

RUNNING HEAD: Fake News

Misinformation and Morality:

Encountering Fake-News Headlines Makes Them Seem Less Unethical to Publish and Share

Daniel A. Effron
London Business School

Medha Raj
Marshall School of Business
University of Southern California

Article in Press at *Psychological Science*
September 15, 2019

Abstract

People may repeatedly encounter the same misinformation when it “goes viral.” Four experiments and a pilot study (two pre-registered; $N = 2,587$) suggest that repeatedly encountering misinformation makes it seem less unethical to spread—regardless of whether one believes it. Seeing a fake-news headline one or four times reduced how unethical participants thought it was to publish and share that headline when they saw it again – even when it was clearly labelled false and participants disbelieved it, and even after statistically accounting for judgments of how likeable and popular it was. In turn, perceiving it as less unethical predicted stronger inclinations to express approval of it online. People were also more likely to actually share repeated (vs. new) headlines in an experimental setting. We speculate that repeating blatant misinformation may reduce the moral condemnation it receives by making it feel intuitively true, and we discuss other potential mechanisms.

[*Word count: 148*]

KEYWORDS: Misinformation, fake news, fluency, familiarity, repetition, moral judgment, intuition, deliberative thinking, ethics, lie, illusory truth effect, mere exposure

Misinformation and Morality:

Encountering Fake-News Headlines Makes Them Seem Less Unethical to Publish and Share

As technology has enabled the rapid spread of “fake news,” some have worried that misinformation has become “the major moral crisis of our times” ([Al-Rodhan, 2017](#)). Yet people sometimes find it morally permissible to spread misinformation – for example, when it supports their political agenda ([Mueller & Skitka, 2018](#)). These moral judgments matter: When people find misinformation permissible, they should be less inclined to take action to stop it ([Effron & Miller, 2012](#); [Skitka & Bauman, 2008](#)), less likely to hold its purveyors accountable ([Effron, 2018](#)), and more likely to spread it themselves. This could help explain why 14% of U.S. adults and 17% of U.K. adults admitted to sharing news that they thought was fake at the time ([Chadwick & Vaccari, 2019](#); [Mitchell, Barthel, & Holcomb, 2016](#)).

The present research investigates what shapes moral judgments of fake news – “articles that are intentionally and verifiably false, and could mislead readers” ([Allcott & Gentzkow, 2017, p. 213](#)). Among fact-checked news, fake articles were more likely than real articles to “go viral” on social media ([Vosoughi, Roy, & Aral, 2018](#)). And when a fake-news article goes viral, people may encounter it multiple times. Prior research raises the concern that people are more likely to believe a fake-news headline if they have seen it before ([Pennycook, Cannon, & Rand, 2018](#)). We propose, however, that regardless of whether one believes a piece of fake news, prior encounters with it can reduce how unethical one thinks it is to spread.

We based this prediction on the idea that previously-encountered information feels more *fluent* (i.e., easier to process) than new information ([Alter & Oppenheimer, 2009](#); [Whittlesea, Jacoby, & Girard, 1990](#)), and that people associate fluency with truth ([McGlone & Tofighbakhsh, 2000](#); [Newman, Garry, Bernstein, Kantner, & Lindsay, 2012](#); [Reber & Schwarz, 1999](#)). This association, we propose, can lend repeated information a “ring of truthfulness”

regardless of whether one believes it. Indeed, information can intuitively feel truthful to a person who explicitly acknowledges its falsity ([Shidlovski, Schul, & Mayo, 2014](#)). Because intuitions drive moral judgments ([Haidt, 2001](#)), people may find it more permissible to spread misinformation they know is false when it intuitively feels true. For example, people are more likely to tell and excuse a blatant falsehood after imagining how it *could have been true* in alternative circumstances ([Effron, 2018](#); [Shalvi, Dana, Handgraaf, & De Dreu, 2011](#)). For these reasons, we predicted that repeated exposure to a fake-news headline would lead people to judge it as less unethical to publish or share, independent of their beliefs in its factual accuracy.

People also judge repeated statements as more accurate – the *illusory-truth effect* (see [Dechêne, Stahl, Hansen, & Wänke, 2010](#); [Hasher, Goldstein, & Toppino, 1977](#)). Even when people know that a claim is false when they first encounter it, they may forget or neglect to draw on this knowledge when encountering the claim again, mistaking fluency for truth ([Fazio, Brashier, Payne, & Marsh, 2015](#); [Skurnik, Yoon, Park, & Schwarz, 2005](#)). Unlike the illusory-truth effect, our studies examine moral judgments of claims whose falsity is highly salient even on subsequent encounters. For example, our studies label fake-news headlines as false just before participants judge their unethicality. Under these conditions, we expect repeated exposure to fake-news headlines to reduce how unethical they seem without affecting how accurate they seem, because the illusory-truth effect does not occur for saliently-false claims ([Pennycook et al., 2018, Study 1](#)). Our studies therefore test for the first time whether repeated exposure to falsehoods can affect moral judgments of these falsehoods, above and beyond accuracy beliefs.

Experiment 1 tested whether four previous encounters with a fake-news headline would make it seem less unethical to publish. Experiment 2 (pre-registered) and its Pilot tested whether a single encounter would suffice. Experiment 3 (pre-registered) tested a boundary condition: If, as we argue, previously-encountered (vs. new) misinformation seems less unethical to spread

because it feels more intuitively true, then encouraging people to think deliberately instead of intuitively should attenuate the effect. Experiment 4 addressed whether repeatedly encountering the headlines could affect moral judgments above and beyond judgments of their accuracy, likeability, and popularity. Also, whereas Experiments 1-3 warned participants they were judging fake news, Experiment 4 omitted this warning to test generalizability. Finally, the experiments examined whether prior encounters with the headlines would increase people's intentions to share and "like" them, reduce their inclination to censure those who posted them, and increase actual sharing behavior. The General Discussion speculates on how intuitive feelings of truthfulness, as well as alternative mechanisms, could explain the results. [*Word count: 753*]

Experiment 1

Method

Experiment 1 and all subsequent studies report all measures, manipulations, and data exclusions. Targeted sample sizes were determined in advance of data collection. Verbatim materials are posted in the Supplemental Online Material (SOM-R, Appendices 7-10).

Participants. In August, 2018, we targeted 150 American participants on Prolific Academic, an online participant panel whose participants are more diverse and less familiar with experimental procedures than Amazon Mechanical Turk workers ([Peer, Brandimarte, Samat, & Acquisti, 2017](#)). Prior to Experiment 1, we ran Experiment 2 and its Pilot, but we present Experiment 1 first for narrative clarity. These prior studies informed Experiment 1's targeted sample size, which was smaller than the prior samples because we expected Experiment 1's manipulation to be stronger. To improve data quality, we prevented participants from beginning the study if they failed a reading-comprehension question, were on a mobile device, or were outside the U.S.

After we excluded people who met our *a priori* exclusion criteria (submitting data from a duplicate IP address, a duplicate Prolific Academic ID, or a non-US IP address) or who did not provide answers to any of the dependent measures, the dataset from this repeated-measures study contained 1,648 observations from 138 people (75 men and 63 women; $M_{\text{age}} = 34$ years, $SD = 13$, range = 18 to 74; 95 people who identified as or leaned Democrat, 22 who identified as or leaned Republican, and 21 who identified with or leaned towards neither political party). A sensitivity analysis suggested that the smallest mean difference this sample could detect with Experiment 1's design at approximately 80% power and with two-tailed $\alpha = .05$ was $|d_z| = .23$, given parameter estimates from the previously mentioned Pilot (see SOM-R, Appendices 1 and 2).

Materials. The stimuli were 12 actual “fake news” headlines about American politics, with accompanying photographs, collected from a fact-checking website ([Pennycook et al., 2018; see SOM-R, Appendix 3](#)). Half the headlines appealed to Republicans (e.g., “Election Night: Hillary Was Drunk, Got Physical With Mook and Podesta”) and half appealed to Democrats (e.g., “Pennsylvania Federal Court Grants Legal Authority to REMOVE TRUMP After Russian Meddling”). Whether the headlines appealed to participants' politics did not moderate the results in any of the studies, so we do not discuss this factor further.

Procedure. The procedure was adapted from [Pennycook et al. \(2018\)](#) and had two phases. In the *familiarization phase*, participants saw 6 of the 12 headlines four times. Each time, they rated the headlines on a different filler item (“How [interesting/engaging/funny/well-written] is this headline?”). Then they completed a distractor task, providing demographic information and answering additional filler questions. Finally, in the *judgment phase*, participants completed the dependent measures for each of the 12 headlines (see below). Half of

the headlines had been shown in the familiarization phase (previously-seen headlines) and half had not (new headlines). We randomized which headlines were shown in the familiarization phase, the order in which participants completed the filler items, and the order in which the headlines appeared in both phases.

Importantly, the following message introduced the judgment phase:

For this part of the study you will be asked to read a series of **fake** news headlines that were recently published online. The information in these headline is not real. Non-partisan fact-checking websites have confirmed that these headlines describe events that did NOT happen.

This message differentiates Experiment 1 from illusory-truth paradigms by eliminating any ambiguity about the headlines' falsity. In illusory-truth studies, repeated falsehoods can appear truthful because their falsity is either unknown to participants ([see Dechêne et al., 2010](#)), or is known to them but not salient in the judgment phase ([e.g., Fazio et al., 2015](#)). For example, in the familiarization phase, [Pennycook et al. \(2018, Study 2\)](#) labelled false headlines “disputed by third-party fact-checkers.” In the subsequent judgment phase, the headlines' truthfulness was ambiguous because this label was omitted and false headlines were interspersed with true headlines. Participants judged familiarized headlines as more accurate than new headlines, presumably because they forgot which headlines were disputed ([for a finding with a similar explanation, see Begg, Anas, & Farinacci, 1992](#)). We expected no illusory-truth effect in Experiment 1 because its participants judged *only* false statements that were presented as *unambiguously false in the judgment phase*.

Measures.

Moral condemnation. For the primary dependent measure, participants moved a slider to indicate how (a) unethical and (b) acceptable (reverse-coded) it would be to publish each

headline (Effron, 2018; 0 = *Not at all*, 100 = *Extremely*; slider's start position = 50; across headlines, $r = .64$, $p < .001$). We averaged the two items.

Intended social-media behaviors. We also tested potential downstream consequences of moral condemnation by asking participants how likely they would be to perform each of four behaviors if an acquaintance posted the headline on social media: “like” it, share it, post a negative comment, and block/unfollow the acquaintance (from 1 = *Not at all* to 7 = *To a great extent*, with 4 = *Somewhat*). We analyzed the items individually because they did not form a reliable scale ($\alpha = .56$ across headlines).

Accuracy beliefs. To test whether the illusory-truth effect could account for our predicted findings, we asked participants to rate each headline's factual accuracy. Following previous research on illusory-truth, we used a 4-point scale (1 = *Not at all accurate*, 2 = *Not very accurate*, 3 = *Somewhat accurate*, 4 = *Very accurate*). (Experiment 4 measured accuracy beliefs on the same 100-point scale as moral judgments).

Comprehension check. After the judgment phase, we checked whether participants understood that they had judged fake-news headlines by asking them whether all headlines were false, true, or whether some were false and some were true. People who chose the latter response were shown the 12 headlines again and were asked to indicate whether each headline was true or false.

Results

Moral condemnation. We submitted the moral condemnation measure to a multilevel regression model, which accounts for the fact that each participant rated 12 headlines. The model included a fixed effect for headline type (1 = *Previously-seen*; 0 = *New*), fixed effects for the specific headlines rated, and random intercepts for participants. The results supported our

hypothesis: People rated the headlines they had seen earlier as less unethical to publish than those they had not, as indicated by a significantly negative coefficient for headline type, $b = -3.83$, $z = 3.86$, $p < .001$, $d_z = -.26$ (see Table 1 for detailed statistics).

Intended social-media behaviors. Participants indicated that they were significantly more likely to “like” and share the previously-seen (vs. new) headlines, $ps < .01$, $d_zs = .24$ and $.27$, and significantly less likely to block/unfollow the person who posted the previously-seen (vs. new) headlines, $p < .001$, $d_z = -.23$. These effects were significantly mediated by moral judgments, $ps < .002$ (see SOM-R, Appendix 4), which is consistent with our expectation that exposure to the headlines would affect intended social-media behaviors by softening moral judgments. However, the standard limitations of mediation analysis prevent strong or causal conclusions (see [Bullock, Green, & Ha, 2010](#)). Participants were not significantly less likely to say they would post a negative comment in response to previously-seen (vs. new) headlines, $p = .092$, $d_z = -.14$.

Accuracy beliefs. Did participants rate previously-seen (vs. new) headlines as less unethical to publish simply because they were more inclined to believe previously-seen headlines? Five findings argue against this possibility. First, the correlation between the unethicality and accuracy measure at the trial level was only modest, $r(1648) = -.40$, $p < .001$, suggesting that each measure tapped a separate construct. Second, participants did not rate previously-seen headlines as significantly more accurate than new ones, $b = .02$, $SE_b = .02$, $z = .80$, $p = .43$, $d_z = .04$ (see Table 1) – i.e., we found no illusory truth effect. Third, participants still rated previously-seen headlines as less unethical to publish than new headlines when we statistically controlled for accuracy beliefs, $b = -3.56$, $SE_b = .94$, $z = 3.80$, $p < .001$. Fourth, the comprehension check showed that participants only misidentified the headlines as true 5.64% of the time. Fifth, all the statistically significant results reported above remained significant at $p <$

.01 when we excluded responses to any headlines a participant misidentified as true (i.e., 5.64% of all observations). Thus, our main results cannot be explained by a tendency to misremember false headlines as true.

Discussion

Experiment 1 suggests that repeatedly encountering a fake-news headline can reduce moral condemnation of publishing it, increase people's inclination to promote it on social media, and decrease their inclination to block/unfollow someone who posted it.

It is unsurprising that we did not replicate the illusory truth effect because our paradigm emphasized that all headlines were false, and people do not display the illusory truth effect with saliently-false stimuli ([Pennycook et al., 2018](#)). Given that we emphasized the headlines' falsity just before participants judged their morality and that participants did not believe repeated headlines any more than new headlines, it is unlikely that repetition reduced moral condemnation because participants forgot or neglected their knowledge that the headlines were false ([Fazio et al., 2015](#); [Skurnik et al., 2005](#)). [*Word count: 126*]

Experiment 2

Experiment 2 tests whether encountering a fake-news headline just *once* would reduce moral condemnation of publishing it. A Pilot Experiment found supportive evidence (see SOM-R, Appendix 2); Experiment 2 was a large-sample, pre-registered replication. [*Word count: 34*]

Table 1

Experiment 1's Results

Measure	Response options	<i>M</i> (previously seen)	95% <i>CI</i> (previously seen)	<i>M</i> (new)	95% <i>CI</i> (new)	<i>M</i> difference	<i>d_z</i>	<i>b</i>	<i>SE(b)</i>	<i>z</i>	<i>p</i>
Moral condemnation	0-100	66.37	[62.49, 70.55]	70.44	[66.32, 74.38]	-4.07	-0.26	-3.83	0.99	3.86	<0.001
“Like”	1-7	1.73	[1.54, 1.89]	1.55	[1.38, 1.73]	0.18	0.24	0.16	0.05	3.18	0.001
Share	1-7	1.62	[1.44, 1.80]	1.49	[1.31, 1.68]	0.13	0.27	0.13	0.04	3.19	0.001
Negative comment	1-7	2.06	[1.83, 2.29]	2.15	[1.92, 2.38]	-0.09	-0.14	-0.09	0.05	1.68	0.092
Block	1-7	2.29	[2.04, 2.55]	2.47	[2.22, 2.72]	-0.17	-0.23	-0.17	0.06	3.06	0.002
Accuracy	1-4	1.45	[1.35, 1.56]	1.43	[1.33, 1.53]	0.02	0.04	0.02	0.02	0.80	0.425

Note. 95% CIs computed from multi-level regression model.

Method

We pre-registered the targeted sample size, measures, hypotheses, and analyses (see <https://aspredicted.org/uj7nd.pdf>).

Participants. We targeted 800 Americans on Amazon's Mechanical Turk (MTurk) in June, 2018. The target sample was informed by the Pilot Experiment, which found significant evidence for a small effect of one exposure with $N = 596$ (see SOM-R, Appendix 2).

Participants could not begin the study if they failed a reading-comprehension question, or responded from a mobile device or non-U.S. IP address. After we excluded people who met our pre-registered exclusion criteria (identical to Experiment 1) and people who did not provide any responses, our dataset contained 9,536 observations from 796 people (467 women, 326 men, and 3 non-binary; 458 who identified as or leaned Democrat, 223 who identified as or leaned Republican, and 115 who did not lean towards either party; $M_{\text{age}} = 34$ years, $SD = 12$, range = 18 to 76). A sensitivity analysis suggested that the smallest mean difference this sample could detect in Experiment 2's design at approximately 80% power was $|d_z| = .10$, with two-tailed $\alpha = .05$ (see SOM-R, Appendix 1). Because we pre-registered one-tailed tests for our directional hypothesis, the smallest detectable mean difference was actually even smaller (the simulation software we used did not allow one-tailed tests).

Procedure and measures. The procedure was identical to Experiment 1, except that the familiarization phase showed each of six headlines once each instead of four times. As the filler item, participants rated how interesting each headline was. In the judgment phase, participants rated all 12 headlines (6 previously seen and 6 new) on a four-item unethicity measure from Effron (2018): how justified (reversed), unethical, and acceptable (reversed) it was to publish the headline, and how much of a lie the headline was (i.e., the two-items from Experiment 1, plus

two more; $\alpha = .85$ across headlines). Participants then completed Experiment 1's comprehension check. We randomized which headlines were previously seen vs. new, and explicitly informed participants at the beginning of the judgment phase that the headlines were false.

Results

As in Experiment 1, participants rated the headlines as less unethical to publish when they had seen them before ($M = 75.74$, 95% CI = [74.28, 77.20]) than when they had not ($M = 77.13$, 95% CI = [75.67, 78.59]), $b = -1.39$, $SE_b = .32$, $z = 4.30$, $p < .001$, $d_z = -.15$, in a mixed-effect regression predicting unethicality ratings from condition (1 = *Previously-seen*, 0 = *New*), fixed effects for the specific headlines, and random intercepts for participants. (As we had a directional hypothesis, we pre-registered a one-tailed significance test). Participants only misidentified the headlines as true 5.68% of the time, and the conclusions were identical even when we excluded these responses, $b = -1.08$, $SE_b = .32$, $z = 3.34$, $p = .001$, $d_z = -.12$.

Discussion

This large-sample, pre-registered experiment suggests that encountering a fake-news article once is sufficient to reduce moral condemnation of publishing this misinformation when it is encountered again. [*Word count: 26*]

Experiment 3

We have argued that previously-encountered fake-news headlines receive less moral condemnation because prior encounters make the headlines feel intuitively true, even when people are currently aware they are false. If this is correct, then shifting the basis of moral judgment from intuition to deliberation should attenuate this effect. Experiment 3 tested this prediction. [*Word count: 53*]

Method

The study design was a 2 (headline type: previously-seen vs. new; within-subjects) X 2 (condition: deliberative thinking vs. intuitive thinking; between-subjects) factorial, with 12 repeated measures. We pre-registered the targeted sample size, method, and analysis plan (see <https://aspredicted.org/3fc65.pdf>).

Participants. We requested 600 complete responses from Prolific Academic in November 2018, limiting sign-ups to participants who had not completed Experiment 1, and preventing people from beginning if they failed a reading-comprehension question or responded from a mobile device or non-U.S. IP address. We targeted this sample size based on an intuitive benchmark: In a design without repeated measures, it would attain the same statistical power as Experiment 1 to detect an interaction that completely attenuates Experiment 1's effect ([Simonsohn, 2014](#)). Like Experiment 1, Experiment 3 showed the repeated headlines four times.

After we excluded people who met our pre-registered criteria (see Experiments 1 and 2) or who did not provide any responses to the dependent measures, the repeated-measures design yielded 8,731 observations from 761 people (407 men, 345 women, and 9 non-binary; $M_{\text{age}} = 33$ years, $SD = 12$, range = 18 to 76; 509 identified as or leaned Democrat, 147 identified as or leaned Republican, and the remainder did not lean towards either party). We obtained more responses than expected because Prolific did not count responses that took > 20 minutes to submit.

Our main hypothesis was that the effect of repeated exposure on moral condemnation would be significantly smaller in the deliberative-thinking condition than in the intuitive-thinking condition. A sensitivity analysis showed that, if the effect of repetition in the intuitive-thinking condition were the same size as what we observed in Experiment 1 ($d_z = -.26$), then the smallest interaction effect this sample could detect in Experiment 3's design with approximately

80% power and $\alpha = .05$ (two-tailed) would reduce the repetition effect to $d_z = -.12$, an interaction coefficient of $b = 1.92$ (see SOM-R, Appendix 1). This analysis gives a conservative estimate of sensitivity because we pre-registered a one-tailed interaction test (our software only allowed two-tailed tests).

Procedure. The first part of the procedure closely followed Experiment 1. In the familiarization phase, participants viewed six headlines (randomly selected from our bank of twelve) four times and provided filler ratings. After a brief distractor task, we introduced the judgment phase by emphasizing that the headlines they were about to rate were “fake news.”

At this point, we introduced Experiment 3’s deliberative-vs.-intuitive-thinking manipulation ([adapted from Nordgren & Dijksterhuis, 2008](#)). Participants randomly assigned to the *deliberative-thinking* condition were instructed to “take time to deliberate,” “think very hard,” “ignore any gut feelings,” and “generate clear reasons” about the ethicality of publishing each headline. After viewing a given headline and before rating it, they had to type two such reasons into the survey. By contrast, participants randomly assigned to the *intuitive-thinking* condition were instructed to rate the headlines’ ethicality or unethicality quickly, based on “their first instinct,” to “pay attention to [their] feelings” and not to “think too hard.” They had no opportunity to type reasons.

For the main dependent measure, participants rated each headline on the two-item moral condemnation measure from Experiment 1 ($r = .61$). To assess downstream consequences, we asked how likely they would be, if an acquaintance posted the headline on social media, to “like” it, share it, and block/unfollow the acquaintance. Because these three items did not correlate highly with each other ($\alpha = .46$), we followed our pre-registered plan of analyzing them separately. After completing these measures for all items, participants responded to the

comprehension check from our previous studies by indicating whether the headlines were true, false, or a mixture of both – and if the latter, which were true and false.

Finally, participants completed the three-item Cognitive Reflection Test (CRT), which assesses individual differences in deliberative thinking ([Frederick, 2005](#)). We pre-registered this measure as exploratory: Although we wondered whether it would moderate Experiment 3's results ([cf. Pennycook & Rand, 2019](#)), we also thought our manipulation of deliberative thinking could swamp any effects of individual differences. CRT scores did not significantly moderate the results (see SOM-R, Appendix 5), so we do not discuss them further.

Results

Because we had strong directional predictions, we pre-registered and report one-tailed significance tests.

Moral condemnation. We hypothesized that previously-seen headlines would receive less moral condemnation than new headlines would when people engaged in intuitive thinking, and that deliberative thinking would attenuate or eliminate this effect. To test this hypothesis, we submitted unethicity ratings to a multi-level regression model with random intercepts for participants, and fixed effects for the specific headlines (dummy-coded), headline type (coded 1 = *Previously-seen*, 0 = *New*), thinking condition (coded 1 = *Deliberative*, 0 = *Intuitive*), and the headline-type X thinking-condition interaction. With this coding, our hypothesis predicts a positive coefficient for the interaction term (and our pre-registered one-tailed test would require retaining the null hypothesis if the coefficient were negative). We did observe a positive coefficient; it was not significant in an initial analysis, $b = 1.39$, $SE_b = .88$, $z = 1.59$, $p = .057$, and significant in a pre-registered robustness check that excluded responses to any headlines a participant had incorrectly labelled as true (5.5% of responses), $b = 1.92$, $SE_b = .89$, $z = 2.15$ $p =$

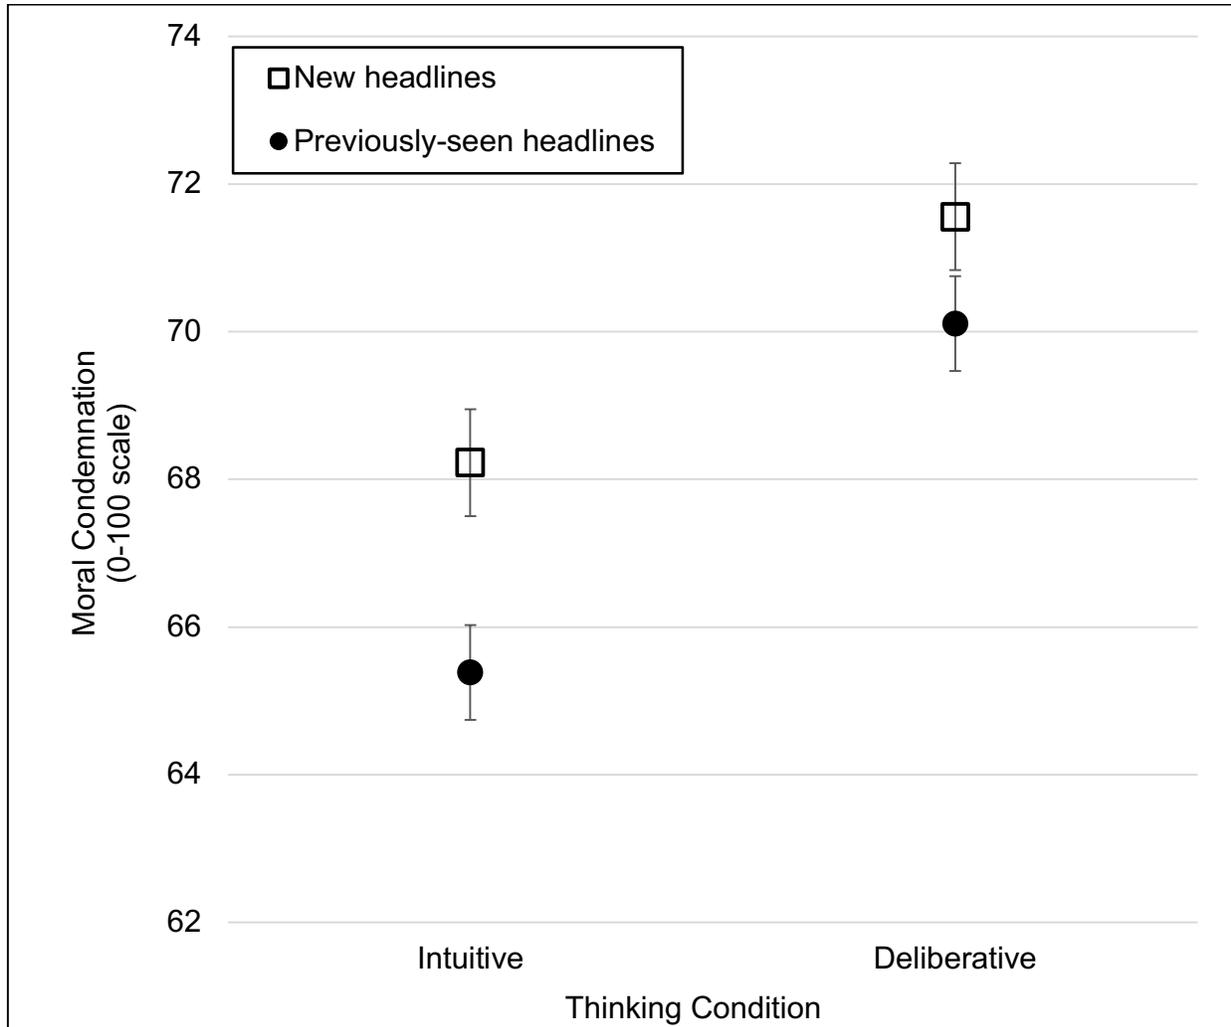
.016 (see Table 2). Recall that the smallest interaction that Experiment 3 could detect at $\alpha = .05$ (two-tailed) with 80% power given our estimated parameters was (coincidentally) $b = 1.92$, so the initial analysis may have been underpowered.

Following up this interaction test with pre-registered simple-slopes analyses showed the predicted pattern (see Figure 1). Because the interaction was not significant in the initial analysis, this pattern should be interpreted with caution. In the intuitive-thinking condition, participants rated the headlines as significantly less unethical to publish when they had (vs. had not) seen them earlier in the experiment, ($M_s = 65.38$ vs. 68.23 , 95% $CIs = [63.10, 67.67]$ and $[65.94, 70.51]$, respectively), with an effect size similar to the one in Experiment 1, $d_z = -.21$, $b = -2.84$, $SE_b = .60$, $z = 4.74$, $p < .001$. This effect was about half the size in the deliberative-thinking condition, ($M_s = 70.11$ vs. 71.56 , 95% $CIs = [67.73, 72.49]$ and $[69.18, 73.94]$, respectively), $d_z = -.11$, $b = -1.44$, $SE_b = .64$, $z = 2.25$, $p = .012$ (d_z was computed with the SD pooled across the two thinking conditions). The results were the same when we excluded responses to headlines categorized as true, except the effect in the deliberative-thinking condition was no longer significant, $p = .071$ (recall we hypothesized that deliberation would either attenuate or eliminate the effect).

Exploratory analyses tested the main effects of each manipulation by removing the interaction term from the mixed regression model. Overall, people expressed significantly less condemnation of previously-seen headlines than new headlines, and significantly more condemnation when induced to think deliberately vs. intuitively (two-tailed tests; see Table 2).

Figure 1

Mean Moral Condemnation in Experiment 3, by Condition



Note. Error bars are within-participant 95% CIs, calculated using the Cousineau-Morey method, around raw means. These error bars facilitate visualization of significant differences among within-participant conditions (i.e., new vs. previously-seen headlines) by removing between-participants variance. As such, they should not be used to make inferences about the between-participants manipulation (i.e., thinking condition). The main text reports between-participant 95% CIs.

Table 2

Results of Mixed Regression Models for Each Dependent Measure in Experiment 3

Measure	Predictor	Initial Analysis		Robustness check	
		Interaction model	Main-effect model	Interaction model	Main-effect model
Moral condemnation					
	Headline type	-2.84***	-2.19***	-2.87***	-1.99***
	Thinking condition	3.33*	4.03*	3.19†	4.14*
	Headline X Thinking	1.39†	.	1.92*	.
"Like"					
	Headline type	.08**	.06**	.08**	.05*
	Thinking condition	-.16*	-.18*	-.14*	-.18**
	Headline X Thinking	-0.04	.	-.07*	.
Share					
	Headline type	0.03	0.03	0.04	0.02
	Thinking condition	-0.16*	-.17*	-.14*	-.16*
	Headline X Thinking	-0.01	.	-0.04	.
Block					
	Headline type	-.19***	-.17***	-.18***	-.16***
	Thinking condition	.29*	.31*	.30*	.32**
	Headline X Thinking	0.03	.	0.05	.

Note. Predictors were dummy-coded (headline type: 1 = *previously-seen*, 0 = *new*; thinking condition: 1 = *deliberative*, 0 = *intuitive*). Thus, their coefficients are only main effects in the models without an interaction term (labelled “main-effect model”). The robustness check excludes participants’ responses to any headlines they misidentified as true. The regression models also included a fixed intercept, random intercepts for participants, and fixed effects for headline. † $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$, with two-tailed significance tests for simple and main effects, and one-tailed tests for pre-registered interactions.

Intended social-media behaviors. We did not find evidence that deliberative thinking moderated the effect of exposure to the headlines on intended social-media behaviors (see Table 2). In the initial analysis, the predicted headline-type X thinking-condition interaction was not significant for intentions to “like,” share, or block/unfollow; respectively, $bs = -.04, -.01,$ and $.03,$ $zs = .96, .34,$ and $.59,$ $ps > .300$. However, pre-registered simple-slopes analysis showed that the intuitive-thinking condition replicated Experiment 1’s finding that people felt more inclined to “like” previously-seen headlines than new headlines, $b = .08, z = 2.81, p = .003, d_z = .13,$ and less inclined to block/unfollow a person who posted previously-seen vs. new headlines, $b = -.19, z = 4.83, p < .001, d_z = -.21$. They did not feel significantly more inclined to share previously-seen vs. new headlines in this study, $b = .03, z = 1.24, p = .107, d_z = .06$. The coefficients were directionally smaller (but not significantly so) for these three measures in the deliberative-thinking condition: $b = .04, z = 1.31, p = .096, d_z = .08$ for “liking,” $b = -.15, z = 3.70, p < .001, d_z = -.15$ for blocking/unfollowing, and $b = .02, z = .69, p = .246, d_z = .03$ for sharing.

The results were the same in the pre-registered robustness check (i.e., excluding headlines misidentified as true), except that the headline-type X thinking-condition interaction was significant for intentions to “like” the headlines, $b = -.07, SE_b = .04, z = 1.69, p = .046$. Simple slopes fit our prediction that the effect of exposure to the headlines would be larger in the intuitive-thinking condition than in the deliberative-thinking condition, respectively, $bs = .08$ and $.01, SE_{bs}$ both = $.03, zs = 2.97$ and $.45, ps = .002$ and $.33, d_zs = .10$ and $.02$.

Exploratory, two-tailed tests of main effects, shown in Table 2, found that when people thought deliberately (vs. intuitively), they were significantly less inclined to “like” and share the headlines, and significantly more inclined to block/unfollow a person who posted the headlines. When judging previously-seen vs. new headlines, people were significantly more

inclined to “like” and less inclined to block/unfollow, but they were not significantly more inclined to share.

Moderated mediation analysis. We also conducted a pre-registered test of whether the indirect effect from repetition to moral condemnation to these intended social-media behaviors, observed in Experiment 1, would be larger in the intuitive-thinking (vs. deliberative-thinking) condition. The results offered only mixed support for this possibility: The difference in indirect effects was only significant in the robustness check, $p = .016$, and not in the initial analysis, $p = .057$. The intuitive-thinking condition, however, robustly replicated Experiment 1’s indirect effect for all three intended social-media behaviors, $ps < .001$ in both the initial analysis and the robustness check (see SOM-R, Appendix 5).

Discussion

Among participants instructed to think intuitively, previous encounters with fake-news headlines made those headlines seem less unethical to publish, which correlated in a mediation analysis with a stronger inclination to “like” and share those headlines, and a weaker inclination to block/unfollow someone who shared them – consistent with our prior studies. However, the evidence was not sufficient to conclude that instructing people to think deliberately attenuated these effects. [*Word count: 67*]

Experiment 4

To distinguish moral judgments from accuracy beliefs, Experiments 1-3 warned participants that the headlines were false. Social media companies have used such warnings, but people usually encounter fake news without warnings. Thus, to increase generalizability, Experiment 4 omitted the warning, but measured and controlled for accuracy beliefs.

Experiment 4 also sought to rule out two alternative explanations for why repetition reduces moral condemnation. Exposure to a fake-news headline could increase how much people like it ([cf. Zajonc, 1968](#)), or how popular they believe it is ([Kwan, Yap, & Chiu, 2015](#); [Weaver, Garcia, Schwarz, & Miller, 2007](#)). We measured and controlled for these variables.

Finally, going beyond Experiment 1 and 3's behavioral intentions measures, Experiment 4 tested whether repeatedly encountering fake-news headlines could increase people's likelihood of actually sharing them in an experimental setting. [*Word count: 133*]

Method

Participants. We posted slots for 300 American Prolific Academic users in March 2019. We selected this sample size based on the intuition that doubling Experiment 1's sample size would provide more than adequate statistical power to detect the same manipulation's effect. Whereas the majority of participants in Experiments 1-3 were Democrats, Experiment 4 used Prolific's prescreen data to request an equal number of Democrats and Republicans. People were unable to begin the study if they failed a reading-comprehension question or a CAPTCHA test to identify bots, or they responded from a mobile device or non-U.S. IP address. After applying the same exclusion criteria as our previous studies, the dataset contained 3,552 observations from 296 participants (147 men, 147 women, 2 non-binary; $M_{age} = 34$ years, $SD = 13$, range 18 to 74; 151 people who identified as or leaned Democrat, 142 who identified as or leaned Republican, and 3 who did not complete the politics measure). A sensitivity analysis suggested that the

smallest mean difference this sample could detect in Experiment 4's design at approximately 80% power, $\alpha = .05$ (two-tailed), was $d_z = .15$ (see SOM-R, Appendix 1).

Procedure. The procedure closely followed our previous studies. In the familiarization phase, participants viewed six fake-news headlines (randomly selected from our bank of 12) four times, rating how engaging, funny, surprising, and interesting they found each (filler items). In the judgment phase (after Experiment 1's distractor task), they rated each of the 12 headlines on the below measures. To better approximate real-world social media experiences, Experiment 4 – unlike Experiments 1-3 – did *not* inform participants the headlines were fake until an end-of-study debriefing.

Measures. Because Experiment 4 assessed more variables than our previous studies, we measured each variable with a single item to minimize participant fatigue.

Moral condemnation. The primary dependent measure was how *unethical* the headline would be to share (0 = *Not at all*, 100 = *Extremely*; responses selected by moving a slider that initially appeared at the midpoint). Whereas Experiment 1-3 asked about the unethicity of *publishing* the headline, Experiment 4 assessed generalizability by asking about *sharing*.

Control variables. We asked how *accurate* participants believed the headline to be and how much they *liked* the headline (slider: 0 = *Not at all*, 100 = *Extremely*), and how *popular* they thought the headline was (i.e., what percentage of American social media users they estimated had seen it before; slider: 0% labelled *No One* and 100% labelled *Everyone*). We also included an exploratory measure of how *well-written* the headline was, which we do not discuss further. We randomized the order of these four measures for each headline. (Note that Experiment 4 measured moral condemnation and accuracy beliefs on the same 100-point scale, in contrast to Experiment 1, which used a 100-point scale for unethicity and a 4-point scale for accuracy to maintain consistency with prior research).

Sharing intentions and behavior. Two measures assessed participants' inclination to share the headlines as potential downstream consequences of moral condemnation. First, as a behavioral intentions measure, we asked how likely they would be to share the headline if an acquaintance had shared it (1 = *Not at all likely*, 4 = *Somewhat likely*, 7 = *Extremely likely*). For each headline, participants answered this item after rating its unethicity, accuracy, how much they liked it, and how popular and well-written it was.

Second, as a behavioral measure, we told participants after they had rated all the headlines, "We would like to run a study where research participants see some of the headlines you've been looking at today. These participants can choose to click on the headlines to read the full article." Then they saw a list of all 12 headlines again (order randomized) and had to "choose four to share with the participants in our next study." We required four selections to avoid floor effects and to reduce error variance by controlling for individual differences in propensity to share headlines. We expected people to select more previously-seen headlines than new headlines.

Results

Dependent measure: moral condemnation. Replicating Experiments 1-3, people rated the headlines they had been shown earlier as less unethical to share than new headlines, as indicated by a significant effect of headline type (1 = *Previously seen*; 0 = *New*) in a multi-level regression model with fixed effects for the specific headlines and random intercepts for participants, $b = -4.64$, $SE(b) = .89$, $z = 5.20$, $p < .001$, $d_z = -.28$ (see Table 3 for detailed statistics). The effect was robust when we added the three control variables to the equation (adjusted $M_s = 56.11$ vs. 54.12 , 95% $CI_s = [53.57, 58.66]$ and $[51.57, 56.66]$, respectively), $b = -2.00$, $SE(b) = .79$, $z = 2.51$, $p = .012$, $d_z = .12$ (d_z calculated from the difference between adjusted means and the SD of the difference in raw means). Thus, the effect of prior exposure on moral

condemnation appears to be independent from its effect on judgments of accuracy, liking, and popularity.

Sharing intentions. People were more inclined to share previously-seen (vs. new) headlines, $p < .001$, $d_z = .34$ (see Table 3) in a multi-level regression model including fixed effects for the specific headlines and random intercepts for participants. Recall that a mean difference in the same direction was significant in Experiment 1, $p = .001$, but not in Experiment 3's intuitive-thinking condition, $p = .107$. Experiment 3 thus increases confidence in this effect's robustness.

As in Experiments 1 and 3, we also predicted that moral condemnation would mediate an indirect effect of repetition on sharing intentions, and it did, $b = .05$, $SE(b) = .010$, $z = 4.97$, $p < .001$. Suggesting that this mediation effect was independent of accuracy, liking, and popularity judgments, it remained significant when we added parallel mediation paths through each of these three control variables (see SOM-R, Appendix 6). This finding reduces concerns about omitted-variable bias in mediation (see [Bullock et al., 2010](#)).

Sharing behavior. Participants shared more headlines that they had repeatedly encountered ($M = 2.25$, $SD = .91$) than headlines they had not ($M = 1.75$, $SD = .91$), paired $t(294) = 4.73$, $p < .001$, $d_z = .27$. (One participant did not complete this measure). Recall that participants were required to share four headlines. Thus, a statistically equivalent way of understanding this effect is that participants shared more repeatedly-encountered headlines ($M = 2.25$) than would be expected by chance ($M = 2.00$), single-sample $t(294) = 4.73$, $p < .001$, $d_z = .27$. (Because the number of shared headlines was measured at the participant level, whereas moral condemnation was measured at the headline level, we could not perform a meaningful mediation analysis).

Table 3

Experiment 4's Results

Measure	Response options	<i>M</i> (previously seen)	95% <i>CI</i> (previously seen)	<i>M</i> (new)	95% <i>CI</i> (new)	<i>M</i> (diff)	<i>d_z</i>	<i>b</i>	<i>SE(b)</i>	<i>z</i>	<i>p</i>
Moral condemnation	0-100	52.80	[50.14, 55.46]	57.43	[54.77, 60.09]	-4.63	-.28	-4.64	.89	5.20	<.001
Accuracy	0-100	28.89	[26.71, 31.08]	24.06	[21.83, 26.20]	4.83	.35	4.88	.72	6.82	<.001
Liking	0-100	28.53	[26.46, 30.59]	25.24	[23.18, 27.31]	3.29	.21	3.28	.86	3.28	<.001
Popularity	0-100	36.29	[33.87, 38.72]	33.45	[31.03, 35.88]	2.81	.27	2.84	.54	5.18	<.001
Sharing intentions	1-7	1.85	[1.71, 1.98]	1.65	[1.51, 1.78]	.20	.34	.20	.03	6.00	<.001
Sharing behavior	0-1	.38	[.36, .39]	.29	[.27, .31]	.08	1.48	.39	.07	5.39	<.001

Note. For the sharing-behavior measure, the effect size is an odds ratio (*OR*) and the results are from logistic regression. For all other measures, the effect size is Cohen's *d_z*, and the results are from linear regression with random intercepts for participants.

Discussion

Extending Experiment 1-3's generalizability, Experiment 4 showed that repeated encounters with a fake-news headline can reduce moral condemnation of sharing it when people are *not* informed the headline is fake. This effect was robust when controlling for known consequences of repetition (judgments of accuracy, liking, and popularity), casting doubt on alternative explanations. Suggesting a relationship between moral judgments and social-media behaviors, a mediation analysis again showed that reduced moral condemnation of repeated (vs. new) headlines correlated with stronger intentions to share the headlines. Beyond intentions, people were more likely to actually share repeated (vs. new) headlines.

[*Word count: 96*]

General Discussion

Four experiments and a pilot suggest that repeatedly encountering a piece of misinformation can make it seem less unethical to spread. Seeing a fake-news headline four times or just once reduced how unethical people thought it was to publish and share the headline when they saw it again, even when it was labelled false. The results also demonstrate for the first time that repeating a fake-news headline can increase the likelihood that people will share it in an experimental context, can strengthen their intentions to promote it on social media, and can reduce their intentions to block/unfollow someone who shares it. Moral judgments correlated with these intentions in mediation analyses.

Amplified across the billions of active social media users, the small effects we observed of repetition on moral judgments (d_z s between .15 and .26) could have meaningful real-world consequences. By weakening moral condemnation, repeatedly encountering misinformation may contribute to its spread and reduce censure of those who spread it.

These findings represent the first evidence that misinformation receives less moral condemnation when it has been repeated. Future research should investigate why; our studies did not directly test mechanism. Repeated stimuli are more fluent than non-repeated stimuli, and perhaps fluency informed participants' moral intuitions. Whereas prior scholarship debates whether reason or emotion is the primary input into moral judgments (see [Monin, Pizarro, & Beer, 2007](#)), we posit that “meta-cognitive experiences” like fluency ([Schwarz, 2004](#)) could be a third, underappreciated input ([Laham, Alter, & Goodwin, 2009](#)).

How might fluency affect moral judgments? We speculate that because people associate fluency with truth, repeated misinformation feels more intuitively true even when people do not believe it, making it seem less unethical to spread. This possibility awaits further tests; Experiment 3's evidence was insufficient to conclude that decreasing intuitive thinking reduced the repetition effect.

Alternatively, because fluency feels good, people could evaluate repeated misinformation more positively overall without drawing on a fluency-truth association ([Reber, Schwarz, & Winkielman, 2004](#)). However, this alternative is less plausible given that repeating misinformation reduced moral condemnation even when controlling for general positivity (i.e., liking, Experiment 4). For the same reason, our findings differ from the mere-exposure effect ([Zajonc, 1968](#)), by which repetition increases liking – although they could be interpreted as the first non-correlational evidence of a “*moral* mere-exposure effect” ([Jarudi, Kreps, & Bloom, 2008](#)).

Fluency is not the only possible mechanism. People experience weaker affective reactions to previously-encountered (vs. new) stimuli ([Dijksterhuis & Smith, 2002](#)). Perhaps this *affective habituation* process leads people to respond with less negative affect – and thus less

moral condemnation ([Haidt, 2001](#)) – to a repeated (vs. initial) encounter with a fake-news headline.

We also considered the possibility that repeatedly encountering a particular headline would make people think it had been more widely shared, and was thus more acceptable to share (*cf.* [Weaver et al., 2007](#)). However, this possibility struggles to entirely account for our findings because repetition reduced moral condemnation even after controlling for judgments of popularity (i.e., how many people had seen the headline; Experiment 4).

Our findings are also distinct from the illusory-truth effect. Repeating fake-news headlines reduced the moral condemnation of spreading them when we eliminated the illusory truth effect by clearly labeling the headlines as false (Experiment 1), when we limited data analysis to headlines that people correctly identified as false (Experiments 1-3), and when we replicated the illusory truth effect but statistically controlled for accuracy beliefs (Experiment 4).

Efforts to curb misinformation typically aim to help people distinguish fact from fiction ([Lazer et al., 2018](#); [Scheufele & Krause, 2019](#)). For example, Facebook has tried informing users when they try to share news that fact-checkers have flagged as false. We suspect that these efforts will be insufficient as long as people find it more morally permissible to share previously-encountered (vs. new) information they *know* is false. Future research is needed to understand whether moral intuitions causally affect sharing behavior in real social media environments.

The wider misinformation spreads, the more likely the same individual will be to encounter it multiple times. And encountering it multiple times could reduce their moral condemnation of it and license them to spread it further. [*GD word count: 687*]

Acknowledgements

We thank Thomas Mussweiler and Scott Wiltermuth for insightful feedback, and Crystal Sun and the LBS Research Lab for research assistance. [*Word count: 21*]

Author Contributions

Both authors developed the research concept, designed the studies, and collected the data. D. Effron analyzed the data and wrote the manuscript, and M. Raj provided critical revisions. Both authors approved the manuscript's final version.

Open Practices

Experiments 2 and 3 were pre-registered at <https://aspredicted.org/uj7nd.pdf> and <https://aspredicted.org/3fc65.pdf>, respectively. The verbatim text of study materials is posted in the Supplemental Online Material (SOM-R). The ethics board overseeing this research withheld permission to post the data online, but the data are available from the lead author upon request.

Total word count, excluding methods and results: 1,996

= 753+126+34+26+53+67+133+96+687+21

References

- Al-Rodhan, N. (2017). Post-truth politics, the fifth estate and the securitization of fake news. *Global Policy*. Retrieved from <https://www.globalpolicyjournal.com/blog/07/06/2017/post-truth-politics-fifth-estate-and-securitization-fake-news>
- Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2), 211-236.
- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychological Review*, 13(3), 219-235. doi:10.1177/1088868309341564
- Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General*, 121(4), 446-458. doi:10.1037/0096-3445.121.4.446
- Bullock, J. G., Green, D. P., & Ha, S. E. (2010). Yes, but what's the mechanism? (Don't expect an easy answer). *Journal of Personality and Social Psychology*, 98(4), 550-558.
- Chadwick, A., & Vaccari, C. (2019). *News sharing on UK social media: Misinformation, disinformation, and correction*. Loughborough: Online Civic Culture Centre, Loughborough University. Retrieved from <https://dspace.lboro.ac.uk/dspace-jspui/bitstream/2134/37720/1/Chadwick%20Vaccari%20O3C-1%20News%20Sharing%20on%20UK%20Social%20Media.pdf>.
- Dechêne, A., Stahl, C., Hansen, J., & Wänke, M. (2010). The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychological Review*, 14(2), 238-257. doi:10.1177/1088868309352251

- Dijksterhuis, A., & Smith, P. K. (2002). Affective habituation: Subliminal exposure to extreme stimuli decreases their extremity. *Emotion, 2*(3), 203-214. doi:10.1037/1528-3542.2.3.203
- Effron, D. A. (2018). It could have been true: How counterfactual thoughts reduce condemnation of falsehoods and increase political polarization. *Personality and Social Psychology Bulletin, 44*(5), 729-745. doi:10.1177/0146167217746152
- Effron, D. A., & Miller, D. T. (2012). How the moralization of issues grants social legitimacy to act on one's attitudes. *Personality and Social Psychology Bulletin, 38*(5), 690-701. doi:10.1177/0146167211435982
- Fazio, L. K., Brashier, N. M., Payne, B. K., & Marsh, E. J. (2015). Knowledge does not protect against illusory truth. *Journal of Experimental Psychology: General, 144*(5), 993-1002. doi:10.1037/xge0000098
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives, 19*(4), 25-42.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review, 108*, 814-834.
- Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior, 16*(1), 107-112. doi:10.1016/S0022-5371(77)80012-1
- Jarudi, I., Kreps, T., & Bloom, P. (2008). Is a refrigerator good or evil? The moral evaluation of everyday objects. *Social Justice Research, 21*(4), 457-469.
- Kwan, L. Y. Y., Yap, S., & Chiu, C.-y. (2015). Mere exposure affects perceived descriptive norms: Implications for personal preferences and trust. *Organizational Behavior and Human Decision Processes, 129*, 48-58. doi:10.1016/j.obhdp.2014.12.002

- Laham, S. M., Alter, A. L., & Goodwin, G. P. (2009). Easy on the mind, easy on the wrongdoer: Discrepantly fluent violations are deemed less morally wrong. *Cognition*, *112*(3), 462-466.
- Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., . . . Rothschild, D. (2018). The science of fake news. *Science*, *359*(6380), 1094-1096.
- McGlone, M. S., & Tofiqbakhsh, J. (2000). Birds of a feather flock conjointly (?): Rhyme as reason in aphorisms. *Psychological Science*, *11*(5), 424-428.
- Mitchell, A., Barthel, M., & Holcomb, J. (2016). Many Americans believe fake news is sowing confusion. *Pew Research Center*. Retrieved from <http://www.journalism.org/2016/12/15/many-americans-believe-fake-news-is-sowing-confusion/>
- Monin, B., Pizarro, D. A., & Beer, J. S. (2007). Deciding vs. reacting: Conceptions of moral judgment and the reason-affect debate. *Review of General Psychology*, *11*(null), 99.
- Mueller, A. B., & Skitka, L. J. (2018). Liars, damned liars, and zealots: The effect of moral mandates on transgressive advocacy acceptance. *Social Psychological and Personality Science*, *9*, 711-718. doi:10.1177/1948550617720272
- Newman, E. J., Garry, M., Bernstein, D. M., Kantner, J., & Lindsay, D. S. (2012). Nonprobative photographs (or words) inflate truthiness. *Psychonomic Bulletin and Review*, *19*(5), 969-974.
- Nordgren, L. F., & Dijksterhuis, A. (2008). The devil is in the deliberation: Thinking too much reduces preference consistency. *Journal of Consumer Research*, *36*(1), 39-46.
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, *70*, 153-163.

- Pennycook, G., Cannon, T., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, *147*, 1865–1880.
doi:10.1037/xge0000465
- Pennycook, G., & Rand, D. G. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, *188*, 39-50.
doi:10.1016/j.cognition.2018.06.011
- Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition*, *8*(3), 338-342.
- Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review*, *8*(4), 364-382.
- Scheufele, D. A., & Krause, N. M. (2019). Science audiences, misinformation, and fake news. *Proceedings of the National Academy of Sciences*, 201805871.
doi:10.1073/pnas.1805871115
- Schwarz, N. (2004). Metacognitive experiences in consumer judgment and decision making. *Journal of Consumer Psychology*, *14*(4), 332-348.
- Shalvi, S., Dana, J., Handgraaf, M. J. J., & De Dreu, C. K. W. (2011). Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior. *Organizational Behavior and Human Decision Processes*, *115*(2), 181-190.
- Shidlovski, D., Schul, Y., & Mayo, R. (2014). If I imagine it, then it happened: The implicit truth value of imaginary representations. *Cognition*, *133*(3), 517-529.
- Simonsohn, U. (2014). No-way interactions. Retrieved from <http://datacolada.org/17>
- Skitka, L. J., & Bauman, C. W. (2008). Moral conviction and political engagement. *Political Psychology*, *29*(1), 29-54.

- Skurnik, I., Yoon, C., Park, D. C., & Schwarz, N. (2005). How warnings about false claims become recommendations. *Journal of Consumer Research*, *31*(4), 713-724.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, *359*(6380), 1146-1151.
- Weaver, K., Garcia, S. M., Schwarz, N., & Miller, D. T. (2007). Inferring the popularity of an opinion from its familiarity: A repetitive voice can sound like a chorus. *Journal of personality and social psychology*, *92*(5), 821-833.
- Whittlesea, B. W., Jacoby, L. L., & Girard, K. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language*, *29*(6), 716-732.
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of personality and social psychology*, *9*, 1-27. doi:10.1037/h0025848