



AARHUS UNIVERSITY



Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the post-print version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Osmundsen, M., Bor, A., Vahlstrup, P. B., Bechmann, A., & Petersen, M. B. (2021). Partisan Polarization Is the Primary Psychological Motivation behind Political Fake News Sharing on Twitter. *American Political Science Review*, 115(3), 999-1015. <https://doi.org/10.1017/S0003055421000290>

Publication metadata

Title: Partisan Polarization Is the Primary Psychological Motivation behind Political Fake News Sharing on Twitter

Author(s): Mathias Osmundsen, Alexander Bor, Peter Bjerregaard Vahlstrup, Anja Bechmann, Michael Bang Petersen

Journal: *American Political Science Review*, 115(3), 999-1015

DOI/Link: <https://doi.org/10.1017/S0003055421000290>

Document version: Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

*Please cite the final version of the paper available at
American Political Science Review*

Partisan polarization is the primary psychological
motivation behind political fake news sharing on Twitter

Mathias Osmundsen^{*1}, Alexander Bor¹, Peter Bjerregaard Vahlstrup², Anja
Bechmann², and Michael Bang Petersen¹

¹Department of Political Science, Aarhus University

²School of Communication and Culture, Aarhus University

Abstract

The rise of “fake news” is a major concern in contemporary Western democracies. Yet, research on the psychological motivations behind the spread of political fake news on social media is surprisingly limited. Are citizens who share fake news *ignorant* and lazy? Are they fueled by sinister motives, seeking to *disrupt* the social status quo? Or do they seek to attack partisan opponents in an increasingly *polarized* political environment? This manuscript is the first to test these competing hypotheses based on a careful mapping of psychological profiles of over 2,300 American Twitter users linked to behavioral sharing data and sentiment analyses of more than 500,000 news story headlines. The findings contradict the ignorance perspective but provide some support for the disruption perspective and strong support for the partisan polarization perspective. Thus, individuals who report hating their political opponents are the most likely to share political fake news and selectively share content that is useful for derogating these opponents. Overall, our findings show that fake news sharing is fueled by the same psychological motivations that drive other forms of partisan behavior, including sharing partisan news from traditional and credible news sources.

*Corresponding author: m.osmundsen@ps.au.dk

With the advent of social media, the circulation of “fake news” has emerged as a major societal concern. Fake news can be defined as “fabricated information that mimics news media content in form but not in organizational process or intent” (Lazer et al. 2018). While some studies suggest most social media users refrain from sharing fake news (Guess, Nagler and Tucker 2019), others argue that false news spread faster and with greater reach than true news (Vosoughi, Roy and Aral 2018). Political observers have linked the spread of fake news to recent key political events, like the “Brexit” vote in Great Britain and the 2016 U.S. presidential election (e.g. Rose 2017, Silverman 2016). Accordingly, media organizations, governments, and academics have poured hefty investments into practices and technologies to stem the tide of fake news on social media. Notable efforts include fact-checking services like FactCheck.org and Snopes, as well as platform-based fake news detection algorithms and bots (Crowell 2017).

In light of these sizable investments, it is surprising how little we know about why people share fake news on social media. Studies suggest the spread of false information cannot be attributed only to bots or “sock puppet” accounts (Guess, Nagler and Tucker 2019, Vosoughi, Roy and Aral 2018). But researchers have largely failed to conduct adequate tests of *why* ordinary citizens take part in the distribution of fake news (Wagner and Boczkowski 2019). This is unfortunate because potential solutions to the spread of fake news rest on psychological assumptions about the root causes of the problem. For example, fact-checking should reliably work only if citizens actually want to share true information but need help to weed out falsehoods. Presently, we do not know whether this portrayal of citizens’ motivations is correct. As summarized in a recent study, “most existing studies rely solely on publicly available social media data and there is little research that seeks to understand people’s motivations for (...) sharing news” (Chadwick, Vaccari and O’Loughlin 2018, 4258).

In this study, we ask: What are the individual-level psychological motivations behind fake news sharing? The first major challenge in addressing this research question is data unavailability (King and Persily 2019, Stier et al. 2019). Data on citizens’ personality and political profiles are rarely available in combination with behavioral data on their sharing activities on social media (for notable exceptions, see Guess, Nagler and Tucker 2019, Guess,

Nyhan and Reifler 2020, Grinberg et al. 2019). Here, we circumvent these data limitations by linking unique behavioral data on news sharing activity and sentiment analyses of more than 500,000 news story headlines to individual-level survey data with detailed psychological measures for a large sample of American Twitter users. This approach allows us to conduct a comprehensive study of competing psychological theories of fake news sharing.

The second challenge is that multiple psychological motivations could potentially influence the sharing of false and factually correct news; these motivations may differ across individuals and across news content. Building on theories about information sharing in general, we first outline these potential psychological motivations before applying them to our research topic: the sharing of political fake news on social media, a domain characterized by intense conflict.¹ Then, we deduce testable implications about relevant individual differences and differences in political news content.

Psychological Motivations Behind News Sharing on Social Media

Cognitive psychologists have demonstrated the existence of two fundamentally different motivational goals that structure how individuals engage with information: Accuracy-oriented and goal-oriented motivations (the latter sometimes referred to as directional motivations) (Kunda 1990). Accuracy-oriented motivations spur individuals to engage with and accept only true and reliable information to form an accurate picture of the world. Goal-oriented motivations, in contrast, prompt people to focus on information that is *useful* given their pre-existing commitments to particular beliefs, groups, or outcomes. For decades, political scientists and communication researchers have applied this dual-motivation framework for studying how citizens engage with legacy media (e.g., newspapers and TV news shows). As audiences had no direct means to influence the content of legacy media channels, the bulk of this work has focused on selective exposure – the idea that political predispositions propel people to read and watch news from sources that share their political beliefs (Stroud 2010)

¹Accordingly, we acknowledge that people may have entirely different reasons for sharing fake news about other issues (e.g., sports, celebrities).

– and the downstream effects it can have on polarizing readers and viewers (Levendusky 2013).

The emergence of the internet, and particularly of social media has fundamentally changed the relationship between media and audience. Whereas once audiences were more or less passive consumers, they are now actively participating in the distribution, and sometimes even the production of news. People today can start their own blogs, live-stream from political events they participate in (e.g. protests), and may even contribute to small media organizations often catering to niche market segments. To be sure, just as with legacy media, online audiences still face a choice of exposure: which news do they choose to consume. But on top of that, they also face a new choice of which news to *share* with their friends and followers. These trends, we contend, have implications for the psychological dynamics underlying the way audiences engage with this media, the novel step of news sharing potentially recruiting new psychological mechanisms.

These trends have pushed to the forefront questions about information veracity. In advanced democracies, strong editorial norms and standards for good reporting meant that citizens could – at least most of the time – trust legacy media outlets to publish stories that accurately reflected the world. For most citizens and for most legacy media, then, it made sense to pay more attention to the political *usefulness* of news stories and, thus, it is not surprising that this news dimension has been critical in driving selective exposure. Yet, on the internet, editorial norms have lost power, and information circulating on social media will to a much larger extent vary on both dimensions of usefulness and veracity. Truth cannot be taken for granted and veracity judgments are thus, potentially, more important than ever.

Yet, these two information dimensions – veracity and usefulness – are not unique for social media users. In everyday conversations when, for example, discussing or relaying a rumor, people must consider both whether they believe it and how useful they think it is. Psychologists have studied the psychology of rumor transmission for decades, highlighting the importance of both accuracy-oriented and goal-oriented motivations (for an overview, see DiFonzo and Bordia 2007). For example, people may share rumors they believe are true to be cooperative, to signal competence, or to get others to affirm the accuracy of the rumors.

Sharing a rumor that turns out to be wrong, in contrast, may damage people’s reputations (Altay, Hacquin and Mercier 2019). Nonetheless, goal-oriented processes may still take precedent, notably under intense social conflict. Some of the most well-studied, and most extreme, examples come from ethnic conflicts, where rumors are heavily used to mobilize the in-group against the out-group. In his comprehensive review of this role of rumors, Horowitz (2000, p. 74) concludes: “Rumors are (...) embedded in the riot situation, because they are satisfying and useful to rioters and their leaders. (...) Rumor is likely to prevail over accurate information. (...) Rumor prevails because it orders and organizes action-in-process.”

Below we apply these general considerations about the psychology of information sharing to the object of our empirical inquiry: The sharing of news – fake or real – about politics on social media.

Accuracy-oriented and goal-oriented motivations in the sharing of fake and real news about politics on social media

The prevalence of false information on social media may strengthen the role of accuracy-oriented motivations and the dimension of veracity for understanding news sharing behavior. People may worry about damaging their reputation by accidentally sharing fake news (Altay, Hacquin and Mercier 2019). To the extent such reputational worries induce accuracy-oriented social media behavior, we should expect that sharing fake news stories is a rare phenomenon relative to sharing real news, but also that the predictors of sharing fake versus real news will differ. Successfully discerning between fake and real news is difficult, first, because fake news is specifically designed to appear real and, second, because the fast-paced social media platforms deplete attention and cognitive resources (Pennycook et al. 2020). Accordingly, individual differences in the ability to detect cues about information veracity should largely determine whether people share fake news (Pennycook and Rand 2019b). In essence, from this perspective, fake news sharing happens because people are *ignorant* that the stories they share are false and, accordingly, below we refer to this argument as the *ignorance theory*.

Goal-oriented motivations, however, may also critically shape sharing of political news on social media. As noted above, goal-oriented motivations affect information sharing during

periods of social conflict. Politics is conflictual, but political conflict in the United States in recent years has been deepened by an increasing political polarization (Mason 2018). Goal-oriented motivations put a premium on the *usefulness* of information. Sharing a political news story (fake or real) can be used to signal one’s political affiliation, by denigrating political opponents or by mobilizing politically like-minded peers. In this regard, fake news stories can be particularly useful: as fake news is not constrained by reality, they can be excessive in their negative portrayal of political opponents (Acerbi 2019). Accordingly, to the extent goal-oriented motivations dominate social media users’ interactions with news stories, the same psychological motivations should drive the sharing of real and fake news. People will share useful news that fits their political goals, and, in the extreme case, they will not pay attention to the veracity dimension of the information at all.

In current American politics, two types of social or political goals could be served by sharing extreme information such as fake news. First, sharing news stories could reflect motivations to help one’s preferred party against opposing parties in the increasingly polarized electoral competition (Iyengar, Sood and Lelkes 2012, Lelkes and Westwood 2017). We will refer to this as *the polarization theory* as it focuses on partisan feelings rather than ideological disagreements between partisans of different stripes. The theory is prominent in current discussions about political events, like elections and referenda, and it has a long research history in political science (Mason 2018, Taber and Lodge 2006). The theory has also informed studies on partisans’ beliefs in political conspiracies (Miller, Saunders and Farhart 2015).

Second, sharing news stories could be fueled by deep-seated discontent with the status quo and serve a desire to disrupt the existing social and political order. For example, Tucker et al. (2017, 53) convincingly argue that “social media can lend a voice to ... anti system forces that actively seek to undermine liberal democracy.” Below, we will refer to this as *the disruption theory*.

Finally, there are also multiple scenarios for the intermediate (and likely) case where both veracity and usefulness considerations are relevant for social media users’ sharing behavior. People may first evaluate the veracity of news and then, conditional on it being accurate,

consider its usefulness. Or people may prioritize usefulness but still pay some attention to the veracity. Hence, the sharing of blatantly false information may, for example, complicate mobilization processes by derailing discussions about whether events really happened (Petersen, Osmundsen and Tooby 2020). Thus, when assessing usefulness, people may factor in veracity and selectively share information that is as extreme as they believe they can get away with.

To elucidate whether and how the ignorance, polarization, and disruption theories account for the sharing of fake news, we use two approaches: First, we examine the role of individual differences, deducing for each theory *who* should be most likely to share fake news and, second, we examine the role of news story content, deducing *what* is most likely to be shared.

The role of individual differences in the sharing of news

The accuracy-oriented *ignorance theory* assumes people want to share accurate information, but end up sharing falsehoods because they lack the cognitive reflection or motivation to discern between true and false information. One line of evidence for the ignorance theory comes from Pennycook and Rand (2019b), who show that people who perform better on a cognitive reflection test can better tell if news story headlines are true or false. Another line of evidence comes from Guess, Nagler and Tucker (2019), who demonstrate that older people are more likely to share fake news on Facebook, presumably because they are less digitally literate than their younger peers. Building on this, we test the *ignorance theory* by focusing on four individual-level predictors of fake news sharing on Twitter: age, cognitive reflection, factual political knowledge, and digital literacy.

The goal-oriented *polarization theory* focuses on partisanship as a key driver of fake news sharing. Currently, there is mixed evidence on how partisan motivations influence fake news sharing. In general, partisans seem to share information from ideologically similar sources (Barberá et al. 2015) and to treat articles shared by political opponents with greater suspicion (Lee, Kim and Coe 2018), suggesting that considerations about the political usefulness of information matters. Guess, Nagler and Tucker (2019, 3) find that partisanship also affects

fake news sharing in that conservatives and Republicans were more likely to share fake news on Facebook in 2016 than liberals and Democrats, a partisan asymmetry attributed to “the overwhelming pro-Trump orientation in both the supply and consumption of fake news during that period.” Pennycook and Rand (2019b), however, tone down the effect of partisanship. They conclude that “people fall for fake news because they fail to think; not because they think in a motivated or identity-protective way” (ibid., 48).

To test the *polarization theory*, we examine whether Democratic partisans are willing to share politically congruent fake news from *pro*-Democratic sources or if the partisanship asymmetry identified by Guess, Nagler and Tucker (2019) and others prevails. Second, an important question in the literature on affective polarization is whether in-party love (Lelkes and Westwood 2017) or out-party hatred (Abramowitz and Webster 2018) sparks polarization. We provide the first test of which of these dynamics explain fake news sharing by relying on two types of individual differences as predictors of news sharing: 1) standard party identification measures and 2) batteries of positive versus negative feelings towards Democrats and Republicans.

Finally, we turn to the goal-oriented *disruption theory*. A central focus in media discussions and psychology has been on so-called “online trolls” who take pleasure in misleading and harassing others online (Buckels, Trapnell and Paulhus 2014). Trolling experts speculated that sharing sensational and often implausible fake news reflects a specific form of *apolitical* online trolling. However, *political* cynicism may also drive fake news sharing. For example, a recent study showed that Americans and Danes who are discontent with political elites and the political system report substantially greater willingness to circulate conspiracy theories online (Petersen, Osmundsen and Arceneaux 2018). Here, we provide the first behavioral test of this theory and zoom in on two individual differences as predictors of fake news sharing: a psychological measure of trolling behavior and a measure of political cynicism.

The role of content in the sharing of news

To elucidate the psychological motivations underlying fake news sharing, we also need to understand what news content people share. Specifically, to adjudicate between the three

theories, the analyses differentiate news stories along two dimensions: (1) Whether a news story originates from a fake news source (i.e., a source containing registered debunked stories) or a real news source (i.e., a generally credible source), and (2) whether a news story comes from a *pro*-Republican or a *pro*-Democratic source. The three theories make very different predictions about which dimension matters for people's sharing preferences. The ignorance theory predicts that the veracity dimension matters most and that people with high cognitive abilities will be less willing to share fake news stories and more likely to share real news stories, irrespective of the political slant of the stories. The latter two theories make the opposite prediction: they posit that people place less value on whether stories are true or false as long as the stories are politically useful. Specifically, the disruption theory implies that disruptive individuals should selectively share as many fake news as possible to sow confusion and uncertainty. For the polarization theory, however, the content-related predictions are more complex, and so it is these predictions we unfold now.

The polarization theory holds that fake news sharing is not an accident caused by ignorance, it is partisan business as usual. The media landscape in the US is politically fragmented, with some news sources catering to Republicans while others offer content that fit Democrats' taste (Groeling 2013). From the perspective of the polarization theory, fake news is an extreme version of traditional partisan news that polarized partisans may find attractive due to its rarity and extremity. Accordingly, partisans motivated to selectively share hyper-partisan content will tend to view the online media landscape of news sources as a one-dimensional continuum reflecting the political usefulness of the stories they provide. For a polarized Republican social media user, the order of the usefulness of news sources goes from *pro*-Republican fake news sources, to *pro*-Republican real news sources over centrist sources to *pro*-Democratic real news sources, and finally, to *pro*-Democratic fake news. For Democrats, the order reverses.

This insight, we contend, leads to three observable implications about the content of shared news. First, partisan social media users should share news sources with similar political leanings, especially near the partisan extremes of the news continuum. Accordingly, we expect that sharing fake news favoring one party over the other should be best predicted by

sharing real news strongly favoring the same party. Conversely, sharing fake news favoring one party should correlate negatively with sharing fake news favoring the other party. Second, if news sources differ systematically in how strongly they cater to partisan motivations, then the association between partisanship and sharing should be stronger for news sources located at the extremes of the news source continuum. Specifically, fake news should show the strongest partisan sorting in terms of who wants to share from these sources. Third, and finally, the negativity of news source content should change as we move across the partisan news continuum. If partisans selectively share stories with content that fits their political tastes, the observable implication is that stories shared from news sources located towards the Democratic extreme on the partisan news continuum should portray Republican elites negatively while stories from the Republican extreme should cast Democratic elites in a negative light.

Data and Methodology

Data

We commissioned the YouGov survey agency to recruit U.S. Twitter users to complete a 20-minute online survey fielded from mid-December 2018 to mid-January 2019. The survey contained questions measuring our set of potential predictors of fake news sharing (see below) and a range of demographic variables. The survey also asked participants to provide their Twitter ID and asked for permission to scrape their publicly available Twitter data and link this data to their confidential survey responses. See SI Section 1 for the invitation letter and link to the consent form. We scraped respondents' Twitter accounts by early March, 2019. Our data includes tweets and retweets posted by the panelists before that period.

YouGov invited over 27,000 U.S. Twitter users from their standing web panel to participate in our study. 8,741 participants accepted the invitation and answered all questions before the final prompt to give consent to link their public Twitter data to the survey. A majority of these respondents (5,725 people) did not share their Twitter IDs at the end of the survey. Of the 2,976 respondents who did attempt to share their Twitter handles, another 639

participants presented IDs that we could not access via Twitter’s Application Programming Interface. Consequently, our final sample size is $N = 2,337$.

Because of high and non-random attrition, our sample is unlikely to perfectly reflect any well-defined population. Still, our study participants’ demographic characteristics resemble the characteristics obtained from participants in other high-quality Twitter studies. Relying on a recent Pew study (Wojcik and Hughes 2019), SM Section 2a demonstrates that our sample is comparable to the Twitter-population on race, gender, income, education, and age. It also shows that participants in our final sample were demographically very similar to the 8,741 participants who accepted to participate in the survey but failed to provide their Twitter handle. In our final sample, 50% of participants are female, 70% are white/Caucasian, the median age is 40 years old ($SD = 14$), the median highest level of completed education is “4-year college degree,” and the median family income level is “\$60,000-\$69,999.”

We scraped Twitter data from all our 2,337 participants. In total, we extracted 2,709,052 tweets and retweets posted by our panelists until mid-January 2019. The median number of (re)tweets for our participants was 488 (Min #of tweets = 1, Max #of tweets = 6433). Importantly, a large majority of the tweets - around 70% - were posted after the 2016 US presidential election; the period that has been the focus of previous work. This allows us to extend earlier findings and test how they hold up in periods characterized by less intense electoral competition.

Fake news and real news

We follow best practices (Guess, Nyhan and Reifler 2020, Grinberg et al. 2019, Lazer et al. 2018) and operationalize “sharing of fake news sources” as when panel members tweet or retweet URLs to external web domains known to publish factually incorrect news. To this end, we cross-referenced all tweeted URLs against a list of 608 fake news sources, constructed by journalists and scholars (for details, see SM Section 10 and Guess et al. 2019, Allcott, Gentzkow and Yu 2019). The list was compiled to document the prevalence of fake news during the 2016 U.S. presidential election but has been updated to account for fake news publishers that have emerged since. To the extent possible, the list creators have coded

the most popular fake news sources for partisan slant: 12 of the most popular were *pro*-Democratic fake news sources while 30 were *pro*-Republican sources. While this number may seem low, note that almost 86% of all shared fake news sources originated from one of these 42 domains for which we have partisan ratings. (See SM Section 6 for the complete list of fake news domains and Figure SM 2b for the 20 fake news domains that were shared the most.) As a robustness test, supplementary analyses discussed below group the fake news sources according to a more fine-grained coding scheme proposed by Grinberg and colleagues (2019) that distinguishes among three classes of fake news sources differing in their likelihood of publishing misinformation. In SM Section 8, we also replicate the main results using a different, but partially overlapping list of fake news sources.

Our measure of fake news sharing records actual online behavior and thus does not suffer from biases that can plague self-reported measures of online activity. Further, while the present study is certainly not the first to couple survey and social media data (e.g., Guess, Nagler and Tucker 2019), it has the most detailed set of psychological measures; see below. Still, it is important to keep in mind when interpreting our results that we cannot tell if the specific stories people share are true or false: Our fake news measure is at the publisher level, not the story level.

To measure the “sharing of real news sources,” we cross-referenced participants’ tweets against a list of real news publishers obtained from the *AllSides* organization. *AllSides* seeks to help citizens navigate the online media landscape by providing “balanced news and civil discourse.” Their web page (www.allsides.com/media-bias/) contains a list of 260 real news publishers with ideological ratings: “Left” (labeled “Strong Democratic” below, $n = 45$), “Lean Left” (“Leaning Democratic”, $n = 63$), “Center/Mixed” ($n = 88$), “Lean Right” (“Leaning Republican”, $n = 29$) or “Right” (“Strong Republican”, $n = 39$). While *AllSides* combines different methods for estimating ideological bias (e.g., editorial reviews, third-party judgments), SM Section 7 shows that their ratings correlate highly – Pearson’s $r = .90$ – with network-based measures of ideological alignment of media sites (Bakshy, Messing and Adamic 2015). See SM Section 7 and Figure SM 2b for the 20 most shared real news domains.

Our measures of real and fake news sharing account for ideological and partisan slant at the

source level and are thus useful for understanding the content people care about when sharing news. To gain further leverage on testing the relevance of content for sharing behavior, we also built two unique data sets of news headlines (see SM Section 11 and 12 for details). We focus on headlines because these are more accessible than full texts and because prior research demonstrates that simply seeing a headline can exert large effects on decisions to share news on social media (Effron and Raj 2019). The first data set was based on the headlines of stories shared by our participants, which gives us a clear representation of the content they deemed sufficiently important to share. To construct the data set, we first retrieved the full URLs from all 85,084 tweeted news links shared by our participants. Based on these URLs, we parsed the news domains’ HTML code and successfully extracted the headlines of 75,560 news articles (89% of all article headlines).

To construct our second data set, we used the Archive.org web page – an indexing system for the Internet’s web pages for historical archiving – to scrape headlines that had appeared on the daily front pages of the most popular (among our participants; see Figure 5 below) news sites in the period 2016-2019 (i.e., covering roughly the same period for which we have tweets from our participants). Compared to our Twitter headlines, this data set contains a much larger number of news story headlines (we successfully scraped a little over 500,000 headlines). More importantly, the front page headlines offer a glimpse of the *universe* of news stories our participants could potentially see and read before deciding which stories to share. In this way, it sidesteps issues of “selection” biases in sharing decisions due to psychologically motivated reasoning; a point we return to during the analyses.

We subjected these headlines to two types of automated analyses. To measure the sentiment of the headlines, we relied on the sentimentR R package, which includes a sophisticated algorithm for calculating text polarity sentiment. This algorithm improves the performance of more traditional dictionary-based methods by calculating sentiment on the sentence level and considering valence shifters (e.g. not, very, hardly). To identify the political affiliations of those mentioned in the headlines, we created two custom dictionaries of prominent Republican and Democratic elites. Our primary dictionary is based on YouGov’s list of the

100 most famous Democrats and Republicans.² We replicate our results with an alternative dictionary listing all members of the US Congress, Senate and the Trump Administration in the supplemental materials. Finally, we matched these dictionaries to the headline data sets and recorded whether a given headline mentions anyone from the Republican or Democratic party.

We use these more detailed content measures in the final part of the analysis, whereas the initial parts focus on the crude distinction between fake and real news sources.

Survey measures: predictors of fake news sharing

All participants responded to the same survey questions, designed to measure the potential predictors of fake news sharing derived from the three theories discussed earlier: the ignorance theory, the disruption theory, and the polarization theory. SM Section 1c lists the questions asked and specific wordings.

The ignorance theory. We included four measures to test the ignorance theory. First, we included the alternate form of the Cognitive Reflection Test (Thomson and Oppenheimer 2016, CRT-2;), a 4-item battery designed to measure the propensity to override intuitive responses that are, upon reflection, wrong (e.g., “If you’re running a race and you pass the person in second place, what place are you in?” Intuitive response: 1st place; Correct response: 2nd place). Better performance on the CRT correlates positively with the ability to discern political fake news from real news (Pennycook and Rand 2019b). Second, we included participant age to test if we could replicate earlier findings that older social media users share more fake news stories than younger users. Finally, we included measures of factual *political* knowledge and digital media literacy to test the claims about why older social media users share more “fake news.” We included a 5-item battery from the American National Election Studies (ANES) to measure factual political knowledge (e.g., “How many members of the U.S. Supreme Court are there?”). To measure digital media literacy in the domain of politics, we created a 10-item battery based on previous research (Feezell 2016, Gil

²Accessible at <https://today.yougov.com/ratings/politics/fame/Democrats/all> and <https://today.yougov.com/ratings/politics/fame/Republicans/all>

de Zúñiga, Molyneux and Zheng 2014) asking participants how often they engage in specific online political behaviors (“Use the Internet to participate in volunteer activities related to a political or social campaign”) ($\alpha = .93$). As this measure also taps political engagement and online civic skills, we adjust for political interest in the analyses. In addition, SM Section 4 replicates the main analysis using a more narrow, validated measure of digital literacy from Hargittai and Hsieh (2011).

The disruption theory. We included two measures of disruptive motivations. To measure apolitical trolling behavior, we solicited the “cyber-bullying” literature and implemented a 4-item trolling scale (Buckels, Trapnell and Paulhus 2014), asking participants whether they agreed or disagreed with statements like “I like to troll people in forums or the comments section of websites” ($\alpha = .86$). We measured political cynicism via an 8-item survey instrument (Dekker and Meijerink 2012), asking participants whether they agreed or disagreed with statements like “Politics in the United States is sick” ($\alpha = .74$).

The polarization theory. We included two measures of political polarization: Partisanship and feelings towards Republicans and Democrats. We used a standard ANES question to measure partisanship on a seven-point scale, ranging from “Strong Democrat” to “Strong Republican” identifier. To measure partisan feelings, we asked participants how strongly they felt various positive (hope, enthusiasm, pride) and negative (anger, frustration, fear) emotions when thinking about Republicans and Democrats. We created two feelings scales to test how in-party love and out-group hate related to news sharing by aggregating responses and taking the averages for each participant (the negative emotion items were flipped). The key models presented below analyze motivations to share news from *pro*-Republican and *pro*-Democratic news sources separately. Accordingly, we coded the “in-party” emotion scale such that higher values always correspond to more positive emotions towards the political side supported by the news source (i.e., higher values indicate the participant has more positive feelings towards Republicans (Democrats) in the analyses of sharing *pro*-Republican (*pro*-Democratic) sources). Analogously, we constructed the “out-party” emotion scale such that higher values indicated more negative feelings towards the political side *opposed* by the news source (i.e., higher values indicate the participant has more nega-

tive feelings towards Democrats (Republicans) in the analyses of sharing *pro*-Republican (*pro*-Democratic) sources).

SM Section 2b presents summary statistics for all main variables.

Results


Do people share news from fake news sources?

Before testing the key hypotheses about the psychological motivations for fake news sharing, we provide descriptive analyses of the data. Overall, we find that news sharing is a small part of our sample’s activity on Twitter. (See Figure SM2c for an overview.) Of the 2.7 million tweets in our data, only 85,344 ($\approx 3\%$) contain links to national news websites, fake or real. Of the tweeted news links, 3,269 ($\approx 4\%$) come from websites with a history of publishing fake news (see Grinberg et al. 2019, for slightly larger estimates during the 2016 US presidential election). Of these stories, a large majority – 2,563 stories – originates from *pro*-Republican fake news publishers. In contrast, over twice as many real news stories come from *pro*-Democratic sources (60%) than from *pro*-Republican sources (23%). Together, these two observations imply that exposure to fake news depends heavily on the ideological leanings of one’s Twitter network. In *pro*-Democratic networks, only a small fraction of the news diet is likely to come from fake news sources, whereas *pro*-Republican networks are likely to display a substantially higher fake-to-real news ratio.

While these aggregate numbers suggest fake news sharing could be widespread within certain social networks, they also mask large heterogeneity in sharing behavior. The left-hand panel in Figure 1 below makes clear that sharing content from fake news sources flourishes mostly among small segments of the population: 1% of the panelists are responsible for sharing about 75% of all links to fake news publishers and 11% of them are responsible for sharing *all* fake news sources. Meanwhile, the top 1% of real news “supersharers” are only responsible for sharing about 30% of all real news sources. Figure 1’s two other panels show similar results when splitting the data based on the partisan slant of the news sources.

Importantly, not all fake news is equally troublesome. SM Section 8 uses Grinberg and colleagues' (2019) fine-grained categorization of fake news sources and find that about 30% of the fake news links come from sources “publishing almost exclusively fabricated stories” (i.e. “black” sites). Another 5% come from sites that “spread falsehoods that clearly reflected a flawed editorial process” (i.e. “red” sites). Finally, two-thirds of the links referred to fake news sources “where annotators were less certain that the falsehoods stemmed from a systematically flawed process” (i.e. “orange” sites). While this means that the number of individual stories containing blatant falsehoods is likely smaller than 3,325 in our data, we note that sharing news stories from websites with a record of disseminating falsehoods can still be problematic. Given the still limited penetration of fact-checking, sharers most likely do not know whether a certain news story is true or false when deciding to share it. Instead, they must make informed guesses based on the reputation of the news source. Thus, decisions to share news from these sources still reflect either ignorance or lack of concern about such matters.

Figure 1: Empirical cumulative distribution functions showing distribution of shared news links by study participants and content type



Note. The x-axes show percentage of participants responsible for sharing a given percentage of all news links (y-axes). The left-hand panel gives results for sharing links to fake (black line) versus real (grey line) news publishers ($N_{\text{Fake News}} = 3,269$; $N_{\text{Real News}} = 78,400$). The middle and right-hand panels display results for sharing links to *pro*-Democratic (blue line) versus *pro*-Republican (red line) fake and real news sources, respectively. For real news sources, we collapse the “Leaning” and “Strong” categories. The dashed black lines denote a hypothetical uniform distribution of news sharing. ($N_{\text{Fake News, pro-Rep}} = 2,563$; $N_{\text{Fake News, pro-Dem}} = 263$; $N_{\text{Real News, pro-Rep}} = 16,081$; $N_{\text{Real News, pro-Dem}} = 46,385$.)

What motivates sharing of fake and real news?

We now turn to the analysis of the psychological motivations for sharing fake news sources. Figure 2 gives average marginal effects (AME)³ from logistic regression models estimating the association between our main predictors and the probability of sharing at least one news story from a fake news (left panel) or a real news source (right panel). To increase power and

³Average marginal effects give the change in predicted probability of sharing one or more news story as our main predictors change from their minimum to maximum values, averaged across participants’ observed values of all other covariates in the model. See SM Section 3 for additional details on the logistic regression models.

simplify the analyses, we here combine the real news sources that *Allsides* rate as strongly partisan and leaning partisan (e.g., we combine news stories from “Strong Democratic” and “Leaning Democratic” real news sources). As our main interest is in whether people share fake news or not, we dichotomize the dependent variables, with 1 indicating that a participant shared at least one news story from a fake or real news source. The figure shows results for models examining the probability of sharing *pro*-Democratic (***pro*-Dem sources**) and *pro*-Republican sources (***pro*-Rep sources**) separately. We estimate separate models for each of our main predictors, where each model controls for gender, income, education, race and political interest. The middle panel, which also gives the probability of sharing stories from fake news sources, additionally control for the (logged) number of shared stories from real news sources. We scale all the predictors to range from 0 to 1. We code all predictors such that positive coefficients in Figure 2 express support for the respective theories.

The upper panels of Figure 2 examine the *ignorance theory*. Like others (e.g., Guess, Nagler and Tucker 2019), we find in the left-hand corner that older panelists are much more likely to share content from both *pro*-Republican and *pro*-Democratic fake news sources. Moving from the youngest (18-22 years) to the oldest (83+ years) age group increases the probability of sharing at least one story from a *pro*-Republican fake news source by 15 percentage points ($AME_{*pro*-Rep\ sources} = .15$, 95% CI: [.10 – .20]) and the probability of sharing a *pro*-Democratic fake news source by 8 percentage points ($AME_{*pro*-Dem\ sources} = .08$, 95% CI: [.4 – .23]); associations that remain robust when controlling for real news sharing in the middle panel. At the same time, as the right-hand panel shows, age also strongly predicts sharing content from both types of real news sources ($AME_{*pro*-Rep\ sources} = .23$, 95% CI: [.15 – .31]; $AME_{*pro*-Dem\ sources} = .08$, 95%: [-.01 – .18]). It appears, then, that older people promiscuously share all types of news on social media.

However, poor reasoning does not seem to explain the “age effect.” While better performance on the cognitive reflection test (CRT) is weakly associated with sharing real news sources, the association between CRT and sharing either type of fake news sources is insignificant at conventional levels and substantially small ($AME_{*pro*-Rep\ sources} = -.02$, 95% CI: [-.06 – .01]; $AME_{*pro*-Dem\ stories} = -.01$, 95% CI: [-.03 – .02]). The other two explanations do not fare bet-

ter. Participants *lowest* in political digital media literacy are the *least* likely to share content from fake news sources ($AME_{*pro*-Rep\ sources} = -.10$, 95% CI: $[-.06 - -.05]$; $AME_{*pro*-Dem\ sources} = -.08$, 95% CI: $[-.12 - -.04]$). Low political awareness also predicts less sharing from fake news sources ($AME_{*pro*-Rep\ sources} = -.04$, 95% CI: $[-.09 - -.01]$; $AME_{*pro*-Dem\ sources} = -.08$, 95% CI: $[-.13 - -.03]$), although the association is reduced considerably when controlling for real news sharing. Collectively, these findings do not support the claim that people fall prey to fake news because they are unable to discern true from false. Instead, fake news sharers are more knowledgeable about politics, regularly use the Internet for political activities, but are neither more nor less reflective than others.

What, then, motivates fake news sharing? The middle panels offer mixed support for the *disruption theory*. We find little evidence that fake news sharing is part of the toolkit of trolls who “just want to have fun” (Buckels, Trapnell and Paulhus 2014): Moving from the lowest to the highest value of trolling behavior in the left-hand panel *decreases* the probability of sharing content from both **pro*-Republican* and **pro*-Democratic* fake news sources by around 7 percentage points (p 's < .05). Trolls are also less inclined to share real news content (p 's < .05). However, we find considerable evidence that fake news sharing reflects disruptive sentiments aimed at the political domain. Changing from the least to the most politically cynical participant increases the probability of sharing both **pro*-Democratic* and **pro*-Republican* fake news sources by about 6-10 percentage points (p 's < .05). Political cynics are also more likely to share content from real news sources, irrespective of the news sources' political bias (p 's < .05). Overall, the results show that apolitical trolling is associated with withdrawal from news sharing while political cynicism is associated with the sharing of information that may contain falsehoods about both Democratic and Republican elites.

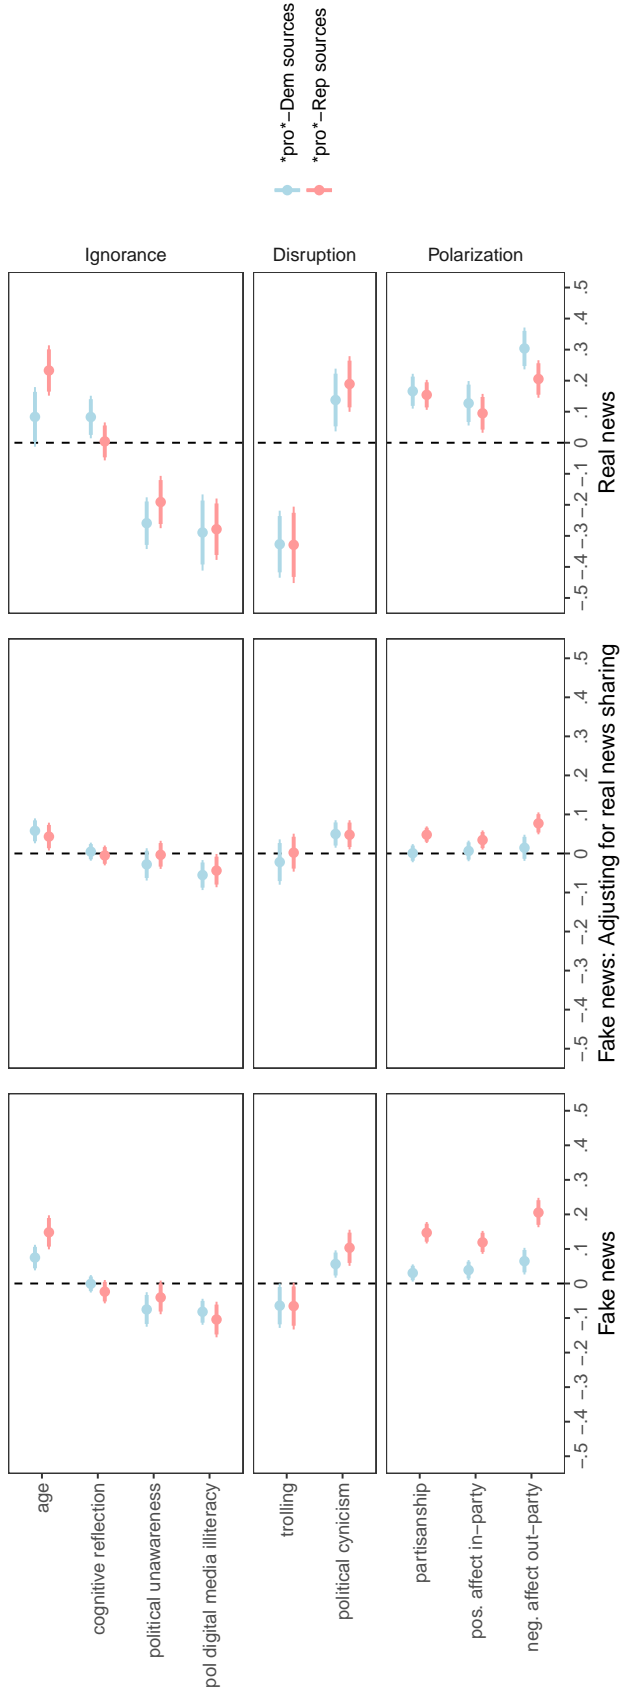
The lower panels of Figure 2 test the *polarization theory*. They suggest that partisan motivations are key to understanding fake news sharing. Turning first to partisanship – coded so that higher values indicate identification with the political party supported by the news source – we find that the probability of sharing fake news sources increases substantially among strong party identifiers. This association is strongest among Republican identifiers linking to **pro*-Republican* fake news sources ($AME_{*pro*-Rep\ sources} = .15$, 95% CI: $[.12 -$

.18]; $AME_{*pro*-Dem\ sources} = .03$, 95% CI: [.01 – .06]); a finding that is consistent with Guess, Nagler and Tucker (2019). Next, we find support for the “negative partisanship” prediction (Abramowitz and Webster 2018). The negative out-party affect coefficients are about twice as large in size as the positive in-party affect coefficients, suggesting that *animus towards political opponents* rather than positive feelings towards the in-party drives sharing of fake news sources. Again, we observe a partisan asymmetry: The association is largest among individuals who dislike Democrats and share content from *pro*-Republican fake news sources ($AME_{*pro*-Rep\ sources} = .21$, 95% CI: [.16 – .25]; $AME_{*pro*-Dem\ sources} = .06$, 95% CI: [.03 – .10]).

As shown in the lower right-hand panel, the partisan factors predict sharing from real news sources in an almost identical fashion. The motivations underlying fake news and real news sharing are, psychologically speaking, highly overlapping. One difference compared to sharing of fake news sources is that our estimates are obtained with greater precision due to the larger number of participants sharing real news sources. Another difference relates to the partisan asymmetry: Concerning real news sharing, we find no asymmetry between Democrats and Republicans; both are equally partisan and strongly motivated by negative feelings towards political opponents.

Altogether, these results fit well with predictions by theories on partisan motivations: Partisans share politically congenial news, primarily because of hostile feelings towards the out-party. Further, the lack of evidence for the ignorance theory suggests that these sharers pay more attention to the political usefulness of news rather the information quality. Below, we put the polarization theory to an even stronger test by considering the role of news content in sharing decisions.

Figure 2: Predictors of news sharing



Note. Average Marginal Effects from logistic regression models. Coefficients give change in probability of sharing 1+ story from news sources. Horizontal bands represent 90% and 95% confidence intervals **Left:** Links to fake news sources. **Middle:** Links to fake news sources, controlling for (logged) real news sharing. **Right:** Links to real news sources. ***pro*-Dem sources:** *pro*-Democratic news publishers only. ***pro*-Rep sources:** *pro*-Republican news publishers only. **partisanship:** High values indicate participant identifies with political party supported by news source. **pos. affect in-party:** High values indicate positive feelings towards political party supported by news source. **neg. affect in-party:** High values indicate negative feelings towards political party opposed by news source. All independent variables range from 0 to 1. All models control for gender, race, education, income, political interest.

Assessing the Robustness of the Findings

For the sake of clarity, our main analysis employs crude binary distinctions between fake and real news as well as between those who share any news and those who do not share news at all. This section tests the robustness of our results with respect to these distinctions. Most importantly, we operationalize the sharing of fake news as the sharing of all links to news sources with a history of publishing falsehoods. However, some fake news sources produce more misinformation than others. It is possible that when it comes to blatant falsehoods, ignorant rather than polarized people are the main culprits. To test this possibility, we rely on the fake news classification of Grinberg et al. (2019) described above and re-run our analysis separately for each of three fake news categories. Our results reported in SM Section 8 show high consistency across the three categories. People who share links from sources that publish “almost exclusively fabricated stories” are also more polarized, but not more ignorant, and only slightly more disruptive than those who do not share any such links.

We also replicate our results relying on expert ratings of the trustworthiness of 60 news sources (20 mainstream news sources, 20 hyperpartisan websites, and 20 fake news sources) based on Pennycook and Rand (2019*a*). Whereas this conceptualization offers an even more fine-grained look into the quality of news sources, it comes at a cost of dropping more than 90% of all news sources from our analysis. Following Pennycook and Rand (2019*a*) we calculate the average trustworthiness of news shared for each individual in our sample and use it as our dependent variable. SM Section 9 demonstrates that the findings reflect those presented in the main text: Political variables – partisanship and feelings towards Republicans and Democrats – correlate strongly with the trustworthiness of shared news sources whereas “ignorance variables” – CRT and political knowledge - do not.

Finally, we replicate our analysis by modeling the actual count of fake and real news shared instead of binary variables of sharing versus not sharing. This involves a change from logistic regression to Quasi-Poisson models. The two models have important substantive differences. After all, sharing a single fake news link amidst dozens of real news links could be an honest mistake. Relying on the count variables, in contrast, put more weight on respondents with higher numbers of shares within a category. SM section 5 indicates this alternative modelling


procedure does not affect our main conclusions. While the effect of political cynicism is significantly reduced, partisan motivation – especially negative affect towards the out-party – remains the primary predictor of fake news sharing.⁴

The Role of News Content: Fake news as the Extreme Ends of a Partisan News Continuum

News source sharing patterns across the partisan news continuum. If partisan motives dictate both fake and real news sharing, social media users should share news sources with similar political leanings, especially near the partisan extremes of the news continuum. To test this, Figure 3 plots a heatmap of correlations based on news sharing across the seven news source types. The emerging patterns firmly favor a partisan story: People who share the most articles from ideologically biased real news sources are also most likely to share fake news sources with a similar political bias. In contrast, the off-diagonal correlations show that people shy away from sharing politically opposed sources. Importantly, the plot also reveals a partisan asymmetry. Whereas sharing *pro*-Republican real news correlates highly with sharing *pro*-Republican fake news, the correlation is substantially weaker between sharing *pro*-Democratic fake and real news. Instead, participants sharing *pro*-Democratic real news often share centrist news as well. Thus, while social media users generally favor partisan news sources, the effect is most pronounced in the Republican end of the continuum.

⁴In SM Section 4, we present a series of additional robustness tests. First, we show that our results hold after removing 61 panelists who followed - or were being followed by - at least one other panelist. Second, we demonstrate that applying YouGov sample-matching weights do not alter our results. Third, we re-run models that control for age and partisanship; this does not change our results. Fourth, we leverage a potentially better measure of digital literacy (Hargittai and Hsieh 2011); using this measure does not change our substantive conclusions. Fifth, we show that our results are robust to an alternative coding scheme that treats *InfoWars* as a *pro*-Republican fake news source rather than a real news source. Further, SM Section 8 shows that our conclusions remain intact when using an alternative list of fake news sources compiled by Grinberg et al. (2019).

Figure 3: Correlations between sharing news from the seven types of sources

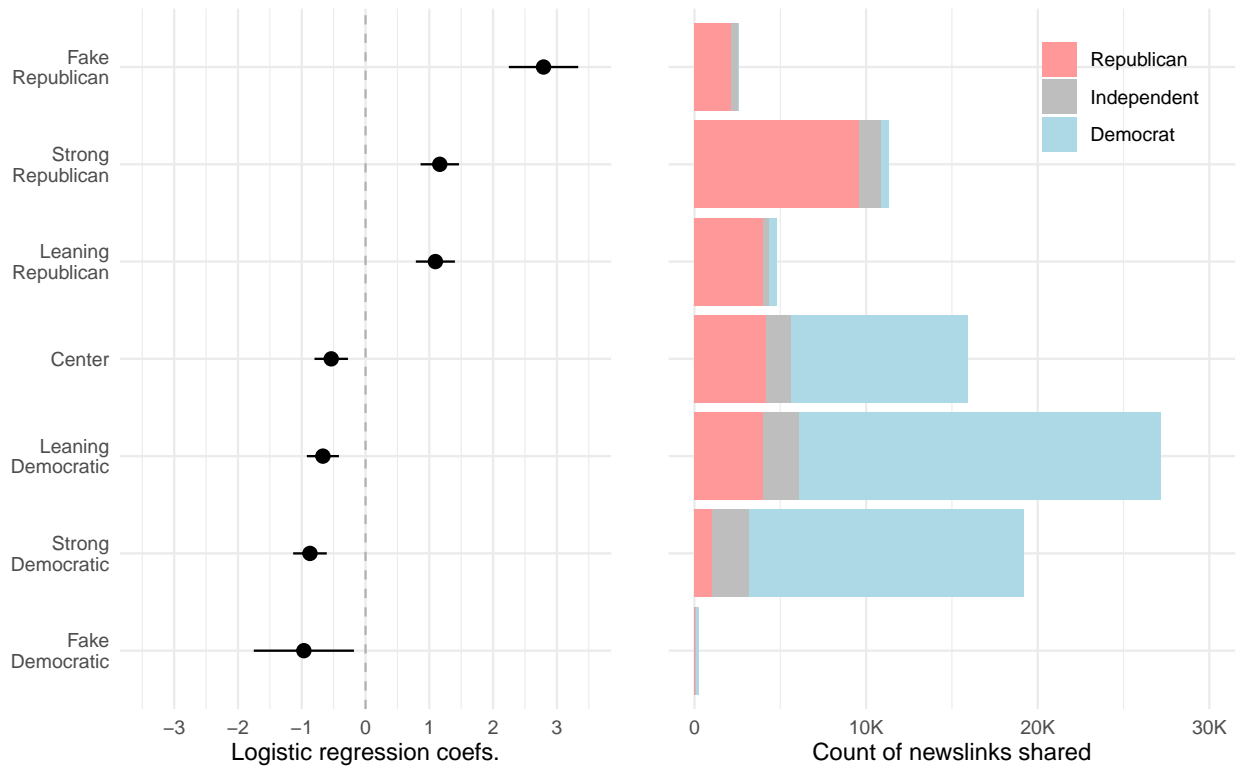


Note. Correlations are based on the logged number of news shared from each source type. In each bivariate correlation, we excluded individuals who have not shared a single piece of news from any category. Figure SM2d in the SM reproduces the plot without these exclusions.

The predictive power of partisanship across the partisan news continuum. If news sources differ systematically in how strongly their content caters to partisan motivations, then the association between partisanship and sharing should be stronger for news sources located at the extremes of the news source continuum. Figure 4 tests this assertion. The left-hand panel plots estimated logistic regression coefficients from models that examine the association between a seven-point partisanship scale (re-scaled to range from 0 = Democratic Identifier to 1 = Republican Identifier) and sharing content from the seven news source types, while the right-hand panel gives raw counts of shared links stacked by participants' partisanship (where, to ease interpretation, we trichotomize the partisanship scale). Overall, Figure 4 provides more evidence in favor of the partisan logic of news sharing: Both panels show that Republicans are more likely than Democrats to share extreme *pro*-Republican news sources and less likely to share extreme *pro*-Democratic news sources, and vice versa. The right-hand panel further shows that Republican and Democratic identifiers are perfectly sorted in terms of fake news sharing: They never share politically incongruent fake news sources.

Finally, Figure 4 also highlights the partisan asymmetry across *pro*-Democratic and *pro*-Republican fake news sharing. Democratic respondents are more likely than Republicans to share centrist, *leaning*-Democratic, *pro*-Democratic real news, as well as *pro*-Democratic fake news. Although the coefficients are monotonously increasing, the differences are small and not statistically significant. Meanwhile, the partisan asymmetry is roughly similar for *leaning* and *pro*-Republican news sources, but much more pronounced for *pro*-Republican fake news.

Figure 4: Relationship between news sharing and partisanship across seven news source types



Note. Left-hand Panel: Estimated logistic regression coefficients for relationship between partisanship (0 = Democratic Identifier ; 1 = Republican Identifier) and sharing links to seven types of news sources. Right-hand panel: Raw counts of news link shared across the seven news source types, stacked by partisanship.

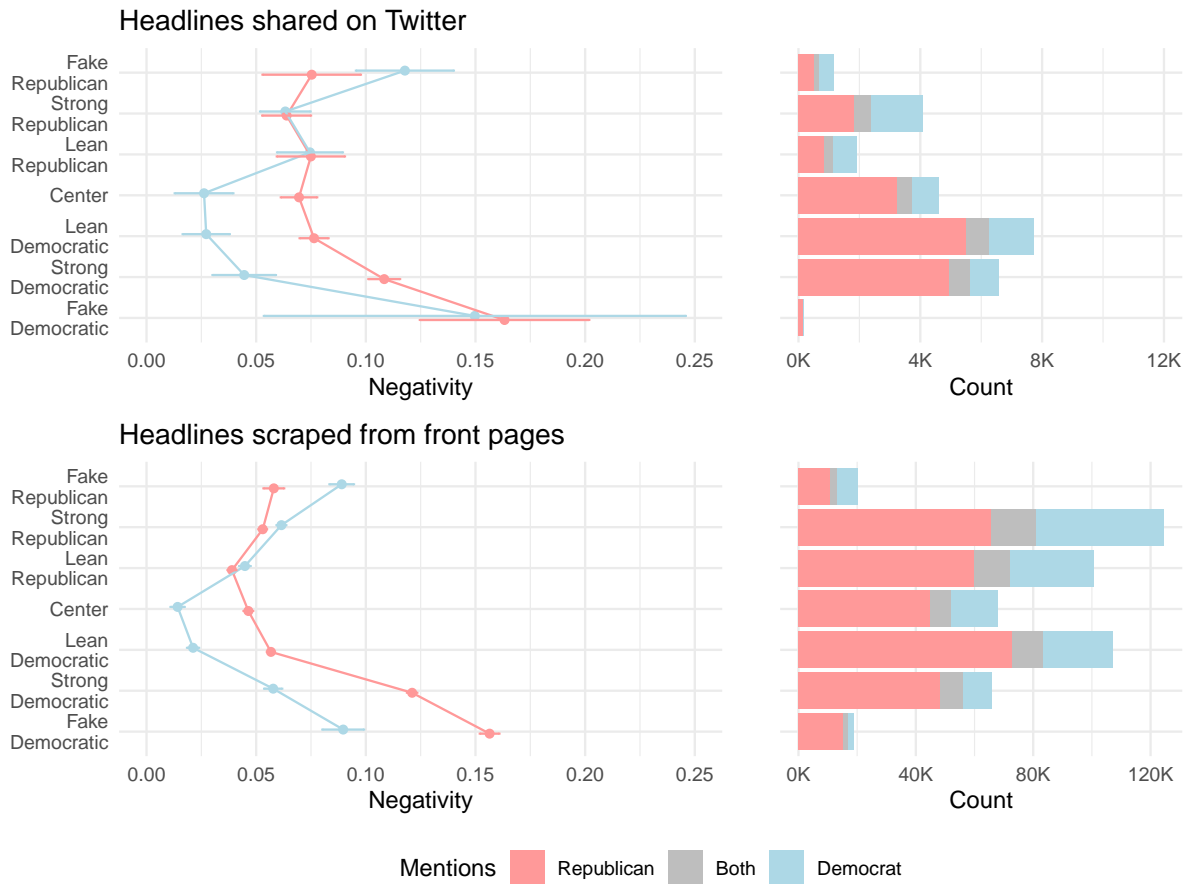
The negativity of news content across the partisan news continuum. Finally, we turn towards the sentiment analyses of headlines shared on Twitter by our panelists and for the average headline from the front web pages of the news sources shared most often by the panelists. The polarization theory implies that stories shared from Democratic news sources should portray Republican elites negatively while stories from Republican news sources should cast Democratic elites in a negative light. To test this, Figure 5 shows in the left-hand panels the negativity of news headlines from the different sources, conditional on whether the headlines from that source mention Republican or Democratic politicians. The right-hand panels show the counts of news story headlines within each news source type, conditional on mentions.

Turning to the left-hand panels, and focusing on the red lines, we observe the expected pattern: Headlines mentioning Republican elites turn more negative as we move towards the *pro*-Democratic end of the news source continuum. This pattern is strikingly similar for the headlines shared by participants (upper panel) and those scraped from the front pages of the news sources (lower panel).

Focusing next on the blue lines – again very similar across the two panels – we observe the expected contrasting trend, with headlines mentioning Democrats generally becoming more negative as we move towards *pro*-Republican news sources. The clear exception here is ostensibly *pro*-Democratic fake news. While headline sentiments from *pro*-Democratic fake news sources are harder to estimate with high precision due to their more limited numbers (cf. the right-hand panels), the findings suggest that these news sources portray Democrats and Republicans equally negative. One, admittedly speculative, interpretation of this finding is that these news sources represent extreme left-wing views critical of the entire political establishment. This interpretation fits well with evidence of similar dynamics operating on the other end of the continuum. Thus, Republican elites are not treated very positively by *pro*-Republican fake news sources either.⁵ Still, the overall pattern provides support for the expected partisan ordering of news sources based on story content: *Pro*-Democratic news sources are negative towards Republicans while *pro*-Republican sources

⁵Note that we categorize babylonbee.com as a *pro*-Republican fake news source (cf. Guess et al. 2019), although it is often referred to as a “satirical” news site, e.g. Grinberg et al. (2019). Excluding news story headlines from this web site does not change our results.

Figure 5: Negativity of news story headlines from seven different types of news sources



Left-hand panels: Headlines shared on Twitter by panelists (upper panel) and headlines scraped from front pages of news sources (lower panel). Red dots: Headlines that mention Republican politicians. Blue dots: Headlines that mention Democratic politicians. In the lower panel, we scraped headlines from the following news sources: **Fake Republican:** babylonbee.com, dailywire.com, ilovemyfreedom.org, theconservativetreehouse.com, iotwreport.com; **Strong Republican:** Breitbart.com, dailycaller.com, freebeacon.com, nypost.com, townhall.com; **Lean Republican:** foxnews.com, hotair.com, telegraph.co.uk, washingtonexaminer.com, washingtontimes.com; **Center:** thehill.com, usatoday.com, npr.com, bbc.com, wsj.com; **Lean Democratic:** politico.com, nbcnews.com, nytimes.com, theguardian.com, washingtonpost.com; **Strong Democratic:** alternet.org, dailykos.com, huffingtonpost.com, rawstory.com, vox.com; **Fake Democratic:** bipartisanreport.com, indiatimes.com, newspunch.com, palmerreport.com, themindunleashed.com. **Right-hand panels:** Headline counts from news sources, conditional on whether headline mentions Republican elites, Democratic elites, or both. Note that headlines mentioning both Republicans and Democrats counted twice in the sentiment analysis. See SM Section 11 and 12 for details on the analysis.

are negative towards Democrats.

We can go further. For partisans, a critical consideration when promoting news sources on social media is not just the level of negativity expressed towards the out-party. Instead, the most politically useful news source is one that is negative towards the out-party *and* positive

towards the in-party. This implies that we should also examine, for each type of news source, the difference in negativity between mentions of Republican and Democratic elites in Figure 5.

Recall that Democrats are more likely to share centrist news sources while Republican partisans prefer fake news sources. The net differences in headline sentiment when mentioning Republican and Democratic elites may explain why. First, like headlines from *pro*-Democratic real news sources, centrist news headlines are significantly more negative towards Republicans than towards Democrats. Second, *pro*-Republican real news headlines paint an almost equally bleak picture of Republican and Democratic elites. Third, and most importantly, stories from *pro*-Republican fake news sources are the only ones consistently publishing news stories that portray Democrats more negatively than Republicans.

We suspect this explains why Republican partisans find fake news more appealing than do Democratic partisans. To get a steady supply of news that caters to their political tastes, Republicans must turn to more extreme news sources, including those known to publish falsehoods. Democrats, in contrast, will have to look no further than to Centrist and *pro*-Democratic real news sources to meet their partisan goals. Importantly, this explanation cannot easily be dismissed by a “self-selection” story in which Republican panelists simply decide to share more extreme stories. Focusing on the front page headlines of stories journalists and editors have deemed sufficiently newsworthy to write and publish leads to the same conclusion. Still, whether this speculation is warranted, the combined observations point to an astonishing overlap in the asymmetries in (1) the sources shared by Democrat and Republican partisans, respectively, and (2) the tone of the stories shared from these sources when covering Democratic and Republican elites, respectively.

Together, these three analyses corroborate that Republicans approach news source types as if they are placed on a one-dimensional partisan continuum. Democratic partisans, to be sure, do too, but with the very important exception that *pro*-Democratic fake news sources are less attractive to highly polarized Democrats when we compare (a) Democrats to Republicans and (b) *pro*-Democratic fake news to traditional *pro*-Democratic news sources. As we turn to below, this overall pattern suggests that the sharing of fake news is

motivated by some combination of accuracy motivations and partisan goals.

Discussion and Conclusions

The 2016 US presidential election and the “Brexit” vote in Great Britain marked the rise in public concern over the spread of fake news on social media. Since then, scholars have documented myriad ways in which fake news can have detrimental effects on citizens’ political beliefs. News organizations and tech companies are beginning to counter fake news with technological safeguards, like fact-checking and blocking of fake news bots. Yet, few studies have tested whether and why citizens themselves participate in the distribution of fake news.

By linking unique data on Twitter news sharing activity to individual-level survey data with detailed psychological measures from over 2,300 American citizens, we tested two types of psychological motivations that can impact the decision to share false political news: accuracy-oriented motivations and goal-oriented motivations. Accuracy motivations imply that people prefer to share the truth but sometimes fail to identify false news due to inattention or digital illiteracy. Like others, we found that older adults were much more likely to share fake news than younger adults, an effect often attributed to ignorance. However, when examining direct measures of ignorance – cognitive reflection, political knowledge, and political digital media literacy – we did not find that fake news sharing reflected an inability to discern whether the information is true or false. This despite the fact that features of the Twitter platform, like the fast pace of tweeting and retweeting, may be especially conducive to reckless sharing behavior.

In contrast, we found strong support for the involvement of goal-oriented motivations. While disruptive goals play some role in fake news sharing, the main finding was that fake news sharing – as well as sharing of real news – reflects partisan goals. Citizens who identified politically with the side supported by the news source were more likely to share stories from that same source. In particular, we found in line with the negative partisanship literature (Abramowitz and Webster 2018) that people share fake news out of animus towards political opponents rather than positive feelings towards their own party.

This suggests that fake news sharing is “business as usual,” dictated by the same logic partisans use to process other types of information, including real news. From a partisan motivated perspective, fake news is not categorically different from other sources of political information. As the analyses of the news content showed, partisans’ decisions to share both fake and real news sources depend on how politically useful they are in derogating the out-party. Under some circumstances, this goal is best achieved with stories from fake news websites. Under other circumstances, this can best be achieved by stories from more credible websites. To the partisan mind, fake news is the extreme end of a news source continuum where the news sources differ in how well they cater to partisan goals. Accordingly, when deciding to share a story from a fake news website, most people do not prioritize whether the story is true or not. This helps explain why cognitive reflection is not related to fake news sharing even though it correlates with the ability to discern fake news from real news. These abilities are not employed in the context of real-life sharing decisions.

At the same time, the sum of findings does suggest a dual involvement of accuracy-oriented and goal-oriented motivations. First, fake news sharing is a relatively rare phenomenon that flourishes only among small segments of the population. Like others, and in contrast to what some commentators have feared, we find that only 11% of our participants shared stories from fake news sources during our period of study while half of the participants shared stories from real news web domains. Second, as discussed below, Republicans were more likely than Democrats to share fake news (e.g., Guess, Nagler and Tucker 2019), even though fake news could help Democrats achieve their partisan goals. This suggests that social media users prioritize the usefulness of information when making sharing decisions but prefer useful information that is true rather than false; they care about the veracity of news as long as it does not compromise its usefulness.

Two observations merit further discussion. First, it is relevant to discuss the finding that negative feelings towards political opponents rather than positive feelings towards the in-party are the strongest predictor of fake news sharing. There is strong evidence that partisans generally prefer helping the in-party rather than hurting the out-party (Lelkes and Westwood 2017). A hostile political climate, however, may shift the balance. Research suggests that the

danger of “symbolic” threats against the moral values of one’s group can make it “necessary to strike back, to undermine the credibility of the opposition, and to ‘hit them where it hurts’” (Amira, Wright and Goya-Tocchetto 2019). Twitter – with its fast diffusion of moralizing and emotion-laden political content (Brady et al. 2017) – may have fostered such a hostile environment, making considerations of hurting the out-party loom larger when deciding which news stories to share.

Second, we found a significant partisan asymmetry: Republicans were more likely than Democrats to share fake news sources. This could reflect that Republicans are less motivated by accuracy-motivations than Democrats. Yet, accuracy motivations have low explanatory power overall and additional analyses presented in the replication code show that the associations between ignorance-related variables and fake news sharing do not consistently differ between Democrats and Republicans. Accordingly, other explanations may be at play. One possibility highlights differences in the supply of news (Grinberg et al. 2019): Democrats and Republicans may be equally motivated by accuracy goals but Republicans end up sharing more falsehoods because their Twitter networks offer a greater supply of fake news sources to share from. Another – complementary – supply-side explanation shifts the focus from accuracy- to goal-oriented motivations: Democrats and Republicans are equally motivated to share information that is *useful* for furthering their political agenda but the supply of useful real news stories is lower for Republicans, propelling them to turn to fake news sites for material. The sentiment analyses of news headlines are consistent with this explanation. Only **pro**-Republican fake news sources were more negative towards Democrats than towards Republicans. This pattern was consistent across news stories shared by the panelists, both partisans and Independents (see SM section 11 for analyses of Independents), and the news stories from the news site front pages. This at least suggests that this asymmetry reflects a difference in the supply of content. A Republican scanning the Internet for materials to share must turn to fake news sources to find news that, on average, satisfy their partisan need to derogate Democrats. Democrats operating in the same manner, in contrast, can stop their search at posts linking to **pro**-Democratic real news websites or even centrist websites.

Without individual-level data on news exposure, it is very difficult to tease apart these explanations (see Eady et al. (2019)) and, most likely, all are at play. Still, if the latter explanations are at least partly correct, the partisan asymmetry in fake news sharing may have less to do with psychological differences between Democrats and Republicans and more to do with differences in the supply of useful information (in a partisan sense). This would dovetail with studies demonstrating that Democrats and Republicans are equally motivated by partisan goals (Brandt et al. 2014). Also, we want to note that it is unclear whether a supply-oriented explanation is evidence of actual media bias. Bias is notoriously difficult to establish. Some prior studies have found similar evidence of a liberal bias in legacy media (Groseclose and Milyo 2005), but the general negative coverage of leading Republicans could also reflect a “bias in reality” (Groeling 2013) during an unusual period in US politics.

Overall, these conclusions have important implications for how to intervene successfully against the circulation of “fake news.” If people care primarily about a story’s ability to hurt political enemies, we should not be surprised if fact-checking fails to reduce sharing. In fact, our findings suggest that offline realities shape online sharing of “fake news.” Fake news sharing relates to the increasingly polarized political debates in the United States, and it may be difficult to fix the problem of fake news without fixing the larger problem of political polarization. Unfortunately, this is much harder to accomplish than adding automatic fact-checks on social media platforms.

This does not imply that it is impossible to demotivate sharing of “fake news.” Evidence is emerging that fake news sharing decreases other people’s trust in the sharer (Altay, Hacquin and Mercier 2019). Reputational concerns may explain the low prevalence of fake news. This implies that successful interventions could work by reminding people to “not make a fool of themselves” when they are about to share information from non-credible sources.

Yet, even if such interventions can drive down sharing of genuine fake news (Pennycook et al. 2019), the underlying problem remains. Polarized partisans can often find plenty of negative content on ordinary news sites. While each individual news story may be true, the overall picture painted by exclusively sharing derogatory news stories may not be. In this perspective, the real problem is the tendency to selectively engage with derogatory news

stories, whether they are true or false.

Supplementary Materials

To view replication materials and the supplementary material for this article, please visit Dataverse at: <https://doi.org/10.7910/DVN/AH6AZK>

References

- Abramowitz, Alan I. and Steven W. Webster. 2018. “Negative Partisanship: Why Americans Dislike Parties But Behave Like Rabid Partisans.” *Political Psychology* 39:119–135.
- Acerbi, Alberto. 2019. *Cultural evolution in the digital age*. Oxford University Press.
- Allcott, Hunt, Matthew Gentzkow and Chuan Yu. 2019. “Trends in the diffusion of misinformation on social media.” *Research & Politics* 6(2):205316801984855.
- Altay, Sacha, Anne-Sophie Hacquin and Hugo Mercier. 2019. “Why do so Few People Share Fake News? It Hurts Their Reputation.”
- Amira, Karyn, Jennifer Cole Wright and Daniela Goya-Tocchetto. 2019. “In-group love versus out-group hate: Which is more important to partisans and when?” *Political Behavior* pp. 1–22.
- Bakshy, E., S. Messing and L. A. Adamic. 2015. “Exposure to ideologically diverse news and opinion on Facebook.” *Science* 348(6239):1130–1132.
- Barberá, Pablo, John T Jost, Jonathan Nagler, Joshua A Tucker and Richard Bonneau. 2015. “Tweeting from left to right: Is online political communication more than an echo chamber?” *Psychological science* 26(10):1531–1542.
- Brady, William J, Julian A Wills, John T Jost, Joshua A Tucker and Jay J Van Bavel. 2017. “Emotion shapes the diffusion of moralized content in social networks.” *Proceedings of the National Academy of Sciences* 114(28):7313–7318.

- Brandt, Mark J, Christine Reyna, John R Chambers, Jarret T Crawford and Geoffrey Wetherell. 2014. “The ideological-conflict hypothesis: Intolerance among both liberals and conservatives.” *Current Directions in Psychological Science* 23(1):27–34.
- Buckels, Erin E., Paul D. Trapnell and Delroy L. Paulhus. 2014. “Trolls just want to have fun.” *Personality and Individual Differences* 67:97–102.
- Chadwick, Andrew, Cristian Vaccari and Ben O’Loughlin. 2018. “Do tabloids poison the well of social media? Explaining democratically dysfunctional news sharing.” *New Media & Society* 20(11):4255–4274.
- Crowell, Colin. 2017. “Our approach to bots and misinformation.” *Twitter Blogpost* .
URL: https://blog.twitter.com/official/en_us/topics/company/2017/Our-Approach-Bots-Misinformation.html
- Dekker, Henk and Frits Meijerink. 2012. “Political Cynicism: Conceptualization, Operationalization, and Explanation.” *Politics, Culture And Society* 3(1):33–48.
- DiFonzo, Nicholas and Prashant Bordia. 2007. *Rumor psychology: Social and organizational approaches*. American Psychological Association.
- Eady, Gregory, Jonathan Nagler, Andy Guess, Jan Zilinsky and Joshua A Tucker. 2019. “How many people live in political bubbles on social media? Evidence from linked survey and Twitter data.” *Sage Open* 9(1):2158244019832705.
- Effron, Daniel A. and Medha Raj. 2019. “Misinformation and Morality: Encountering Fake-News Headlines Makes Them Seem Less Unethical to Publish and Share.” *Psychological Science* 31(1):75–87.
- Feezell, Jessica T. 2016. “Predicting Online Political Participation: The Importance of Selection Bias and Selective Exposure in the Online Setting.” *Political Research Quarterly* 69(3):495–509.
- Gil de Zúñiga, Homero, Logan Molyneux and Pei Zheng. 2014. “Social media, political expression, and political participation: Panel analysis of lagged and concurrent relationships.” *Journal of Communication* 64(4):612–634.

- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson and David Lazer. 2019. “Fake news on Twitter during the 2016 U.S. presidential election.” *Science* 363(6425):374–378.
- Groeling, Tim. 2013. “Media bias by the numbers: Challenges and opportunities in the empirical study of partisan news.” *Annual Review of Political Science* 16.
- Groseclose, Tim and Jeffrey Milyo. 2005. “A measure of media bias.” *The Quarterly Journal of Economics* 120(4):1191–1237.
- Guess, Andrew, Benjamin Lyons, Jacob M Montgomery, Brendan Nyhan and Jason Reifler. 2019. “Fake news, Facebook ads, and misperceptions.” *Democracy Fund* .
- Guess, Andrew, Jonathan Nagler and Joshua Tucker. 2019. “Less than you think: Prevalence and predictors of fake news dissemination on Facebook.” *Science advances* 5(1):eaau4586.
- Guess, Andrew M., Brendan Nyhan and Jason Reifler. 2020. “Exposure to untrustworthy websites in the 2016 US election.” *Nature Human Behaviour* 4(5):472–480.
- Hargittai, Eszter and Yuli Patrick Hsieh. 2011. “Succinct Survey Measures of Web-Use Skills.” *Social Science Computer Review* 30(1):95–107.
- Horowitz, Donald L. 2000. *Ethnic groups in conflict, updated edition with a new preface*. Univ of California Press.
- Iyengar, Shanto, Gaurav Sood and Yphtach Lelkes. 2012. “Affect, not ideology a social identity perspective on polarization.” *Public opinion quarterly* 76(3):405–431.
- King, Gary and Nathaniel Persily. 2019. “A New Model for Industry-Academic Partnerships.” *PS: Political Science and Politics* .
- Kunda, Ziva. 1990. “The case for motivated reasoning.” *Psychological bulletin* 108(3):480.
- Lazer, David M. J., Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson,

- Duncan J. Watts and Jonathan L. Zittrain. 2018. “The science of fake news.” *Science* 359(6380):1094–1096.
- Lee, Tae Kyoung, Youngju Kim and Kevin Coe. 2018. “When Social Media Become Hostile Media: An Experimental Examination of News Sharing, Partisanship, and Follower Count.” *Mass Communication and Society* 21(4):450–472.
- Lelkes, Yphtach and Sean J. Westwood. 2017. “The Limits of Partisan Prejudice.” *The Journal of Politics* 79(2):485–501.
- Levendusky, Matthew S. 2013. “Why do partisan media polarize viewers?” *American Journal of Political Science* 57(3):611–623.
- Mason, Liliana. 2018. “Uncivil Agreement: How Politics Became Our Identity.” 21(2):214–219.
- Miller, Joanne M., Kyle L. Saunders and Christina E. Farhart. 2015. “Conspiracy Endorsement as Motivated Reasoning: The Moderating Roles of Political Knowledge and Trust.” *American Journal of Political Science* 60(4):824–844.
- Pennycook, Gordon and David G. Rand. 2019a. “Fighting misinformation on social media using crowdsourced judgments of news source quality.” *Proceedings of the National Academy of Sciences* 116(7):2521–2526.
- Pennycook, Gordon and David G. Rand. 2019b. “Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning.” *Cognition* 188:39–50.
- Pennycook, Gordon, Jonathon McPhetres, Yunhao Zhang, Jackson G Lu and David G Rand. 2020. “Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention.” *Psychological science* 31(7):770–780.
- Pennycook, Gordon, Ziv Epstein, Mohsen Mosleh, Antonio Alonso Arechar, Dean Eckles and David Gertler Rand. 2019. “Understanding and reducing the spread of misinformation online.”

- Petersen, Michael Bang, Mathias Osmundsen and John Tooby. 2020. “The Evolutionary Psychology of Conflict and the Functions of Falsehood.”
- Petersen, Michael Bang, Mathias Osmundsen and Kevin Arceneaux. 2018. “A “Need for Chaos” and the Sharing of Hostile Political Rumors in Advanced Democracies.” *Working paper* .
- Rose, Jonathan. 2017. “Brexit, Trump, and Post-Truth Politics.” *Public Integrity* 19(6):555–558.
- Silverman, Craig. 2016. “This Analysis Shows How Fake Election News Stories Outperformed Real News On Facebook.” *Buzzfeed* .
URL: <https://www.buzzfeed.com/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook>
- Stier, Sebastian, Johannes Breuer, Pascal Siegers and Kjerstin Thorson. 2019. “Integrating Survey Data and Digital Trace Data: Key Issues in Developing an Emerging Field.” *Social Science Computer Review* 38(5):503–516.
- Stroud, Natalie Jomini. 2010. “Polarization and partisan selective exposure.” *Journal of communication* 60(3):556–576.
- Taber, Charles S. and Milton Lodge. 2006. “Motivated Skepticism in the Evaluation of Political Beliefs.” *American Journal of Political Science* 50(3):755–769.
- Thomson, Keela S and Daniel M Oppenheimer. 2016. “Investigating an alternate form of the cognitive reflection test.” *Judgment and Decision making* 11(1):99.
- Tucker, Joshua A, Yannis Theocharis, Margaret E Roberts and Pablo Barberá. 2017. “From liberation to turmoil: social media and democracy.” *Journal of democracy* 28(4):46–59.
- Vosoughi, Soroush, Deb Roy and Sinan Aral. 2018. “The spread of true and false news online.” *Science* 359(6380):1146–1151.
- Wagner, María Celeste and Pablo J. Boczkowski. 2019. “The Reception of Fake News: The Interpretations and Practices That Shape the Consumption of Perceived Misinformation.” *Digital Journalism* 7(7):870–885.

Wojcik, Stefan and Adam Hughes. 2019. "Sizing Up Twitter Users."

URL: <https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/>