Article

new media & society
2024, Vol. 26(8) 4847–4866

DOI: 10.1177/14614448221130422
journals.sagepub.com/home/nms
SAGE

Not who you think? Exposure and vulnerability to misinformation

Nicolas M Anspach
York College of Pennsylvania, USA

Taylor N Carlson
Washington University in St. Louis, USA

Abstract

Is exposure to false information necessary for misbelief? In this article, we consider the possibility that certain individuals hold misinformed beliefs without encountering misinformation, thus questioning for whom exposure to "fake news" is most deleterious. Using a pre-registered experiment on a diverse sample of 1079 US respondents, we find that the young, those with low information literacy, and those with high trust in government tend to hold mistaken beliefs, even without exposure to misinformation. Because these groups are already misinformed, eventual exposure to fake news does little to influence their misperceptions. Instead, misinformation exposure affects the elderly, those with high information literacy, and those with low trust in mainstream media the most. These results suggest that research focused on correcting misperceptions should avoid studying how certain characteristics correlate with misbelief only in misinformation's presence.

Keywords

Experiment, misbelief, misinformation, political knowledge, social media

Introduction

With the growth of social media, misinformation and fake news have received increased scholarly attention. Broadly speaking, studies of misinformation often fall into one of

Corresponding author:
Nicolas M Anspach, York College of Pennsylvania, York, PA 17403-3651, USA.
Email: nanspach@ycp.edu

two categories: those that study misinformation *exposure* (Who sees misinformation? How much misinformation is out there? Where does misinformation come from and spread?) and those that study *misbelief* (Who believes misinformation? Why?). Both research agendas have given us important insights about one of the most prominent problems facing public discourse today, but they often fail to examine the extent to which exposure and individual characteristics interact to jointly contribute to misbelief. Implicit in both strands of research is the assumption that exposure to misinformation is necessary for misbelief.

In this article, we consider the possibility that certain individuals hold misinformed beliefs without even encountering misinformation. Simply put, people can be wrong for many reasons; exposure to misinformation is just one mechanism of misbelief. An alternate explanation of misbelief may be the fact that many Americans do not follow the news and are therefore uninformed about the issues (Allen et al., 2020; Altay, 2022; Delli Carpini and Keeter, 1996; Prior, 2005). When asked about these issues, these individuals may offer their best guesses—guesses which may or may not be factually correct. Thus, exposure to misinformation is the key difference between being *un*informed (lacking knowledge about an issue) and being *mis*informed (having encountered and internalized incorrect information about an issue). We argue that serious study of misbelief must account for misinformation exposure to distinguish between the two.

If misinformation scholarship fails to establish baseline misbelief levels in the absence of misinformation, researchers may mistakenly attribute misbelief's causality to misinformation exposure. Although previous research has measured baseline levels of misbelief in control conditions of experiments designed to correct misinformation (e.g. Lewandowsky et al., 2012; Seifert, 2014; Walter and Tukachinsky, 2020), most of this work does not treat these baseline levels as a quantity of interest. Since much of this work primarily investigates the different aspects of misinformation or its corrections (see Walter and Tukachinsky's (2020) meta-analysis), examining audience characteristics may help us better understand both baseline levels of misbelief and the effects of misinformation upon exposure.

We address this gap in the literature by examining the intersection between exposure to and belief in misinformation. Specifically, we test whether and how individual characteristics (age, education, information literacy, trust in government and the mass media, and preference for clickbait content) influence misbelief in both factual and misinformative information environments. Rather than focusing primarily on characteristics of the content or source of the misinformation, we instead examine the extent to which these individual characteristics moderate misinformation's effect on misbelief.

We test our pre-registered hypotheses using an experimental design in which we randomly assign participants to either a factual or a misinformative social media feed, thereby randomizing exposure to misinformation. Our analyses indicate that the young, those with low information literacy, and those with high trust in government are more likely to give factually incorrect responses without even being exposed to misinformation. Demonstrating that certain individuals are mistaken even before exposure to misinformation, these results suggest that encountering misinformation online may be more harmful for some groups than others. Indeed, we find that the elderly, those with high

information literacy, and those with low trust in the mainstream media have the greatest increases in misbelief when exposed to fake news.

These findings demonstrate the need for scholars to move toward more complete theories of misbelief that integrate both the information environment and individual characteristics. Beyond the contribution to the growing body of knowledge about misinformation exposure and misbelief on social media, our results speak to the efficacy of removing misinformation from online platforms. As social media companies, politicians, and the public consider ways to combat misinformation, our results suggest that the removal of misleading content, while generally effective, will benefit certain individuals more than others.

Misinformation and social media

Social media platforms provide ample opportunity to learn about politics (Anspach et al., 2019), especially for those with low levels of political interest (Feezell, 2018). This learning can result from purposive, news-seeking behavior (Gil de Zúñiga et al., 2017; Weeks et al., 2017) or from incidental exposure to news contained in users' feeds (Fletcher and Nielsen, 2018; Karnowski et al., 2017). Despite this potential to inform, many Americans are concerned about the presence of misinformation online. A majority of Americans believe the prevalence of misinformation has a detrimental impact on citizens' confidence both in government and in each other (Mitchell et al., 2019: 3). It is perhaps unsurprising, then, that a majority of Americans view fake news as a very big problem, with 79% believing steps should be taken to restrict its spread (Mitchell et al., 2019: 3–4).

However, there is scholarly disagreement over just how regularly exposure to misinformation occurs. One estimate found that during the 2016 presidential campaign, one in four Americans visited a fake news website (Guess et al., 2020). Yet, when taken in context of an individual's entire media diet, exposure to misinformation may actually be quite small, as a recent study found that misinformation represents approximately 0.15% of Americans' daily media consumption (Allen et al., 2020).

Given the debate regarding the prevalence of misinformation, it is important that scholars consider alternate explanations of misbelief. Widespread concerns about misinformation have pressured social media companies to take action to reduce exposure to misinformation on their platforms. While social media giants such as Facebook and Twitter have increased their efforts to remove and flag false content on their platforms, the task has proven challenging—technically and politically. On the technical side, it is an incredible undertaking to sift through millions of posts to quickly detect and flag or remove false content. Politically, policing content breeds concerns over free speech and censorship. It may be unrealistic to expect the complete removal of misinformation from social media.

Though it is unlikely that social media companies could efficiently remove all misinformation from their platforms, it is not unreasonable to suggest that platforms can do more to police and remove misinformative content. Indeed, some have already made strides in that direction. In May 2020, Twitter made the bold step of fact-checking a misleading tweet by President Trump regarding mail-in voting. In October of the same

year, Facebook announced that it was banning all QAnon (a popular conspiracy theory that traffics in misinformation) accounts across its platforms. After the 2020 US general election, Facebook shut down "stop the steal" Facebook groups that it considered to be sharing misinformation about the election outcome and Twitter went so far as ban President Trump from its platform. However, recent probes into Facebook's misinformation teams suggest that the company heavily scaled back its misinformation protections after the election (e.g. Timberg et al., 2021). Some current and former Facebook employees consider this to be a defining feature of Facebook's role in the January 6th attack at the Capitol.

As the world tries to make sense of the misinformation that spread on social media in the wake of the 2020 US presidential election and the COVID-19 pandemic, there is increased demand for companies to remove false content. However, the question of whether misinformation's removal ameliorates misbelief remains. While there is anecdotal evidence from comments from Facebook employees that the amount of misinformation increased on its platform when it scaled back its precautions after the 2020 election, this is hardly causal. It is possible that removing false content from a social media platform does little to reduce misbelief if exposure is not a necessary condition for misbelief. Moreover, some people might be misinformed even without exposure to specific false claims on social media. If certain individuals are misinformed even without encountering misinformation, then efforts to curb exposure to fake news may have limited results.

The determinants of misbelief

To understand whether exposure is a necessary condition for misbelief, we must first unpack what we know about who believes misinformation and why. Much of the scholarship that investigates the psychological and demographic determinants of misbelief does so by presuming or forcing exposure to false information. For example, some studies do not randomize exposure to misinformation at all and simply measure the individual characteristics that correlate with beliefs in misinformation or conspiracy theories (e.g. Miller et al., 2016). In this context, exposure to misinformation is presumed, but not measured or manipulated experimentally—only the beliefs are measured.

Other previous research has measured baseline levels of misbelief in control groups, but this measurement is almost exclusively used for comparison purposes in studies designed to measure the effectiveness of corrections (e.g. Lewandowsky et al., 2012; Nyhan and Reifler, 2010; Seifert, 2014), rather than considering it to be a meaningful quantity of interest. For example, experiments investigating the continued influence effect often randomly assign participants to receive nothing (pure control), misinformative content, or misinformative content and a correction. Researchers then measure the gap in misbelief between the two treatment conditions and control. Walter and Tukachinsky (2020) conducted a meta-analysis of studies that "report on a comparison between a misinformation and correction condition and a no-misinformation control group" (p. 163). While the meta-analysis identifies 32 studies that include a no-misinformation control condition, the quantity of interest in these studies is the effectiveness of corrections rather than whether exposure to misinformation causes misbelief. More importantly, when researchers have examined moderators within the continued

influence effect literature, they have primarily examined features of the misinformation or features of the corrections, rather than characteristics of the individual (Walter and Tukachinsky, 2020).

Recent evidence suggests that people encounter fewer untrustworthy websites than originally thought (Guess et al., 2020), which means that understanding the sources of misbelief beyond exposure is all the more important. If survey data suggest that misbeliefs are relatively high in the United States, for example, but exposure (as measured by behavioral web-tracking data) is relatively low, we need to further investigate misbelief absent exposure. For example, individuals can have mistaken priors due to belief bias (Evans et al., 1983), false memories (Loftus and Pickrell, 1995), motivated reasoning (Kunda, 1990), or simply being uninformed (Allen et al., 2020; Altay, 2022; Delli Carpini and Keeter, 1996; Prior, 2005). People can also provide incorrect doorstep responses or simply guess incorrectly (Berinsky, 2004). All of these mechanisms may lead to misbelief without exposure to misinformation.

Studies that report on misbelief only when exposed to misinformation, rather than detailing misbelief among those who were not exposed, could run the risk of misidentifying the groups most susceptible to fake news. Such designs fail to consider whether these individuals were misinformed *before* exposure to misinformation. If so, exposure to misinformation has a negligible effect for them; they are already wrong. Indeed, recent evidence suggests that certain groups traditionally considered vulnerable to fake news—the elderly and the less digitally literate—are, in fact, *less* affected by misinformation exposure (Munger et al., 2021: 10). Clearly, more investigation is warranted.

In this article, we examine which individual characteristics are associated with misbelief without exposure to misinformation, reporting these results as important quantities of interest in and of themselves. We then assess how those characteristics interact with exposure to affect misbelief. We focus primarily on characteristics that previous scholarship has identified as being correlated with either misbelief (after exposure) or with exposure to misinformation.

We begin with age, as research indicates that elderly individuals are more engaged with misinformation. Munger et al. (2021) summarize previous research, finding that individuals over 65 share fake news articles substantially more than younger individuals—7 times more on Facebook (Guess et al., 2019). The authors argue that these patterns are likely driven by variation in digital literacy by age, though additional evidence also indicates that older individuals are more trusting of online content (Liao and Fu, 2014).

There are two possible explanations for the association between age and misbelief found by previous scholarship. One possibility is that the elderly, on average, are simply more misinformed than younger people. Another possibility is that because younger people have less engagement with politics and lower levels of political knowledge (Pew Research Center 2018), they may be more likely to report incorrect information even without exposure to misinformation. In this situation, exposure to misinformation has a negligible effect for younger people; they are already wrong. As the research discussed above is dependent on exposure to misinformation, we address the question of how age is correlated with misbelief in misinformation's absence. Given the emphasis on the positive correlation between age and misbelief in previous scholarship, we test against

our pre-registered predictions that the elderly will report higher levels of misbelief when misinformation is absent (H1a) and that age's effect will increase when interacted with misinformation's presence (H1b).

Scholars have also shown that higher levels of education are associated with lower levels of misbelief (Converse, 1964; Delli Carpini and Keeter, 1996; Flynn et al., 2017). It could be the case that those with higher levels of education follow political news and have developed stronger critical thinking skills, allowing them to distinguish between true and false content. If this is the case, we would expect education to be negatively correlated with misbelief, regardless of exposure to misinformation. In our pre-registered hypotheses, we predict that those with higher levels of education are more likely to hold accurate beliefs in factual environments (H2a) and that education's effect on misbelief will increase when interacted with exposure to misinformation (H2b).

In addition to examining the effects of education on belief in misinformation, we also examine a related concept—information literacy. Information literacy encompasses reflective discovery, use, and evaluation of information (Association of College & Research Libraries 2016). These skills might very well be a function of education, but it is important to note that many curricula do not necessarily focus on information literacy. However, the theory underpinning information literacy's hypothesized effect is similar to that of education. Like education, information literacy may help people distinguish between true and false content more accurately. In addition, those with high information literacy may be able to weigh evidence objectively without engaging in motivated reasoning or (dis)confirmation biases. Because of the similarities in the theoretical mechanisms between education (H2) and information literacy (H3), we include both in our analyses to test whether one, the other, or both hold predictive power.

Next, scholars have shown that misbelief is also associated with trust in the two institutions arguably most responsible for the dissemination of political information: the mainstream media and the government itself. Specifically, individuals with lower levels of trust in the mainstream media tend to consume more alternative or misinformative content (Tsfati and Cappella, 2003). When individuals distrust the mainstream media, they are more likely to opt for fringe sources that they believe to be more trustworthy. Even in the absence of misinformation, those with low trust in the media may be more likely to engage in conspiratorial thinking (De Coninck et al., 2021; Mari et al., 2022), thus leading them to adopt misbeliefs that run counter to factual news reports.

Similarly, individuals with lower levels of trust in the government also tend to believe misinformation more, especially with respect to conspiracy theories (Miller et al., 2016). Such individuals may be more likely to trust misinformative claims or, in the absence of misinformation, form misbeliefs in opposition to official government narratives. Yet, it is important to note that the relationship between trust in government and misbelief could be context-dependent. Although trust in government is typically measured generally (e.g. "How much do you trust the government to tell the truth?"), how people interpret "the government" could vary. Some might consider governmental institutions, while others may consider the current presidential administration. This has particularly important implications given that we fielded our experiment during Donald Trump's administration, a fact that we address in our results section.

The last individual characteristic we consider is whether preference for questionable content correlates with misbelief. The nature of (mis)information exposure on social media remains a subject of debate in communication research. On one hand, some argue that because people may see content on their feeds that they would not have sought out independently, inadvertent exposure to (mis)information on social media is a common occurrence (Mitchell et al., 2013). On the other hand, many platforms' algorithms have discerned users' preferences, enabling them to prioritize content consistent with those preferences. If someone prefers conspiratorial content, for example, Facebook's algorithm will likely display more conspiratorial content in their News Feed. Because of this, user preference—even in the social media context—remains an important element of an information environment, suggesting that scholars would be remiss in dismissing user preferences completely. In addition, preference for questionable content may also act as a preference for conspiracy theories. If so, then one may not even need to be exposed to questionable content to be misinformed (Uscinski et al., 2022). We therefore examine whether preference predicts misbelief independent of the other individual characteristics described above.

In the following analyses, we investigate how each individual characteristic interacts with the (mis)information environment to contribute to misinformed views.[1] In our pre-analysis plan (see Supplemental Appendix B), we pre-registered hypotheses regarding misinformation exposure and misbelief, largely to replicate previous work and contextualize the interactions between our individual characteristics and exposure to misinformation. However, we only present the findings most central to our research question in this article, and fully report results from the rest of the pre-analysis plan in Supplemental Appendix H. Though we are open to how each characteristic interacts with misinformation's presence, we formally test against the findings of much of the literature:

H1a. Absent exposure to misinformation, age will have a positive effect on misbelief.

H1b. Age's effect on misbelief will increase when interacted with misinformation's presence.

H2a. Absent exposure to misinformation, education will have a negative effect on misbelief.

H2b. Education's effect on misbelief will increase when interacted with misinformation's presence.

H3a. Absent exposure to misinformation, information literacy will have a negative effect on misbelief.

H3b. Information literacy's effect on misbelief will increase when interacted with misinformation's presence.

H4a. Absent exposure to misinformation, trust in government will have a negative effect on misbelief.

H4b. Trust in government's effect on misbelief will increase when interacted with misinformation's presence.

H5a. Absent exposure to misinformation, trust in mainstream media will have a negative effect on misbelief.

H5b. Trust in mainstream media's effect on misbelief will increase when interacted with misinformation's presence.

H6a. Absent exposure to misinformation, preference for questionable content will have a positive effect on misbelief.

H6b. Preference's effect on misbelief will increase when interacted with misinformation's presence.

Research design

We recruited a nationally diverse sample of 1200 respondents, quota-matched to Census demographics, from Lucid (see Coppock and McClellan, 2019 for a validation) in September 2020.[2] We include descriptive statistics and balance tests between our experimental conditions in Supplemental Appendix E, finding no statistically significant differences on observables between treatment groups.

We test our theories using an experiment in which we randomly exposed participants to social media information environments, manipulating the presence of misinformation. In the survey (see Supplemental Appendix D), we measure our individual characteristics of interest pre-treatment, beginning with basic demographics. We then measure preference for clickbait using a selection exercise, a common practice for measuring media preferences (Arceneaux and Johnson, 2013; Luca et al., 2022; Prior, 2005). We presented respondents with a list of five headlines that varied in the extent to which they resembled mainstream or clickbait content. Participants chose which headline they would most like to read. On the next screen, they were presented with the remaining four headlines and again asked to choose which they would like to read, repeating the procedure until no headlines remained. Respondents who ranked the clickbait COVID-19 headline higher than the mainstream COVID-19 headline were considered to have a preference for clickbait. The headlines were

- CDC updates timeline for COVID-19 vaccine (factual)
- What the "experts" won't tell you: the truth behind how COVID-19 spreads (clickbait)
- The 10 most-anticipated shows coming to Netflix in the fall (entertainment)
- Will they play? Current status of all major sports leagues (entertainment)
- The best albums of 2020 (so far) (entertainment)

After indicating their preferences via the headline selection exercise, participants answered questions that allow us to measure information literacy (adapted from O'Connor et al., 2001) and trust in government and media. We also included Munger

Table 1. News feed article descriptions by treatment group.

	Factual condition	Misinformative condition
Factual article previews		
Russian Interference (*New York Times*)	X	X
Russian Interference (*ABC News*)	X	X
Mail-In Voting (*Pew Research Center*)	X	–
Mail-In Voting (*Brookings*)	X	–
Misinformative article previews		
EU Conspiracy (*US Patriot News)*	–	X
EU Conspiracy (*ABCNews.com.co*)	–	X
Entertainment article previews		
Coronavirus Movie Delays (*Vulture*)	X	X
Quarantine Book List *(Oprah Magazine*)	X	X

et al.'s (2021) measure of digital literacy and political knowledge measures for robustness tests (see Supplemental Appendix G).

At this point, participants began the experimental portion of the survey where we randomly assigned participants to view one of two Facebook News Feeds. The factual condition (coded as 0) featured six posts: four factual political news posts and two entertainment posts, with no misinformation. The misinformative condition (coded as 1) also featured six posts, but included a mix of two factual political news posts, two misinformative posts, and two entertainment posts. Table 1 displays the specific posts included in each condition. Images of the News Feeds are available in Supplemental Appendix C.

Participants in all conditions were shown the following instructions:

> In the next exercise, you will see a mock Facebook News Feed. You can read as few or as many posts as you wish—just like on the actual Facebook. When you click a post, you will be taken to the full article. After reading the article, you will have the option to return to the News Feed to select another article.
>
> When you are done selecting posts and wish to move on to the next portion of the survey, select "I am ready to move on" from the bottom of the News Feed.

Participants were then presented with their randomly assigned News Feed. The posts included in the News Feed were designed to look like news article previews. Each included an image, a headline, and a brief summary of the article itself. Importantly, all of the article previews in the factual condition included true information, whereas the two misinformative article previews in the misinformative condition included false information directly in the preview itself.

If respondents clicked on a post, they were taken directly to the article associated with that post. For the factual articles, we include the actual text as published by the news outlets, with light editing to keep the length and emphasis consistent across articles. For the misinformative posts, we wrote two fake articles, modeled closely after articles we observed in

Table 2. Belief in misinformation measures.

	Factual answer	Fake news answer
1. Which of the following poses the greatest threat to the integrity of the 2020 election?	Russia (2 posts, 2 articles)	The EU (2 posts, 2 articles)
2. How are foreign adversaries allegedly trying to disrupt the 2020 election?	Cyberattacks and Misinformation (1 post, 2 articles)	Bribery (1 post, 2 articles)

our own searching. Participants in all conditions were free to select up to six articles and they could read an article more than once. This means that although the starting point in their information environment is randomly assigned, respondents still had some degree of choice over their exposure to the full articles' content. This approach preserves the study's external validity, as the selection processes in our design mirror those on actual Facebook. In addition, this design allows us to identify the causal effect of exposure to misinformation within context of the News Feed because the misinformation treatment was included both in the articles and in their Feed previews. However, the "dose" of exposure could vary based on individual preferences because individuals could choose which articles to click on to read more. We address the implications of this further in the discussion section.

Finally, participants answered two questions to test their factual knowledge about the (mis)information included in the News Feeds. Table 2 includes information on the answers to each question based on the factual and fake news article previews. We measure our dependent variable (misbelief) by counting how many misinformed responses participants gave on our two questions, so that the score ranges from 0 to 2.

For these questions, both the factual and misinformative conditions included posts addressing their respective topics (threat and method of subversion). However, the conditions differed in whether misinformation about those topics is present. The factual condition contained only factual information—that is, that the FBI and FEC have identified Russia as the greatest threat to the election, and that Russia is employing cyberattacks and misinformation campaigns to undermine electoral integrity. The misinformative condition included this same factual information, but also contained misinformation regarding a conspiracy in which European Union bankers bribed state election officials to change vote tallies. Note that we fielded this experiment before the 2020 US Presidential election.

After completing the post-treatment questions, respondents received a debriefing statement. Because two of the News Feed posts and related articles contained fake news, we informed participants of the false content and supplied them with links to the FEC website and to a primer on how to spot misinformation.

Results

Article selection and responses

Before we formally test our hypotheses, we present an exploratory analysis of how people interacted with the content in our faux Facebook News Feeds. Specifically, Figure 1

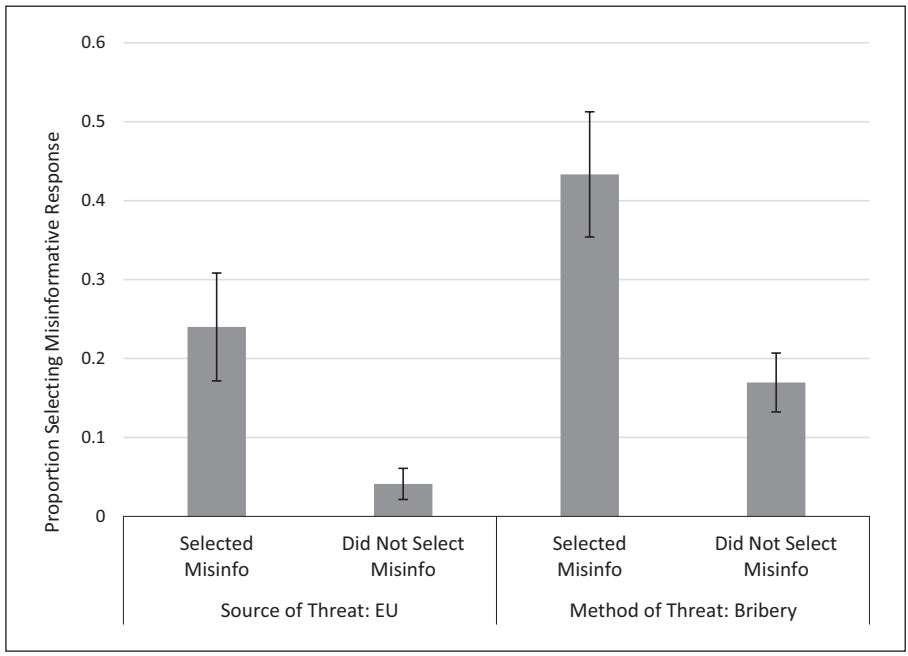

Figure 1. Responses by misinformation selection (misinformative condition only).

analyzes those assigned for the misinformative treatment, displaying the differences in responses between those who clicked to read a misinformative article and those who did not. Though we randomly assigned participants to either a factual or misinformative News Feed, participants could decide for themselves whether to click on the Feeds' posts to read the respective article. Because this decision to click a post is not randomly assigned, Figure 1 is not meant to be interpreted causally.

Of those assigned to the misinformative condition ($n = 540$), 150 (27.8%) clicked at least one misinformative post. This selection is positively correlated with responding to our two dependent variable questions with the specific misinformation embedded within these articles. In other words, those who clicked through to the full misinformative articles were significantly more likely than their non-clicking counterparts to answer (1) that the European Union posed the greatest threat to the integrity of the 2020 election (24.0% vs 4.1%, $z = 7.0$, $p < .05$) and (2) that foreign adversaries were bribing election officials to disrupt the election (43.3% vs 17.0%, $z = 6.4$, $p < .05$). Additional analysis and discussion of the role of selection in the experimental News Feeds is located in Supplemental Appendix I.

These exploratory analyses make two key points that help contextualize the hypothesis tests that follow. First, many participants assigned to the misinformative condition clicked to read at least one misinformation article. This gives us more confidence that our treatment stimuli were taken seriously by participants. Second, the more participants engaged with the misinformative content, the more likely they were to provide answers consistent with their exposure. This suggests that the results we present in the next section—which focus

on exposure only to the randomly assigned News Feed—could underestimate treatment effects, particularly for those who choose to engage more deeply with posted content. However, the choice to engage with some content over others is likely correlated with one's underlying preferences that could confound estimates of misbelief, so we do not read too much into these exploratory analyses and instead put more confidence in our estimates of the average treatment effects, as presented in the next section.

Exposure and vulnerability to misinformation

In this section, we examine the extent to which individual characteristics interact with misinformation exposure to influence misbelief. We test our hypotheses using Poisson regression models in which the dependent variable is the number of misinformed responses to our post-treatment questions (0–2). We first focus on each individual characteristic's effect on misbelief in the factual condition and then examine the interaction terms to determine whether those effects change in misinformation's presence.

The results, presented in Figure 2 (full models in Supplemental Appendix F), indicate mixed support for our hypotheses regarding who is most vulnerable to misinformation. Looking first at age, we observe that older individuals are *less* likely to hold misinformed beliefs when they are not exposed to misinformation, rejecting H1a. This is somewhat surprising, given that much of the misinformation scholarship implicitly or explicitly associates increased age with misbelief (e.g. Guess et al., 2019; Liao and Fu, 2014; though see Munger et al., 2021). We also reject H1b, as we find that age *loses* its predictive power in misinformative environments; when exposed to misinformation, the young and the elderly display statistically indistinguishable levels of misbelief. Together, these results demonstrate the need for scholars to consider misbelief in both factual and misinformative environments. Without that comparison, we could be ineffectively designing and targeting corrective efforts to the wrong groups.

Contrary to H2a and H2b's expectations, we find no evidence of a relationship between education and misbelief. Despite the appearance of a negative slope in Figure 2, the effect of education on misbelief is statistically insignificant in both experimental conditions. Emphasizing the distinction between education and information literacy, however, we do find a significant relationship between information literacy and misbelief. As predicted by H3a, higher levels of information literacy are associated with lower levels of misbelief in the factual condition. But when interacted with misinformation exposure, information literacy—like age—also loses its explanatory power. This is contrary to H3b's expectation and again illustrates the importance of measuring misbelief without exposure to misinformation. Because those with the lowest levels of information literacy are misinformed even without exposure to fake news, exposure influences them less. Instead, it is those with the highest levels of information literacy that have the most to lose from encountering misinformation.

Next, we examine how trust in government and in the mainstream media influence misbelief. Contrary to H4a and H4b's predictions, we find that trust in government's influence on misbelief is consistent across experimental conditions: as trust in government increases, so does misbelief. This surprising result, which runs counter to previous research on the topic (e.g. Miller et al., 2016), warrants further investigation. Perhaps

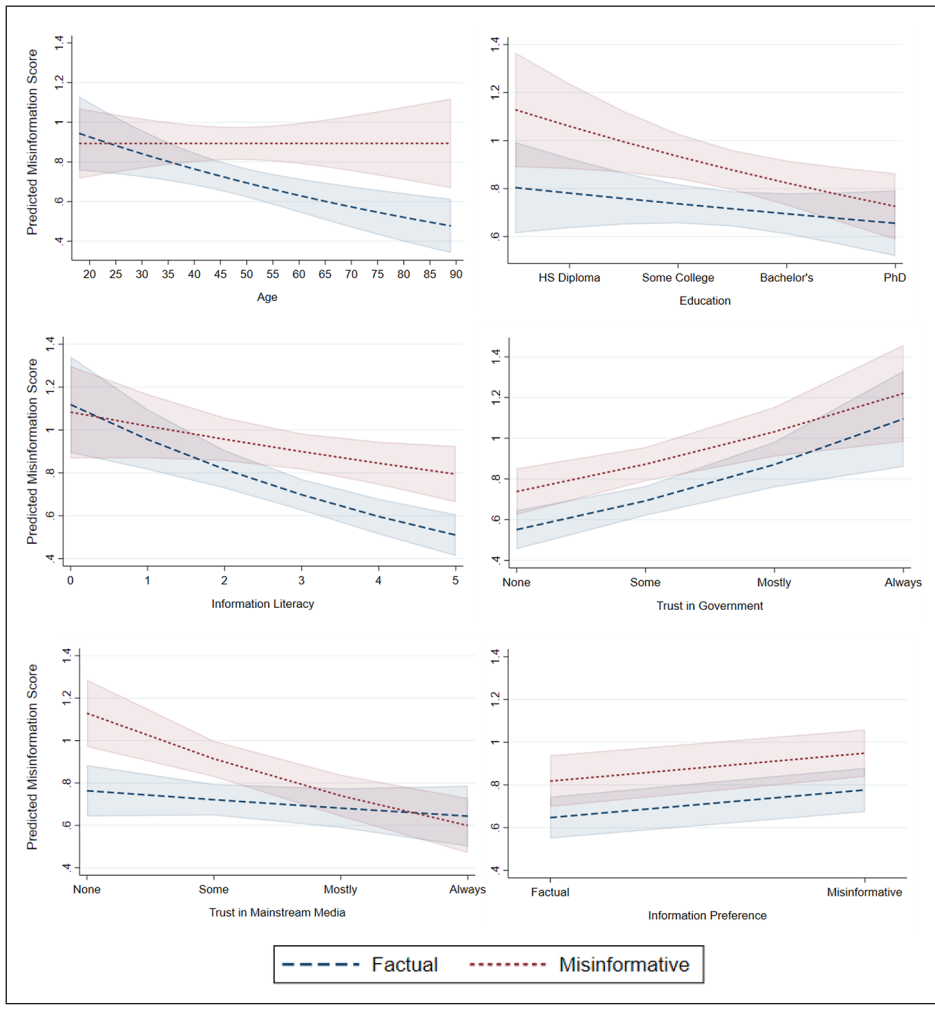

Figure 2. Interacted treatment effects.

different interpretations of the word "government" may explain our divergent results; we might expect one's trust in the democratic institutions of American government to differ from one's trust in the Trump administration, for example. Indeed, we find evidence that identification with the Republican Party is positively associated with trust in government ($r = .28$)—a finding that runs counter to traditional conservative principles.

Similarly, we do not find support for H5a, as trust in mainstream media has no effect on misbelief in the factual condition. However, we find evidence for H5b, as media trust gains predictive power when interacted with misinformation exposure. In the factual condition, trust in mainstream media has no effect on misbelief, perhaps because the factual condition did not include content from questionable sources. Since mainstream

outlets were the only sources available, even those with low trust in mainstream media had no choice but to use them. However, the misinformative condition included two fringe sources, allowing those with low levels of trust to eschew mainstream media in favor of more conspiratorial outlets, likely explaining the negative relationship shown in Figure 2. This particular finding is noteworthy, as it is the only interaction effect that works as predicted by previous research (e.g. Tsfati and Cappella, 2003).

Finally, we also examine the relationship between preference for questionable content and misbelief. Despite the appearance of a positive slope in Figure 2, preference for questionable content's effect is only suggestive ($p < .10$) in the factual condition. Though H6a and H6b predicted that such a preference would increase misbelief and gain explanatory power when interacted with misinformation, Figure 2 provides no evidence for either hypothesis. This pattern holds for the stricter definition of preference (see Supplemental Appendix G.5).

Discussion

In this article, we examined how individual characteristics, misinformation exposure, and interactions between the two influence misbelief. Our results suggest that misinformation exposure on social media does not influence everyone equally. Instead, such exposure has the most deleterious effects for the elderly, those with high information literacy, and those with low trust in the mainstream media. We found that misinformation on social media does not necessarily harm those groups traditionally thought to be most vulnerable to misinformation, likely because they are already misinformed before exposure to fake news.

While the results from our study are important, they do come with some limitations. First, to obtain causal identification on the information environment, we created treatment News Feeds. In so doing, we stripped away some features of the real Facebook environment that could still be influential. For example, participants viewed posts without any social cues, such as likes and comments, which previous research has shown to be influential (Anspach and Carlson, 2020), especially when they come from close friends (Anspach, 2017). Moreover, participants were perusing a News Feed within the context of a survey, not in their real lives. The incentives with which individuals log in to Facebook vary and could affect the information that they retain (Settle, 2018). Finally, we only examined one misinformation topic, one that we created ourselves and that we assume to have been wholly novel to our subjects.

Second, although our study included an experiment, there are two points about causal identification we should consider carefully. First, we randomly assigned participants to view a purely factual News Feed or a News Feed that included two posts that included misinformation. This gives us causal identification on exposure to misinformation relative to no exposure, but we have no guarantee that participants in the misinformative condition actually read the misinformative posts. We still maintain causal identification with this intent to treat (ITT) effect, but we acknowledge that some participants in the misinformative treatment group may not have been "treated" with exposure if they did not read the false information provided to them. Similarly, participants were free to engage with the content as they would while browsing social media, meaning those who

clicked on the fake news article posts were exposed to a stronger dose of the treatment than those who chose not to read the articles. This choice was not randomly assigned, so we do not have causal identification on exposure to misinformation in the *article* content.

The second point to consider with respect to causal identification involves the individual characteristics themselves. We drew these characteristics from the literature, but they are impossible to manipulate experimentally. Kam and Trussler (2017) warn that observing characteristics of interest rather than randomly assigning them can introduce confounds when determining causality. However, following their recommendations, we control for potential confounding factors in our models. Moreover, we measured our moderators pre-treatment to avoid post-treatment bias (Montgomery et al., 2018), which is further justified by recent evidence that measuring moderators pre-treatment does not threaten inferences from the experiment (Sheagley and Clifford, n.d.).

Third, we need to be cautious in interpreting our results within the context of a survey experiment. As with all survey data, there is bound to be measurement error. In our study, one concern might be that individuals in the factual condition are not actually misinformed without exposure, but are instead clicking at random, guessing, or otherwise answering the questions in ways that do not reflect their true underlying (mis)beliefs. Moreover, when it comes to considering the extent to which individual characteristics correlate with misbelief, it could be the case that certain characteristics lead people to interpret and respond to survey questions differently, rather than having different levels of misbelief. This may especially be the case with our "trust in government" variable, which some respondents may have interpreted as trust in the Trump administration, while others may have interpreted the question as government more generally.

Finally, our sample was diverse and quota matched to census records, but it was not a national probability sample. This means that while our results can speak to a broad group of American adults, they are not necessarily representative because there could be unobserved features that led our respondents to opt in to Lucid's survey panel. That said, Lucid has recently been used extensively by academic researchers, especially within the context of experiments (Coppock and McClellan, 2019), and weekly opinion polling via the Democracy Fund and UCLA Nationscape surveys (Tausanovitch et al., 2019). We are confident that our results are meaningful, but acknowledge that we have not established true "baseline" levels of misinformation among various demographic groups in the United States. Rather, we view these results as indicative of patterns comparing one group against another.

Despite these limitations, these results have important implications for our understanding of misinformation, its belief, and potential solutions to the effects of fake news. Chiefly, exposure to misinformation is not the only determinant of misbelief, as individuals can hold misinformed opinions even when misinformation is absent. Although previous research has measured similar quantities, our results suggest that future research should engage more critically with misbelief absent exposure to misinformation (i.e. in the control groups in many experimental designs). This suggests that research focused on correcting misperceptions should avoid studying how certain characteristics correlate with misbelief only in misinformation's presence, lest scholars attribute causality to the wrong mechanism.

Conclusion

In conclusion, this article makes contributions to the misinformation and misbelief literatures. First, we uniquely draw attention to misbelief *absent exposure* to misinformation. Although previous research has measured rates of misbelief without exposure in experimental control groups, it has generally not given much attention to this quantity as theoretically important. In their meta-analysis of misinformation correction experiments, Walter and Tukachinsky (2020) argue that the key comparison of interest should be between "individuals who were exposed to corrected misinformation and those who were not exposed to misinformation in the first place. This contrast allows to determine the extent to which correction interventions revert the individuals' attitudes and beliefs back to baseline" (p. 157). The authors consider the baseline levels of misinformation to be important for comparative purposes, but do not do more to consider why this quantity is independently important.

We argue that understanding the rates of misbelief without exposure to relevant misinformation is important in and of itself because many scholars, policymakers, and pundits suggest that better policing of misinformative content on social media will cure misbelief. Our analyses demonstrate that exposure to misinformation is but one mechanism of misbelief among many. Drawing attention to misbelief absent exposure allows us to better understand these alternate mechanisms, such as the individual characteristics studied in this article. Though we are not the first to examine the correlations between such characteristics and misbelief, the focus of much of the previous research is on the characteristics' interaction with misinformation exposure rather than as independent contributors to misbelief.

Indeed, some of our findings provide nuance to the narratives surrounding these individual characteristics. For example, previous work has shown that age is positively correlated with misbelief (Brashier and Schacter, 2020; Chia et al., 2021). While we find the same, it is perhaps for a reason overlooked by much of the literature. The differences in misbelief between the young and the old are not because the young are better able to identify and reject misinformation. Instead, it is because the young, on average, are misinformed even when unexposed to misinformation. The elderly, however, are better informed in misinformation's absence, thus making them more susceptible to its exposure.

Fourth, a finer contribution made in this article is that information literacy is an individual characteristic that future research should take care to examine. The concept remains a slippery one, with scholars examining both media (Lyons et al., 2021) and digital (Munger et al., 2021) literacies. Though we complicate the picture further by introducing O'Connor et al.'s (2001) measure of informational literacy, we find the contribution to be a worthwhile one. Indeed, we find that information literacy is negatively correlated with misbelief, but Munger et al.'s (2021) measure of digital literacy to be uncorrelated (see Supplemental Appendix G.1). This suggests that the skills associated with critically evaluating information might be a relevant consideration in future research, above and beyond educating people about how the Internet works (digital literacy) and media bias (media literacy). Future research could build on our results to not only refine measurement of information literacy, but also consider interventions based on increasing information literacy to reduce misperceptions.

Taken together, our findings demonstrate that we must consider both the information environment and individual characteristics to understand misbelief. By examining the extent to which that exposure can dampen or exacerbate the effects of individual characteristics associated with misbelief, we hope to offer scholars a more complete picture of misbelief, and social media platforms a greater understanding of the implications of their decisions to police and remove misinformative content.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Nicolas M Anspach https://orcid.org/0000-0001-9326-5486

Supplemental material

Supplemental material for this article is available online.

Notes

1. Our pre-analysis plan can be found at https://osf.io/jnf42/?view_only=6dd50689b3af42c39d64f6a30d812d6a. As we describe in detail in Supplemental Appendix A, we reworded and renumbered our hypotheses based on reviewer feedback. Supplemental Appendix A details all deviations from the pre-analysis plan and Supplemental Appendix H presents all results strictly following the pre-analysis plan. Hypothesis predictions and results presented in the article are substantively the same as those in the pre-analysis plan.
2. Scholars have recently questioned Lucid because of respondent inattentiveness (Aronow et al., 2020). Approximately 12% of our respondents failed the pre-treatment attention check. Following our pre-analysis plan, we omitted these respondents from the main analyses. We conduct supplemental analyses that include these respondents in Supplemental Appendix G.4.

References

Allen J, Howland B, Mobius M, et al. (2020) Evaluating the fake news problem at the SCAE of the information ecosystem. *Science Advances* 6: eaay3539.

Altay S (2022) How effective are interventions against misinformation? Available at: https://psyarxiv.com/sm3vk/

Anspach NM (2017) The new personal influence: how our Facebook friends influence the news we read. *Political Communication* 34(4): 590–606.

Anspach NM and Carlson TN (2020) What to believe? Social media commentary and belief in misinformation. *Political Behavior* 42(3): 697–718.

Anspach NM, Jennings JT and Arceneaux K (2019) A little bit of knowledge: Facebook's news feed and self-perceptions of knowledge. *Research & Politics* 6(1): 1–9.

Arceneaux K and Johnson M (2013) *Changing Minds or Changing Channels? Media Effects in the Era of Expanded Choice*. Chicago, IL: University of Chicago.

Aronow P, Kalla J, Orr L, et al. (2020) Evidence of rising rates of inattentiveness on lucid in 2020. SocArXiv. 15 September 15. DOI: 10.31235/osf.io/8sbe4.

Association of College Research Libraries (2016) *Framework for Information Literacy in Higher Education*. Chicago, IL: American Library Association.

Berinsky A (2004) *Silent Voices*. Princeton, NJ: Princeton University Press.

Brashier NM and Schacter DL (2020) Aging in an era of fake news. *Current Directions in Psychological Science* 29(3): 316–323.

Chia SC, Lu F and Sun Y (2021) Tracking the influence of misinformation on elderly people's perceptions and intention to accept COVID-19 vaccines. *Health Communication*. Epub ahead of print 23 September. DOI: 10.1080/10410236.2021.1980251.

Converse PE (1964) The nature of belief systems in mass publics. In: Apter DE (ed.) *Ideological and Discontent*. New York: Free Press, pp. 206–261.

Coppock A and McClellan OA (2019) Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents. *Research & Politics* 6(1): 1–14.

De Coninck D, Frissen T, Matthijs K, et al. (2021) Beliefs in conspiracy theories and misinformation about COVID-19: comparative perspectives on the role of anxiety, depression and exposure to and trust in information sources. *Frontiers in Psychology* 12: 646394.

Delli Carpini MX and Keeter S (1996) *What Americans Know about Politics and Why It Matters*. New Haven, CT: Yale University Press.

Evans JStBT, Barston JL and Pollard P (1983) On the conflict between logic and belief in syllogistic reasoning. *Memory & Cognition* 11(3): 295–306.

Feezell JT (2018) Agenda setting through social media: the importance of incidental news exposure and social filtering in the digital era. *Political Research Quarterly* 71(2): 482–494.

Fletcher R and Nielsen RK (2018) Are people incidentally exposed to news on social media? A comparative analysis. *New Media & Society* 20(7): 2450–2468.

Flynn DJ, Nyhan B and Reifler J (2017) The nature and origins of misperceptions: understanding false and unsupported beliefs about politics: nature and origins of misperceptions. *Advances in Political Psychology* 38: 127–150.

Gil de Zúñiga H, Weeks B and Ardèvol-Abreu A (2017) Effects of the news-finds-me perception in communication: social media use implications for news seeking and learning about politics. *Journal of Computer-Mediated Communication* 22(3): 105–123.

Guess A, Nagler J and Tucker J (2019) Less than you think: prevalence and predictors of fake news dissemination on Facebook. *Science Advances* 5(1): 1–8.

Guess AM, Nyhan B and Reifler J (2020) Exposure to untrustworthy websites in the 2016 US election. *Nature Human Behaviour* 4(5): 472–480.

Kam CD and Trussler MJ (2017) At the nexus of observational and experimental research: theory, specification, and analysis of experiments with heterogeneous treatment effects. *Political Behavior* 39(4): 789–815.

Karnowski V, Kümpel AS, Leonhard L, et al. (2017) From incidental news exposure to news engagement. How perceptions of the news post and news usage patterns influence engagement with news articles encountered on Facebook. *Computers in Human Behavior* 76: 42–50.

Kunda Z (1990) The case for motivated reasoning. *Psychological Bulletin* 108(3): 480–498.

Lewandowsky S, Ecker UKH, Seifert CM, et al. (2012) Misinformation and its correction: continued influence and successful debiasing. *Psychological Science in the Public Interest* 13(3): 106–131.

Liao QV and Fu W-T (2014) Age differences in credibility judgment of online health information. *ACM Transactions on Computer-Human Interaction* 21(1): 1–23.

Loftus EF and Pickrell JE (1995) The formation of false memories. *Psychiatric Annals* 25(12): 720–725.

Luca M, Munger K, Nagler J, et al. (2022) You won't believe our results! But they might: heterogeneity in beliefs about the accuracy of online media. *Journal of Experimental Political Science* 9: 267–277.

Lyons BA, Montgomery JM, Guess AM, et al. (2021) Overconfidence in news judgments is associated with false news susceptibility. *Proceedings of the National Academy of Sciences* 118(23): 1–10.

Mari S, Gil de Zúñiga H, Suerdem A, et al. (2022) Conspiracy theories and institutional trust: examining the role of uncertainty avoidance and active social media use. *Political Psychology* 43(2): 277–296.

Miller JM, Saunders KL and Farhart CE (2016) Conspiracy endorsement as motivated reasoning: the moderating roles of political knowledge and trust. *American Journal of Political Science* 60(4): 824–844.

Mitchell A, Gottfried J, Guskin E, et al. (2013) *The Role of News on Facebook*. Pew Research Center. Available at: https://www.journalism.org/wp-content/uploads/sites/8/2013/10/FINAL-Facebook-Survey-Report-10-22-13-1.pdf

Mitchell A, Gottfried J, Stocking G, et al. (2019) *Many Americans Say Made-Up News Is a Critical Problem That Needs to Be Fixed*. Pew Research Center. Available at: https://www.journalism.org/wp-content/uploads/sites/8/2019/06/PJ_2019.06.05_Misinformation_FINAL-1.pdf

Montgomery JM, Nyhan B and Torres M (2018) How conditioning on posttreatment variables can ruin your experiment and what to do about it. *American Journal of Political Science* 62(3): 760–775.

Munger K, Gopal I, Nagler J, et al. (2021) Accessibility and generalizability: are social media effects moderated by age or digital literacy? *Research & Politics* 8(2): 1–16.

Nyhan B and Reifler J (2010) When corrections fail: the persistence of political misperceptions. *Political Behavior* 32(2): 303–330.

O'Connor LG, Radcliff CJ and Gedeon JA (2001) Assessing information literacy skills: developing a standardized instrument for institutional and longitudinal measurement. In: Thompson HA (ed.) *Crossing the Divide: Proceedings of the Tenth National Conference of the Association of College and Research Libraries*. Chicago, IL: Chicago University Press, pp. 163–174.

Pew Research Center (2018) *The Public, the Political System and American Democracy*. Washington, DC: Pew Research Center. Available at: https://www.pewresearch.org/politics/wp-content/uploads/sites/4/2018/04/4-26-2018-Democracy-release-1.pdf

Prior M (2005) News vs. entertainment: how increasing media choice widens gaps in political knowledge and turnout. *American Journal of Political Science* 49(3): 577–592.

Seifert CM (2014) The continued influence effect: the persistence of misinformation in memory and reasoning following correction. In: Rapp DN and Braasch JLG (eds) *Processing Inaccurate Information: Theoretical and Applied Perspectives from Cognitive Science and the Educational Science*. Cambridge, MA: MIT Press, pp. 39–71.

Settle J (2018) *Frenemies: How Social Media Polarizes America*. Cambridge: Cambridge University Press.

Sheagley G and Clifford S (n.d.) No need for a panel study: measuring moderators doesn't alter treatment effects. Available at: http://scottaclifford.com/wp-content/uploads/2012/02/SheagleyClifford_NoNeedforPanelStudy.pdf

Tausanovitch C, Vavreck L, Reny T, et al. (2019) *Democracy Fund+ UCLA Nationscape: Methodology and Representativeness Assessment*. Voter Study Group. Available at: https://www.voterstudygroup.org/uploads/reports/Data/NS-Methodology-Representativeness-Assessment.pdf (accessed 31 March 2021).

Timberg C, Dwoskin E and Albergotti R (2021) Inside Facebook, Jan. 6 violence fueled anger, regret over missed warning signs. *Washington Post*, 22 October. Available at: https://www.wash

ingtonpost.com/technology/2021/10/22/jan-6-capitol-riot-facebook/ (accessed 10 January 2022).
Tsfati Y and Cappella JN (2003) Do people watch what they do not trust?: exploring the association between news media skepticism and exposure. *Communication Research* 30(5): 504–529.
Uscinski J, Enders AM, Klofstad C, et al. (2022) Cause and effect: on the antecedents and consequences of conspiracy theory beliefs. *Current Opinion in Psychology* 47: 101364.
Walter N and Tukachinsky R (2020) A meta-analytic examination of the continued influence of misinformation in the face of correction: how powerful is it, why does it happen, and how to stop it? *Communication Research* 47(2): 155–177.
Weeks BE, Lane DS, Kim DH, et al. (2017) Incidental exposure, selective exposure, and political information sharing: integrating online exposure patterns and expression on social media. *Journal of Computer-Mediated Communication* 22(6): 363–379.

Author biographies

Nicolas M Anspach is Associate Professor of Political Science at York College of Pennsylvania.

Taylor N Carlson is Assistant Professor of Political Science at Washington University in St. Louis.