rumor spreading" or "not rumor spreading") is accurate, ranging from 0 (no confidence) to 1 (total confidence). Using these estimates, we selected only those tweets that were labeled as "rumor spreading" with a confidence probability greater than .75.

Based on this procedure, we removed 115,459 tweets that were classified as "not rumor spreading" (~21.9% of the total). This left us with a corpus of 411,423 rumor-specific messages that were sent by 141,670 individual Twitter users whose

ideological position we were able to estimate based on the same method of classification described in Study 1. Despite the fact that liberals were more likely than conservatives to tweet about Epstein in general, more than two-thirds of the rumor-spreaders were classified as conservative (67.8%, $n = 96,105$), and only 32.2% as liberal ($n = 45,565$). Furthermore, 81.7% of the tweets containing rumors were sent by conservatives ($n = 335,957$), as compared with 18.3% sent by liberals ($n = 75,466$). Thus, conservative users in our final sample sent 3.50 rumor-spreading tweets on average, compared to 1.66 sent by liberal users.

Identification of fake news domains. We also investigated how users disseminated fake news, as this type of content can be utilized to spread unsubstantiated rumors. Specifically, we compared the fake news sharing habits of liberals and conservatives by analyzing links included in rumor-spreading tweets. For this purpose, we used Grinberg et al.'s (2019) list of 505 well-known fake news sites, which comprise (a) "black" sites ($n = 382$) that were taken from a list compiled by experts; (b) "red" sites ($n = 61$), such as infowars.com or dailystormer.com that frequently disseminate major falsehoods, generally aligned with a specific ideological agenda; and (c) and

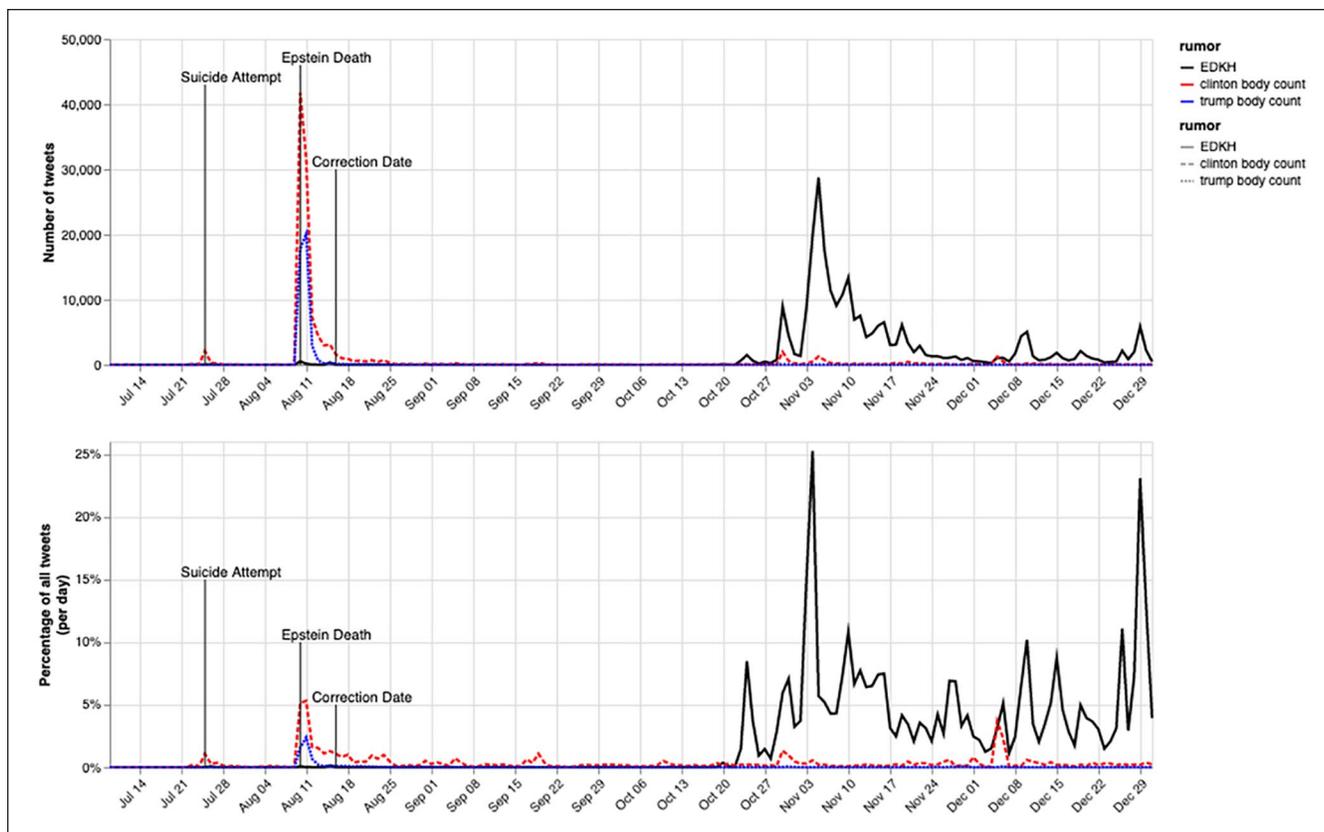


Figure 3. Number (top) and percentage (bottom) of rumor-specific tweets per day about the death of Jeffrey Epstein.

Note: The top panel shows the raw daily frequencies of tweets spreading each of the three major rumors. The bottom panel shows the percentage of the tweets about each of the three rumors, relative to the total number of tweets in the Epstein collection sent on the same day. Vertical lines indicate especially salient dates (Epstein’s unsuccessful suicide attempt, his death, and the official announcement of the cause of death).

“orange” sites ($n = 47$) that, according to the fact-checking website Snopes.com, also spread false information, but appear to do so less often or in a way that may or may not indicate a “systematically flawed” editorial process (see Grinberg et al., 2019, p. 374). To match the domains included in our corpus of tweets to fake news sites, we extracted the top-level domain from individual tweets by utilizing a Python package (*urlExpander*) that expands and identifies shortened URLs (Yin, 2018).

Results

The daily numbers and proportions of tweets about rumors pertaining to the Jeffrey Epstein case are shown in Figure 3. There is a clear spike in rumor-sharing (especially with respect to the CBC rumor) in the immediate aftermath of Epstein’s death and an apparent reduction in rumor-spreading following the official correction issued 6 days later. However, there is a consistent but low level of rumor-spreading before and after the correction date, with several additional spikes concerning the EDKH rumor showing up months later. Conservative Twitter users were much more likely than liberals to spread the EDKH and CBC rumors,

both in terms of raw numbers (see Table 4) and as a percentage of total tweets sent about Epstein within each ideological group (see Figure 4). This was especially true during the post-correction period.

Before the correction was made, liberals in the sample spread the EDKH rumor 125 times, the CBC rumor 3,596 times, and the TBC rumor 40,144 times. During the same period, conservatives spread the EDKH rumor 953 times, the CBC rumor 87,393 times, and the TBC rumor 1,676 times. Overall, conservatives were slightly more than twice as likely to spread rumors about Epstein’s death. It is important to keep in mind that conservatives did not tweet more about the Epstein case in general. As noted above, nearly 60% of the messages about Epstein in our sample were sent by liberals. Thus, liberals were more likely to tweet about the Epstein case overall, but conservatives were disproportionately likely to engage in rumor-spreading activity.

At the same time, the overall effect of ideology masks considerable variability with respect to the contents of the rumors that were spread. Concerning the most general (EDKH) rumor, conservatives were 7.6 times more likely than liberals to spread it before the correction, but there were very few total cases during this period. Conservatives were responsible

Table 4. Rumor-Specific Tweets Sent by Liberal, Conservative, and Total Users Before and After the Official Correction Took Place (Study 2).

	Pre-correction period	Post-correction period
Liberal Twitter users (<i>n</i> = 45,565)		
Tweets spreading EDKH rumor	125	28,912
No. of users who spread EDKH rumor	122	23,498
Tweets per user	1.02	1.23
Liberal Twitter users (<i>n</i> = 45,565)		
Tweets spreading CBC rumor	3,596	1,331
No. of users who spread CBC rumor	2,757	649
Tweets per user	1.30	2.05
Liberal Twitter users (<i>n</i> = 45,565)		
Tweets spreading TBC rumor	40,144	1,182
No. of users who spread TBC rumor	21,938	749
Tweets per user	1.83	1.58
Conservative Twitter users (<i>n</i> = 96,105)		
Tweets spreading EDKH rumor	953	220,689
No. of users who spread EDKH rumor	947	72,704
Tweets per user	1.01	3.04
Conservative Twitter users (<i>n</i> = 96,105)		
Tweets spreading CBC rumor	87,393	21,992
No. of users who spread CBC rumor	40,657	13,640
Tweets per user	2.15	1.61
Conservative Twitter users (<i>n</i> = 96,105)		
Tweets spreading TBC rumor	1,676	82
No. of users who spread TBC rumor	1,245	61
Tweets per user	1.35	1.34
Total users who could be ideologically classified (<i>n</i> = 141,670)		
Tweets spreading EDKH rumor	1,078	249,601
No. of users who spread EDKH rumor	1,069	96,202
Tweets per user	1.01	2.59
Total users who could be ideologically classified (<i>n</i> = 141,670)		
Tweets spreading CBC rumor	90,989	23,323
No. of users who spread CBC rumor	43,414	14,289
Tweets per user	2.10	1.63
Total users who could be ideologically classified (<i>n</i> = 141,670)		
Tweets spreading TBC rumor	41,820	1,264
No. of users who spread TBC rumor	23,183	810
Tweets per user	1.80	1.56

EDKH = Epstein Didn't Kill Himself; CBC = Clinton body count; TBC = Trump body count.

for spreading the CBC rumor 24 times more than liberals were, whereas liberals were responsible for spreading the TBC rumor 24 times more than conservatives were. This suggests a symmetrical pattern in which both liberals and conservatives were spreading ideologically congenial rumors (e.g., Crawford, 2012), but it is important to bear in mind that even before the official correction conservatives spread the CBC rumor 2.2 times more than liberals spread the TBC rumor.

Following the correction, liberals' overall rumor-spreading behavior decreased by nearly 28.35% (from 43,865 to 31,425 rumor-spreading tweets). By contrast, conservatives' rumor-spreading increased by 2.7 times (from 89,722 to 242,763 rumor-spreading tweets). Thus, whereas conservatives spread twice as many rumors before the correction, they were

responsible for spreading 7.7 times as many rumors as liberals after the correction. Nearly, all of the ideological divergence in rumor-spreading pertained to the EDKH rumor (conservatives spread this rumor 7.6 times more than liberals) and the CBC rumor (conservatives spread this rumor 16.5 times more than liberals). Consistent with the idea that people are more likely to spread ideologically congenial (vs. uncongenial) rumors, liberals spread the TBC rumor 14.4 times more than conservatives did. At the same time, it is important to keep in mind that there were very few cases of this rumor being spread after the official correction (see Figure 5). Moreover, consistent with the hypothesis of ideological asymmetry, the CBC rumor was shared 18.6 times more by conservatives than the TBC rumor was shared by

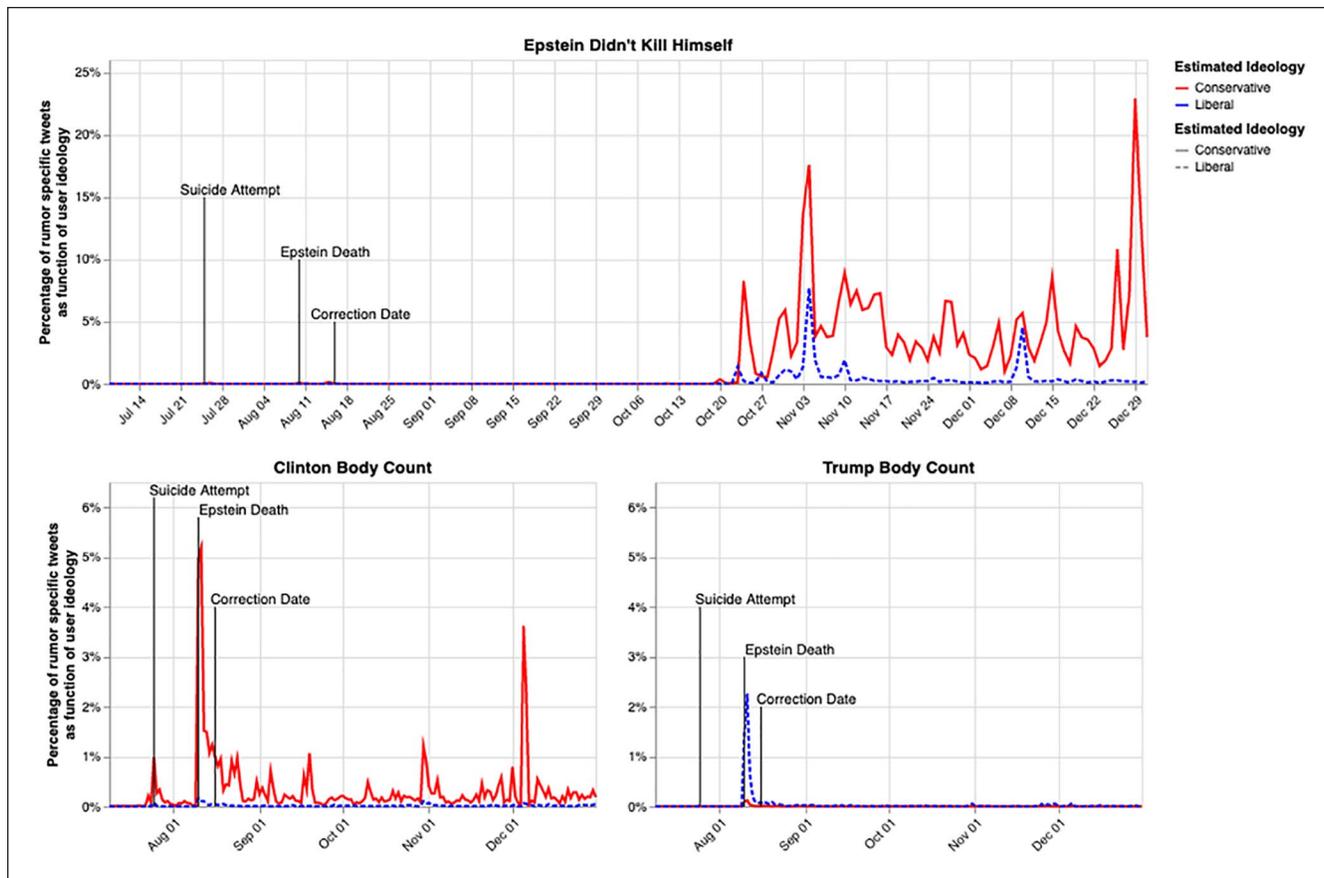


Figure 4. Percentage of rumor-spreading tweets per day for each of the three rumors about the death of Jeffrey Epstein for liberal and conservative social media users separately.

Note: For each of the three major rumors (shown in separate panels), percentages reflect the number of tweets sent by users who were classified as liberal or conservative, relative to total number of tweets in the Epstein collection sent by liberal and conservative users, respectively, on the same day. Vertical lines indicate especially salient dates (Epstein’s unsuccessful suicide attempt, his death, and the official announcement of the cause of death).

liberals during the post-correction period.⁴ These patterns are illustrated in Figure 5.

We also explored ideological differences in the sharing of links to fake news domains. In the corpus of tweets for Study 2, there were 71,246 tweets that included a total of 73,390 URLs (excluding twitter.com). Of these, 9,115 tweets contained links to at least one fake news site; there were a total of 9,249 fake news sites shared. Pooling “black,” “red,” and “orange” sites, we observed that 96.11% of the fake news sites (8,889 out of 9,249) were shared by conservative Twitter users, and only 3.89% (360 out of 9,249) were shared by liberal users. With respect to the EDKH rumor conservatives shared 98.10% of the fake news sites that were shared (2,375 out of 2,421), and with respect to the CBC rumor they shared 97.51% of the fake news sites (6,471 out of 6,636). With respect to the TBC rumor, liberals shared 77.60% of the fake news sites that were shared (149 out of 192), but there were far fewer fake news sites shared overall. Whereas conservatives shared 6,471 fake news sites concerning the CBC rumor, liberals shared only 149 fake news sites concerning the TBC rumor; this is a ratio of more than 40 to 1.

General Discussion

We have revisited a classic issue in social psychology—namely, the diffusion of rumor under circumstances of ambiguity and emotional salience—in the contemporary context of vast social media networks. Using supervised learning techniques to analyze the contents of over 30,000 tweets in Study 1 and over 400,000 tweets in Study 2, we have moved the study of political misinformation—and its ideological basis—from the confines of the laboratory to online communication with “real world” consequences.

In both studies, we observe evidence that suggests there was an ideological difference in rumor-spreading behavior prior to official corrections—and these differences were exacerbated during the post-correction period. With respect to the Boston Marathon bombings, rumors were spread more than twice as often by conservatives as liberals before the correction period, and they were spread nearly 10 times as often by conservatives after the attempt at correction was made. With respect to the death of Jeffrey Epstein, rumors were again spread twice as often by conservatives as liberals

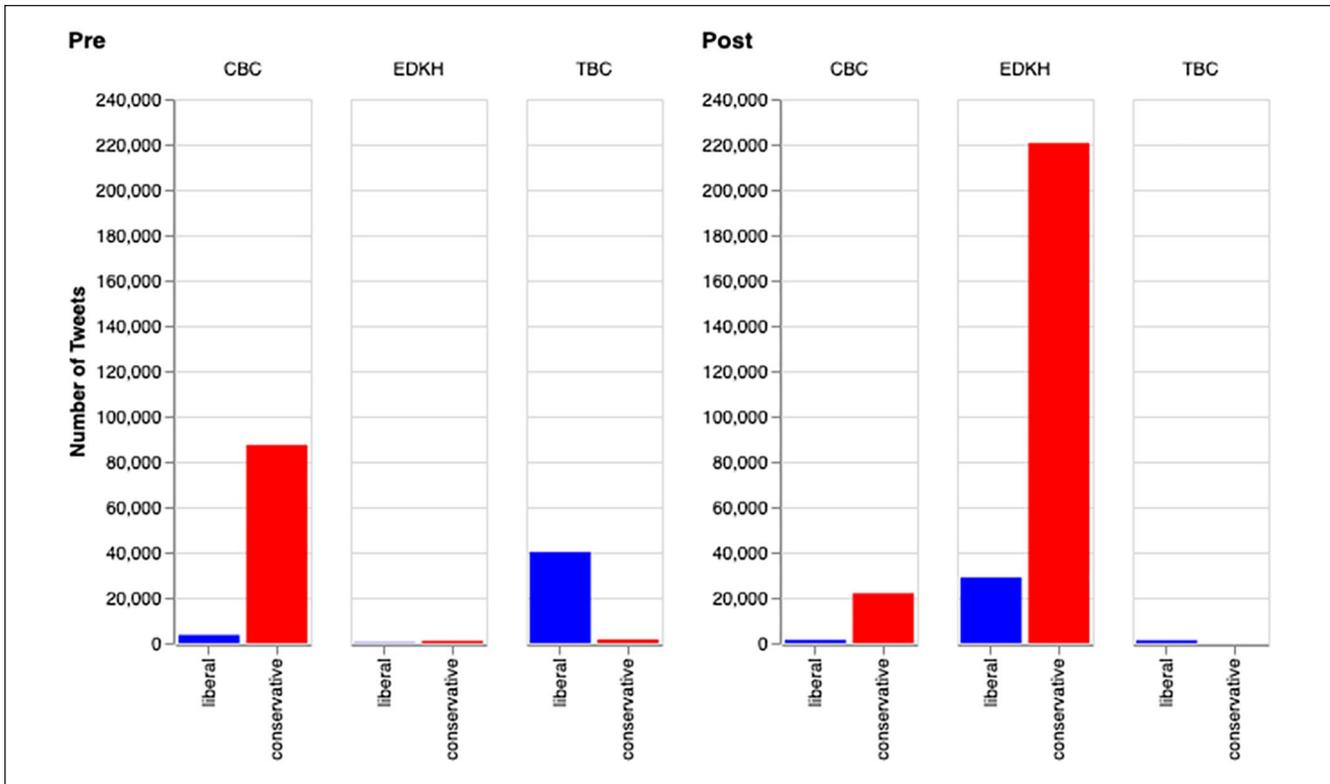


Figure 5. Rumor-spreading behavior by liberal and conservative Twitter users before and after the official correction for each of the three rumors about the death of Jeffrey Epstein.

Note: This figure shows the raw frequencies of tweets spreading each of the three major rumors sent by liberal and conservative Twitter users before and after the official correction.

before the correction period, and they were spread nearly eight times as often by conservatives after the attempt at correction was made. In both cases, conservatives—but not liberals—appeared to share rumors even more enthusiastically after they were officially debunked than before. Liberals, on the other hand, exhibited a more rational style of belief updating, sharing rumors less after the corrections were issued (Baron & Jost, 2019; Sinclair et al., 2020).

We do not mean to imply that liberals and leftists would never be highly motivated to engage in rumor-mongering or that conservatives and rightists are always more likely to engage in rumor-mongering. There are surely some issues cloaked in informational ambiguity that are higher in emotional intensity for leftists (such as John F. Kennedy's assassination or alleged Russian hacking of voting machines during the 2016 election). It is also crucial to consider the limitations of our data, which are mainly rooted in the common trade-offs that accompany conducting natural experiments on social media platforms.

To begin, it is well known that Twitter users are younger and more likely to be Democrats than the general public, so it is unclear how well our results generalize to the larger population (Pew Research Center, 2021). In addition, these data are extracted from a complex and constantly changing

information ecosystem outside of our control. As a result, it is possible that factors related to the ecosystem, and not ideological asymmetry, could be contributing to our results. For example, there is some evidence supporting the idea that the “market” for political misinformation may be ideologically skewed to the right in American politics (Allcott & Gentzkow, 2017). Thus, there may be an ideological asymmetry in exposure to corrective information (Grinberg et al., 2019), rather than motivated resistance to counter-attitudinal information in general. To ensure that corrections were widely observed by a broad cross-section of Twitter users, we selected two of the most salient rumors in American life over the past decade; however, it is difficult to know whether our case selection worked as intended.

On the other hand, it also may be possible that there is an ideological imbalance with respect to the influential suppliers of misinformation within this market. For example, then-President Donald Trump infamously retweeted a post containing conspiracy theories about the Clinton's being involved in Epstein's death (Vasquez, 2019). While this single incident took place before the official correction was provided about Epstein's suicide, our data does not allow us to rule out the possibility that such events might influence our findings.

It is also worth noting that, given the nature of this work, we were unable to control individuals' baseline propensity to spread the rumors we observed. For example, conspiracy theories surrounding the Clinton family have existed for many years ("Fact Check: Clinton Body Bags," 1998), so it may be plausible that conservatives would more readily adopt the CBC Epstein rumor. That said, this does not explain why conservatives more willingly spread the general EDKH rumor after it was debunked.

Importantly, these uncontrollable asymmetries may exist, at least in part, because of the psychological characteristics of conservative elites and their followers (Jost, 2021). That is, the market for misinformation may be ideologically skewed because of demand factors as well as issues of supply (Gries et al., 2022). These market differences, as well as any asymmetrical preexisting tendencies to spread rumors, may be driven by the fact that conservatives possess a more "intuitive" thinking style, are less prone to cognitive reflection and more prone to self-deception, in comparison with liberals (Jost, 2017). These differences in cognitive style may help to explain why conservatives are heavier sharers and consumers of "fake news," rumors, and conspiracy theories (Fessler et al., 2017; Guess et al., 2019, 2020; Mosleh et al., 2021; Sinclair et al., 2020; Van der Linden et al., 2021).

Another limitation is that we cannot be certain that tweeting behavior corresponds closely to endorsing specific rumors or, conversely, that a reduction in tweeting behavior reflects rational belief updating. It is possible that many people are simply tweeting about other events and have not corrected their initial misconceptions. Thus, the ideological asymmetry may have more to do with the thresholds for sharing behavior—and liberal and conservative social norms that influence sharing behavior—than with ideological differences in epistemic motivation or ability *per se*. The asymmetrical patterns we observed may stem from ideological divergence in gullibility, motivated reasoning processes, or willingness to pass on information that individuals know could be unreliable (or even false)—or to some combination of these processes (and perhaps others).

Finally, an obvious limitation of our research is that it is based on an analysis of rumors pertaining to only two major events in recent history. No doubt there will be ideological, cultural, and other sources of variability in the types of rumors that are spread and events that elicit rumor-spreading behavior. However, it is important to note that neither of the rumors we investigated was chosen because of their ideological significance. Instead, they were selected due to their widespread societal relevance, regardless of one's ideology. Future research on ideological symmetries and asymmetries would do well to expand the number and types of rumors under investigation.

Notwithstanding these limitations, our research highlights the ease with which inaccurate information can spread through social media platforms and the difficulties

associated with correcting misinformation in a real-world setting. Our studies illustrate the clear persistence of rumor and misinformation online even after official attempts at correction, and even when people may appear to update their beliefs initially. Public officials should continue to debunk political rumors. But this alone is not enough; citizens themselves must also be willing and able to change their minds and, more to the point, their online behavior. Not everyone, it would appear, is equally inclined—or equipped—to do so.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by a National Science Foundation award (#SES-1248077-001).

ORCID iDs

Matthew R. DeVerna  <https://orcid.org/0000-0003-3578-8339>

John T. Jost  <https://orcid.org/0000-0002-2844-4645>

Supplemental Material

Supplemental material is available online with this article.

Notes

1. A few other rumors that were frequently mentioned in addition to the three categories we focused on, but we did not analyze them separately because they fit the broader "Epstein Didn't Kill Himself" category. These included the "Prince Andrew" rumor (claiming that Epstein was killed by the Duke of York), the "Mossad Agent" rumor (that he was killed because was a secret agent for the Israeli national intelligence services), and the "Body Swap" rumor (that the body removed from the prison was not Epstein's).
2. Two research assistants agreed upon the labeling of 50 tweets concerning each of three rumors before labeling the remaining 4,500 tweets. Thus, they were able to discuss and resolve ambiguities in the coding procedure to label tweets in a consistent, principled manner.
3. <https://doi.org/10.7910/DVN/TYCTGN>
4. According to Wikipedia, the "Clinton Body Count" conspiracy theory has been in circulation since the 1990s, whereas the "Trump Body Count" is much more recent. The fact that the former conspiracy theory has been around longer than the latter could help to explain its greater popularity overall, but it is important to keep in mind that (a) we sampled tweets based on keywords pertaining to the Epstein case, not CBC, so that, the sample is very unlikely to contain many tweets pertaining to CBC that were unrelated to rumors about Epstein's death and (b) EDKH and the other rumors we investigated in both studies were at least as recent as TBC at the time we investigated them.

References

- Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, *31*, 211–236.
- Allport, G. W., & Postman, L. (1946). An analysis of rumor. *Public Opinion Quarterly*, *10*, 501–517.
- Azevedo, F., & Jost, J. T. (2021). The ideological basis of anti-scientific attitudes: Effects of authoritarianism, conservatism, religiosity, social dominance, and system justification. *Group Processes & Intergroup Relations*, *24*(4), 518–549.
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, *26*(10), 1531–1542.
- Baron, J., & Jost, J. T. (2019). False equivalence: Are liberals and conservatives in the U.S. equally “biased”? *Perspectives on Psychological Science*, *14*, 292–303.
- Benkler, Y., Faris, R., Roberts, H., & Zuckerman, E. (2017). Study: Breitbart-led right-wing media ecosystem altered broader media agenda. *Columbia Journalism Review*, *1*, 7. <https://www.cjr.org/analysis/breitbart-media-trump-harvard-study.php>
- Calvillo, D. P., Ross, B. J., Garcia, R. J. B., Smelter, T. J., & Rutchick, A. M. (2020). Political ideology predicts perceptions of the threat of COVID-19 (and susceptibility to fake news about it). *Social Psychological and Personality Science*, *11*(8), 1119–1128.
- Carney, D. R., Jost, J. T., Gosling, S. D., & Potter, J. (2008). The secret lives of liberals and conservatives: Personality profiles, interaction styles, and the things they leave behind. *Political Psychology*, *29*, 807–840.
- Crawford, J. T. (2012). The ideologically objectionable premise model: Predicting biased political judgments on the left and right. *Journal of Experimental Social Psychology*, *48*, 138–151.
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., & Quattrociochi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, *113*(3), 554–559.
- Di Fonzo, N., & Bordia, P. (2007). *Rumor psychology: Social and organizational approaches*. American Psychological Association.
- Ditto, P. H., Liu, B. S., Clark, C. J., Wojcik, S. P., Chen, E. E., Grady, R. H., Celniker, J. B., & Zinger, J. F. (2019). At least bias is bipartisan: A meta-analytic comparison of partisan bias in liberals and conservatives. *Perspectives on Psychological Science*, *14*(2), 273–291.
- Ecker, U. K., & Ang, L. C. (2019). Political attitudes and the processing of misinformation corrections. *Political Psychology*, *40*(2), 241–260.
- Evans, A., Slegers, W., & Mlakar, Ž. (2020). Individual differences in receptivity to scientific bullshit. *Judgment and Decision Making*, *15*(3), 401.
- Fact Check: Clinton Body Bags. (1998). Snopes.com. <https://www.snopes.com/fact-check/clinton-body-bags/>
- Fessler, D. M., Pisor, A. C., & Holbrook, C. (2017). Political orientation predicts credulity regarding putative hazards. *Psychological Science*, *28*(5), 651–660.
- Festinger, L., Cartwright, D., Barber, K., Fleischl, J., Gottsdanker, J., Keysen, A., & Leavitt, G. (1948). A study of a rumor: Its origin and spread. *Human Relations*, *1*, 464–486.
- Frimer, J. A., Skitka, L. J., & Motyl, M. (2017). Liberals and conservatives are similarly motivated to avoid exposure to another’s opinions. *Journal of Experimental Social Psychology*, *72*, 1–12.
- Graham, J., Haidt, J., & Nosek, B. (2009). Liberals and conservatives use different sets of moral foundations. *Journal of Personality and Social Psychology*, *96*, 1029–1046.
- Gries, T., Müller, V., & Jost, J. T. (2022). The market for belief systems: A formal model of ideological choice. *Psychological Inquiry*, *33*, 65–83.
- Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on Twitter during the 2016 US presidential election. *Science*, *363*(6425), 374–378.
- Gross, D. (2013, April 17). 5 viral stories about Boston attacks that aren’t true. *CNN*. <https://www.cnn.com/2013/04/16/tech/social-media/social-media-boston-fakes/index.html>
- Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, *5*(1), eaau4586.
- Guess, A., Nyhan, B., & Reifler, J. (2020). Exposure to untrustworthy websites in the 2016 US election. *Nature Human Behaviour*, *4*(5), 472–480.
- Hodson, G., & Busseri, M. A. (2012). Bright minds and dark attitudes: Lower cognitive ability predicts greater prejudice through right-wing ideology and low intergroup contact. *Psychological Science*, *23*, 187–195.
- Jerit, J., & Zhao, Y. (2020). Political misinformation. *Annual Review of Political Science*, *23*(1), 77–94.
- Jost, J. T. (2017). Ideological asymmetries and the essence of political psychology. *Political Psychology*, *38*, 167–208.
- Jost, J. T. (2021). *Left & right: The psychological significance of a political distinction*. Oxford University Press.
- Kraft, P. W., Lodge, M., & Taber, C. S. (2015). Why people “don’t trust the evidence” motivated reasoning and scientific beliefs. *The ANNALS of the American Academy of Political and Social Science*, *658*(1), 121–133.
- Lambert, A. J., & Chasteen, A. L. (1997). Perceptions of disadvantage versus conventionality: Political values and attitudes toward the elderly versus blacks. *Personality and Social Psychology Bulletin*, *23*, 469–481.
- Lazer, D., Baum, M., Grinberg, N., Friedland, L., Joseph, K., Hobbs, W., & Mattsson, C. (2017). Combating fake news: An agenda for research and action. <https://shorensteincenter.org/combating-fake-news-agenda-for-research/>
- Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, *13*(3), 106–131.
- Losee, J. E., Smith, C. T., & Webster, G. D. (2020). Politics, personality, and impulsivity can color people’s perceptions of—and responses to—hurricane threats of varying severity. *Personality and Social Psychology Bulletin*, *47*, 1435–1451. <https://doi.org/10.1177/0146167220969021>
- Marwick, A., & Lewis, R. (2017). Media manipulation and disinformation online. *Data & Society Research Institute*. https://datasociety.net/pubs/oh/DataAndSociety_MediaManipulationAndDisinformationOnline.pdf
- Miller, J. M., Saunders, K. L., & Farhart, C. E. (2016). Conspiracy endorsement as motivated reasoning: The moderating roles of political knowledge and trust. *American Journal of Political Science*, *60*(4), 824–844.

- Mosleh, M., Pennycook, G., Arechar, A. A., & Rand, D. G. (2021). Cognitive reflection correlates with behavior on Twitter. *Nature Communications*, *12*, 921.
- Nyhan, B., & Reifler, J. (2015). Displacing misinformation about events: An experimental test of causal corrections. *Journal of Experimental Political Science*, *2*(1), 81–93.
- Pennycook, G., & Rand, D. G. (2021). The psychology of fake news. *Trends in Cognitive Sciences*, *25*(5), 388–402.
- Pew Research Center. (2021). Public trust in government, 1958–2021. <https://www.pewresearch.org/politics/2021/05/17/public-trust-in-government-1958-2021/>
- Pfattheicher, S., & Schindler, S. (2016). Misperceiving bullshit as profound is associated with favorable views of Cruz, Rubio, Trump and conservatism. *PLOS ONE*, *11*, e0153419.
- Rosnow, R. L. (1991). Inside rumor: A personal journey. *American Psychologist*, *46*, 484–496.
- Rosnow, R. L., & Fine, G. A. (1976). *Rumor and gossip: The social psychology of hearsay*. Elsevier.
- Sasahara, K., Chen, W., Peng, H., Ciampaglia, G. L., Flammini, A., & Filippo, M. (2021). Social influence and unfollowing accelerate the emergence of echo chambers. *Journal of Computational Social Science*, *4*, 381–402. <https://doi.org/10.1007/s42001-020-00084-7>
- Schapiro, R., & Dienst, J. (2019). Jeffrey Epstein's former cellmate cleared of wrongdoing in incident that preceded his death. *NBC News*. <https://www.nbcnews.com/news/crime-courts/jeffrey-epstein-s-former-cellmate-cleared-wrongdoing-incident-preceded-his-n1043341>
- Sinclair, A. H., Stanley, M. L., & Seli, P. (2020). Closed-minded cognition: Right-wing authoritarianism is negatively related to belief updating following prediction error. *Psychonomic Bulletin & Review*, *27*, 1348–1361.
- Starbird, K., Maddock, J., Orand, M., Achterman, P., & Mason, R. M. (2014). Rumors, false flags, and digital vigilantes: Misinformation on twitter after the 2013 Boston marathon bombing. *Iconference 20 Proceedings*, *14*, 654–662.
- Sterling, J. L., Jost, J. T., & Pennycook, G. (2016). Are neoliberals more susceptible to bullshit? *Judgment and Decision Making*, *11*, 352–360.
- Swire, B., Berinsky, A. J., Lewandowsky, S., & Ecker, U. K. (2017). Processing political misinformation: Comprehending the Trump phenomenon. *Royal Society Open Science*, *4*(3), 1–21.
- Thorson, E. (2016). Belief echoes: The persistent effects of corrected misinformation. *Political Communication*, *33*(3), 460–480.
- Van der Linden, S., Panagopoulos, C., Azevedo, F., & Jost, J. T. (2021). The paranoid style in American politics revisited: An ideological asymmetry in conspiratorial thinking. *Political Psychology*, *42*(1), 23–51.
- Van der Linden, S., Roozenbeek, J., & Compton, J. (2020). Inoculating against fake news about COVID-19. *Frontiers in Psychology*, *11*, 2928.
- Vasquez, C. (2019). Trump amplifies Epstein conspiracy theories on social media. *Politico*. <https://www.politico.com/story/2019/08/10/trump-retweets-epstein-conspiracy-theories-1456305>
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, *359*(6380), 1146–1151.
- Vraga, E. K. (2015). How party affiliation conditions the experience of dissonance and explains polarization and selective exposure. *Social Science Quarterly*, *96*, 487–502.
- Yin, L. (2018). SMAPPNYU/urlExpander: Initial release.
- Zmigrod, L., Eisenberg, I., Bissett, P., Robbins, T. W., & Poldrack, R. A. (2021). The cognitive and perceptual correlates of ideological attitudes: A data-driven approach. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *376*, 20200424.