Twitter's disputed tags may be ineffective at reducing belief in fake news and only reduce intentions to share fake news among Democrats and Independents

Jeffrey Lees^{1*}, Abigail McCarter², Dawn M. Sarno³

¹ Media Forensics Hub, Clemson University, & Andlinger Center for Energy and the Environment, Princeton University, ORCID: [0000-0001-6030-4207](https://orcid.org/0000-0001-6030-4207)

² Department of Political Science, Clemson University, ORCID: [0000-0003-4767-8188](https://orcid.org/0000-0003-4767-8188)

³ Department of Psychology, Clemson University, ORCID: [0000-0001-5605-5957](https://orcid.org/0000-0001-5605-5957)

* Corresponding Author: Jeffrey Lees (jeffrey.m.lees@gmail.com)

Published at the Journal of Online Trust and Safety

Cite as:

Lees, J., McCarter, A., & Sarno, D. M. (2022). Twitter's Disputed Tags May Be Ineffective at Reducing Belief in Fake News and Only Reduce Intentions to Share Fake News Among Democrats and Independents. *Journal of Online Trust and Safety*, 1(3).
<https://doi.org/10.54501/jots.v1i3.39>

Keywords: misinformation, Twitter, partisanship, fake news, cognitive reflection

Open Data & Materials: <https://osf.io/h65nv/>

This Preprint: <https://psyarxiv.com/cujdn>

Open Access Published Version: <https://tsjournal.org/index.php/jots/article/view/39/26>

DISPUTED TAGS

Abstract

Throughout the 2020 US elections, one of Twitter’s defenses against misinformation was its “This claim has been disputed” tags. The utility of such tags, however, remains unclear. A survey-based experiment, meant to simulate the Twitter environment, with a convenience sample of 318 US participants found that while disputed tags reduced the sharing of misinformation among Democrats and Independents, they had no effect on the sharing habits of Republicans and did not reduce belief in fake news for any group. We also found that higher scores on the Cognitive Reflection Test (a measure of analytical rather than intuitive thinking) correlated with lower belief in fake news, but had no relationship with sharing habits. Further, conservatism positively correlated with belief in and sharing intentions for tagged false headlines, but not untagged false headlines or true headlines. Our results suggest that the tags employed by Twitter to combat the spread of fake news may have been ineffective at reducing belief in fake news, and may only have attenuated fake news sharing among Democrats and Independents.

1. Introduction

The spread of misinformation about political (Jerit & Zhao, 2020), scientific (Maertens et al., 2020), and medical (Romer & Jamieson, 2020) issues has caused serious concerns among the public, policymakers, and scientists alike. The rapid proliferation of misinformation and fake news (false news headlines presented as legitimate) on social media platforms (Cinelli et al., 2020; Pennycook & Rand, 2021) has led many to argue that social media companies bear some responsibility for reducing the spread of misinformation on their platforms. For example, both Facebook (Guynn, 2017) and Twitter (Ortutay, 2021; Roth & Pickles, 2020) implemented, then ceased, using tags like “This claim about election fraud is disputed” or “This is disputed” (i.e., “disputed” tags) to identify posts and headlines deemed false or misleading (see Figure 1). Facebook cited a lack of effectiveness of “disputed” tags and concerns over “backfire” effects as reasons for shifting their misinformation mitigation efforts (Smith, 2017), despite ample research suggesting that fact-checking and other lightweight interventions reduce the sharing of misinformation (Lewandowsky & van der Linden, 2021; Pennycook, McPhetres, et al., 2020; Walter et al., 2020), and little evidence that backfire effects from fact-checking are real (Nyhan, 2021; Swire-Thompson et al., 2020; Wood & Porter, 2019). Twitter cited the need for “more context” as a reason for no longer using “disputed” tags (Twitter Support, 2021), choosing instead to label them with ostensibly clearer labels such as “misleading.”

Figure 1: Example disputed tags from Facebook, Twitter, and the present study.

A dominant view among psychological scientists is that a *lack of attentiveness*, not political bias, is what drives the sharing of misinformation on social media platforms (Epstein et al., 2021; Pennycook & Rand, 2021, 2019b). According to this account, partisans are able to discern false from true information independent of whether it is concordant or discordant with their ideology, yet a lack of attentiveness when it comes to *sharing* behaviors leads people to inadvertently share information they would correctly identify as false otherwise. Since people are motivated to be accurate in their social judgments (Pennycook, Epstein, et al., 2021), this account suggests that nudging people out of their inattentiveness and toward a reflective mindset will reduce the sharing of misinformation. For example, a recent large-scale experiment performed on Twitter found that priming Twitter users with an accuracy-mindset nudge reduced the sharing of political misinformation (Pennycook, Epstein, et al., 2021).

Given the evidence for the effectiveness of accuracy nudges, and that disputed tag warnings are enough to modestly reduce misinformation sharing on Facebook (Clayton et al., 2020), one could reasonably hypothesize that Twitter’s retired “disputed” tags successfully increased attentiveness and reduced misinformation sharing when it was deployed on Twitter. However, scholars have noted that these interventions are not equally effective across parties.

DISPUTED TAGS

Experiments examining accuracy nudges tend to find significantly lower levels of truth discernment among Republicans compared to Democrats (Gawronski, 2021). Moreover, a reanalysis of COVID-19-based accuracy nudge data from Pennycook et al. (2020) suggests that the effects are driven largely by Democrats, and that the nudges are least effective on Republicans (see Letter to the Editor in Pennycook, McPhetres, et al., 2021), although the original authors have disputed this claim directly (see response to Letter to the Editor in Pennycook & Rand, 2022). These debates parallel larger bodies of evidence suggesting significant asymmetries in partisans' fake news belief and consumption, where conservatives are generally more likely to believe and share fake news and misinformation relative to liberals (Garrett & Bond, 2021; Jost et al., 2018; Pereira et al., 2021).

In the present work we seek to examine the effectiveness of Twitter's "disputed tags" in reducing belief in and sharing of fake news headlines. In an online survey we asked Democrats, Independents, and Republicans to respond to a series of headlines in tweet form, which they viewed in random order. Ten of the headlines were true, 10 of the headlines were false and had a "this claim is disputed" tag, and 10 were untagged false headlines. For each headline, we asked participants if they would be willing to share the headline and whether they believed the headline was accurate. We also measured participants' levels of cognitive reflectiveness, impulsivity, and ideology. We included an impulsivity measure because users who demonstrate a lack of cognitive reflectiveness are likely to also be impulsive, as this relationship has been demonstrated in the phishing domain (Kumaraguru et al., 2007). We hypothesized that the disputed tags would be effective at reducing belief in and the willingness to share fake news, relative to untagged fake news, and that these effects would be greater for unreflective and impulsive participants.

DISPUTED TAGS

Contrary to our hypotheses, there was no difference in the perceived accuracy of untagged and tagged false headlines, suggesting that the presence of a “disputed” tag on the headline did not reduce belief in the accuracy of the headline. We did observe an effect of the tags on sharing intentions, such that participants were less likely to say they would share tagged versus untagged headlines. However, this effect only appeared for Democrats and Independents, but not for Republicans. We also found that participants' ideology, but not their levels of cognitive reflectiveness, interacted with the effect of the tags. Participants who were more conservative were more likely to share and believe false headlines with tags—relative to liberal participants—but no such correlation with ideology arose for true untagged headlines or false untagged headlines.

2. Method

2.1 Open Science

All materials, data, and analysis scripts needed to replicate this study are available online (<https://osf.io/h65nv/>). Data collection was preregistered (<https://osf.io/vj35r/>), as was a series of analyses testing confirmatory hypotheses. However, many of the results presented in this paper are from non-preregistered analyses, including all findings related to party identification and ideology. Details on the preregistered and non-preregistered analyses can be found in Section 2.6 below.

2.2 Participants

Data were collected June 11-15, 2021, using the Prolific survey platform (Palan & Schitter, 2018). Using Prolific's prescreen functions, we collected a convenience sample with an equal number of self-identified Democrats, Republicans, and Independents. Participants were not

DISPUTED TAGS

prescreened based on usage of Twitter, and the base rate of willingness to share in the experiment was 16% (the base rate of willingness to share these and other headlines from Pennycook et al., 2021 varied ~15-35%). We collected 328 responses, and 10 participants failed the attention check, leaving a final $N = 318$; $M_{\text{age}} = 33.3$, 34.6% Democrat, 32.7% Republican, 32.7% Independent, 49.7% female, 47.2% male, 3.1% non-binary/unlisted, 73.0% White. Self-reported ideology (1-10 Likert, “Very Liberal” - “Very Conservative”) was $M = 4.35$, $SD = 3.13$, indicating a slight liberal skew but overall well distributed across the ideological spectrum.

2.3 Statistical Power

We conducted Monte Carlo simulations, using the *simr* R package (Green & MacLeod, 2016), of the regression model testing the standardized effect of the False Tagged condition relative to the False Untagged condition on sharing intention, as this is of primary interest in the present research. The simulations were performed on the exact model reported below in the results, namely the main effects model without controls or interactions for party identification. Each sensitivity estimate was based on 1,000 simulations and an $\alpha = 0.05$. We found that the sharing intentions model was sensitive enough to detect a standardized effect of $\beta = -0.18$ with 81.2% power, 95% CI = [76.6%, 83.6%], and sensitive enough to detect a standardized effect of $\beta = -0.21$ with 89.7% power, 95% CI = [87.6%, 91.5%].

2.4 Procedure

This study utilized a 3x1 repeated measures within-subjects design, where each participant was exposed to all experimental conditions (tagged false, untagged false, and true headlines), and for each headline asked the same set of questions (i.e., repeated measures). Judd, Westfall, & Kenny (2017) provide a useful typology for understanding this design, which

DISPUTED TAGS

delineates participants, targets (here, headlines) and condition (see Table 2 therein). First, participants and condition are crossed (C), meaning participants are exposed to the tagged false, untagged false, and true headlines conditions. Second, headlines are nested (N) within condition, meaning that each unique headline is fixed to a condition. Third, participants and headlines are crossed (C), meaning that participants see all headlines. Judd, Westfall, & Kenny (2017) refer to this as a “CNC” design.

After providing informed consent and reading basic instructions, all participants responded to a combined 7-item numerical and non-numerical Cognitive Reflection Test (CRT) (Frederick, 2005; Thomson & Oppenheimer, 2016) and Version 11 of the 30-item Barrett Impulsivity Scale (BIS) (Patton et al., 1995) to measure participants’ levels of reflective thinking and impulsivity. Both of these measures assess an individual’s tendency to think critically about information and avoid rash decisions. CRT and BIS-11 order was counterbalanced and item ordering was randomized. Participants then responded to the attention check “Please, in the box below, write out the answer to the following math problem, capitalizing the first letter of your answer (e.g. ‘Eight’, not ‘eight’ or ‘8’). What does one plus three equal?”.

Afterward, participants read and rated the perceived reliability of 30 headlines in Tweet format, in randomized order. Ten headlines were true and untagged, 10 headlines were false and untagged, and 10 headlines were false and tagged “(!) This claim has been ‘disputed’” in accordance with Twitter’s tags (see Figure 2). Note “This claim has been disputed” is a genericized version of Twitter’s disputed tags, which varied widely in their real-world implementation (i.e., Twitter’s tags sometimes directly referenced the topic at hand, and were sometimes in the present tense).

Our 30 true and false headlines (see Appendix) were a randomly selected subset of a

DISPUTED TAGS

larger corpus of left-right balanced 76 true and 70 false political headlines from Pennycook et al. (2021). We generated our stimuli using Tweetgen (<https://www.tweetgen.com/create/tweet.html>). We then randomly determined which 10 of the 20 false headlines we would tag. For each headline participants were asked “To the best of your knowledge, how accurate is the claim in the above headline?” (1-6 Likert, “Extremely Inaccurate” - “Extremely Accurate”) and “If you were to see the above article on social media, how likely would you be to share it? (1-6 Likert, “Extremely Unlikely” - “Extremely Likely”), both adopted from Pennycook et al. (2021).

Figure 2: Example a. True headline, b. Tagged False Headline, and c. Untagged False Headline, all with the accompanying measures of perceived accuracy and sharing intentions.

After responding to all 30 headlines, participants provided basic demographic information and proceeded to a debrief slide. In the debrief, participants were informed that many of the headlines they saw were false, and participants were shown the images of the true headlines. Participants had the option to also view the false headlines. 30.5% of participants opted to view the false headlines before leaving the survey. By party, the proportion of participants who chose to view the false headlines was 39.1% of Democrats, 30.8% of

DISPUTED TAGS

Independents, and 21.2% of Republicans, and the pairwise contrast between Democrats and Republicans was statistically significant, $p_{tukey} = 0.012$. Viewing the false headlines was not associated with sharing intentions of the false (tagged and untagged) headlines ($b = 0.02$, $p = 0.024$), but was positively associated with belief in the false headlines ($b = 0.19$, $p = 0.029$), meaning that participants who were more likely to believe fake news were also more likely to view the list of false headlines in the debrief.

2.5 Analyses

All analyses were conducted using restricted maximum likelihood linear mixed-effects modeling with crossed random intercepts for participants and headline to model the hierarchical dependencies within the data, as all judgments are nested within-participant and nested within-headline. We take a maximal random structure approach to modeling random slopes, and per recent guidelines (Barr et al., 2013; Brauer & Curtin, 2018) we model random slopes for all within-unit predictors, namely random slopes for condition within participant, and random slopes for individual difference measures within stimuli. If models would not converge under any available optimizers then we began removing secondary random slopes. All p -values were derived through Welch-Satterthwaite degrees of freedom approximation using the *lmerTest* R package (Kuznetsova et al., 2017), and all post-hoc tests utilized the Tukey method for p -value correction. Perceived accuracy, sharing intention, CRT, and BIS-11 responses were z -scored to allow for comparisons of standardized effects. The fixed effect for within-subjects Condition (True, False Untagged, False Tagged) and the fixed effect for party identification (Republican, Democrat, Independent) were sum coded.

2.6 Preregistered Hypotheses and Non-preregistered Analyses

All hypotheses were confirmatory and preregistered (<https://osf.io/vj35r>). Hypotheses 1a-b predicted that perceived accuracy (1a) and willingness to share (1b) would be highest for True Headlines and lowest for Tagged False Headlines, and Non-Tagged False Headlines would be in between, with all contrasts being significantly different. Hypotheses 2a-b predicted that the main effects predicted by H1a-b would be significantly moderated by cognitive reflectiveness, such that the difference in perceived accuracy (2a) and willingness to share (2b) between Untagged and Tagged False Headlines would be greater for participants lower in cognitive reflectiveness relative to those higher in cognitive reflectiveness. Hypotheses 3a-b predicted that main effects predicted by H1a-b would be significantly moderated by impulsivity, such that the difference in perceived accuracy (3a) and willingness to share (3b) between Untagged and Tagged False Headlines would be greater for participants higher in impulsivity relative to those lower in impulsivity.

Any and all analyses that are not stated as numbered hypotheses were not preregistered. Notably, none of the analyses below investigating party identification or ideology were preregistered. Such analyses were pursued in response to a lack of support for many of the stated hypotheses. Additionally, the preregistration erroneously neglected to specify the modeling of random intercepts for headline, and erroneously stated that participants see 15 headlines, rather than 30.

3 Results

3.1 Belief

In partial support for Hypothesis 1a, across all partisan groups we found that participants

DISPUTED TAGS

perceived both tagged ($\beta = -0.80, p < 0.001$) and untagged ($\beta = -0.76, p < 0.001$) as less accurate than true headlines; however, there was no difference between untagged and tagged false headlines ($\beta = -0.04, p = 0.766$). To test Hypothesis 2a we interacted participants' performance on the Cognitive Reflection Test (CRT) with condition and observed a significant interaction ($p = 0.008$). This interaction indicated that the CRT-belief slope was different by condition. We found that CRT scores negatively correlated with belief in tagged ($\beta = -0.07, p = 0.011$), but not associated with belief in untagged headlines ($\beta = -0.05, p = 0.072$) or belief in the true headlines ($\beta = 0.05, p = 0.057$). At no level of cognitive reflectiveness (± 1 SD) did we observe a difference in belief between tagged and untagged fake news. Rather, the gap between belief in true and false headlines grew as participants' levels of cognitive reflectiveness increased; see Figure 3 below. This overall suggests that cognitive reflectiveness's role in the effect of disputed tags was minimal. We also found no support for Hypothesis 3a, as self-reported impulsivity (BIS rating) was not associated with belief ($\beta = 0.04, p = 0.061$), nor did it interact with condition (all $ps > 0.160$). The full regression tables for these analyses, and all analyses presented below, can be found in Appendix B.

Figure 3: Predicted mean values, with 95% confidence intervals, of belief in headline by score on the cognitive reflection task. Dotted black line at 3.5 represents the mean threshold of perceived as accurate (above 3.5) or inaccurate (below 3.5).

The main effect of condition on belief was qualified by two significant interactions with party identification ($p = 0.039$ and $p = 0.009$). These interactions did not indicate that the tags were effective for any political subgroup; instead, they indicated that the tags induced partisan correlations within condition. All pairwise contrasts utilized Tukey-adjusted p -values. Among tagged false headlines, Republicans were more likely to perceive them as accurate relative to Democrats ($\beta = 0.26, p < 0.001$) and Independents ($\beta = 0.25, p = 0.002$); however, no such differences were observed for untagged false headlines between Republicans and Democrats ($\beta = 0.09, p = 0.843$) or Independents ($\beta = 0.02, p = 0.100$), suggesting the tags induced politically motivated patterns of belief. See Figure 4 for a plotting of the belief effects by party.

The same pattern of politically motivated belief was observed for the effects of

DISPUTED TAGS

self-reported ideology on belief. Ideology interacted with the main effect of tags ($p = 0.002$), and ideology was only associated with belief for tagged false headlines ($\beta = 0.16, p = 0.005$) such that individuals who identified as left-leaning were less likely to believe tagged false headlines than right-leaning individuals. No such linear relationship was observed for untagged false headlines ($\beta = -0.02, p = 0.664$) or true headlines ($\beta = -0.06, p = 0.244$), indicating that the presence of the tags induced politically motivated belief, similar to the pattern we observe with belief; see Figure 5 below for simple slopes. For no party or level of ideology was the contrast between belief in untagged vs. tagged false headlines significant.

Figure 4: Predicted mean values, with 95% confidence intervals, of belief in headline by party identification. Dotted black line at 3.5 represents the mean threshold of perceived as accurate (above 3.5) or inaccurate (below 3.5).

Figure 5: Simple slopes, with 95% confidence intervals, of belief in headline and ideology. Higher values on ideology measure mean more conservative, lower values more liberal. Data points are jittered.

3.2 Sharing

Across all partisan groups we observed a main effect of disputed tags on sharing intentions. In partial support for Hypothesis 1b, participants were significantly less likely to share tagged false headlines compared to true headlines ($\beta = -0.24, p = 0.002$) and untagged true false headlines compared to true headlines ($\beta = -0.13, p = 0.0495$), but there was no difference between untagged and tagged headlines ($\beta = 0.11, p = 0.219$). However, the main effect of tags was qualified by a significant interaction of party identification ($p < 0.001$). All pairwise contrasts utilized Tukey-adjusted p -values. Tags reduced sharing intentions relative to true headlines for Democrats ($\beta = -0.37, p < 0.001$) and Independents ($\beta = -0.29, p < 0.001$), but not Republicans ($\beta = -0.05, p = 0.802$). When comparing untagged vs tagged false headlines,

DISPUTED TAGS

Democrats were less likely to share false tagged headlines ($\beta = 0.20, p = 0.014$), but not Independents ($\beta = 0.15, p = 0.078$) or Republicans ($\beta = -0.03, p = 0.871$). See Figure 3 for a plotting of the sharing effects by party.

In examining the effects of cognitive reflectiveness and impulsivity, we find no support for Hypotheses 2b or 3b. There was no association of CRT performance with sharing ($\beta = 0.07, p = 0.079$), nor did CRT performance interact with condition effects (all $ps > 0.270$). Similarly, there was no association of BIS rating with sharing ($\beta = 0.03, p = 0.480$), nor did BIS interact with condition effects (all $ps > 0.201$).

Figure 6: Predicted mean values, with 95% confidence intervals, of sharing intentions by party identification. Dotted black line at 3.5 represents the mean threshold of willingness to share a headline (above 3.5) or unwillingness (below 3.5).

To examine the role of self-reported left-right ideology, we interacted ideology with

DISPUTED TAGS

condition and found that the main effect of tags was qualified by a significant interaction ($p = 0.031$). Ideology was only associated with sharing intentions for tagged false headlines ($\beta = 0.15$, $p = 0.014$), such that individuals who identified as left-leaning were less likely to say they would share tagged false headlines than right-leaning individuals. No such relationship was observed for untagged false headlines ($\beta = 0.04$, $p = 0.524$) or true headlines ($\beta = -0.01$, $p = 0.831$), suggesting that the presence of the disputed tags caused participants to respond in a fashion aligned with their ideology—namely that liberals were less likely to express sharing intentions and conservatives more likely. We interpret this as a form of politically motivated sharing intentions. Conservatives are less trusting of social media companies than liberals (Pennycook & Rand, 2019a), and we theorize that this difference in trust leads to differing evaluations of how Twitter’s disputed tags may be pro-liberal/anti-conservative, leading to the observed correlation between ideology and sharing intention in the Tagged False Headline condition. See Figure 7 for simple slopes.

DISPUTED TAGS

Figure 7: Simple slopes, with 95% confidence intervals, of sharing intentions and ideology. Higher values on ideology measure mean more conservative, lower values more liberal. Data points are jittered.

In summary, while tags reduced sharing intentions relative to true (but not untagged false) among Democrats and Independents, these findings are consistent with the interpretation that disputed tags induce politically motivated intentions to share.

3.3 Headline Bias

One concern regarding our findings is that the randomized selection of headlines from the larger corpus of fake and real news headlines from Pennycook et al. (2021) may have inadvertently introduced confounding partisan biases. For example, if by chance the headlines in the False Tagged Condition were more pro-Republican, on average, than the headlines in the False Untagged Condition, that could explain observed differences in partisans' sharing

DISPUTED TAGS

intentions across the conditions, rather than said differences being the result of the tag itself. To address this potential confound, we examined the partisanship scores generated by Pennycook et al. (2021) in their testing of their corpus of headlines. Specifically, we used the “partisanship combined” score of their Democratic and Republican participants’ response to the question “Assuming the above headline is entirely accurate, how favorable would it be to Democrats versus Republicans?” with the response options being “1. More favorable for Democrats, 2. Moderately more favorable for Democrats, 3. Slightly more favorable for Democrats, 4. Slightly more favorable for Republicans, 5. Moderately more favorable for Republicans, 6. More favorable for Republicans.” As such, a value greater than 3.5 means the headline was viewed as pro-Republican on average, and a value less than 3.5 means the headline was viewed as pro-Democrat on average. The partisanship score for each headline used in this experiment can be found in Appendix A.

We then took these scores and added them as fixed-effects to every model reported in this manuscript, to examine if our results were robust to controlling for whether the headlines were pro-Democrat or pro-Republican. All models and statistical tests regarding the partisanship of the headlines are reported in Appendix B.

In summary, our results were fully robust to controlling for headline partisanship. In no model was headline partisanship significantly associated with belief or sharing intentions. In no model did the addition of the headline partisanship predictor significantly improve model fit, and in no model did the inclusion of headline partisanship meaningfully change the other observed relationships (e.g., it never caused a previously significant coefficient to become non-significant). As such, the reported findings in this manuscript are robust to the variance in partisanship across the text of the headlines.

4. Discussion

The present findings suggest that Twitter's disputed tags may have been relatively ineffective for many Twitter users. Although participants perceived all false headlines as less accurate than true headlines, the presence of the disputed tag did not reduce belief in false headlines.

Sharing intentions presented a slightly different story. Participants were less likely to say they would share tagged false headlines relative to true headlines, but disputed tags only reduced sharing intention relative to untagged false headlines among liberal participants.

Additionally, while performance on the cognitive reflectiveness test was related to participants' ability to identify tagged false headlines, cognitive reflectiveness was not associated with how participants responded to untagged false headlines, nor was there any evidence the tags were more effective on those low in cognitive reflectiveness vs. those high in cognitive reflectiveness. Because individuals with low cognitive reflectiveness tend to need the most help in identifying misinformation (Pennycook & Rand, 2019b), ideal interventions might target those individuals. Twitter's "disputed" tags did not appear to do this. Notably, cognitive reflectiveness and self-reported impulsivity were not observably associated with sharing intentions for any headlines, false or true.

Republicans were more likely to perceive tagged false headlines as accurate relative to Democrats and Independents, and right-leaning individuals were more likely to believe the tagged false headlines compared to their left-leaning counterparts. These associations between partisanship/ideology and belief in/sharing of the headlines *only* arose for tagged false headlines. No such relationships were observed for true headlines or untagged false headlines. We interpret this as evidence for politically motivated sharing intentions and belief in response to the presence

DISPUTED TAGS

of a tag, and theorize that it is rooted in differing levels of trust toward social media platforms between liberals and conservatives (Pennycook & Rand, 2019a). For a broader discussion of how “politically motivated” processes can be operationalized in the context of fake news research, see Pennycook & Rand (2021). Taken together, these findings suggest that the disputed tags may provide modest benefits for reducing sharing of false headlines for Democrats and Independents, but do not influence Republicans.

Overall, the disputed tags appeared to only minimally reduce sharing intentions and in contrast with other work (e.g., Pennycook et al., 2020, 2021; Yaquib et al., 2020) interacted with ideology rather than reflectiveness to produce those effects. Specifically, disputed tags only reduced sharing intentions for Democrats and Independents, not Republicans. Ideology demonstrated a similar association with sharing intentions, with left-leaning participants less likely to share tagged false headlines than right-leaning participants. Similar patterns were found for perceived accuracy for the headlines. Notably, we do replicate the overall pattern of past work: participants' ability to discern true from false headlines does not translate into differences in sharing intentions.

4.1 Contributions and Limitations

A notable contribution of the present work is the use of a fully within-subject design, compared to most past work, which has used between-subjects manipulations (e.g., Clayton et al., 2020; Pennycook et al., 2021). Such past work has also used stronger stimuli (i.e., more information, visually bolder) compared to the more subtle Twitter tags. As such, the present study possessed a level of verisimilitude that past work often lacked, which may explain our divergence from past findings. For example, Clayton et al. (2020) used a between-subjects experiment to investigate the effect of Facebook’s disputed tags on the perceived accuracy of

DISPUTED TAGS

fake news, and observed a moderate and significant standardized effect of 0.26. Conversely, using our within-subjects design, we observed a very small and insignificant standardized effect of 0.05, and were powered at 90% to detect standardized effects of 0.21 or greater. Most Twitter users do not view a single fact-check, then a single tweet of which they then make a judgment. Rather, they view a stream of tweets in which any potential fact-check (or other intervention) is embedded. Our results suggest that between-subject experiments where participants do not have to discern between tagged and untagged stimuli may be inflating the effect sizes of tags, and similar warnings, on belief in and willingness to share fake news. Nonetheless, our experiment was conducted in a controlled and contrived setting, so caution is needed in generalizing to how individuals respond to misinformation tags “naturally” on Twitter itself.

Despite a relatively consistent pattern of results, we caution against strong inferences to the null. While our experiment was statistically powered to reliably detect small effects by conventional definitions (i.e., $d = 0.2$), we were unable to detect *exceptionally* small effects, such as $< d = 0.10$, which have been observed in some past studies on fact-checking (e.g., Clayton et al., 2020). Defenders of interventions with such small effects have provided compelling evidence that even when an effect is very small, if that intervention is distributed to millions of individuals at low costs then it is still beneficial (e.g., Bond et al., 2012). That logic would certainly apply to Twitter’s disputed tags (and their current “misleading” tags), and we do not rule out the possibility that the tags do have a very small beneficial effect, as our study was not powered to detect the small effect sizes observed in Bond et al. (2012). However, our results suggest that any such positive effects may occur alongside an induction of politically motivated processes. Nonetheless, our findings related to party identification and ideology were from non-preregistered analyses, which were pursued largely because many of our stated hypotheses

DISPUTED TAGS

were not supported. As such, future work should seek to replicate and extend the present findings via confirmatory tests.

Another limitation of the current work is that it may have artificially reduced base rates of sharing intentions. In choosing to use previously published stimuli, we necessarily used tweets that may have seemed old to participants, as all tweets were dated to September 2020 and the true headlines were from that month, yet participants partook in June of 2021. Also, the accuracy nudge literature finds that simply asking people to think about accuracy reduces sharing intentions (for an overview, see Pennycook & Rand, 2021), and therefore asking both perceived accuracy and sharing intentions simultaneously may have reduced sharing intentions. While these design decisions may have reduced base rates of sharing intentions, we argue that it is unlikely they would alter the relationship between sharing intentions and cognitive reflectiveness. It is also unlikely that such a potential effect would have varied across political subgroups in our sample (e.g., it absorbed the effect of the tags for Republicans only), as the literature from which our stimuli and dependent variables were drawn has consistently argued that the effectiveness of accuracy nudges is invariant to political ideology (see Pennycook & Rand, 2022). Moreover, the possibility that floor effects are concealing an association between sharing intention and cognitive reflection is unlikely, because we do observe linear associations between ideology and sharing intentions.

Similarly, a concern with our fully within-subjects design is that the first presence of the tag effectively acts as an accuracy nudge, which then affects responses to all subsequent headlines, such that it would suppress the effect of subsequent tags relative to untagged false headlines on participant responses. Probabilistically, 92.4% of participants saw a disputed tag within the first five (of 30) headlines, and because headline order was completely randomized,

DISPUTED TAGS

robustly examining possible order effects of the first tag participants see was not plausible. While this design concern is reasonable, work on the implied truth effect would suggest the opposite effect. For example, attaching warning labels to some fake news headlines can lead to unlabeled fake news to be seen as more accurate (Pennycook, Bear, et al., 2020), which we do not observe. More broadly, we argue that our main effects are largely consistent with the accuracy nudge literature, and the ways in which our results diverge from the accuracy nudge literature do not suggest that the tags are suppressing effects one would predict based on past accuracy nudge findings. Congruent with the accuracy nudge literature, we found more evidence that our tags affect sharing intentions more than perceived accuracy (for overview, see Pennycook & Rand, 2022). Unlike the accuracy nudge literature, we found that party identification and ideology were more associated with sharing intentions and perceived accuracy of tagged fake news than cognitive reflection. It is unlikely this latter finding is the result of spillover effects from participants having seen repeated tags. Rather, we interpret this finding as evidence that conservatives' decreased trust in social media platforms relative to liberals (Pennycook & Rand, 2019a) led to politically motivated sharing intentions and belief in response to headlines containing the tag.

Twitter's removal of disputed tags echoes the present findings, that "this claim has been disputed" tags may have been ineffective at reducing belief in fake news. Twitter is now testing labels such as "stay informed" or "misleading" with colored icons indicating the severity of the misinformation (e.g., red exclamation point for dangerous misinformation) (Ortutay, 2021). These tags may be more effective at reducing fake news sharing due to increased saliency and specific messaging (Braun et al., 1995). A lack of understanding was thought to be a critical issue with the disputed tags, with users confused on who "disputed" the information (Ortutay,

DISPUTED TAGS

2021). Nonetheless, it is important not to exaggerate the differences between “this tweet is misleading” and “this tweet is disputed.” While it is plausible that this small change in language decreased confusion, and therefore increased the effectiveness of the tags, it is much less plausible that the “misleading” tags are interacting with different psychological mechanisms than the “disputed” tags, given their otherwise similar characteristics. If both styles of tags work through the same psychological mechanisms, the findings presented herein are likely generalizable to the “misleading” tags Twitter is now using.

The present study only explored headlines from fake news websites, yet Twitter users can also post misinformation unattached to a headline (e.g., Twitter placed “disputed” tags on several posts of then U.S. President Donald Trump). Our results may not generalize to contexts where the source of the fake news is perceived as (in)congruent with one’s ideology (Traberg & van der Linden, 2022), such as when fake news is shared by a politician (Swire-Thompson et al., 2020). Moreover, given the growing evidence that fake news is overall rare in the social media ecosystem (Allen et al., 2020; Guess et al., 2019), and that elite rhetoric may be a primary driver of belief in misinformation (Clayton et al., 2021), future research should explore how the source of misinformation may interact with disputed tags.

In the phishing context, similar warnings regarding website authenticity are more effective when the user is trained on the rationale behind the warning (Yang et al., 2017). Thus, Twitter users may require more formal training to understand the purpose of the new tags rather than simply relying on participants to be more reflective when a tag is present. Additionally, given the partisan response to tagged headlines, it is possible that interventions will be ineffective for Republicans, who may be suspicious of any attempts to label misinformation and perceive such efforts as anti-conservative. Lastly, users may begin to ignore tags over time,

DISPUTED TAGS

especially in situations of low reliability (i.e., incorrectly tagging true information) (Parasuraman & Riley, 1997). This may have contributed to the failure of the disputed tags, and presents a challenge for future interventions, regardless of their initial efficacy.

In conclusion, we find mixed evidence for the effectiveness of Twitter's disputed tags, and all potential benefits lie in their effect on Democrats and Independents. For Republicans, tags did not affect sharing intentions or belief in false headlines, and instead we found evidence that tags led liberals and conservatives to respond in a politically motivated fashion. We also found little evidence that cognitive reflectiveness or impulsivity played a meaningful role in the impact the disputed tags had on participants.

References

- Allen, J., Howland, B., Mobius, M., Rothschild, D., & Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14), eaay3539. <https://doi.org/10.1126/sciadv.aay3539>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415), 295–298. <https://doi.org/10.1038/nature11421>
- Brauer, M., & Curtin, J. J. (2018). Linear mixed-effects models and the analysis of nonindependent data: A unified framework to analyze categorical and continuous independent variables that vary within-subjects and/or within-items. *Psychological Methods*, 23(3), 389–411. <https://doi.org/10.1037/met0000159>
- Braun, C. C., Mine, P. B., & Clayton Silver, N. (1995). The influence of color on warning label perceptions. *International Journal of Industrial Ergonomics*, 15(3), 179–187. [https://doi.org/10.1016/0169-8141\(94\)00036-3](https://doi.org/10.1016/0169-8141(94)00036-3)
- Cinelli, M., Quattrocioni, W., Galeazzi, A., Valensise, C. M., Brugnoli, E., Schmidt, A. L., Zola, P., Zollo, F., & Scala, A. (2020). The COVID-19 social media infodemic. *Scientific Reports*, 10(1), 16598. <https://doi.org/10.1038/s41598-020-73510-5>
- Clayton, K., Blair, S., Busam, J. A., Forstner, S., Glance, J., Green, G., Kawata, A., Kovvuri, A., Martin, J., Morgan, E., Sandhu, M., Sang, R., Scholz-Bright, R., Welch, A. T., Wolff, A. G., Zhou, A., & Nyhan, B. (2020). Real solutions for fake news? Measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Political Behavior*, 42(4), 1073–1095. <https://doi.org/10.1007/s11109-019-09533-0>
- Clayton, K., Davis, N. T., Nyhan, B., Porter, E., Ryan, T. J., & Wood, T. J. (2021). Elite rhetoric can undermine democratic norms. *Proceedings of the National Academy of Sciences*, 118(23), e2024125118. <https://doi.org/10.1073/pnas.2024125118>
- Epstein, Z., Berinsky, A. J., Cole, R., Gully, A., Pennycook, G., & Rand, D. G. (2021). Developing an accuracy-prompt toolkit to reduce COVID-19 misinformation online. *Harvard Kennedy School Misinformation Review*. <https://doi.org/10.37016/mr-2020-71>
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4), 25–42. <https://doi.org/10.1257/089533005775196732>
- Garrett, R. K., & Bond, R. M. (2021). Conservatives' susceptibility to political misperceptions. *Science Advances*, 7(23), eabf1234. <https://doi.org/10.1126/sciadv.abf1234>
- Gawronski, B. (2021). Partisan bias in the identification of fake news. *Trends in Cognitive Sciences*, 25(9), 723–724. <https://doi.org/10.1016/j.tics.2021.05.001>
- Green, P., & MacLeod, C. J. (2016). simr: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493–498. <https://doi.org/10.1111/2041-210X.12504>
- Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, 5(1), eaau4586. <https://doi.org/10.1126/sciadv.aau4586>
- Guyonn, J. (2017, December 21). Facebook fail: Social network scraps “disputed” flags on “fake

- news.” *USA Today*.
<https://www.usatoday.com/story/tech/2017/12/21/facebook-fail-social-network-scraps-disputed-flags-fake-news/973153001/>
- Jerit, J., & Zhao, Y. (2020). Political misinformation. *Annual Review of Political Science*, 23(1), 77–94. <https://doi.org/10.1146/annurev-polisci-050718-032814>
- Jost, J. T., van der Linden, S., Panagopoulos, C., & Hardin, C. D. (2018). Ideological asymmetries in conformity, desire for shared reality, and the spread of misinformation. *Current Opinion in Psychology*, 23, 77–83. <https://doi.org/10.1016/j.copsyc.2018.01.003>
- Judd, C. M., Westfall, J., & Kenny, D. A. (2017). Experiments with more than one random factor: Designs, analytic models, and statistical power. *Annual Review of Psychology*, 68(1), 601–625. <https://doi.org/10.1146/annurev-psych-122414-033702>
- Kumaraguru, P., Rhee, Y., Sheng, S., Hasan, S., Acquisti, A., Cranor, L. F., & Hong, J. (2007). Getting users to pay attention to anti-phishing education: Evaluation of retention and transfer. *Proceedings of the Anti-Phishing Working Group's 2nd Annual ECrime Researchers Summit*, 70–81.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1–26. <https://doi.org/10.18637/jss.v082.i13>
- Lewandowsky, S., & van der Linden, S. (2021). Countering misinformation and fake news through inoculation and prebunking. *European Review of Social Psychology*, 1–38. <https://doi.org/10.1080/10463283.2021.1876983>
- Maertens, R., Anseel, F., & van der Linden, S. (2020). Combatting climate change misinformation: Evidence for longevity of inoculation and consensus messaging effects. *Journal of Environmental Psychology*, 70, 101455. <https://doi.org/10.1016/j.jenvp.2020.101455>
- Nyhan, B. (2021). Why the backfire effect does not explain the durability of political misperceptions. *Proceedings of the National Academy of Sciences*, 118(15), e1912440117. <https://doi.org/10.1073/pnas.1912440117>
- Ortutay, B. (2021, July 1). What’s in a tag? Twitter revamps misinformation labels. *The Associated Press*.
<https://apnews.com/article/health-coronavirus-pandemic-election-2020-misinformation-technology-37cee761758f1072f91e1a362f769f5e>
- Palan, S., & Schitter, C. (2018). Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 17, 22–27. <https://doi.org/10.1016/j.jbef.2017.12.004>
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253.
- Patton, J. H., Stanford, M. S., & Barratt, E. S. (1995). Factor structure of the barratt impulsiveness scale. *Journal of Clinical Psychology*, 51(6), 768–774.
- Pennycook, G., Bear, A., Collins, E. T., & Rand, D. G. (2020). The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings. *Management Science*, 66(11), 4944–4957. <https://doi.org/10.1287/mnsc.2019.3478>
- Pennycook, G., Binnendyk, J., Newton, C., & Rand, D. G. (2021). A practical guide to doing behavioral research on fake news and misinformation. *Collabra: Psychology*, 7(1), 25293. <https://doi.org/10.1525/collabra.25293>

- Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., & Rand, D. G. (2021). Shifting attention to accuracy can reduce misinformation online. *Nature*. <https://doi.org/10.1038/s41586-021-03344-2>
- Pennycook, G., McPhetres, J., Zhang, Y., Lu, J. G., & Rand, D. G. (2020). Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention. *Psychological Science*, 31(7), 770–780.
- Pennycook, G., McPhetres, J., Zhang, Y., Lu, J. G., & Rand, D. G. (2021, December 24). *sj-docx-3-pss-10.1177_0956797620939054 – Supplemental material for Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for a Scalable Accuracy-Nudge Intervention* [Letter to the Editor]. <https://doi.org/10.25384/SAGE.12594110.v2>
- Pennycook, G., & Rand, D. (2021). The psychology of fake news. *Trends in Cognitive Sciences*, 15. <https://doi.org/10.1016/j.tics.2021.02.007>
- Pennycook, G., & Rand, D. G. (2019a). Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*, 116(7), 2521–2526. <https://doi.org/10.1073/pnas.1806781116>
- Pennycook, G., & Rand, D. G. (2019b). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188, 39–50. <https://doi.org/10.1016/j.cognition.2018.06.011>
- Pennycook, G., & Rand, D. G. (2022). Accuracy prompts are a replicable and generalizable approach for reducing the spread of misinformation. *Nature Communications*, 13(2333). <https://doi.org/10.31234/osf.io/v8ruj>
- Pereira, A., Harris, E., & Bavel, J. J. V. (2021). Identity concerns drive belief: The impact of partisan identity on the belief and dissemination of true and false news. *Group Processes*, 24.
- Romer, D., & Jamieson, K. H. (2020). Conspiracy theories as barriers to controlling the spread of COVID-19 in the U.S. *Social Science & Medicine*, 263, 113356. <https://doi.org/10.1016/j.socscimed.2020.113356>
- Roth, Y., & Pickles, Ni. (2020, May 11). *Updating our approach to misleading information*. Twitter Blog. https://blog.twitter.com/en_us/topics/product/2020/updating-our-approach-to-misleading-information
- Smith, J. (2017, December 20). Designing against misinformation. *Design at Meta*. <https://medium.com/designatmeta/designing-against-misinformation-e5846b3aa1e2>
- Swire-Thompson, B., DeGutis, J., & Lazer, D. (2020). Searching for the backfire effect: Measurement and design considerations. *Journal of Applied Research in Memory and Cognition*, 9(3), 286–299. <https://doi.org/10.1016/j.jarmac.2020.06.006>
- Swire-Thompson, B., Ecker, U. K. H., Lewandowsky, S., & Berinsky, A. J. (2020). They might be a liar but they're my liar: Source evaluation and the prevalence of misinformation. *Political Psychology*, 41(1), 21–34. <https://doi.org/10.1111/pops.12586>
- Thomson, K. S., & Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. *Judgment and Decision Making*, 11(1), 15.
- Traberg, C. S., & van der Linden, S. (2022). Birds of a feather are persuaded together: Perceived source credibility mediates the effect of political bias on misinformation susceptibility. *Personality and Individual Differences*, 185, 111269. <https://doi.org/10.1016/j.paid.2021.111269>

DISPUTED TAGS

- Twitter Support, @TwitterSupport. (2021, July 1). Last year, we started using labels to let you know when a Tweet may include misleading information. For some of you on web, we'll be testing a new label design with more context to help you better understand why a Tweet may be misleading [Tweet]. *Twitter*.
<https://twitter.com/TwitterSupport/status/1410646566885601280>
- Walter, N., Cohen, J., Holbert, R. L., & Morag, Y. (2020). Fact-checking: A meta-analysis of what works and for whom. *Political Communication*, 37(3), 350–375.
<https://doi.org/10.1080/10584609.2019.1668894>
- Wood, T., & Porter, E. (2019). The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior*, 41(1), 135–163.
<https://doi.org/10.1007/s11109-018-9443-y>
- Yang, W., Xiong, A., Chen, J., Proctor, R. W., & Li, N. (2017). Use of phishing training to improve security warning compliance: Evidence from a field experiment. *Proceedings of the Hot Topics in Science of Security: Symposium and Bootcamp on - HoTSoS*, 52–61.
<https://doi.org/10.1145/3055305.3055310>
- Yaqub, W., Kakhidze, O., Brockman, M. L., Memon, N., & Patil, S. (2020). Effects of credibility indicators on social media news sharing intent. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–14.
<https://doi.org/10.1145/3313831.3376213>

Authors

Jeffrey Lees was a Visiting Assistant Professor in the Department of Economics at Clemson University at the time this research was conducted, and is currently an Associate Research Scholar at Princeton University's Andlinger Center for Energy and the Environment. **Abigail McCarter** was an undergraduate student at Clemson University at the time this research was conducted, and is currently the Development and Impact Operations Manager at the Truman Center for National Policy and Truman National Security Project, and a master's candidate at Johns Hopkins School of Advanced International Studies. **Dawn Sarno** is an Assistant Professor of Psychology at Clemson University.

Acknowledgements

We thank all participants in the Spring 2021 Clemson University course *Phishing for Trolls, Bots, and Hackers: The Generalized Analysis of Online Inauthenticity* for their feedback on study design, Dave Rand for feedback on early versions of the manuscript, and attendees to the Watt Family Center's April Fools Symposium at Clemson University and the Symposium on Uncommon Yet Consequential Online Harms at the Stanford University's Internet Observatory for their helpful comments and questions on our findings.

Data Availability Statement

All data, analysis code, and study materials are publicly available on the Open Science Framework at <https://osf.io/h65nv/>. Jeffrey Lees prepared all the open materials, code, and data; do not hesitate to reach out to them with any questions.

Funding Statement

This project was supported by the Clemson University Creative Inquiry program, through the Clemson University Media Forensics Hub. All funds were used to compensate participants.

Competing Interests

The authors declare no competing interests.

Ethical Standards

This research was approved by Clemson University's Institutional Review Board. All participants provided informed consent to participate. No deception was utilized in this research, and all participants were debriefed at the end of the study where the false nature of the headlines was made explicitly clear.

Appendix A: Headline Texts

News Headlines	Headline Reference # ¹	Partisanship Score ²
<i>True Headlines</i>		
Facts did not justify wiretap of Trump aide	01	3.96
DC sues inaugural committee for overpaying.	02	2.635
Amid signs of trouble for Trump, Republican mega-donors scramble to retain the senate	03	3.05
A Virginia House candidate campaigns by attacking a man dressed as a woman.	04	3.025
Biden leads Trump in three key battleground states	05	2.47
Apple, Facebook, and other major companies denounce Trump visa restrictions.	06	2.615
Biden urged not to debate Trump	07	3.195
Trump tax break under investigation.	08	2.565
Conservative groups aggressively oppose socialized health care	09	3.515
Trump allies handing out cash to Black voters.	10	3.225
<i>False Tagged Headlines</i>		
New black panthers come dangerously close to President Trump	11	3.64
Malia praises Trump	12	4.085
Walmart playing audio about BLM	13	2.91
Haiti official who exposed the Clinton Foundation found dead.	14	4.25
Trump organization wins lucrative contract to rebuild Syrian airport	15	3.335
Maine House Democrats vote to allow female genital mutilation	16	4.01
CNN shocks America, Fires journalist.	17	3.69
White House chef quits because Trump has only eaten fast food for 6 months	18	2.89
Hillary Clinton accepted \$30,000 donation from child sex cult.	19	4.425
Trump bringing back the draft.	20	3.69
<i>False Untagged Headlines</i>		
Trump wants to deport American Indians to India	21	2.6
Journalist who filmed Trump sobbing in garden still in custody	22	2.805
Baltimore mayor to spite Trump with taco trucks on every corner	23	2.98
Malia Obama among 10 arrested	24	4.265
Rubio "Rape victims should be in custody if there is suspicion that they're	25	2.87

DISPUTED TAGS

planning abortion”		
Trump reveals which Dem president was also KKK member	26	4.215
Trump advisor denies moon landing	27	3.24
Michigan House passes human microchipping legislation	28	3.51
Denzel Washington supports Trump.	29	4.32
Republican senator unveils plan to send teachers to marine bootcamp	30	3.11

¹ Headline Reference # refers to the arbitrary number assigned to each headline in the data file available on the OSF (<https://osf.io/3x874>).

² Partisanship score was generated by Pennycook, Binnendyk, Newton, & Rand (2021) during the creation and testing of the large corpus of fake and real news headlines from which the headlines used herein were randomly chosen. Specifically, this score is the “partisanship combined” score of their Democratic and Republican participants’ response to the question “Assuming the above headline is entirely accurate, how favorable would it be to Democrats versus Republicans?” with the response options being “1. More favorable for Democrats, 2. Moderately more favorable for Democrats, 3. Slightly more favorable for Democrats, 4. Slightly more favorable for Republicans, 5. Moderately more favorable for Republicans, 6. More favorable for Republicans”. As such, a value greater than 3.5 means the headline was viewed as pro-Republican on average, and a value less than 3.5 means the headline was viewed as pro-Democrat on average. We collected these values from the dataset Pennycook et al. made publicly available online (<https://osf.io/uny58>).

Appendix B: Regression Output

Table A: Predicting Belief

Predictors	Base Model			Bade Model w/ Headline Partisanship		
	Std. Beta (SE)	95% CI	p	Std. Beta (SE)	95% CI	p
(Intercept)	0.00 (0.04)	-0.09 – 0.09	1.000	0.00 (0.04)	-0.09 – 0.09	1.000
Condition [False Untagged]	-0.01 (0.04)	-0.08 – 0.07	0.835	-0.01 (0.04)	-0.08 – 0.07	0.866
Condition [False Tagged]	-0.11 (0.04)	-0.19 – -0.04	0.004	-0.09 (0.04)	-0.17 – -0.01	0.024
Headline Partisanship				-0.04 (0.03)	-0.10 – 0.01	0.125
Random Effects						
σ^2	0.52			0.52		
τ_{00}	0.43 ID			0.43 ID		
	0.02 Headline_Num			0.02 Headline_Num		
τ_{11}	0.02 ID.Condition_sumFalse_Untagged			0.02 ID.Condition_sumFalse_Untagged		
	0.02 ID.Condition_sumFalse_Tagged			0.02 ID.Condition_sumFalse_Tagged		
ρ_{01}	-0.02 ID.Condition_sumFalse_Untagged			-0.02 ID.Condition_sumFalse_Untagged		
	-0.25 ID.Condition_sumFalse_Tagged			-0.25 ID.Condition_sumFalse_Tagged		
ICC	0.48			0.48		
N	318 ID			318 ID		
	30 Headline_Num			30 Headline_Num		
Observations	9540			9540		
Marginal R ² / Conditional R ²	0.009 / 0.482			0.011 / 0.483		

Table A: Linear mixed-effects models predicting belief in the headline. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 0.21$, $p = 0.645$.

Table B: Predicting Belief

Predictors	BIS Model			BIS Model w/ Headline Partisanship		
	Std. Beta (SE)	95% CI	p	Std. Beta (SE)	95% CI	p
(Intercept)	-0.00 (0.07)	-0.13 – 0.13	1.000	-0.00 (0.07)	-0.13 – 0.13	1.000
Condition [False Untagged]	-0.24 (0.09)	-0.42 – -0.06	0.012	-0.24 (0.09)	-0.42 – -0.05	0.014
Condition [False Tagged]	-0.28 (0.09)	-0.47 – -0.10	0.003	-0.27 (0.10)	-0.47 – -0.07	0.011
BIS Score	0.04 (0.02)	-0.00 – 0.08	0.061	0.04 (0.02)	-0.00 – 0.08	0.061
Condition	0.02	-0.01 – 0.05	0.156	0.02	-0.01 – 0.05	0.156

DISPUTED TAGS

[Untagged] x BIS	(0.01)			(0.01)		
Condition [Tagged] x BIS	0.00 (0.02)	-0.03 – 0.04	0.885	0.00 (0.02)	-0.03 – 0.04	0.885
Headline Partisanship				-0.03 (0.07)	-0.18 – 0.12	0.670
Random Effects						
σ^2	0.58			0.58		
τ_{00}	0.13 ID			0.13 ID		
	0.11 Headline_Num			0.12 Headline_Num		
τ_{11}	0.02 ID.Condition_sumFalse_Untagged			0.02 ID.Condition_sumFalse_Untagged		
	0.06 ID.Condition_sumFalse_Tagged			0.06 ID.Condition_sumFalse_Tagged		
ρ_{01}	0.43 ID.Condition_sumFalse_Untagged			0.43 ID.Condition_sumFalse_Untagged		
	0.02 ID.Condition_sumFalse_Tagged			0.02 ID.Condition_sumFalse_Tagged		
ICC	0.34			0.34		
N	318 ID			318 ID		
	30 Headline_Num			30 Headline_Num		
Observations	9540			9540		
Marginal R ² / Conditional R ²	0.136 / 0.431			0.136 / 0.433		

Table B: Linear mixed-effects models predicting belief in the headline. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 0.21$, $p = 0.645$.

Table C: Predicting Belief

<i>Predictors</i>	CRT Model			CRT Model w/ Headline Partisanship		
	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>
(Intercept)	-0.00 (0.07)	-0.13 – 0.13	1.000	-0.00 (0.07)	-0.13 – 0.13	1.000
Condition [False Untagged]	-0.24 (0.09)	-0.42 – -0.06	0.012	-0.24 (0.09)	-0.42 – -0.05	0.014
Condition [False Tagged]	-0.28 (0.09)	-0.47 – -0.10	0.003	-0.27 (0.10)	-0.47 – -0.07	0.011
CRT Score	-0.03 (0.02)	-0.07 – 0.02	0.252	-0.03 (0.02)	-0.07 – 0.02	0.252
Condition [Untagged] x CRT	-0.03 (0.01)	-0.05 – 0.00	0.053	-0.03 (0.01)	-0.05 – 0.00	0.053
Condition [Tagged] x CRT	-0.05 (0.02)	-0.08 – -0.01	0.008	-0.05 (0.02)	-0.08 – -0.01	0.008
Headline Partisanship				-0.03 (0.07)	-0.18 – 0.12	0.670
Random Effects						
σ^2	0.58			0.58		
τ_{00}	0.14 ID			0.14 ID		

DISPUTED TAGS

	0.11	Headline_Num	0.12	Headline_Num
τ_{11}	0.02	ID.Condition_sumFalse_Untagged	0.02	ID.Condition_sumFalse_Untagged
	0.06	ID.Condition_sumFalse_Tagged	0.06	ID.Condition_sumFalse_Tagged
ρ_{01}	0.44	ID.Condition_sumFalse_Untagged	0.44	ID.Condition_sumFalse_Untagged
	0.01	ID.Condition_sumFalse_Tagged	0.01	ID.Condition_sumFalse_Tagged
ICC	0.34		0.34	
N	318	ID	318	ID
	30	Headline_Num	30	Headline_Num
Observations	9540		9540	
Marginal R ² / Conditional R ²	0.138 / 0.431		0.138 / 0.433	

Table C: Linear mixed-effects models predicting belief in the headline. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 0.21$, $p = 0.645$.

Table D: Predicting Belief

<i>Predictors</i>	Party ID Model			Party ID Model w/ Headline Partisanship		
	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>
(Intercept)	0.00 (0.07)	-0.13 – 0.13	0.999	0.00 (0.07)	-0.13 – 0.13	0.999
Condition [False Untagged]	-0.24 (0.09)	-0.42 – -0.06	0.011	-0.24 (0.09)	-0.42 – -0.05	0.013
Condition [False Tagged]	-0.28 (0.09)	-0.46 – -0.10	0.004	-0.27 (0.10)	-0.47 – -0.07	0.009
Party ID [Democrat]	-0.01 (0.04)	-0.08 – 0.07	0.888	-0.01 (0.04)	-0.08 – 0.07	0.888
Party ID [Independent]	-0.00 (0.03)	-0.07 – 0.06	0.922	-0.00 (0.03)	-0.07 – 0.06	0.922
Condition[Untagged] x Party ID [Dem]	0.06 (0.04)	-0.02 – 0.14	0.121	0.06 (0.04)	-0.02 – 0.14	0.121
Condition[Tagged] x Party ID [Dem]	-0.09 (0.04)	-0.17 – -0.00	0.039	-0.09 (0.04)	-0.17 – -0.00	0.039
Condition[Untagged] x Party ID [Indep]	-0.01 (0.02)	-0.06 – 0.04	0.588	-0.01 (0.02)	-0.06 – 0.04	0.588
Condition[Tagged] x Party ID [Indep]	-0.08 (0.03)	-0.13 – -0.02	0.009	-0.08 (0.03)	-0.13 – -0.02	0.009
Headline Partisanship				-0.02 (0.06)	-0.15 – 0.11	0.770
Random Effects						
σ^2	0.56			0.56		
τ_{00}	0.14	ID		0.14	ID	
	0.11	Headline_Num		0.12	Headline_Num	
τ_{11}	0.02	ID.Condition_sumFalse_Untagged		0.02	ID.Condition_sumFalse_Untagged	
	0.05	ID.Condition_sumFalse_Tagged		0.05	ID.Condition_sumFalse_Tagged	
	0.02	Headline_Num.Party_ID_sumDemocrat		0.02	Headline_Num.Party_ID_sumDemocrat	

DISPUTED TAGS

	0.00	Headline_Num.Party_ID_sumIndependent	0.00	Headline_Num.Party_ID_sumIndependent
ρ_{01}	0.46		0.46	
	0.02		0.02	
	-0.02		-0.06	
	0.57		0.54	
ICC	0.36		0.36	
N	318	ID	318	ID
	30	Headline_Num	30	Headline_Num
Observations	9540		9540	
Marginal R ² / Conditional R ²	0.141 / 0.447		0.141 / 0.449	

Table D: Linear mixed-effects models predicting belief in the headline. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 0.05$, $p = 0.824$.

Table E: Predicting Belief

<i>Predictors</i>	Ideology Model			Ideology Model w/ Headline Partisanship		
	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>
(Intercept)	-0.00 (0.07)	-0.13 – 0.13	1.000	-0.00 (0.07)	-0.14 – 0.14	1.000
Condition [False Untagged]	-0.24 (0.09)	-0.42 – -0.06	0.012	-0.24 (0.09)	-0.42 – -0.05	0.014
Condition [False Tagged]	-0.28 (0.09)	-0.46 – -0.10	0.003	-0.29 (0.10)	-0.50 – -0.09	0.006
Ideology [Conservatism]	0.03 (0.04)	-0.05 – 0.10	0.481	0.03 (0.04)	-0.05 – 0.10	0.481
Condition[Untagged] x Ideology	-0.05 (0.04)	-0.13 – 0.03	0.239	-0.05 (0.04)	-0.13 – 0.03	0.239
Condition[Tagged] x Ideology	0.14 (0.04)	0.05 – 0.22	0.002	0.14 (0.04)	0.05 – 0.22	0.002
Headline Partisanship				0.02 (0.07)	-0.12 – 0.16	0.776
Random Effects						
σ^2	0.55			0.55		
τ_{00}	0.14	ID		0.14	ID	
	0.11	Headline_Num		0.12	Headline_Num	
τ_{11}	0.02	ID.Condition_sumFalse_Untagged		0.02	ID.Condition_sumFalse_Untagged	
	0.05	ID.Condition_sumFalse_Tagged		0.05	ID.Condition_sumFalse_Tagged	
	0.02	Headline_Num.Ideology_std		0.02	Headline_Num.Ideology_std	
ρ_{01}	0.48	ID.Condition_sumFalse_Untagged		0.48	ID.Condition_sumFalse_Untagged	
	-0.02	ID.Condition_sumFalse_Tagged		-0.02	ID.Condition_sumFalse_Tagged	
	-0.17	Headline_Num		-0.21	Headline_Num	
ICC	0.36			0.37		
N	318	ID		318	ID	

DISPUTED TAGS

	30 Headline_Num	30 Headline_Num
Observations	9540	9540
Marginal R ² / Conditional R ²	0.144 / 0.454	0.144 / 0.458

Table E: Linear mixed-effects models predicting belief in the headline. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 0.04$, $p = 0.837$.

Table F: Predicting Sharing Intentions

Predictors	Base Model			Bade Model w/ Headline Partisanship		
	Std. Beta (SE)	95% CI	<i>p</i>	Std. Beta (SE)	95% CI	<i>p</i>
(Intercept)	0.00 (0.04)	-0.09 – 0.09	1.000	0.00 (0.04)	-0.09 – 0.09	1.000
Condition [False Untagged]	-0.01 (0.04)	-0.08 – 0.07	0.835	-0.01 (0.04)	-0.08 – 0.07	0.866
Condition [False Tagged]	-0.11 (0.04)	-0.19 – -0.04	0.004	-0.09 (0.04)	-0.17 – -0.01	0.024
Headline Partisanship				-0.04 (0.03)	-0.10 – 0.01	0.125
Random Effects						
σ^2	0.52			0.52		
τ_{00}	0.43 ID			0.43 ID		
	0.02 Headline_Num			0.02 Headline_Num		
τ_{11}	0.02 ID.Condition_sumFalse_Untagged			0.02 ID.Condition_sumFalse_Untagged		
	0.02 ID.Condition_sumFalse_Tagged			0.02 ID.Condition_sumFalse_Tagged		
ρ_{01}	-0.02 ID.Condition_sumFalse_Untagged			-0.02 ID.Condition_sumFalse_Untagged		
	-0.25 ID.Condition_sumFalse_Tagged			-0.25 ID.Condition_sumFalse_Tagged		
ICC	0.48			0.48		
N	318 ID			318 ID		
	30 Headline_Num			30 Headline_Num		
Observations	9540			9540		
Marginal R ² / Conditional R ²	0.009 / 0.482			0.011 / 0.483		

Table F: Linear mixed-effects models predicting sharing intentions. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 2.70$, $p = 0.101$.

Table G: Predicting Sharing Intentions

Predictors	BIS Model			BIS Model w/ Headline Partisanship		
	Std. Beta (SE)	95% CI	<i>p</i>	Std. Beta (SE)	95% CI	<i>p</i>
(Intercept)	0.00 (0.04)	-0.09 – 0.09	1.000	0.00 (0.04)	-0.09 – 0.09	1.000

DISPUTED TAGS

Condition [False Untagged]	-0.01 (0.04)	-0.08 – 0.07	0.835	-0.01 (0.04)	-0.08 – 0.07	0.864
Condition [False Tagged]	-0.11 (0.04)	-0.19 – -0.04	0.004	-0.09 (0.04)	-0.17 – -0.01	0.022
BIS Score	0.03 (0.04)	-0.05 – 0.10	0.480	0.03 (0.04)	-0.05 – 0.10	0.480
Condition [Untagged] x BIS	0.02 (0.02)	-0.01 – 0.05	0.200	0.02 (0.02)	-0.01 – 0.05	0.200
Condition [Tagged] x BIS	0.00 (0.02)	-0.03 – 0.04	0.833	0.00 (0.02)	-0.03 – 0.04	0.833
Headline Partisanship				-0.04 (0.03)	-0.10 – 0.02	0.144
Random Effects						
σ^2	0.52			0.52		
τ_{00}	0.43 ID			0.43 ID		
	0.02 Headline_Num			0.02 Headline_Num		
τ_{11}	0.01 ID.Condition_sumFalse_Untagged			0.01 ID.Condition_sumFalse_Untagged		
	0.02 ID.Condition_sumFalse_Tagged			0.02 ID.Condition_sumFalse_Tagged		
	0.00 Headline_Num.BIS_std			0.00 Headline_Num.BIS_std		
ρ_{01}	-0.02 ID.Condition_sumFalse_Untagged			-0.02 ID.Condition_sumFalse_Untagged		
	-0.25 ID.Condition_sumFalse_Tagged			-0.25 ID.Condition_sumFalse_Tagged		
	0.23 Headline_Num			0.08 Headline_Num		
ICC	0.48			0.48		
N	318 ID			318 ID		
	30 Headline_Num			30 Headline_Num		
Observations	9540			9540		
Marginal R ² / Conditional R ²	0.010 / 0.484			0.012 / 0.485		

Table G: Linear mixed-effects models predicting sharing intentions. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 2.22$, $p = 0.136$.

Table H: Predicting Sharing Intentions

<i>Predictors</i>	CRT Model			CRT Model w/ Headline Partisanship		
	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>
(Intercept)	0.00 (0.04)	-0.09 – 0.09	1.000	0.00 (0.04)	-0.09 – 0.09	1.000
Condition [False Untagged]	-0.01 (0.04)	-0.08 – 0.07	0.835	-0.01 (0.04)	-0.08 – 0.07	0.870
Condition [False Tagged]	-0.11 (0.04)	-0.19 – -0.04	0.004	-0.09 (0.04)	-0.17 – -0.01	0.028
CRT Score	-0.07 (0.04)	-0.14 – 0.01	0.079	-0.07 (0.04)	-0.14 – 0.01	0.079
Condition [Untagged] x CRT	-0.01 (0.01)	-0.03 – 0.02	0.589	-0.01 (0.01)	-0.03 – 0.02	0.589

DISPUTED TAGS

Condition [Tagged] x CRT	-0.02 (0.01)	-0.04 – 0.01	0.270	-0.02 (0.01)	-0.04 – 0.01	0.270
Headline Partisanship				-0.05 (0.03)	-0.10 – 0.01	0.085
Random Effects						
σ^2	0.52			0.52		
τ_{00}	0.43 _{ID}			0.43 _{ID}		
	0.02 _{Headline_Num}			0.02 _{Headline_Num}		
τ_{11}	0.02 _{ID.Condition_sumFalse_Untagged}			0.02 _{ID.Condition_sumFalse_Untagged}		
	0.02 _{ID.Condition_sumFalse_Tagged}			0.02 _{ID.Condition_sumFalse_Tagged}		
	0.00 _{Headline_Num.CRT_std}			0.00 _{Headline_Num.CRT_std}		
ρ_{01}	-0.02 _{ID.Condition_sumFalse_Untagged}			-0.02 _{ID.Condition_sumFalse_Untagged}		
	-0.26 _{ID.Condition_sumFalse_Tagged}			-0.26 _{ID.Condition_sumFalse_Tagged}		
	0.06 _{Headline_Num}			-0.23 _{Headline_Num}		
ICC	0.48			0.48		
N	318 _{ID}			318 _{ID}		
	30 _{Headline_Num}			30 _{Headline_Num}		
Observations	9540			9540		
Marginal R ² / Conditional R ²	0.014 / 0.484			0.016 / 0.484		

Table H: Linear mixed-effects models predicting sharing intentions. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 2.93$, $p = 0.087$.

Table I: Predicting Sharing Intentions

<i>Predictors</i>	Party ID Model			Party ID Model w/ Headline Partisanship		
	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>
(Intercept)	-0.00 (0.04)	-0.09 – 0.09	0.993	-0.00 (0.04)	-0.09 – 0.09	0.993
Condition [False Untagged]	-0.01 (0.04)	-0.08 – 0.07	0.828	-0.01 (0.04)	-0.08 – 0.07	0.858
Condition [False Tagged]	-0.11 (0.04)	-0.19 – -0.04	0.004	-0.09 (0.04)	-0.17 – -0.01	0.026
Party ID [Democrat]	0.02 (0.05)	-0.08 – 0.12	0.700	0.02 (0.05)	-0.08 – 0.12	0.700
Party ID [Independent]	-0.14 (0.05)	-0.24 – -0.03	0.010	-0.14 (0.05)	-0.24 – -0.03	0.010
Condition[Untagged] x Party ID [Dem]	0.02 (0.02)	-0.02 – 0.05	0.299	0.02 (0.02)	-0.02 – 0.05	0.299
Condition[Tagged] x Party ID [Dem]	-0.08 (0.02)	-0.11 – -0.04	<0.001	-0.08 (0.02)	-0.11 – -0.04	<0.001
Condition[Untagged] x Party ID [Indep]	0.01 (0.02)	-0.02 – 0.05	0.497	0.01 (0.02)	-0.02 – 0.05	0.497
Condition[Tagged] x Party ID [Indep]	-0.03 (0.02)	-0.07 – 0.00	0.059	-0.03 (0.02)	-0.07 – 0.00	0.059

DISPUTED TAGS

Headline Partisanship		-0.04 (0.03)	-0.10 – 0.01	0.125
Random Effects				
σ^2	0.52	0.52		
τ_{00}	0.42 ID	0.42 ID		
	0.02 Headline_Num	0.02 Headline_Num		
τ_{11}	0.01 ID.Condition_sumFalse_Untagged	0.01 ID.Condition_sumFalse_Untagged		
	0.02 ID.Condition_sumFalse_Tagged	0.02 ID.Condition_sumFalse_Tagged		
ρ_{01}	0.00 ID.Condition_sumFalse_Untagged	0.00 ID.Condition_sumFalse_Untagged		
	-0.36 ID.Condition_sumFalse_Tagged	-0.36 ID.Condition_sumFalse_Tagged		
ICC	0.47	0.47		
N	318 ID	318 ID		
	30 Headline_Num	30 Headline_Num		
Observations	9540	9540		
Marginal R ² / Conditional R ²	0.023 / 0.484	0.025 / 0.484		

Table I: Linear mixed-effects models predicting sharing intentions. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 2.70$, $p = 0.100$.

Table J: Predicting Sharing Intentions

<i>Predictors</i>	Ideology Model			Ideology Model w/ Headline Partisanship		
	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>	<i>Std. Beta (SE)</i>	<i>95% CI</i>	<i>p</i>
(Intercept)	0.00 (0.04)	-0.09 – 0.09	1.000	0.00 (0.04)	-0.09 – 0.09	1.000
Condition [False Untagged]	-0.01 (0.04)	-0.08 – 0.07	0.835	-0.01 (0.04)	-0.09 – 0.07	0.835
Condition [False Tagged]	-0.11 (0.04)	-0.19 – -0.04	0.003	-0.12 (0.04)	-0.20 – -0.04	0.007
Ideology [Conservatism]	0.06 (0.05)	-0.03 – 0.15	0.201	0.06 (0.05)	-0.03 – 0.15	0.201
Condition[Untagged] x Ideology	-0.02 (0.04)	-0.10 – 0.06	0.634	-0.02 (0.04)	-0.10 – 0.06	0.634
Condition[Tagged] x Ideology	0.09 (0.04)	0.01 – 0.18	0.031	0.09 (0.04)	0.01 – 0.18	0.031
Headline Partisanship				0.01 (0.03)	-0.05 – 0.06	0.785
Random Effects						
σ^2	0.50			0.50		
τ_{00}	0.43 ID			0.43 ID		
	0.02 Headline_Num			0.02 Headline_Num		
τ_{11}	0.02 ID.Condition_sumFalse_Untagged			0.02 ID.Condition_sumFalse_Untagged		
	0.01 ID.Condition_sumFalse_Tagged			0.01 ID.Condition_sumFalse_Tagged		
	0.02 Headline_Num.Ideology_std			0.02 Headline_Num.Ideology_std		

DISPUTED TAGS

ρ_{01}	-0.00 ID.Condition_sumFalse_Un tagged	-0.00 ID.Condition_sumFalse_Un tagged
	-0.38 ID.Condition_sumFalse_Tagged	-0.38 ID.Condition_sumFalse_Tagged
	-0.43 Headline_Num	-0.45 Headline_Num
ICC	0.50	0.50
N	318 ID	318 ID
	30 Headline_Num	30 Headline_Num
Observations	9540	9540
Marginal R ² / Conditional R ²	0.018 / 0.508	0.018 / 0.508

Table J: Linear mixed-effects models predicting sharing intentions. All condition effects are sum coded, and all other predictors are z-scored. *P*-values are derived using Welch-Satterthwaite approximation, and *p*-values less than 0.05 are bolded. The model with Headline Partisanship did not significantly improve model fit, $\chi^2(1) = 0.22$, $p = 0.881$.