RUNNING HEAD: FAKE NEWS GUIDE

In press, Collabra

A practical guide to doing behavioral research on fake news and misinformation

Gordon Pennycook^{1,2*}, Jabin Binnendyk², Christie Newton¹, & David G. Rand^{3,4,5†}

¹Hill/Levene Schools of Business, University of Regina, ²Department of Psychology, University of Regina

³Sloan School of Management, Massachusetts Institute of Technology, ⁴Institute for Data, Systems, and Society, Massachusetts Institute of Technology, ⁵Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

*Corresponding authors: gordon.pennycook@uregina.ca

Date: 7/12/21

Abstract

Coincident with the global rise in concern about the spread of misinformation on social media, there has been influx of behavioral research on so-called “fake news” (fabricated or false news headlines that are presented as if legitimate) and other forms of misinformation. These studies often present participants with news content that varies on relevant dimensions (e.g., true v. false, politically consistent v. inconsistent, etc.) and ask participants to make judgments (e.g., accuracy) or choices (e.g., whether they would share it on social media). This guide is intended to help researchers navigate the unique challenges that come with this type of research. Principle among these issues is that the nature of news content that is being spread on social media (whether it is false, misleading, or true) is a moving target that reflects current affairs in the context of interest. Steps are required if one wishes to present stimuli that allow generalization from the study to the real-world phenomenon of online misinformation. Furthermore, the selection of content to include can be highly consequential for the study’s outcome, and researcher biases can easily result in biases in a stimulus set. As such, we advocate for pretesting materials and, to this end, report our own pretest of 225 recent true and false news headlines, both relating to U.S. political issues and the COVID-19 pandemic. These headlines may be of use in the short term, but, more importantly, the pretest is intended to serve as an example of best practices in a quickly evolving area of research.

Keywords: misinformation; fake news; news media; truth discernment; social media

Introduction

Coincident with the huge increase in concern about the spread of false and misleading content online (primarily through social media), there has been an influx of research across the social sciences on the topic. In this guide, we will explain the steps involved in an increasingly common sub-category of this body of work: Online or lab survey studies that present people with *actual* news headlines of varying quality and that take some measurements relating to those headlines (e.g., whether people believe them to be accurate or are willing to share them on social media). The present guide is of primary relevance for the psychology of fake news and misinformation (for a review, see Pennycook & Rand, 2021), but may be useful for other areas as well.

By accident of history, the studies that use this design tend to focus on the phenomenon of “fake news” (i.e., news headlines that are fabricated but presented as if legitimate and spread on social media; Lazer et al., 2018); however, one could use the same guide to create stimuli to test hypotheses using only true content from mainstream sources, or for false content that comes in other forms, etc. Thus, although this is a practical guide to doing research on fake news and misinformation, it’s really a guide to doing behavioral research on “news” (or news-adjacent; e.g., memes) content that has ecological validity because it is taken from the real world (as opposed to being constructed by researchers). Furthermore, the guide is most relevant for studies of news *headlines* as opposed to full articles. The reason for this is that having participants read full news articles poses a host of methodological difficulties (e.g., do you allow people to decide how much of an article to read or try to enforce full attention?) and most people on social media do not typically read past the full headline anyway (Gabelkov et al., 2016). Nonetheless, aspects of this guide could easily be adapted to accommodate studies of full articles.

Background

The first published empirical study that used actual fake news headlines from social media as stimuli (see Figure 1 for representative examples) investigated the role of exposure (and, specifically, repetition) on belief in true and false news content (Pennycook et al., 2018). Participants were presented with a set of political headlines that varied on two dimensions: 1) They were either “fake news” (i.e., examples of headlines that have been determined to be false by fact-checkers) or “real news” (i.e., examples of true headlines that come from mainstream news sources) and 2) They were either good for the Democratic party in the United States (Pro-Democrat) or good for the Republican party (Pro-Republican), based on a pretest (discussed subsequently). The key experimental manipulation was exposure: Some headlines were initially presented in a familiarity phase where participants were asked about familiarity or their willingness to share the headline on social media and, thus, when participants reached the critical assessment phase (where they judged the accuracy of *each* headline) some of the headlines had been previously seen in the context of the experiment. The critical finding is that a single prior exposure to the fake news headlines increased later belief in the headlines – and that this was true even for those headlines that were *inconsistent* with people’s political partisanship (i.e., repetition increased belief even for headlines that were Pro-Republican among Democrats, and vice versa for Pro-Democratic headlines among Republicans). This brought new light to the long established “illusory truth effect” (Fazio et al., 2015; Hasher et al., 1977), at least as argued by the authors.

Figure 1. Examples of fake/false (top) and real/mainstream/true (bottom) headlines that are either Pro-Republican (left) or Pro-Democratic (right).

This basic paradigm of presenting participants with news headlines and varying aspects of the presentation or changing the questions that are asked of participants has been used in a number of published papers since this initial publication (see Pennycook & Rand, 2021). Although many studies have focused on people's belief in fake versus real news (e.g., Allcott & Gentzkow, 2017; Pennycook & Rand, 2019; Vegetti & Mancosu, 2020), others have focused more on people's willingness to share the content on social media (e.g., Altay et al., 2020; Osmundsen et al., 2021; Pennycook et al., 2021). However, an important caveat for these studies is that asking about *both* accuracy and social media sharing undermines inferences that one can make about the latter judgment. Specifically, recent work shows that people often do not consider whether a headline is accurate when they make judgments about whether they would be willing to share it (Pennycook et al., 2021; Pennycook, McPhetres, et al., 2020); thus, asking about accuracy prior to sharing biases sharing judgments to be, in essence, too reasonable. Importantly, hypothetical judgments about social media sharing do correspond with actual sharing decisions on social media (specifically, the same headlines that people report a greater willingness to share in online surveys are actually shared more frequently on social media; Mosleh, Pennycook, & Rand, 2020), so long as these sharing decisions are not being influenced by other questions in the study.

Other work focuses on interventions against fake news. Using actual true and false content from social media (and that is representative of the broader categories) is particularly important if one

wants to make the argument that their favored intervention is likely to have an impact if implemented in the real world. This research has looked at factors such as fact-checking (Brashier et al., 2021; Clayton et al., 2019; Pennycook, Bear, et al., 2020), emphasizing news sources/publishers (Dias et al., 2020), digital media literacy tips (Guess, Lerner, et al., 2020), inoculation and other educational approaches (van der Linden et al., 2020), and subtly prompting people to think about accuracy to improve sharing decisions (Pennycook et al., 2021; Pennycook, McPhetres, et al., 2020).

Since research of this nature may have an impact on policy, it is particularly important for researchers to uphold a high standard for testing. This requires a strong understanding of the type of content that the interventions do and do not impact. To give a stylized and perhaps extreme example, one could easily devise an “intervention” that teaches people to be wary of a set of idiosyncratic features of fake news, such as the sole use of capital letters or the presence of spelling errors (these do, in fact, come up every once in a while). In testing this intervention, if only headlines that contain these features are used, it will make the intervention appear to be extremely effective at helping people distinguish between true and false news. However, if the intervention were tested on a random sample of fake news headlines (in which full caps and spelling errors are not particularly common overall), the researcher would realize the limits of the intervention as it would likely have no impact on people’s ability to discern between true and false news outside of those particular tactics. This is a fundamental issue with all digital media literacy, inoculation, and (more broadly) educational approaches: Teaching people about the specific tactics that are used, or what features to look for when identifying fake news (and related) can only be effective insofar as the tactics continue to be used and the features continue to be present in the content that people come across in everyday life.

We will next outline a step-by-step guide to doing research of this nature based on the approach taken by the present authors. Of course, there are several ways to do such studies and alternative approaches will be discussed throughout. As discussed in Step 1, there are a number of assumptions that have to be made when doing research of this nature and so a major focus of this guide is to highlight these nuances.

Step 1 – Selecting your misinformation

One of the most critical things to keep in mind when researching fake news and misinformation is that these are not natural categories. When one makes reference to these things, it is actually a reference to whatever happens to exist as the present form of misinformation (or “true content”) that is occurring in the world. What this means is that using stimuli in a study (for the purpose of generalizing beyond the study) *has to* be informed by what is happening in the world.

If you wish to create your own fake news stimuli (as has been done in some studies; Pereira et al., 2020), you need to take extra care to make sure that your content is similar to what is seen in everyday life – otherwise your results cannot be taken to generalize (i.e., you won’t be researching fake news per se, but your own version of “fake news” which may or may not be representative of the larger category). Furthermore, many false headlines that sprout up on the internet are rarely shared; and, thus, the headlines that spread sufficiently widely to warrant a fact-check are pre-

selected (in a sense) to be the type of content that contain the features that facilitate their spread on social media. For these reasons, it is often better (in our view) to obtain stimuli that actually has been spread on social media¹. Creating your own headlines means that you don't (and can't) know if they are the sort that would actually spread on social media, although there may be good theoretical reasons to do so. Nonetheless, it is critical that any investigation of fake news (and related) is sensitive to these issues and discusses them directly.

Whether or not you decide to create your own “fake news”, it is important to get a sense of what fake news looks like in the context that you are interested in. To state the obvious, if you would like to run a study on misinformation in Brazil, it would not make sense to use examples of fake news headlines from the United States. Relatedly, if you are reading this paper in 2025 and would like to run a study in the United States, the headlines mentioned in the pretest reported below are not going to be sufficiently up-to-date.

Now that you have some sense of the broader category of misinformation that you want to draw from, you can start searching for stimuli. Fortunately, there are numerous fact-checking organizations around the world that keep records of false or misleading content that is spread on the internet. A list of the signatories of the Poynter International Fact-Checking Network can be found here: <https://ifcncodeofprinciples.poynter.org/signatories>. Common fact-checking sources in North America are <https://www.snopes.com/>, <https://www.factcheck.org/>, and <https://www.politifact.com/>. Another major fact-checking organization (internationally) is <https://factcheck.afp.com/>. Importantly, many falsehoods on these websites are not examples of “fake news” – again, it is up to you to decide what type of misinformation that you would like to research. For example, snopes.com uses “junk news” to refer specifically to news headlines that are false or misleading (“fake news”). However, they also fact-check false memes, widespread conspiracy theories, claims from politicians, etc.

You've found an example of misinformation that you would like to include in a study, what next? It depends on the class of misinformation. The fact-checking website will sometimes have an image of the falsehood with their explanation of why it is incorrect – if so, you're off! Just save the image and you have your stimuli. If not, you will have to decide how you want the falsehood to be presented. For example, in the context of fake news headlines, it is often desirable to present the image in “Facebook format” – i.e., in the way that the headline would look if it were shared on Facebook (which is a particularly common source for fake news, at least in North America; Allcott et al., 2019; Del Vicario et al., 2016; Guess, Nyhan, et al., 2020). For this, we find the *original* URL for the fake news headline (sometime this requires a search outside of the fact-checking website and sometimes the link cannot be found – see Figure 2 for a representative process) and literally put it into Facebook as if it was going to be shared with friends and family. [Important:

---

¹ There are, of course, cases where this does not hold. For example, McPhetres, Rand, and Pennycook (2020) investigated the specific role of “character deprecation” in the spread of fake news and found that, using real-world stimuli, Republicans preferred headlines that called the character of Democrats into question (relative to vice versa for Democrats). However, when we made up our own headlines that were perfectly matched in an experimental design (the headlines were inspired by “real” fake news headlines), the same result did not emerge. Thus, it appears that character deprecation is more common and salient in right wing news, but not because Republicans have a particularly strong predilection for it.

Don't actually share the headline.] Facebook then provides a "preview" of what the shared headline would look like. Take a screenshot of this and edit it down to a single image (see Figure 2 for further details). Of note, Facebook seems to change slight aspects of this formatting from time-to-time, although we doubt this is particularly important. Also, when taking screenshots, the quality of your image will depend on the display resolution of your screen; thus, try to increase the display resolution for your laptop or monitor as high as possible.

A)

Fact Checks · Junk News

Did AOC Say 'It's Okay to Eat Pets' if Food Runs Scarce?

Inflammatory, fabricated quotes have been making their way around social media.

BETHANIA PALMA
PUBLISHED 14 MAY 2020

Image via Tom Williams/CO-Roll Call, Inc via Getty Images

archive.today Saved from <https://www.news.bizwiznews.com/ocasio-cortez-if-food-runs-scarce-its-okay-to-eat-pets/> 14 May 2020 15:58:00 UTC

B)

Origin

On May 9, 2020, the BIZ News website published a [story](#) reporting that U.S. Rep. Alexandria Ocasio-Cortez, D-N.Y., stated that it was "okay to eat pets" if food runs scarce.

The story circulated on Facebook, prompting Snopes readers to ask whether Ocasio-Cortez, often called by her initials AOC, had really made the remark quoted in the headline: "If Food Runs Scarce, It's Okay To Eat Pets."

The story also contained another purported quote from Ocasio-Cortez:

"Look, I don't know why, if we're closing everything and being all scared and piling up food and toilet paper, we should have to worry about some dumb dogs and cats too. Like, last time I checked, they didn't do anything except sit around. And if food does start running short, I mean, sorry Charlie, right? Gawd these glazed are to die for."

Ocasio-Cortez did not make these comments, which have not been quoted or referenced in any legitimate news sources. We therefore rate this claim false.

BIZ News is a junk news website with a track record of fabricating inflammatory quotes and attributing them to Democratic political figures, including a fake quote [attributed](#) to House Speaker Nancy Pelosi stating that senior citizens were "not entitled to Social Security."

C)

BIZ News

Ocasio-Cortez: 'If Food Runs Scarce, It's Okay To Eat Pets'

admin · May 9, 2020 · Leave a comment

D)

Create post

Gord Pennycook

Friends

<https://www.news.bizwiznews.com/ocasio-cortez-if-food-runs-scarce-its-okay-to-eat-pets/>

URL Blocked

Could not scrape URL because it has been blocked.

OK

Add to your post

Post

Figure 2. Step-by-step process of obtaining an image of a fake news headline (Note: This is just an example – results will vary). A) Find a suitable headline (in this case, the headline came from Snopes: <https://www.snopes.com/fact-check/aoc-pets/>). B) Find link to original headline source in fact-check article (if you can't find it, you may have to try searching online). C) Copy URL. D) Put URL in Facebook to generate a preview (important: Do not actually share the headline). In this case, Facebook blocked the URL. (In many cases, Step D is the final step because Facebook has not blocked the URL.). E) Search online (via Google or whatever your favored search engine) for the content of the headline and find a non-fact-checked reference to it. F) Put that link into Facebook to generate a preview. Take a screenshot. G) Put the screenshot into picture editing software (e.g., MS Paint, Photoshop), remove the "X" in the top right corner (or whatever Facebook happens to generate when you're performing this task – the "X" is not displayed when the headline is shared and should ideally be removed from the image) and save the image. Note: You may have to do some minimal editing of the post. In this case, the source "DailyWorldUpdate" places a subtle "satire" warning on the bottom left of the image, presumably so Facebook will not block their URLs. We simply removed this part of the image. What remains is a fairly close approximation of the originally fact-checked fake news article. The source has changed, but people are not familiar with these sources and, anyway, sources have very little impact on judgments for headlines of this sort (Dias et al., 2020). The overall process is typically easier than this example.

Step 2 – Selecting true content

For most studies, it won't be enough to simply have misinformation. Often, for example, determining how people distinguish between true and false content is central (see Batailler et al., 2021; Pennycook & Rand, 2021 for a discussion of this issue). If you're only interested in testing misinformation and not using a "true content" baseline or comparison, skip this step.

In a certain sense, finding misinformation is the easy part. This is because there is a clear way to operationalize "misinformation" – namely, those things that fact-checkers have identified as being false or misleading. When it comes to true content, there is a much larger world of possibilities. In our past work on fake news, we made an attempt to find true content that was similar in form to the false content. That is, we go to reliable mainstream sources such as *The New York Times* or *Washington Post* (which sources are reliable will depend on the country that you are running your study in) and find headlines that are *not* likely to be highly familiar to the average person. Note that, when choosing a source for true content, it is important for it to be reliable for the content to actually be true. There are many possible sources to choose from (and this number may increase in the future), so care is required.

In terms of what sort of stories to choose, very significant news events, such as a major natural disaster, may not make for very good content for studies that focus on belief because they are obviously true. They would also probably not be very good for studies that focus on sharing because people would not likely to be particularly willing to share them (assuming the study was run in the days or weeks following the news story, making it "yesterday's news"). When searching for content, it is also important to sample from as many different sources as possible (preferably across political lines), so long as they are reliable.

An important consideration for true content (that is less a problem for false content²) is that particularly salient or topical news headlines may quickly become outdated. Our favored approach is to select headlines that would make sense for someone to share several months after the headline was published because they describe events that could happen at any time. The pretest, described subsequently, has many examples of such headlines. At any rate, the literal process for obtaining the images for true content is more straightforward than for fake news: One merely has to put the URL into Facebook and a preview will generate.

As a brief aside, we suggest *against* only using true content that has been verified as such by fact-checkers (as done, for example, by Vosoughi, Roy, & Aral, 2018) if you want to generalize to "true news" more broadly. This is because the type of true content that is fact-checked is not *remotely* representative of the larger category of "true" news: The only headlines for which fact-checking is necessary have (at least somewhat) ambiguous truth-claims. Many, if not most, "true" headlines are not the sort that even need fact-checking. For example, headlines that report on the results of an election (under normal circumstances), or headlines that report on the results of an election being disputed (under less than normal circumstances), or headlines that report on the

---

² False content is not constrained by reality, and hence does not as readily become outdated by well-known events. For example, false headlines about the Clinton's being responsible for murders will be political salient for as long as the Clintons are known political figures in the United States.

spread of a global pandemic, or headlines that report on a direct statement made by a political figure or major event with numerous witnesses, etc.

Step 3 – Consider completing a pretest

Depending on the research that you wish to do, a pretest of the materials selected in Steps 1 and 2 may be a good idea. The fundamental issue is that individual biases from the researcher – not to mention pure randomness – will surely influence the headlines that are selected for the study in Steps 1 and 2. For this reason, if there is a dimension of the content that you have a particular interest in but that may vary across selected headlines, it may be a good idea to complete a pretest to quantify this variability and ensure that you use headlines in your study that vary in the way that you want them to.

To make this advice more concrete, we will present one of our own pretests. The impetus for this pretesting is largely to match headlines on political partisanship. For example, if one set out, with good intentions but without a pretest, to select (say) 20 headlines that are Pro-Democratic and 20 that are Pro-Republican, it is quite possible that the Pro-Democratic headlines end up being *more* Pro-Democratic than the Pro-Republican headlines are Pro-Republican (or vice versa), depending on who selects them (or even random chance). If this is central to the study – for example, because the goal is to investigate whether political ideology correlates with whether people fall for fake news – then a larger set of headlines need to be put into a pretest to confirm their partisanship. This requires extra time and money, but it is better than running a study and wondering if an interesting set of results is because of idiosyncratic choices in what to present people in the study.

For the sake of illustration, and also because the actual data and content may be of use to people, we will now describe a recent pretest that we ran in the United States on both political and COVID-19 true (“real”) and false (“fake”) news headlines. For this, we collected 225 news headlines of the following categories: 30 false COVID-19, 49 true COVID-19, 70 false political, 76 true political. The full collection of headlines can be found here: <https://osf.io/xyq4t/>. The headlines are U.S.-centric, though many of the COVID-19 headlines may be appropriate for other countries. Importantly, the headlines are true and false *at the time that they were collected*. If you run a study, it is important to re-verify that the true/false labels continue to be accurate.

In past pretests, we have used an identical number across categories. In this case we just happened to collect this many good headlines and decided to run with what we had. Some of the false political headlines, in particular, are actually years old but have remained relevant (e.g., the racist “BLM Thug Protests President Trump With Selfie... Accidentally Shoots Himself In The Face” was fact-checked as false by Snopes.com on November 12th, 2016 but continues to be something that is representative of the type of false content that, unfortunately, is shared on the political right in the U.S. - and, in fact, the image used in the headline is from, at least, 2012: <https://www.snopes.com/fact-check/blm-thug-shoots-himself-taking-selfie-with-gun-in-protest-of-trump/>).

The study was run on Lucid, which is a polling firm that collects U.S. samples that are quota-matched to national demographics on age, gender, region, and ethnicity (Coppock & McClellan, 2019). We had a target of 2,000 participants. This large of a sample was needed in this case because

we only gave each participant 10 of the headlines (randomly selected from the full set³) – see below for further discussion of methodological considerations. We chose 10 headlines in this case because the study was completed online (as is true of every other study on fake news cited herein), so we tried to keep the length of the study as short as possible. Furthermore, participants were asked 9 different questions for each presented headline (from Table 1 and listed in order of presentation for the pretest: d, e, f, g, h, k, l, m, a). It is not necessary to ask so many questions, of course. Participants also answered several demographics questions and a couple of attention checks. Full materials are available on OSF: <https://osf.io/xyq4t/>. The pretest was completed using Qualtrics survey software – the ‘.qsf’ file is also available on OSF.

Table 1. *Some of the questions we have asked people in our past work on fake news and misinformation.*

a) If you were to see the above article on social media, how likely would you be to share it? 1) Extremely unlikely, 2) Moderately unlikely, 3) Slightly unlikely, 4) Slightly likely, 5) Moderately likely, 6) Extremely likely
b) To the best of your knowledge, is the claim in the above headline accurate? 1) Not at all accurate, 2) Not very accurate, 3) Somewhat accurate, 4) Very accurate
c) To the best of your knowledge, is the claim in the above headline accurate? 1) Yes, 2) No
d) What is the likelihood that the above headline is true? 1) Extremely unlikely, 2) Moderately unlikely, 3) Slightly unlikely, 4) Slightly likely, 5) Moderately likely, 6) Extremely likely
e) Assuming the above headline is entirely accurate, how favorable would it be to Democrats versus Republicans? 1) More favorable for Democrats, 2) Moderately more favorable for Democrats, 3) Slightly more favorable for Democrats, 4) Slightly more favorable for Republicans, 5) Moderately more favorable for Republicans, 6) More favorable for Republicans
f) Assuming the above headline is entirely accurate, how important would this news be? 1) Extremely unimportant, 2) Moderately unimportant, 3) Slightly unimportant, 4) Slightly important, 5) Moderately important, 6) Extremely important
g) To what extent is this headline anger provoking? 1) Not at all, 2) Slightly, 3) Somewhat, 4) Moderately, 5) Very much, 6) Extremely
h) To what extent is this headline anxiety provoking? 1) Not at all, 2) Slightly, 3) Somewhat, 4) Moderately, 5) Very much, 6) Extremely
i) How exciting is this headline? 1) Not at all, 2) Slightly, 3) Somewhat, 4) Moderately, 5) Very much, 6) Extremely
j) How worrying is this headline? 1) Not at all, 2) Slightly, 3) Somewhat, 4) Moderately, 5) Very much, 6) Extremely
k) How funny is this headline?

³ In other pretests, we randomly selected headlines from *within* the subsets, meaning that participants only (for example) saw false or true headlines. There are positives and negatives with both approaches: Selecting from within sets is justified if one only wishes to compare headlines from within the sets (e.g., by selecting only the most political false headlines and, separately, the most political true headlines). In this case, we fully randomized across the full set because judgments about, for example, how political the COVID-19 headlines are should be informed by the presence of other headlines that were selected explicitly to be political.

1) Not at all, 2) Slightly, 3) Somewhat, 4) Moderately, 5) Very much, 6) Extremely
l) How informative is this headline? 1) Not at all, 2) Slightly, 3) Somewhat, 4) Moderately, 5) Very much, 6) Extremely
m) Are you familiar with the above headline (have you seen or heard about it before)? 1) Not at all, 2) Slightly, 3) Somewhat, 4) Moderately, 5) Very much, 6) Extremely
n) Are you familiar with the above headline (have you seen or heard about it before)? 1) Yes, 2) Unsure, 3) No

To counteract low quality responding, we created an initial attention check and participants who failed were not allowed to continue the study. Specifically, participants were asked “Puppy is to dog as kitten is to?” and any response that didn’t approximate “cat” was considered an attention check fail. Of the 2,697 participants who made it past the consent screen, 474 failed this attention check (although some of these were genuine spelling errors that were not caught with our code) and a further 9 gave incorrect responses that were coded as correct with our code. Finally, an additional 201 participants quit at some point in the survey. This left us with data for 2,013 participants in total.

To maximize the amount of data that we used, we had participants choose whether they preferred the Democratic or Republican party ($N = 1107$ Democrats, 902 Republicans). A further 4 participants did not answer this question and were removed from the data set. In the resulting item-level report for the 225 headlines, each question has three columns: Democrat ratings, Republican ratings, and Overall ratings (equally weighing Democrat and Republican ratings). Due to randomization and missing data (skipped questions), the number of overall ratings that went into each mean for the item-level report varied. As an example, N ’s varied from 31-77 for Democrats ($Mean\ N = 49$) and 22-61 for Republicans ($Mean\ N = 40$) for the likelihood question. The item-level norm data and the full participant-level data set are both available on OSF: <https://osf.io/xyq4t/>. Note that this was put in an Microsoft Excel file for easy access and visualization. Caution is necessary if you intend on analyzing these data using Excel (the file can easily be converted to other formats for analysis).

We will now walk through the item-level data to facilitate use by other research groups. The excel file (see the above OSF link) has 5 tabs: 1) The full list of items listed in order of the corresponding item number in the participant-level data file (there is no meaning behind the order, apart from being alphabetical based on file names within the sets), 2) Only the COVID_Fake items, 3) Only the COVID_True items, 4) Only the Political_Fake items, 5) Only the Political_True items. The sets are separated into individual sheets because this facilitates the selection of items. Naturally, it is entirely up to the individual researcher to decide what characteristics to use for selection. Recall that we asked 9 questions for each item, which means that if you’re (for example) looking to run a study with particularly funny or anxiety provoking items, then you can use this data to find such items.

The data includes summaries of the actual content of the headlines (but check the image files for the full headlines) and columns for each of the 9 recorded variables. As noted, each numerical value represents the mean for either Democrats or Republicans (“Combined” is the mean of these

two means, placing equal weights on Democrat vs. Republican). Each variable also includes a “Diff” column which is the difference between the mean ratings for Democrats and Republicans; this is helpful for identifying points of disagreement in, for example, perceived likelihood/plausibility. [For the sake of curiosity, the false headline that had the greatest partisan disagreement was (verbatim) “THE ‘OBAMA FOUNDATION’ JUST BROKE ITS FIRST FEDERAL LAW” from “weaponstricks.com” – mean likelihood was 4.42 (slightly-to-moderately likely) for Republicans and 2.8 (slightly unlikely) for Democrats. The overall correlation, at the item-level, between likelihood ratings between Democrats and Republicans was nonetheless quite high, $r(223) = .60$.]

To demonstrate the process that one might use to select headlines, we have placed a secondary item-level data file that sorts and organizes the data based on partisanship. Imagine the goal of the study is to investigate whether people engage in motivated reasoning when being presented with false and true news headlines (Pennycook & Rand, 2019); for this, it would be important to select headlines that are politically contentious. This is the purpose of the partisanship question (‘e’ in Table 1), where ‘1’ means the headline, if true, would be quite positive for the Democratic Party whereas ‘6’ means it would be quite positive for the Republican Party. To select headlines, in this case, we use the combined Democrat/Republican score because the correlation between the partisanship ratings for Democrats and Republicans is quite high, $r(223) = .76$. We then simply sort the data based on this combined partisanship score.

The data file also includes a “baseline” calculation, which is the difference between the rating and scale midpoint. This allows for a direct comparison of *relative* partisanship across both sides of the aisle, as it were. We then, for simplicity, colour-code the items that are nominally below the scale midpoint for partisanship in blue (for Pro-Democratic) and those above the midpoint in red (for Pro-Republican). Of course, some of the items are quite neutral on this measure; for example, “Coconut oil’s history in destroying viruses, including Coronaviruses”, a false headline, was rated as a combined 3.45 on the scale (midpoint = 3.5).

On whatever measure that you happen to be interested in, it is important to keep in mind that it is likely that selecting the items that are highest/lowest on the measure will create imbalance. For example, the mean difference from baseline in partisanship for the five most Pro-Democratic true political headlines is 1.06 whereas the same difference for the five most Pro-Republican true political headlines is .82. If there was no pretest and we, somehow by sheer luck, selected the 10 headlines that are the strongest examples of Pro-Democratic and Pro-Republican headlines and put them in a study, we would nonetheless have headlines that are *more* Pro-Democratic than they are Pro-Republican. What this means, then, is if we want to balance on this element, we need to swap out some headlines.

As a tip, a tactic we use is to take the smaller value as a baseline (e.g., in this case, the true Pro-Republican headlines are at .82, which means we need to revise the selection of Pro-Democratic headlines so that the mean is also $\sim .82$ – and, ideally, we can find headlines that match the *distribution* of scores in the more constrained category). Our selection of five true political headlines that are equally partisan across party lines can be found in Table 2. This, of course, is only an example: we went with five headlines of each category for mere demonstrative simplicity.

You may want more headlines, or headlines of a different quality. Furthermore, these headlines are *only* matched on partisanship. Indeed, Republicans rated the five Pro-Republican headlines that were selected as more likely to be true (*Mean* = 4.37) relative to how Democrats rated the five Pro-Democratic headlines on likelihood (*Mean* = 3.95). By chance, importance was roughly similar (Democrats rating Pro-Democratic headlines, *Mean* = 4.24; Republicans rating Pro-Republican headlines, *Mean* = 4.26). Naturally, other variables differ as well – and it may not be possible to find a set of headlines that is perfectly matched across all the measured variables. It is up to the individual researcher to decide what variables are important to balance across (and there are, of course, elements that will differ that haven’t been measured here!). The excel file with the items from Table 2 placed in bold (“partisanship selection example”) can be found on OSF: <https://osf.io/xyq4t/>.

Table 2. An example set of true political headlines that are balanced on partisanship. A partisanship score of 1 = More favorable for Democrats, 6 = More favorable for Republicans

Partisan Lean	Headline	Combined Partisanship Score	Difference from Baseline
Pro-Democratic	“Top Democrats say postmaster general acknowledged new policies that workers say are delaying mail” (washingtonpost.com)	2.55	0.95
	“District of Columbia Sues Inaugural Committee For ‘Grossly Overpaying’ at Trump Hotel” (npr.org)	2.635	0.865
	“Facebook removes Trump ads with symbols once used by Nazis” (apnews.com)	2.66	0.84
	“Watchdog: ICE doesn’t know how many veterans it has deported” (nbcnews.com)	2.745	0.755
	“Republican anxiety grows as Democratic Senate challengers outraise incumbents” (cnn.com)	2.8	0.7
	<i>Mean</i>	2.68	0.822
Pro-Republican	“Trump signs ‘buy American first’ pharma executive order” (cnn.com)	4.2	0.7
	“Plant a trillion trees: Republicans offer fossil-friendly climate fix” (reuters.com)	4.205	0.705
	“Trump welcomes ‘The Walking Marine’ to White House” (apnews.com)	4.36	0.86
	“Black GOP House candidate praises Trump in convention speech” (apnews.com)	4.36	0.86
	“Trump gets endorsement of NYC police union, warns ‘no one will be safe in Biden’s America’” (nbcnews.com)	4.495	0.995
	<i>Mean</i>	4.32	0.824

Step 4 – Creating your experiment

We're now reaching the end of the practical utility of this guide as the way in which you create your experiment will depend on what you wish to experiment on. Nonetheless, we will offer some general tips and pointers.

We'll begin by returning to a point that we've repeated a few times now: The selection of your content is absolutely central to the validity of your study. The selection should be done with the experiment or study in mind. For example, in an early study (Pennycook & Rand, 2019), we used fake news to test a popular theory about motivated reasoning ("identity-protective cognition", see Kahan, 2013; Kahan, 2017), in which it is argued, contrary to much work in the dual-process tradition (Evans & Stanovich, 2013; Kahneman, 2011), that reasoning *facilities* partisan bias and higher cognitive capacities should therefore be associated with increased ideological polarization in terms of fake news belief. Contrary to this account, we found that people who are more reflective reasoners are better at discerning between true and false news content *regardless of whether it was consistent or inconsistent with their political ideology*. For this study, it was essential to have headlines that were highly political as those are the type of headlines that should facilitate partisan bias. We therefore used a pretest to confirm this (see above for an example of how).

How many headlines/pieces of content should be presented? Again, this will depend on your experiment. Our only suggestion here is to run smaller batches of participants initially to time out your experiment so that it is not too long or fatiguing. Naturally, the more headlines you have the more precision you will get. One thing worth considering is, instead of selecting some subset of headlines (e.g., 20 in total), one could randomize across a large set of headlines while keeping the number presented low. That is, the content for your experiment may be 40 headlines, but participants are only shown a random selection of 20. This will improve your ability to generalize (since your results will be based on twice as many headlines sampled from the world), although it will likely add some noise to your experiment. If you are committed to studying, for example, fake news *per se*, then this may be a good strategy. If you are simply using fake news to test other hypotheses (as in the motivated reasoning example above), then selecting 20 really good headlines may be a better strategy.

Another central question is what you plan on asking participants about. For reference, I've listed questions that we've used in the past in Table 1. One central consideration for studies of this nature is that the order that questions are asked is very important. For example, asking about accuracy changes subsequent judgments about (hypothetical) sharing on social media (Pennycook et al., 2021; Pennycook, McPhetres, et al., 2020).

In terms of what instructions to give people, this again will depend on your study. For generic studies where people are simply presented with a set of headlines and asked some sort of question(s) about them, we use variations on the following (depending on what people are asked – these instructions were used in the pretest outlined above): "You will be presented with a series of actual news headlines. There are [N] in total. We are interested in your opinion about the following: 1) Do you think that the headline is likely to be true? 2) Is the headline more favorable for Democrats or Republicans? 3) Is the headline important? 4) Is the headline anxiety or anger provoking? 5) Have you ever seen or heard about the news story before? 6) Would you consider sharing the headline on social media?"

A key question is where to recruit participants. we have used six sources in the past few years for online studies that we'll list here for simplicity: Amazon's Mechanical Turk (<https://www.mturk.com/>), Prolific (<https://www.prolific.co/>), Qualtrics panels (<https://www.qualtrics.com/research-services/online-sample/>), Lucid (<https://luc.id/>), Dynata (<https://www.dynata.com/>), and YouGov (<https://yougov.co.uk/>). Data quality considerations, such as obtaining nationally representative samples, are not unique to research on misinformation and there are many helpful resources on these topics elsewhere.

Step 5 – Analyzing the data

Your analyses will depend on your experiment. As a brief note, however, it has been common practice in cognitive and social psychology to collapse across items of the same category (e.g. true versus false headlines) and compare average responses across categories (e.g., an ANOVA predicting average accuracy ratings using the interaction between headline veracity and experimental condition). However, it turns out that this averaging approach can inflate the likelihood of false positives, as it obscures potential variation that exists across items (Barr et al., 2013; Judd et al., 2012). Thus, it is important to conduct the analysis at the level of the individual rating, rather than averaging. Psychologists typically do such analyses using multi-level models, with crossed random effects for subject and headline; economists typically use regressions with standard errors clustered on subject and headline.

Another issue relates to whether one focuses on “discernment” (otherwise referred to as “sensitivity” or “overall accuracy”) versus “bias” (otherwise referred to as “overall belief”) (Batailler et al., 2021; Pennycook & Rand, 2021), or both (e.g., Bronstein et al., 2019). This distinction is a classic issue in signal detection theory (Wickens, 2002): Someone who (for example) does not believe any fake news headlines may have the appearance of being highly discerning in their belief; however, if they also don't believe any *true* news headlines than they can be said to have a bias toward rejection or disbelief and are actually weak at discerning what is or it not true. In studies that have focused on analytic thinking (e.g., Pennycook & Rand, 2019), the tendency has been to focus on people's ability to discern between true and false news (i.e., having a large difference in belief between true and false news, calculated simply as true minus false news⁴). However, one may also be interested in the overall tendency for people to believe claims regardless of whether they are true or false (e.g., because of an interest in political consistency; see Batailler et al., 2021).

Other considerations

As with other areas of academic inquiry, it is important to maintain open science practices. In the case of fake news research, this will often involve sharing of materials that are taken from the internet. One may wish to discuss copyright issues with the relevant experts at their own institution prior to sharing materials (e.g., by putting them on the Open Science Framework's website, as we've done). Our understanding is that sharing screenshots of content via the OSF that could have

---

⁴ Note that if one wishes to compute d' (overall accuracy in signal detection theory; akin to “discernment” as used in-text), it is important to first center belief in true and false news belief before computing their difference. This ensures that both will be equally influential in the discernment score.

otherwise been shared openly on social media clearly falls into the realm of “fair use” in the United States or “fair dealing” in Canada, but please do not consider this legal advice. One benefit of storing materials on OSF is that they can be easily removed if issues emerge surrounding copyright. Finally, when publishing the screenshots in an academic paper (as we have done here), the journal may have policies on copyright that they will surely discuss with you.

Another issue to consider are the relatively unique ethical considerations that come with doing research on misinformation and fake news. First, it is important to be aware of the costs and benefits of showing participants examples of misinformation in studies. This is obviously necessary for us to learn about the psychology of misinformation, which is something that could produce a long term benefit for people. On the other hand, as discussed above, a single prior exposure to a fake news headline does increase later belief in the headline (Pennycook et al., 2018); thus, there may be consequences to exposing people to misinformation. One mitigation strategy is to debrief people at the end of the experiment by explaining which headlines were accurate. It is important to note that most people do *not* believe most of the fake news headlines that have been used in the published work so far (Pennycook & Rand, 2021). Still, it is possible that particular content will be especially convincing and impactful, and this is something that researchers need to bear in mind.

A second issue is that misinformation may target vulnerable groups or may be offensive in some way. Again, it is important to be mindful of what content people are being presented and one must avoid harm. Related to this, researchers who are tasked with themselves finding examples of misinformation for use in studies will likely encounter offensive content. This risk needs to be considered deeply when doing research of this type and mental health supports may be necessary.

A final consideration is that this guide largely focuses on misinformation spread on Facebook (e.g., by producing examples in a “Facebook format”; note that this is similar to Twitter). There are, of course, many other sources of misinformation (both on social media and otherwise), such as WhatsApp (Garimella & Eckles, 2020; Machado et al., 2019), TikTok (Basch et al., 2021), YouTube (Donzelli et al., 2018; Hussein et al., 2020), Reddit (Achimescu & Chachev, 2021), and Parler (Baines et al., 2021), among others. Nonetheless, many of the most substantial considerations considered herein will be relevant regardless of the source of the misinformation. In any case, investigating misinformation holistically without putting undue focus on particular forms or sources of misinformation is an important goal.

Conclusion

Online misinformation is not likely to go away, unfortunately. It is therefore absolutely necessary for social scientists to form a better understanding of why and how falsehood has spread so easily on social media platforms. Fortunately, as outlined in this practical guide, there are steps that can be taken to build robust surveys and experiments using “real world” examples of news content. There are challenges that come with doing research on a topic that is inexorably linked to current affairs, but if we can build a collective knowledge-base around the issue, perhaps it will not take a miracle for meaningful interventions to have an impact on the quality of content that people share on social media.

Acknowledgement

The authors gratefully acknowledge funding from the Ethics and Governance of Artificial Intelligence Initiative of the Miami Foundation, the William and Flora Hewlett Foundation, the Omidyar Network, the John Templeton Foundation, the Canadian Institutes of Health Research, and the Social Sciences and Humanities Research Council of Canada.

References

- Achimescu, V., & Chachev, P. D. (2021). Raising the flag: Monitoring user perceived disinformation on reddit. *Information (Switzerland)*, 12(1), 1–24. <https://doi.org/10.3390/info12010004>
- Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. In *Journal of Economic Perspectives* (Vol. 31, Issue 2, pp. 211–236). American Economic Association. <https://doi.org/10.1257/jep.31.2.211>
- Allcott, H., Gentzkow, M., & Yu, C. (2019). Trends in the diffusion of misinformation on social media. *Research & Politics*, 6(2), 205316801984855. <https://doi.org/10.1177/2053168019848554>
- Altay, S., de Araujo, E., & Mercier, H. (2020). “If this account is true, it is most enormously wonderful”: Interestingness-if-true and the sharing of true and false news. *PsyArXiv*, 1–29. <https://doi.org/10.31234/osf.io/tdfh5>
- Baines, A., Ittefaq, M., & Abwao, M. (2021). #Scamdemic, #plandemic, or #scaredemic: What parler social media platform tells us about covid-19 vaccine. *Vaccines*, 9(5), 421. <https://doi.org/10.3390/vaccines9050421>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Basch, C. H., Meleo-Erwin, Z., Fera, J., Jaime, C., & Basch, C. E. (2021). A global pandemic in the time of viral memes: COVID-19 vaccine misinformation and disinformation on TikTok. *Human Vaccines and Immunotherapeutics*. <https://doi.org/10.1080/21645515.2021.1894896>
- Batailler, C., Brannon, S. M., Teas, P. E., & Gawronski, B. (2021). A signal detection approach to understanding the identification of fake news. *Perspectives on Psychological Science*.
- Brashier, N. M., Pennycook, G., Berinsky, A. J., & Rand, D. G. (2021). Timing matters when correcting fake news. *Proceedings of the National Academy of Sciences of the United States of America*, 118(5). <https://doi.org/10.1073/pnas.2020043118>
- Bronstein, M. V., Pennycook, G., Bear, A., Rand, D. G., & Cannon, T. D. (2019). Belief in fake news is associated with delusionality, dogmatism, religious fundamentalism, and reduced analytic thinking. *Journal of Applied Research in Memory and Cognition*, 8(1), 108–117. <https://doi.org/10.1016/j.jarmac.2018.09.005>
- Clayton, K., Blair, S., Busam, J. A., Forstner, S., Glance, J., Green, G., Kawata, A., Kovvuri, A., Martin, J., Morgan, E., Sandhu, M., Sang, R., Scholz-Bright, R., Welch, A. T., Wolff, A. G.,

- Zhou, A., & Nyhan, B. (2019). Real Solutions for Fake News? Measuring the Effectiveness of General Warnings and Fact-Check Tags in Reducing Belief in False Stories on Social Media. *Political Behavior*, 1–23. <https://doi.org/10.1007/s11109-019-09533-0>
- Coppock, A., & McClellan, O. A. (2019). Validating the Demographic, Political, Psychological, and Experimental Results Obtained from a New Source of Online Survey Respondents. *Research & Politics*. https://alexandercoppock.com/papers/CM_lucid.pdf
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., & Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3), 201517441. <https://doi.org/10.1073/pnas.1517441113>
- Dias, N., Pennycook, G., & Rand, D. G. (2020). Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media. *Harvard Kennedy School Misinformation Review*, 1(1). <https://doi.org/10.37016/mr-2020-001>
- Donzelli, G., Palomba, G., Federigi, I., Aquino, F., Cioni, L., Verani, M., Carducci, A., & Lopalco, P. (2018). Misinformation on vaccination: A quantitative analysis of YouTube videos. *Human Vaccines and Immunotherapeutics*, 14(7), 1654–1659. <https://doi.org/10.1080/21645515.2018.1454572>
- Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8(3), 223–241. <https://doi.org/10.1177/1745691612460685>
- Fazio, L. K., Brashier, N. M., Payne, B. K., & Marsh, E. J. (2015). Knowledge Does Not Protect Against Illusory Truth. *Journal of Experimental Psychology. General*, 144(5), 993–1002. <https://doi.org/http://dx.doi.org/10.1037/xge0000098>
- Gabrielkov, M., Ramachandran, A., & Chaintreau, A. (2016). Social Clicks: What and Who Gets Read on Twitter? *ACM SIGMETRICS 2016*. <http://dl.acm.org/citation.cfm?id=2901462>
- Garimella, K., & Eckles, D. (2020). Images and Misinformation in Political Groups: Evidence from WhatsApp in India. *Harvard Kennedy School Misinformation Review*. <http://arxiv.org/abs/2005.09784>
- Guess, A. M., Lerner, M., Lyons, B., Montgomery, J. M., Nyhan, B., Reifler, J., & Sircar, N. (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*, 201920498. <https://doi.org/10.1073/pnas.1920498117>
- Guess, A. M., Nyhan, B., & Reifler, J. (2020). Exposure to untrustworthy websites in the 2016 US election. *Nature Human Behaviour*, 4(5), 472–480. <https://doi.org/10.1038/s41562-020-0833-x>
- Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 16(1), 107–112. [https://doi.org/10.1016/S0022-5371\(77\)80012-1](https://doi.org/10.1016/S0022-5371(77)80012-1)
- Hussein, E., Juneja, P., & Mitra, T. (2020). Measuring Misinformation in Video Search Platforms: An Audit Study on YouTube. *Proceedings of the ACM on Human-Computer*

- Interaction*, 4(CSCW1). <https://doi.org/10.1145/3392854>
- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69. <https://doi.org/10.1037/a0028347>
- Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, 8(4), 407–424. <https://doi.org/10.2139/ssrn.2182588>
- Kahan, Dan M. (2017). Misconceptions, Misinformation, and the Logic of Identity-Protective Cognition. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.2973067>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Lazer, D., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., & Zittrain, J. L. (2018). The science of fake news. *Science*, 9(6380), 1094–1096. <https://doi.org/10.1126/science.aao2998>
- Machado, C., Kira, B., Narayanan, V., Kollanyi, B., & Howard, P. N. (2019). A study of misinformation in whatsapp groups with a focus on the brazilian presidential elections. *The Web Conference 2019 - Companion of the World Wide Web Conference, WWW 2019*, 1013–1019. <https://doi.org/10.1145/3308560.3316738>
- McPhetres, J., Rand, D. G., & Pennycook, G. (2020). Character Deprecation in Fake News: Is it in Supply or Demand? *Group Processes & Intergroup Relations*. <https://doi.org/10.31234/osf.io/2agqt>
- Mosleh, M., Pennycook, G., & Rand, D. (2020). Self-reported willingness to share political news articles in online surveys correlates with actual sharing on Twitter. *PLOS One*, 15, e0228882. <https://doi.org/10.1371/journal.pone.0228882>
- Osmundsen, M., Bor, A., Vahlstrup, P. B., Bechmann, A., & Petersen, M. B. (2021). Partisan polarization is the primary psychological motivation behind “fake news” sharing on Twitter. *American Political Science Review*. <https://doi.org/10.31234/osf.io/v45bk>
- Pennycook, G., Bear, A., Collins, E., & Rand, D. G. (2020). The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Stories Increases Perceived Accuracy of Stories Without Warnings. *Management Science*. <https://doi.org/10.1287/mnsc.2019.3478>
- Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior Exposure Increases Perceived Accuracy of Fake News. *Journal of Experimental Psychology: General*. <https://doi.org/10.1037/xge0000465>
- Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., & Rand, D. (2021). Shifting attention to accuracy can reduce misinformation online. *Nature*, 1–63. <https://doi.org/10.1038/s41586-021-03344-2>
- Pennycook, G., McPhetres, J., Zhang, Y., Lu, J. G., & Rand, D. G. (2020). Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy nudge intervention. *Psychological Science*. <https://doi.org/10.31234/OSF.IO/UHBK9>

- Pennycook, G., & Rand, D. G. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188, 39–50. <https://doi.org/10.1016/j.cognition.2018.06.011>
- Pennycook, G., & Rand, D. G. (2021). The Psychology of Fake News. *Trends in Cognitive Sciences*, 1–29. <https://doi.org/10.1016/j.tics.2021.02.007>
- Pereira, A., Van Bavel, J. J., & Harris, E. (2020). Identity concerns drive belief: The impact of partisan identity on the belief and dissemination of true and false news. *PsyArXiv Working Paper*. <https://doi.org/10.31234/OSF.IO/7VC5D>
- van der Linden, S., Roozenbeek, J., & Compton, J. (2020). Inoculating Against Fake News About COVID-19. *Frontiers in Psychology*, 11, 2928. <https://doi.org/10.3389/fpsyg.2020.566790>
- Vegetti, F., & Mancosu, M. (2020). The Impact of Political Sophistication and Motivated Reasoning on Misinformation. *Political Communication*, 1–18. <https://doi.org/10.1080/10584609.2020.1744778>
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>
- Wickens, T. (2002). *Elementary signal detection theory*. Oxford University Press.