Believe It or Not!

The Impact of Exposure to Misinformation of Varying Levels of Perceived Truthfulness on Political Attitudes and Beliefs

Kevin James Berry
April 11, 2022

A Senior Thesis presented to the Faculty of the Princeton School of Public and International Affairs in partial fulfillment of the requirements for the degree of Bachelor of Arts

Acknowledgements

Writing this thesis proved to be an undertaking unlike anything else I have managed to produce erstwhile in my time at Princeton University. I believe some thanks are in order; in no particular order:

Thank you to my thesis advisor, Professor Markus Prior, for your advice and feedback throughout this process, as well as your tolerance of some of my, let's say, less-than-ideal academic habits.

Thank you to my junior paper advisor from the Spring of 2020, Professor Emily Pronin, for your continued guidance and getting me the head start I desperately needed in thinking about this topic, and this thesis.

Thank you to Princeton's School of Public and International Affairs for awarding me the funding to make this research possible, as well as all the support during my time at Princeton.

Thank you to you, the readers of this thesis. I can only hope you garner some enjoyment out of it.

Table of Contents

ABSTRACT

In recent years, misinformation—frequently referred to by the popular moniker “fake news”—has become increasingly relevant to the American political sphere, particularly in the wake of Russian interference in the 2016 presidential election, claims of fraud in the 2020 election, and vaccine hesitancy resulting from the circulation of COVID-19 misinformation. In an age where the ideological divide between the American political parties runs deeper than ever before, it was thought that exposure to fake news articles may result in an increase in polarization, and a decrease in trust in media and government, even if the people who read these articles do not believe the information contained within to be true.

To test this prediction, an experiment was run in which subjects were asked to listen to an audio clip that may be plausible or implausible and pro-Republican or pro-Democrat, or a control group unrelated to politics. It was expected that participants would be more willing to accept plausible politically congenial articles and dismiss implausible ones as obviously fake and thus harmless. Politically uncongenial articles, on the other hand, were expected to provoke feelings of resentment toward the opposing party.

Unfortunately, the manipulation check for the experiment failed. While the mean perceived accuracy for implausible audio was lower than plausible audio, the difference was too subtle to be statistically significant. However, all experimental audio was perceived as significantly less accurate than the control audio, allowing for some interpretation of the effects more generally.

These effects were limited. There was no statistically significant effect of exposure to the treatment groups on affective polarization, nor trust in media, nor political institutions such as the White House, Congress, or the election system. Most of the difference in scoring for these values was captured by control variables such as sentiment toward Former President Donald Trump or the social networking website Facebook, and perceived and actual political knowledge.

There was, however, a significant difference recorded in the presumed influence of misinformation. Participants rated partisans as significantly less likely to believe or share politically uncongenial misinformation, but just as likely to believe or share politically congenial information as truthful information.

This appears to suggest that the purported effects of misinformation are more perceived than actual. For this reason, more subtle interventions aimed at reducing effective spread, rather than raising awareness, may be advisable.

I. INTRODUCTION

> "I believe that nothing in the newspapers is ever true," said Madame Phoebus.
> "And that is why they are so popular," added Euphrosyne; "the taste of the age being so decidedly for fiction."
>
> *Benjamin Disraeli, Lothair (1870)*

Misinformation (false, inaccurate, or misleading information) and disinformation (a subset of misinformation that is created intentionally to deceive), often together referred to under the umbrella term of "fake news," have together begun to play an increasing role in the political sphere in the past few years. Though misinformation is not a new concept, this recent surge in interest began during the 2016 US Presidential election, where Russian interference may have contributed to Donald Trump's victory. During this cycle, fake news consumption was estimated to account for 6% of *all* news consumption on Twitter, which should serve as an indication of the sheer volume of misinformation flooding social media platforms (Grinberg et al.).

Since then, the influence of fake news has only increased. In the leadup to the 2020 election, troll farms—"professionalized groups that work in a coordinated fashion to post provocative content, often propaganda, to social networks"—were reaching 140 million Americans each month on Facebook (Hao). These groups are intentionally provocative and even a form of strategic foreign policy—one former employee of Russia's most well-known troll farm, the Internet Research Agency, summarized their mission statement thus: "Our task was to set Americans against their own government, to provoke unrest and discontent" (Sheth).

Interestingly, the primary goal of these Russian attacks is not necessarily to promote fake information as true. They aim not to prove a falsehood, but rather to cast doubt on the truth (Pierre, "Illusory Truth, Lies, and Political Propaganda" 2). In doing so, the campaigns would "exploit societal fractures, blur the lines between reality and fiction, erode trust in media entities and in democracy itself" (Bail et al.). By hastening political polarization in the United States through exacerbating divisive issues in its social media campaigns, Russia hopes to weaken the relationship between the American electorate and its government and alter the public's voting behavior—and believes this can be done with misinformation that people may even recognize to be false (Bail et al.).

Indeed, in the Information Age, playing the information (or misinformation) game appears to be a prerequisite for being a global superpower, or for those vying for geopolitical hegemony. China's so-called "50 Cent Army" is a suspected group of as many as 2 million people hired by the Chinese government for "astroturfing" campaigns, posting large amounts of Internet comments defending the government's side in political debates as though they were the genuine opinions of everyday Chinese citizens. It has been estimated that this group "fabricates and posts about 448 million social media comments a year" in an enormous cheerleading operation for the Chinese Communist Party (King et al.). Clearly, an investment of time and money of this size indicates a real belief in the measure's efficacy.

Unfortunately, the United States is not exempt in this regard. The source of misinformation can often be domestic rather than foreign, produced and spread by partisans within the United States. Following the 2020 election, for example, Donald Trump's endorsement of false claims of election fraud cast serious doubt on the results of this democratic process, culminating in a riot at the Capitol on January 6, 2021 by a number of his supporters

who believed the election had been rigged or stolen. Even for Americans not rioting, research shows that exposure to misinformation can lower trust in media (Ognyanova et al., "Misinformation in Action"), increase affective polarization between the political parties (Suhay et al., "The Polarizing Effects of Online Partisan Criticism"), and damage satisfaction in democracy (Nisbet et al., "The Presumed Influence of Election Misinformation on Others Reduces Our Own Satisfaction with Democracy").

Although it is often politicized, misinformation and its contents are not always inherently political, per se. Although it has been incorporated into the larger culture war, misinformation related to the COVID-19 pandemic—particularly regarding its severity and spread, as well as the efficacy and purported purpose of masking and vaccination efforts—is a different breed of misinformation related to medical science. Indeed, this flood of misinformation is so prevalent that the World Health Organization has called it an "infodemic" with serious implications for global health.

Even some of the more implausible conspiracy theories have gained traction, such as the idea that COVID vaccines contain tracking microchips that connect to 5G—a claim that 20% of Americans rated as "probably" or "definitely" true in one survey, and that is linked to over a hundred arson attacks on telecommunications equipment in the United Kingdom (Schoolov). Additionally, the Pew Research Center reported that over 70% of Americans had heard of the conspiracy theory that the COVID-19 outbreak was intentionally planned by powerful people, and 25% thought this to be "probably" or "definitely" true—a number that rises to 37% when looking only at conservative Republicans, an increase likely related to the claim that this was done by a "Deep State" to decrease satisfaction with Donald Trump's presidency in the leadup to the 2020 election (Schaeffer). It is worth noting that this particular conspiracy theory is tightly

linked to a 26-minute long documentary-style video entitled “Plandemic” that was widely circulated on YouTube and Facebook in May 2020, prior to its removal (Fichera et al.).

Increase in the belief of false claims is not the only impact here, however. It translates into real-world behavior. Exposure to pandemic-related misinformation on social media has been linked to increased non-compliance with masking and social-distancing, as well as increased vaccine hesitancy, a result that has directly prolonged the duration and spread of COVID-19 (Bridgman et al.). The main takeaway here: fake news, misinformation, and disinformation on social media, beyond simply propagating misperceptions, also have a direct, measurable effect on the actions of the American people.

While these are only some of the most high-profile instances of misinformation, the reality is that encountering false information online has become an everyday occurrence, and one with very tangible and far-reaching consequences. One of the central tenets necessary to a functioning democracy is to have a well- and *accurately*-informed electorate. A 2019 survey from the Pew Research Center suggests that Americans themselves are aware of this: half of the respondents to the survey rated made-up news as a critical problem in the country today; 79% of those surveyed agreed that steps should be taken to restrict made-up news (Mitchell et al.). Due to these effects, it is important to identify effective ways to combat misinformation. This study aims to better understand some aspects of the effects of misinformation so as to be better equipped to combat it.

The hypothesis of this study is as follows:

H_1: The effects of misinformation are maintained (in whole or in part) regardless of whether the consumer perceives the misinformation as accurate or not.

This will be tested by running an experiment on Amazon's Mechanical Turk (MTurk) platform. MTurk workers who choose to participate in the study will be assigned to one of five groups. There will be four treatment groups plus a control group. Each participant will answer a few demographic questions, listen to a short (1 minute) audio clip of a neutral voice speaking generated by a text-to-speech tool (an audio "deep fake"), then answer some more questions about their political attitudes and beliefs. The audio clip for the control group will be a neutral clip unrelated to politics. The audio clips for the experimental groups will be a fake news clip coded as of one of the four following types: Plausible and Pro-Democrat; Implausible and Pro-Democrat; Plausible and Pro-Republican; Implausible and Pro-Republican. This data can then be analyzed to see whether it supports the hypothesis, and provide greater insight into more nuanced effects of misinformation.

While there are many scholarly sources available about misinformation, many are longitudinal observational studies tracking incidental Internet activity, or else more epidemiological in nature; that is, focused more on how misinformation spreads. There are comparatively fewer experimental designs examining misinformation and its effects. Better understanding the effect misinformation has on people, and how that effect is reached, can provide insight into how to better counteract it.

Specifically pertaining to this study, if many of the effects of misinformation are independent from the misinformation's perceived truthfulness and rather inherent to some other quality of misinformation—as this experiment aims to determine—the current prevailing practice

of identifying misinformation and flagging it as false would not be the most effective countermeasure. A result of this sort from the study would indicate other approaches may be more useful, such as a narrative corrective that explains the motivation behind the creator of the misinformation.

II. LITERATURE REVIEW

> "It is a melancholy truth, that a suppression of the press could not more completely deprive the nation of its benefits, than is done by its abandoned prostitution to falsehood […] The man who never looks into a newspaper is better informed than he who reads them; inasmuch as he who knows nothing is nearer to truth than he whose mind is filled with falsehoods & errors."
>
> *Thomas Jefferson, Letter to John Norvell (1807)*

EXPOSURE & ASSIMILATION

What happens following exposure to misinformation? How is it processed and stored by the brain? What role might partisanship play in this processing?

Biased Assimilation and Attitude Polarization

Increased polarization, which is seen in today's political climate in the United States, reasonably entails more strongly-held opinions among partisans. These stronger opinions can in turn create a feedback loop and further increase polarization itself. Lord, Ross, and Lepper (1979) found that people tend to derive higher value from, and take greater stock in, information that supports opinions that they already hold, even if these opinions are based on little more than preconceptions and assumptions. In their experiment, two groups, one in favor of and one against the death penalty, were exposed to the same objective, relevant, empirical—and most importantly, mixed—evidence on the matter.

This inconclusive evidence, rather than leading to more moderate opinions in each group, led to the opposite: participants did not process the data impartially, instead seizing upon the

evidence that supported their viewpoint while searching for flaws to justify eschewing that which went against it, emerging from the experiment with increasingly polarized views (Lord et al.).

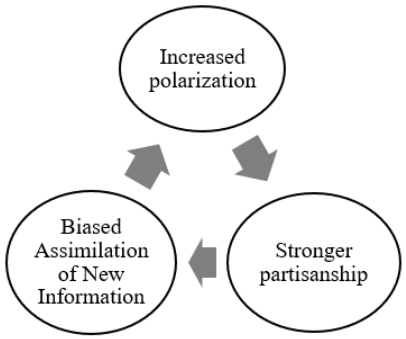

Figure 1: Diagram of the Feedback Loop that occurs due to Biased Assimilation in the modern American political environment

However, Lord et al. qualify their findings, acknowledging a normative issue at hand. The sin committed by their subjects was using evidence processed in a biased manner to make their hypotheses unfalsifiable. But when an "objective truth" is known, studies that reflect this truth (for example, corroborating the speed of light) are reasonably given greater weight than studies that do not (Lord et al.).

The Role of Media & Group Influence

Consider situations where a group's biases conflict with an "objective truth"—as in the case of conspiracy theorists, like flat-earthers or anti-vaxxers, or consumers of fake news. Are these the result of the same mechanism of biased assimilation, or another psychological phenomenon entirely? Tribalism may be to blame. A 2020 Hungarian study "found that partisanship predicted belief in political fake news more strongly than conspiracy mentality"

(Faragó et al.). This is a troubling finding; if partisanship alone is the most significant predictor of susceptibility to fake news articles, and not a rarer condition like conspiracy mentality, then it follows that large swaths of the population are therefore vulnerable to disinformation campaigns.

The conclusion reached in the Hungarian study aligns with Geoffrey Cohen's findings regarding the dominating impact of group influence on political beliefs. As determined in Lord et al., people process information differently based on a subjective lens—often a partisan lens. Perception is colored by biases, ego, and prior beliefs. Party alignment is included in these determinants. In Cohen's study, participants, regardless of the actual content of the policies with which they were presented, sided with the one that they were told their party supported. The subjects derived social meaning based off what they were told about a group with which they identified, and this heavily influenced their opinions on policy matters. Even more concerning is the revelation that participants were unaware of this influence exerted on them (G. L. Cohen).

There are a number of problems related to media encountered by individuals. There is also the hostile media phenomenon, derived from biased perception. Partisans tend to view even-handed presentation of an issue as biased against their own side; they are wont to believe that media gives their own side too much scrutiny, while giving the opposition too little scrutiny (Vallone and Ross). This is to say, charges of media bias are often unfounded, stemming from a desire for preferential treatment for one's viewpoint and the imputation of conflicting motives on the media (Hastorf and Cantril).

What about cases when media actually is hostile? The obvious example here is disinformation, the dissemination of information that is incorrect by intent (as contrasted with misinformation, information that is incorrect by accident). Disinformation, depending on the perceptions, attitudes, and beliefs of the person exposed to it, can demonstrate remarkable

persistence and resistance to correction, making it all the more dangerous as a weapon of propaganda (Lewandowsky, Stritzke, et al.).

In the age of information, disinformation is increasingly abused in this manner. Controlling information is a critical component of national strategy; "victory will be defined more in terms of capturing the psycho-cultural rather than the geographical high ground" (Lewandowsky, Stritzke, et al.). Indeed, disinformation has already been implemented toward this end. One study found that during the 2016 election cycle, "5.0% of aggregate exposures to political URLs were from fake news sources" (Grinberg et al.).

Inattention, Repeat Exposure, & the Illusory Truth Effect

Consumption of social media is often referred to as a "mindless" activity. Scrolling through one's feed or timeline is often done in conjunction with other tasks. It is an everyday activity, characterized by the same limitations as many other everyday undertakings: done while busy, distracted, on autopilot. Monitoring of cognitive operations is relatively light in this state, and reading posts on social media is done with an uncritical eye and often goes unquestioned, meaning that information seen and assessments made of their truthfulness may appear sensible enough, but would not hold up under more discerning inspection. Consumers on social media, as a result, are highly susceptible to manipulation. Bad-faith actors who produce disinformation and disseminate it on social media are aware of the distracted nature of social media use, and exploit it to their advantage.

The natural assessments made by social media users while browsing serve as the foundation for the later application of heuristics: intuitive procedures and strategies, deliberate or not, that rely on natural assessments to reduce a complex task of estimation or prediction to a

simpler judgmental operation. The availability heuristic, in particular, helps to explain the danger of exposure to misinformation on social media. In using the availability heuristic, people judge the probability or frequency of something by the ease with which instances come to mind. These biases are borne from vividness and retrievability from memory. An obvious conclusion to draw from this is that media exposure makes certain claims more "available" than others, and thus perceived as more likely, even if this perception goes against the reality. For example, a study by Johnson and Tversky found that subjects exposed to purported newspaper coverage on a subject vastly overestimated its incidence, rating, for example, death from fire as more likely than death from diabetes, when in actuality the inverse was overwhelming true. They further found that the correlation between people's estimates of these frequencies and the frequencies in news coverage was almost .9, indicating high correlation (Johnson and Tversky). Many users who are exposed to misinformation on social media face this exposure numerous times, making the claims posited more available to them.

Logic suggests that this relationship between news coverage and perception would hold regardless of the truth of the claim put forth by said coverage. Research agrees with this conclusion, supporting the idea of an illusory truth effect, which is the tendency to believe false information to be correct after repeated exposure. Arguably more importantly, this effect may hold even for headlines that contradict an individual's party affiliation, and even when people have prior knowledge that they are false (Pierre, "Illusory Truth, Lies, and Political Propaganda").

Some studies suggest that repetition is not even a necessary component of this: one found that even a single exposure to a fake news headline pulled from Facebook was enough to increase subsequent perceptions of its accuracy, both within the same session and in a session a

week later; and this effect held despite the low level of believability of the claims, even when the stories are labeled as contested by fact-checkers, or when it is uncongenial to the reader's political ideology (Pennycook, Cannon, et al.).

Misinformation that is perceived as truthful, it follows, is more likely to be shared, which in turn increases the number of individuals exposed and the overall propagation of the false claims, in a manner that could be considered epidemiological.

The Presumed Influence of Misinformation

The same concept of being unaware of influence is expressed in the third person effect: the idea that individuals overestimate the influence that mass communications and media have on the attitudes and behavior of others, while simultaneously underestimating its effect on themselves. The third person effect is applicable to responses to fake news articles as well: most people are confident in their ability to discern fake news, while dubious of that ability in others (Calvert).

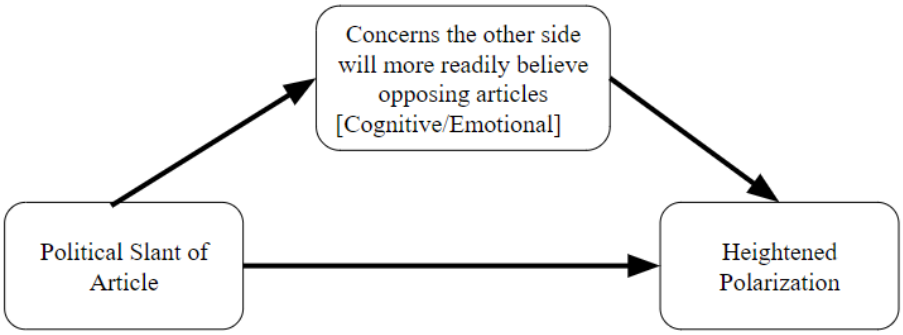

Figure 2: Diagram showing the expected relationship between the experimental variables and the mediator of expected degree of belief across party lines

A 2020 survey found that Democrats and Independents who paid more attention to political news had higher presumed influence of misinformation (PIM) in others. It also found that those (regardless of political identification) who believed that misinformation influenced others more than themselves were significantly less satisfied with American democracy (Nisbet et al., "The Presumed Influence of Election Misinformation on Others Reduces Our Own Satisfaction with Democracy"). This is a severe outcome, and just one of many possible impacts of exposure to misinformation!

IMPACTS

Why does exposure to misinformation matter? What impacts has exposure been suggested to have on political attitudes and beliefs?

Belief in False Claims & Belief Perseverance

Unsurprisingly, exposure to misinformation can increase a consumer's belief in the article's false central claim; interestingly, a study found that this increase was especially apparent among Democrats reading a politically congenial article (Guess et al.).

Prior studies have demonstrated that "beliefs can survive the complete subtraction of the critical formative evidence on which they were initially based," and that "entrenched beliefs can also survive the addition of non-supportive evidence" (Lord et al.). Research into belief perseverance suggests that attitudinal shifts may remain even following explicit revelation that articles read by consumers of media contained misleading, false information: "initial beliefs may persevere in the face of a subsequent invalidation of the evidence on which they are based, even when this initial evidence is itself weak and inconclusive" (Anderson et al.).

As frivolous as some of the claims put forth in articles of misinformation may seem, even the most absurd and unsubstantiated—or even disproven—can still have some far-reaching and long-lasting effects. Consider, for instance, the assertion that former president Barack Obama is set to claim Kenyan citizenship to avoid being tried for treason in the United States. This allegation, baseless and illogical as it is, was in fact made by radical right-wing conspiracy theorist Mark Taylor (Mantayla). Indeed, the "birther" movement that it stems from, which claimed that Obama was born in Kenya rather than Hawaii and therefore ineligible for the presidency, found significant reach despite being promoted primarily—at least initially—by fringe theorists; a CNN poll conducted in 2010 found that over a quarter of Americans surveyed

had doubts as to whether the president was in fact born in the United States (*CNN Poll*). This was two years *after* a copy of his official Hawaiian birth certificate had been released, pre-election, in 2008. It was not until his long-form birth certificate was released in April of 2011 that this proportion of "birthers" dropped significantly, to roughly 10% (J. Cohen). To further illustrate the scope of the birther movement, it is worth acknowledging that one of its most prominent promoters was Donald Trump, who would later succeed Obama as President of the United States (Barbaro).

Affective Polarization

Over the past several decades, polarization has increased markedly in the American electorate. Republicans and Democrats are more ideologically divided now than at any previous point in recent history. Overlap of ideals between the two parties has shrunk considerably, and as a result partisan antipathy has grown in kind. Significant portions of each party hold highly negative views of the opposing side—in 2014, 38% of Democrats viewed the Republican Party very unfavorably, while 43% of Republicans viewed the Democratic party very unfavorably. These negative sentiments often extend to the point of perceiving the policy agenda of the opposition as a threat to the nation's well-being (Dimock et al.).

The ramifications of this growing divide are far-reaching. Capitol Hill faces increasing gridlock, and political negotiations frequently end in stalemate, as both liberals and conservatives see an ideal agreement as, perhaps unsurprisingly, one in which their side gets more of what it wants, typically at the expense of the other party.

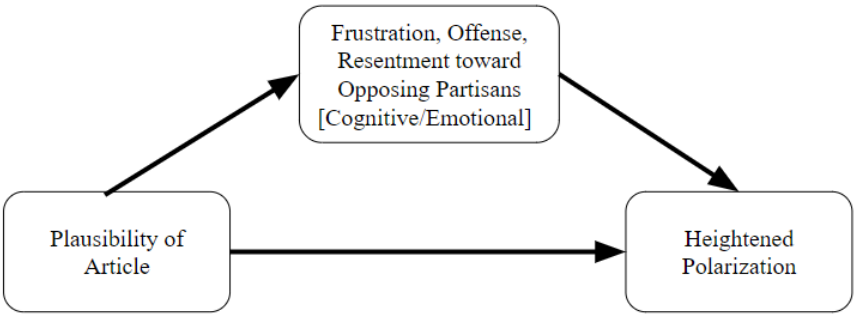

Figure 3: Diagram showing the expected relationship between the experimental variables and the mediator of negative feelings toward opposing partisans

Online partisan criticism (and uncivil discourse) contributes to affective and social polarization. In a 2018 study, participants read a news article discussing an issue debate between the two parties, then read reader comments containing vitriolic partisan criticism (e.g., calling Republicans "greedy bastards" or Democrats "Moochers"). Afterwards, participants were asked to rate Democrats and Republicans on a Feeling Thermometer. The study found that partisan criticism that is derogatory toward political opponents increases affective polarization (Suhay et al., “The Polarizing Effects of Online Partisan Criticism”).

Most misinformation is partisan in nature, and often very caustic. For this reason, it is not surprising to observe increased polarization as a result. It functions in the same way as invectives that do carry some truth. Following this logic, this would also suggest that both plausible and implausible misinformation would accomplish the same result.

Trust in Media & Government

A 2018 longitudinal survey utilizing a browser extension keeping record of online behavior; thus, serving as a view of organic, rather than laboratory exposure to misinformation—found that online

misinformation was widely linked to lower trust in mainstream media. Interestingly, it also found an overall increase in trust in political institutions (including the White House, Congress, Supreme Court, military, justice system), which was comprised of a decline in trust for strong liberals and an increase in trust for moderates and conservatives (Ognyanova et al., "Misinformation in Action"). Worth noting, however, is that most fake-news content is right-leaning, and that Republicans controlled all branches of government at the time of the survey. A potential explanation they give for this asymmetry is that liberals exposed to right-leaning misinformation reject it and mistrust the Republican government, while conservatives exposed to the same right-leaning misinformation take it at face value and increase their trust in the Republican-held political institutions. A similar line of reasoning to this contributed in part to the hypothesis of this experiment: misinformation feels like a bigger, more serious deal if it serves to advance a political stance that you do not agree with.

INTERVENTIONS

What are current methods employed by social media companies for dealing with the propagation of misinformation on their platforms?

Warnings

Current interventions to the spread of misinformation that have been implemented by social media platforms primarily focus on fact-checking efforts, labeling disputed posts as such. These warnings comprise the bulk of social media companies' strategies for combatting misinformation.

Although it has been suggested that the presence of such warnings might lead to a backfire effect, where users share and believe flagged headlines more through some sort of spite response, most studies agree that warning labels work, at least for users posting because they believe the claim (Rivero). It is also worth noting here, however, that some more partisan-motivated users may post even with the knowledge that what they are sharing is false; this proportion has been estimated to be as high as 30% of all fake news sharers (Chadwick et al., *The Amplification of Exaggerated and False News on Social Media*).

So, while they do work to some capacity, these warnings are an incomplete intervention, and they fall victim to the very same attention deficit that leads to the spread of fake news in the first place: invasive warnings hamper the user experience and may lead to frustration; non-invasive, smaller warnings run the risk of going unnoticed by a distracted consumer. While this effort is rooted in psychological research—providing "warnings at the time of the initial exposure" has been identified as an effective way to combat disinformation—it is far from perfect (Lewandowsky, Ecker, et al.). Warnings decrease in effectiveness after repeated exposure if not used sparingly, and when one considers that most disinformation is encountered by the

same users repeatedly, warnings may prove to be ineffective for the group most susceptible to disinformation.

More important, however, is the acknowledgment that the algorithms that detect and flag fake news are imperfect. The determination of the criteria for how to separate fake news from true news is vulnerable to the same problem that plagues all of signal detection theory: while we can try to eliminate the overlap between these two distributions, error is inevitable. Due to the sheer volume of content posted to social media sites, manual moderation is impossible and false positives and false negatives are a surety. False positives, flagging true posts as false, would decrease the apparent value in the warnings as perceived by readers; it makes it obvious that they do not work as intended. However, there is risk in false negatives (false posts that are not flagged as such) as well: the existence of warnings that users are aware of can lead to the belief that anything not flagged has been tacitly verified. This "implied truth effect" can lend additional (unfounded) legitimacy to misinformation that has fallen through the cracks of the detection algorithm (Pennycook, Bear, et al.).

Trusting users to discern true news from fake news, in the current environment, is not a viable solution either. People's assessments of their own ability to recognize misinformation are characterized by overconfidence; their confidence in this ability is not well correlated with their actual accuracy. In general, people believe others to be more susceptible to misinformation than themselves; there is a high presumed influence of misinformation in the population (Baek et al.). Part of this overconfidence can be attributed to systematic factors, particularly a lack of feedback. Corrections to misinformation have been shown to spread less than the initial misinformation itself (Song et al.). As a result, users may never learn that they have been

exposed to or even have shared false information. In the absence of this knowledge, users have difficulty being more aware of their own fallibility, as well as calibrating their media sensibilities to better identify inaccuracies they may encounter in the future.

Therefore, a more effective countermeasure agenda against misinformation would improve upon these factors: less reliance on warning systems that are inherently, by design, imperfect; increased insulation from bad actors who intentionally spread disinformation; and increasing attention paid by users to news headlines and informing them of their own vulnerability in the face of this problem.

For those who will still share misinformation anyway, outside of labeling it as such, Facebook has introduced a countermeasure that harnesses the concept of choice architecture. The tactics used in this measure could be effectively implemented on other platforms that have yet to adopt it. Rather than removing certain fake news from feeds entirely, or merely putting a warning label on it, Facebook has chosen to alter the way fake news appears in a user's feed.

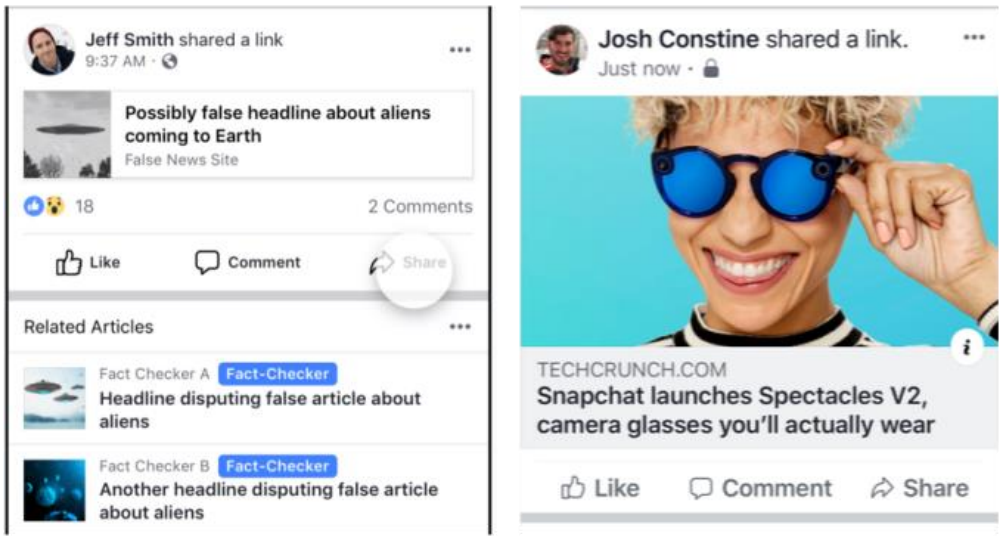

Figure 4 - A comparison of how fake news and true news may be displayed on Facebook (Constine)

News determined to be true is displayed much more prominently than fake news. It is given a full-size thumbnail and larger text size. Fake news, in comparison, is relegated to a much smaller space, with a thumbnail one-tenth of the size (and thus much less eye-catching), and a headline of a smaller text size. Additionally, the fake news post is accompanied by related headlines that directly contradict and debunk its false claim, meaning that users who still see it will be introduced to a corrective at the very same time ("Facebook Shrinks Fake News after Warnings Backfire"). This subtle, yet ingenious change makes it so that fake news is much less likely to be seen than real news—users are more likely to simply scroll past it in their feed, or pay less attention to it, or see it alongside disproving information. Even someone who might have reposted the link before knowing that it was false, may now unknowingly scroll past it due to its less prominent display and never get the chance to do so.

Inoculation Theory, or "Pre-Bunking"

Ultimately, however, such algorithms cannot entirely solve the problem, as was discussed earlier. This is why inoculation theory is a growing strategy for combatting misinformation. In essence, inoculation theory states that "if the mind is exposed to weakened doses of disinformation, it is later able to recognize it and therefore handle the information in a different way" ("Bad News"). In practice, this may take the shape of games such as *Bad News*, developed as a collaboration between researchers at Cambridge University and a Dutch organization against the spread of misinformation. The game puts players into the role of a fake news-monger, along the way teaching them the tactics used by real-world developers of fake news and exposing them to anodyne examples, so they can recognize real ones later. The same group responsible for *Bad News* also has two other games: *Harmony Square*, which highlights the use of disinformation as

a way of fomenting unrest; and *Go Viral!*, which teaches players about COVID-19 misinformation. These sorts of narrative correctives—understanding why you may have encountered fake news in the first place, as well as the tactics used to manipulate you—are some of the most effective ways to stop misinformation before it spreads, via "pre-bunking." Making them short and entertaining, like these games are, can help increase their reach. Facebook already has a Games section on its website. Adding these inoculation games to the collection and advertising them to users could help equip them with better tools to recognize and prevent misinformation before further propagation.

Counternarratives

One of the most effective methods of combating disinformation and correcting misbelief is by telling "an alternative story that fills the coherence gap otherwise left by the retraction" (Lewandowsky, Ecker, et al.). Correcting disinformation may be most effective when victims of misbelief are provided with an explanatory narrative of why they were presented with false information in the first place. From this, it follows that one of the most effective methods of combating, for example, Russian disinformation campaigns, may be to increase public awareness of the campaigns themselves and the rationale behind them.

THE PRESENT EXPERIMENT

It has been observed that there is a myopic focus of literature on the technological dimensions of disinformation and the technological dimensions of combating its distribution and propagation, at the expense of understanding of the human factors and psychological dimensions at play (Nisbet and Kamenchuk). The present experiment, then, seeks to help fill this gap in understanding the human and psychological dimension, looking at whether the observable effects are directly tied to the inaccuracy of misinformation, or whether another aspect of misinformation might be responsible. Gaining a better understanding of these factors could help policymakers and fact-checkers be better equipped to deal with disinformation swiftly and effectively, to slow or even halt its spread.

The present experiment is explained in more detail in the Data and Methodology section of this paper.

III. DATA & METHODOLOGY

"Don't believe everything you read on the Internet."

Abraham Lincoln, oft misattributed ("QuotePark")

Are the effects of misinformation on political attitudes and beliefs (e.g., affective polarization, trust in media and government) contingent on belief in the false claims?

H_1: The effects of misinformation are maintained (in whole or in part) regardless of whether the consumer perceives the fake news as truthful or not.

EXPERIMENTAL DESIGN

This thesis aims to approach an answer to the above research question by testing the above hypothesis via an experiment.

There are four experimental conditions and one control condition, to which subjects are randomly assigned based on which of four fake news (or in the case of the control, one real news) audio clips they are given to listen to. The available clips contain content presented in a style that is intended to be perceived as either Plausible or Implausible, and endorse a viewpoint that is either Pro-Democrat or Pro-Republican. The experiment is in the form of a between-subjects, 2x2 design (Implausible/Plausible X Pro-Republican/Pro-Democrat).

For the full transcripts of the audio clip stimuli, see Appendix B. In an attempt to minimize confounds, it was decided that it would be best to keep the same general topic for each of the experimental stimuli. For the same reason of ensuring parity, they were kept roughly the same length—1 minute long—and had very similar paragraph structure. When picking this

general topic for the stimuli, COVID initially seemed like it could be a good choice for subject matter, given its current relevance at the time of writing.

However, it ended up proving difficult to find a liberal analogue to something like the conservative conspiracy that the pandemic was planned by the "Deep State" so as to ensure Trump's loss in the 2020 election (Jamison et al.). Making up a conspiracy was briefly considered—something like claiming Trump planned to expose the ailing Ruth Bader Ginsburg to COVID-19, based on a real joke made by Republican Congressman Paul Gosar—but thought that it would hurt the external validity to have one stimulus sourced from a real-world conspiracy theory that has significant traction and another cobbled together just for this study and thus heretofore unheard of (Caralle). Besides that, there was also some uncertainty as to whether something like vaccine hesitancy, for example, would track as neatly along the political aisle as other subjects might.

Given this, I opted for an election-related theme, which I felt was still relevant and valuable, but had more examples to draw from on both sides of the aisle. Initially, the "implausible" stimuli were more on the level of conspiracy theory and the "plausible" ones were more run-of-the-mill, incorrect information, such as incorrect quote attributions, that would not as readily arouse suspicion to an untrained or undiscerning eye.

However, we realized that with this model, the content of the treatments involved entirely different claims for plausible and implausible versions. The key comparison of the research—do people react differently when they believe, as opposed to doubt, the veracity of a claim—is between these two versions. They were therefore altered to be more similar, to rule out as many confounds as possible by virtue of randomization. The claims were kept from the implausible versions, and then a few sentences were changed in an attempt to make them appear more

plausible. Based on findings about trust in news, the plausible versions were a bit more professionally worded, less obviously biased, and cite some supposedly “credible" experts to back up their claims (“What Makes News Sources or Topics Trustworthy?”). The implausible versions, by contrast, are more aggressively partisan, informal, meandering, and make no external appeal to ethos.

The claims that are made in the four experimental groups are, to be clear, unsubstantiated. There has not been credible evidence supporting the claim that Donald Trump is a cultivated Russian agent, as is claimed in the Pro-Democrat stimuli; nor the claim, as is made in the Pro-Republican stimuli, that the 2020 presidential election was rigged by an insidious Deep State (Cunliffe; Watts; Evon).

Subjects were recruited for the study via Amazon’s Mechanical Turk (MTurk) Marketplace. Further detail about this procedure can be found in the Participants section of this report.

The dependent variables—affective polarization, trust in media, and trust in government—were measured via survey questions answered by participants after listening to their assigned audio clip. Further detail about this procedure is discussed in the Methods section of this report.

USE OF DEEPFAKE AUDIO IN STIMULI

An additional unique aspect of this experiment is the use of deepfake audio to present the stimuli. Most existing experiments about misinformation administer the experimental treatment as an altered tweet or print news article excerpt. Opting for the output of a synthetic speech tool was a way to provide a unique, interesting, and innovative delivery mechanism of the misinformation. I suspected that people are a little more wary of, or have wizened to, how easily plaintext can be manipulated, whereas they might not be as suspicious of something like this—or if they are, it could work as a way of evading some demand characteristics, as people might think the experiment is about deepfakes specifically rather than misinformation more generally.

The online text-to-speech vocoder artificial intelligence "FakeYou" was used to create the stimuli (*FakeYou*). This website had modules to produce audio files mimicking a wide array of recognizable voices, ranging from actors such as Morgan Freeman, to cartoon characters such as SpongeBob SquarePants, to famous politicians—which is part of why it is an important field to study.

Initially, I intended to spoof the voice of conservative political commentator Ben Shapiro for the Pro-Republican articles, and *Last Week Tonight* host John Oliver for the Pro-Democrat articles, assuming these to be recognizable voices that could pair with the political slant of the respective stimuli in a believable way. However, it was decided that it was not really important to the experimental design for subjects to recognize the voice, and that it would introduce unnecessary heterogeneity to the sample for some to recognize the voice and others to not.

An alternative, to eliminate this heterogeneity, could have been to use better known voices. The website had the technology to imitate Donald Trump, Joe Biden, and Barack Obama, to name a few, who any participant in the study would surely have heard speak before.

However, making a deepfake of a former or sitting president, especially if it was rather convincing, might do real harm to participants that debriefing may not neutralize. Since the deep fake of a known voice is not central to the research design, this option was not the right choice, simply for the ethical concerns it raised.

Ultimately, the same voice was used for all five audio clips. That voice was of the late television and radio host, Larry King, who was politically involved but not aggressively partisan (Lloyd).

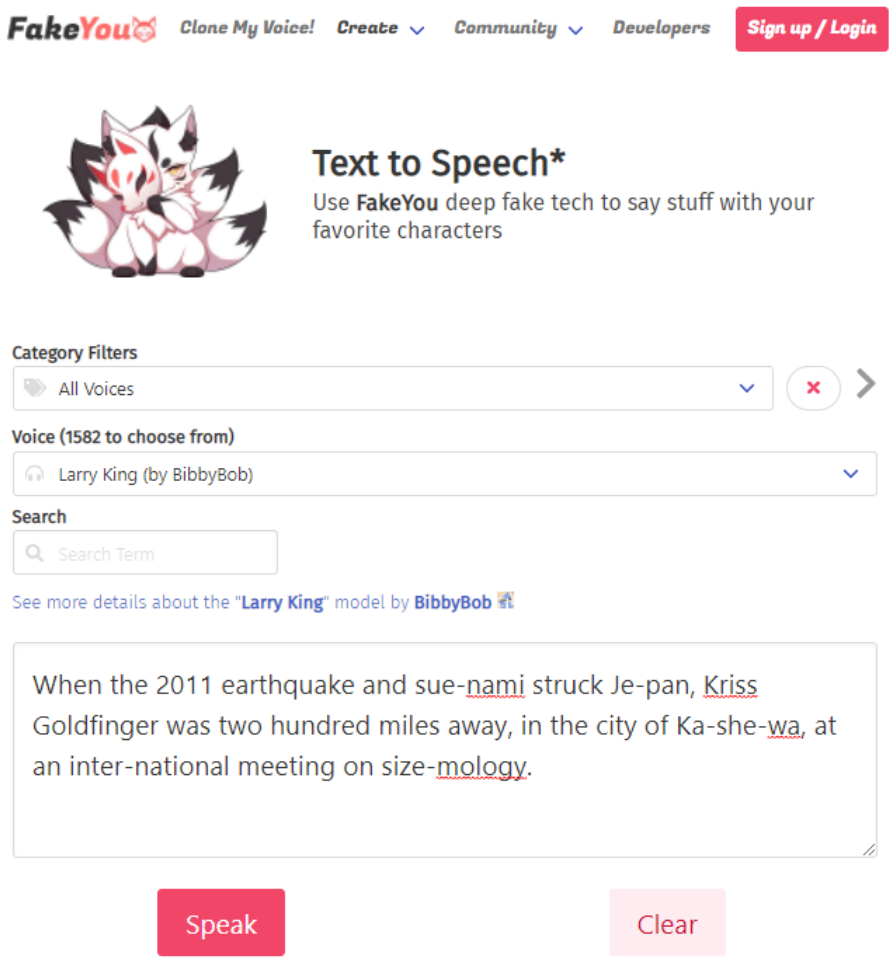

Figure 5: A screenshot of FakeYou.com, the text-to-speech tool used to create the deepfake audio for the stimuli.

Creating the audio clips was relatively simple, even for someone with none of the requisite training or knowledge to create such a tool themselves. The script for each stimulus was fed into the site's text entry box line by line, as it handles small snippets of text better. These outputs were then stitched together into a single audio file using the audio-editing software Audacity. Occasionally, the text fed into the website was altered to be spelled more phonetically, so that it could be processed better by the text-to-speech tool— "tsunami," for example, became "sue-nami" to help with parsing.

The resulting audio clip is not perfectly humanoid at all times, but can be eerily convincing at points. Captions were also included in-time with the audio, to promote understanding during potential audio imperfections, as well as increasing attention to the stimulus by using more than one sensory input. A 2020 study on video deepfakes notes that due to the realism heuristic, individuals are more likely to make false perceptions from misleading audiovisual stimuli than from misleading text, and that while participants were more likely to be uncertain about its authenticity than outright misled, exposure to the deepfake still led to reduced trust in political news on social media (Vaccari and Chadwick).

PARTICIPANTS

A total of 430 subjects were recruited from Amazon's Mechanical Turk (MTurk) online crowdsourcing marketplace. The use of MTurk allows for a more diverse pool of subjects in demographic categories such as age or geographic location than could be found if participants were to be sourced from the university undergraduate population (McDuffie). It also allows for quick data collection on a large scale—several workers can complete the survey at the same time, using their own computer and Internet connection, without assistance or monitoring by the researcher. Data collection was completed in merely a few hours.

As is the nature of MTurk, participants were self-selecting workers—the study employed a convenience sample (Gerlich et al.). The study appeared in a list of Human Intelligence Tasks (HIT) that workers on the platform were able to choose to participate in based on the displayed title and reward (i.e., payment amount). The study title participants saw was "Answer a survey about United States politics and media (5-10 min)," with the description "For those eligible to vote in the US (not required to be registered): Listen to a short audio clip (1 minute) and answer questions about your beliefs and attitudes regarding politics and media. Please actually listen to the audio clip! It is not long!"

Subjects were asked to be voting-age Americans; this was the only selection criterion. They did not need to be registered to vote. To help ensure that this condition was met, the HIT was constructed to only be open to workers whose location was in the United States. MTurk does not enable United States citizens residing outside the United States to use the platform, so this was not a population that was excluded by the use of this restriction—they were already, unfortunately, being excluded by the platform itself (*Amazon Mechanical Turk Worker Help*).

MTurk also requires that its workers be at least 18-years-old, so it should not have been possible for people who were below the voting age to participate in the study, either.

Risk associated with participation was minimal, and participants were compensated monetarily for their participation, receiving $3.34 after their completion of the survey. Though this compensation rate may appear low upon first glance, it is consistent with the highest state minimum wages for hourly work and significantly higher than the mean level of compensation given for most Human Intelligence Tasks on MTurk, which studies have suggested could be lower than $2 per hour (Tompkins and Swift). The suggested pricing given by MTurk itself is only $0.10 to $0.20 per minute.

The average elapsed time for workers in this study between accepting the HIT on MTurk and submitting it back to the researcher was around 12 minutes, which would equate to $16.70 per hour. Many workers completed the task in significantly shorter time frames, as well: an aggregation of twenty-two reviews of the HIT on the MTurk worker forum TurkerView placed the average equivalent wage garnered at $53.36 per hour, due to the very low completion times recorded by those users (*Kjberry | TurkerView MTurk Requester Profile*).

Mechanical Turk collects 40% in fees on top of the compensation to the workers, bringing the total spending per participant to $4.68.

Following completion of the survey, participants were debriefed and informed that the audio clips they listened to (if they were in an experimental group) contained false information; they were apprised of some resources that fact-check, discuss, and combat misinformation that they may consult if they wish to learn more or investigate further. For additional questions or concerns, they were able to contact the researcher or Institutional Review Board via email. For the full text of the debrief, see the end of the survey in Appendix A.

Additionally, one current theory of combatting misinformation is "pre-bunking" (as compared to debunking), or inoculation: giving people the tools to identify fake news to help stop it before it goes viral. By showing participants examples of misinformation, and then explaining during debriefing that it was misinformation, participants can become better at spotting and questioning fake news the next time they encounter it in the wild. Research by the Cambridge Social Decision-Making Lab suggests that this sort of short inoculation--though it does decay over time--can reduce susceptibility to misinformation for a period of months (Lewsey). In the best case, participating in the survey could, in theory, provide an additional benefit to the participants beyond the monetary compensation.

The demographic makeup of the participants who were recruited for the study is discussed, in part, in the Data Cleaning & Limitations section.

MEASURES

For the full text of the survey used, see Appendix B. The survey was designed and administered using the survey builder on the Qualtrics XM Platform.

Prior to exposure to the stimuli, participants were asked to provide a self-reported rating of their political knowledge, and then asked a series of political knowledge questions to assign an objective rating. These two measures could then be combined to create a measure of a subject's overconfidence. The questions of the political knowledge battery were modeled off those used to measure political knowledge by the American National Election Studies, or ANES ("ANES 2019 Pilot Study Questionnaire Specifications"). Previous studies seemed to indicate that these questions were sufficiently difficult to adequately stratify the knowledge levels of MTurk workers, even though they tend to be more knowledgeable than the general population (Berinsky et al.). They were also asked at this time to provide a party identification using the ANES branching question format, which generates a 7-point scale ranging from 1, a Strong Democrat, to 7, a Strong Republican.

To ensure that the experimental manipulations were successful, participants were asked to rate how accurate they perceived the claim made by their audio clip to be, in order to check that the articles of the "implausible" experimental condition were rated less believable than those of the "plausible" condition and that the independent variable of perceived truthfulness had indeed been properly manipulated. This gives insight as to whether the stimuli actually induced variation in whether subjects believed the claims. Placing this check before the dependent variable measures was determined to be problematic because it might prime subjects to think about believability, but putting it at the end could interfere in other ways. As a compromise, the

placement of the manipulation check was randomized to appear evenly before and after the measures of the dependent variables.

To measure the dependent variables, participants were asked a number of questions. They were asked to rate a number of political leaders and institutions on a "Feelings Thermometers." The ratings of institutions like the White House, Congress, and the election system could be used as a measure of the participant's trust in government. The difference between the ratings for the Democratic and Republican parties could be used to create a measure of affective polarization (Guess et al.). Media trust was measured using questions borrowed from the Pew Research Center, which ask about trust in information sourced from national news organizations, local news organizations, and social networking sites ("2021 Pew Research Center's American Trends Panel - Trust in Media").

To assess a potential causal source of any shifts in attitude—that is, possible mediators—subjects were asked to indicate how likely they perceived members of both political parties to believe the claim made in the audio clip, as well as how likely these partisans would be to share the audio clip.

Finally, participants were asked a few demographic questions: their gender, race, state of residence, and where they were born.

Several questions were also included for data cleaning purposes. These questions existed to facilitate the removal of inattentive participants during data analysis. Participants were asked to indicate once they have finished listening to their assigned audio clip in order to proceed with the survey, and asked again to verify that they did not experience any technical difficulties in doing so. An attention check asked the participant to identify the person who was the main subject of the audio clip from a list of four choices, which should not have presented any sort of

problem for those who actually listened to their assigned stimulus. One final question was included at the end, asking the participant to admit if they had responded to questions by random clicking at any point during the survey.

DATA CLEANING AND LIMITATIONS

Data collection was fast and painless, a benefit of using Amazon's Mechanical Turk (MTurk). Within a few hours of posting the Human Intelligence Task (HIT), several hundred participants had been recruited, and their responses recorded. However, the data that was collected also needed to be cleaned extensively before it could be used in analysis. All MTurk workers were compensated for completing the survey regardless of the content of their answers, but the quality of a significant number of the responses was, frankly, a bit disappointing.

There were 470 responses recorded by Qualtrics before the survey was closed. This was higher than the 430 subjects who were recruited on MTurk, but the majority of these extra responses were from workers who began, but did not complete the HIT. First, partial responses were removed—any response that Qualtrics recorded as having been started but not finished, with progress less than 100%. This resulted in 439 remaining responses.

Next, respondents who had not participated faithfully in the political knowledge battery were removed. This included both respondents who rejected the request beforehand to answer the questions to the best of their ability without outside sources because it was "very important to our research," as well as respondents who admitted to cheating when asked after the questions if they had used outside sources. While only three participants outright refused the request for integrity ahead of time, an astoundingly high amount fessed up to cheating after the fact: 97 respondents.

This is baffling for a number of reasons. First, each of these 97 respondents had agreed to avoid asking others for help or looking up the answers ahead of time, and then reneged on that agreement. It had been made clear that they were not expected to know the correct answers, and

stressed that their honest effort was critical to the experiment's success. Second, there was nothing to be gained from performing better on the political knowledge battery, besides perhaps a sense of superiority from scoring higher, or a fulfillment of one's own curiosity. Certainly, there was no financial incentive or bonus. This dovetails nicely with the third reason for confusion: it takes more time to look up the answer than to just guess. For an MTurk worker, whose primary goal on the platform is typically to complete as many HITs as possible in as short a time as possible, so as to garner as much money as possible, taking the time to look up the answer to a bit of political trivia seems to be somewhat an odd decision.

Some of this may have been due to the random clicking that surely plagues any online survey. A higher proportion may have been due to not carefully reading the questions and answer choices. The desired response to the question before the political knowledge battery ("Do you agree to answer the following questions to the best of your ability without help from outside sources?") was "Yes," as in "Yes, I agree to answer without help from outside sources." By contrast, the desired answer to the after question ("In answering the last few items, did you look up any of the questions online or ask someone for help?") was "No," as in "No, I did not use any outside sources." A participant paying particularly little attention may have assumed they were meant to respond "Yes" again and defaulted to selecting that answer choice. Regardless, it was deemed that both explanations, either inattention or a demonstrated inability to faithfully engage with the study, constituted grounds for removal.

Next, one participant who indicated they had not successfully been able to listen to the audio clip that they had been assigned was removed. Nineteen participants who admitted to responding randomly at some point during the survey were also removed.

Finally, respondents who could not successfully identify the subject of their assigned audio clip were removed. This identification should have been an unfailingly simple task for anyone who had actually listened to even part of the clip they had been asked to. Each of the experimental groups' audio clips shared the same subject—Donald Trump—and name him within the first 20 seconds of their (already short) one-minute runtime. The control group's subject, Seismologist Chris Goldfinger, is named within 3 seconds! Even so, nearly 10% of all survey respondents responded that the main subject of their assigned audio clip was Former President Barack Obama, despite him not being named, referenced, or alluded to at any point of any of the audio clips. Twenty-two additional participants were removed here for failing this attention check, bringing the final count of respondents that had passed through all these cleaning procedures to 297.

Interestingly, the cleaning procedure removed subjects somewhat asymmetrically across groups. The control group was left with 48 remaining respondents, the Implausible Pro-Republican group was left with 60, and the rest each had 63.

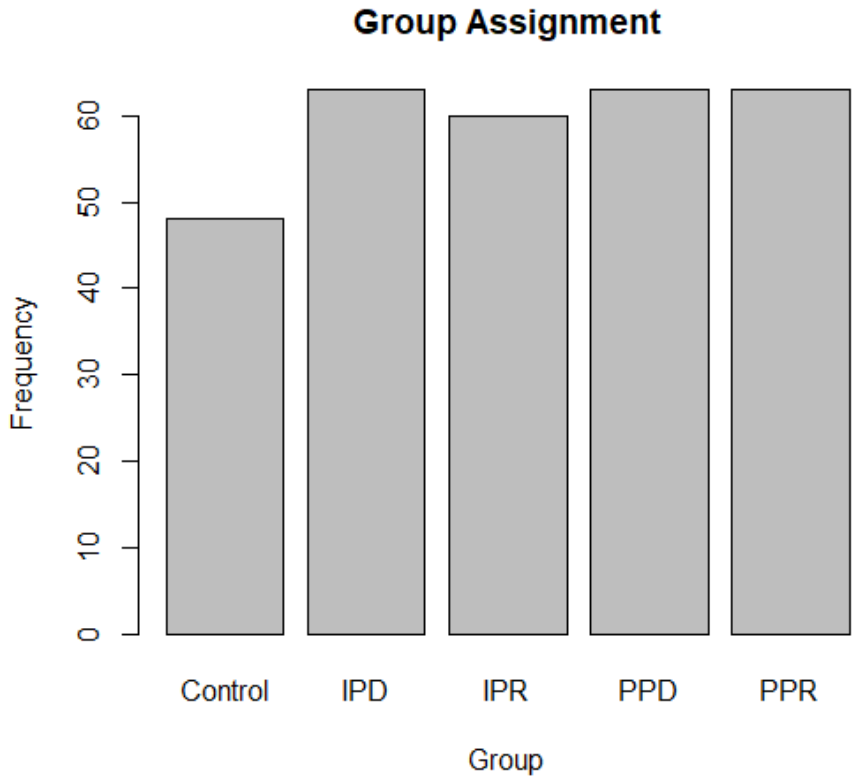

Figure 6: Bar graph showing the distribution of study participants across groups following data cleaning. The control group ended up noticeably smaller.

This asymmetry is despite Qualtrics assigning study participants to each group evenly. The lower number of remaining Control subjects, specifically, likely points to an inadequacy of the screening question used in this attention check. Participants knew that they would be participating in a politically-themed survey. The HIT appeared to them in MTurk with the title "Answer a survey about United States politics and media (5-10 min)." In their pursuit of completing the HIT as quickly as possible, some workers assuredly will have skipped listening to the audio entirely and have checked the "I'm done listening" box immediately to move on with the survey. The view counts on the stimuli after the survey was first administered corroborated this assumption, as they were initially all slightly lower than the number of subjects assigned to each group.

Following this premature advancement, workers would have then been presented with the attention check, asking the participant to identify the main subject of their assigned audio clip from four possible options: two politicians (Former Presidents Obama and Trump) and two persons unrelated to politics (one being the seismologist discussed in the control clip).

Unsure of what the correct answer was (due to not actually engaging with the stimulus), a participant in this situation would be likely to roll the dice between the two most politically-relevant choices and hope they had picked the right one; notably, the option unrelated to politics that was not the correct answer for the control group was chosen by zero respondents. This coin-flip tactic failed in choosing both Obama and Trump for subjects who had been assigned to the control category, hence the higher number of failures of the attention check for that group. However, roughly half of these random guessers would have succeeded in choosing Trump for the experimental groups, despite not having listened to the audio at all.

This suggests that the true number of faithful participants for the IPD, IPR, PPD, and PPR conditions should be closer to the 48 remaining in the control. Combining this difference in group sizes and the 10% of survey respondents who answered Obama for the attention check, it can be deduced that roughly 20% of the participants remaining in each of the treatment conditions did not listen to their audio clip, but merely had to good luck to guess correctly. Unfortunately, these respondents are indistinguishable from those who participated faithfully, and ergo impossible to remove from analysis.

Inattention appeared to be a common theme among several of the MTurk workers. A handful complained that they had not received their unique completion code to enter into MTurk so as to prove their participation and be awarded payment, despite it being at the top of the debriefing page in **bold font**, indicating that they may have been skimming for much of the

survey. Two such workers lodged this complaint on TurkerView, a forum used by anonymous MTurk workers to identify worthwhile HITs to complete as a community. These two workers, in their reviews, also both logged completion times of well-under five minutes, suggesting that, in a survey containing a one-minute video, they had not spent much time reading at all. This hastiness was a not uncommon theme of the reviews. While all praised the study for being straightforward with above-average pay, one user still commented, "Cons: 1-minute timed video" (*Kjberry | TurkerView MTurk Requester Profile*).

At any rate, the data that passed through the cleaning process revealed the same limitations as a sample that is typical of MTurk. While the ideal sample for a political survey is one that is perfectly reflective of the American electorate, MTurk as a platform reflects a group that is not completely representative of the general population. It skews younger (around 30 years old), white, over-educated, and liberal (DeSoto).

These traits were reflected in the data collected in this study. The vast majority of participants (82.5%) identified as White, which is much higher than the 61.6% of the United States population recorded by the 2020 Census. By contrast, the proportion of participants who identified as Black, for example, was only 6.4%, only slightly above half the 12.4% recorded by the 2020 Census (Jones et al.).

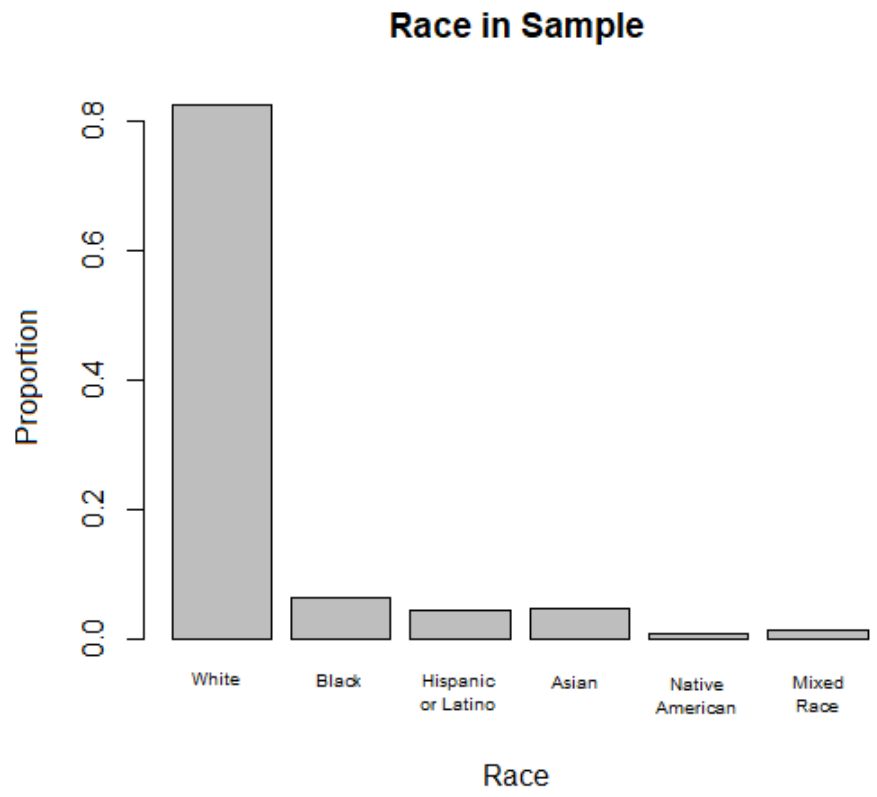

Figure 7: Bar graph depicting the proportion of participants' race in the sample of data collected

The gender and age of participants were closer to representative. The sample skewed somewhat male (59%), but not overwhelmingly so. The median age of the sample was 36, which is close to that of the United States population as a whole—2020 estimates put the population median at 38.6 years old ("Median Age of the U.S. Population 2020"). It is worth noting here, however, that the median age of the electorate tends to be older than that of the population; in 2019, the median age among all registered voters was 50 years old (Gramlich). This is understandable when factoring in the minimum voting age and the higher likelihood of older citizens to be registered to vote; voter turnout increases with age as well, with 76.0% of voters age 65-74 voting in the 2020 presidential election, compared to only 51.4% of 18-24 year-olds

("2020 Presidential Election Voting and Registration Tables Now Available"). When considering the characteristics of MTurk for workers—an online, gig-economy platform that is not the most user-friendly—it is not all that surprising that the userbase would lean more techno-literate, and thus also younger.

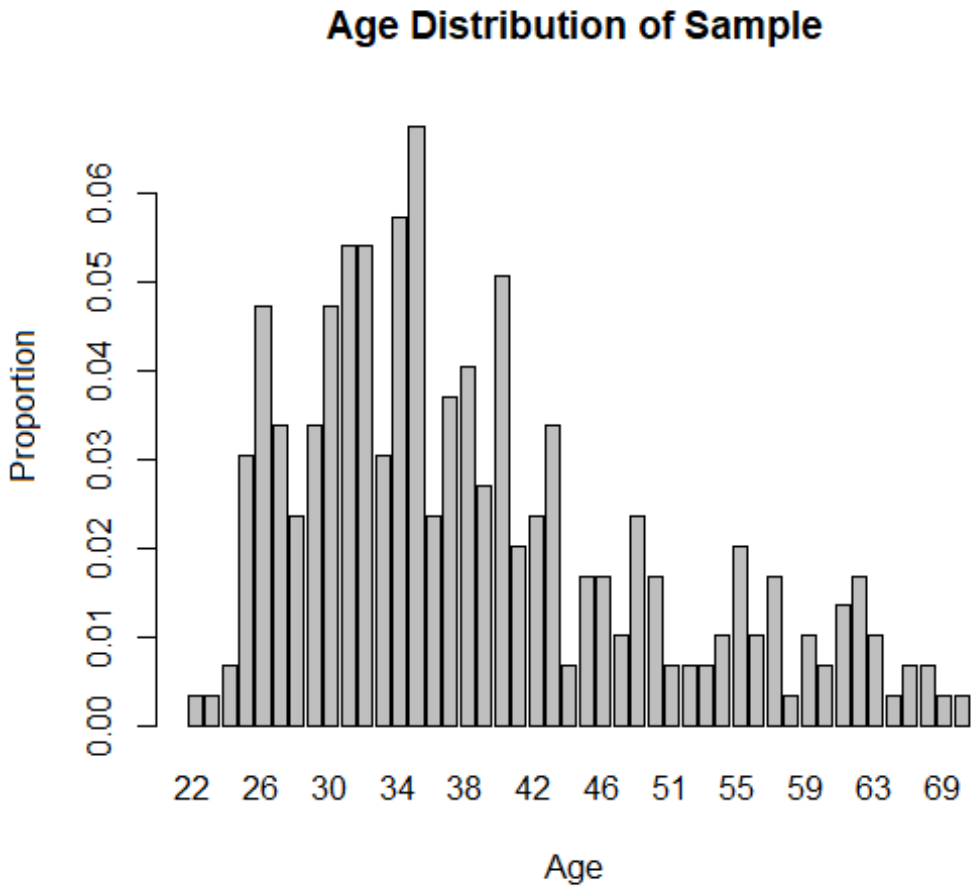

Figure 8: A bar graph depicting the distribution of participants' age in the sample. The mean age was 38.8, while the median age was 36.

The education level of the sample was rather high. More users reported having a 4-year college degree (57.2% of respondents) than any of the other options combined. In comparison, a 2019 Pew Research Center survey showed that only 36% of registered votes in the United States

had a college degree. Only 10% of survey participants reported having only a high school education or less, compared to 32% in the 2019 Pew poll (Gramlich).

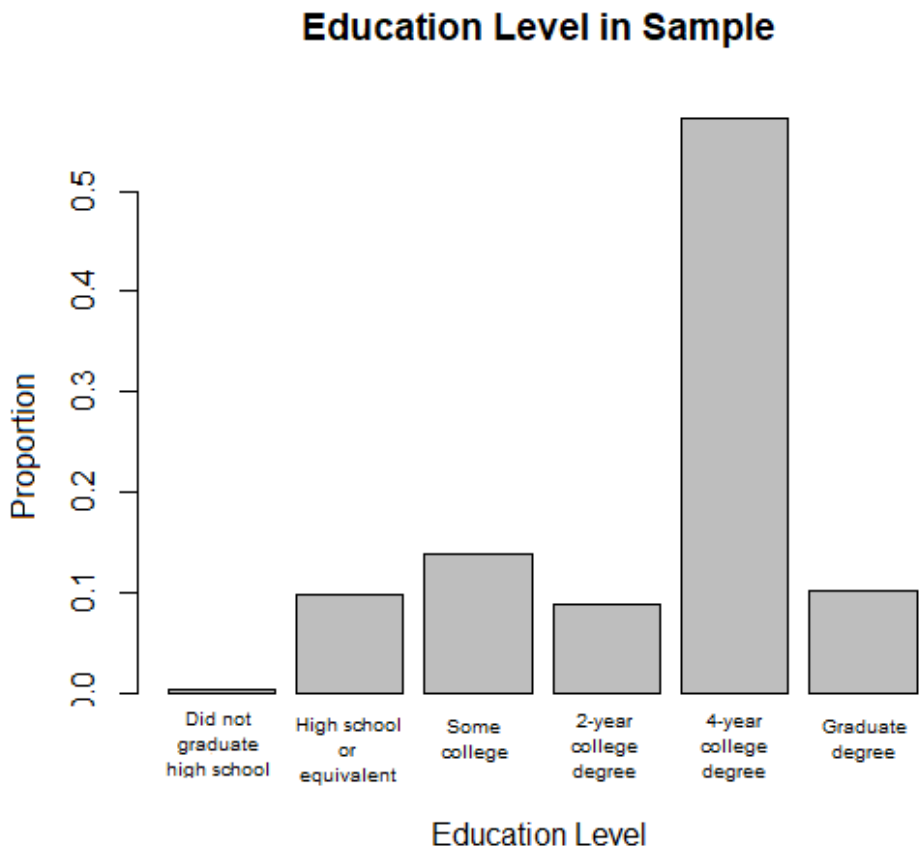

Figure 9: A bar graph showing the distribution of educational attainment among the sample. Overwhelmingly, participants had at least some college experience.

The sample also skewed fairly heavily Democrat. Two-thirds (66.3%) of respondents identified as Democrat or lean-Democrat, as compared to only 30% who identified as Republican or lean-Republican. While Democrats do have a slight edge in their proportion of registered voters, it is nowhere near this one-sided. Pew reported 49% of all registered voters in 2019 were

Democrat or lean-Democrat, compared to 44% of registered voters who were Republican or lean-Republican.

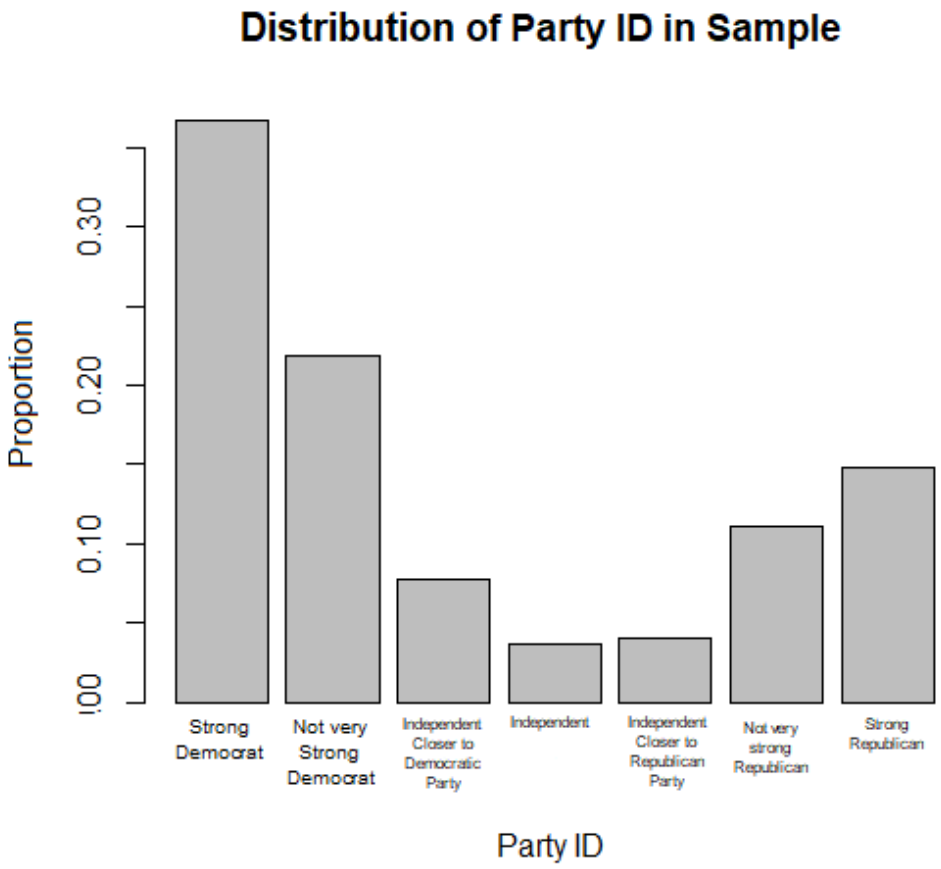

Figure 10: A bar graph showing the distribution of party self-identification among participants. The sample notably leaned heavily Democrat.

These quirks of the demographics of the sample altogether point to another limitation of the MTurk data, particularly as it relates to this experiment. Research has shown that right-wing voters are more frequently exposed to disinformation than left-wing voters (Grinberg et al.). It has also been shown that older, technologically illiterate American adults are the most susceptible to misinformation campaigns (Grinberg et al.). Regrettably, this means that the exact

subset of the population most-commonly found on MTurk (and in this sample)—Democrat, college-educated, younger, technologically literate—is likely to yield participants that may be the least susceptible to the effects of misinformation, and best equipped to recognize it from the start.

IV. RESEARCH FINDINGS

"So much Fake News. Never been more voluminous or more inaccurate. But through it all, our country is doing great!"

Donald Trump, Twitter, now removed (2018)

MANIPULATION CHECK

First, the manipulation check was performed to determine whether the implausible conditions had been measurably perceived as less believable than the plausible conditions.

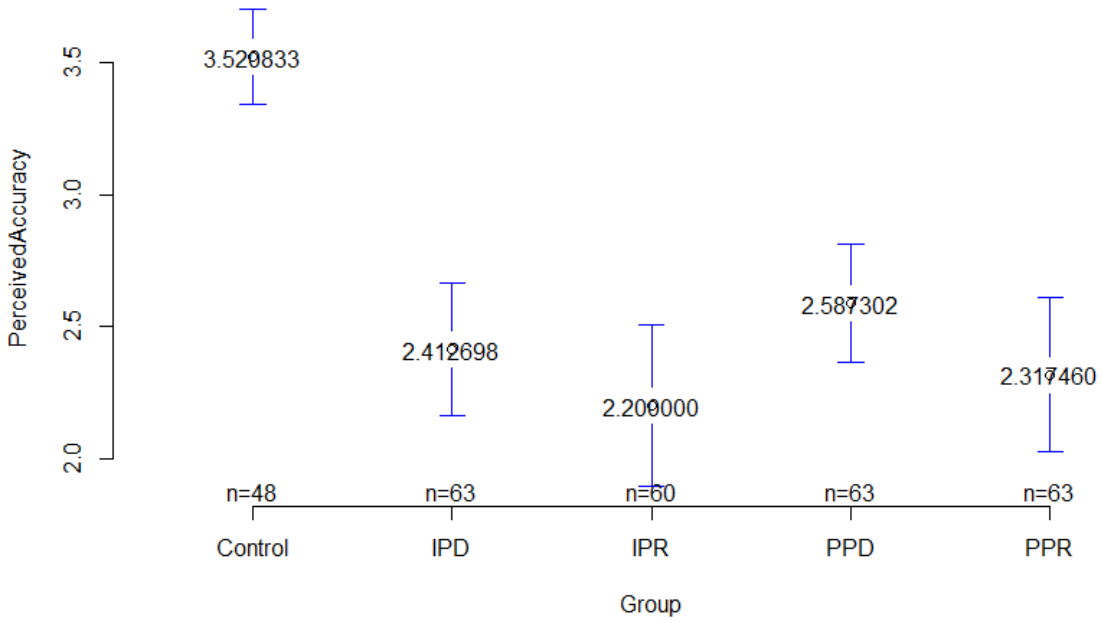

Figure 11: A plot of the means of the Perceived Accuracy of each of the groups of the experiment

As you can see in the above graph, the control group was perceived by far to be the most accurate. Additionally, the Implausible, Pro-Democrat group was perceived less accurate than

the Plausible, Pro-Democrat group; and the Implausible, Pro-Republican group was perceived les accurate than the Plausible, Pro-Republican group. These findings are in line with the expected results.

However, performing one-tailed t-tests (H: $\mu_{implausible} < \mu_{plausible}$) on the means of the perceived accuracy between the implausible and plausible groups revealed that the difference was not significant.

	n	Mean	t	df	p
Implausible, Pro-Democrat	63	2.413	-1.0379	124	.1507
Plausible, Pro-Democrat	63	2.587			

Table 1: t-test results for manipulation check perceived accuracy of Pro-Democrat articles

	n	Mean	t	df	p
Implausible, Pro-Republican	60	2.200	-0.0554	121	.2904
Plausible, Pro-Republican	63	2.317			

Table 2: t-test results for manipulation check perceived accuracy of Pro-Republican articles

While there appeared to be some variation perceived in the means, it must have been too subtle to result in a significant difference. Either the scale for perceived accuracy could have been increased in size to encourage a greater spread of ratings, or the distinguishing characteristics between plausible and implausible groups could have been made more noticeable, or both. Regardless, this is an unfortunate finding, as it means that the key comparison of the study—whether people react differently when they believe, as opposed to doubt, the veracity of a

claim—cannot be statistically determined. Any significant differences found between groups may be due to alternative explanations than the variable that was attempted to be manipulated.

At the very least, all the treatments groups were perceived as significantly less accurate than the control group. Table 3 shows the results of one-tailed t-tests (H: $\mu_{misinformation} < \mu_{control}$) comparing the means of the perceived accuracy of the experimental groups with that of the control group.

	n	Mean	t	df	p
Control	48	3.521	-	-	-
Implausible, Pro-Republican	60	2.200	7.432	92.339	5.283e-11***
Plausible, Pro-Republican	63	2.317	7.0199	98.7	2.846e-10***
Implausible, Pro-Democrat	63	2.413	7.2049	105.19	9.129e-11***
Plausible, Pro-Democrat	63	2.587	6.5068	108.15	2.461e-09***

Table 3: t-test results comparing mean perceived accuracy of treatment groups to control groups

On the whole, both the Plausible and Implausible Pro-Democrat groups were perceived as more accurate than the Pro-Republican groups, which makes sense given the Democratic-skew of the sample.

DEPENDENT VARIABLES

Affective Polarization

Exposure to the misinformation did not appear to have a significant impact on affective polarization. Feelings toward Trump, political knowledge, and overconfidence in political knowledge were the biggest predictors of this. See Table 4 on the following page for regression results.

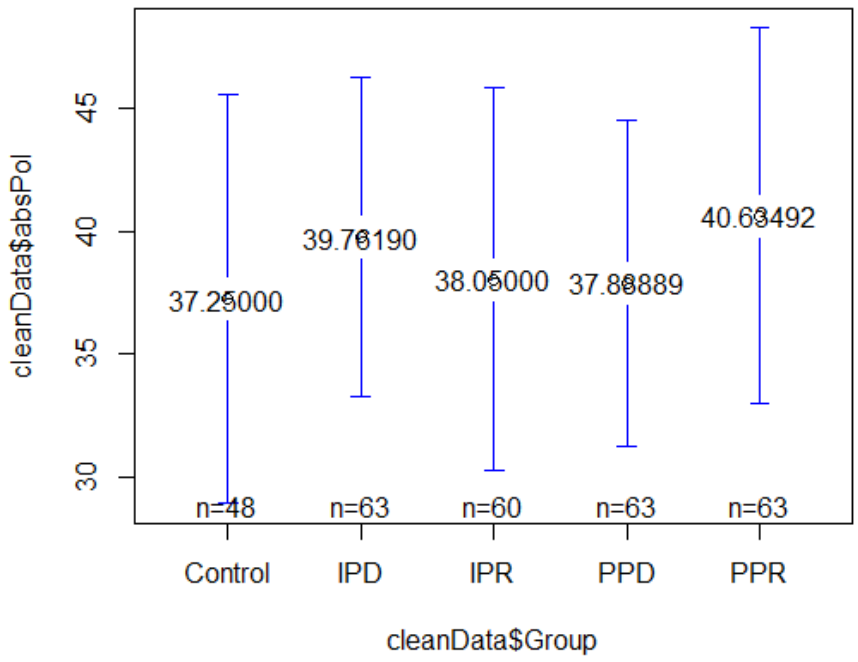

Figure 12: Plot showing the mean affective polarization of participants across all groups. Affective polarization here is defined as the absolute value of the difference between a participant's ratings of the Democratic and Republican parties

	Affective Polarization
Exposure to IPD	4.586 (5.293)
Exposure to IPD x Democrat	11.662 (10.865)
Exposure to PPD	-1.309 (5.226)
Exposure to PPD x Democrat	13.549 (11.171)
Exposure to IPR	2.18411 (5.292)
Exposure to IPR x Democrat	-10.721 (11.233)
Exposure to PPR	3.471 (5.252)
Exposure to PPR x Democrat	9.233 (10.870)
Democrat	2.0173 (4.150)
Political Knowledge	8.524** (2.964)
Overconfidence in Political Knowledge	6.202* (2.684)
Bachelor's Degree	-0.507 (3.679)
Feelings toward Trump	-0.235*** (0.057)
Age	0.226 (0.156)
Nonwhite	-4.760 (4.419)
Constant	10.487 (12.467)
Observations	286
R^2	0.1451
Adjusted R^2	0.09761
Residual Std. Error	26.33 (df=270)
F Statistic	3.055*** (df = 15, 270)

Table 4: Affective polarization, *** p < 0.001, ** p < .01, * p < .05. OLS model.

Media Trust

No significant difference was found between the groups and their ratings of trust in different types of media. Feelings towards Trump and Facebook, as well as overconfidence in political knowledge were greater predictors of media trust.

See the results of the regressions on Table 5 on the following page.

	Trust in National News	Trust in Local News	Trust in News from Social Networks
Exposure to IPD	-0.077 (0.147)	-0.128 (0.144)	-0.167 (0.149)
Exposure to IPD x Democrat	0.084 (0.304)	0.217 (0.296)	-0.011 (0.307)
Exposure to PPD	-0.254 (0.146)	-0.165 (0.142)	-0.113 (0.148)
Exposure to PPD x Democrat	-0.046 (0.313)	0.241 (0.306)	-0.025 (0.317)
Exposure to IPR	0.017 (0.147)	0.025 (0.144)	-0.110 (0.149)
Exposure to IPR x Democrat	-0.126 (0.314)	-0.026 (0.307)	-0.046 (0.318)
Exposure to PPR	-0.026 (0.146)	-0.161 (0.143)	-0.217 (0.148)
Exposure to PPR x Democrat	-0.066 (0.304)	0.397 (0.296)	0.130 (0.308)
Democrat	0.430*** (0.122)	-0.008 (0.119)	0.283* (0.123)
Political Knowledge	0.222** (0.082)	0.185* (0.080)	0.193* (0.083)
Overconfidence in Political Knowledge	0.258*** (0.074)	0.189* (0.073)	0.273*** (0.075)
Bachelor's Degree	0.277* (0.109)	0.053 (0.106)	0.017 (0.110)
Feelings toward Trump	-0.004** (0.001)	-0.005*** (0.001)	0.006*** (0.001)
Feelings toward Facebook	0.009*** (0.001)	0.006*** (0.001)	0.014*** (0.001)
Age	0.010* (0.004)	0.010* (0.004)	-0.006 (0.004)
Nonwhite	-0.012 (0.123)	0.036 (0.120)	-0.044 (0.125)
Constant	1.158*** (0.347)	1.932*** (0.339)	1.274*** (0.352)
Observations	286	286	286
R^2	0.3101	0.1433	0.421
Adjusted R^2	0.2691	0.09234	0.3866
Residual Std. Error	0.7348 (df = 269)	0.7171 (df = 269)	0.7442 (df = 269)
F Statistic	7.557*** (df = 16; 269)	2.812*** (df = 16; 269)	12.23*** (df = 16; 269)

Table 5: Media trust, *** p < 0.001, ** p < .01, * p < .05. OLS models.

Government Trust

Overall, feelings toward Facebook proved to be the greatest predictor of feelings toward government institutions, with more positive feelings toward Facebook being associated with more positive feelings toward the government as well.

Interestingly, there was a significant negative effect on feelings toward the Supreme Court found among Democrats exposed to the Plausible, Pro-Democrat and Implausible, Pro-Republican conditions, but not the other two. Dissatisfaction with the Supreme Court is much more common for Democrats than Republicans; perhaps exposure to the plausible, Pro-Democrat condition affirmed group identity and these concerns, while exposure to the Implausible, Pro-Republican condition exacerbated it.

See Table 6 on the following page for regression results.

	Feelings toward White House	Feelings toward Congress	Feelings toward Supreme Court	Feelings toward Election System
Exposure to IPD	-5.123 (4.133)	-1.240 (3.847)	-1.845 (4.341)	-0.153 (4.334)
Exposure to IPD x Democrat	10.276 (8.588)	-1.906 (7.912)	-3.355 (9.022)	8.331 (9.006)
Exposure to PPD	-6.734 (4.090)	-4.361 (3.809)	-5.878 (4.297)	-7.719 (4.289)
Exposure to PPD x Democrat	6.054 (8.845)	-7.857 (8.170)	-22.009* (9.292)	-4.396 (9.276)
Exposure to IPR	-0.579 (4.133)	-1.434 (3.849)	0.647 (4.342)	-1.292 (4.334)
Exposure to IPR x Democrat	0.262 (8.872)	-9.121 (8.189)	-18.955* (9.320)	4.331 (9.304)
Exposure to PPR	0.747 (4.108)	0.884 (3.819)	-2.886 (4.315)	0.692 (4.308)
Exposure to PPR x Democrat	6.730 (8.613)	-2.143 (7.919)	-13.038 (9.048)	10.271 (9.033)
Democrat	0.925 (3.408)	4.960 (3.182)	-7.143* (3.580)	2.176 (3.574)
Political Knowledge	3.229 (2.300)	-3.584 (2.153)	-0.711 (2.416)	3.443 (2.412)
Overconfidence in Political Knowledge	4.247* (2.088)	-1.056 (1.951)	0.259 (2.194)	4.426* (2.190)
Bachelor's Degree	0.202 (3.031)	6.133* (2.839)	-2.164 (3.184)	2.045 (3.179)
Feelings toward Trump	-0.169*** (0.048)	0.015 (0.045)	0.020 (0.051)	-0.033 (0.051)
Feelings toward Facebook	0.444*** (0.050)	0.367*** (0.047)	0.380 (0.053)	0.399*** (0.052)
Age	0.357** (0.121)	0.220 (0.113)	0.099 (0.127)	0.340** (0.127)
Nonwhite	3.296 (3.427)	4.583 (3.212)	0.891 (3.600)	2.467 (3.594)
Constant	21.240* (9.662)	28.826** (9.057)	46.407*** (10.150)	17.593 (10.132)
Observations	284	286	284	284
R^2	0.332	0.3317	0.2605	0.2974
Adjusted R^2	0.292	0.292	0.2162	0.2553
Residual Std. Error	20.4 (df=267)	19.13 (df = 269)	21.43 (df = 267)	21.39 (df = 267)
F Statistic	8.295*** (df=16;267)	8.345***(df=16;269)	5.88***(df=16;267)	7.064***(df=16;267)

Table 6: Government trust, *** p < 0.001, ** p < .01, * p < .05. OLS models.

PRESUMED INFLUENCE

Here, the results were visually clear. Survey participants rated partisans as significantly less likely to share and believe politically uncongenial material than politically congenial material. This was so pronounced that partisans were rated just as likely to believe and share congenial misinformation than the control, which contains objective fact.

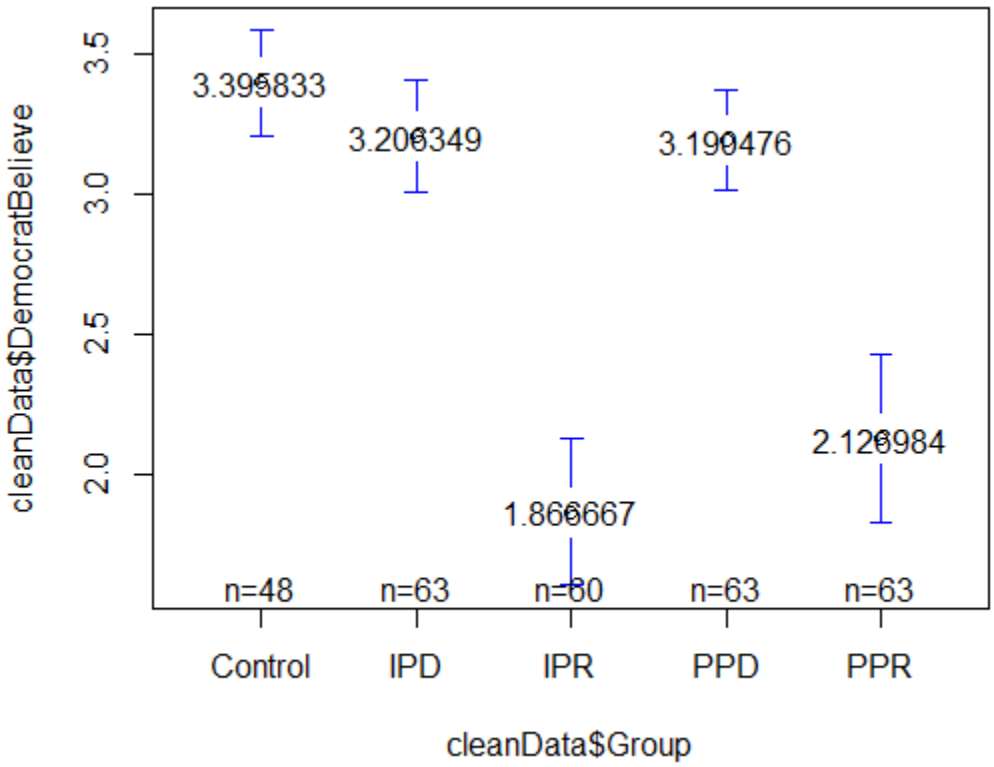

Figure 13: Plots of means of perceived likelihood for a Democrat to share each audio clip

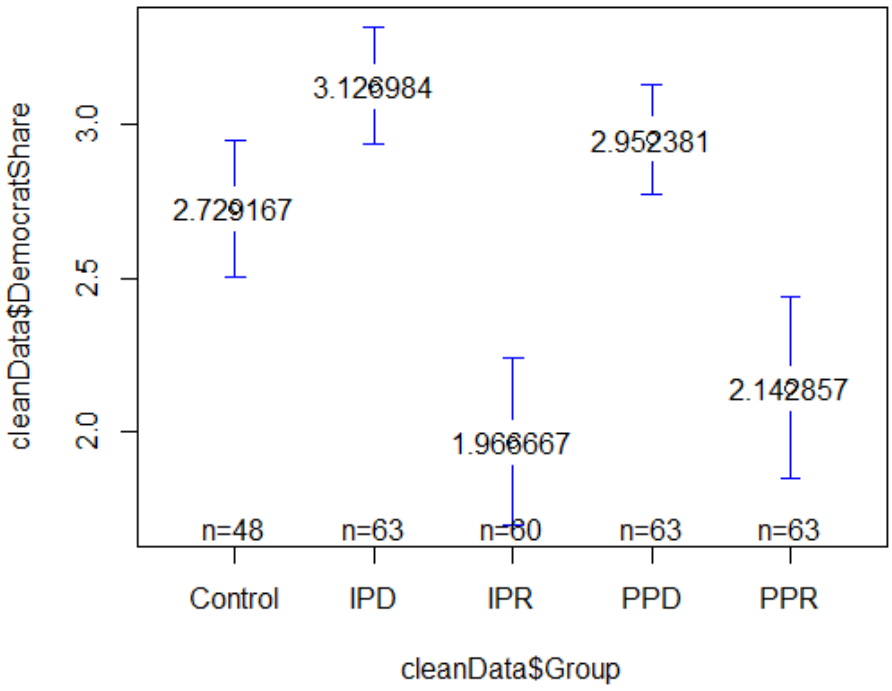

Figure 14: Plot of perceived likelihood of a Democrat to share each audio clip

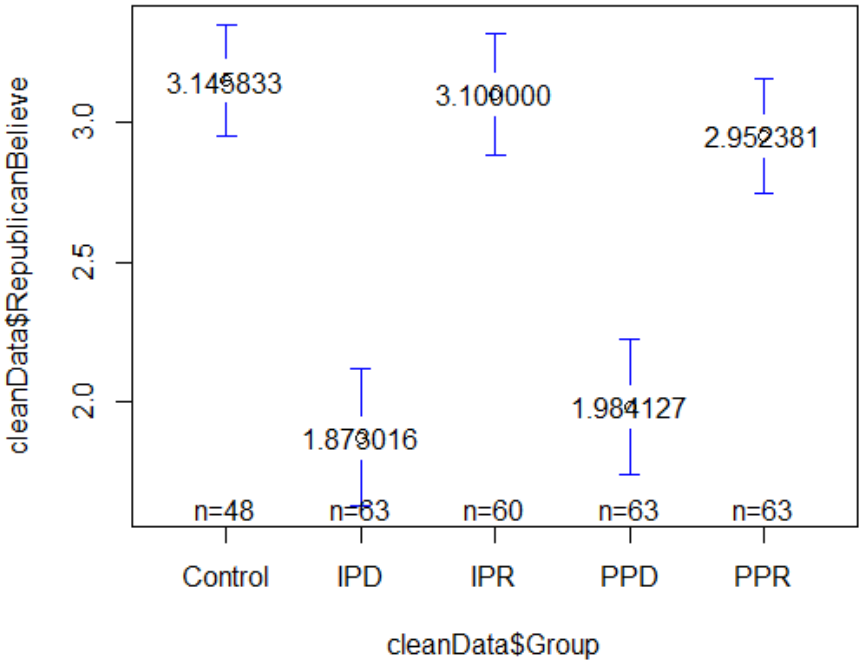

Figure 15: Plot of perceived likelihood of a Republican to believe each audio clip

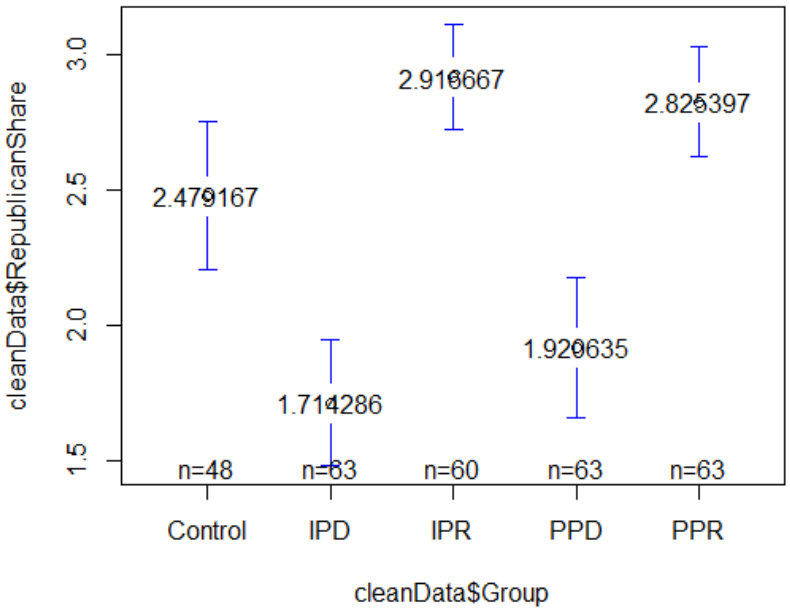

Figure 16: Plot of means of perceived likelihood for a Republican to share each audio clip

Some of this admittedly may have been due to demand characteristics; participants would have easily been able to determine the political bias of the audio clip, and from there assumed the answers we were expecting to find.

Interestingly, beyond the consistent trend of partisans being rated less likely to share or believe politically uncongenial material, Republicans were rated significantly more likely to share the Implausible, Pro-Republican audio, but *not* significantly more likely to believe it. This suggests the impression among participants that Republicans are sharing misinformation that they *know to be untrue*.

These results were highly statistically significant. For results of the regression analyses, see Table 7 on the following page.

	Likelihood Democrat Believes	Likelihood Democrat Shares	Likelihood Republican Believes	Likelihood Republican Shares
Exposure to IPD	-0.273 (0.175)	0.334 (0.172)	-1.213*** (-0.172)	-0.745*** (0.171)
Exposure to IPD x Democrat	-0.470 (0.360)	-0.486 (0.353)	-0.280 (0.354)	-0.050 (0.353)
Exposure to PPD	-0.265 (0.173)	0.208 (0.170)	-1.177*** (0.170)	-0.554** (0.170)
Exposure to PPD x Democrat	-0.696 (0.371)	-0.609 (0.365)	-0.368 (0.365)	-0.075 (0.364)
Exposure to IPR	-1.591*** (0.175)	-0.781*** (0.172)	-0.0581 (0.172)	0.453** (0.171)
Exposure to IPR x Democrat	-0.507 (0.372)	-0.718 (0.366)	0.450 (0.366)	0.655 (0.365)
Exposure to PPR	-1.308*** (0.173)	-0.602*** (0.170)	-0.233 (0.171)	0.289 (0.170)
Exposure to PPR x Democrat	-0.711* (0.360)	-0.536 (0.354)	-0.077 (0.354)	0.630 (0.353)
Democrat	0.226 (0.144)	0.133 (0.142)	0.0378 (0.142)	0.088 (0.142)
Political Knowledge	0.230* (0.098)	0.191* (0.096)	0.0027 (0.096)	0.046 (0.096)
Overconfidence in Political Knowledge	0.236** (0.088)	0.218* (0.087)	0.0606 (0.087)	0.188* (0.087)
Bachelor's Degree	0.132 (0.129)	0.026 (0.126)	0.207 (0.127)	0.038 (0.126)
Feelings toward Trump	0.006** (0.002)	0.007*** (0.002)	-6.684×10^{-5} (0.002)	0.002 (0.002)
Feelings toward Facebook	-0.001 (0.002)	0.003 (0.002)	0.00373 (0.002)	0.005** (0.002)
Age	-0.004 (0.005)	-0.0004 (0.005)	-0.00624 (0.005)	-0.002 (0.005)
Nonwhite	-0.130 (0.146)	-0.184 (0.143)	-0.0667 (0.143)	0.133 (0.143)
Constant	2.454*** (0.412)	1.696*** (0.404)	3.103*** (0.405)	2.088*** (0.404)
Observations	286	286	286	286
R^2	0.3972	0.3213	0.3454	0.3388
Adjusted R^2	0.3614	0.2809	0.3065	0.2994
Residual Std. Error	0.8706 (df=269)	.8551 (df =269)	0.8563 (df=269)	0.8536 (df=269)
F Statistic	11.08***(df=16;269)	7.959***(df=16;269)	8.871***(df=16;269)	8.614***(df=16;269)

Table 7: Presumed influence, *** p < 0.001, ** p < .01, * p < .05. OLS model.

V. CONCLUSIONS

> "Credulity is a greater evil in the present day than it ever was before, because, owing to the growth of education, it is much easier than it used to be to spread misinformation, and, owing to democracy, the spread of misinformation is more important than in former times to the holders of power."
>
> *Bertrand Russell, Free Thought and Official Propaganda (1922)*

Unfortunately, due to the manipulation check returning as not statistically significant, it is difficult to draw clear conclusions from the data, especially pertaining specifically to the research question. Is the lack of a significant difference between the Plausible and Implausible treatment groups due to them having roughly the same effects, as was hypothesized, or is it because they are effectively the same group due to lack of variation during experimentation?

Setting this aside for the moment, what *do* the data show? The treatment groups were all perceived as significantly less accurate than the control group, so they can be used to discuss misinformation more generally. That perceptual difference itself suggests that participants were able to effectively discern misinformation from true information; or at least, they were able to distinguish these instances of misinformation. Given their basis in real-life, high-profile examples, it is certainly possible that participants had already been exposed to both the stories and fact-checking of them.

Regardless, the effects of exposure to the misinformation audio were rather limited. Beyond a slight increase in negative feelings toward the Supreme Court among Democrats, there was no significant difference found for experimental groups in reported political polarization;

trust in news from national, local, or social network sources; or feelings toward the White House, Congress, and the election system. In a way, this almost supports the experiment's hypothesis: implausible and plausible misinformation conditions both had the same effects on the

This finding of limited effects of exposure to misinformation is actually in line with some of the most recent literature; a 2020 study by Guess et al. found that one-time exposure to misinformation did not produce much in the way of results beyond the increase in the belief of the central claim of the article (Guess et al.). Of course, repeated consumption of misinformation would have different effects; but then, the population that regularly consumes fake news is self-selecting, and is likely different from the citizenry at large in a number of ways.

It would be naïve to say that misinformation is completely anodyne. One needs to look no further than vaccine hesitancy to see that propagation of false information does, indeed, have real effects on the actions of the American electorate. However, the effects do not appear to extend to the electorate's political attitudes.

This may suggest a mis-ordering of the relationship. Perhaps, misinformation does not lead to or exacerbate polarization; rather, it is the very politicization of misinformation in a polarized society that makes misinformation seem dangerous. Much of the misinformation present in the media sphere is interpreted through a partisan lens, and believed based on partisan cues. Misinformation about mask mandates, vaccine safety, and hydroxychloroquine, for example, would likely have seen much less acceptance and spread in the absence of endorsement by Trump and other Republican officials.

This partisan influence is seen in the experimental results for the presumed influence of misinformation: partisans were rated unlikely to believe or share politically uncongenial information, and just as likely to share congenial information as the truth. Misinformation is

becoming more recognized as a tool (or weapon) in the political arsenal. In the results of this experiment, Republicans were actually rated as more likely to share an Implausible, Pro-Republican article than to believe it. This sharing of congenial misinformation one knows to be untrue is actually supported by the literature, as well (Chadwick et al., *The Amplification of Exaggerated and False News on Social Media*).

In a way, this is a comforting finding. Misinformation has fewer and less dire effects than previously thought; they are mostly due to the existing political climate. Obviously, however, there is an issue associated with false information being allowed to spread freely. But if people do not care that misinformation may be false, the most effective method of stopping spread may simply be to make it less visible. This practice has already seen some traction, with Facebook reducing the size of shared links from dubious sources. Implementing it on a wider scale may be the best next step.

But, ultimately, getting people up in arms about misinformation and raising its awareness may only serve to increase stress and partisan anxiety, while exposing instances of it to more people who may be incentivized to spread it, regardless of whether or not they believe it. Misinformation may be a symptom, not the disease itself. Treatment, then, is suppression of this symptom, but acknowledgement that the illness runs deeper is necessary.

APPENDIX A: SURVEY

Survey Flow

EmbeddedData
Random ID = ${rand://int/100000:999999}

Standard: Consent (2 Questions)
Block: Self-assessment of political knowledge (1 Question)
Standard: Political Identification (5 Questions)
Standard: Political knowledge battery (7 Questions)

BlockRandomizer: 1 - Evenly Present Elements

Standard: Treatment - Control (2 Questions)
Standard: Treatment - Implausible, Republican (2 Questions)
Standard: Treatment - Plausible, Republican (2 Questions)
Standard: Treatment - Implausible, Democrat (2 Questions)
Standard: Treatment - Plausible, Democrat (2 Questions)

Standard: Attention Check (2 Questions)

Branch: New Branch
If
If Were you able to successfully listen to the audio clip? No Is Selected

Block: Debriefing (1 Question)

EndSurvey:

BlockRandomizer: 2 - Evenly Present Elements

Standard: Manipulation Check - randomized before/after dependent variables (1 Question)
Standard: Dependent variables (9 Questions)

Standard: Demographic information (6 Questions)
Standard: Data cleaning (1 Question)
Standard: Debriefing (1 Question)

Page Break

Start of Block: Consent

Consent info
Please read the following information, and then, if you are willing, provide your consent to participate at the bottom of the page.

TITLE OF RESEARCH: The Relationship Between Media Consumption and Political Attitudes and Beliefs
PRINCIPAL INVESTIGATOR'S DEPARTMENT: School of Public and International Affairs & Department of Politics

Key information about the study: Your informed consent is being sought for research. Participation in the research is voluntary.

The purpose of the research: to determine the relationship between media consumption and political attitudes and beliefs

Participation should not take more than 5-10 minutes.

You will be asked a few questions about politics. You will then listen to a short audio news clip. Finally, you will answer some more questions about your political attitudes and beliefs, and demographics.

You may be exposed to viewpoints that frustrate you or that you disagree with. You will be asked a few knowledge-based questions that you may not know the answers to.

The benefits to the subject or to others, e.g., society that may reasonably be expected from the research: a better understanding of the relationship between media consumption and political attitudes and beliefs, so as to make more informed policy decisions

The alternative procedures, if any, that might be advantageous to the subject: You may choose not to participate in the research.

Additional information about the study:
Confidentiality: All records from this study will be kept confidential. Your responses will be kept private, and we will not include any information that will make it possible to identify you in any report we might publish. Research records will be stored securely in a locked cabinet and/or on password-protected computers. The research team will be the only party that will have access to your data.
Compensation: You will be paid $3.34 following your participation.

If you have questions regarding your rights as a research subject, or if problems arise which you do not feel you can discuss with the Investigator, please contact the Institutional Review Board.

Summary:
I understand the information that was presented and that:
My participation is voluntary.
Refusal to participate will involve no penalty or loss of benefits to which I am otherwise entitled.
I may discontinue participation at any time without penalty or loss of benefits.
I do not waive any legal rights or release Princeton University or its agents from liability for negligence.

X→

Consent Do you give your consent to be a subject of the research? You may discontinue your participation at any time.

- ○ Yes, I hereby give my consent to be a subject of the research and participate in the survey (1)
- ○ No, I do not consent to participating in the survey (0)

Skip To: End of Survey If Do you give your consent to be a subject of the research? You may discontinue your participation... = No, I do not consent to participating in the survey

End of Block: Consent

Start of Block: Self-assessment of political knowledge

SelfAssessment In general, how would you rate your political knowledge?

- ○ Not at all knowledgeable (1)
- ○ Not very knowledgeable (2)
- ○ Somewhat knowledgeable (3)
- ○ Very knowledgeable (4)

End of Block: Self-assessment of political knowledge

Start of Block: Political Identification

PoliticalID Generally speaking, do you usually think of yourself as a Democrat, a Republican, and independent, or what?

- ○ Democrat (1)
- ○ Republican (2)
- ○ Independent (3)
- ○ Something else (4)

Display This Question:

If Generally speaking, do you usually think of yourself as a Democrat, a Republican, and independent... = Something else

if "something else" What is that?

_____

Display This Question:

If Generally speaking, do you usually think of yourself as a Democrat, a Republican, and independent... = Democrat

X→

ifDemocrat Would you call yourself a strong Democrat or a not very strong Democrat?

- ○ Strong (1)
- ○ Not very strong (2)

Display This Question:

If Generally speaking, do you usually think of yourself as a Democrat, a Republican, and independent... = Republican

ifRepublican Would you call yourself a strong Republican or a not very strong Republican?

- Strong (7)
- Not very strong (6)

Display This Question:

If Generally speaking, do you usually think of yourself as a Democrat, a Republican, and independent... = Independent

Or Generally speaking, do you usually think of yourself as a Democrat, a Republican, and independent... = Something else

And And What is that? Text Response Is Displayed

ifIndependent Do you think of yourself as closer to the Republican Party or to the Democratic Party?

- Closer to the Republican Party (5)
- Closer to the Democratic Party (3)
- Neither (4)

End of Block: Political Identification

Start of Block: Political knowledge battery

Instructions The following questions are intended to help us learn what aspects of political information are commonly known by the public. Please answer these questions to the best of your ability. It is very important to our research that you do not ask anyone or look up the answers. It is common for many people to not know the answers to these questions, but please do your best to respond even if you are unsure of the correct answer.

Do you agree to answer the following questions to the best of your ability without help from outside sources? This is necessary to our research.

- ○ Yes, I agree to answer without help from outside sources (1)
- ○ No, I do not agree to answer without help from outside sources (0)

Page Break

Senate What is the length of a U.S. Senate term?

- ○ 2 years (0)
- ○ 4 years (0)
- ○ 6 years (1)
- ○ 10 years (0)

Roberts What job or political office is now held by John Roberts?

- ○ Chief Justice of the Supreme Court (1)
- ○ President pro tempore of the US Senate (0)
- ○ Speaker of the House (0)
- ○ US Senate Minority Whip (0)

Scholz What job or political office is now held by Olaf Scholz?

- Chancellor of Germany (1)
- Secretary-General of the United Nations (0)
- Secretary-General of NATO (0)
- US Ambassador to Germany (0)

HouseSeats Which party has a majority of seats in the U.S. House of Representatives?

- Democratic Party (1)
- Republican Party (0)

Spending On which of the following does the U.S. federal government currently spend the least?

- Foreign aid (1)
- Medicare (0)
- National defense (0)
- Social Security (0)

Page Break

Cheat In answering the last few items, did you look up any of the questions online or ask someone for help?

- ○ No, I did not use any outside sources (1)
- ○ Yes, I used outside sources (0)

End of Block: Political knowledge battery

Start of Block: Treatment - Control

C Please listen to the 1-minute audio clip below.

[Embedded video: https://youtu.be/1ibIUFwQtgY]

Control Once you have finished listening, check this box to continue with the survey

- ○ I'm done listening! (1)

End of Block: Treatment - Control

Start of Block: Treatment - Implausible, Republican

IPR Please listen to the 1-minute audio clip below.

[Embedded video: https://youtu.be/dS8dFZQHpOQ]

IPR Once you have finished listening, check this box to continue with the survey

- ○ I'm done listening! (1)

End of Block: Treatment - Implausible, Republican

Start of Block: Treatment - Plausible, Republican

PPR Please listen to the 1-minute audio clip below.

[Embedded video: https://youtu.be/M5Ra1UUE-d0]

PPR Once you have finished listening, check this box to continue with the survey

- ○ I'm done listening! (1)

End of Block: Treatment - Plausible, Republican

Start of Block: Treatment - Implausible, Democrat

IPD Please listen to the 1-minute audio clip below.

[Embedded video: https://youtu.be/WYHxLMbcd30]

IPD Once you have finished listening, check this box to continue with the survey

- ○ I'm done listening! (1)

End of Block: Treatment - Implausible, Democrat

Start of Block: Treatment - Plausible, Democrat

PPD Please listen to the 1-minute audio clip below.

[Embedded video: https://youtu.be/pln3od7ao5E]

PPD Once you have finished listening, check this box to continue with the survey

- I'm done listening! (1)

End of Block: Treatment - Plausible, Democrat

Start of Block: Attention Check

AudioWorks Were you able to successfully listen to the audio clip?

- Yes (1)
- No (0)

Skip To: End of Block If Were you able to successfully listen to the audio clip? = No

Page Break

AttnCheck Who was the primary subject of the clip that you listened to?

- Former President Donald Trump (1)
- Chef Gordon Ramsay (0)
- Former President Barack Obama (0)
- Seismologist Chris Goldfinger (2)

End of Block: Attention Check

Start of Block: Manipulation Check - randomized before/after dependent variables

PerceivedAccuracy How accurate would you consider the claim put forth by the audio clip you listened to?

- Not at all accurate (1)
- Not very accurate (2)
- Somewhat accurate (3)
- Very accurate (4)

End of Block: Manipulation Check - randomized before/after dependent variables

---

Start of Block: Dependent variables

RepublicanBelieve How likely would you expect the average Republican to *believe* the information in the clip you listened to?

- Not at all likely (1)
- Not very likely (2)
- Somewhat likely (3)
- Very likely (4)

RepublicanShare How likely would you expect the average Republican to *share* the clip you listened to?

- Not at all likely (1)
- Not very likely (2)
- Somewhat likely (3)
- Very likely (4)

DemocratBelieve How likely would you expect the average Democrat to *believe* the information in the clip you listened to?

- Not at all likely (1)
- Not very likely (2)
- Somewhat likely (3)
- Very likely (4)

DemocratShare How likely would you expect the average Democrat to *share* the clip you listened to?

- Not at all likely (1)
- Not very likely (2)
- Somewhat likely (3)
- Very likely (4)

Page Break

FeelingsThermometer Please rate your feelings toward the following political leaders and institutions that often appear in the news. A score of 0-49 would indicate cold or unfavorable feelings (0 being very cold), while a score of 51-100 would indicate warm or favorable feelings (100 being very warm). A score of 50 would indicate no particular feelings, favorable or unfavorable (neutrality). If you do not recognize a name or institution, please choose "not applicable".

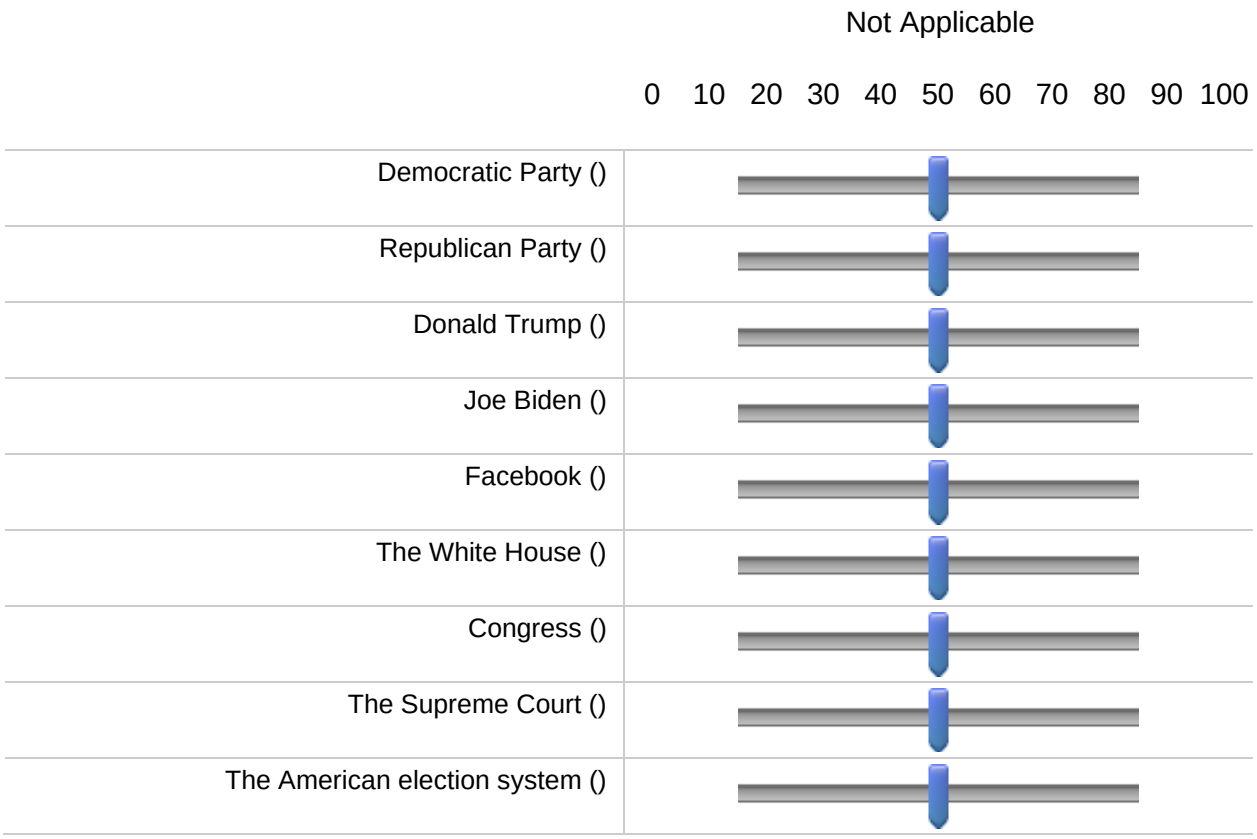

Page Break

Media attitude **How much, if at all, do you trust the information that comes from each of the following sources?**

NationalNews National news organization

- ○ Not at all (1)
- ○ Not too much (2)
- ○ Some (3)
- ○ A lot (4)

LocalNews Local news organizations

- ○ Not at all (1)
- ○ Not too much (2)
- ○ Some (3)
- ○ A lot (4)

SocialNetworks Social networking sites, such as Facebook and Twitter

- ○ Not at all (1)
- ○ Not too much (2)
- ○ Some (3)
- ○ A lot (4)

End of Block: Dependent variables

Start of Block: Demographic information

BirthYear What is the year you were born?

▼ 2004 (18) ... 1905 (117)

Gender What best describes your gender?

- ○ Male (1)
- ○ Female (0)
- ○ Non-binary / other

Race What racial or ethnic group best describes you?

- White (1)
- Black or African-American (2)
- Hispanic or Latino (3)
- Asian (4)
- Native American (5)
- Middle Eastern (6)
- Mixed Race (7)
- Other (8)

Education What is the highest level of education you have completed?

- Did not graduate high school (1)
- High school or equivalent (2)
- Some college (3)
- 2-year college degree (4)
- 4-year college degree (5)
- Graduate degree (6)

Ideology Which best describes your political ideology?

- ○ Very liberal (1)
- ○ Somewhat liberal (2)
- ○ A little liberal (3)
- ○ Neither liberal nor conservative (4)
- ○ A little conservative (5)
- ○ Somewhat conservative (6)
- ○ Very conservative (7)

State Which state do you live in?

▼ Alabama (1) ... Wyoming (50)

End of Block: Demographic information

Start of Block: Data cleaning

DataCleaning At any point during the survey (aside from the initial questions about political facts, where it was okay to guess), have you responded to questions randomly? Please be honest!

- ○ No, I have not responded randomly (0)
- ○ Yes, I have responded randomly (1)

End of Block: Data cleaning

Start of Block: Debriefing

Debrief **TO FINISH RECORDING YOUR RESPONSES, PLEASE CLICK THE ARROW AT THE BOTTOM OF THE PAGE.**

Here is your unique completion code: ${e://Field/Random%20ID}
Please copy and paste this into MTurk to receive your payment.

Thank you for participating in our study. In this study, we wanted to see whether the perceived truthfulness of a claim would alter the effects of exposure to misinformation on political attitudes and beliefs – the effects under observation being primarily trust in media and government, and political polarization.

We predict that regardless of whether the participant perceives the fake news as truthful or not, the effects of exposure to fake news will be maintained (in whole or in part).

In order to test this prediction, we asked each participant to listen to one of five possible audio clips. One was an excerpt of a New Yorker article about earthquakes unrelated to politics, acting as a control group. The other four conditions—framed as either pro-Democrat or pro-Republican, and either more or less plausible—put forth false claims that were adapted from existing misinformation and fake news.

The claims made in these four experimental groups are, to be clear, unsubstantiated. There has not been credible evidence supporting the claim that Donald Trump is a cultivated Russian agent, nor the claim that the 2020 presidential election was rigged by an insidious Deep State.

Additionally, the voice you heard in your assigned audio file was actually synthesized using a text-to-speech tool -- this one designed to sound like the late television host Larry King. The same technology can be used to produce audio files mimicking many recognizable voices, ranging from actors, such as Morgan Freeman, to politicians, such as former President Donald Trump (which is part of why it is important to study)!

Luckily, there are resources available to counteract misinformation. You may visit websites such as Snopes, PolitiFact, and FactCheck.org that compile and fact-check claims made by political actors or circulating online. Most social media sites have systems in place that, while not perfect, will flag disputed claims with warnings and provide resources for more information. If you are interested in learning to better recognize misinformation online, you might consider playing the Bad News Game, which provides an interactive way to learn some of the tactics employed by creators of misinformation.

We ask that you not discuss the details of this study online or with others who may participate in the future. If you have any questions or comments, you may contact the Principal Investigator. If you have concerns that you feel you cannot discuss with the researchers, you may contact the Institutional Review Board. Please click the arrow below to finish recording your responses.

End of Block: Debriefing

APPENDIX B: STIMULI TRANSCRIPTS

Implausible, Pro-Democrat (IPD)

(Audio available at https://youtu.be/WYHxLMbcd30)

Everybody knows that Russia interfered with the 2016 presidential election in the United States. But what they didn't tell you is that Donald Trump was specifically selected, groomed, courted, and used by Russia to undermine American democracy. When Trump first visited Russia in 1987, now-Russian President Vladimir Putin, who used to be a part of the Russian intelligence agency, the KGB, began scouting the American businessman. Trump was apparently targeted for being intellectually vulnerable and prone to flattery. Vlad sure got something right! Between then and Trump's election victory, Russia compiled several tapes of *kompromat*, or compromising information, that they have been using to blackmail President Trump ever since. This, at least, would give some explanation for Trump's weirdly forgiving stance on foreign relations with the Kremlin—he has been acting in their interest as a cultivated asset all along! ("The Manchurian Candidate Theory Will Never Die") (Davidson) ("Did Ex-KGB Spy Say Russia Cultivated Trump as an 'Asset' for 40 Years?") (Unger)

Plausible, Pro-Democrat (PPD)

(Audio available at https://youtu.be/pln3od7ao5E)

It has, at this point, been well-documented that Russia attempted to interfere with the 2016 presidential election in the United States. However, according to ex-Russian intelligence agent Yuri Shvets, the interference may have run deeper and started earlier than previously known. Shvets claims that Donald Trump was specifically selected, groomed, courted, and used

by Russia to undermine American democracy. Per Shvets, when Trump first visited Russia in 1987, now-Russian President Vladimir Putin, then part of the Russian intelligence agency, the KGB, began scouting the American businessman. Trump was reportedly targeted for being intellectually vulnerable and prone to flattery. Between then and Trump's election, Shvets alleges, Russia compiled several tapes of *kompromat*, or compromising information, that they have been using to blackmail President Trump. This would help to explain Trump's forgiving stance on foreign relations with the Kremlin—he has been acting in their interest as a cultivated asset. ("The Manchurian Candidate Theory Will Never Die") (Davidson) ("Did Ex-KGB Spy Say Russia Cultivated Trump as an 'Asset' for 40 Years?") (Unger)

Implausible, Pro-Republican (IPR)

(Audio available at https://youtu.be/dS8dFZQHpOQ)

The 2020 election was completely stolen. The results were rigged by the same Deep State that attempted to undermine Donald Trump's entire presidency. The Deep State, for those not yet in the know, is a super-secret and nefarious group of military, intelligence, and government officials who manipulate the government from the inside. Trump's great willingness to criticize these unelected officials, unlike all of his predecessors, led to them working to take down his administration. The whole fake Russia investigation was perhaps the most glaring example of this. That politically-motivated investigation was based on a bunk dossier that made claims that were both ludicrous and frivolous, like that Russian President Putin was in possession of a video in which Trump watched strippers pee on a hotel bed. This was all specifically done in a concerted effort to damage the Trump reelection campaign! (Pengelly) (Fox News) (Lewandowski and Bossie)

Plausible, Pro-Republican (PPR)

(Audio available at https://youtu.be/M5Ra1UUE-d0)

The truth of the results of the 2020 election has been under scrutiny for some time now. According to many, such as political expert Corey Lewandowski, the results were rigged by the same Deep State that attempted to undermine Donald Trump's presidency. The Deep State, says Lewandowski, is a clandestine group of military, intelligence, and government officials who secretly manipulate the government. Trump's apparent willingness to criticize these unelected officials allegedly led to them working to sow doubt against his administration. In his new book, Lewandowski holds up the Russia investigation as perhaps the greatest example of this. The politically-motivated investigation, led by scorned FBI director James Comey, attempted to legitimize a dossier that made a number of dubious claims, such as the assertion that Russian President Putin was in possession of a video in which Trump watched strippers urinate in a hotel room. This was all done, per Lewandowski, in a concerted effort to damage the Trump reelection campaign. (Pengelly) (Fox News) (Lewandowski and Bossie)

Control

(Unrelated to politics – an excerpt from "The Really Big One" by Kathryn Schulz

Audio available at https://youtu.be/1ibIUFwQtgY)

"When the 2011 earthquake and tsunami struck Tohoku, Japan, Chris Goldfinger was two hundred miles away, in the city of Kashiwa, at an international meeting on seismology. As the shaking started, everyone in the room began to laugh. Earthquakes are common in Japan—that one was the third of the week—and the participants were, after all, at a seismology conference. Then everyone in the room checked the time. Seismologists know that the length of an

earthquake is a decent proxy for its magnitude. The 1989 earthquake in California, which killed sixty-three people and caused six billion dollars' worth of damage, lasted about fifteen seconds and had a magnitude of 6.9. A thirty-second earthquake generally has a magnitude in the mid-sevens. A minute-long quake is in the high sevens, a two-minute quake has entered the eights, and a three-minute quake is in the high eights. By four minutes, an earthquake has hit magnitude 9.0" (Schulz).

WORKS CITED

"2020 Presidential Election Voting and Registration Tables Now Available." *Census.Gov*, 29 Apr. 2021, https://www.census.gov/newsroom/press-releases/2021/2020-presidential-election-voting-and-registration-tables-now-available.html.

"2021 Pew Research Center's American Trends Panel - Trust in Media." *Pew Research Center*, https://www.pewresearch.org/wp-content/uploads/2021/08/Topline-trust-in-media-FINAL.pdf. Accessed 11 Apr. 2022.

Amazon Mechanical Turk Worker Help. https://www.mturk.com/worker/help. Accessed 10 Apr. 2022.

Anderson, Craig A., et al. "Perseverance of Social Theories: The Role of Explanation in the Persistence of Discredited Information." *Journal of Personality and Social Psychology*, vol. 39, no. 6, 1980, pp. 1037–49. *DOI.org (Crossref)*, https://doi.org/10.1037/h0077720.

"ANES 2019 Pilot Study Questionnaire Specifications." *American National Election Studies*, 2019, https://electionstudies.org/wp-content/uploads/2020/02/anes_pilot_2019_questionnaire.pdf.

"Bad News." *Bad News*, https://www.badnews.eu. Accessed 13 Dec. 2021.

Baek, Young Min, et al. "Fake News Should Be Regulated Because It Influences Both 'Others' and 'Me': How and Why the Influence of Presumed Influence Model Should Be Extended." *Mass Communication and Society*, vol. 22, no. 3, May 2019, pp. 301–23. *Taylor and Francis+NEJM*, https://doi.org/10.1080/15205436.2018.1562076.

Bail, Christopher A., et al. "Assessing the Russian Internet Research Agency's Impact on the Political Attitudes and Behaviors of American Twitter Users in Late 2017." *Proceedings of the National Academy of Sciences*, vol. 117, no. 1, Jan. 2020, pp. 243–50. *DOI.org (Crossref)*, https://doi.org/10.1073/pnas.1906420116.

Barbaro, Michael. "Donald Trump Clung to 'Birther' Lie for Years, and Still Isn't Apologetic." *The New York Times*, 16 Sept. 2016. *NYTimes.com*, https://www.nytimes.com/2016/09/17/us/politics/donald-trump-obama-birther.html.

Berinsky, Adam J., et al. *Using Mechanical Turk as a Subject Recruitment Tool for Experimental Research*. 17 Sept. 2010, https://qipsr.as.uky.edu/sites/default/files/Berinsky.Using%20Mechanical%20Turk%20as%20a%20Subject%20Recruitment%20Tool%20for%20Experimental%20Research.pdf.

Bridgman, Aengus, et al. "The Causes and Consequences of COVID-19 Misperceptions: Understanding the Role of News and Social Media." *Harvard Kennedy School Misinformation Review*, vol. 1, no. 3, June 2020. *misinforeview.hks.harvard.edu*, https://doi.org/10.37016/mr-2020-028.

Brown, Brendan. *Search on Trump Twitter Archive*. //www.thetrumparchive.com. Accessed 11 Apr. 2022.

Calvert, Clay. *Fake News, Censorship & the Third-Person Effect: You Can't Fool Me, Only Others! | HuffPost*. 16 Feb. 2017, https://www.huffpost.com/entry/fake-news-censorship-the-third-person-effect-you_b_58a4ccb3e4b026a89a7a25e4.

Caralle, Katelyn. "Republican Quarantined over Coronavirus Jokes about Meeting with RBG." *Mail Online*, 11 Mar. 2020, https://www.dailymail.co.uk/news/article-8100975/Congressman-quarantined-coronavirus-exposure-jokes-meeting-Ruth-Bader-Ginsburg.html.

Chadwick, Andrew, et al. *The Amplification of Exaggerated and False News on Social Media: The Roles of Platform Use, Motivations, Affect, and Ideology*. Mar. 2021. *repository.lboro.ac.uk*, https://repository.lboro.ac.uk/articles/journal_contribution/The_amplification_of_exaggerated_and_false_news_on_social_media_the_roles_of_platform_use_motivations_affect_and_ideology/14223083/1.

---. *The Amplification of Exaggerated and False News on Social Media: The Roles of Platform Use, Motivations, Affect, and Ideology*. Mar. 2021. *repository.lboro.ac.uk*, https://repository.lboro.ac.uk/articles/journal_contribution/The_amplification_of_exaggerated

_and_false_news_on_social_media_the_roles_of_platform_use_motivations_affect_and_ideology/14223083/1.

CNN Poll: Quarter Doubt Obama Was Born in U.S. – CNN Political Ticker - CNN.Com Blogs. 24 Apr. 2011, https://web.archive.org/web/20110424025320/http://politicalticker.blogs.cnn.com/2010/08/04/cnn-poll-quarter-doubt-president-was-born-in-u-s/.

Cohen, Geoffrey L. "Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs." *Journal of Personality and Social Psychology*, vol. 85, no. 5, 2003, pp. 808–22. *DOI.org (Crossref)*, https://doi.org/10.1037/0022-3514.85.5.808.

Cohen, Jon. "Poll: Number of 'Birthers' Plummets." *The Washington Post*, 5 May 2011, https://web.archive.org/web/20170913232144/https://www.washingtonpost.com/blogs/behind-the-numbers/post/number-of-birthers-plummets/2011/05/04/AF3GAZxF_blog.html?utm_term=.5f52db663add.

Constine, Josh. "Facebook Shrinks Fake News after Warnings Backfire." *TechCrunch*, 13 Dec. 2021, https://social.techcrunch.com/2018/04/27/facebook-false-news/.

Cunliffe, Kyle. "Donald Trump: More Likely Useful Idiot Than Putin's Agent." *Snopes*, https://www.snopes.com/news/2021/02/05/donald-trump-more-likely-useful-idiot-than-putins-agent/. Accessed 11 Apr. 2022.

Davidson, Adam. "A Theory of Trump Kompromat." *The New Yorker*, July 2018. *www.newyorker.com*, https://www.newyorker.com/news/swamp-chronicles/a-theory-of-trump-kompromat.

DeSoto, K. Andrew. "Under the Hood of Mechanical Turk." *APS Observer*, vol. 29, no. 3, 3, Feb. 2016. *www.psychologicalscience.org*, https://www.psychologicalscience.org/observer/under-the-hood-of-mechanical-turk.

"Did Ex-KGB Spy Say Russia Cultivated Trump as an 'Asset' for 40 Years?" *Snopes.Com*, https://www.snopes.com/fact-check/kgb-spy-russia/. Accessed 9 Apr. 2022.

Dimock, Michael, et al. “Political Polarization in the American Public.” *Pew Research Center*, June 2014, p. 124.

Evon, Dan. “Trump’s Election Fraud Claims: A Collection.” *Snopes*, https://www.snopes.com/collections/trumps-election-fraud-claims/. Accessed 11 Apr. 2022.

“Facebook Shrinks Fake News after Warnings Backfire.” *TechCrunch*, https://social.techcrunch.com/2018/04/27/facebook-false-news/. Accessed 13 Dec. 2021.

FakeYou. https://fakeyou.com. Accessed 11 Apr. 2022.

Faragó, Laura, et al. “We Only Believe in News That We Doctored Ourselves: The Connection between Partisanship and Political Fake News.” *Social Psychology*, vol. 51, no. 2, 2020, pp. 77–90. 2019-57441-001, *EBSCOhost*, https://doi.org/10.1027/1864-9335/a000391.

Fichera, Angelo, et al. “The Falsehoods of the ‘Plandemic’ Video.” *FactCheck.Org*, 8 May 2020, https://www.factcheck.org/2020/05/the-falsehoods-of-the-plandemic-video/.

Fox News. *Deep State Undermined Trump Presidency: Glenn Greenwald*. 2021. *YouTube*, https://www.youtube.com/watch?v=7xuceo6OCcc.

Gerlich, R. Nicholas, et al. *MECHANICAL TURK: IS IT JUST ANOTHER CONVENIENCE SAMPLE?* no. 1, 2018, p. 11.

Gramlich, John. “What the 2020 Electorate Looks like by Party, Race and Ethnicity, Age, Education and Religion.” *Pew Research Center*, https://www.pewresearch.org/fact-tank/2020/10/26/what-the-2020-electorate-looks-like-by-party-race-and-ethnicity-age-education-and-religion/. Accessed 10 Apr. 2022.

Grinberg, Nir, et al. “Fake News on Twitter during the 2016 U.S. Presidential Election.” *Science*, vol. 363, no. 6425, Jan. 2019, pp. 374–78. *science-sciencemag-org.ezproxy.princeton.edu*, https://doi.org/10.1126/science.aau2706.

Guess, Andrew M., et al. "'Fake News' May Have Limited Effects beyond Increasing Beliefs in False Claims." *Harvard Kennedy School Misinformation Review*, vol. 1, no. 1, 1, Jan. 2020. *misinforeview.hks.harvard.edu*, https://doi.org/10.37016/mr-2020-004.

Hao, Karen. "Troll Farms Reached 140 Million Americans a Month on Facebook before 2020 Election, Internal Report Shows." *MIT Technology Review*, https://www.technologyreview.com/2021/09/16/1035851/facebook-troll-farms-report-us-2020-election/. Accessed 18 Sept. 2021.

Hastorf, Albert H., and Hadley Cantril. *CASE REPORTS THEY SAW A GAME: A CASE STUDY*. 1952, p. 6.

Jamison, Amelia M., et al. "Not Just Conspiracy Theories: Vaccine Opponents and Proponents Add to the COVID-19 'Infodemic' on Twitter." *Harvard Kennedy School Misinformation Review*, vol. 1, no. 3, Sept. 2020. *misinforeview.hks.harvard.edu*, https://doi.org/10.37016/mr-2020-38.

Jefferson, Thomas. "Letter to John Norvell, June 11, 1807." *Library of Congress, Washington, D.C. 20540 USA*, 11 June 1807, https://www.loc.gov/resource/mtj1.038_0592_0594/?sp=2&st=text.

Johnson, Eric J., and Amos Tversky. "Affect, Generalization, and the Perception of Risk." *Journal of Personality and Social Psychology*, vol. 45, no. 1, 1983, pp. 20–31. *APA PsycNet*, https://doi.org/10.1037/0022-3514.45.1.20.

Jones, Nicholas, et al. "2020 Census Illuminates Racial and Ethnic Composition of the Country." *Census.Gov*, 12 Aug. 2021, https://www.census.gov/library/stories/2021/08/improved-race-ethnicity-measures-reveal-united-states-population-much-more-multiracial.html.

King, Gary, et al. "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument." *American Political Science Review*, vol. 111, no. 3, Aug. 2017, pp. 484–501. *DOI.org (Crossref)*, https://doi.org/10.1017/S0003055417000144.

Kjberry | TurkerView MTurk Requester Profile. https://turkerview.com/requesters/A1TPVEE1IUZAHC-kjberry. Accessed 10 Apr. 2022.

Lewandowski, Corey R., and David N. Bossie. *Trump's Enemies: How the Deep State Is Undermining the Presidency*. Center Street, 2018.

Lewandowsky, Stephan, Ullrich K. H. Ecker, et al. "Misinformation and Its Correction: Continued Influence and Successful Debiasing." *Psychological Science in the Public Interest*, vol. 13, no. 3, Dec. 2012, pp. 106–31. *SAGE Journals*, https://doi.org/10.1177/1529100612451018.

Lewandowsky, Stephan, Werner G. K. Stritzke, et al. "Misinformation, Disinformation, and Violent Conflict: From Iraq and the 'War on Terror' to Future Threats to Peace." *American Psychologist*, vol. 68, no. 7, Oct. 2013, pp. 487–501. 2013-35832-001, *EBSCOhost*, https://doi.org/10.1037/a0034515.

Lewsey, Fred. "Cambridge Game 'Pre-Bunks' Coronavirus Conspiracies." *University of Cambridge*, 11 Oct. 2020, https://www.cam.ac.uk/stories/goviral.

Lloyd, Robert. "Larry King, Digital Pioneer at 84, on Life, Work, Trump's Phone Calls and His Newest Emmy Nominations." *Los Angeles Times*, 24 Apr. 2018, https://www.latimes.com/entertainment/tv/la-et-st-larry-king-conversation-20180424-htmlstory.html.

Lord, Charles G., et al. *Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence*. 1979, p. 12.

Mantayla, Kyle. "Mark Taylor: Obama Will Claim Kenyan Citizenship to Keep From Being Tried for Treason." *Right Wing Watch*, 16 Apr. 2019. *www.rightwingwatch.org*, https://www.rightwingwatch.org/post/mark-taylor-obama-will-claim-kenyan-citizenship-to-keep-from-being-tried-for-treason/.

McDuffie, Danielle. "Student Notebook: Benefits, Drawbacks, and Suggestions for Using Amazon's Mechanical Turk." *APS Observer*, vol. 32, no. 2, Jan. 2019. *www.psychologicalscience.org*,

https://www.psychologicalscience.org/observer/using-amazons-mechanical-turk-benefits-drawbacks-and-suggestions.

"Median Age of the U.S. Population 2020." *Statista*, https://www.statista.com/statistics/241494/median-age-of-the-us-population/. Accessed 10 Apr. 2022.

Mitchell, Amy, et al. "Many Americans Say Made-Up News Is a Critical Problem That Needs To Be Fixed." *Pew Research Center's Journalism Project*, 5 June 2019, https://www.pewresearch.org/journalism/2019/06/05/many-americans-say-made-up-news-is-a-critical-problem-that-needs-to-be-fixed/.

Nisbet, Erik C., et al. "The Presumed Influence of Election Misinformation on Others Reduces Our Own Satisfaction with Democracy." *Harvard Kennedy School Misinformation Review*, Mar. 2021. *misinforeview.hks.harvard.edu*, https://doi.org/10.37016/mr-2020-59.

---. "The Presumed Influence of Election Misinformation on Others Reduces Our Own Satisfaction with Democracy." *Harvard Kennedy School Misinformation Review*, Mar. 2021. *misinforeview.hks.harvard.edu*, https://doi.org/10.37016/mr-2020-59.

Nisbet, Erik C., and Olga Kamenchuk. "The Psychology of State-Sponsored Disinformation Campaigns and Implications for Public Diplomacy." *The Hague Journal of Diplomacy*, vol. 14, no. 1–2, 1–2, Apr. 2019, pp. 65–82. *brill-com.ezproxy.princeton.edu*, https://doi.org/10.1163/1871191X-11411019.

Ognyanova, Katherine, et al. "Misinformation in Action: Fake News Exposure Is Linked to Lower Trust in Media, Higher Trust in Government When Your Side Is in Power." *Harvard Kennedy School Misinformation Review*, June 2020. *misinforeview.hks.harvard.edu*, https://doi.org/10.37016/mr-2020-024.

---. “Misinformation in Action: Fake News Exposure Is Linked to Lower Trust in Media, Higher Trust in Government When Your Side Is in Power.” *Harvard Kennedy School Misinformation Review*, June 2020. *misinforeview.hks.harvard.edu*, https://doi.org/10.37016/mr-2020-024.

Pengelly, Martin. “Trump’s Obsession with Deep State Conspiracy ‘Delusional’, John Boehner Says.” *The Guardian*, 6 Apr. 2021. *The Guardian*, https://www.theguardian.com/us-news/2021/apr/06/john-boehner-book-trump-obsession-deep-state-conspiracy.

Pennycook, Gordon, Tyrone D. Cannon, et al. “Prior Exposure Increases Perceived Accuracy of Fake News.” *Journal of Experimental Psychology: General*, vol. 147, no. 12, 2018, pp. 1865–80. *APA PsycNet*, https://doi.org/10.1037/xge0000465.

Pennycook, Gordon, Adam Bear, et al. “The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings.” *INFORMS*, Nov. 2020. *dspace.mit.edu*, https://dspace.mit.edu/handle/1721.1/130380.

Pierre, Joe. “Illusory Truth, Lies, and Political Propaganda: Part 1.” *Psychology Today*, 22 Jan. 2020. *www.psychologytoday.com*, https://www.psychologytoday.com/blog/psych-unseen/202001/illusory-truth-lies-and-political-propaganda-part-1.

---. “Illusory Truth, Lies, and Political Propaganda: Part 2.” *Psychology Today*, 23 Jan. 2020. *www.psychologytoday.com*, https://www.psychologytoday.com/blog/psych-unseen/202001/illusory-truth-lies-and-political-propaganda-part-2.

“QuotePark.” *Quotepark.Com*, http://quotepark.com/quotes/1755767-abraham-lincoln-dont-believe-everything-you-read-on-the-internet/. Accessed 11 Apr. 2022.

Rivero, Nicolás. “The Risk of Putting Warning Labels on Election Misinformation.” *Quartz*, https://qz.com/1925272/the-risk-of-putting-warning-labels-on-election-misinformation/. Accessed 13 Dec. 2021.

Russell, Bertrand. *Free Thought and Official Propaganda*. https://www.gutenberg.org/files/44932/44932-h/44932-h.htm. Accessed 11 Apr. 2022.

Schaeffer, Katherine. "A Look at the Americans Who Believe There Is Some Truth to the Conspiracy Theory That COVID-19 Was Planned." *Pew Research Center*, https://www.pewresearch.org/fact-tank/2020/07/24/a-look-at-the-americans-who-believe-there-is-some-truth-to-the-conspiracy-theory-that-covid-19-was-planned/. Accessed 13 Dec. 2021.

Schoolov, Katie. "Why It's Not Possible for the Covid Vaccines to Contain a Magnetic Tracking Chip That Connects to 5G." *CNBC*, 1 Oct. 2021, https://www.cnbc.com/2021/10/01/why-the-covid-vaccines-dont-contain-a-magnetic-5g-tracking-chip.html.

Schulz, Kathryn. "The Really Big One." *The New Yorker*, 13 July 2015, https://www.newyorker.com/magazine/2015/07/20/the-really-big-one.

Sheth, Sonam. "'Our Task Was to Set Americans against Their Own Government': New Details Emerge about Russia's Trolling Operation." *Business Insider*, 17 Aug. 2017. *www.businessinsider.com*, https://www.businessinsider.com/former-troll-russia-disinformation-campaign-trump-2017-10.

Song, Yunya, et al. "The 'Parallel Pandemic' in the Context of China: The Spread of Rumors and Rumor-Corrections During COVID-19 in Chinese Social Media." *American Behavioral Scientist*, vol. 65, no. 14, Dec. 2021, pp. 2014–36. *SAGE Journals*, https://doi.org/10.1177/00027642211003153.

Suhay, Elizabeth, et al. "The Polarizing Effects of Online Partisan Criticism: Evidence from Two Experiments." *The International Journal of Press/Politics*, vol. 23, no. 1, Jan. 2018, pp. 95–115. *SAGE Journals*, https://doi.org/10.1177/1940161217740697.

---. "The Polarizing Effects of Online Partisan Criticism: Evidence from Two Experiments." *The International Journal of Press/Politics*, vol. 23, no. 1, 1, Jan. 2018, pp. 95–115. *SAGE Journals*, https://doi.org/10.1177/1940161217740697.

"The Manchurian Candidate Theory Will Never Die." *National Review*, 11 Nov. 2019, https://www.nationalreview.com/2019/11/the-manchurian-candidate-theory-will-never-die/.

Tompkins, Kelley A., and Joshua K. Swift. "An MTurk Primer for Psychotherapy Researchers." *Society for the Advancement of Psychotherapy*, 2019, https://societyforpsychotherapy.org/an-mturk-primer-for-psychotherapy-researchers/.

Unger, Craig. *American Kompromat: How the KGB Cultivated Donald Trump, and Related Tales of Sex, Greed, Power, and Treachery*. Dutton, 2021.

Vaccari, Cristian, and Andrew Chadwick. "Deepfakes and Disinformation: Exploring the Impact of Synthetic Political Video on Deception, Uncertainty, and Trust in News." *Social Media + Society*, vol. 6, no. 1, 1, Jan. 2020, p. 2056305120903408. *SAGE Journals*, https://doi.org/10.1177/2056305120903408.

Vallone, Robert P., and Lee Ross. *The Hostile Media Phenomenon: Biased Perception and Perceptions of Media Bias in Coverage of the Beirut Massacre*. 1985, p. 9.

Watts, Clint. "Is Trump Russia's Manchurian Candidate? No. Here's Why - FPRI." *Foreign Policy Research Institute*, 6 Mar. 2017, https://www.fpri.org/2017/03/trump-russias-manchurian-candidate-no-heres/.

"What Makes News Sources or Topics Trustworthy?" *American Press Institute*, 17 Apr. 2016, https://www.americanpressinstitute.org/publications/reports/survey-research/news-trust-source-topic/.

This thesis represents my own work in accordance with University Regulations.

/s/ Kevin Berry