



Overconfidence in news judgments is associated with false news susceptibility

Benjamin A. Lyons^{a,1}, Jacob M. Montgomery^b, Andrew M. Guess^c, Brendan Nyhan^d, and Jason Reifler^e

^aDepartment of Communication, University of Utah, Salt Lake City, UT 84112; ^bDepartment of Political Science, Washington University in St. Louis, St. Louis, MO 63130; ^cDepartment of Politics, Princeton University, Princeton, NJ 08544; ^dDepartment of Government, Dartmouth College, Hanover, NH 03755; and ^eDepartment of Politics, University of Exeter, Exeter EX4 4RJ, United Kingdom

Edited by Peter John Loewen, University of Toronto, Toronto, Canada, and accepted by Editorial Board Member Margaret Levi April 13, 2021 (received for review September 22, 2020)

We examine the role of overconfidence in news judgment using two large nationally representative survey samples. First, we show that three in four Americans overestimate their relative ability to distinguish between legitimate and false news headlines; respondents place themselves 22 percentiles higher than warranted on average. This overconfidence is, in turn, correlated with consequential differences in real-world beliefs and behavior. We show that overconfident individuals are more likely to visit untrustworthy websites in behavioral data; to fail to successfully distinguish between true and false claims about current events in survey questions; and to report greater willingness to like or share false content on social media, especially when it is politically congenial. In all, these results paint a worrying picture: The individuals who are least equipped to identify false news content are also the least aware of their own limitations and, therefore, more susceptible to believing it and spreading it further.

overconfidence | misinformation | social media

Concern about public susceptibility to false news is widespread. However, though Americans believe confusion caused by false news is extensive, relatively few indicate having seen or shared it (1)—a discrepancy that suggests that members of the public may not only have a hard time identifying false news, but also fail to recognize their own deficiencies at doing so (2–5). Such overconfidence may make individuals more likely to inadvertently expose themselves to misinformation and to participate in its spread. If people incorrectly see themselves as highly skilled at identifying false news, they may unwittingly be more likely to consume, believe, and share it, especially if it conforms to their worldview.

Overconfidence plays a key role in shaping behavior, at least in some domains (e.g., refs. 6–10). However, we know very little about its potential role in the spread of false news. Even basic descriptive data on the phenomenon of overconfidence in news discernment (the ability to distinguish false from legitimate news) are yet to be established. How pervasive is overconfidence? Is overconfidence related to false news exposure? Are overconfident individuals actually more likely to hold misperceptions or share false stories? We currently lack answers to these questions.

In this paper, we examine the relationship between perceived and actual ability to distinguish between false and legitimate information, drawing on a theoretical framework for understanding biased self-perception (4). In two large, nationally representative samples ($N = 8,285$), respondents completed a discernment task evaluating the accuracy of a series of headlines as they appear on Facebook. They were further asked to rate their own abilities in discerning false news content relative to others. We use these two measures to assess overconfidence among respondents and how it is related to beliefs and behaviors.

Our results paint a worrying picture. The vast majority of respondents (about 90%) reported that they are above aver-

age in their ability to discern false and legitimate news headlines, meaning that many Americans substantially overestimate their abilities. Accordingly, people's self-perceptions are only weakly correlated with actual performance. Further, using data measuring respondents' online behavior, we show that those who overrate their ability more frequently visit websites known to spread false or misleading news. These overconfident respondents are also less able to distinguish between true and false claims about current events and report higher willingness to share false content, especially when it aligns with their political predispositions. Although discernment ability is a strong predictor of these outcomes, an alternative analysis using a "residualized" measure of overconfidence net of actual ability also explains additional variance in these behaviors.

In the next section, we review existing research on overconfidence and how we expect it to operate for news discernment. *Materials and Methods* describes our research design, including a task assessing respondents' news-discernment abilities. Our results show that overconfidence is both common and associated with a range of undesirable media-related behaviors. Although our design does not allow us to identify the causal effect of overconfidence, these findings suggest that the mismatch between one's perceived ability to spot false stories and people's actual abilities may play an important and previously unrecognized role in the spread of false information online.

Significance

Although Americans believe the confusion caused by false news is extensive, relatively few indicate having seen or shared it—a discrepancy suggesting that members of the public may not only have a hard time identifying false news but fail to recognize their own deficiencies at doing so. If people incorrectly see themselves as highly skilled at identifying false news, they may unwittingly participate in its circulation. In this large-scale study, we show that not only is overconfidence extensive, but it is also linked to both self-reported and behavioral measures of false news website visits, engagement, and belief. Our results suggest that overconfidence may be a crucial factor for explaining how false and low-quality information spreads via social media.

Author contributions: B.A.L., J.M.M., A.M.G., B.N., and J.R. designed research; B.A.L. analyzed data; B.A.L. and J.M.M. wrote the paper; and B.A.L., J.M.M., A.M.G., B.N., and J.R. revised the paper.

The authors declare no competing interest.


This article is a PNAS Direct Submission. P.J.L. is a guest editor invited by the Editorial Board.

Published under the PNAS license.

¹ To whom correspondence may be addressed. Email: ben.lyons@utah.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2019527118/-/DCSupplemental>.

Published May 28, 2021.



Who Spreads False News? Which individuals are more likely to engage with, believe, and spread dubious news? One body of research emphasizes the role of partisan predispositions or motivated reasoning in the assessment of news content (11) and exposure to it and sharing of it (12, 13). A second literature considers how improving individuals' information evaluation and digital literacy skills can reduce their vulnerability to false information online (14, 15). Finally, other studies focus on the role of purposeful reasoning processes in reducing individual vulnerability to misinformation. People who think more analytically or are more deliberative in their evaluation of news claims rate false news as less accurate (16, 17). Conversely, people who tend to rely on emotion as they process information or wrongly claim familiarity with nonexistent entities are more likely to see false headlines as accurate (18, 19).

Our research builds on cognitive style accounts by examining the disparity between people's ability to spot false news and their beliefs about their skill in doing so. This approach is intended to assess the contribution of cognition as well as metacognition to engagement behaviors. As we argue below, overconfidence in one's ability to distinguish between legitimate and false news may help account for whether and how individuals engage with false or dubious online content (e.g., liking or sharing). To put the point more directly, some portion of the public is likely to be especially vulnerable to false information precisely because they do not realize that they are, in fact, vulnerable to false information. As a result, these individuals may be more likely to unknowingly consume, believe, and share false news.

The Dunning–Kruger Effect for News Discernment. Building on prior studies of perceptual bias in self-assessments, we test for a Dunning–Kruger effect (DKE) in false news discernment. The DKE describes a general tendency of poor performers in social and intellectual domains to be unaware of their own deficiency (4). By contrast, the most competent performers slightly underestimate their own ability relative to others due to a form of false consensus effect in which they assume others are performing more similarly to themselves than they really are (20). This pattern arises whether researchers elicit comparative self-evaluations (ratings of performance relative to peers) or self-evaluations using absolute scales (5).

DKE research contends that poor performers suffer from a double bind: Not only does a lack of expertise produce errors in the first place, it also prevents recognition of these errors and awareness of others' capabilities. In studies of perception and performance, people in the bottom quartile of performers have tended to provide the most upwardly distorted self-perceptions. For instance, Anson (21) finds that individuals who perform worst on a quiz measuring basic political knowledge rate their own performance the same or even better than high performers.

The reported overconfidence of underperformers is not erased by financial or social incentives (6) and is corroborated by real-world behavior [e.g., in (not) selecting insurance for examination performance (7)]. These studies suggest that low performers genuinely believe in their own abilities and are not simply making face-saving expressions of self-worth. Further, past research shows that overconfidence is more common when people have reason to see themselves as knowledgeable or competent—i.e., if the subject is not arcane and is prevalent in everyday life (5). Given its familiarity, judgments of news accuracy are likely to fit the DKE pattern, as do knowledge about either politics (21) or vaccines (22). We therefore propose the following research question:*

*We filed a preregistration for this project prior to accessing the data. We report a "populated preanalysis plan" (23) that details our preregistered hypotheses and analysis plan and identifies which main text findings are preregistered in [SI Appendix E](#).

Research Question 1. *To what extent will people who are least accurate at distinguishing between legitimate and false news overrate their ability to distinguish mainstream from false news?*

Importantly, the DKE predicts that low performers will not recognize how poorly they performed in relative terms, not that low performers will think they perform best. We therefore do not expect that low performers will think that they are the best at our task of distinguishing between legitimate and false news. Instead, we will examine the extent to which poor performers do not recognize that they are worse than most others at the task.

Does Overconfidence Matter? Importantly, the DKE may have downstream effects on behavior. Because overconfident individuals fail to recognize their own poor performance, they are less able to improve their domain-specific skills. For instance, several studies find that overconfident individuals learn the least in classroom settings (24).[†] We therefore expect that overconfidence in news discernment will be associated with a variety of tendencies including exposure to false news, belief in its accuracy, and sharing it with others.

To begin, we expect a positive association between overconfidence and visits to false news websites. The DKE implies less ability to discern which news stories are false when an individual is exposed (e.g., on a social media platform) combined with lesser awareness of this discernment deficiency, which would lead to greater incidental exposure to false news stories. Similarly, overconfidence may be seen as a form of invulnerability bias in which assumed mastery leads people to feel little need to take preventative actions (e.g., to be cautious or engage in deliberate thinking about which sites one visits), which may produce additional exposure to questionable media messages (28, 29). We therefore propose the following research question:

Research Question 2. *Is overconfidence in one's ability to distinguish mainstream from false news positively related to false news exposure?*

In addition, overconfidence may make people less likely to question a dubious news story's veracity, as high confidence is associated with less reflection (30, 31). As a result, people who are overconfident may be more willing to accept false claims and to engage with false content in the form of liking or sharing these stories [similarly, recent work suggests that people generally lack awareness of their susceptibility to inaccurate "general knowledge" claims they come across when reading works of fiction (32)]. Further, previous research indicates that individuals are generally more likely to believe false claims when they are consistent with their own prior political beliefs (33). Therefore, we would expect that the relationship between overconfidence and beliefs and engagement will be strongest when the content involved aligns with respondents' partisan preferences.

Research Question 3. *1) Is overconfidence positively related to holding misperceptions on specific topics? 2) Is this relationship stronger when the claim is politically congenial?*

Research Question 4. *1) Is overconfidence positively related to self-reported willingness to like or share false content? 2) Is this relationship stronger when the claim is politically congenial?*

[†]A related literature details the confidence with which individuals hold political misperceptions (25–27). This work shows that many people are somewhat aware of their ignorance, and, therefore, many misperceptions are not confidently held (27), and these individuals are more likely than the confidently wrong to update their beliefs in response to corrections (26).

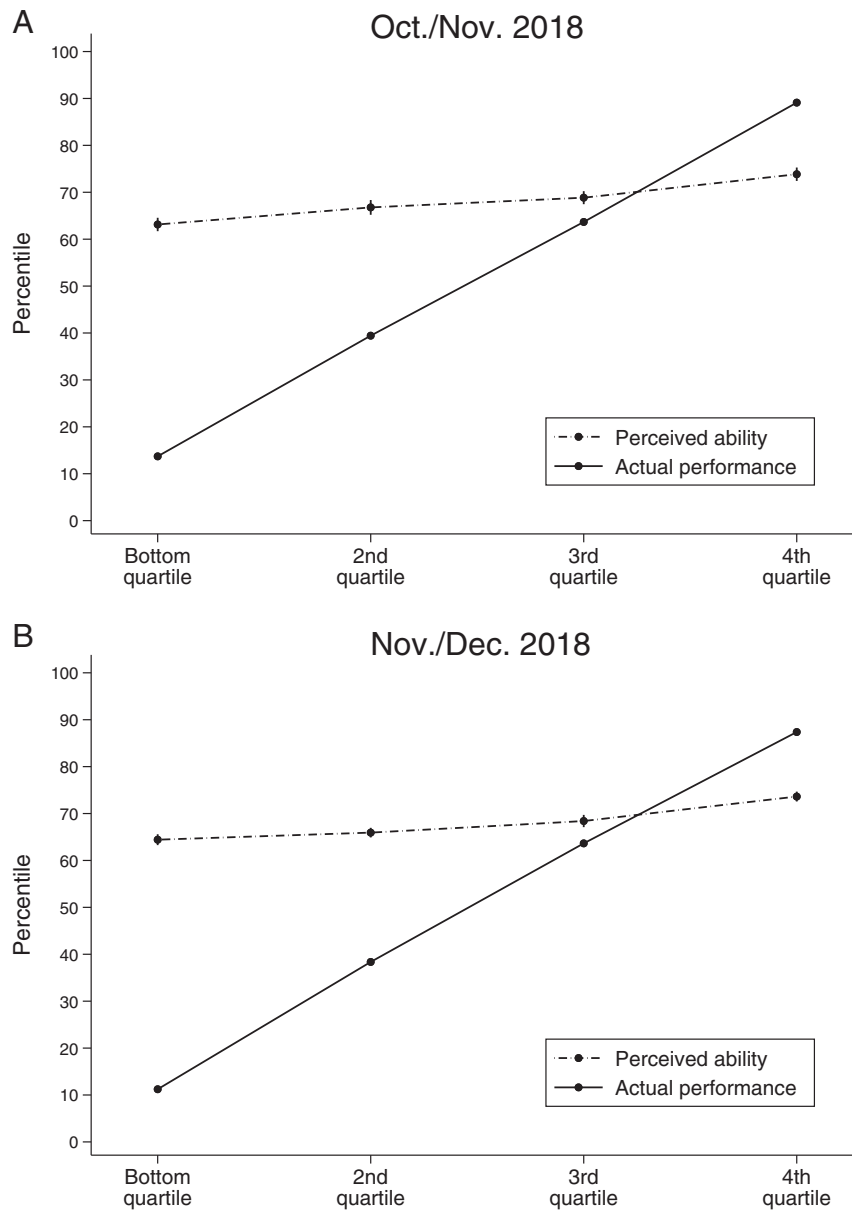


Fig. 1. Perceived false news detection ability for respondents grouped by actual performance. Notes: Gaps depict miscalibration between actual and self-assessed percentile of performance for quartile groups based on actual performance with 95% CIs (note: CIs are smaller than the markers for actual performance and thus not visible). Oct./Nov., $N = 2,855$; Nov./Dec., $N = 4,150$.

Results

We first describe the DKE in our data. We divide the sample into four quartiles based on respondents' actual performance in our discernment task. For each of these four groups, we calculate the mean score for both actual and perceived ability (percentiles ranging from 1 to 100), which we present in Fig. 1. As expected, actual performance closely tracks the idealized 45° line when we plot the mean performance score in each quartile. However, for perceived ability, we see a much flatter line. Perceived ability increases modestly across our measure of actual ability. The mean self-reported percentile for individuals in the bottom quartile in actual ability (i.e., the 1st to 25th percentile) is 63 in the October (Oct.)/November (Nov.) survey and 64 in the Nov./December (Dec.) survey. This quantity rises to only 74 for the top quartile in both surveys. In other words, those who are in the bottom quartile in actual performance rate themselves as being in about the 63rd/64th percentile, a vast overestimate of

their own performance. While those in the top quartile of actual performance rate their perceived ability higher than those in the bottom quartile do, they underestimate where they rank in actual ability.[‡]

In general, performance is only weakly associated with perceived ability (Oct./Nov., $r = 0.08$; Nov./Dec., $r = 0.10$), as shown in *SI Appendix, Fig. B3*. Moreover, average self-reported percentile (69th) is well above 50 (one-sample t -test, $P < 0.005$),

[‡]Critics of DKE analyses like the one presented in Fig. 1 argue that it does not reflect the proposed mechanism—metacognitive differences (i.e., perception accuracy) between high and low performers—and is instead the result of systematic bias or measurement error; e.g., regression to the mean and the better-than-average effect (34, 35). Proposed alternative accounts of the DKE have led to vigorous theoretical and empirical debates (5, 6, 36–41). Although no consensus has emerged, recent work suggests that metacognitive differences, general biases in self-estimation, and statistical artifacts each contribute to the DKE (41).

Table 1. Overconfidence and news exposure (binary measures)

	Oct./Nov.		Nov./Dec.		Pooled	
	False	Mainstream	False	Mainstream	False	Mainstream
Overconfidence	0.0609** (0.0231)	-0.0450 (0.0505)	0.0003 (0.0003)	-0.0007 (0.0006)	0.0569*** (0.0186)	-0.0415 (0.0411)
Constant	-0.0815* (0.0354)	0.4645*** (0.1010)	-0.0225 (0.0498)	0.3735*** (0.1199)	-0.0715* (0.0298)	0.4419*** (0.0799)
Control variables	✓	✓	✓	✓	✓	✓
R ²	0.17	0.11	0.08	0.17	0.11	0.11
N	1,780	1,780	767	767	2,547	2,547

Cell entries are OLS coefficients estimated using survey weights. The overconfidence measure subtracts the respondent's actual percentile from their self-rated percentile and is rescaled to range from -1 to 1. False news exposure is coded as one if the respondent visited any such domain and zero otherwise. Mainstream news exposure is coded as one if the respondent visited any such domain and zero otherwise. All models include controls for Democrat, Republican, college education, gender, non-White racial background, age, and media diet slant. * $P < .05$, ** $P < .01$, *** $P < .005$ (two-sided).

indicating that many people are overconfident. As Fig. 1 illustrates, this overconfidence is concentrated most heavily among individuals in the bottom quartile. That is, the individuals whose performance is objectively at the lowest level are the most overconfident in their abilities.

In line with prior work, male respondents display more overconfidence (7, 8, 42), and overconfidence is negatively associated with general political knowledge. There is no association with age (43), despite age-based disparities in exposure to false news (12, 13). Finally, Republicans are more overconfident than Democrats (44), which is not surprising given the lower levels of media trust they report (see *SI Appendix C*, which shows that mass media trust and media affect are both negatively associated with overconfidence). We report preregistered analyses regarding demographics in greater depth in a separate manuscript (see *SI Appendix E* for details).

False News Exposure. We next examine whether visits to false news websites are associated with overconfidence. Building on prior research examining the difference between subjective self-perceptions and objective performance (8–10, 45, 46), we measure this concept as the difference between self-reported relative performance and our objective measure of relative performance. As Parker and Stone (47) argue, the difference score measure we employ here is appropriate when the theoretical mechanism of interest is overconfidence rather than self-assessed ability per se (i.e., controlling for ability). We are interested in the miscalibration between these components because the DKE relies on the double bind of low ability paired with a lack of awareness.

We estimate ordinary least squares (OLS) models using binary measures of false and mainstream news exposure for both surveys (Oct./Nov. and Nov./Dec.). In each of these models, which are estimated using survey weights, we include a set of standard covariates as well as a measure of the ideological orientation of respondents' news diet. Finally, we rescale our measure of overconfidence to range from -1 to 1 rather than from -100 to 100 to aid in interpretation. Results are shown in Table 1 and Fig. 2. The baseline exposure rate to false news in the Oct./Nov. survey was 6.5%. We find that overconfidence is associated with greater rates of exposure in that survey ($\beta = 0.06$, $SE = 0.02$, $P < .01$). Specifically, respondents at the 95th percentile of overconfidence were about 6 percentage points more likely to have been exposed to false news in the postsurvey period than those at the 5th percentile, conditional on demographics. Similarly, those at the maximum value of overconfidence were about 11 percentage points more likely to have been exposed than those at the minimum. The relationship is not statistically significant for the Nov./Dec. model, but the sample size for that

survey is significantly reduced ($n = 767$). When we instead pool the data, the results are nearly identical to the results for the Oct./Nov. survey (see *SI Appendix, Table F1* for results from logit models, which are substantively identical). One concern is that overconfident individuals may simply be more (or less) likely to visit online news websites in general. To test for this, we also estimate identical regressions with mainstream news exposure as the dependent variable. We find that overconfidence is not associated with our binary measure of mainstream news exposure after accounting for demographics.

Topical Misperceptions. Next, we examine the association between overconfidence and ability to distinguish between true and false claims about political events that were topical at the time the surveys were fielded. Here we examine a misperceptions battery from the October/November survey measuring beliefs in claims related to Brett Kavanaugh's Supreme Court nomination. These regression models are again estimated using survey weights and our set of standard covariates. We also again rescale our measure of overconfidence to range from -1 to 1 rather than from -100 to 100 to aid in interpretation. The main results of this analysis are shown in Table 2. The first

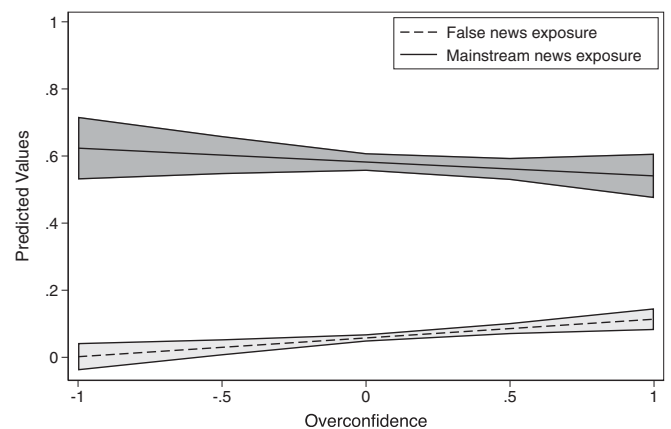


Fig. 2. Overconfidence and news exposure. Notes: Predictive margins with 95% CIs, based on full model, all other variables held constant. The overconfidence measure subtracts respondents' actual percentile from their self-rated percentile and is rescaled to range from -1 to 1. False news exposure is coded as one if the respondent visited any such domain and zero otherwise. Mainstream news exposure is coded as one if the respondent visited any such domain and zero otherwise. Data come from the pooled model ($N = 2,547$), which pools data from Oct./Nov. ($N = 1,780$) and Nov./Dec. ($N = 767$) surveys.

column shows the results for the two false statements provided to respondents. We include fixed effects for each statement to account for their differing baseline levels of plausibility and cluster at the respondent level to account for correlations between their ratings across headlines. We find no main effect of overconfidence on belief in these false claims in isolation (and no evidence that this relationship is moderated by congeniality). The second column, however, shows results for difference scores (discernment), which subtract the perceived accuracy of false claims from that of true claims. Higher discernment scores reflect greater belief in true statements relative to false ones. The negative coefficient ($\beta = -0.37$, $SE = 0.06$, $P < 0.005$) thus indicates that overconfidence is negatively associated with discernment ability on these topical claims.

Self-Reported Engagement. We next turn to our measure of self-reported engagement intention (intent to like or share a post on social media). These regression models again include a set of standard covariates, and we rescale our measure of overconfidence to range from -1 to 1 , rather than from -100 to 100 , to aid in interpretation. We use survey weights in all models. These results are shown in Table 3. The first and third columns show results at the headline level, where the outcome is a four-point scale of intention to either share or like a false story. Overconfidence has a clear positive relationship with liking or sharing false stories (Research Question [RQ] 4.1). Moreover, this relationship varies as a function of partisan congeniality (RQ4.2). Likewise, the second and fourth columns of Table 3 use the average difference in engagement intention across true and false headlines as a measure of discernment. The results show that overconfidence is negatively related to discernment in which stories respondents would engage with. Overconfident individuals are thus not merely more likely to engage with news content in general, but instead are specifically more inclined to share false stories versus mainstream ones relative to respondents who are less overconfident.

Alternative Specifications. We construct our primary independent variable above as: Overconfidence = (Perceived ability – Actual ability). We view this measurement strategy as appropriate for two reasons. First, this approach is consistent with how overconfidence has been measured in related studies of its behavioral

Table 2. Overconfidence and topical misperceptions

	False	Difference score
Overconfidence	0.1123 (0.0966)	–0.3667*** (0.0604)
Congeniality	0.8557*** (0.0475)	
Overconfidence × congeniality	–0.0390 (0.1296)	
Constant	1.9550*** (0.1056)	–0.2223* (0.0932)
Control variables	✓	✓
Statement fixed effects	✓	
R ²	0.16	0.15
N (statement)	4,872	
N (respondent)	2,444	2,904

Cell entries are OLS coefficients. Respondents rated the accuracy of four statements regarding the Kavanaugh appointment on four-point scales. The first model's outcome variable is perceived accuracy of false statements only. The second model's outcome variable is the difference in the mean perceived accuracy of true and false statements. The overconfidence measure subtracts the respondent's actual percentile from their self-rated percentile and is rescaled to range from -1 to 1 . Controls: Democrat, Republican, college education, gender, non-White racial background, and age. * $P < 0.05$; *** $P < 0.005$ (two-sided).

Table 3. Overconfidence and engagement intention

	Oct./Nov.		Nov./Dec.	
	False	Diff. score	False	Diff. score
Overconfidence	0.6690*** (0.0691)	–0.4051*** (0.0193)	0.7610*** (0.0466)	–0.3984*** (0.0283)
Congeniality	0.1893*** (0.0224)		0.1902*** (0.0151)	
Overconf × congenial	0.2394*** (0.0721)		0.2692*** (0.0511)	
Constant	1.0620*** (0.0242)	0.0520 (0.0270)	1.1020*** (0.0139)	0.0600 (0.0369)
Control variables	✓	✓	✓	✓
Headline fixed effects	✓		✓	
R ²	0.11	0.15	0.13	0.13
N (headline)	10,194		14,720	
N (respondent)	2,549	2,566	3,680	3,717

Cell entries are OLS coefficients. Dependent variables are based on self-reported intention to “like” or “share” each of the articles in the headline task on Facebook (four-point scales; 1 = not at all likely, 4 = very likely). These questions are asked only of respondents who report using Facebook. The first model's outcome variable is engagement intent for false headlines only. The second model's outcome variable is the difference in the mean engagement intent for mainstream and false headlines. The overconfidence measure subtracts the respondent's actual percentile from their self-rated percentile and is rescaled to range from -1 to 1 . Controls: Democrat, Republican, college education, gender, non-White racial background, and age. *** $P < .005$ (two-sided). Diff., difference score.

effects (8–10, 45, 46). Second, and more importantly, our theory is explicitly about the difference between perceived and actual ability, and not about the independent role of either component. Thus, the main regressions of interest are (broadly) structured as:

$$y_i = \gamma_0 + \gamma_1(\text{Perceived ability} - \text{Actual ability}) + \epsilon_i, \quad [1]$$

where y_i represents the outcome of interest and ϵ_i is our error term.

We use this specification because we have no theoretical expectations about the independent role of these predictors—our theory is about the mismatch between actual and perceived ability (47). However, it is still worthwhile to consider whether one of these factors (perceived or actual ability) is responsible for our results. In particular, it is important to try to isolate the effects of perceived ability given prior findings showing that false news belief and exposure are related to individual-level differences in analytical thinking skills and reasoning ability (13, 14, 16, 17).

We therefore consider two alternative model specifications that seek to estimate the direct association between perceived ability and our outcome measures independent of its relationship to actual ability below. First, we attempt to “residualize” perceived ability, an approach that has become standard in personality and social psychology (48). Second, we disaggregate the two components and include them independently in a regression. (These approaches are mathematically quite similar, but we include them both for the sake of completeness.[§]) To account for the fact that these three approaches each have unique weaknesses, recent work has suggested that all three be employed (48).

[§]Indeed, the results would be identical if we used the full set of covariates from the disaggregated regression in our residualization process.

Residualizing Perceived Ability. We begin by following the strategy outlined in Anderson et al. (49), which uses a residualized measure of perceived ability. Specifically, we first fit the regression

$$\text{Perceived}_i = \text{Actual}_i + \delta_i, \quad [2]$$

where δ_i is the residual error term. Assuming this model is correct, we can then use the estimated residual error $\hat{\delta}_i$ as a measure of perceived ability that is unrelated to actual ability. We then fit a model such as

$$y_i = \beta_0 + \beta_1 \hat{\delta}_i + \epsilon_i, \quad [3]$$

where β_1 is intended to represent the independent relationship between (residualized) perceived ability and the outcome.

With this residualization approach, we start with news exposure in *SI Appendix, Table F2*. We find a positive correlation between residualized perceived ability and exposure, but it is only statistically significant for the pooled model (Oct./Nov.: $\beta = 0.07, P > 0.05$; Nov./Dec.: $\beta = 0.02, P > 0.05$; pooled: $\beta = 0.06, P < 0.05$). This result differs from our primary analysis only in that the coefficient for the Oct./Nov. survey is not statistically significant, though as in the primary analysis, it is similar to the pooled coefficient. Turning to topical misperceptions, *SI Appendix, Table F3* shows that there is a significant interaction between residualized perceived ability and congeniality ($\beta = 0.59, P < 0.05$), indicating that overconfident individuals are more likely to believe in false statements that are consistent with their prior beliefs. This result is more favorable for our theory than the one reported in the primary analysis. However, unlike the primary results, residualized perceived ability is not significantly associated with decreased discernment between true and false claims. Finally, *SI Appendix, Table F4* shows that residualized perceived ability is positively associated with liking or sharing false stories (Oct./Nov.: $\beta = 0.33, P < 0.01$; Nov./Dec.: $\beta = 0.47, P < 0.005$). These relationships are strongest for congenial stories (Oct./Nov.: $\beta = 0.41, P < 0.005$; Nov./Dec.: $\beta = 0.47, P < 0.005$). However, there is, again, no significant association with discernment between mainstream and false news in either wave (Oct./Nov.: $\beta = -0.10, P > 0.05$; Nov./Dec.: $\beta = 0.03, P > 0.05$).

Disaggregating Overconfidence. Our second approach is to include perceived and actual ability as two separate independent variables in our model. Although our theory focuses on overconfidence, one might expect the coefficient for perceived ability to be positive and the coefficient for actual ability to be negative for the outcome measures we consider. To illustrate this idea, we simulate data according to formulas that assume a data-generating process in which overconfidence is linearly associated with some outcome measure per our theory (*SI Appendix, Table F5*). These results, which show that perceived ability is positively associated with the outcome and actual ability is negatively related to the outcome, suggest that the disaggregation approach will provide the correct conclusion.

However, it is important to emphasize a few important limitations before presenting our disaggregated results. First, this approach assumes that the component measures do not affect one another, despite the fact that self-perception and performance likely do so (5, 48). Second, this strategy is more difficult to interpret because both coefficients relate to the theory of interest. Increased perceived ability (controlling for actual ability) is an indicator for overconfidence, but so, too, is decreasing actual ability (controlling for perceived ability). Interpreting either coefficient in isolation with respect to our theory is therefore difficult, especially in more complex models [i.e., those that include interaction terms; see Parker and Stone (47) for more extensive discussion of this point]. Third, the simulated results are based on the assumption of constant levels of measurement

error between the perceived and actual ability variables. If measurement error varies between them, however, it may appear as if only one of the two variables is important, despite the fact that both are equally weighted in the true data generating process. To illustrate this point, we conduct a version of the simulation described above, but now add additional measurement error to the observed perceived ability variable included in the disaggregated regression. We thus assume the same data-generating process where overconfidence drives our results, but now add differential measurement error for perceived ability. The results in *SI Appendix, Table F6* now show a null result for perceived ability and a significant negative association with actual ability. Researchers who failed to consider the possibility of differential measurement error might mistakenly infer that it is only actual ability that drives these results. This scenario seems empirically plausible. A priori, we would not expect equal rates of measurement error between these components. Specifically, actual ability is measured via a series of 12 objective evaluation tasks that are combined into an aggregate score. By contrast, perceived ability is measured as the average of two self-assessment survey items. Standard psychometric theory would suggest higher rates of measurement error for the perceived ability indicator.

With these caveats, we turn to our disaggregated results below. First we re-examine RQ1, which predicts that overconfidence will be related to differential rates of exposure to false news websites. The disaggregated models are shown in *SI Appendix, Table F7*. Consistent with the extrapolation from our theory described above, the perceived and actual ability coefficients are signed in opposite directions, but only the actual ability coefficients are significant for the Oct./Nov ($\beta = -0.06, P < 0.05$) and pooled samples ($\beta = -0.06, P < 0.05$) (perceived ability: Oct./Nov. $\beta = 0.06$, pooled $\beta = 0.05$, both not significant). These results suggest either that actual ability is more important than perceived ability or is measured with less error, per our discussion above. As in the primary analysis, neither is significant for the Nov./Dec. sample. Next, we turn to the topical misperceptions results (*SI Appendix, Table F8*). In the primary analysis, we find no main effects or interactions with congeniality, but do find a main effect for the difference outcome. When we disaggregate, we do find main effects for the actual ability measure in both analyses. The interaction terms, however, tell a complicated story. There is a positive significant interaction between perceived ability and congeniality ($\beta = 0.62, P < 0.01$), indicating that more overconfident individuals are more likely to believe false claims that are congenial to their prior beliefs. However, there is also a positive significant coefficient for the interaction with actual ability ($\beta = 0.26, P < 0.05$), which suggests that overconfidence (decreased actual ability controlling for perceived ability) increases belief in false stories only when they are not congenial. Finally, we turn to the results for engagement intentions (*SI Appendix, Table F9*). For the headline-level analyses, the results again mirror the findings in the primary analysis, with both the perceived ability and actual ability coefficients being significant (but signed in opposite directions). The interactions with headline congeniality are also both significant and correctly signed. For the difference-score analysis, the results are more mixed. The actual ability coefficient is significant and positive for the Oct./Nov. sample ($\beta = 0.60, P < 0.005$) and the Nov./Dec. sample ($\beta = 0.61, P < 0.005$). However, the perceived ability coefficient is not significant for the Oct./Nov. sample ($\beta = 0.03, P > 0.05$) and significant, but incorrectly signed, for the Nov./Dec. analysis ($\beta = 0.19, P < 0.005$).

In all, these additional tests provide a somewhat mixed picture. While our results certainly show that not all results are driven purely by the actual ability measure, some of the evidence suggests that actual ability could be playing a crucial role for some of our results. However, we cannot rule out the

possibility that these differences are attributable to differential measurement error. In several cases, the point estimates for perceived and actual ability are quite similar, and the main differences in our inferences are the result of our estimates of perceived ability being more imprecise (e.g., *SI Appendix, Table F7*).

Discussion

We find that respondents tend to think they are better than the average person at news discernment, and perceived ability is only weakly associated with actual ability, with the worst performers also being the most overconfident. Importantly, overconfidence is associated with a range of normatively troubling outcomes, including visits to false news websites in online behavior data. The overconfident also express greater willingness to share false headlines and are less able to discern between true and false statements about contemporaneous news events. Notably, the overconfident are particularly susceptible to congenial false news. These results suggest that overconfidence may be a crucial factor for explaining how false and low-quality information spreads via social media.[¶] Many people are simply unaware of their own vulnerability to misinformation. Targeting these overconfident individuals could be an important step toward reducing misinformation on social media sites, though how best to do so remains an open question. Other research finds that the behavioral effects of high confidence and weak performance include resistance to help, training, and corrections (5, 26, 50). An incorrect view of one's ability to detect false news might reduce the influence of new information about how to assess media items' credibility, as well as willingness to engage with digital literacy programs. For this reason, it may be important to better understand the roots of overconfidence, from demographics (51) to domain involvement (52) to social incentives (49, 53), and how they apply in the case of perceptions of news discernment.

These results should also be understood in the context of their limitations. Most critically, our analyses are correlational and, thus, face concerns about endogeneity. In this context, we have specific *ex ante* reasons to suspect that the relationship between overconfidence and our outcome measures is at least partially endogenous. For instance, habitual exposure to false news might lead to poorly calibrated estimates of one's ability to detect it, especially given the tendency for people to treat incoming information as true and the subsequent effects this can have on feelings of fluency (54). Overconfidence and false news engagement could even mutually reinforce one another over time (55). Future work must determine the extent to which overconfidence plays a causal role in the behaviors with which we show it is associated. One possible direction would be to experimentally manipulate overconfidence by informing respondents about their relative performance. Another approach might be to manipulate individuals' self-perception by randomly assigning them a competency score, although such a study would require careful ethical consideration.

Beyond issues of endogeneity, it is important to carefully interpret the associations we detect. Based on prior work regarding the role of purposeful reasoning (16, 17) and literacy skills (14) in individual vulnerability to misinformation, we would assume discernment ability itself—from which our overconfidence measure is in part derived—drives engagement with this content. Unsurprisingly, we find that people who are worse at discerning between legitimate and false news in the context of a survey

are worse at doing so in their browsing habits. Further, actual ability is a stronger predictor than perceived ability, though the effect sizes are similar; as noted, this discrepancy may be a reflection of greater measurement error in our measure of perceived ability. However, our results also show that inflated perceptions of ability are independently associated with engaging with misinformation, suggesting that perceived ability net of actual ability may be a further source of vulnerability (i.e., an additional, metacognitive component). Specifically, when residualized, perceived ability net of actual ability is associated with dubious news site exposure, misperceptions, and sharing intent. It is not our goal here to argue that overconfidence supersedes ability itself as the key predictor or cause of vulnerability to misinformation, nor do our findings support this interpretation. Indeed, our results lend further support to work that shows ability deficits are a serious issue in this domain. Further, because excess confidence is associated with less reflection (30, 31), the ways that discernment ability and overconfidence influence engagement with dubious information may be linked. Ultimately, adjudicating between these accounts would require further improvements to the measurement of overconfidence, which remains a complicated endeavor in all research contexts (5, 47, 48). We rely on overconfidence as measured by the difference between actual and self-assessed performance on a news-discernment task. Future research should explore different approaches to measuring overconfidence in this domain and assess how they relate to who views, believes, and spreads false news content. In particular, scholars should consider how to measure perceived ability with more precision and/or seek to directly manipulate these concepts in isolation to understand their independent effects.

Finally, although we replicate our results in multiple samples, further efforts to demonstrate that the relationship we observe holds in other contexts would be valuable. First, work should validate these results with mobile data and with data that allow us to observe actual sharing behavior in addition to self-reported sharing [although they appear to correspond at least to some extent (56)]. Likewise, our data come from the American context, though based on cross-national findings regarding the pervasive nature of overconfidence (57), it is reasonable to believe the outcomes are not unique to the United States and may be even more worrisome elsewhere.

Ultimately, our results provide evidence of an important potential mechanism by which people may fall victim to misinformation and disseminate it online using survey and behavioral data from multiple large national samples. Understanding overconfidence may be an important step toward better understanding the public's vulnerability to false news and the steps we should take to address it.

Materials and Methods

To answer our research questions, we draw on data from two two-wave survey panels conducted by the survey company YouGov during and after the 2018 US midterm elections, allowing us to replicate our analyses across time and samples:

- A two-wave panel study fielded October 19–26 (wave 1; $N = 3,378$) and October 30–November 6, 2018 (wave 2; $N = 2,948$);
- A two-wave panel study fielded November 20–December 27, 2018 (wave 1; $N = 4,907$) and December 14, 2018–January 3, 2019 (wave 2; $N = 4,283$).

Respondents were selected by YouGov's matching and weighting algorithm to approximate the demographic and political attributes of the US population (*SI Appendix A*). Participants were ineligible to take part in more than one study. Both surveys in this research were approved by the institutional review boards of the University of Exeter, the University of Michigan, Princeton University, and Washington University in St. Louis. All subjects

[¶]*SI Appendix C* also explores whether and how overconfidence is related to trust in the media. We show that overconfidence is negatively associated with trust in the mainstream media, but positively associated with trust in information seen on Facebook.

gave informed consent to participate in each survey. The preanalysis plans are available at <https://osf.io/fr4k5> and <https://osf.io/r2jvb>.[#]

Measuring Discernment Ability: News Headline Rating Task. In each survey, we asked respondents to evaluate the accuracy of a number of headlines on a four-point scale ranging from “Not at all accurate” (1) to “Very accurate” (4). The articles, all of which appeared during the 2018 midterms, were published by actual mainstream and false news sources and were balanced within each group in terms of their partisan congeniality. In total, we selected four mainstream news articles that were congenial to Democrats and four that were congenial to Republicans (each split between low- and high-prominence sources) and two pro-Democrat and two pro-Republican false news articles. We define high prominence mainstream sources as those that more than 4 in 10 Americans reported recognizing in recent polling by Pew (60). False news stories were verified as false by at least one third-party fact-checking organization.^{||} To the extent possible, we chose stories that were balanced in their face validity. The complete listing of all stories tested is provided in *SI Appendix A*.

The stories were formatted exactly as they appeared in the Facebook news feed at the time the study was designed. This format replicated the decision environment faced by everyday users, who frequently assess the accuracy of news stories given only the content that appears in social media feeds.^{**} Respondents rated 12 stories provided in randomized order during the second wave of each survey.^{††}

We then calculated their measured ability to discern mainstream from false news. We did this by taking the difference in the mean perceived accuracy between true and false news headlines (i.e., mean perceived mainstream news accuracy – mean perceived false news accuracy). We used a difference score, rather than perceived accuracy of false news alone, to account for respondents who may tend to rate all news as mostly accurate (i.e., are highly credulous) or all news as mostly inaccurate (i.e., are indiscriminately skeptical). This approach has been frequently used in past studies (e.g., refs. 14 and 16).

SI Appendix, Table B1 shows descriptive statistics for the mean perceived accuracy of mainstream and false headlines, as well as the difference score in each wave. The results show that, on average, respondents did find mainstream stories to be more credible. For instance, the average rating for mainstream articles in the Oct./Nov. wave was 2.68, while it was 1.90 for false headlines. Although these differences are statistically distinguishable, the difference—our measure of discernment—is less than one point on the four-point scale (0.78 for Oct./Nov. and 0.62 for Nov./Dec.). In other words, respondents rated a mainstream headline as less than one point more accurate on our four-point accuracy scale compared to false news headlines. The ranges of values we observed for discernment are –1.5 to 2.88 in Oct./Nov. and –1.38 to 2.75 in Nov./Dec. Although our inferences regarding overconfidence are based on a relatively small number of news headlines ($k = 12$), these headlines appear to be comparable to the large set of political headlines in Pennycook et al. (61) ($k = 146$). After rescaling all outcomes to range from zero to one, the average accuracy rating for our mainstream headlines was 0.67 and 0.66 in our two surveys, and the average rating for false headlines was 0.48/0.50. These mean values are highly similar to Pennycook et al., who found an average rating of 0.63 for mainstream headlines and 0.49 for false headlines.

With our discernment measure, we then then ordered respondents and calculated their percentile. That is, each respondent was scored on a scale ranging from 1 to 100 based on their performance, where a score of 1 means

that 99% of respondents performed better and a score of 99 means that they performed better than 99% of respondents. In the Oct./Nov. survey, the 25th percentile score for discernment was 0.38, the 50th was 0.88, the 75th was 1.25, and the 99th was 2.38. Similarly, in Nov./Dec., the 25th percentile discernment score was 0.13, the 50th was 0.63, the 75th was 1.00, and the 99th was 2.25.

Accuracy of Perceptions of Relative Ability (Overconfidence). After the headline-rating task, we asked two questions in wave 2 of each survey that directly measure differences in perceived ability to detect false news compared to the public:

1. “How do you think you compare to other Americans in your general ability to recognize news that is made up? Please respond using the scale below, where 1 means you’re at the very bottom (worse than 99% of people) and 100 means you’re at the very top (better than 99% of people),”
2. “How do you think you compare to other Americans in how well you performed in this study at recognizing news that is made up? Please respond using the scale below, where 1 means you’re at the very bottom (worse than 99% of people) and 100 means you’re at the very top (better than 99% of people).”

For each question, respondents could use a slider to indicate a number between 1 and 100. These measures (“general ability”/“in this study”) are highly correlated ($r = 0.73$ in both the Oct./Nov. 2018 and Nov./Dec. 2018 surveys), so we take their average as our measure of perceived relative ability. On the resulting scale, the mean self-assessed relative ability was in the 69th percentile for both surveys (Oct./Nov., $M = 69.46$, $SD = 18.59$; Nov./Dec., $M = 69.43$, $SD = 17.8$). In both surveys, fewer than 12% of respondents placed themselves below the 50th percentile. The full distributions are shown in *SI Appendix, Fig. B1*.

We then combine these variables to compute the overconfidence measure as the difference between people’s self-reported ability and their actual performance. The result is a scale that can range from –100 to 100. We show the distribution of overconfidence in *SI Appendix B2*. (*SI Appendix, Fig. B3* shows the fairly weak relationship between self-rating and actual ability underlying the overconfidence measure.) In both surveys, 73% of respondents were at least somewhat overconfident, with an average overconfidence score of 21.76 in Oct./Nov. and 21.7 in Nov./Dec. ($SD = 30.77$ to 30.43), meaning that the average respondent placed themselves about 22 percentiles higher than their actual score warranted. About 20% of respondents in each survey rated themselves 50 or more percentiles higher than their discernment score warranted.

One potential concern is that our measure of overconfidence may be driven by differences in people’s ability to recognize one type of stories rather than how well they can differentiate between them per se. *SI Appendix, Fig. B4* therefore disaggregates these components. The figure shows that overconfident respondents perceived mainstream news as less accurate than their counterparts and, to an even greater extent, perceived false news as more accurate than their counterparts.

Outcomes and Behaviors of Interest. To answer RQ2–RQ4, we also create measures of visits to false news websites, topical misperceptions, and self-reported engagement intentions (sharing/liking). We describe our measures for each in turn.

News Exposure Data. News exposure is measured using behavioral data on respondents’ web visits collected unobtrusively with their informed consent. Data are available from users’ laptop or desktop computers. Web visits are collected anonymously with users’ permission through a mix of browser plug-ins, proxies, and VPNs. The provider of these passive metering data is the firm Reality Mine, whose technology underlies the YouGov Pulse panel from which survey respondents were sampled. Our measures of news exposure come from a period immediately following the survey. The lists we used to code each type of media are below:

- Mainstream news visit: One of AOL, ABC News, CBSNews.com, CNN.com, FiveThirtyEight, FoxNews.com, Huffington Post, MSN.com, NBCNews.com, NYTimes.com, Politico, RealClearPolitics, Talking Points Memo, The Weekly Standard, WashingtonPost.com, WSJ.com, or Wikipedia.
- False news visit: Any visit to one of the 673 domains identified as a false news producer as of September 2018 (62), excluding those with print versions (including, but not limited to, *Express*, the British tabloid) and also domains that were previously classified (63) as a source of hard news. In addition, we exclude sites that predominantly feature user-generated content (e.g., online bulletin boards) and political interest groups.

[#]Participants received an orthogonal treatment related to media literacy in both surveys. Due to a programming error, all respondents received the treatment in the Oct./Nov. survey. The results of this study are reported in ref. 14. Other orthogonal studies embedded in these surveys are reported in refs. 58 and 59.

^{||} Respondents also rated the accuracy of four hyperpartisan news headlines, which are technically factual, but present slanted facts in a deceptive manner. We do not include these articles in this analysis due to the inherent ambiguity as to whether they are truthful. These headlines were included as part of a separate study reported in Guess et al. (14).

^{**}Due to Facebook’s native formatting, the visual appearance of the false article previews differed somewhat from those of the mainstream articles—see *SI Appendix, Figs. A1 and A2*.

^{††}In each survey’s first wave, respondents were randomly assigned to evaluate one of the two stories that fall into each of the six categories (e.g., pro-Republican false news, pro-Democrat high-prominence mainstream news, etc.) for a total of six headline evaluations. In the second wave, respondents evaluated all 12 stories using the same approach. We focus only on the wave 2 measures.

Duplicate visits to webpages were not counted if they were successive (i.e., a page that was reloaded after first opening it). URLs were cleaned ofreferrer information and other parameters before deduplication. [For more details, see the processing steps described in Guess et al. (13).]

We first created a binary measure of whether respondents made one or more visits to false news sites.^{††} Our binary measure of *false news exposure* is coded as one if the respondent visited any of the domains in our list (Oct./Nov.: 7%; Nov./Dec.: 6%) and zero otherwise. We also created a binary measure of *mainstream news exposure* that is coded as one if the respondent visited any such domain in our list (Oct./Nov.: 60%; Nov./Dec.: 52%) and zero otherwise. We use the latter measure to account for the possibility that overconfident individuals may simply be more likely to be exposed to news online.

In addition to false and mainstream news exposure, we also measure the overall ideological slant of respondents' total information diet, which we divide into deciles from most liberal (decile 1) to most conservative (decile 10) using the method presented by Guess (64). We use this measure in our analysis of news exposure to control for the general ideological orientation of respondents' news diets.

Importantly, not all respondents who were part of our survey chose to provide behavioral data. Thus, our sample sizes using these data decrease, especially in the Nov./Dec. wave, in which only 22% of respondents also provided online traffic data (versus 63% in Oct./Nov.). The decline between surveys reflects the lack of available respondents who 1) participated in the YouGov Pulse panel and 2) did not participate in our earlier waves of data collection. The result is that analyses using news-exposure data have less power (we also consider pooled analyses across surveys for this reason).

Topical Misperceptions, Engagement, and Congeniality. In the Oct./Nov. survey, we included a battery of questions asking respondents about their beliefs in specific claims related to the confirmation hearings for Justice Brett Kavanaugh, which occurred shortly before the survey was fielded.^{§§} Respondents were shown two true and two false statements that they rated on a four-point accuracy scale, ranging from "Not at all accurate" to "Very accurate."

These statements were balanced in terms of partisan orientation, so one true and one false statement was congenial to Democrats and one true and one false statement was congenial to Republicans. Both the true and false statements were highly visible on social media during the hearings.^{¶¶} We measured potential engagement (liking/sharing) with false news stories during the headline-rating task. For each headline, respondents were asked to self-report their intention to like or share each article (1 = not at all likely,

4 = very likely). This question was asked only of respondents who report using Facebook.^{##} It should be noted that perceived accuracy questions appeared immediately before our sharing intent questions in the survey, which may prime accuracy concerns among respondents and thereby alter self-reported sharing behavior (65).

For both misperceptions and engagement, we analyze the data in two ways. First, we create a *difference score*. For the topical misperceptions, for instance, we subtracted the perceived accuracy of false statements from the perceived accuracy of true statements to create a measure of discernment. We calculated mean responses of intentions to like/share mainstream stories, false stories, and the difference using an identical procedure. Descriptive statistics for these measures are shown in *SI Appendix, Table B1*.

Finally, we also examine our results at the headline or statement level using only false statements/headlines so that we can test whether the relationship between overconfidence and beliefs or behavior varies by partisan congeniality. Congeniality is coded at the headline or statement level for partisans to indicate that a story or statement is consistent with the respondents' partisan leanings (e.g., a Democrat evaluating a story that is favorable to a Democrat). To determine the partisanship of respondents in the US survey, we used the standard two-question party-identification battery (which includes leaners) to classify respondents as Democrats or Republicans.

Additional Covariates. Our statistical models include a series of standard covariates, including dichotomous indicators of Democrat and Republican party affiliation (including leaners), college education, gender, non-White racial background, and dichotomous indicators of membership in age groups (30–44, 45–59, and 60+; 18–29 is the omitted category). Complete descriptions of all survey items and measures are included in *SI Appendix A*.

Our October/November 2018 respondents are 57% female, 80% White, median age 55, 37% hold a four-year college degree or higher, 49% identify as Democrats (including leaners), and 34% identify as Republicans (including leaners). Our November/December 2018 respondents are 55% female, 68% White, median age 50, 32% hold a 4-y college degree or higher, 46% identify as Democrats (including leaners), and 36% identify as Republicans (including leaners).

Data Availability. Data files and scripts necessary to replicate the results in this article have been made available at the following Open Science Framework repository (<https://osf.io/xygwvt/>) (66).

ACKNOWLEDGMENTS. This work was supported by Democracy Fund; the European Research Council under the European Union's Horizon 2020 research and innovation program (Grant 682758); the Nelson A. Rockefeller Center at Dartmouth College; Carnegie Corporation of New York; and the Weidenbaum Center on the Economy, Government, and Public Policy at Washington University in St. Louis. We thank Rick Perloff, Ye Sun, Matt Motta, Mike Wagner, and seminar participants at the University of Gothenburg for helpful comments and Sam Luks and Marissa Shih at YouGov for survey assistance. All conclusions and any errors are our own.

^{##}We only observe self-reported behavioral intentions. However, self-reported sharing intention for political news articles has been shown to correlate with aggregate observed sharing behavior on Twitter at $r = 0.44$ (56).

^{††}The distribution was highly skewed; 93 to 94% of respondents visited zero false news sites, and the distribution among nonzero respondents had a long right tail (Oct./Nov., $M = 0.43$, $SD = 3.24$, $\min = 0$, $\max = 75$; Nov./Dec., $M = 0.21$, $SD = 1.37$, $\min = 0$, $\max = 25$).

^{§§}The confirmation hearings where Kavanaugh and Christine Blasey Ford testified took place in late September 2018. The final Senate vote took place on October 6.

^{¶¶}This battery was included in both waves of the Oct./Nov. survey. We focus only on the wave 2 results as the first wave preceded our collection of the overconfidence measure, but *SI Appendix D* shows that our results replicate fully when using the wave 1 topical misperception battery.

1. M. Barthel, A. Mitchell, J. Holcomb, Many Americans believe fake news is sowing confusion. *Pew Res. Center* 15, 12 (2016).
2. W. P. Davison, The third-person effect in communication. *Publ. Opin. Q.* 47, 1–15 (1983).
3. Y. Sun, L. Shen, Z. Pan, On the behavioral component of the third-person effect. *Commun. Res.* 35, 257–278 (2008).
4. J. Kruger, D. Dunning, Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *J. Pers. Soc. Psychol.* 77, 1121 (1999).
5. D. Dunning, The Dunning-Kruger effect: On being ignorant of one's own ignorance. *Adv. Exp. Soc. Psychol.* 44, 247–296 (2011).
6. J. Ehrlinger, K. Johnson, M. Banner, D. Dunning, J. Kruger, Why the unskilled are unaware: Further explorations of (absent) self-insight among the incompetent. *Organ. Behav. Hum. Decis. Process.* 105, 98–121 (2008).
7. P. J. Ferraro, Know thyself: Competence and self-awareness. *Atl. Econ. J.* 38, 183–196 (2010).
8. P. Ortoleva, E. Snowberg, Overconfidence in political behavior. *Am. Econ. Rev.* 105, 504–35 (2015).
9. L. Sheffer, P. Loewen, Electoral confidence, overconfidence, and risky behavior: Evidence from a study with elected politicians. *Polit. Behav.* 41, 31–51 (2019).
10. R. J. Kovacs, M. Lagarde, J. Cairns, Overconfident health workers provide lower quality healthcare. *J. Econ. Psychol.* 76, 102213 (2020).

11. F. Vegetti, M. Mancosu, The impact of political sophistication and motivated reasoning on misinformation. *Polit. Commun.* 37, 678–695 (2020).
12. N. Grinberg, K. Joseph, L. Friedland, B. Swire-Thompson, D. Lazer, Fake news on Twitter during the 2016 US presidential election. *Science* 363, 374–378 (2019).
13. A. M. Guess, B. Nyhan, J. Reifler, Exposure to untrustworthy websites in the 2016 US election. *Nat. Hum. Behav.* 4, 472–480 (2020).
14. A. M. Guess et al., A digital media literacy intervention increases discernment between mainstream and false news in the United States and India. *Proc. Natl. Acad. Sci. U.S.A.* 117, 15536–15545 (2020).
15. J. Roozenbeek, S. Van Der Linden, Fake news game confers psychological resistance against online misinformation. *Palg. Commun.* 5, 65 (2019).
16. G. Pennycook, D. G. Rand, Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition* 188, 39–50 (2018).
17. B. Bago, D. G. Rand, G. Pennycook, Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines. *J. Exp. Psychol. Gen.* 149, 1608–1613 (2020).
18. G. Pennycook, D. G. Rand, Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *J. Pers.* 88, 185–200 (2018).
19. C. Martel, G. Pennycook, D. G. Rand, Reliance on emotion promotes belief in fake news. *Cognit. Res. Principles Implications* 5, 47 (2020).

20. L. Ross, D. Greene, P. House, The 'false consensus effect': An egocentric bias in social perception and attribution processes. *J. Exp. Soc. Psychol.* **13**, 279–301 (1977).
21. I. G. Anson, Partisanship, political knowledge, and the Dunning-Kruger effect. *Polit. Psychol.* **39**, 1173–1192 (2018).
22. M. Motta, T. Callaghan, S. Sylvester, Knowing less but presuming more: Dunning-Kruger effects and the endorsement of anti-vaccine policy attitudes. *Soc. Sci. Med.* **211**, 274–281 (2018).
23. E. Duflo *et al.*, In praise of moderation: Suggestions for the scope and use of pre-analysis plans for RCTs in economics (NBER Working Paper 26993, National Bureau of Economic Research, Cambridge, MA 2020).
24. S. Pazicni, C. F. Bauer, Characterizing illusions of competence in introductory chemistry students. *Chem. Educ. Res. Pract.* **15**, 24–34 (2014).
25. J. Pasek, G. Sood, J. A. Krosnick, Misinformed about the Affordable Care Act? Leveraging certainty to assess the prevalence of misperceptions. *J. Commun.* **65**, 660–673 (2015).
26. J. Li, M. W. Wagner, The value of not knowing: Partisan cue-taking and belief updating of the uninformed, the ambiguous, and the misinformed. *J. Commun.* **70**, 646–669 (2020).
27. M. H. Graham, Self-awareness of political knowledge. *Polit. Behav.* **42**, 305–326 (2020).
28. K. M. Douglas, R. M. Sutton, Right about others, wrong about ourselves? Actual and perceived self-other differences in resistance to persuasion. *Br. J. Soc. Psychol.* **43**, 585–603 (2004).
29. E. M. Hansen, K. Yakimova, M. Wallin, L. Thomsen, "Can thinking you're skeptical make you more gullible? The illusion of invulnerability and resistance to manipulation" in *The Individual and the Group: Future Challenges. Proceedings from the 7th GRASP Conference*, C. Jacobsson, M. R. Ricciardi, Eds. (University of Gothenburg, Gothenburg, Sweden, 2010), 52–64.
30. V. A. Thompson, J. A. P. Turner, G. Pennycook, Intuition, reason, and metacognition. *Cognit. Psychol.* **63**, 107–140 (2011).
31. G. Pennycook, R. M. Ross, D. J. Koehler, J. A. Fugelsang, Dunning-Kruger effects in reasoning: Theoretical implications of the failure to recognize incompetence. *Psychon. Bull. Rev.* **24**, 1774–1784 (2017).
32. N. A. Salovich, D. N. Rapp, Misinformed and unaware? Metacognition and the influence of inaccurate information. *J. Exp. Psychol. Learn. Mem. Cognit.*, <https://doi.org/10.1037/xlm0000977> (2020).
33. D. J. Flynn, B. Nyhan, J. Reifler, The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics. *Polit. Psychol.* **38**, 127–150 (2017).
34. K. A. Burson, R. P. Larrick, J. Klayman, Skilled or unskilled, but still unaware of it: How perceptions of difficulty drive miscalibration in relative comparisons. *J. Pers. Soc. Psychol.* **90**, 60 (2006).
35. J. Krueger, R. A. Mueller, Unskilled, unaware, or both? The better-than-average heuristic and statistical regression predict errors in estimates of own performance. *J. Pers. Soc. Psychol.* **82**, 180 (2002).
36. J. Krueger, D. Dunning, Unskilled and unaware—but why? A reply to Krueger and Mueller (2002). *J. Pers. Soc. Psychol.* **82**, 189–192 (2002).
37. J. Feld, J. Saueremann, A. De Grip, Estimating the relationship between skill and overconfidence. *J. Behav. Exp. Econ.* **68**, 18–24 (2017).
38. T. Schlösser, D. Dunning, K. L. Johnson, J. Krueger, How unaware are the unskilled? Empirical tests of the "signal extraction" counterexplanation for the Dunning-Kruger effect in self-evaluation of performance. *J. Econ. Psychol.* **39**, 85–100 (2013).
39. G. E. Gignac, M. Zajenkowski, The Dunning-Kruger effect is (mostly) a statistical artefact: Valid approaches to testing the hypothesis with individual differences data. *Intelligence* **80**, 101449 (2020).
40. J. E. Miller, P. D. Windschitl, T. A. Treat, A. M. Scherer, Unhealthy and unaware? Misjudging social comparative standing for health-relevant behavior. *J. Exp. Soc. Psychol.* **85**, 103873 (2019).
41. R. D. McIntosh, E. A. Fowler, T. Lyu, S. D. Sala, Wise up: Clarifying the role of metacognition in the Dunning-Kruger effect. *J. Exp. Psychol. Gen.* **148**, 1882–1897 (2019).
42. M. Niederle, L. Vesterlund, Do women shy away from competition? Do men compete too much? *J. Econ.* **122**, 1067–1101 (2007).
43. J. P. Prims, D. A. Moore, Overconfidence over the lifespan. *Judgm. Dec. Making* **12**, 29–41 (2017).
44. P. Ortoleva, E. Snowberg, Are conservatives overconfident? *Eur. J. Polit. Econ.* **40**, 333–344 (2015).
45. K. Bregu, Overconfidence and (over) trading: The effect of feedback on trading behavior. *J. Behav. Exp. Econ.* **88**, 101598 (2020).
46. J. Lambert, V. Bessière, G. N'Goala, Does expertise influence the impact of overconfidence on judgment, valuation and investment decision? *J. Econ. Psychol.* **33**, 1115–1128 (2012).
47. A. M. Parker, E. R. Stone, Identifying the effects of unjustified confidence versus overconfidence: Lessons learned from two analytic methods. *J. Behav. Decis. Making* **27**, 134–145 (2014).
48. P. Belmi, M. A. Neale, D. Reiff, R. Ulfe, The social advantage of miscalibrated individuals: The relationship between social class and overconfidence and its implications for class-based inequality. *J. Pers. Soc. Psychol.* **118**, 254–282 (2020).
49. C. Anderson, S. Brion, D. A. Moore, J. A. Kennedy, A status-enhancement account of overconfidence. *J. Pers. Soc. Psychol.* **103**, 718–735 (2012).
50. O. J. Sheldon, D. Dunning, D. R. Ames, Emotionally unskilled, unaware, and uninterested in learning more: Reactions to feedback about deficits in emotional intelligence. *J. Appl. Psychol.* **99**, 125–137 (2014).
51. J. J. Mondak, M. R. Anderson, The knowledge gap: A reexamination of gender-based differences in political knowledge. *J. Polit.* **66**, 492–512 (2004).
52. R. M. Perloff, Ego-involvement and the third person effect of televised news coverage. *Commun. Res.* **16**, 236–262 (1989).
53. J. T. Cheng *et al.*, The social transmission of overconfidence. *J. Exp. Psychol. Gen.* **150**, 157–186 (2020).
54. N. M. Brashier, E. J. Marsh, Judging truth. *Annu. Rev. Psychol.* **71**, 499–515 (2020).
55. M. D. Slater, Reinforcing spirals: The mutual influence of media selectivity and media effects and their impact on individual behavior and social identity. *Commun. Theor.* **17**, 281–303 (2007).
56. M. Mosleh, G. Pennycook, D. G. Rand, Self-reported willingness to share political news articles in online surveys correlates with actual sharing on Twitter. *PLoS One* **15**, e0228882 (2020).
57. L. Stankov, J. Lee, Overconfidence across world regions. *J. Cross Cult. Psychol.* **45**, 821–837 (2014).
58. A. M. Guess *et al.*, "Fake news" may have limited effects beyond increasing beliefs in false claims. Harvard Kennedy School Misinformation Review, 10.37016/mr-2020-002 (2020).
59. N. Berlinksi *et al.*, The effects of unsubstantiated claims of voter fraud on confidence in elections. *Journal of Experimental Political Science*. <https://cpb-us-e1.wpmucdn.com/sites.dartmouth.edu/dist/5/2293/files/2021/03/voter-fraud.pdf>. Accessed 17 May 2021.
60. A. Mitchell, J. Gottfried, J. Kiley, K. E. Matsa. Political polarization & media habits (2014). Pew Research Center. <https://www.pewresearch.org/wp-content/uploads/sites/8/2014/10/Political-Polarization-and-Media-Habits-FINAL-REPORT-7-27-15.pdf> (Accessed March 21, 2019).
61. G. Pennycook, J. Binnendyk, C. Newton, D. Rand, A practical guide to doing behavioural research on fake news and misinformation (2020). *PsyArXiv* [Preprint]. <https://psyarxiv.com/g69ha> (Accessed 21 January 2021).
62. H. Allcott, M. Gentzkow, C. Yu, Trends in the diffusion of misinformation on social media. *Res. Polit.* **6**, 2053168019848554 (2019).
63. E. Bakshy, S. Messing, L. A. Adamic, Exposure to ideologically diverse news and opinion on Facebook. *Science* **348**, 1130–1132 (2015).
64. M. A. Guess, (Almost) everything in moderation: New evidence on Americans' online media diets. *Am. J. Polit. Sci.*, <https://doi.org/10.1111/ajps.12589> (2021).
65. G. Pennycook *et al.*, Shifting attention to accuracy can reduce misinformation online (2020). *PsyArXiv* [Preprint]. <https://psyarxiv.com/3n9u8/> (Accessed 21 January 2021).
66. B. Lyons, J. Montgomery, A. M. Guess, B. J. Nyhan, J. Reifler, Data from "Overconfidence in news judgments is associated with false news susceptibility." Open Science Framework (OSF). <https://osf.io/krz7f/>. Deposited 4 May 2021.