“Fact-checking” videos reduce belief in misinformation and improve the quality of news shared on Twitter

Alexander Bor^{*1,3}, Mathias Osmundsen¹, Stig Hebbelstrup Rye Rasmussen¹, Anja Bechmann², and Michael Bang Petersen¹

¹Department of Political Science, Aarhus University

²School of Communication and Culture, Aarhus University

³Democracy Institute, Central European University

April 11, 2023

Abstract

Misinformation is widely acknowledged as an important challenge for Western democracies. Yet, surprisingly little effort has been devoted to measuring the effects of various counter-strategies. We address this void by running a field experiment analyzing the causal effects of popular fact-checking videos on both believing and sharing misinformation among a diverse sample of American Twitter users (N = 1,600). We offer three empirical contributions: first, we demonstrate that fact-checking videos improve the ability to discern between mainstream and false news stories. Second, we show that they also improve the quality of news shared on Twitter. Third, we highlight the similarities and differences in the psychological antecedents of sharing and believing misinformation. Our paper aids the effort to design interventions for preventing or mitigating its negative effects.

Keywords: False news; Misinformation; Digital literacy; Motivated reasoning; Twitter; social media

Word count: 8,211

This working paper has not been peer-reviewed yet.

Acknowledgements We are grateful to the Social Observatory for Disinformation and Social Media Analysis (SOMA) expert panel for their help in selecting educational materials, to Glenn Kessler for sharing the Washington Post’s video with us, to the audience at the Central European University for their insightful comments. This research has been supported by the Carlsberg Foundation’s grant CF18_1108 to MBP.

*Corresponding author: bora@ceu.edu

Introduction

The prevalence of false and misleading news stories is perceived as a major problem in Western democratic politics. Just recently, misinformation about COVID-19 had created unique challenges during the pandemic, and many platforms struggle to contain pro-Kremlin disinformation about the war in Ukraine (Theocharis et al. 2021; Golovchenko 2020). While false news are often created by elite actors for financial or political gain and often rely on bot networks for initial promotion, lay people are instrumental in facilitating their spread. As false news stories are often both novel and emotionally engaging, social media users share them “faster and further” than accurate news stories (Vosoughi, Roy and Aral 2018). While the overall prevalence of false news stories is low (Guess, Nagler and Tucker 2019), research shows that single stories can attract significant public attention and that false information may be particularly likely to receive circulation in situations where it may be most dangerous: at times of high social polarization (Horowitz 2001).

Recent years have witnessed an explosion in collaborative efforts to stem the tide of misinformation. Computer scientists and tech companies like Facebook are exploring ways to remove bots and “sock puppets” from their user base and to adjust their page-ranking algorithms to filter out fabricated news stories (Weedon, Nuland and Stamos 2017). Data scientists are developing machine learning techniques that rely on source and text classification, and network analysis to detect and stop misinformation before they go viral on social media (Shu, Bernard and Liu 2019), while educators and journalists are constructing training programs and other resources that teach schoolchildren to evaluate information credibility (Jones 2017). The European Union has invested millions of euros in projects like the European Digital Media Observatory (EDMO) that monitors and evaluates counter-strategies against the spread and impact of online misinformation, and the East Stratcom Task Force that helps civil society build resilience against Russian disinformation campaigns. Perhaps most prominently, media companies across the world are investing heavily in fact-checking programs (Graves 2016) that, according to communication research, can hinder the spread of misinformation (Nyhan and Reifler 2015).

The promise of these advances notwithstanding, the proliferation of misinformation poses a societal challenge. Tech companies and scientists are in an arms race against misinformation

producers who use increasingly ingenious and sophisticated AI-powered tools to disseminate misinformation. Constraints on time and resources put pressure on professional fact-checkers to keep up with the latest false news stories before they hit social media. People who have been exposed to misinformation rarely visit fact-checking sites afterwards (Guess, Nyhan and Reifler 2020). And even when fact-checking reaches its intended audience, its effectiveness is disputed (for a discussion, see Nyhan et al. 2020). Thus, a recent meta-analysis of 30 studies found that while fact-checking delivers some protection against misinformation, this effect is muted by prior political beliefs and knowledge. Worse still, the effects “gradually become negligible the more the study design resembles a real-world scenario of exposure to fact-checking” (Walter et al. 2020, 357; see Aslett et al. (2022) for a recent example of the limited impact of an ecologically valid news credibility intervention).

For these reasons, social scientists increasingly emphasise the importance of *prebunking* to complement *debunking* strategies (Roozenbeek et al. 2022; van der Linden 2019; Guess et al. 2020). These approaches come in various flavors, ranging from inoculation against misinformation by playing a “Bad News Game” (Roozenbeek, Traberger and van der Linden 2022), over brief reminders encouraging people to make “accurate” sharing decisions (Pennycook, McPhetres, Zhang, Lu and Rand 2020), to so-called “metacognitive prompts” that make people aware of their own propensity to be affected by false information (Salovich and Rapp 2021). While the effectiveness of these approaches is still debated (e.g., Roozenbeek, Freeman and van der Linden 2021; Pennycook and Rand 2021), they all share a common goal of equipping citizens with the confidence and a set of broadly applicable tools to recognize misinformation in order to *prevent* belief and sharing. Moreover, the interventions are, in principle at least, scalable, and could be cost-effectively administered to large numbers of citizens.

In this article, we examine one particular prebunking strategy that enjoys widespread popularity among media organizations, government agencies and researchers: Minutes-long entertaining educational videos (Moyer-Gusé 2008) that seek to turn citizens into fact-checkers by teaching the core principles of spotting misinformation. These literacy campaigns encourage reading beyond the headline, googling the headline or main claims, and verifying the credibility of the source and the author. Multiple stakeholders, like governmental and non-governmental institutions, and major news sites have created their own versions of such video-based programs,

and they have been watched by millions (see below). Educational videos and other teaching materials are also at the heart of the EU’s strategy to raise awareness about misinformation in civil society and to enable citizens to navigate the online news environment safely (EC 2018).

Video-based information is often more persuasive than text-based messages (Wittenberg et al. 2021). Yet, despite high hopes for the role they can play in fighting misinformation, the educational interventions lack rigorous program evaluation: We simply do not know how effective they are (although, see Rasmussen, Lindekilde and Petersen 2022). In the present manuscript, we contribute by simultaneously estimating the effectiveness of educational video interventions in improving people’s ability to discern false and mainstream news, *and* in decreasing people’s actual sharing of misinformation on a social media platform, Twitter. In achieving these ends, the article also introduces a novel measure of news quality based on state-of-the-art machine learning algorithms, and offers a nuanced comparison of the psychological foundations of believing in and behavioral sharing of misinformation.

The psychology of believing and sharing misinformation

Empowering citizens to avoid falling for misinformation necessitates an understanding of the psychological factors shaping the decision to share false news stories. Convincing someone to stop sharing misinformation would require different tools depending on whether they believe these news and share them in good faith, whether they find them funny and post them to entertain their (potentially unassuming) followers, or whether they share to smear dirt on their partisan opponents. Educational fact-checking interventions assume (often implicitly) that people are motivated by a particular set of motivations, which social psychologists call *accuracy motivations* (Pennycook, Bear, Collins and Rand 2020). Accuracy motivations imply that people are motivated to hold accurate beliefs about the world (Chaiken and Trope 1999). In the context of social media, these accuracy motivations imply that the average user wants to share accurate news but may not be motivated or competent enough to assess the veracity of news (Pennycook, Bear, Collins and Rand 2020). From this perspective, educational interventions work because they make citizens better able to determine whether news stories are true or false.

In line with these arguments, we know that individuals who invest more cognitive resources in problem-solving (Pennycook and Rand 2019b) and who enjoy analytical thinking (Pennycook

and Rand 2020) are less likely to fall prey to misinformation. Elderly social media users – arguably due to shortfalls in digital media literacy and a failure to discern false from true online – are among the worst offenders when it comes to sharing misinformation on multiple social media platforms (Guess, Nagler and Tucker 2019; Osmundsen et al. 2021). Guess et al. (2020) recently demonstrated that exposure to Facebook’s brief “Tips to Spot False News,” improved the ability to discern between false and mainstream news story headlines in both US American and Indian samples.

That said, caution is warranted when it comes to the generalizability of the positive effects of educational interventions. Guess and colleagues 2020 found positive effects only in two samples already familiar with the internet. In their final, low-literacy sample, no treatment effect was found. Similarly, Badrinathan (2021) found that a multi-hour, face-to-face digital literacy intervention did not improve truth discernment among a sample of Indians who were recruited from one of the most illiterate regions in the country. In fact, the intervention has backfired among supporters of the government.

To the best of our knowledge, no prior study offers direct evidence that the inability to discern false from true news contributes to misinformation *sharing*. Hence, the key psychological premise underlying fact-checking and related interventions remains untested. Buchanan (2020) demonstrates that people report lower *intention* to share news they do not believe, but Effron and Raj (2020) show that people think it is fine to share false news, if many others have done it already. There is also increasing evidence that when people process politically relevant information, they are less concerned about the veracity of news and more about their usefulness for derogating disliked groups (Mercier 2020; Acerbi 2020). For example, people tend to believe more in conspiracy theories, if they fit their ideological or partisan dispositions, and these effects are larger among those most knowledgeable about politics (Miller, Saunders and Farhart 2016). Furthermore, people are motivated to share news that portray disliked groups negatively, in some cases despite knowing that the news are false (Petersen, Osmundsen and Arceneaux 2022). Indeed, the very fact that false news are freed from the constraints of reality may enable them to contain a multitude of cues appealing to such partisan (and more broadly social) considerations (Acerbi 2019). If people are undisturbed by the mendacity of the news they share online, efforts to increase their ability or motivation to establish veracity may be

ineffective. In fact, Guess et al. (2020) show that these concerns are valid: While their media literacy intervention improved respondent’s ability to detect false news, it did not reduce their likelihood of consuming or intention to share misinformation.

On these grounds, our aim is to conduct a direct behavioral test of whether prominent educational fact-checking videos reduce false news sharing on Twitter by improving the ability to discern between information and misinformation. In doing so, we make two main contributions. First, we test the effects of educational fact-checking videos on *both* discernment ability and news sharing under real-world conditions. This is important as the bulk of previous studies limit their inquiry to truth discernment, and rarely explore whether they also propel social media users to actually *share* better quality news on social media. As a result, we know little about how helpful these interventions are in limiting the flow of misinformation online. Second, we explore the relationship between truth discernment and news sharing, and assess the psychological antecedents of both behaviors. In combination, these analyses illuminate the psychological mechanisms underlying the potential of fact-checking interventions.

Data and Methods

Our analyses are based on a pilot study on Amazon’s Mechanical Turk ($N = 686$) and a pre-registered field experiment in a highly naturalistic environment, using a unique panel of Twitter users ($N = 1,600$). In both studies, respondents in the treatment arm were randomly assigned to watch one of six professionally produced fact-checking videos. Further, by relying on these specific videos, we make sure that the experimental materials have maximal ecological validity, while also testing the effectiveness of the materials actually available to social media users at scale. These short (<10 min) fact-checking videos highlighted widely acknowledged anti-misinformation practices like the importance of considering source credibility and searching the Internet for supporting evidence. We examined in both the MTurk and the Twitter samples whether respondents who watched the videos were better at assessing the veracity of news headlines in a news quiz than participants who did not watch the videos. Relying on the Twitter sample, we examined the real-world implications of the intervention by testing whether the educational videos improved the quality of news shared on Twitter.

We pre-registered our hypotheses and analyses both with regards to the causal effects of the

fact-checking videos on truth discernment and news sharing, and for the relationship between truth discernment and news sharing. However, as we describe below, our analyses on news sharing deviate from this pre-registration. Our pre-registrations are posted at the end of the online appendix and they are also available at https://aspredicted.org/EEH_YRT and https://aspredicted.org/BHW_DXM.

Sampling and participants

A diverse but non-representative sample of US Americans were recruited for our pilot via MTurk ($N = 686$), mean (SD) age = 40.5 (12.4), 57% male. As usual with MTurk samples, they are more educated than the national population (67% completed higher-education) and skew towards the left (54% self-identified Democrats and leaners, 10% Independents, 35% Republicans and leaners). Participants were reimbursed with a pay commensurate with the national hourly minimum wage.

Our participants in the field experiment are Twitter users recruited by YouGov survey agency to be part of a multi-wave panel survey. We chose to run our study on Twitter because it was the largest and most relevant social network (in the Western world), which provided researchers sufficient access to its data to implement a field experiment on false news sharing. Admittedly, the overwhelming focus on Twitter is a weakness of the misinformation literature, but our ambition here was not to broaden the scope of our knowledge by looking at other platforms but to deepen our understanding of psychological procedures that have been studied previously. As such, we consider it a strength of our study that its findings speak directly to previous works in the literature.

We fielded the first wave of the panel survey between December 2018 and January 2019, where participants gave informed consent to scraping their Twitter feeds based on the Twitter usernames they disclosed and to linking it to their survey responses. We embedded the present experiment in Wave 2 of the panel, which was fielded between November 2019 and December 19, 2019. 1,870 people participated in our experiment, but 270 respondents gave usernames in Wave 1 that we could not match to a Twitter account, which prevented us from scraping their feeds and analysing their behavior. Our final sample size is therefore $N = 1,600$ (799 in the control, 801 in the treatment group). Our experimental groups appear well balanced across

demographic variables. The sample resembles the Twitter population in terms of partisanship, income and gender but our respondents are older, whiter and more educated than the average user (see OA Section B.1 on detailed demographic information).

Materials

Our survey first asked respondents a series of political questions, then randomly assigned them to either a treatment condition where they watched an educational video on misinformation or to a control condition where they did not watch any videos. Following a number of filler questions, all respondents participated in a news quiz where they rated the perceived accuracy of eight news headlines.

Our analyses on the psychological foundations of believing and sharing misinformation rely on a comprehensive set of self-reported survey instruments. With the exception of digital literacy, these were measured in the first wave of the panel, 10-12 months before treatment and the false news quiz. This greatly reduces concerns about common source bias. We measured a number of individual differences highlighted in prior research on misinformation and problematic online behavior. Specifically, we measure factual political knowledge and partisan identity using standard instruments from the American National Election Study (ANES). We measure political cynicism using a scale from Dekker and Meijerink (2013). We measure cognitive reflection with the CRT2 scale (Thomson and Oppenheimer 2016). We measure disruptive motivations with the Global Assessment of Internet Trolling scale from Buckels, Trapnell and Paulhus (2014). Finally, in the present survey, we also included a validated measure of self-reported digital literacy from Hargittai and Hsieh (2012).

Educational videos

To avoid building our analysis on a single stimulus, we sought to find the best short educational videos publicly available and estimate their average effect. Beyond searching the internet ourselves, we solicited candidate videos through the Social Observatory for Disinformation and Social Media Analysis (SOMA). SOMA was an international, interdisciplinary project founded by the European Commission to fight false news with a diverse expert panel. Based on the recommendations of the experts and our own results, we compiled a list of 18 videos that ful-

filled our criteria: 1) professionally produced, 2) in English, 3) under 10 minutes long, and 4) aimed at an adult audience. We chose the six best videos based on our assessment of production quality and the specificity of recommendations for spotting misinformation.

Table 1: The six educational videos employed in our experiment

Title	Author	Link	Views
How to choose your news – Damon Brown	Ted ED	youtu.be/q-Y-z6HmRgI	1100K
Five ways to spot fake news	Quartz	youtu.be/y7eCB2F89K8	150K
How to Spot Fake News - FactCheck.org	Factcheck.org	youtu.be/AkwWcHekMdo	550K
The Fact Checker’s guide for detecting fake news	Washington Post	youtu.be/SoWCDJAMk2Y	26K
Helping Students Identify Fake News with the Five C’s of Critical Consuming	John Spencer	youtu.be/xf8mjbVRqao	363K
Four ways to tell if something is true online - Break the Fake	MediaSmarts	youtu.be/E-049KTrYBg	10K

Note: View counts as of 2022. December 6.

Table 1 describes the details of the six videos, which are produced by respectable governmental and non-governmental organizations, educators and media companies. The six videos combined have been viewed over two million times only on YouTube.¹

All six videos recommend simple techniques for establishing the veracity of online news stories such as reading beyond the headline, considering whether the source of the information is likely to have sinister or benevolent motives, double-checking website URLs, and cross-checking with other sources and experts. Section A.1 in the Online Appendix (OA) details specific recommendations for each educational video. These recommendations can be situated broadly in the rich literature exploring how certain types of *literacy* can help media users critically assess information and inoculate against the harmful effects of misinformation. Thus, the recommendations can – to varying degrees – be seen as ways for improving both knowledge about the inner workings of the news media press (i.e., *procedural news literacy*, see Amazeen and Bucy 2019), the ability to adapt to new technologies and search the online environment for useful

¹The Washington Post’s video has been distributed through the newspaper’s website. We are grateful to Glenn Kessler at the Washington Post for uploading their video to YouTube to aid our research.

information (i.e., *digital literacy*, see Hargittai and Hsieh 2012), and broadly applicable skills for understanding, finding, evaluating, and using information as a tool for acting with reasonable accuracy in the world (i.e., *information literacy*, see Jones-Jang, Mortensen and Liu 2019).

False news quiz

We built on previous research by Pennycook and colleagues 2019*b*; 2020 to develop our false news quiz, which consists of 8 items (4 false and 4 mainstream; 4 pro-Republican and 4 pro-Democratic, fully crossed) presented in a random order. Each item consists of a news headline, a lead (one-sentence summary of the news) and a picture, mimicking the display of news links on social media. The items are displayed in OA Section A.2 along with additional information. The eight items were selected from a larger pool in a pilot study on Lucid ($N = 997$) using item response theory models to minimize measurement error (for details see Section C in the OA).

Scraping Twitter and assessing news quality

Our respondents’ tweets were scraped weekly starting from March 2019 (i.e. eight months prior the present experiment was fielded) using Twitter’s API’s “user timeline” endpoint. Continuous scraping ensures that tweets that are removed days or weeks after posting (whether by the user or the platform) are less likely to be missing from our data. We identify the study period as the 52 weeks pre-treatment and 5 weeks post-treatment for each user, covering the time between November 2018 and January 2020. We extracted URLs from tweets posted within this time-frame, using the “entities urls” section of each tweet. As many URLs were shortened, we expanded URLs with Python’s `urlexpander` package. We tested the robustness of the the expansion of URLs using Python’s `urllib` and `http` modules. Once all URLs were expanded, we extracted the domain from each URL also using the `urlexpander` package. We match these URLs to a list of 903 news domains containing 251 mainstream news outlets (identified by AllSides.com’s Media Bias database), and 652 media outlets with a history of sharing misinformation (identified by Guess et al. (2019), and Grinberg et al. (2019)). Thereby, we identify 57,231 newlinks in our data (52,023 from mainstream news sites, 5,208 from false news sites).

While our pre-registered analysis strategy specified false news counts as our primary outcome measure, these analyses proved inconclusive due to the properties of this measure (see below, and

in OA Section E). Accordingly, we rely on two measures with better measurement properties, both tapping the average quality of shared news. First, we use professional fact-checkers’ assessment of the trustworthiness of 60 news domains (39 of which are found in our data) (Pennycook and Rand 2019a). This approach assigns a quality score to 41% of news links shared by our participants on Twitter (23K newslinks). Second, to estimate news quality for all news domains, we employed a super-unsupervised approach Rasmussen et al. (2021). Our strategy builds on recent evidence demonstrating that laypeople are highly accurate in evaluating the quality of news sources (Pennycook and Rand 2019a). However, instead of inviting people to rate news domains on perceived trustworthiness, here, we explore patterns in users’ actual news sharing behavior on Twitter. Specifically, we take a large corpus of US Twitter data (730M links from 890K users) and group domains by who (which user) and when (which month) shared them. We train word vectors using the Word2Vec algorithm on this corpus, mapping how likely domains are to be shared together with other domains. Using these word vectors, we calculate how “close” each news domain is to the most trustworthy news sites according to professional fact-checkers i.e. to the New York Times and the Washington Post see (for more details, see Rasmussen et al. 2021).

We predicted and preregistered that this proximity score will be correlated with expert ratings of news quality both at the domain level, and at the user level https://aspredicted.org/H23_1TM. As predicted, we find a very strong evidence for convergent validity at the domain level Spearman’s $\rho(56) = 0.85$ between expert ratings and super-unsupervised news quality scores. Yet, this test treats all domains equal, although arguably, we care more about domains which are shared more often by respondents. Accordingly, as a second test, we calculated an average news quality score for respondents who shared at least one of the expert-rated newssites ($N = 586$). Again, we find that expert-ratings and super unsupervised scores are highly similar (Spearman’s $\rho(586) = 0.91$, see plots in OA Section F.1).

The benefit of super unsupervised scores is that they are available for all 264 news domains, not only the 39 with expert ratings, yielding considerable increases in the number of newslinks (23K vs 57K) and users (586 vs 723) analyzed. In the followings, we report estimates using both the expert ratings and our novel super-unsupervised news quality scores. We transform both measures into z-scores with higher values indicating more trustworthy news sites.

Analysis strategy

Truth discernment is modeled in a linear multilevel model with a dummy for treatment group and varying slopes for the six videos. This model estimates the sample average treatment effect of the intervention, while taking into account that this effect is conditional on the specific stimulus. Given that we have no information whether participants actually received the treatment, all our causal estimates are intend-to-treat effects.

Our original, pre-registered analysis strategy for the sharing of misinformation relied on estimating effects on the average number of misinformation stories shared by respondents. After accessing the data, we realised this strategy was sub-optimal for two reasons. First, links to verified misinformation domains are rarely shared. Even the most comprehensive list of misinformation domains identifies merely 5,788 misinformation stories shared by 256 respondents during the study period. The scarcity of misinformation severely limits the statistical power of our analyses of the intervention’s effects. Indeed, post-hoc power simulations (reported in OA Section D) indicate that unless the treatment eliminated 40% of all false news shared our sample, our models would be always consistent with a null effect. Second, as identified by prior research (Guess, Nagler and Tucker 2019; Osmundsen et al. 2021) misinformation sharing is not only rare but highly skewed with a small minority (“supersharers”) sharing a large majority of false news. This creates a further complication, because our original strategy weighted all respondents equally, irrespective of the number of news they shared. Consequently, even if our intervention successfully reduced the number of misinformation stories shared by supersharers, it would have been washed away by the sea of respondents who shared few (if any) news stories pre- and post-intervention alike. To remedy this complication, we use news links as our unit of analysis, thereby implicitly assigning equal weight to each newslink in our network. While this is an important step in the right direction compared to user-focused analyses, it remains a conservative estimate. Our data shows that users who post more newslinks also have more followers (Spearman’s $\rho=0.43$), meaning that the links of supersharers are likely to generate more views, and could be justifiably assigned more weight. Given that we lack good data on tweet-level impressions generated, we do not pursue this idea further. We report our original, pre-registered analyses and evidence underpinning our decision to abandon them in OA Sections D and E.

We proceed with exploratory analyses using expert ratings and super-unsupervised scores as measures of news quality at the link level. The effect of our intervention is modelled in difference-in-differences models with standard errors clustered at the user-level. Specifically, we regress news quality scores on a dummy for the treatment group, a dummy that switches on the day after the respondent participated in our study, and an interaction term between the two dummy variables. This interaction term can be interpreted as the change in average news quality in the treatment group from what would have been expected given the pre-treatment difference in news quality between the groups, and the time trend in the control group. Values beyond the confidence interval around this point estimate are values unlikely given our data and model. We interpret these as the lower and upper bounds of our estimates.

Finally, for our inquiry on the relationship between believing and sharing misinformation, we use data from the 713 respondents for whom we have data on both variables. We use only news links shared in the year *prior to* exposure to the treatment and calculate an average news quality score for each respondent. We use the super-unsupervised scores, as they offer more fine-grained measures on more respondents than the expert-ratings. We regress the quality of news shared on truth discernment scores using OLS regressions, 1) without covariates, 2) with demographic covariates (age, gender, and education), or 3) with psychological covariates (partisan identity, cognitive reflection, trolling, political cynicism, and knowledge). We also report simple Spearman’s rank-order correlations between each of our psychological measures, truth discernment and quality of news shared.

Results

Educational videos improve truth discernment

First, we find that the educational videos work as intended; they increase the ability to differentiate between false and mainstream news. Compared to participants in the control group, Twitter: $M(SD) = 0.37 (0.24)$; MTurk: $M(SD) = 0.30 (0.23)$, participants who were exposed to the educational videos, Twitter : $M(SD) = 0.42 (0.25)$; MTurk: $M(SD) = 0.38 (0.25)$ had significantly higher truth discernment scores, Twitter: $\Delta M = 0.05$, 95% CI = $(0.02; 0.08)$, $p < 0.01$, Cohen’s $d = 0.2$; MTurk: $\Delta M = 0.07$, 95% CI = $(0.02; 0.13)$, $p < 0.05$, Cohen’s $d = 0.3$. Although these effects may appear modest, it is worth noting that

the average difference between participants with and without higher education in our Twitter sample is not much larger ($\Delta M = 0.08$, 95% CI = (0.06; 0.11), $t(991) = 6.8$, $p < 0.001$, Cohen’s $d = 0.36$).

Figure 1: Educational videos improve truth discernment ability both among Twitter (blue) and MTurk users (red).

Note: The figure shows coefficient estimates from multilevel models regressing discernment on a dummy on treatment with random slopes for videos. Errorbars denote 90% and 95% confidence intervals.

Figure 1 plots these pooled estimates along with estimates for each of our six videos. It demonstrates that there is little systematic variation across the videos. While all of them have similar, statistically significant effects (with a single exception), some of them had above average effects in the MTurk sample, others in the Twitter sample. John Spencer and MediaSmarts produced the two videos with slightly above average effects in both samples.

As a post-hoc analysis, we disaggregate discernment scores into a measure of “gullibility” – the average perceived accuracy of false news – and “trust in mainstream media” – the average perceived accuracy of mainstream news. Regressing our treatments on these measures shows that the educational materials have the intended effects on both attributes, although the decrease in gullibility is somewhat larger (at 0.04 points) than the increase in mainstream media trust (between 0.01 and 0.03 points). Detailed results are reported in Section F.2

Educational videos improve the quality of news shared

Next, we turn to the effect of educational videos on news sharing on Twitter. Our main model, (displayed in Panel A of Figure 2) relies on 3 months of pre-treatment and a 1 week post-treatment period. Both of our measures indicate that our intervention improved the average news quality (expert ratings: $\beta = 0.27$, 95%CI [0.03, 0.51]; super-unsupervised scores: $\beta = 0.19$, 95%CI [-0.03, 0.40]) albeit both negligible and large effects are consistent with the data and the model (i.e., the 95% confidence intervals are close to or cross 0). To get a sense of the substantive size of this effect, we note that it is comparable to the difference between the average news quality posted by users who completed higher education versus users who did not (expert ratings: $\beta = 0.28$, 95%CI [0.10, 0.46]; super unsupervised scores: $\beta = 0.23$, 95%CI [0.07, 0.39]).

Zooming in on three months pre-treatment, and one week post-treatment is an arbitrary decision. Our post-hoc analyses show that generally, longer pre-treatment periods yield more precise estimates. Meanwhile, the selection of post-treatment periods is less straightforward, because although more weeks means more newlinks, any effect from the treatment may wane in a few weeks time. Accordingly, panel B of Figure 2 we report the estimated treatment effects for several combinations of pre-treatment (facet columns) and post-treatment (x-axis) periods. The estimates are relatively robust to the length of pre-treatment, consistently showing small positive effects that hover at around the threshold of statistical significance. The figure also demonstrates that any effect from the treatment gradually wanes and approach zero by 5 weeks of post-treatment data.

What mechanism underlies this improvement the average news quality? Do people stop sharing false news, or do they share more mainstream news? We estimate the intervention's effect conditional on the average news quality of users in the pre-treatment period using super unsupervised scores. We find that the intervention is very effective amongst users sharing above average quality news ($\beta = 0.17$, 95%CI [0.07, 0.28]). Conversely, amongst users sharing mostly low quality news, the intervention is much less effective: the point estimate is negative, the confidence interval is wide, and includes zero ($\beta = -0.12$, 95%CI [-0.38, 0.15]). Testing the formal significance of this subgroup-difference with a three-way interaction returns a just significant result ($\beta = 0.29$, 95%CI [0.003, 0.57]). However, while the estimated treatment

effect among users sharing better (versus lower) quality news on average is consistently higher, confidence intervals are very wide and sometimes include 0 once we explore alternative model specifications.

A related question is if the intervention changes the news shared by the users, or merely changes the composition users sharing (more) news. Put differently, do users follow the instructions and start to post better quality news than they usually do, or do they simply adjust the frequency of posting depending on whether they were encouraged or discouraged by the videos? Adding respondent fixed effects to our models offers some insights on this dilemma. First, it limits our inquiry to users who posted newlinks both pre- and post-treatment. Second, it de-means quality scores at the user level, comparing each user’s news quality to what they post on average (given the specific window). These models zoom in on the change in news quality within respondent, and ignore changes in average news quality due to some users posting more or less. Figure OA7 in the OA reports treatment effects across all pre- and post-treatment periods. They indicate that around half of the change in average news quality could be attributed to respondents posting higher quality news after they were exposed to fact checking videos. That said, given the smaller sample size, even our most optimistic models (with long pre-treatment and short post-treatment periods) remain consistent with both null and medium sized effects. For example, using the same window as above (3 months pre-, 1 week post-treatment) the estimated effect for expert ratings is $\beta = 0.15$, 95%CI [-0.01, 0.31]; and for super unsupervised scores is $\beta = 0.09$, 95%CI [-0.01, 0.18]).

Similarities and differences between sharing and believing misinformation

Our results so far show that watching a fact checking video increases truth discernment and likely improves the quality of news shared on Twitter. How important is truth discernment ability for sharing high(er) quality news on Twitter? First, we regress the quality of news shared (using the super-unsupervised scores) on truth discernment scores. We find a moderate relationship: respondents with a standard deviation higher truth discernment share 22% of a standard deviation higher quality news. This difference is similar if we adjust for demographic covariates, but halved by adjusting for all psychological covariates simultaneously. Additional analyses, reported in Section F.4 in the OA indicate that respondents better at discernment

Figure 2: Treatment effect on average news source quality by various pre-treatment periods (vertical panels) and post-treatment periods (x axis). All coefficients are based on the interaction terms in a difference-in-differences model

Note: Error bars denote 95% confidence intervals.

share somewhat less false news, but much more mainstream news. That said, discernment ability alone explains but 0.04 of the variance in the quality of news shared (i.e., Pearson's $r = 0.2$). This leaves considerable room for unique psychological profiles for sharing and believing.

Accordingly, next, we turn to the psychological predictors of sharing and believing false news. Figure 3 offers simple descriptive comparisons between factors known to shape truth discernment (digital literacy and cognitive reflection) and factors known to shape false news sharing (age and Republican identity) as well as additional variables.² When it comes to truth discernment (in blue), factual political knowledge is by far the strongest predictor. This is hardly surprising as familiarity with the political system should directly translate to evaluating the veracity of news. People who have completed higher education, and who self-report higher digital literacy are also better able to distinguish misinformation from mainstream news. We replicate the finding that cognitive reflection correlates with truth discernment, albeit less strongly than the three factors mentioned above. Two variables appear to predict worse discernment: a internet trolling and

²As this is a post-hoc analysis, and we analyze the psychological antecedents of false news sharing with partially overlapping data reported elsewhere (Osmundsen et al. 2021), we limit our analysis to simple descriptive comparisons here.

Republican partisan identity (using standard 7-point scale). Finally, political cynicism and age does not appear to make a substantial difference in truth discernment.

Figure 3: Believing and sharing misinformation – their relationship (A) and their correlates (B)

Turning now to the predictors of misinformation sharing (in red), Republican identity is the strongest predictor, and indeed the only variable in our data with a strong correlation with both sharing and believing misinformation. Political knowledge and higher education are much weaker, albeit significant predictors of the quality of news shared. Meanwhile, political cynicism does not correlate with sharing either. The remaining variables underpin the distinct psychological profiles of those who believe and share low quality information. Age did not matter for discernment, but correlates with poorer quality news shared (replicating Guess, Nagler and Tucker 2019). Finally, digital literacy, trolling, (and to a lesser extent) cognitive reflection were important predictors of discernment, yet they fail to distinguish between those who share higher versus lower quality news on Twitter.

Discussion and Conclusions

Stemming the tide of false news propagated on social media is a daunting task. Solving it requires extensive, interdisciplinary efforts. Here, we have taken the first stab at evaluating

one counter-strategy that media organizations and researchers have enthusiastically embraced: educational videos aimed at empowering citizens to become their own fact-checkers.

As anticipated by past work arguing that the failure to think analytically drives beliefs in misinformation (Pennycook and Rand 2020), we found that the educational materials on how to spot misinformation improved participants’ ability to assess the credibility of news story headlines. These findings suggested that the fact-checking videos helped participants build immediate resistance against misinformation. This is an important and desirable finding: Fact-checking videos may serve as an effective intervention to reduce belief in misinformation. Our article thus contributes to a budding literature on how various illiteracy interventions reduce belief in misinformation (Guess et al. 2020; Badrinathan 2021).

Are fact-checking videos an effective tool for increasing truth discernment? The smaller effect we identify (Cohen’s $d=0.2$, among Twitter users) is larger than those found by most prior studies and indeed, is equal to the most successful intervention – to the best of our knowledge – in the literature by Guess and colleagues Guess et al. (2020). The effect we identify among MTurkers in our pilot study surpassed this by 50% (Cohen’s $d=0.3$). That said, our intervention is also somewhat demanding by internet-standards. While millions of people are willing to watch educational videos, they undoubtedly constitute a minority of internet users.

Our results also indicate that fact-checking videos improve the average quality of news shared on Twitter in the first few weeks after exposure. Our results show tentative evidence that the intervention improves the quality of the news shared by users, and does not merely discourage posting among those who fall for false news. We also see that the improvements are larger among those who post better quality news on average. This has bittersweet implications: on the one hand, it means that users most likely to voluntarily watch a video about fact-checking (more sophisticated, more interested) are actually those who benefit from them the most (Guess, Nyhan and Reifler 2020). On the other hand, it also means these videos may (further) reduce the weight of false news on social media, but are unlikely to eliminate them. That said, these effects are imprecisely estimated and hover at the border of statistical significance.

Our study is subject to a number of limitations. First, we significantly deviate from our pre-registration when it comes to estimating the effects of the intervention on news sharing on Twitter, and much of our results should be considered as exploratory. Our experience should be

an important methodological reminder for future attempts at implementing field experiments on Twitter intending to study news sharing behavior that even samples conventionally considered large (1,600 respondents, 57,231 news links spanning months pre- and post-treatment) could yield noisy estimates once the difficulties of merging survey and Twitter data, and the rarity of misinformation sharing are factored in. At the same time, our paper also introduces an important methodological innovation – estimating news quality using super-unsupervised scores, which can augment previous methods relying on a very limited subset of expert rates news sources. Future studies should either employ even larger sample sizes than we did, or should target users with a history of sharing misinformation (see Munger 2017; Siegel and Badaan 2020).

Another limitation of the present study is that while it provides important proof-of-concept for the beneficial effects of fact-checking videos, our data is too noisy to identify more or less effective videos. While the industry of designing educational materials to reduce beliefs and sharing of misinformation is still in its infancy, as the field matures, more attention should be spent to find out which interventions work for which audiences. This is all the more important as – despite our encouraging results – the effectiveness of these interventions cannot be taken for granted (Badrinathan 2021).

Finally, our study highlights the lack of a nuanced understanding of the psychological mechanisms through which fact-checking videos improve news sharing on Twitter. In light of the modest correlation between truth discernment and news sharing quality, the impact of our intervention is too large to be entirely driven by improvements in truth discernment. It is plausible that fact-checking videos not only offer heuristics to discern false and mainstream news, but also reinforce the norms of accuracy. Our data implies that plenty of respondents who share low quality news are digitally literate, politically knowledgeable individuals. For these people, interventions which emphasise norms, or highlight the reputational costs of low quality news sharing could be more effective (Altay, Hacquin and Mercier 2022). Sharers of low quality news on social media are mainly identifiable by their political allegiance. One possibility is thus that effective interventions need to help sharers realize that the social and personal benefits of giving in to political biases are offset by other countervailing social costs.

References

- Acerbi, Alberto. 2019. “Cognitive attraction and online misinformation.” *Palgrave Communications* 5(1):1–7.
- Acerbi, Alberto. 2020. *Cultural evolution in the digital age*. Oxford University Press.
- Altay, Sacha, Anne-Sophie Hacquin and Hugo Mercier. 2022. “Why do so few people share fake news? It hurts their reputation.” *New Media & Society* 24(6):1303–1324.
- Amazeen, Michelle A and Erik P Bucy. 2019. “Conferring resistance to digital disinformation: The inoculating influence of procedural news knowledge.” *Journal of Broadcasting & Electronic Media* 63(3):415–432.
- Aslett, Kevin, Andrew M. Guess, Richard Bonneau, Jonathan Nagler and Joshua A. Tucker. 2022. “News credibility labels have limited average effects on news diet quality and fail to reduce misperceptions.” *Science Advances* 8(18).
- Badrinathan, Sumitra. 2021. “Educative interventions to combat misinformation: Evidence from a field experiment in India.” *American Political Science Review* 115(4):1325–1341.
- Buchanan, Tom. 2020. “Why do people spread false information online? The effects of message and viewer characteristics on self-reported likelihood of sharing social media disinformation.” *Plos one* 15(10):e0239666.
- Buckels, Erin E, Paul D Trapnell and Delroy L Paulhus. 2014. “Trolls just want to have fun.” *Personality and Individual Differences* 67:97–102.
- Chaiken, Shelly and Yaacov Trope. 1999. *Dual-process theories in social psychology*. Guilford Press.
- Dekker, Henk and Frits Meijerink. 2013. “Political cynicism: Conceptualization, operationalization, and explanation.” *PCS–Politics, Culture and Socialization* 3(1+ 2):9–10.
- EC. 2018. “Joint Communication to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions: Action Plan against Disinformation (JOIN(2018) 36 final).”.

URL: <https://ec.europa.eu/digital-single-market/en/news/action-plan-against-disinformation>

- Effron, Daniel A and Medha Raj. 2020. "Misinformation and morality: encountering fake-news headlines makes them seem less unethical to publish and share." *Psychological Science* 31(1):75–87.
- Golovchenko, Yevgeniy. 2020. "Measuring the scope of pro-Kremlin disinformation on Twitter." *Humanities and Social Sciences Communications* 7(1):1–11.
- Graves, Lucas. 2016. *Deciding what's true: The rise of political fact-checking in American journalism*. Columbia University Press.
- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson and David Lazer. 2019. "Fake news on Twitter during the 2016 U.S. presidential election." *Science* 363(6425):374–378.
- Guess, Andrew, Benjamin Lyons, Jacob M Montgomery, Brendan Nyhan and Jason Reifler. 2019. "Fake news, Facebook ads, and misperceptions." *Democracy Fund Report* .
- Guess, Andrew, Jonathan Nagler and Joshua Tucker. 2019. "Less than you think: Prevalence and predictors of fake news dissemination on Facebook." *Science advances* 5(1):eaau4586.
- Guess, Andrew M, Brendan Nyhan and Jason Reifler. 2020. "Exposure to untrustworthy websites in the 2016 US election." *Nature Human Behaviour* 4(5):472–480.
- Guess, Andrew M., Michael Lerner, Benjamin Lyons, Jacob M. Montgomery, Brendan Nyhan, Jason Reifler and Neelanjan Sircar. 2020. "A digital media literacy intervention increases discernment between mainstream and false news in the United States and India." *Proceedings of the National Academy of Sciences* p. 201920498.
- Hargittai, Eszter and Yuli Patrick Hsieh. 2012. "Succinct Survey Measures of Web-Use Skills." *Social Science Computer Review* 30(1):95–107.
- Horowitz, Donald L. 2001. *The deadly ethnic riot*. University of California Press.
- Jones, Carolyn. 2017. "Bill would help California schools teach about 'fake news,' media literacy." *EDSource* .

URL: <https://edsource.org/2017/bill-would-help-california-schools-teach-about-fake-news-media-literacy/582363>

- Jones-Jang, S. Mo, Tara Mortensen and Jingjing Liu. 2019. "Does Media Literacy Help Identification of Fake News? Information Literacy Helps, but Other Literacies Don't." *American Behavioral Scientist* 65(2):371–388.
- Lazer, David M. J., Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts and Jonathan L. Zittrain. 2018. "The science of fake news." *Science* 359(6380):1094–1096.
- Mercier, Hugo. 2020. *Not born yesterday: The science of who we trust and what we believe*. Princeton University Press.
- Miller, Joanne M, Kyle L Saunders and Christina E Farhart. 2016. "Conspiracy endorsement as motivated reasoning: The moderating roles of political knowledge and trust." *American Journal of Political Science* 60(4):824–844.
- Moyer-Gusé, Emily. 2008. "Toward a theory of entertainment persuasion: Explaining the persuasive effects of entertainment-education messages." *Communication Theory* 18(3):407–425.
- Munger, Kevin. 2017. "Tweetment effects on the tweeted: Experimentally reducing racist harassment." *Political Behavior* 39(3):629–649.
- Nyhan, Brendan, Ethan Porter, Jason Reifler and Thomas J Wood. 2020. "Taking fact-checks literally but not seriously? The effects of journalistic fact-checking on factual beliefs and candidate favorability." *Political Behavior* 42(3):939–960.
- Nyhan, Brendan and Jason Reifler. 2015. "The effect of fact-checking on elites: A field experiment on US state legislators." *American Journal of Political Science* 59(3):628–640.
- Osmundsen, Mathias, Alexander Bor, Peter Bjerregaard Vahlstrup, Anja Bechmann and Michael Bang Petersen. 2021. "Partisan polarization is the primary psychological motivation behind political fake news sharing on Twitter." *American Political Science Review* 115(3):999–1015.

- Pennycook, Gordon, Adam Bear, Evan T Collins and David G Rand. 2020. “The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings.” *Management Science* 66(11):4944–4957.
- Pennycook, Gordon and David G. Rand. 2019*a*. “Fighting misinformation on social media using crowdsourced judgments of news source quality.” *Proceedings of the National Academy of Sciences* 116(7):2521–2526.
- Pennycook, Gordon and David G. Rand. 2019*b*. “Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning.” *Cognition* 188:39–50.
- Pennycook, Gordon and David G Rand. 2020. “Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking.” *Journal of Personality* 88(2):185–200.
- Pennycook, Gordon and David G Rand. 2021. “The psychology of fake news.” *Trends in Cognitive Sciences* 25(5):388–402.
- Pennycook, Gordon, Jonathon McPhetres, Yunhao Zhang, Jackson G Lu and David G Rand. 2020. “Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention.” *Psychological Science* 31(7):770–780.
- Petersen, Michael Bang, Mathias Osmundsen and Kevin Arceneaux. 2022. “The “Need for Chaos” and Motivations to Share Hostile Political Rumors.” *forthcoming in American Political Science Review* .
URL: psyarxiv.com/6m4ts
- Rasmussen, Jesper, Lasse Lindekilde and Michael Bang Petersen. 2022. “Public health communication decreases false headline sharing by boosting self-efficacy.” *PsyArXiv* .
URL: <https://psyarxiv.com/8wdfp/>
- Rasmussen, Stig Hebbelstrup Rye, Alexander Bor, Mathias Osmundsen and Michael Bang Petersen. 2021. “Super-unsupervised text classification for labeling online political hate.” *PsyArXiv* .
URL: <https://psyarxiv.com/8m5dc/>

- Roozenbeek, Jon, Alexandra L. J. Freeman and Sander van der Linden. 2021. “How Accurate Are Accuracy-Nudge Interventions? A Preregistered Direct Replication of Pennycook et al. (2020).” *Psychological Science* 32(7):1169–1178.
- Roozenbeek, Jon, Cecilie S Traberg and Sander van der Linden. 2022. “Technique-based inoculation against real-world misinformation.” *Royal Society Open Science* 9(5):211719.
- Roozenbeek, Jon, Sander van der Linden, Beth Goldberg, Steve Rathje and Stephan Lewandowsky. 2022. “Psychological inoculation improves resilience against misinformation on social media.” *Science Advances* 8(34):eabo6254.
- Salovich, Nikita A and David N Rapp. 2021. “Misinformed and unaware? Metacognition and the influence of inaccurate information.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 47(4):608.
- Shu, Kai, H Russell Bernard and Huan Liu. 2019. Studying fake news via network analysis: detection and mitigation. In *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*. Springer pp. 43–65.
- Siegel, Alexandra A and Vivienne Badaa. 2020. “# No2Sectarianism: Experimental approaches to reducing sectarian hate speech online.” *American Political Science Review* 114(3):837–855.
- Theocharis, Yannis, Ana Cardenal, Soyeon Jin, Toril Aalberg, David Nicolas Hopmann, Jesper Strömbäck, Laia Castro, Frank Esser, Peter Van Aelst, Claes de Vreese et al. 2021. “Does the platform matter? Social media and COVID-19 conspiracy theory beliefs in 17 countries.” *new media & society* p. 14614448211045666.
- Thomson, Keela S and Daniel M Oppenheimer. 2016. “Investigating an alternate form of the cognitive reflection test.” *Judgment & Decision Making* 11(1).
- van der Linden, Sander. 2019. “Countering science denial.” *Nature Human Behaviour* 3(9):889–890.
- Vosoughi, Soroush, Deb Roy and Sinan Aral. 2018. “The spread of true and false news online.” *Science* 359(6380):1146–1151.

- Walter, Nathan, Jonathan Cohen, R Lance Holbert and Yasmin Morag. 2020. “Fact-checking: A meta-analysis of what works and for whom.” *Political Communication* 37(3):350–375.
- Weedon, Jen, William Nuland and Alex Stamos. 2017. “Information Operations and Facebook.” *Facebook blogpost* .
- URL:** <https://fbnewsroomus.files.wordpress.com/2017/04/facebook-and-information-operations-v1.pdf>
- Wittenberg, Chloe, Ben M. Tappin, Adam J. Berinsky and David G. Rand. 2021. “The (minimal) persuasive advantage of political video over text.” *Proceedings of the National Academy of Sciences* 118(47).

Online Appendix

Contents

A	Additional information on materials and methods	27
A.1	Specific recommendations by educational videos	27
A.2	False news quiz	28
B	Descriptive statistics	31
B.1	Twitter sample demographics	31
B.2	News Quality Measures	32
C	Lucid pilot fine-tuning the false news quiz	33
D	Post-hoc power analyses	36
E	Pre-registered effects of videos on false news sharing	36
F	Detailed results	41
F.1	Validating Super-Unsupervised news quality scores	41
F.2	Educational videos and truth discernment	41
F.3	Educational videos and news sharing on Twitter	44
F.4	Truth discernment and news sharing	45

A Additional information on materials and methods

A.1 Specific recommendations by educational videos

“How to choose your news” - TED-Ed

1. Get original news cutting out the middlemen
2. In chaotic times, do not follow news constantly
3. Check multiple sources
4. Be aware of anonymous sources
5. Verify before spreading

“Five ways to spot fake news” - Quartz

1. Consider where the information is coming from
2. Consider if the headline sound neutral
3. Consider who wrote it
4. Consider what the resources are
5. Consider if the images are accurate

“How to Spot Fake News” - Factcheck.org

1. Consider the source
2. Read beyond headline
3. Check the author
4. Consider what the support is for a claim
5. Check the date
6. Consider if it could be satire
7. Check biases
8. Consult experts

“The Fact Checker’s guide for detecting fake news” - Washington Post

1. Double check the url
2. Consider if the photo seems unrealistic
3. Check the sources
4. Use dedicated plug-ins

“Fake News with the Five C’s of Critical Consuming” - John Spencer

1. Context
2. Credibility

3. Construction
4. Corroboration
5. Compare

“Four ways to tell if something is true online - Break the Fake” - MediaSmarts

1. Use fact checking tools
2. Find the source
3. Verify the source
4. Check other sources

A.2 False news quiz

We built on previous research by Pennycook and colleagues 2019*b*; 2020 to develop our false news quiz, which consists of 8 items (4 false and 4 mainstream; 4 pro-Republican and 4 pro-Democratic, fully crossed). Each item consists of a news headline, a lead (one-sentence summary of the news) and a picture, mimicking the display of news links on social media. The items are displayed in OA Section A.2. The items were presented in a random order, and participants were asked the following question: “To the best of your knowledge, does this headline accurately describe an event that actually happened?” The answer options included 1) Definitely not, 2) Probably not, 3) Probably yes, and 4) Definitely yes.

We calculate truth discernment scores for each respondent by subtracting the average accuracy ratings of false news from the average accuracy ratings of mainstream news and dividing it by three. The truth discernment scale has a theoretical range from -1 to 1, where 0 represents no ability to discern between false and mainstream news (answering as if at random), 1 represents an excellent ability to discern between false and mainstream news, and -1 corresponds to a person who believes every false news story and disbelieves every mainstream news story. The single lowest score in our sample was -0.32, and most other negative scores received by 4% of our sample are close to 0, likely reflecting unlucky guesswork rather than delusion.

Our objective was to rely on established measures of truth discernment and to reduce measurement error as much as possible, despite the hard space constraints in our survey. We reanalyzed open data from Pennycook, Bear, Collins and Rand (2020) and ran a pilot study on

Lucid ($N = 997$) performing item response theory models to identify the eight items with the highest discrimination score (for details see Section C in the OA).

Gauging truth discernment with such a false news quiz is an established and efficient method, yet it is not without shortcomings. Specifically, some of the techniques recommended in the educational videos cannot be readily applied without access to the full text of a news story. We consider it an important task for future research to probe the external validity of our findings with alternative measures of discernment.

Participants saw the following instructions prior to starting the false news quiz:

In the next section, we will present you 8 headlines from news stories that appeared on the internet. Your task will be to indicate to the best of your knowledge, whether each headline accurately describes an event that actually happened on the following scale: 1 (definitely not), 2 (probably not), 3 (probably yes), and 4 (definitely yes).

Note that we are NOT interested whether you agree with a claim someone (allegedly) made. Instead, we would like to know whether you think the claim was actually made by that person. Similarly, we are NOT interested whether you are happy about an event that (allegedly) happened. Instead, we want to know whether you think that the event actually happened or not.

Please note that the images may take a moment to load.

Table OA1: The eight quiz items in our survey

	false news	mainstream news
Pro-Dem.	<p>Journalist Who Filmed Trump Sobbing in Rose Garden Still in Custody WASHINGTON, D.C. - A freelance photojournalist remains in custody nearly a week after he was detained by officers from the Secret Service's Uniformed... REALNEWSRIGHTNOW.COM</p>	<p>New plagiarism row for Melania Trump as pamphlet bears uncanny likeness to 2014 version Booklet on children's online safety for first lady's 'Be Best' campaign seems almost identical to one first published by FTC four years ago THEGUARDIAN.COM</p>
Pro-Dem.	<p>Trump: "Why would I care about the climate? I'll be dead in 10 years anyway" Donald Trump has once again managed to snub the international community. The USA pulling out of the Paris climate agreement rends 20 years of tough negotiations useless. Now, in an... THE-POSTILLON.COM</p>	<p>Trump once said pleading the Fifth was for 'the mob' and now he might do it Donald Trump has derided the Fifth Amendment as the refuge of mobsters, and Richard Nixon famously told aides, "I want you all to stonewall it, let them plead the Fifth Amendment,..." MSN.COM</p>
Pro-Rep.	<p>Obama Stole \$1 BILLION From Social Security, Gave It To THIS Group! Unlawful migration. Controlling it prompts taking care of such a large number of different issues in our nation. Among those issues are violations like medication trafficking, human... USACONSERVATIVEREPORT.COM</p>	<p>Darryl Strawberry says Donald Trump 'is a great man' Darryl Strawberry may want athletes to "stick to sports," but clearly he sees himself as the exception. NYDAILYNEWS.COM</p>
Pro-Rep.	<p>President Trump Fires All 14 Muslim Federal Judges And for good reason. ASAMERICANASAPPLEPIE.ORG</p>	<p>Massachusetts Democrat charged in bribery, kickback schemes A former state senator from Massachusetts was charged last week with collecting about \$1 million in bribes and kickbacks that he allegedly laundered through his company. FOXNEWS.COM</p>

B Descriptive statistics

B.1 Twitter sample demographics

Table OA2 displays the demographic characteristics of our sample. It demonstrates first that the two treatment groups are similar in terms of demographics. It also displays the PEW Research Center’s estimates of the U.S. Twitter population.³ It shows that our sample is somewhat older, more white, and more educated than the Twitter population, while it is similar in terms of gender, income and partisan identity.

Table OA2: Sample characteristics by treatment group

	Control group	Video group	PEW benchmark
Age: 18-29	0.10	0.13	0.29
Age: 30-49	0.46	0.44	0.44
Age: 50-64	0.33	0.31	0.19
Age: 65+	0.11	0.11	0.08
Black	0.08	0.08	0.11
Hispanic	0.09	0.10	0.17
White	0.78	0.75	0.60
No high school	0.01	0.01	0.04
High school	0.31	0.31	0.54
College+	0.68	0.68	0.42
Female	0.52	0.50	0.50
<30K USD	0.18	0.21	0.23
30K-75K USD	0.41	0.39	0.36
75K+ USD	0.41	0.40	0.41
Republican	0.29	0.27	0.35
Democrat	0.59	0.61	0.60
Independent	0.12	0.12	NA
N	799	801	—

³“Sizing up Twitter users” by Stefan Wojcik and Adam Hughes <https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/> last accessed 2020-06-06

B.2 News Quality Measures

Figure OA1 demonstrates that the super-unsupervised (SU) news-quality scores offer insights about many more newslinks, covering a much wider range of news, particularly on the lower end of the scale.

Figure OA1: Distribution of z-scored news quality measures.

C Lucid pilot fine-tuning the false news quiz

Developing a precise measure of truth discernment ability has been an important challenge for this study, as it is featured prominently both in the analyses of the video intervention and the of relationship between sharing and believing false news . While ideally we could have relied on a long battery with dozens of items to get a nuanced measure, this was not possible in the Twitter sample due to limits in respondent time (and more broadly, resources). Instead, we had to make sure that the items we rely on offer the least noise estimates possible. Thus, we ran a pilot study on Lucid, a large marketplace for online survey panels, relying on a sample of 997 US respondents quota sampled to resemble the marginal distribution of population demographics in terms of age, gender, education, race and region.

The ambition of this pilot was to get as much data on the performance of false news quiz items as possible. Therefore, we took the largest publicly available inventory of quiz items containing 14 False Pro-Democratic, 26 False Pro-Republican, 15 Mainstream Pro-Democratic, and 20 Mainstream Pro-Republican items (Pennycook, Bear, Collins and Rand 2020). We invited participants in our pilot to rate the accuracy of 32 items each, randomly sampling 8 items from each of the 4 categories, presented in a random order.⁴

We employ a simple, two parameter item response theory (IRT) model to identify the best items tapping into the latent trait of truth discernment. IRT is a psychometric method developed for designing and improving tests. The two parameter model estimates an item *difficulty* parameter i.e. how easy or hard an item is, and a *discrimination* parameter i.e. how good the item is at distinguishing between people with similar ability. As a rule of thumb, the best tests have items with varying difficulty and high discrimination.

Importantly, IRT models assume that the underlying latent construct is uni-dimensional. Truth discernment does not satisfy this criterion, because dispositions such as cynicism (a general disbelief in all news) will have opposite effects on truth discernment when it comes to false news items (cynicism improves performance) and “mainstream news” items (cynicism undermines performance). Therefore, we analyse each of our 4 categories separately. Although participants answer the items on a 4 point scale (as described in Section A.2), we transformed these into a dummy of correctly / incorrectly perceiving news accuracy.

⁴Ideally, all participants would have rated all items, but Lucid poses a 15 minute time limit on surveys and respondents may have experienced fatigue sooner than that.

Figure OA2: Truth discernment item difficulty and discrimination parameter estimates from a 2-factor IRT model

Note: Error bars denote standard errors. Blue dots mark the parameter estimates for the 8 (+ 1) items, which are subsequently employed in our MTurk and Twitter samples.

Figure OA2 summarises our findings plotting item difficulty (x axis) against discrimination (y axis) along with error bars denoting standard errors. Several conclusions emerge. First, all false news items and pro-Republican mainstream news items show high variance both on difficulty and discrimination. Although the data is too noisy to be sure that the best items are significantly better than the most other items at 95%, they are efficient items. Second, mainstream pro-Democratic items are much more difficult and on average substantially worse at discriminating between underlying ability. Shedding light on this interesting asymmetry is an interesting avenue for future research.

The attentive reader may notice that 3 items are highlighted with blue in the mainstream Pro-Republican panel. This is because one of the items (“Foxnews.com: Deep state has

‘weaponized’ security clearances against Trump, conservative Pentagon official’s lawyer says”) has proved to be an outlier in the MTurk pilot, and has been swapped for the next best item (“NYDailynews.com: Darryl Strawberry says Donald Trump ‘is a great man’”).

D Post-hoc power analyses

The power analysis was run in R using monte carlo simulations. We have used the same time period as in the pre-registration a three months period before the intervention and a two months period afterwards. This window is the closest we can approximate the pre-registered intention to study a 3-months pre-post period. The estimated models correspond to those pre-registered. For the measure of false news we thus run a negative binomial regression. All models are also differences-in-differences models, where we investigate the significance of the interaction term in the estimated models. For all models we use a significance level of .05. The power in our simulation studies simply corresponds to the proportion of statistically significant interaction terms in our simulated datasets. We ran 1000 simulations for each effect size for all models.

We have used the `rnegbin` function from the `MASS` library to simulate the number of false news using the dispersion parameter from a negative binomial regression on the actual data. We then randomly draw realizations from this distribution varying the effect size, i.e. the reduction in the number of false news before and after the intervention. The mean number of false news shared per user-week in the pre-period is 0.0252. An effect size of -0.0252 would thus correspond to the intervention removing *all* false news in the post period. From figure OA3 we can see that 80 percent power is obtained at an effect size of -.01 which corresponds to removing about 40% of all false news in the post period.

E Pre-registered effects of videos on false news sharing

This section reports pre-registered analyses estimating the causal effect of our intervention on false news shared on Twitter. In testing the effects of educational videos on false news sharing on Twitter, we followed Guess et al. (2019); Grinberg et al. (2019); Lazer et al. (2018) and identified false news on the source level. In other words, we operationalized false news sharing as (re)tweeting URLs from websites with a known history of publishing factually incorrect news. The lists of false news sites are based on previous work by (Guess et al. 2019; Grinberg et al. 2019) and are available in the project’s OSF repository.

Our analyses depart from our pre-registrations in three ways: 1) We excluded a single respondent in the control group who shared a very large number of false news stories in the pre-

Figure OA3: Power for false news sharing

Note: Power for various effect sizes for false news

treatment period but then substantially reduced their false news sharing activity a few weeks *before* the experiment was fielded. 2) Whereas we planned to run multilevel models investigating the unique treatment effect of the six educational videos, we limit these to analyses of truth discernment. The models of news sharing quality are too noisy to yield meaningful video-level effects. 3) While we planned to observe respondents for up to 3 months post-treatment, our study period is limited to 9 weeks after exposure. In retrospect, we have been overly optimistic about the decay effects, as similar interventions are often found to persist for a few weeks at most.

We model number of false news stories shared in negative binomial models. We model the

binary outcome (did the respondent share any false news in a given week?) and the proportion of false to mainstream news in binomial regression models. We run all models on false news sharing in a Bayesian setting relying on `stan` and the `brms` R package. We defined weakly informative priors centered at 0, which do not affect our estimates but help with model convergence.

Our respondents shared 2,319 false news in the study period, and only 12% of participants (189 users) shared any false news in the study period. We employ three tools to improve the accuracy of our estimates in our data: 1) we employ a difference-in-differences design, which allows us to ask whether the treatment makes a difference taking into account differences between the treatment and control groups in the pre-treatment period and the behavior change from pre- to post-treatment in the control group; 2) we rely on three different conceptualizations of false news sharing behavior, each mitigating different aspects of methodological issues; 3) finally, we aggregate our data to user-weeks, instead of more granular units (e.g. user-days), which reduces the variance of our data substantially.

Our main model contrasts the number of false news shared in the three months prior to and nine weeks after the intervention. Our main pre-registered model, (3 months pre-treatment, and maximal, 9 weeks post-treatment period) implies that the treatment group shared 55% *more* false news than expected based on pre-treatment differences and time trend. That said, the estimate is very noisy, ranging from a small (7%) improvement to more than 150% increase ($\beta = 1.55$; 95%*CI* : [0.93 – 2.58]).

We report several additional models in Figure OA4. First, we explore both shorter and longer pre-treatment periods from one year to two weeks (depicted as vertical facets). Second, we explore shorter post-treatment periods from one to nine weeks (depicted on x axis). Third, we test the effect for two additional outcome measures: 1) the probability that the respondent shares at least one false news story in a given week, and 2) the proportion of false news to mainstream news shared on a given week. Although most models' point estimate finds that in the initial weeks post-exposure participants posted fewer false news, the credible intervals are so wide across all models that neither huge improvements, nor large setbacks in false news sharing cannot be ruled out.

We investigated (post-hoc) if a more comprehensive list of false news domains from Grinberg et al. (2019) improves the precision of our estimates. Alas, despite three times more links

Figure OA4: Treatment effect on sharing false news by various pre-treatment periods (vertical panels) and post-treatment periods (x axis). All coefficients are odds ratios based on the interaction terms in difference-in-differences models

Note: Number of false news modeled in a negative binomial model, Probability of sharing one or more false news and Proportion of false news are modeled as binomial models. Error bars denote 90% and 95% credible intervals.

categorised as false news and 40% more respondents sharing at least one this specification remains inconclusive too (see Figure OA5).

Figure OA5: Treatment effect on sharing false news by various pre-treatment periods (vertical panels) and post-treatment periods (x axis). All coefficients are odds ratios based on the interaction terms in difference-in-differences models

Note: Number of false news modeled in a negative binomial model, Probability of sharing one or more false news and Proportion of false news are modeled as binomial models. Error bars denote 90% and 95% credible intervals.

F Detailed results

F.1 Validating Super-Unsupervised news quality scores

Figure OA6: Predicted news quality scores based on super-unsupervised method (y) and expert ratings (x) are highly correlated.

Note: Points denote news domains (left) or users (right) Blue lines denote OLS regression line.

F.2 Educational videos and truth discernment

Table OA3: Main effects from MLM models

	<i>Dependent variable:</i>	
	Truth discernment	
	Mturk pilot	Twitter sample
	(1)	(2)
ATE	0.07* (0.02, 0.13)	0.05** (0.02, 0.08)
Control group	0.30*** (0.25, 0.35)	0.37*** (0.35, 0.39)
Observations	686	1,600

Note: *p<0.05; **p<0.01; ***p<0.001

Table OA4: Random effects for the six videos from MLM models

	Group	ATE	SE	Data
1	Pooled estimate	0.07	0.03	MTurk pilot
2	FactCheck.org	0.06	0.02	MTurk pilot
3	John Spencer	0.08	0.02	MTurk pilot
4	MediaSmarts	0.10	0.02	MTurk pilot
5	Quartz	0.09	0.02	MTurk pilot
6	TED-Ed	0.04	0.02	MTurk pilot
7	Washington Post	0.07	0.02	MTurk pilot
8	Pooled estimate	0.05	0.02	Twitter sample
9	FactCheck.org	0.04	0.02	Twitter sample
10	John Spencer	0.06	0.02	Twitter sample
11	MediaSmarts	0.06	0.02	Twitter sample
12	Quartz	0.02	0.02	Twitter sample
13	TED-Ed	0.07	0.02	Twitter sample
14	Washington Post	0.07	0.02	Twitter sample

Table OA5: Main effects from MLM models

	<i>Dependent variable:</i>			
	Gullibility		Trust in MSM	
	Mturk pilot	Twitter sample	Mturk pilot	Twitter sample
	(1)	(2)	(3)	(4)
ATE	−0.04 (−0.09, 0.01)	−0.04*** (−0.07, −0.02)	0.03 (−0.01, 0.07)	0.01 (−0.01, 0.03)
Control group	0.26*** (0.21, 0.30)	0.24*** (0.23, 0.26)	0.56*** (0.53, 0.59)	0.61*** (0.60, 0.62)
Observations	686	1,600	686	1,600
<i>Note:</i>			*p<0.05; **p<0.01; ***p<0.001	

F.3 Educational videos and news sharing on Twitter

Table OA6: Detailed regression results from our main models

	<i>Dependent variable:</i>			
	Expert ratings		Super Unsupervised scores	
	OLS (1)	Fixed Effects (2)	OLS (3)	Fixed Effects (4)
Treatment group	-0.44*** (-0.77, -0.11)		-0.30 (-0.68, 0.07)	
Post-treatment period	-0.07 (-0.21, 0.07)	-0.03 (-0.15, 0.09)	-0.07 (-0.16, 0.02)	-0.02 (-0.08, 0.05)
Treatment $\times$ Post-treatment	0.27*** (0.03, 0.51)	0.15* (-0.01, 0.31)	0.19* (-0.03, 0.40)	0.09* (-0.01, 0.18)
Constant	0.22** (0.05, 0.39)		0.16 (-0.06, 0.39)	
Observations	7,340	7,340	17,468	17,468
Adjusted R ²	0.05	0.36	0.02	0.46
Residual Std. Error	0.98 (df = 7336)	0.80 (df = 7014)	0.97 (df = 17464)	0.72 (df = 17039)
<i>Note:</i>				
*p<0.1; **p<0.05; ***p<0.01				

Figure OA7: Estimating the within-user effect of the intervention across several pre- and post-treatment periods.

Note: Coefficient estimates are the interaction term in a DiD model with user fixed effects. Errorbars are 95% confidence intervals.

F.4 Truth discernment and news sharing

Here we report full models for the relationship between truth discernment and false news sharing.

Table OA7: Regressing Shared News Quality on Discernment, Demographics, and Psychological Covariates

	<i>Dependent variable:</i>		
	Quality of Shared News (SU Scores)		
	(1)	(2)	(3)
Truth Discernment	0.22*** (0.14, 0.29)	0.21*** (0.13, 0.28)	0.11** (0.02, 0.19)
Age		−0.60*** (−0.93, −0.27)	−0.36** (−0.71, −0.02)
Female		−0.15** (−0.30, −0.01)	−0.09 (−0.23, 0.06)
Higher education		0.17** (0.01, 0.33)	0.08 (−0.07, 0.24)
Republican ID			−0.96*** (−1.17, −0.75)
CRT			0.01 (−0.05, 0.07)
Trolling			−0.002 (−0.08, 0.08)
Cynicism			−0.08** (−0.15, −0.02)
Knowledge			0.03 (−0.04, 0.10)
Constant	−0.03 (−0.10, 0.04)	0.19* (−0.01, 0.39)	0.72*** (0.25, 1.18)
Observations	715	715	702
R ²	0.04	0.07	0.18
<i>Note:</i>		*p<0.1; **p<0.05; ***p<0.01	

Table OA8: Regressing Shared News Quality on Discernment etc using expert ratings

	<i>Dependent variable:</i>		
	Quality of Shared News (experts)		
	(1)	(2)	(3)
Truth Discernment	0.33*** (0.25, 0.41)	0.32*** (0.24, 0.40)	0.23*** (0.14, 0.32)
Age		-0.76*** (-1.12, -0.40)	-0.63*** (-1.00, -0.25)
Female		-0.01 (-0.16, 0.15)	-0.01 (-0.17, 0.15)
Higher education		-0.003 (-0.07, 0.06)	0.01 (-0.06, 0.07)
Republican ID			-0.68*** (-0.91, -0.45)
CRT			0.05 (-0.01, 0.12)
Trolling			0.08 (-0.02, 0.17)
Cynicism			-0.07* (-0.14, 0.001)
Knowledge			0.07* (-0.003, 0.14)
Constant	-0.06 (-0.13, 0.02)	0.29** (0.07, 0.51)	0.29 (-0.24, 0.82)
Observations	575	575	564
Adjusted R ²	0.10	0.12	0.18
<i>Note:</i>		*p<0.1; **p<0.05; ***p<0.01	

Table OA9: Negative binomial regression of # of fake news shared on discernment and co-variates

	Model 1	Model 2	Model 3
Intercept	2.02*	-1.03*	-2.50*
	[1.78; 2.28]	[-1.72; -0.30]	[-4.07; -0.92]
Truth Discernment	-0.46*	-0.09	-0.06
	[-0.77; -0.14]	[-0.34; 0.18]	[-0.36; 0.23]
Age		5.71*	4.32*
		[4.55; 6.85]	[3.16; 5.50]
Female		0.91*	0.09
		[0.43; 1.39]	[-0.44; 0.61]
Higher education		-0.43	-0.53*
		[-0.98; 0.10]	[-1.05; -0.02]
White		-0.42	-0.34
		[-1.08; 0.17]	[-0.96; 0.24]
Party ID			2.67*
			[2.00; 3.35]
CRT			-0.49*
			[-0.67; -0.30]
Trolling			0.11
			[-0.16; 0.39]
Cynicism			0.23
			[-0.00; 0.45]
Knowledge			0.21
			[-0.06; 0.46]
N	715	715	702

* Null hypothesis value outside 95% credible interval (reported in square brackets).

Table OA10: Negative binomial regression of # of real news shared on discernment and covariates

	Model 1	Model 2	Model 3
Intercept	3.15*	1.76*	0.71
	[2.89; 3.44]	[1.26; 2.26]	[−0.18; 1.62]
Truth Discernment	2.18*	2.37*	1.91*
	[1.60; 2.75]	[1.86; 2.85]	[1.35; 2.47]
Age		3.13*	3.05*
		[2.54; 3.72]	[2.44; 3.66]
Female		−0.04	−0.40*
		[−0.31; 0.23]	[−0.67; −0.12]
Higher education		−0.07	−0.37*
		[−0.37; 0.20]	[−0.65; −0.09]
White		−0.23	−0.18
		[−0.59; 0.10]	[−0.51; 0.13]
Party ID			0.41*
			[0.01; 0.81]
CRT			−0.38*
			[−0.49; −0.28]
Trolling			0.31*
			[0.16; 0.47]
Cynicism			−0.06
			[−0.18; 0.05]
Knowledge			0.46*
			[0.34; 0.58]
N	715	715	702

* Null hypothesis value outside 95% credible interval (reported in square brackets).

Truth discernment and sharing of fake and real news on Twitter (#31178)

Created: 11/14/2019 12:35 PM (PT)

Shared: 01/31/2020 02:16 AM (PT)

This pre-registration is not yet public. This anonymized copy (without author names) was created by the author(s) to use during peer-review. A non-anonymized version (containing author names) will become publicly available only if an author makes it public. Until that happens the contents of this pre-registration are confidential.

1) Have any data been collected for this study already?

It's complicated. We have already collected some data but explain in Question 8 why readers may consider this a valid pre-registration nevertheless.

2) What's the main question being asked or hypothesis being tested in this study?

H1a. The higher a person's ability to discern fake news from real news, the fewer fake news stories they will share on Twitter.

H1b. The higher a person's ability to discern fake news from real news, the more fake news stories they will share on Twitter.

H2. The higher a person's ability to discern fake news from real news, the more real news stories they will share on Twitter.

3) Describe the key dependent variable(s) specifying how they will be measured.

The key DV is the number of fake news stories shared on Twitter from verified fake news websites in the period between 2018 November (12 months prior to data collection) and 2020 February (3 months after data collection). The verified fake news website list comes from Guess, Andrew, Jonathan Nagler, and Joshua Tucker. "Less than You Think: Prevalence and Predictors of Fake News Dissemination on Facebook." Science Advances 5, no. 1 (2019). <https://doi.org/10.1126/sciadv.aau4586>.

4) How many and which conditions will participants be assigned to?

This is a cross-sectional study. Participants will be participating in two experiments unrelated to the present pre-registration.

Our main independent variable is truth discernment. Participants take a "fake news quiz" seeing 8 news headlines, (4 fake and 4 real). They are asked to rate each headline whether it accurately describes an event that actually happened, truth discernment is calculated as the average accuracy ratings of real news minus average accuracy ratings of fake news divided by 3. Higher values indicate better discernment.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.

Our main analysis concerns the correlation between the number of fake news shared and the truth discernment score. We will run three regression models, regressing number of fake news shared on truth discernment score 1) without any covariates, 2) with demographic covariates (age, gender, and education), 3) with psychological covariates (partisan identity, cognitive reflection score, trolling, political cynicism, and political knowledge).

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.

We exclude participants who provided two different yet valid Twitter ID in Wave 1 and Wave 2.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.

This is wave 2 of a panel study of Twitter users, who consented to scrape their feeds. In the first wave, 3016 participants gave their consent by entering a Twitter ID in the appropriate field. Of these, we could identify and scrape 2336 users. All 3016 participants will be invited to participate in the survey, but naturally we can include in the study only participants who provided a valid Twitter ID in at least one of the two surveys.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)

Truth discernment will be measured with a quiz implemented in wave 2 of our panel study, which is yet to be run. However, we will rely on fake news counts from participants' scraped Twitter feed. Parts of this data we already analyzed for a different study.

We consider it useful to distinguish between three periods when it comes to participants' Twitter behavior. 1) prior to enrolling in wave 1 of our study (– 2019 February), 2) period between the two waves (2019 Feb – 2019 Nov), 3) period following Wave 2 (2019 Nov –). Experimental demand effects and Twitter's efforts to reduce fake news on the network affect each period differently. Consequently, we will run robustness checks splitting the DV to these three periods.

The individual-level causal effects of popular anti-“fake news” videos (#31174)

Created: 11/14/2019 12:23 PM (PT)

Shared: 01/31/2020 02:17 AM (PT)

This pre-registration is not yet public. This anonymized copy (without author names) was created by the author(s) to use during peer-review. A non-anonymized version (containing author names) will become publicly available only if an author makes it public. Until that happens the contents of this pre-registration are confidential.

1) Have any data been collected for this study already?

It's complicated. We have already collected some data but explain in Question 8 why readers may consider this a valid pre-registration nevertheless.

2) What's the main question being asked or hypothesis being tested in this study?

A) Does exposure to an educational anti fake news video improve truth discernment (the ability to distinguish between fake news and real news)?

B) Does exposure to an educational anti fake news video reduce the number of fake news shared on Twitter?

3) Describe the key dependent variable(s) specifying how they will be measured.

A) Participants will take a "fake news quiz", seeing 8 news headlines, (4 fake and 4 real). They are asked to rate each headline whether it accurately describes an event that actually happened using the following scale: 1 (definitely not), 2 (probably not), 3 (probably yes), and 4 (definitely yes).

Our DV is the truth discernment score i.e. the average accuracy ratings of real news minus average accuracy ratings of fake news divided by 3. Higher values indicate better discernment.

B) The number of links shared from verified fake news websites in the 3 months before and after treatment.

4) How many and which conditions will participants be assigned to?

Control: No video. 50% of sample

Treatment: Asked to watch an educational videos on how to spot fake news. 50% of sample. Participants in the treatment group are randomly allocated to one of six videos. All videos are available on Youtube.

1. Four ways to tell if something is true online – MediaSmarts, <https://youtu.be/E-049KTrYBg>,
2. How to choose your news – Damon Brown – Ted ED – <https://youtu.be/q-Y-z6HmRgl>,
3. Helping Students Identify Fake News with the Five C's of Critical Consuming – John Spencer – <https://youtu.be/xf8mjbVRqao>
4. Five ways to spot fake news – Quartz – <https://youtu.be/y7eCB2F89K8>
5. How to Spot Fake News - FactCheck.org – <https://youtu.be/AkwWcHekMdo>
6. The Fact Checker's guide for detecting fake news – Washington Post – <https://youtu.be/SoWCDJAMk2Y>

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.

A) Simple independent samples t-test (control vs treatment) with truth discernment scores.

B) Difference-in-differences estimator of the number of fake news stories shared on Twitter with 3 months prior to exposure as the baseline period and 3 months after the exposure as the treatment period.

For both analyses we will be running multilevel models estimating the effect of the specific videos.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.

We exclude participants who fail an audio/video equipment test or provide two different yet valid Twitter ID in Wave 1 and Wave 2 of our survey.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.

This is wave 2 of a panel study of Twitter users, who consented to share data from their feeds. In the first wave, 3016 participants gave their consent by entering a Twitter ID in the appropriate field. Of these, we could identify and scrape 2336 users. All 3016 participants will be invited to participate in the survey, but naturally we can include in the study only participants who provided a valid Twitter ID in at least one of the two surveys.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)

We already possess, although have not analyzed the data constituting the baseline of fake news sharing behavior. We have *not* run the experiment yet. We will run robustness checks varying the size of the window around the treatment to identify fake news sharing prior and after participant in the experiment.