

## Is Journalistic Truth Dead? Measuring How Informed Voters Are about Political News<sup>†</sup>

By CHARLES ANGELUCCI AND ANDREA PRAT\*

*To investigate general patterns in news information in the United States, we combine a protocol for identifying major political news stories, 11 monthly surveys with 15,000 participants, and a model of news discernment. When confronted with a true and a fake news story, 47 percent of subjects confidently choose the true story, 3 percent confidently choose the fake story, and the remaining half are uncertain. Socioeconomic differences are associated with large variations in the probability of selecting the true news story. Partisan congruence between an individual and a news story matters, but its impact is up to an order of magnitude smaller. (JEL D72, D83, L82)*

The news media play a key role in modern democracy by providing citizens with the information they need to keep government accountable. This effect has been documented in numerous contexts by a growing body of literature (see Eisensee and Strömberg 2007; Ferraz and Finan 2008; Gerber, Karlan, and Bergan 2009; Snyder and Strömberg 2010; Enikolopov, Petrova, and Zhuravskaya 2011; Banerjee et al. 2012; Kendall, Nannicini, and Trebbi 2015; Arias et al. 2018, 2022; Labonne, Cruz, and Keefer 2021; Chen and Yang 2019; Knight and Tribin 2022).<sup>1</sup>

In recent years, the increased awareness of the importance of the news media has been accompanied by widespread concerns about voters' information. A key issue has been the spread of political fake news (Lazer et al. 2018). Misinformation spread through social media has been identified as an important factor in elections

\*Angelucci: MIT, Sloan School of Management (email: [cangeluc@mit.edu](mailto:cangeluc@mit.edu)); Prat: Columbia University (email: [ap3116@gsb.columbia.edu](mailto:ap3116@gsb.columbia.edu)). Sylvain Chassang was the coeditor for this article. We thank Roland Bénabou, Alessandro Bonatti, Laurent Bouton, Julia Cagé, Teresa Esteban Casanelles, Simone Galperti, Leo Goldman, Andy Guess, Emeric Henry, Amit Khandelwal, José L. Montiel Olea, Suresh Naidu, Jacopo Peregó, Jett Pettus, Vincent Pons, Silvio Ravaoli, Miklos Sarvary, Andrey Simonov, Anya Schiffrin, Carlo Schwarz, Alistair Wilson, Noam Yuchtman, Sevgi Yuksel, and seminar audiences at Bocconi University, Brown University, UC Berkeley, UBC, Columbia University, CSEF, the University of Exeter, Harvard University, IESE, IIOC 2021, Indiana University, Johns Hopkins University, MIT Sloan, Northwestern University (MEDS), the NYC Media Seminar, the University of Oslo, Université Paris Saclay, Sciences Po Paris, the Stockholm School of Economics, University of Toronto Rotman, the fourth Summer Workshop in Political Economy at EIEF, and the Rational Inattention and Political Economy Conference at UC San Diego for useful suggestions and comments. We are grateful to Michel Gutmann and Szymon Sacher for outstanding research assistance. We gratefully acknowledge funding from the Media and Technology Program at Columbia Business School. This paper formerly circulated under the title “Measuring Voters’ Knowledge of Political News.” We obtained approval from the Columbia University IRB under protocol AAAS3426 and from the MIT COUHES (Exempt ID E-3899).

<sup>†</sup>Go to <https://doi.org/10.1257/aer.20211003> to visit the article page for additional materials and author disclosure statements.

<sup>1</sup>See Strömberg (2015) for a survey. Media bias also affects political outcomes (on this and related topics see, for instance, Mullainathan and Shleifer 2005; DellaVigna and Kaplan 2007; Gentzkow, Shapiro, and Stone 2015; Martin and Yurukoglu 2017).

around the world (Levitin 2016; Stengel 2019; Howard 2020). The potential effect of fake news is not limited to the spread of falsehoods. Vosoughi, Roy, and Aral (2018)<sup>2</sup> document how on Twitter “false news reached more people than the truth; the top 1 percent of false news cascades diffused to between 1,000 and 100,000 people, whereas the truth rarely diffused to more than 1,000 people.”<sup>3</sup> The spread of fake news could also have the indirect effect of blocking the distribution of accurate information (Meyer 2018). As information overwhelms voters, individual attention is the binding constraint, and we may be unable to acquire accurate information because we are drowning in an ocean of irrelevant or false information (Gleick 2011). These “flooding” tactics can even be used on purpose to suppress inconvenient truths (Tufekci 2017).

References to the “death of truth” and the “post-truth world” have become commonplace and are often accompanied by calls for immediate action to counter this risk (d’Ancona 2017; Ball 2018; Kakutani 2018; Baghini 2017). As Barack Obama put it, “One of the biggest challenges that we have to our democracy is the degree to which we don’t share a common baseline of facts.”<sup>4</sup> Indeed, a number of initiatives have been launched or proposed, including media literacy efforts, engagement programs, fact-checking platforms, and software solutions to block false statements, as well as legal reform and public funding of journalism (Schiffrin 2020; Lazer et al. 2018; Guess et al. 2020b). Constitutional scholars have suggested that we rethink the role of the First Amendment to restrict the dissemination of falsehoods (Bollinger and Stone 2018; Sunstein 2019).<sup>5</sup> Wu (2018, p. 556) argues that technology should make us reassess freedom of speech: the First Amendment was designed at a time when information was scarce, but this is no longer the case (“Gone are the days when the CBS evening news might reach the nation automatically, or whatever made the front cover of the *New York Times* was known to all.”); in a world where information flood is an important risk, the First Amendment may be “obsolete.”

Given the stakes involved, collecting as much evidence as possible about the “post-truth” phenomenon is crucial. Lately, the literature has focused on gathering systematic evidence of fake news’ impact on voters’ information (see Allcott and Gentzkow 2017; Barrera et al. 2020; Nyhan 2020; Guess et al. 2020a; Nyhan, Porter, and Reifler 2020). A number of papers have found that fake news may not be as widely circulated as initially feared (see Grinberg et al. 2019; Allen et al. 2020; Guess, Nagler, and Tucker 2019; Guess, Nyhan, and Reifler 2020) and that people can recognize their falseness (see Pennycook et al. 2021a,b; Pennycook and Rand 2019; Pennycook et al. 2020; Allen et al. 2021).

<sup>2</sup> Vosoughi, Roy, and Aral (2018, p. 1) define as news “any story or claim with an assertion in it,” and they restrict attention to the subset of these news that were classified as either true or false by six fact-checking organizations. As will become clear below, the set of true and fake news they analyze is different from that considered in this article.

<sup>3</sup> Quoted from *Science*’s online summary at <https://www.science.org/doi/10.1126/science.aap9559>.

<sup>4</sup> Interview with David Letterman in January 2018. Obama added, “We are operating in completely different information universes. If you watch Fox News, you are living on a different planet than you are if you, you know, listen to NPR.”

<sup>5</sup> Outside the United States, a number of regimes have invoked the death-of-truth narrative to pass laws against fake news that amount to censorship. In 2019, Russia passed a law to block social media that spread factually inaccurate information (as determined by the state). Turkey passed a similar law in 2020. On the role of information in autocracies, see Guriev and Treisman (2019).

Perhaps surprisingly, less systematic evidence on real news has been collected recently. How accurate is Obama's statement that we live in different political information universes and do not share a common baseline of facts? Given a set of recent important political events, how many voters are informed about them and can distinguish them from fake news? As we argue in the literature review, our work is unique because of its scale, because it combines best practices from a number of existing papers, and because it involves the estimation of a model of news discernment. As a result, our work provides a previously unavailable picture of mainstream political news.

Our methodology comprises three steps. The first step consists of identifying the "journalistic truth" on which individuals are tested. In our protocol, a panel of mainstream journalists select the three major domestic news stories of the past month among the set of Reuters wires related to federal politics. In addition, we employ two methods to select three fake domestic political news stories that could have happened within the same time frame as the real news stories. The first method consists of identifying three widely circulated fake news stories about the federal government by using Snopes (snopes.com), a leading fact-checking website. For the second method, we rely on our panel of journalists to write false statements about the federal government. We make no claim that the true and false news stories are selected or written objectively. Our only goal is transparency. The subjectivity in our protocol can be ascribed to a well-defined set of actors, namely news organizations like Reuters and Snopes, as well as a panel of professional journalists. Our exercise should be seen as an attempt to measure how discerning about real and false political news different segments of voters are, when employing the criteria used by mainstream journalists.

Second, we run surveys on the news selected according to the protocol above. We repeat the survey for 11 months on 14 different panels of approximately 1,000 US voters. In our main sample, respondents are selected by YouGov, a polling company, to produce a nationally representative sample of US adult citizens. As part of the survey, respondents take multiple financially incentivized quizzes. In each quiz, respondents are given three true statements and three false statements about recent events related to the federal government; they are asked to select the three statements they believe are most likely to be true. The three true statements correspond to the three most important news stories of the month according to our panel of journalists, and the three false statements correspond to the three fake news stories selected either according to Snopes or our panel of journalists.

In the third step, we use the data to estimate the parameters of a news discernment model. In our model, when an individual is confronted with a news story, she forms a *belief* about the story, defined—in a standard Bayesian way—as the probability she assigns to the story being true. The belief is a continuous variable with a value between zero and one that depends on (i) features of the story, like the easiness with which it can correctly be identified as true or false and its partisanship, and (ii) features of the individual, like her degree of discernment and her partisanship. The individual uses these probabilities to select the three stories she thinks are most likely to be true. Our methodology relies on techniques inspired by item response theory (see Bock 1972). The main adjustment is that while that theory accounts for only a "vertical" parameter (individual discernment), our response rates are also affected

by a “horizontal” parameter (partisan congruence between a respondent and a statement). The model yields a discrete choice specification that can be estimated with standard Maximum Likelihood Estimation methods. We estimate individual-level discernment parameters, but we also obtain estimates for the relative ease with which each news story can be correctly classified as true or false as well as other parameters of the model, such as the effect of time passing on discernment and the strength of partisan congruence. We use the model’s parameter estimates to analyze individuals’ ability to confidently distinguish true from fake news stories.

Our findings can be split into aggregate and disaggregate. Let us begin with the aggregate results, which paint a sobering picture. On average, we estimate that 82 percent of individuals will correctly identify the true news story when presented with a typical pair of recent true and fake news stories.<sup>6</sup> However, this figure does not provide a complete understanding. Therefore, we also estimate the percentage of individuals who have a high level of confidence in the true news story, a high level of confidence in the fake news story, and those who fall in between these two confidence levels. Most notably, we estimate that 47 percent of individuals assign 9 to 1 odds in favor of the true news story, 3 percent of individuals assign 9 to 1 odds in favor of the fake news story, and the remaining half are uncertain. When we lower the level of confidence, we estimate that 66 percent of individuals assign 3 to 1 odds in favor of the true news story, 8 percent of individuals assign 3 to 1 odds in favor of the fake news story, and the rest are uncertain.

These results suggest that journalistic truth is not dead. A majority of US voters confidently identify the news stories that journalists consider most important and confidently reject the widely circulated or well-crafted fake news included in our quizzes. However, there is also a sizable minority of people who are quite distant from journalistic truth because they are unable to confidently distinguish real and fake news.

At the disaggregate level, we try to understand which individual variables predict whether a subject will confidently identify real stories. First, according to the “parallel universe” hypothesis discussed above, ideological polarization should play a role. Partisan voters should be more likely to identify as real stories that are favorable to their side. Second, there could be socioeconomic variables that predict the ability of individuals to identify real stories in general. This is an important policy question because political accountability theory predicts that less informed segments of society are treated worse by officials with reelection concerns (Strömberg 2004), and evidence for this effect has been found in a number of contexts (see Snyder and Strömberg 2010).

Regarding the parallel universe hypothesis, we find that partisan congruence between a news story and an individual matters significantly in determining information levels. In a quiz containing a typical pair of true and fake news stories, we predict that a partisan individual is about 2 percentage points more likely to select a true news story if the news story reflects favorably on his or her preferred political

<sup>6</sup>In our main analysis, we do not distinguish between the fake news stories written by the panel of journalists and the fake news stories identified with Snopes. Our conclusions do not depend on the type of fake news stories used, as shown in online Appendix F.5.

party rather than unfavorably. The same individual is also about 4 percentage points more likely to assign 3 to 1 odds in favor of the true story.

Regarding socioeconomic variables, we find that the probability that a subject selects a true story is predicted—in order of size—by their age (above median), education (college or more), gender (male), income (above \$60,000), and ethnicity (White). These effects are highly robust to adding information about the subjects' political preferences or media consumption habits. Each of these effects is individually larger than the parallel universe effect measured above. For instance, the probability that a subject over 52 years of age selects a real story is 6 percentage points higher than that of a subject under 52, and the probability that she assigns 3 to 1 odds in favor of the true story is 9 percentage points higher.

When taken jointly, the overall effect of socioeconomic variables is almost one order of magnitude larger than the partisanship effect. The probability that an older, high-income, college-educated White man identifies a true story is 18 percentage points greater than that of an individual with the 5 complementary characteristics. That is nine times the parallel universe effect measured above.

In sum, our results indicate that the starkest pattern about the ability of US voters to identify major news stories is not a generalized death of truth or polarization along ideological lines. The starkest pattern appears to be inequality along socioeconomic lines. In the conclusion, we discuss how this finding calls into question some of the policies that have been proposed to combat the death of truth and instead suggests that researchers should explore why a significant proportion of the American public appears to be very distant in terms of information from mainstream journalism.

We perform a number of additional exercises in various extensions. We study the effect of time: every month that passes since news stories were written reduces by 2–3 percentage points the probability that an individual successfully completes a news quiz containing a typical pair of true and fake news stories and by about 4 percentage points the probability that he or she assigns 3 to 1 odds or higher in favor of the true story. We also use our model to predict how well individuals will perform on quizzes that feature micro-targeted fake news tailored to their socioeconomic and partisan characteristics for increased plausibility. We find that the effect of fake news targeting is significant, but it does not alter our main conclusions. Further, we exploit a survey we launched immediately before the 2020 presidential election to look at information differences between decided and undecided voters. We document that undecided voters are significantly less discerning about mainstream political news, raising concerns about the effects of misinformation campaigns targeted at pivotal voters. Further, as most of the news stories in our quizzes are related to the Donald Trump presidency, we perform an external validity exercise by looking at news stories about the Democratic Party primaries in the run-up to the 2020 presidential elections. Democratic primary voters are more likely to be young, female, minority, and low-income compared to presidential election voters (Kamarck and Podkul 2018). We again document large information inequalities. To get a benchmark outside of politics, we perform a similar exercise for news about sports and entertainment. Sports and entertainment compete for individuals' attention (Prior 2005) and are vulnerable to fake news (Pew 2019). Our estimates suggest that Americans are just as discerning about national politics as they are about sports and entertainment. Finally, we replicate our methodology using Amazon Mechanical Turk (MTurk)

and Ipsos. We show that the findings we obtain when relying on these alternative platforms line up with our main analysis.

Our results are subject to important caveats, which we discuss in depth in the paper. In particular, we only measure people's factual knowledge of important news stories, not how they interpret them. Also, our methodology allows us to conclude that most individuals seem able to confidently identify real news mainstream journalists consider important when the alternatives are widely circulated generic fake news or well-crafted but made-up generic fake news. We do not measure individuals' ability to identify less important real news and, similarly, although our findings are reassuring, we cannot rule out that some segments of the population are unable to identify the falsity of the fake news stories they are exposed to in their daily lives (and that were not included in our data). Finally, while we find that many individuals today are able to confidently identify important true news stories, we are unable to measure whether news discernment has decreased or increased in recent decades. In the conclusion we suggest possible avenues for future research.

The rest of the paper is structured as follows. Section II reviews the news-generating process and the survey design. Section III describes the model as well as our estimation approach. Section IV reports our main results. Section V presents various extensions of our analysis as well as robustness checks. Section VI states our conclusions.

## I. Literature Review

There exists a sizable literature on voters' information about political news that spans at least three decades. A partial list includes Price and Zaller (1993); Delli Carpini and Keeter (1996); Prior and Lupia (2008), and the references therein. Polling organizations regularly report survey results on voters' information (Pew 2017; Eurobarometer 2017).<sup>7</sup> Media outlets also sometimes feature news quizzes (e.g., the *New York Times*' News Quiz). Our survey differs from the literature in a number of important ways.

First, to the best of our knowledge, nearly all scholarly and commercial surveys are subject to the same limitation (we discuss some recent exceptions below): the true and fake news stories subjects are quizzed on are selected and usually written by the researcher and the pollster according to unspecified criteria. This observation—which should not be taken as a criticism of a literature that typically has different objectives—makes it difficult to use the results to answer questions about general news information levels.

We illustrate this challenge with a pioneering piece of work in this area, Price and Zaller (1993), which measures recall of 16 news stories. Restricting attention to political news, only 14 percent of participants remembered a congressional debate on catastrophic health insurance, while 65 percent recalled a trip by President Bush to Europe. This large variance affects the answer to the question posed above: “If a

<sup>7</sup>In turn, this literature is situated within a larger literature on voters' information not just about news stories but also about other political facts. For instance, Delli Carpini and Keeter (1996) also studies information about institutions and processes (e.g., “What is the purpose of NATO?” or “Name the three branches of government?”), information about key actors (e.g., “What is the name of the Secretary of State?”), information about statistical facts (“what is the percentage of population below the poverty line?”), knowledge of geography, and knowledge of social and political history. However, the present paper focuses exclusively on information about news stories.

major political event occurs, what percentage of Americans become aware of it?" If a researcher chooses stories like the congressional debate on catastrophic health insurance, the percentage is low. If the same researcher chooses stories like Bush's trip to Europe, the answer is much more encouraging. The percentage might be even higher, or lower, with other political stories that Price and Zaller (1993) did not include in their survey.

In a nutshell, the problem is that the set of possible news stories on which subjects can be tested is basically unbounded, and the outcome of the test depends on which true and fake news stories are selected. Without some discipline on the selection criteria, we can get extremely low response rates, extremely high response rates, and everything in between. This problem cannot be addressed after the survey is run because we do not know what other news stories the researcher considered but did not include, and we do not know whether and how those stories differed from the ones that ended up in the survey.

The news selection issue does not affect just the interpretability of the *absolute* value of the results. It also affects *comparisons* across different groups. What can we conclude if Democrats are more likely than Republicans to identify news stories selected by a particular researcher or pollster with criteria we do not know? Instead, in our setting, that finding would be interpreted as the average Democrat being more discerning when distinguishing between true and fake news stories selected by mainstream journalists. There would still be a crucial subjective component (maybe there is a liberal media bias) but one that can now be ascribed to a specific set of actors, can be further investigated, and can be compared to other similarly measured subjective biases.

Naturally, our approach is not the only method for creating a transparent protocol for selecting which news to quiz survey respondents on. Other recent studies on news discernment have used a combination of true news stories from mainstream sources and fake news stories from fact-checking websites (see Allcott and Gentzkow 2017; Pennycook and Rand 2019; Allcott et al. 2020; Mosquera et al. 2020; Pennycook et al. 2021a). For example, in their influential study of the 2016 US presidential election, Allcott and Gentzkow (2017) rely on the *Guardian*, Snopes, PolitiFact, BuzzFeed, and the *Wall Street Journal* for their selection of true and fake news stories. Our methodology similarly aims to identify important news stories as determined by mainstream media, and we use Reuters and Snopes for selecting true and fake news stories. However, we also utilized a panel of journalists to determine which news stories Reuters covered in the previous month were the most important and to write synthetic fake news stories in the same journalistic style as the true news stories. In Section II, we delve deeper into the role played by the panel of journalists. Additionally, we discuss the challenges of establishing an objective criterion for identifying fake news, which we addressed by utilizing two methods to generate fake news stories.

We also add to the existing literature in other important ways. Our research is distinctive in its scope, as we conducted 11 surveys over a period of nearly 3 years, which we administered to a total of about 15,000 individuals through 3 different platforms (YouGov, Ipsos, and MTurk). As already mentioned, we employed financial

incentives to promote accurate responses in our survey quizzes<sup>8</sup>, and we used two different methods to generate fake news stories to ensure that our findings were not specific to one type of fake news. Lastly, we estimate a model of news discernment to disentangle individual discernment from the effect played by partisan congruence and news story plausibility. Thanks to this set of features, our work provides insights about voters' information in the United States today that are not available elsewhere in the literature and allows us to assess the relative importance of factors that are associated with different information levels.

## II. Design

We review the protocol employed to generate the true and fake news stories we include in our quizzes, and we describe the information collected through the surveys.

### A. News Selection Process

Our objective is to determine how discerning about real political news US citizens are. To measure this knowledge using online surveys, we design quizzes in which real news stories are presented alongside fake news stories and in which participants are rewarded for identifying the real news stories.

*Real Political News.*—Given that the set of real political news is extremely large, how do we decide which real news to include in our quizzes? We design a protocol to identify, each month, the three most important news stories about the US federal government according to mainstream media. First, we rely on Reuters' publicly available wire stories about US national politics to approximate the universe of relevant mainstream news stories.<sup>9</sup> Second, we employ a panel of three professional journalists recruited through the Columbia School of Journalism.<sup>10</sup>

To avoid recency bias, we ask each journalist to choose, on a weekly basis, the top five wire stories they consider to be the most important from an editorial perspective. Specifically, journalists are provided with each wire story's headline, brief summary, and URL to the longer article. Because multiple wire stories can deal with the same underlying event, or "meta-story," we ask the journalists to select only one wire story per meta-story. In their weekly selection, we rely on journalists' subjective

<sup>8</sup>On the role of partisan congruence and incentives to recall information accurately, see Prior, Sood, and Khanna (2015) and Bullock et al. (2015). Both papers show that monetary incentives lead to less party cheerleader behavior in answering survey questions. On the effects of monetary incentives in surveys that measure political information, see also Prior and Lupia (2008).

<sup>9</sup>Reuters's wires dedicated to US national politics can be found at <https://www.reuters.com/news/archive/politicsNews>. There are approximately 80 wire stories a week. Because Reuters eventually moved behind a pay-wall, we relied on the Associated Press for our last survey in March 2022.

<sup>10</sup>In the course of the project, we have worked with four journalists in total. All journalists (three women and one man) are US citizens in their late 20s and early 30s who are currently working for mainstream media companies. An alternative to our "human" protocol would be an algorithmic approach, perhaps based on rankings produced by news aggregators such as Google News. However, any such approach would also ultimately build on the subjective views of the users whose data generate the ranking, and it would be less transparent, as neither the ranking algorithm nor the users' characteristics are known. If we chose stories on that basis, we would not know exactly whose subjective judgment we are relying upon.



assessment of whether two Reuters wire stories deal with the same underlying event. At the end of every month, we take the four previous weeks' selected wire stories and filter out the wire stories that do not cover the federal government.<sup>11</sup> We select a journalist to pool the remaining wire stories into their relevant meta-stories. We then present each meta-story and associated wire stories to our panel and ask them to select and rank the five most important meta-stories of the month. The choices are aggregated to produce the top three stories of the month. In online Appendix C, we compute a measure of interrater reliability and show that our journalists were typically in relatively strong agreement regarding their selection of the most important news of the month. Once the three stories are selected, a short statement about each story is written (e.g., *The US Senate acquitted Trump of impeachment charges*). To avoid confusing survey participants, journalists were asked to write primarily in the past tense and to avoid using numbers and figures.

*Fake Political News.*—In our quizzes, survey participants are asked to select three out of six news stories. Three of these news stories correspond to the three true news stories described in the previous paragraphs. The remaining three news stories are false short statements about the federal government.

What protocol should one employ to select the fake news stories? For the real news, we opted to focus on the news that mainstream journalists consider to be important for the general audience. No symmetric criterion exists to select fake news stories. We adopt two approaches based on two very different criteria.

In the first method, we relied on our panel of journalists to produce three plausible but entirely made-up short statements. One advantage of this approach is that we were able to instruct the journalists to write false statements of roughly equal length as the true statements and in the same journalistic style. We also instructed them to avoid writing negations of events that really took place or statements that could be perceived as related to the real statements. We refer to these stories as “synthetic” fake news stories.

In the second method, we rely on Snopes to provide us with three fake news stories about the federal government that they fact-checked during the four weeks that preceded any survey. Journalists at Snopes identify which rumors and questionable claims to fact-check and classify each as either false, mostly false, mixture, mostly true, or true. We asked Snopes to provide us exclusively with stories classified as false. The advantage of this second approach is that the corresponding false statements are more similar (in terms of tone, topics, etc.) to the fake news individuals might actually be exposed to in their daily lives. We refer to these stories as “actual” fake news stories.

To summarize, our method to select which real news stories and which fake news stories to include in our quizzes is geared toward measuring individuals' ability to identify mainstream journalistic truth, when the alternative to this truth is false statements generated using two distinct methods.

<sup>11</sup> We adopt the US definition of the “Federal Government” as being composed of the legislative, executive, and judicial branches.

TABLE 1—SURVEY PARTICIPANTS CHARACTERISTICS

<i>Panel A. Socioeconomic characteristics</i>		
Statistic	YouGov	ACS 2020
Median age	50	52
Percent Black	11	10
Percent White	70	73
Percent female	53	52
Percent 4-yr college degree	31	31
Percent married	47	53
Percent family income $\geq$ 60k	44	67
<i>Panel B. Party affiliations</i>		
Party affiliation	YouGov	Pew 2018
Percent Republican	25	26
Percent Democrat	35	33
Percent Independent	29	37
Percent other	11	4

*Note:* Full sample of YouGov participants who completed quizzes about the federal government (*Observations* = 8,437).

## B. Survey Design

This paper exploits data gathered from 11 online monthly surveys we conducted through YouGov, Ipsos, and MTurk. In total, 14,954 individuals participated in our surveys. The first survey took place in June 2019 and the last survey in March 2022. We focus on the nine surveys we ran through YouGov in our main analysis, and this section describes this sample only.

For each survey, YouGov enrolled a representative sample of the US citizen adult population.<sup>12</sup> All surveys were administered to about 1,000 individuals. We instructed YouGov to avoid enrolling individuals who participated in prior editions of the survey. This restriction was partly lifted from the third survey onward. Overall, 9,641 individuals participated in our 9 YouGov surveys (and 8,437 completed quizzes about the federal government). Our survey took respondents on average five to six minutes to complete. Participants received about \$1.9 on average (paid via gift cards). Payments included a \$0.50 show-up fee and bonuses worth \$1 for each quiz correctly answered.

YouGov provides a wide array of background information concerning each respondent, with the information collected months before our surveys. Table 1, panel A provides basic descriptive statistics regarding the socioeconomic characteristics of the survey participants we recruited through YouGov. It also reports the corresponding statistics for the population of US adult citizens according to the 2020 American Community Survey (ACS) of the Census Bureau (US Census Bureau 2020). All dimensions appear broadly aligned with the general population, with the exception

<sup>12</sup>To construct the sample, YouGov employs a two-step procedure. In the first step, a random sample is drawn from the population (using either census information or the American Community Survey). This sample is referred to as the target sample. In the second step, a matching technique is utilized to match each member of the target sample with members of YouGov's pool of respondents. For further details, see [https://smpa.gwu.edu/sites/g/files/zaxdzs4161/files/downloads/YG\\_Matching\\_and\\_weighting\\_basic\\_description.pdf](https://smpa.gwu.edu/sites/g/files/zaxdzs4161/files/downloads/YG_Matching_and_weighting_basic_description.pdf).

of family income.<sup>13</sup> Table 1, panel B reports information on the party affiliation of our survey respondents and compares it with the statistics provided by Pew (2018). The proportions are roughly comparable, with the exception of Independents, who appear somewhat underrepresented in the YouGov sample. For simplicity, in our main analysis we pool survey respondents who think of themselves as Independent with those who are “not sure” about the party with which they identify and those who identify with “something else” rather than the main parties. In the main analysis, we do not weight the survey data. In online Appendix F.1, however, we present our findings when weighting individual observations with the weights provided to us by YouGov to improve the sample’s representativeness.

Our survey was composed of two main parts: (i) a series of standard questions about media consumption habits (see online Appendix B.1 for details) and (ii) a series of questions about recent news. Online Appendix H presents the language used in our survey questions.

All surveys contained one or two quizzes on news stories from the previous month related to the federal government, as we described in the selection process. Survey respondents were told that the list contained exactly three true statements and three false statements “about recent events related to the Federal Government.” Respondents were asked to select which three, to the best of their knowledge, were the true statements. They were not told that the true statements had been selected by mainstream journalists nor that the true statements were chosen because of their editorial importance. News quizzes either included only synthetic fake news stories or only actual fake news stories (i.e., we never mixed synthetic and actual fake news stories). When surveys included both types of quizzes, the quizzes included the same three true news stories. Finally, in some surveys, we repeated prior editions of the news quizzes to study the effect of time passing. For a discussion of the advantages and disadvantages of our quiz design, see online Appendix G.

To prevent individuals from obtaining information elsewhere, respondents were given 60 seconds to make their selection.<sup>14</sup> We offered \$1 (paid via a gift card) to all respondents who selected all three true statements. All survey respondents were shown the correct answers once they took the quiz. Presumably because of the 60-second limit, 17 percent of respondents ended up selecting fewer than or more than 3 statements. The vast majority of these respondents selected strictly fewer than three statements. We exclude these respondents from our main analysis. As a robustness check, in online Appendix F.4 we reestimate the model by including respondents who selected fewer than or more than three statements.

Table 2 provides descriptive statistics at the statement level. On average, true news stories about the federal government were selected by 75 percent of respondents. The true news stories that were ranked as first news story of the month by our journalists were selected by 82 percent of respondents. By contrast, on average, both synthetic and actual fake news stories were selected by 25 percent of respondents.

<sup>13</sup>To the best of our understanding, the discrepancy between both family income figures is due to the categorical nature of the family income variable in the YouGov data.

<sup>14</sup>Imposing a time limit may lead us to underestimate respondents’ discernment (see Bago, Rand, and Pennycook 2020).

TABLE 2—NEWS STORY-LEVEL SUMMARY STATISTICS

	Mean	SD	Min	Max	Quizzes	<i>N</i>
Share of true news stories selected	0.75	0.15	0.36	0.97	12	23,688
Share of first true news stories selected	0.82	0.18	0.46	0.97	12	7,896
Share of synthetic fake stories selected	0.25	0.15	0.07	0.63	9	19,419
Share of actual fake stories selected	0.25	0.09	0.15	0.42	3	4,269

*Notes:* The table reports the average share of individuals who select each type of statement when completing the news quizzes about federal politics. These numbers are computed using our main sample of YouGov respondents who selected exactly three statements. The table distinguishes between true news stories, true news stories that were ranked as first news stories of the month by our panel of journalists, synthetic fake news stories, and actual fake news stories.

*Partisan Score.*—After completing a news quiz and being shown which statements were true and false, the respondents were asked to indicate how favorably each true statement reflected on the Republican Party, based on their opinion. Similarly, for each false statement, respondents were asked how favorably, in their opinion, the statement would have reflected on the Republican Party had it been true. Specifically, respondents were asked to select one option from the following scale: 1, very unfavorable; 2, unfavorable; 3, neither unfavorable nor favorable; 4, favorable; and 5, very favorable. Across all quizzes about federal politics, the average true statement had an average partisan score of 2.91 (standard deviation: 0.43). The average synthetic fake news story had an average partisan score of 2.92 (standard deviation: 0.40). Lastly, the average actual fake news story had an average partisan score of 3.13 (standard deviation: 0.36). In our analysis, we use the average partisan score  $b_j$  for each statement  $j$ , computed utilizing the full sample of YouGov survey participants. We standardize these scores using the mean and standard deviation derived from the  $b_j$  parameters of all true and all fake news stories included in our main sample.

*Extensions.*—We ran four surveys on samples of respondents recruited through MTurk and one survey on a sample of respondents recruited through Ipsos. Additionally, in five surveys we included news quizzes about the Democratic Party primaries, and in two surveys we included news quizzes about sports and entertainment. We provide greater detail in the relevant extensions. See also Table B.2 in online Appendix B.2 for descriptive statistics at the quiz and survey platform level, distinguishing between quizzes about the federal government, quizzes about sports and entertainment, and quizzes about the Democratic Party presidential primaries.

### III. Model

To make sense of the data we collect, we develop a theory of choice among news stories. Subjects are confronted with a set of stories and asked to select a subset of them. The model should allow for heterogeneity across subjects and across stories. It should allow for heterogeneity along an ideological dimension. In Section IIIA, we situate our model within the literature on Item Response Theory. In Section IIIB, we develop our model. In Section IIIC, we discuss estimation.

### A. Literature

Our model is related to Item Response Theory (IRT), a set of statistical models used to analyze test results with the objective of inferring the difficulty of the test questions and the traits of the test takers (Van der Linden and Hambleton 1997). IRT has been applied to analyze data from voter information surveys by, for instance, Baek and Wojcieszak (2009) and Anderson, Verkuilen, and Peyton (2010).

In standard IRT applications such as the Rasch model (Rasch 1960), the researcher can rank alternatives a priori (usually because an answer can only be right or wrong). Here, instead, we cannot a priori rank different statement bundles that contain different subsets of true statements. Suppose that A, B, and C are true statements and D, E, and F are false statements: it is not ex ante clear whether choosing, say, (A, B, D) is better than choosing (A, C, E). We are closest to an extension of IRT called Nominal Response Model (NRM), developed by Bock (1972), which allows items to be ranked in a partially unknown manner (see also Anderson, Verkuilen, and Peyton 2010).

One modification is necessary. We cannot use NRM directly because we are interested in measuring two factors: the underlying “skill” of our respondents (how discerning they are about journalistic truth) and the effect of partisan congruence. The latter effect is not salient in educational testing where only skill is measured. We therefore augment Bock (1972) by developing a model where individuals have two traits, discernment and partisanship, and news stories have two characteristics, difficulty to correctly classify as true or false and partisanship.

### B. A Discrete Choice Model

In the data, we will observe a set of subjects being asked to select three stories (henceforth, “statements”) out of a set of six stories. The model is developed in four steps: General belief formation; Selecting one statement out of a set of statements; Selecting three statements out of six; Implied choice probabilities.

First, we ask in general terms what is the probability that a subject assigns to a particular statement being true. Suppose agent  $i$  is trying to determine whether statement  $j$  is true ( $\tau_j = 1$ ) or false ( $\tau_j = 0$ ). Based on her information  $I_{ij}$ , she forms a belief about the truth of the statement  $q_{ij} = \Pr(\tau_j = 1 | I_{ij})$ . In what follows, we use a monotonic transformation of this belief, the log odds:  $z_{ij} = \ln(q_{ij}/(1 - q_{ij}))$ . The log odds  $z_{ij}$  is a random variable that we assume can be written as

$$(1) \quad z_{ij} := \underbrace{\gamma_j \theta_i \delta^t + \alpha b_j p_i + \lambda_i}_{\mu_{ij}} + \eta_{ij},$$

where  $\eta_{ij}$  has a standard Gumbel CDF.

The location of agent  $i$ ’s distribution of log odds (from now on, we refer to the log odds simply as “beliefs”) is determined by  $\mu_{ij}$ , which, in turn, depends on how discerning agent  $i$  is (captured by  $\theta_i \in \mathbb{R}^+$ ) and on how easily identifiable the truth or falsity of statement  $j$  is (captured by  $\gamma_j \in \mathbb{R}$ , which can be positive, usually for a true statement, or negative, usually for a false statement). Anticipating that  $\delta > 0$ , we note that larger values of  $\gamma_j$  shift agent  $i$ ’s distribution of beliefs to

the right, and the more so, the higher  $\theta_i$  is. The parameter  $\gamma_j$  thus also captures how discriminating statement  $j$  is: higher absolute values of  $\gamma_j$  lead to distributions of beliefs  $z_{ij}$  of agents with different levels of discernment  $\theta_i$  that lie farther apart.<sup>15</sup> Further, the joint effect of  $\theta_i$  and  $\gamma_j$  on agent  $i$ 's distribution of beliefs depends on the number of months  $t$  since statement  $j$  was written, via the parameter  $\delta$ . If  $0 < \delta < 1$ , the truth or falsity of statement  $j$  becomes less easily identifiable as time passes.

In addition, the location of agent  $i$ 's distribution of beliefs  $\mu_{ij}$  depends on the partisan congruence term  $\alpha b_j p_i$ , where  $b_j$  is statement  $j$ 's partisanship and  $p_i$  is agent  $i$ 's partisanship.<sup>16</sup> If  $\alpha > 0$ , agent  $i$ 's distribution of beliefs is shifted to the right (respectively, to the left) if statement  $j$  is congruent (respectively, incongruent). Lastly,  $\lambda_i$  is a parameter that shifts agent  $i$ 's distribution of beliefs independently of statement  $j$ 's characteristics. For instance,  $\lambda_i$  captures agent  $i$ 's systematic skepticism or credulity toward mainstream news.

Second, suppose now that agent  $i$  is asked to choose a statement  $j$  belonging to a set  $J$  of statements (all written  $t$  months ago) that is the most likely to be true according to (1). Assuming that  $\eta_{ij}$  is i.i.d. across statements, the probability agent  $i$  selects statement  $j$  as the most likely to be true among the set  $J$  of statements is<sup>17</sup>

$$(2) \quad \pi_{ij \in J} = \frac{e^{\mu_{ij}}}{\sum_{k \in J} e^{\mu_{ik}}} = \frac{e^{\gamma_j \theta_i \delta^t + \alpha b_j p_i}}{\sum_{k \in J} e^{\gamma_k \theta_i \delta^t + \alpha b_k p_i}}$$

Expression (2) is the probability that agent  $i$ 's belief in the truth of statement  $j$  is higher than those associated with the remaining statements. Accordingly, this probability is strictly increasing in  $\mu_{ij}$  and strictly decreasing in  $\mu_{ij'}$  (for  $j' \neq j$ ). Note that  $\lambda_i$  does not enter expression (2) because the parameter shifts the distributions of beliefs associated with each statement identically.

Third, our survey instrument is actually slightly more complicated than the choice problem in (2). In our survey quizzes, respondents read six statements, and they are given the additional information that exactly three statements are true. They are rewarded if they successfully select the three true statements. The solution to the respondents' problem involves the iterated application of expression (2). To see this, let  $\tau \equiv (\tau_1, \tau_2, \tau_3, \tau_4, \tau_5, \tau_6) \in \{0, 1\}^6$  denote a possible "truth vector" and  $T$  the set of all possible truth vectors before being told that exactly three statements are true. Let also  $T_3$  denote the subset of truth vectors such that three statements are true and

<sup>15</sup>In the model we estimate,  $\theta_i$  is allowed to be negative. If an agent is the opposite of discerning, larger values of  $\gamma$  for true statements and smaller values of  $\gamma$  for false statements make agent  $i$ 's inference problem more difficult instead of easier.

<sup>16</sup>Recall that we interpret  $b_j \in \mathbb{R}$  as the partisanship of the news story: a high (low)  $b_j$  denotes a story that reflects favorably (unfavorably) on the Republican Party. Similarly,  $p_i \in \{-1, 0, 1\}$  denotes agent  $i$ 's partisanship, where  $p_i = 1$  ( $p_i = -1$ ) means that agent  $i$  identifies with the Republican Party (Democratic Party) and  $p_i = 0$  means that agent  $i$  identifies as Independent. The term  $b_j p_i$  captures the congruence ( $b_j p_i > 0$ ) or incongruence ( $b_j p_i < 0$ ) between an individual's partisanship and a news story's partisanship. The parameter  $\alpha$  measures the strength of the effect of partisan congruence.

<sup>17</sup>The expression above holds under the assumption that the random variable  $\eta_{ij}$  is independent across statements. In practical terms, this means that the statements are not related in ways that make their plausibility value correlated. An obvious violation occurs when two statements refer to related stories "President Trump visited France" and "President Trump met with President Macron." We believe the independence condition is satisfied in practice, as both the true stories and the fake stories are designed to belong to distinct meta-stories (see Section II).

$\Pr(\tau \in T_3)$  denote the prior total probability that an individual assigns to three statements being true before being told that exactly three statements are true. If the subject selects statements  $j', j'', j'''$  as true and  $k', k''$ , and  $k'''$  as false, the subjective probability she wins the reward is equal to  $q_{ij'}q_{ij''}q_{ij'''}(1 - q_{ik'}) (1 - q_{ik''}) (1 - q_{ik'''}) / \Pr(\tau \in T_3)$ . This probability is maximized if and only if the agent chooses the statements with the three highest  $q_{ij}$ s.<sup>18</sup> Given our logit specification, the probability of selecting statements  $\{j, j', j''\}$  in this exact order is given by

$$(3) \quad \pi_{ij \in J} \cdot \pi_{ij' \in J \setminus \{j\}} \cdot \pi_{ij'' \in J \setminus \{j, j'\}}.$$

Let  $S(j', j'', j''')$  denote the set of all six possible permutations of statements  $j', j''$ , and  $j'''$  and  $s$  a typical element of  $S(j', j'', j''')$ . Let also  $\pi_{is}$  denote the probability that  $i$  selects the statements  $j', j''$ , and  $j'''$  in the exact order  $s$ . The total probability of selecting the statements  $j', j''$ , and  $j'''$  is equal to  $\Pi(j', j'', j''') \equiv \sum_{s \in S(j', j'', j''')} \pi_{is}$ . Our survey instrument generates draws over  $\Pi$ .

Fourth, although in our quizzes respondents were asked to select three statements out of six, we will use our parameter estimates to predict individuals' performance in counterfactual quizzes with only one true statement and one false statement. This approach will simplify the interpretation of our findings. Relatedly, we can use the expressions above to determine the probability that agent  $i$ —if given statements  $j$  and  $j'$  to read and told that exactly one statement is true—is willing to bet  $x$  against 1 dollars on statement  $j$  being true, which we denote  $\rho_{ij}(x)$ . Specifically, agent  $i$  is willing to bet  $x$  against 1 dollars on statement  $j$  in case the probability  $q_{ij}(1 - q_{ij'})$  she assigned to the event “ $j$  is true and  $j'$  is false” before being told that only one statement is true is  $x$  times larger than the probability  $q_{ij'}(1 - q_{ij})$  she assigned to the event “ $j$  is false and  $j'$  is true,” which, in turn, is equal to the probability that the difference between  $z_{ij}$  and  $z_{ij'}$  is equal to or larger than  $\ln x$ .<sup>19</sup> Therefore,

$$(4) \quad \rho_{ij}(x) = \frac{e^{\theta_i(\gamma_j - \gamma_{j'}) + \alpha p_i(b_j - b_{j'}) - \ln x}}{e^{\theta_i(\gamma_j - \gamma_{j'}) + \alpha p_i(b_j - b_{j'}) - \ln x} + 1}.$$

By extension, if given the freedom to choose which statement to bet  $x$  against 1 dollar on, agent  $i$  enters the bet with probability  $\rho_{ij}(x) + \rho_{ij'}(x) \leq 1$ . She does not enter the bet if she does not assign a high enough relative probability of truth to either statement. In this sense, we will rely on  $\rho_{ij}(x)$  to discuss our findings in terms of individuals' *confidence* in the truth of a statement versus another. We note that the parameter  $\lambda_i$  does not enter (4) for the same reasons as in expression (2): in principle, a discerning individual can assign a much larger probability of truth to a true statement than to a false statement and be skeptical or credulous toward both statements at the same time.

We end by noting that the model is silent on how beliefs are formed and why some subjects are more discerning than others. Obviously, some subjects may have

<sup>18</sup>This is equivalent to picking the statement with the highest  $q$  out of 6, then the statement with the highest  $q$  out of the remaining 5, and finally the statement with the highest  $q$  out of the remaining 4.

<sup>19</sup>To see this, note that  $\Pr(q_{ij}(1 - q_{ij'})/q_{ij'}(1 - q_{ij}) \geq x) = \Pr(\ln(q_{ij}/(1 - q_{ij})) - \ln(q_{ij'}/(1 - q_{ij'})) \geq \ln x) = \Pr(\eta_{ij} - \eta_{ij'} \geq \ln x - (\mu_{ij} - \mu_{ij'}))$ .

a higher  $\theta$  because they are more informed about the news. However, one can also imagine that general knowledge or general cognitive skills can help certain subjects assess the truth or falsity of the statements. The model is also silent as to why some statements might be more or less easy to identify as true or false. How easily identifiable as true or false a statement is may depend on how much media coverage the events described in the statement have received but possibly also on some of the statement's underlying characteristics (e.g., some of its institutional details or the importance of the events being described).

### C. Estimation

Due to well-known issues in identification of choice problems (e.g., only differences in  $\mu_{ij}$ 's matter; see, for instance, Train 2009), we must normalize the scale and the variance of  $\theta_i$ . Because the  $\gamma_j$ 's are unobservable, we also normalize, within each quiz, one  $\gamma$  to 1; we arbitrarily normalize the parameter associated with the most selected statement. Our results are presented in terms of choice probabilities, and these normalizations are therefore innocuous.

In online Appendix E, we show that the distribution of theta is nonparametrically identified subject to some technical conditions. However, for the purposes of estimation, we restrict  $\theta_i$  to be of the form  $\theta_i = 1 + \beta \mathbf{X}_i + \epsilon_i$ , where  $\mathbf{X}_i$ 's include socioeconomic covariates and  $\epsilon$  is normally distributed with mean zero and a variance  $\sigma^2$  to be estimated. The intercept of 1 is the normalization that addresses the fact that the scale and the variance of theta are otherwise unidentified. We assume that the population parameters  $\beta$ ,  $\sigma^2$ ,  $\delta$ , and  $\alpha$  are constant throughout our time period. In our main model, we include gender, age (lower than versus greater than or equal to median age of 52 across US adults), family income (lower than versus greater than or equal to \$60,000), education (bachelor's degree or more versus not), and ethnicity (White versus nonwhite) as socioeconomic variables. In the course of the analysis, we will present findings when estimating alternative versions of this model in which other socioeconomic characteristics are included in  $\mathbf{X}_i$ . We estimate the parameters of the model (namely,  $\beta$ ,  $\sigma^2$ ,  $\gamma$ 's,  $\delta$ ,  $\alpha$ ) by the method of Maximum Simulated Likelihood. In Section IVC we show that the model's predicted probabilities fit the data well.

*Discussion.*—We briefly outline the variation in the data that enables us to pin down the model's parameters. Essentially, each  $\beta$  parameter is determined by systematic differences in quiz performance across socioeconomic groups, while the variance of  $\epsilon$  captures quiz performance differences that the model cannot explain using predetermined socioeconomic characteristics. Additionally, the free  $\gamma$  parameters are pinned down by differences in selection shares across statements. The partisan congruence effect parameter  $\alpha$  is determined by differences between partisan and nonpartisan individuals when choosing between statements with different partisan scores. Finally, the  $\delta$  parameter is determined by differences in quiz performance for quizzes repeated over time.

In the theoretical model of Section IIIB, the parameters  $\theta_i$  and  $\gamma_j$  enter multiplicatively in equation (1). Allowing for  $\theta_i$  and  $\gamma_j$  to also enter additively would result in similar expressions for the choice probabilities, with slight changes in the



interpretation of these parameters.<sup>20</sup> The main model assumes that the strength of the partisan congruence effect is identical across all partisan individuals. Alternatively, we could have allowed Democrats and Republicans, or moderate partisans and strong partisans, to have different  $\alpha$  parameters. Similarly, we could have allowed for different  $\alpha$  parameters for true news stories and for fake news stories (and further distinguish between synthetic and actual fake news). Estimating these alternative specifications leads to very similar results. Further, the main model assumes that the  $\gamma$  parameters are identical across all individuals. We allow the  $\gamma$  parameter to vary by partisan and socioeconomic group in an extension presented in Section VA. We find that there exists strong agreement about the relative plausibility of true and fake news stories across groups, suggesting that the assumption we make in the main model is opportune. For simplicity, the model assumes that individuals are equally discerning across the various topics that make up mainstream political news. In extensions, however, we estimate individuals' degree of discernment about news that cover sports and entertainment as well as the Democratic Party primaries. By and large, the information patterns we uncover for these alternative topics mirror closely those we see for mainstream political news about the federal government. Although we cannot be certain, these findings suggest that we would obtain very similar results if we were to allow the  $\theta$  parameter to vary by topic within our main analysis.

## IV. Analysis

### A. Overview

We estimate the main model using the 12 quizzes (9 with synthetic fakes news stories and 3 with actual fake news stories) about the federal government that we administered through the 9 YouGov surveys. This approach, in turn, means that we need to estimate 63 statement-level  $\gamma_j$  parameters, in addition to 8 population-level parameters: the 5  $\beta$  parameters associated with the 5 socioeconomic factors that enter the individual discernment parameter  $\theta$ , the standard deviation  $\sigma$  of the random component  $\epsilon$  that also enters  $\theta$ , the partisan congruence parameter  $\alpha$ , and the time parameter  $\delta$ .<sup>21</sup> Estimating the main model yields that  $\sigma = 0.6$  (SE: 0.034),  $\alpha = 0.14$  (SE: 0.015), and  $\delta = 0.85$  (SE: 0.018). The last columns in Tables 6 and 7 list all the  $\gamma$  parameters, and the first column in Table 9 lists all the  $\beta$  parameters. Because the magnitude of these estimates is not easy to interpret, in what follows, we present our main findings by relying on the  $\pi$  and  $\rho$  functions discussed in

<sup>20</sup>To see this, suppose  $z_{ij} = \theta_i \gamma_j + \gamma_j + (2\tau_j - 1)\theta_i + \lambda_i + \alpha p_i b_j + \eta_{ij}$ . The expression for  $z_{ij}$  can be rewritten more compactly as  $z_{ij} = \tilde{\theta}_i \tilde{\gamma}_j - (2\tau_j - 1) + \lambda_i + \alpha p_i b_j + \eta_{ij}$ , where  $\tilde{\theta}_i = \theta_i + 1$  and  $\tilde{\gamma}_j = \gamma_j + (2\tau_j - 1)$ . It follows that, if  $j$  is a true statement and  $j'$  is a false statement, then  $z_{ij} - z_{ij'} = \theta_i(\tilde{\gamma}_j - \tilde{\gamma}_{j'}) - 2 + \alpha p_i(b_j - b_{j'}) + \eta_{ij} - \eta_{ij'}$ . In other words, allowing for main effects of  $\theta_i$  and  $\gamma_j$  would lead to identical functional forms (up to a constant) after appropriate changes of variables. Also, as noted above, the model we estimate is such that  $z_{ij} = \theta_i \gamma_j + \lambda_i + \alpha p_i b_j + \eta_{ij}$ .

<sup>21</sup>For the true statements that appear simultaneously in variants with synthetic fake news stories and in variants with actual fake news stories, we impose that their associated  $\gamma_j$  parameter is identical across the two variants. If we take the opposite approach (i.e., if we allow a true news story to have different  $\gamma$  parameters across variants), the  $\gamma$  parameters we estimate are very similar across variants.

TABLE 3—PROBABILITY OF SELECTING TRUE STORY

	All	First
$\bar{\pi}(\text{true} \mid 1 \text{ true, 1 false})$	0.82	0.85
$\bar{\pi}(\text{true} \mid 1 \text{ true, 3 false})$	0.63	0.70

*Notes:* The first row reports the probability that individuals select a true news story when faced with a typical pair of true and fake news stories. The second row reports the corresponding probability when individuals are faced with one typical true news story and three typical fake news stories. In the first column (All), the true news story is ranked as either first, second, or third news story of the month by the journalists. In the second column (First), the true news story is ranked as first news story of the month.

Section III. The methodology we employ to produce all the tables and figures shown below is described in online Appendix A.

### B. Aggregate Information Levels

We begin by analyzing individuals' ability to distinguish true and fake news stories. The first row in Table 3 reports the predicted probability  $\bar{\pi}$  that, on average, individuals select the true news story when given a typical pair of true and fake news stories (both less than a month old). On average, 82 percent of individuals select the true news story. This probability increases to 85 percent if the true news story was ranked as first news story of the month by the panel of journalists. The second row in Table 3 reports the predicted probability that individuals successfully complete a more difficult version of the quiz, in which only one out of four news stories is true. We find that, on average, 63 percent of individuals select a typical true news story and 70 percent select a typical first news story of the month.<sup>22</sup> In other words, as expected, individuals find it more difficult to identify the true news story when presented alongside multiple false statements.

Next, we investigate how confident individuals are when assessing the truth of true versus fake news stories. We again imagine that individuals are given a typical pair of true and fake news stories to read (both less than a month old). Respondents are told that exactly one statement is true and are given the opportunity to bet  $x$  against 1 dollars on a statement of their choice being true. We refer to the ratio of the probability that statement  $j$  is true to the probability that statement  $j'$  is true as the *odds of truth* in favor of statement  $j$ . Only individuals who assign  $x:1$  or higher odds of truth in favor of one of the two statements are willing to accept the bet.<sup>23</sup>

Table 4 reports, for various values of  $x$ , the probability  $\bar{\rho}_{\text{true}}(x)$  that an average individual assigns odds of truth equal to  $x:1$  or higher in favor of the true news story, the probability  $\bar{\rho}_{\text{false}}(x)$  that she assigns odds of truth equal to  $x:1$  or higher in favor of the fake news story, and the probability  $1 - \bar{\rho}_{\text{true}}(x) - \bar{\rho}_{\text{false}}(x)$  that she does not assign odds of truth equal to  $x:1$  or higher to either news story. To begin,

<sup>22</sup> As noted earlier, the methodology used to construct Table 3 and all other tables and figures is described in online Appendix A.

<sup>23</sup> For simplicity, we assume that individuals are risk neutral. We note, however, that the estimation of the model's parameters does not require us to take a stand on individuals' degree of risk aversion.

TABLE 4—PROBABILITY OF ASSIGNING FAVORABLE ODDS TO TRUE STORY, FALSE STORY, OR NEITHER STORY

Odds	Story rank	$\bar{p}_{true}$	$\bar{p}_{false}$	$\bar{p}_{no\ bet}$
9:1	All	0.47	0.03	0.5
	First	0.56	0.03	0.42
3:1	All	0.66	0.08	0.26
	First	0.72	0.07	0.22
2:1	All	0.72	0.11	0.17
	First	0.77	0.09	0.14

*Notes:* The table assumes that individuals are given a typical pair of true and fake news stories to read. It reports the probability  $\bar{p}_{true}(x)$  that individuals assign  $x:1$  or higher odds of truth in favor of the true news story, the probability  $\bar{p}_{false}(x)$  that they assign  $x:1$  or higher odds of truth in favor of the fake news story, and the probability  $1 - \bar{p}_{true}(x) - \bar{p}_{false}(x)$  that they do not assign  $x:1$  or higher odds of truth in favor of either news story. Three values of  $x$  are considered: 2, 3, and 9. For each value of  $x$ , the probabilities are reported assuming (i) a typical pair of true and fake news stories where the true news story is ranked as either first, second, or third news story of the month by the journalists (All); and (ii) a typical pair of true and fake news stories where the true news story is ranked first news story of the month by the journalists (First).

consider a relatively extreme bet that specifies  $x = 9$ . If given a typical pair of true and fake news stories, the model predicts that 47 percent of individuals will bet on the true news story, 3 percent will bet on the fake news story, and 50 percent of individuals will refuse to enter the bet. If the true news story is a typical first news story of the month, the corresponding probabilities become 56 percent, 3 percent, and 42 percent. If the bet instead specifies  $x = 3$ , the model predicts that 66 percent of individuals will bet on the true news story, 8 percent will bet on the fake news story, and 26 percent of individuals will refuse to enter the bet. If the true news story is a typical first news story of the month, the corresponding probabilities are 72 percent, 7 percent, and 22 percent. Finally, if the bet specifies  $x = 2$ , 72 percent of individuals will bet on the true news story, 11 percent will bet on the fake news story, and 17 percent of individuals will refuse to enter the bet. If the true news story is a typical first news story of the month, the corresponding probabilities are 77 percent, 9 percent, and 14 percent.

Regardless of the value of  $x$  we consider, a majority or close to a majority of individuals are willing to bet on the true news story. As expected, the share of individuals willing to enter the bet increases as  $x$  decreases, and a majority of the individuals who progressively enter the bet, bet on the true news story. For simplicity, in the remainder of the analysis, we will focus on bets that specify  $x = 3$  when presenting our findings in terms of  $\bar{p}$ . None of our qualitative conclusions depend on this particular choice of  $x$ .

Next, we look at heterogeneity across individuals. Specifically, we rank individuals by their level of discernment  $\theta_i$  and present results separately for the average individual in each tier of the distribution. Table 5, panel A returns to the hypothetical quiz with a typical pair of recent true and fake news stories. The average individual in the top third of the distribution is 13 percent more likely to select the true news story relative to the average individual in the bottom third. If the true news story is ranked as the first news story of the month, this difference in probabilities becomes

TABLE 5—HETEROGENEITY ACROSS DISCERNMENT TIERS

Story rank	Discernment tier		
	Lower	Middle	Higher
<i>Panel A. Probability of selecting true story</i>			
All	0.77	0.83	0.87
First	0.8	0.86	0.89
<i>Panel B. Probability of assigning favorable odds to true story</i>			
All	0.58	0.67	0.73
First	0.63	0.73	0.79

*Notes:* The top row (All) of panel A reports the probability that individuals in various tiers of the discernment distribution select a true news story when faced with a typical pair of true and fake news stories. The bottom row (First) of panel A reports the corresponding probabilities when the true news story is ranked as first news story of the month by the journalists. The top row (All) of panel B reports the probability that individuals in various tiers of the discernment distribution assign 3:1 or higher odds of truth in favor of a typical true news story, when the alternative is a typical fake news story. The bottom row (First) of panel B reports the corresponding probabilities when the true news story is ranked as first news story of the month by the journalists.

11 percent. Similarly, Table 5, panel B returns to the hypothetical bet with a typical pair of true and fake news stories and a loss/win ratio  $x$  equal to 3. The average individual in the top third of the distribution is about 26 percent more likely to bet on a typical true news story than the average individual in the bottom third of the distribution. The corresponding difference in probabilities when the true news story is ranked as the first news story of the month is 25 percent.

To summarize, we have shown that a sizable share of individuals confidently identify a recent true news story that mainstream journalists consider important when the alternative is a typical fake news story. Nonetheless, a significant share of individuals are somewhat unsure about the truth of true news stories, and a small but nonnegligible share of individuals find the fake news stories to be much more plausible than the real news stories.

### C. Heterogeneity across News Stories

Next, we explore heterogeneity across news stories. Tables 6 and 7 list, for each quiz separately, all the true and fake news stories that were included. The table distinguishes between synthetic and actual fake news when relevant. For each news story, the table reports the share of survey respondents who selected the statement when completing the quiz (Share), the standardized partisan score  $b$  given by the average respondent, the predicted  $\gamma_j$  parameter, and the predicted share of respondents who—according to our model's estimates—will select the statement when completing the quiz. In addition, the tables report the predicted probability  $\bar{\rho}(3)$  that an average respondent assigns 3:1 or higher odds of truth in favor of the statement (where the alternative is a false statement chosen at random from the same quiz if the news story is true or conversely, a true statement chosen at random from the same quiz if the news story is false).

As expected, within each quiz there exists significant heterogeneity across news stories. Some statements were selected by virtually all our respondents, and others

TABLE 6—NEWS QUIZZES JUNE 2019 TO MAY 2020

First survey date	Fake stories	Statement	Share	<i>b</i>	$\gamma$	$\bar{\pi}$	$\bar{\rho}$
June 2019	Synthetic	Alabama’s governor signed a bill to ban nearly all abortions in the state.	0.9	0.59	1.0	0.9	0.82
		Mexico agreed to take more migrants seeking asylum in the United States while they await adjudication of their cases.	0.7	1.38	0.28	0.71	0.58
		President Trump proposed plan to make US immigration more merit-based.	0.65	1.35	0.14	0.65	0.53
		US Border Patrol facility admitted to measles outbreak among migrant children in custody.	0.42	-0.18	-0.28	0.41	0.11
		Attorney General Barr released text message from Special Counsel prosecutor Robert Mueller: “We’re taking down Trump.”	0.19	-0.26	-0.99	0.19	0.05
		Trump administration to continue to allow US research using fetal tissue from abortions.	0.14	0.16	-1.34	0.14	0.04
Oct. 2019	Synthetic	Whistle-blower report complains of White House cover-up on Trump-Ukraine scandal.	0.9	-1.72	1.0	0.91	0.78
		Supreme Court granted a request by President Trump’s administration to fully enforce a new rule that would curtail asylum applications by immigrants at the US-Mexico border.	0.69	0.08	0.18	0.7	0.47
		At a closed-door meeting at the White House, top envoy to China delivered evidence of rising Farm Belt frustration over bio-fuel policy.	0.36	-0.71	-0.48	0.36	0.26
		Vaping case to make its way to Supreme Court.	0.45	0.44	-0.32	0.44	0.16
		President Trump announces he will resume peace talks with Iran at UN General Assembly.	0.37	0.98	-0.47	0.36	0.14
		China blacklists Apple and Microsoft amid escalating trade war.	0.23	-0.73	-0.84	0.24	0.1
Nov. 2019	Synthetic	A whistleblower filed a complaint against President Trump, leading to an impeachment inquiry.	0.91	-1.63	1.0	0.91	0.87
		Republican lawmakers in the House of Representatives condemned President Trump’s decision to withdraw troops from Syria.	0.69	0.38	-0.02	0.71	0.58
		The Trump administration credited cooperation from Mexico and Central American countries in cracking down on migrants.	0.62	0.91	-0.17	0.64	0.53
		China and the United States agreed on a new comprehensive trade deal.	0.52	1.73	-0.45	0.49	0.13
		ISIS beheaded three Americans in response to Al-Baghdadi’s death.	0.16	-0.94	-1.68	0.17	0.05
		President Trump’s Tax Returns showed billions given to various charities.	0.09	0.89	-2.65	0.09	0.03
Feb. 2020	Synthetic	The US Senate acquitted Trump of impeachment charges.	0.95	0.19	1.0	0.93	0.92
		Attorney General William Barr said that President Trump’s attacks on prosecutors, the judge and jurors in the trial of Roger Stone undermined the Justice Department’s work.	0.83	-0.76	0.05	0.83	0.69
		The House of Representatives passed legislation seeking to rein in President Trump’s ability to deploy US forces to fight abroad.	0.8	-0.27	-0.07	0.8	0.66
		President Trump took a week-long break from campaigning to deal with coronavirus outbreak.	0.24	0.85	-1.34	0.23	0.05
		Mitt Romney decided to run for president against Trump in the 2020 race after breakout role in impeachment.	0.12	-0.33	-2.18	0.12	0.03
		A tape surfaced of President Trump supporting abortion.	0.07	-1.19	-2.72	0.09	0.03
Apr. 2020	Synthetic	President Trump declared coronavirus a national emergency.	0.92	1.58	1.0	0.9	0.84
		President Trump notified Congress he is firing the inspector general of US intelligence community.	0.78	-1.01	0.37	0.79	0.64
		US Supreme Court allowed President Trump’s “Remain in Mexico” asylum policy.	0.64	0.29	0.08	0.67	0.54
		Agriculture trade group marched in Washington to draw attention to export problems.	0.3	-0.8	-0.59	0.3	0.08
		Nancy Pelosi under investigation by Justice Department over alleged insider trading during coronavirus outbreak.	0.22	0.37	-0.92	0.21	0.06
		President Trump fired coronavirus advisor Dr. Anthony Fauci.	0.14	-2.07	-1.43	0.14	0.04

(continued)

TABLE 6—NEWS QUIZZES JUNE 2019 TO MAY 2020 (*continued*)

First survey date	Fake stories	Statement	Share	$b$	$\gamma$	$\bar{\pi}$	$\bar{p}$
May 2020	Synthetic	President Trump said he would address national debt if re-elected.	0.71	1.18	1.0	0.69	0.58
		In win for President Trump, US Supreme Court made deporting immigrants for crimes easier.	0.66	0.79	0.97	0.67	0.57
		Senior US House members vowed to pass major defense bill despite pandemic.	0.64	0.37	0.87	0.62	0.53
		President Trump's campaign saw steep rise in donations after press conferences.	0.63	1.02	0.92	0.65	0.24
		Around 20% of IRS stimulus checks bounced.	0.18	-1.53	-0.25	0.2	0.05
		President Trump announced his tax returns will be released by Mid-May.	0.17	0.44	-0.33	0.17	0.04

*Notes:* The table lists, for each quiz separately, all true and fake news stories. For each quiz, the three top statements correspond to the true statements. For each news story, the table reports the share of survey respondents who selected the statement when completing the quiz (Share), the standardized average partisan score ( $b$ ), the predicted  $\gamma_j$  parameter ( $\gamma$ ), the predicted share of respondents who select the statement when completing the quiz ( $\bar{\pi}$ ), and the predicted probability that an average respondent assigns 3:1 or higher odds of truth in favor of the statement ( $\bar{p}$ ).

were selected only by a tiny share of respondents. Likewise, there exists substantial heterogeneity in the partisan scores of news stories. Many news stories reflected more favorably on the Republican Party than on the Democratic Party, and vice versa (recall that in the model, what matters is the difference between the partisan scores of news stories rather than the actual level of the scores).

Further, recall that the  $\gamma_j$  parameter captures how easily identifiable the truth or falsity of statement  $j$  is. What the tables suggest is that some true statements are much more easily detectable as true by discerning respondents than others. Similarly, some false statements are much more easily detectable as false by discerning respondents than others. Next, the tables report, for each statement, the model's predicted share of respondents who select it when completing the quiz (taking into account the characteristics of the remaining five statements included in the same quiz). As suggested by the numbers, our model approximates the actual data well, irrespective of whether a statement was chosen by few or many respondents.

Finally, there exists significant heterogeneity across both true and fake news stories in terms of respondents' probability of assigning 3:1 or higher odds of truth. Looking at real news, the probability that individuals assign odds of truth equal to 3:1 or higher ranges from 0.26 to 0.92. For example, the model predicts that 92 percent of respondents would assign odds of truth equal to 3:1 or higher to the (true) story "The US Senate acquitted Trump of impeachment charges." By contrast, 47 percent of individuals would assign odds of truth equal to 3:1 or higher to the true news story "Supreme Court granted a request by President Trump's administration to fully enforce a new rule that would curtail asylum applications by immigrants at the US-Mexico border" (despite 69 percent of our sample selecting the statement when completing the quiz). This last news story—with its relatively large difference between the share of respondents who select the statement and the share of respondents who assign it odds of truth of 3:1 or higher—illustrates how our approach takes into account the various properties of all the news stories included in the quiz

TABLE 7—NEWS QUIZZES OCTOBER 2020 TO MARCH 2022

First survey date	Fake stories	Statement	Share	<i>b</i>	$\gamma$	$\bar{\pi}$	$\bar{\rho}$
Oct. 2020	Synthetic	Trump Supreme Court pick Amy Coney Barrett pledged to follow law, not personal views.	0.89	1.32	1.0	0.88	0.76
		Second US presidential debate officially canceled after Trump balked.	0.62	-1.35	0.19	0.65	0.46
		Mitch McConnell avoided White House, citing laxity on masks, COVID-19 precautions.	0.5	-0.69	-0.08	0.5	0.37
		White House to host election night viewing party, Fauci calls it "potential disaster."	0.47	-1.56	-0.18	0.46	0.15
		President Trump tweeted about Black Lives Matter protests taking place in front of Mar-a-Lago.	0.36	-0.89	-0.38	0.35	0.12
		Kanye West called for special prosecutor if Biden elected.	0.16	-0.68	-1.11	0.17	0.06
		Trump Supreme Court pick Amy Coney Barrett pledged to follow law, not personal views.	0.87	1.32	1.0	0.88	0.74
	Actual	Second US presidential debate officially canceled after Trump balked.	0.61	-1.35	0.19	0.62	0.42
		Mitch McConnell avoided White House, citing laxity on masks, COVID-19 precautions.	0.48	-0.69	-0.08	0.47	0.32
		While speaking about Violent Crime Control and Law Enforcement Act of 1994, Joe Biden referred to Black Americans as "super-predators."	0.42	0.9	-0.15	0.43	0.15
		President Trump said: "The doctors said they've never seen a body kill the Coronavirus like my body. They tested my DNA and it wasn't DNA. It was USA."	0.35	-1.07	-0.35	0.34	0.12
		Democratic US presidential nominee Joe Biden said that he grew up in section 8 housing during town hall debate.	0.27	0.22	-0.57	0.26	0.1
		Joe Biden sworn in as US president.	0.94	-1.89	1.0	0.94	0.89
		US Senate Republican leader McConnell said Trump "provoked" January 6 riot.	0.74	-0.97	-0.14	0.76	0.58
Feb. 2021	Synthetic	Joe Biden said US coronavirus death toll to probably top 500,000 by end of February.	0.72	0.34	-0.23	0.71	0.56
		Biden in favor of temporarily barring guests from Capitol and other federal buildings.	0.39	0.38	-0.78	0.39	0.1
		Mike Pence revealed bombshell allegations in impeachment trial.	0.13	-0.66	-1.93	0.13	0.05
		Biden team's Twitter handle under fire after mistakenly reposting anti-Trump tweets.	0.07	1.71	-2.71	0.07	0.04
		Joe Biden sworn in as US president.	0.97	-1.89	1.0	0.95	0.9
		US Senate Republican leader McConnell said Trump "provoked" January 6 riot.	0.81	-0.97	-0.14	0.79	0.59
		Joe Biden said US coronavirus death toll to probably top 500,000 by end of February.	0.72	0.34	-0.23	0.75	0.56
	Actual	US Rep. Marjorie Taylor Greene said "If English was good enough for Jesus, it's good enough for us."	0.19	0.01	-1.39	0.2	0.06
		As of late January 2021, Donald Trump had started a new US political party called the "Patriot Party."	0.16	0.21	-1.57	0.17	0.05
		CNN issued a correction that read, "Sen. Ted Cruz was seen wearing a pin featuring a QAnon symbol. It was later discovered that this was not a QAnon pin, but a Doritos snack chip stuck to his suit."	0.15	0.8	-1.67	0.15	0.05

(continued)

when measuring respondents' confidence about the truth or falsehood of the stories they are given to read.

For fake news, there do not appear to be systemic differences across synthetic and actual fake news stories. Consistent with our earlier findings, whereby a large share of individuals can confidently discern true statements from false ones, no fake news story would have had more individuals assigning it 3:1 or higher odds of truth

TABLE 7—NEWS QUIZZES OCTOBER 2020 TO MARCH 2022 (*continued*)

First survey date	Fake stories	Statement	Share	$b$	$\gamma$	$\bar{\pi}$	$\bar{p}$
Mar. 2022	Synthetic	Biden nominates Jackson, first Black woman, to Supreme Court.	0.92	-1.34	1.0	0.92	0.86
		Zelenskyy pleads to US Congress: “We need you right now.”	0.82	0.14	0.58	0.85	0.74
		New Biden pandemic plan: Closer to normal for the nation.	0.65	-0.77	-0.06	0.63	0.52
		Harris celebrated Women’s History Month with girls at US/Mexico Border.	0.3	0.01	-0.6	0.3	0.07
		January 6 trials come to a halt amid Ukraine crisis.	0.18	0.9	-1.08	0.18	0.05
		Biden signed bill to mandate climate change curriculum in all K-8 classrooms.	0.13	0.06	-1.46	0.12	0.03
	Actual	Biden nominates Jackson, first Black woman, to Supreme Court.	0.9	-1.34	1.0	0.91	0.84
		Zelenskyy pleads to US Congress: “We need you right now.”	0.85	0.14	0.58	0.83	0.71
		New Biden pandemic plan: Closer to normal for the nation.	0.58	-0.77	-0.06	0.59	0.48
		Congress members awarded themselves a pay raise in 2022.	0.31	2.21	-0.53	0.31	0.08
		In March 2022, US Representative Paul Gosar defended himself from criticism with a tweet in which he indicated he had been called “stupid” for his whole life.	0.2	0.51	-0.92	0.2	0.05
		Former President Donald Trump’s “Truth Social” platform will cost users \$4.99 a week.	0.16	0.3	-1.11	0.16	0.04

*Notes:* The table lists, for each quiz separately, all true and fake news stories. For each quiz, the three top statements correspond to the true statements. For each news story, the table reports the share of survey respondents who selected the statement when completing the quiz (Share), the standardized average partisan score ( $b$ ), the predicted  $\gamma_j$  parameter ( $\gamma$ ), the predicted share of respondents who select the statement when completing the quiz ( $\bar{\pi}$ ), and the predicted probability that an average respondent assigns 3:1 or higher odds of truth in favor of the statement ( $\bar{p}$ ).

than any true news story would. Finally, the probability that individuals assign 3:1 or higher odds of truth to a fake news story ranges from 0.03 to 0.24. In Section VA, we exploit this heterogeneity to predict individuals’ performance when completing quizzes in which only the most plausible fake news stories are inserted.

#### D. Partisan Congruence and Time Passing

Do individuals exhibit a tendency to believe in or be better informed about news stories that reflect favorably on their preferred political party? If so, to what extent? When focusing on the most important political news of the month, how accurate is Obama’s statement that voters live in different political information universes and that they do not share a common baseline of facts? Individuals can hold partisan beliefs for a variety of reasons, including motivated beliefs (Bénabou and Tirole 2002, 2006) or selective exposure to news (see, for instance, Gentzkow and Shapiro 2011; Flaxman, Goel, and Rao 2016; and Sunstein 2018 on partisan news diets and echo chambers). Individuals can also (rationally) rely on their partisan priors when they are unsure about the accuracy of competing news sources, as in our news quizzes with true and false statements.

Table 8 considers an average partisan individual (either Republican or Democrat) and a pair of news stories composed of one nonneutral true news story and one neutral fake news story. The news stories are otherwise typical. The first column reports the average probabilities  $\bar{\pi}$  and  $\bar{p}(3)$  for various percentiles in the distribution of the true news story’s associated partisan score  $b_j$  (tenth, twenty-fifth, fiftieth, seventy-fifth, and ninetieth) by considering news stories that are less than four



TABLE 8—PARTISAN CONGRUENCE AND TIME PASSING

	Months passed					
	(a) $t = 0$		(b) $t = 1$		(c) $t = 2$	
	$\bar{\pi}$	$\bar{\rho}$	$\bar{\pi}$	$\bar{\rho}$	$\bar{\pi}$	$\bar{\rho}$
Story favorability						
Very unfavorable	0.8	0.63	0.77	0.58	0.74	0.54
Unfavorable	0.81	0.64	0.78	0.6	0.76	0.56
Neutral	0.82	0.66	0.8	0.62	0.77	0.58
Favorable	0.83	0.68	0.81	0.64	0.79	0.6
Very favorable	0.84	0.69	0.82	0.66	0.8	0.62

Notes: The first column ( $t = 0$ ) of the table reports the average probability  $\bar{\pi}$  that a partisan individual selects the true statement when faced with one true and one false statement (both less than one month old) by varying the favorability toward the individual's preferred party of the true statement and assuming a neutral false statement. The first column also reports the corresponding probabilities  $\bar{\rho}(3)$  of assigning 3:1 or higher odds of truth in favor of the true statement. The second ( $t = 1$ ) and third ( $t = 2$ ) columns report the same probabilities when the news stories are 5 to 8 weeks old and 9 to 12 weeks old, respectively.

weeks old.<sup>24</sup> Consider first *median* (i.e., typical) levels of congruence/noncongruence: news stories that reflect either favorably on the Democratic Party (twenty-fifth percentile in the distribution of  $b_j$ ) or favorably on the Republican Party (seventy-fifth percentile in the distribution of  $b_j$ ).<sup>25</sup> On average, partisan individuals are about 2 percent more likely to select the true news story and 6 percent more likely to assign it 3:1 or higher odds of truth if the true news story reflects favorably on their preferred party rather than unfavorably. Now consider *extreme* levels of congruence/noncongruence: news stories that reflect either very favorably on the Democratic Party (tenth percentile in the distribution of  $b_j$ ) or very favorably on the Republican Party (ninetieth percentile in the distribution of  $b_j$ ). On average, partisan respondents are about 5 percent more likely to select the true news story and 10 percent more likely to assign it 3:1 or higher odds of truth if the true news story reflects very favorably on their preferred party rather than very unfavorably.

The model allows time passing to matter in determining individuals' discernment about the news, via the decay parameter  $\delta$ . Intuitively, time may matter if individuals have limited memory and older news stories receive less media coverage. The last two columns of Table 8 repeat the same thought experiment as above by assuming that one month and two months have passed since both the true and the fake news stories were written. Time plays a significant role: looking, for instance, at neutral true news stories, every month that passes reduces the probability that individuals select the true statement by about 2–4 percent and the probability that they assign it odds of truth equal to 3:1 or higher by about 6 percent. Also, the strength of the partisan congruence effect increases with time passing. Considering true news stories with extreme levels of congruence/noncongruence, we find that, on average, a partisan individual is about 6 percent more likely to select the true news story and 14 percent more likely to assign it 3:1 or higher odds of truth if the story reflects very favorably compared to very unfavorably on their preferred party when the stories are

<sup>24</sup>Recall that the methodology used to construct the tables and figures is described in online Appendix A.

<sup>25</sup>Throughout, we rely on the bipartisan nature of American politics to assume that a story that reflects favorably on the Republican Party reflects unfavorably on the Democratic Party. Similarly, we assume that a story that “neither reflects favorably nor unfavorably” on the Republican Party is neutral in that it does not reflect either favorably or unfavorably on the Democratic Party either.

between 5 and 8 weeks old. The corresponding differences in probabilities are 8 percent and 15 percent when the news stories are between 9 and 12 weeks old. As time passes and information decays, individuals increasingly rely on partisan congruence when assessing news stories' truth or falsehood.

We have shown that partisan congruence plays a significant role in determining individuals' ability to confidently identify mainstream journalistic truth. These findings are in line with a rapidly growing literature documenting individuals' partisan perceptions of factual information related to topics ranging from income inequality to social mobility (see Kuziemko et al. 2015; Alesina, Stantcheva, and Teso 2018; Alesina, Miano, and Stantcheva 2020). Below, we show that differences along socioeconomic lines play an arguably even larger role in determining individuals' discernment about political news.

### E. Inequalities

There exists a literature documenting the relationship between media coverage and voters' information and, in turn, the relationship between voters' information and the attention voters receive from politicians. One channel through which this accountability channel operates is voting. If voters are informed about the policies implemented by politicians, the latter have greater incentives to cater to voters' preferences to increase their chances of reelection. Investigating how information levels vary across socioeconomic groups is therefore of interest: as politicians are likely aware of the link between information and voting, they have incentives to skew their policies toward the better-informed voters.<sup>26</sup> The literature on voters' information reviewed in Section I has highlighted the importance of socioeconomic factors in determining knowledge (see Verba, Lehman Schlozman, and Brady 1995; Nie, Junn, and Stehlik-Barry 1996; Delli Carpini and Keeter 1996; Prior 2007; Pew 2007). This section uses our methodology to quantify differences in discernment across socioeconomic groups.

Figure 1, panel A reports the association between the predicted probability  $\bar{\pi}$  of choosing the true news story when faced with a typical pair of true and fake news stories and our five key socioeconomic characteristics: age, gender, ethnicity, education, and family income (as well as various combinations of these five variables). Figure 1, panel B instead focuses on the predicted probability  $\bar{\rho}(3)$  of assigning 3:1 or higher odds of truth to the true news story.<sup>27</sup> Each single variable is associated with large differences in  $\bar{\pi}$  and  $\bar{\rho}(3)$ . For example, an average White individual is 6 percent more likely to identify the true news story and 13 percent more likely to assign it odds of truth equal to 3:1 or higher than an average minority individual. Similarly, an average individual aged 52 or more is 8 percent more likely to identify the true news story and 15 percent more likely to assign it odds of truth equal to 3:1 or higher than an average individual aged 51 or less. Combining these variables to

<sup>26</sup>In online Appendix D, we develop a model of retrospective voting in which various groups of voters differ in their policy preferences  $u_g(\cdot)$ , their size  $s_g$ , and information levels  $\bar{\rho}_g$ . An incumbent politician seeking reelection has incentives to allocate weights equal to  $(\bar{\rho}_g/\bar{\rho})s_g$  on the various groups of voters, where  $\bar{\rho}$  denotes the average voter's level of information, in contrast to the weights  $s_g$  a utilitarian social planner would use.

<sup>27</sup>To isolate the effect played by differences in discernment across socioeconomic groups, we shut down the partisan congruence channel by supposing that both the true and the fake news stories are neutral when constructing Figure 1, panels A and B. This choice explains why the population averages differ slightly with those presented in Tables 3 and 4.




FIGURE 1. SOCIOECONOMIC INEQUALITY AND PARTISAN CONGRUENCE

Notes: Panel A reports the probability that individuals belonging to various subgroups of the population select the true news story when faced with a typical pair of true and fake news stories. It also reports the probability that partisan individuals select the true news story when it is politically congruent and when it is politically noncongruent. Panel B reports the probability that individuals belonging to various subgroups of the population assign 3:1 or higher odds of truth in favor of the true news story when faced with a typical pair of true and fake news stories. It also reports the probability that partisan individuals assign 3:1 or higher odds of truth in favor of the true news story when it is politically congruent and when it is politically noncongruent. In both figures, the vertical bar corresponds to the population average.

create more granular socioeconomic groups leads to even starker differences. Both figures report the least and most discerning groups when we create socioeconomic groups by progressively combining two, three, four, and five variables. In the most extreme case, we find that older, high-income, White, college-educated men are 25 percent more likely to identify the true news story and 54 percent more likely

TABLE 9—SOCIOECONOMIC FACTORS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Democrat		0.097 (0.096, 0.098)	-0.022 (-0.023, -0.021)	0.163 (0.162, 0.165)	0.015 (0.013, 0.016)	0.01 (0.008, 0.012)	0.009 (0.006, 0.012)	0.064 (0.061, 0.067)
Republican		0.026 (0.025, 0.027)	-0.114 (-0.115, -0.112)	-0.135 (-0.136, -0.134)	-0.27 (-0.272, -0.268)	-0.286 (-0.288, -0.284)	-0.337 (-0.341, -0.334)	-0.25 (-0.254, -0.246)
Strong Partisan			0.21 (0.209, 0.211)		0.218 (0.217, 0.22)	0.202 (0.2, 0.204)	0.22 (0.217, 0.223)	0.112 (0.109, 0.115)
News Interest								0.734 (0.729, 0.738)
Age ≥ 52	0.368 (0.367, 0.37)			0.402 (0.4, 0.403)	0.389 (0.388, 0.39)	0.376 (0.374, 0.377)	0.447 (0.444, 0.45)	0.277 (0.274, 0.28)
Inc. ≥ 60k	0.198 (0.197, 0.199)			0.214 (0.213, 0.215)	0.209 (0.208, 0.211)	0.211 (0.209, 0.213)	0.305 (0.303, 0.308)	0.216 (0.213, 0.218)
College+	0.274 (0.272, 0.275)			0.273 (0.272, 0.275)	0.276 (0.275, 0.278)	0.303 (0.301, 0.305)	0.365 (0.362, 0.368)	0.277 (0.274, 0.28)
Female	-0.234 (-0.235, -0.233)			-0.264 (-0.265, -0.262)	-0.259 (-0.26, -0.258)	-0.266 (-0.267, -0.265)	-0.393 (-0.396, -0.391)	-0.343 (-0.345, -0.34)
White	0.264 (0.263, 0.266)			0.33 (0.329, 0.332)	0.326 (0.324, 0.327)	0.413 (0.411, 0.415)	0.527 (0.524, 0.531)	0.406 (0.402, 0.41)
Sources 3+						0.393 (0.391, 0.395)	0.696 (0.692, 0.701)	0.539 (0.534, 0.543)
Total Time (hrs)						0.022 (0.022, 0.022)	0.032 (0.032, 0.032)	0.028 (0.027, 0.028)
Observations	6,177	7,060	7,060	6,177	6,177	5,850	5,850	5,715
Extra media controls							X	X

*Notes:* The table reports the estimated  $\beta$  parameters as well as their associated 95 percent confidence intervals. Column 1 corresponds to the estimates obtained in the main model. The remaining columns correspond to variants of the model in which individuals' discernment parameter  $\theta$  is allowed to depend on various socioeconomic, partisan, and news consumption characteristics. See online Appendix A for a description of how confidence intervals are constructed. Strong Partisan is a dummy variable taking value 1 if individual  $i$  reports being either a Strong Republican or a Strong Democrat. Sources 3+ is a dummy variable taking value 1 if individual  $i$  reports relying on three or more news media outlets during the previous seven days. Total Time is the number of hours dedicated to consuming national news during the previous seven days reported by individual  $i$ . News Interest is a dummy variable taking value 1 if individual  $i$  reports being interested in general politics. Extra media controls include voter registration; indicators for using TV, print, online, and radio as a news source; as well as dummies for the ten biggest news sources interacted with using at least three sources. Media consumption questions were not included in every survey. See online Appendix B.1 for a description of news media consumption variables.

to assign it odds of truth equal to 3:1 or higher compared to young, low-income, minority women without a college education.

Information inequality across socioeconomic groups can also be quantified by examining the five estimated beta parameters that correspond to the five socioeconomic variables that enter an individual's degree of discernment,  $\theta_i$ . These parameters are reported in column 1 of Table 9. The parameter for age is equal to 0.368, which corresponds to roughly a 61 percent of a standard deviation positive shift in the mean of the discernment parameter,  $\theta_i$ .<sup>28</sup> The parameters for college education and income correspond to 46 percent and 33 percent of a standard deviation positive shifts, respectively. Finally, the parameter for being White corresponds to a 44 percent of a standard deviation positive shift, while the parameter for gender corresponds to a 39 percent of a standard deviation negative shift.<sup>29</sup>

<sup>28</sup> Recall that  $\theta_i$  is normally distributed with an estimated standard deviation of approximately 0.6.

<sup>29</sup> According to Lizotte and Sidman (2009), gender differences in measured political knowledge can be largely attributed to varying levels of risk aversion during quiz taking. However, our study is not affected by this factor, as all participants are required to select three statements.

Next, we report the estimated  $\beta$  parameters corresponding to versions of the model in which we progressively let the discernment parameter  $\theta_i$  depend on additional characteristics. Column 4 in Table 9 adds affiliations with the Democratic and the Republican Parties (the excluded category is Independents), and column 5 adds a variable measuring whether an individual feels either strongly Republican or strongly Democrat. The inclusion of these additional characteristics does not affect much the parameters associated with our five key socioeconomic variables. Focusing on column 5, the parameters associated with being a Democrat and a Republican are equal to 0.015 (2.5 percent of a standard deviation positive shift) and  $-0.27$  (45 percent of a standard deviation negative shift), and the parameter associated with being strongly partisan is equal to 0.218 (36 percent of a standard deviation positive shift). These estimates suggest that, on average, and other things being equal, Democrats have higher discernment about mainstream political news than Independents, who in turn have higher discernment than Republicans. In addition, being strongly partisan is positively associated with discernment.

We add media consumption habits (see online Appendix B.1 for details about these underlying variables) in columns 6 and 7. In both columns the self-reported number of news outlets and time usage (in hours) are positively associated with discernment about mainstream political news. In particular, relying on 3 or more news sources corresponds to about a 90 percent of a standard deviation positive shift in the mean of the discernment parameter  $\theta_i$ , and each additional hour of news consumption corresponds to a 5 percent of a standard deviation positive shift. Finally, column 8 includes self-reported interest in political news. Previous research has identified interest in politics as an important factor in news literacy (Prior 2007), and our results support this finding. The coefficient for general interest in political news is 0.734, corresponding to a 122 percent of a standard deviation positive shift in the mean of  $\theta_i$ . The strong association observed between self-reported news consumption and discernment about political news underscores the significance of the mainstream political news stories selected by our team of journalists.

#### *F. Taking Stock: Ideological Polarization versus Socioeconomic Inequality*

Our analysis has found that both ideological polarization and socioeconomic inequality explain the information patterns we observe. The rest of this section compares the magnitude of both factors. Ideological polarization could affect voter information about the news in three ways. First, it could be that the average Republican is systematically more discerning about mainstream news than the average Democrat, or vice versa. Column 2 in Table 9 reports the values of the parameters associated with being a Democrat and a Republican (the excluded category is Independents) in a version of the model in which the discernment parameter  $\theta_i$  depends exclusively on these two characteristics. We find that Democrats exhibit somewhat higher discernment than Republicans (the difference in means corresponds to about 12 percent of a standard deviation), who in turn exhibit a level of discernment only slightly higher than Independents (about 4 percent of a standard deviation). Although these effects are significant, we note that they are smaller than those associated with any of our five key socioeconomic variables (see column 4).

Second, it could be that more extreme voters are less informed than less extreme voters.<sup>30</sup> However, we actually find that the opposite is true. As discussed above, we find that individuals who feel either strongly Republican or strongly Democrat exhibit higher levels of discernment about mainstream political news (compared both to Independents and to individuals who feel only moderately Republican or Democrat).

Third, it could be that polarization takes the form suggested by many commentators, including Obama: voters on different sides of the political spectrum have different information. Table 8 has provided evidence of such a partisan congruence effect. For ease of comparison, Figure 1 replicates the main findings from Table 8. Specifically, Figure 1, panel A reports the probability that an average partisan individual (either Democrat or Republican) selects the true news story when faced with a typical pair of true and fake news stories, assuming that the true news story exhibits either a median level of congruence or a median level of noncongruence. Figure 1, panel B repeats the same exercise by focusing on the probability  $\bar{\rho}(3)$  of assigning 3:1 or higher odds of truth to the true news story. On average, partisan respondents are about 2 percentage points more likely to select the true news story if it reflects favorably rather than unfavorably on their preferred party and about 4 percentage points more likely to assign it odds of truth equal to 3:1 or higher. Although large, the magnitude of these partisan congruence effects is smaller than the effect associated with any socioeconomic factor in isolation. The effect of partisan congruence is also significantly smaller than the effect produced when combining socioeconomic characteristics to create more granular socioeconomic groups. Taking the most extreme example, the increases in  $\bar{\pi}$  and  $\bar{\rho}(3)$  from belonging to the most versus the least informed socioeconomic groups (older, high-income, White, college-educated men versus young, low-income, minority women without a college education) are, respectively, 9 and 6.75 times larger than the effect played by partisan congruence.

To summarize, our results indicate that socioeconomic differences appear to play as large, if not larger, a role as ideological polarization in explaining factual news information. We discuss some policy implications in Section VI.

## V. Extensions and Robustness Checks

### A. Fake News Targeting

We have previously assumed that the fake news stories included in the hypothetical quizzes and gambles were randomly selected. In practice, fake news spreaders, both human and automated, exploit social media algorithms and network structures to spread misinformation. They manipulate these tools to promote the most popular and convincing fake news stories and to target specific audiences with their lies (Howard 2020). This section delves into these issues by using two counterfactual exercises, whose results appear in Table 10.

To facilitate comparison, the first row in Table 10 reproduces our main aggregate results from Section IV: it reports the probabilities  $\bar{\pi}$ ,  $\bar{\rho}_{true}$ ,  $\bar{\rho}_{false}$ , and  $1 - \bar{\rho}_{true} - \bar{\rho}_{false}$

<sup>30</sup>On this issue, the existing evidence—which tends to measure political knowledge rather than knowledge of political news specifically—is mixed (Palfrey and Poole 1987; Delli Carpini and Keeter 1996; Prior 2007).

TABLE 10—FAKE NEWS TARGETING

	$\bar{\pi}$	$\bar{\rho}_{true}$	$\bar{\rho}_{false}$	$\bar{\rho}_{no\ bet}$
Random fake news	0.82	0.66	0.08	0.26
Untargeted hardest fake news	0.74	0.53	0.12	0.35
Targeted hardest fake news	0.73	0.52	0.12	0.36

Notes: This table considers quizzes that include one true and one fake news story. It reports the probability  $\bar{\pi}$  that an average individual selects the true news story, the probability  $\bar{\rho}_{true}$  that they assign 3:1 odds of truth or higher to the true news story, the probability  $\bar{\rho}_{false}$  that they assign 3:1 odds of truth or higher to the fake news story, and the probability  $1 - \bar{\rho}_{true} - \bar{\rho}_{false}$  that they do not assign 3:1 odds of truth or higher to either news story. In the first row, true and fake news stories are chosen at random. In the second row, the true news story is chosen at random, but only the most plausible fake news story is included. In the third row, the true news story is chosen at random, and the fake news story is individually targeted on the basis of socioeconomic and partisan characteristics.

when individuals are given a random pair of true and fake news stories. Table 10’s second row shows the corresponding probabilities under our first counterfactual scenario, where individuals are faced with a random true news story and its most credible fake counterpart written in the same month, as determined by the highest estimated  $\gamma_j$  parameter. Considering the most plausible fake news story as opposed to a random fake news story has a large impact. For example, this type of targeting leads to a 10 percent decrease in the average probability of selecting the true news story and a 20 percent decrease in the probability of assigning it odds of truth equal to 3:1 or higher.

The third row of Table 10 shows the results from our second counterfactual scenario, where individuals are faced with micro-targeted fake news stories. Specifically, we estimate a variant of our model with  $\gamma$  parameters that vary by predetermined socioeconomic and partisan group. For each individual in the sample, we compute the  $\pi$  and  $\rho$  functions when presented with a randomly chosen true news story and the predicted most plausible fake news story for this individual, based on the estimated  $\gamma$  parameters as well as partisan congruence.<sup>31</sup> Somewhat strikingly, exploiting individuals’ partisan and socioeconomic characteristics has an impact only slightly larger than when simply including the overall most plausible fake news stories: for instance, it reduces the probabilities  $\bar{\pi}$  and  $\bar{\rho}_{true}$  by only 1 extra percentage point. This relatively small effect occurs because strong agreement exists across socioeconomic and partisan groups about the relative plausibility of our fake news stories.

To summarize, when only the most plausible fake news stories are considered, we find that the population is almost exactly split into two halves: 52 percent of individuals confidently identify mainstream journalistic truth, and 48 percent either struggle or fail to do so. We also find that, at least in our data, targeting fake news on the basis of socioeconomic characteristics and partisan congruence does only marginally better than simply selecting the overall most plausible fake news. Although some caution is needed, this finding suggests that limiting fake news purveyors’

<sup>31</sup> To separately identify the individual parameter  $\theta_i$  and the group parameter  $\gamma_{g(i)}$  (where  $g(i)$  is individual  $i$ ’s group), the characteristics used to define a group are distinct from the five socioeconomic variables that enter  $\theta_i$ . See online Appendix A for more details.

micro-targeting abilities, such as making it harder to track individual engagement data, may only have limited effectiveness.

### *B. Information about Political News and Voting Intentions*

In this section, we focus on the period leading up to the 2020 presidential election to document differences in discernment about mainstream political news between (i) individuals who planned to vote in the election and those who did not and (ii) individuals who had a preference for the candidates and those who did not. To do so, we take advantage of questions about voting intentions we inserted in a survey run only a few days before the 2020 election.<sup>32</sup>

Among our respondents, 56 percent indicated that they had already voted either by mail or in person by the time they completed our survey, and 23.5 percent claimed that they were “definitely” going to vote. Figure 2, panel A shows the relationship between our five key socioeconomic variables and self-reported voting intentions.<sup>33</sup> Comparing Figure 2, panel A and Figure 1 reveals that the factors that predict an individual’s discernment about political news also predict their decision whether to vote. In other words, those who intend to vote tend to be more discerning about mainstream political news than the general population, while those who abstain tend to be less discerning.

Next, we investigate the relationship between individuals’ level of discernment about political news and whether they have a preference over the running candidates. Exploiting cross-country survey data, Le Pennec and Pons (2023) document that possibly as many as 29 percent of voters are undecided whom to vote for 60 days before an election. They also present evidence suggesting that many undecided voters eventually make up their minds by gathering information in the final stages of campaigns and that this process has a potentially significant large impact on final vote shares. For this reason, undecided voters are likely to be pivotal, and their level of discernment about political news is of particular interest. Among the respondents who had not voted yet at the time of our survey, 20.1 percent reported that they were unsure which candidate they would vote for if they were to vote. Figure 2, panel B reports the association between our five key socioeconomic variables and the probability of expressing a preference over the candidates.<sup>34</sup> Comparing Figure 2, panel B

<sup>32</sup>The 2020 United States elections were held on Tuesday, November 3, 2020. The survey was run between October 28, 2020 and November 2, 2020 through YouGov, and it included 1,600 participants in total. See online Appendix B.3 for further details and summary statistics.

<sup>33</sup>We code an individual as being a voter if they self-report having already voted or claim that they are definitely going to vote. As in most surveys, even such a conservative approach yields a predicted turnout rate that considerably exceeds the actual turnout rate (in our case, 79.5 percent versus 68 percent). The difficulty lies in that survey questions about voting intentions cannot easily be incentivized. Our findings would remain essentially unchanged if we were to define voters differently (e.g., also include those who state that they are “likely” going to vote).

<sup>34</sup>In Figure 2, panel B, we include all respondents who state that they have not voted at the time of the survey (and for whom information about our five socioeconomic variables is available). We are confident that our findings apply to the subset of individuals who actually choose to vote for three reasons. First, the relationship between each socioeconomic factor and vote certainty is strikingly similar to that found in Le Pennec and Pons (2023), who instead consider only voters (see Table A.6 in the online Appendix to Le Pennec and Pons (2023) for a direct comparison). Second, as a robustness check, we replicated our analysis using the 2020 Cooperative Election Study (Schaffner, Ansolabehere, and Luks 2021), which was administered by YouGov on a nationally representative sample of 61,000 individuals. We find very similar results even when restricting attention to voters. Lastly, we obtain similar results even if we restrict the sample to the individuals who state that they are “definitely” going to vote.



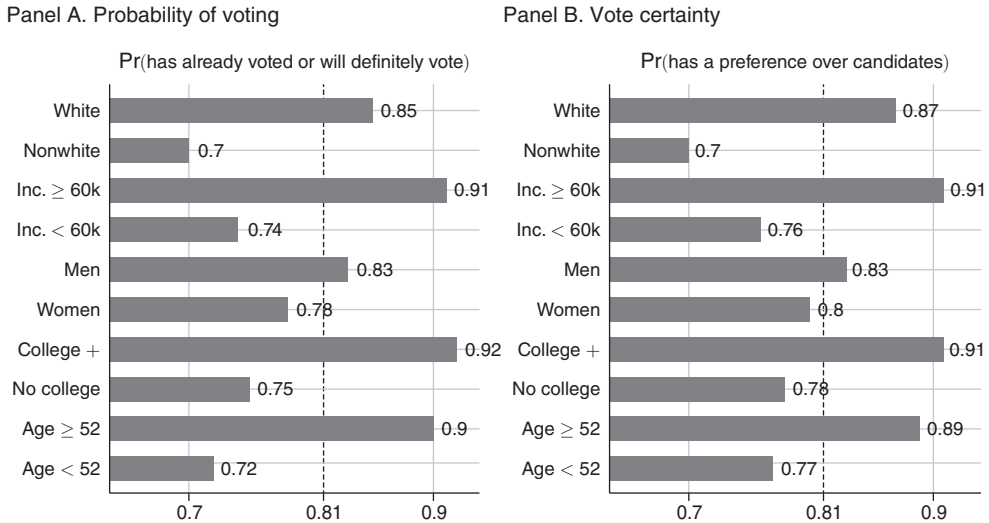


FIGURE 2. VOTING

Notes: Panel A reports the probability that individuals belonging to various subgroups of the population vote in the 2020 presidential election. We define as voters survey participants who either state that they have already voted at the time they complete the survey or state that they are “definitely” going to vote. Panel B restricts attention to survey participants who do not report having voted at the time they complete the survey and reports the probability that individuals belonging to various subgroups of the population express a preference for a candidate.

to Figure 1 reveals that the factors predicting individuals’ level of discernment about political news also strongly predict whether individuals have a preference over the candidates.

Noting that most of our surveys were conducted outside of electoral periods, our findings suggest that voters who have clear preferences for candidates early on in a campaign tend to consistently show high levels of discernment about political news, while those who are uncertain tend to be less discerning. This finding is noteworthy because uncertain voters may be more susceptible to targeted misinformation campaigns and, as noted, can often be pivotal in elections (Howard 2020).

### C. Sports and Entertainment

The “post-truth” phenomenon is often discussed with respect to political news. We sought a benchmark outside politics and included news quizzes about sports and entertainment in two surveys (four quizzes in total, with synthetic fake news stories only). Online Appendix F.2 presents our main aggregate results when estimating the model on these data exclusively. US citizens appear just as discerning about mainstream political news as they are about sports and entertainment.

### D. Democratic Party Presidential Primaries

Most of our news quizzes about the federal government included news directly related to the Donald Trump presidency. We cannot exclude the possibility that the

sizable differences in information levels across socioeconomic groups that we find are somehow driven by his four-year tenure in the White House. In five surveys, we included news quizzes devoted to news about the Democratic Party presidential primaries. Noting that Democratic primary voters are more likely to be young, female, minority, and low-income compared to presidential election voters (Kamarck and Podkul 2018), this extension allows us to perform an external validity exercise. In online Appendix F.3, we present our findings when estimating the main model using the quizzes about the Democratic Party presidential primaries exclusively. We again find evidence of large information inequalities along socioeconomic lines, with a much lower effect played by partisan congruence.

### *E. Robustness Checks*

In our main analysis, we excluded the 17 percent of respondents who selected fewer than or more than 3 statements when completing the quizzes. This exclusion could potentially bias our results. In online Appendix F.4, we replicate our analysis by including these respondents. We impute the missing choices of the respondents who selected fewer than three statements by randomly choosing unselected statements. For those who selected more than three statements, we remove at random some of the selected statements to arrive at three choices. As expected, this leads to a slight decrease in aggregate discernment about mainstream journalistic truth. However, our main conclusions remained unchanged. Further, for simplicity, the main analysis combined quizzes with synthetic and actual fake news stories. In online Appendix F.5, we estimate the model separately on both types of quizzes and show that our findings do not depend on the type of fake news employed.

Finally, we took several steps to alleviate concerns that the YouGov pool of participants is not representative of the US population. First, in online Appendix F.1 we present our main findings when aggregating observations using the weights provided by YouGov. While weighting observations changes the composition of our sample slightly, it has only minor impacts on the information patterns we document. Second, we present our findings when we estimate the model using data from the four surveys administered through MTurk and find that the results are largely in line with those from the YouGov sample (see online Appendix F.6). Third, we replicate our analysis using data from a survey administered through Ipsos in March 2022 (see online Appendix F.7). Ipsos uses a different method to construct its panel of respondents (probability-based sampling) compared to YouGov, but our results are very similar to those with the YouGov data.

## **VI. Concluding Remarks**

The debate on political news has centered around “the death of truth” and the existence of “parallel universes,” sometimes leading to urgent calls for drastic reforms. Our work casts doubt on this narrative. The death of truth is far from a universal phenomenon. A majority of Americans have a solid grasp of real news stories and can distinguish them confidently from well-crafted or widely circulated fake news. Regarding parallel universes, ideological polarization matters, but it is not—by a margin—the major explanatory factor. Our findings indicate that the starkest pattern

about the ability of voters to identify major news stories is its unequal distribution along socioeconomic lines. A sizable minority of US voters have low information levels; these voters are much more likely to be found among the young, the poor, women, the less educated, and minorities.

Our work creates a bridge with an older sociology literature that documents a strong heterogeneity in informational access of citizens. Gaxie (1978, p. 45–46) speaks of a “hidden class-based disenfranchisement,” as poorer, less educated citizens are de facto excluded from the political debate because of their lack of information. Our data also confirm that the conclusions of Delli Carpini and Keeter’s (1996) extensive analysis of US voters’ information are still valid in the 2020s and with newer methodologies. Socioeconomic variables—age, gender, income, education, race—still determine who the informed and the uninformed are in America.

Informational inequality along socioeconomic lines poses crucial policy issues. As discussed above, politicians with reelection concerns have incentives to neglect the preferences of the less informed segments of society (see Strömberg 2004; Snyder and Strömberg 2010, and the references listed in the introduction). In addition, the evidence presented in Section VB, together with findings from Le Pennec and Pons (2023), suggests that undecided voters, who often play a crucial role in elections, belong to the same socioeconomic groups that tend to be less discerning about political news, making them vulnerable to targeted misinformation campaigns.

Our paper raises a fundamental question about voter information: Where do these large inequalities come from? One can advance three possible hypotheses. First, it could be that we chose a political topic—domestic politics—that is much more interesting to certain segments. However, this explanation is not promising because we find the same inequality patterns when we investigate news on the Democratic Party primaries. Second, it could be an economics-driven access explanation. Certain segments of the population have less disposable income to spend and are therefore less likely to consume costly news sources. However, this is unlikely to be the whole explanation because inequality patterns survive after controlling for income.

A more plausible explanation is that mainstream journalistic truth about domestic politics—at least in its current form—is more appealing to certain socioeconomic groups than to others. In turn, this explanation may be due to demand and supply factors. A possible demand factor is that the opportunity cost of devoting time to political information is higher for some socioeconomic groups. On the supply side, a factor that deserves further scrutiny is that news producers tend to belong to the high-information groups: the newsroom staff at major national newspapers is overwhelmingly male (WMC 2019) and White (Arana 2018).

Our findings suggest that any policy reform should address informational inequality along socioeconomic lines. Examining the role that organizations, such as political parties, unions, and nongovernment organizations, can play in mitigating this problem would be beneficial.<sup>35</sup> Further, exploring the extent to which recent initiatives aimed at solving political disenfranchisement in the United States also tackle

<sup>35</sup>For example, see PEN America’s Media Literacy Program, which, among other initiatives, offers community leaders workshops about media literacy and misinformation defense (visit <https://pen.org/report/the-impact-of-community-based-digital-literacy-interventions-on-disinformation-resilience/>). See also the News Literacy Project, which instead targets mostly educational organizations (visit <https://newslit.org/>).

the issue of unequal access to news knowledge can provide a comprehensive understanding of the situation.<sup>36</sup> However, further research is necessary to thoroughly understand the root causes of these inequalities and to assess the effectiveness of policies designed to promote equal access to news sources, such as changes to the funding of news media (Cagé 2016).

## REFERENCES

- Alesina, Alberto, Armando Miano, and Stefanie Stantcheva.** 2020. "The Polarization of Reality." *American Economic Review Papers and Proceedings* 110 (2): 324–28.
- Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso.** 2018. "Intergenerational Mobility and Support for Redistribution." *American Economic Review* 108 (2): 521–54.
- Allcott, Hunt, and Matthew Gentzkow.** 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31 (2): 211–36.
- Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow.** 2020. "The Welfare Effects of Social Media." *American Economic Review* 110 (3): 629–76.
- Allen, Jennifer, Baird Howland, Markus Mobius, David Rothschild, and Duncan J. Watts.** 2020. "Evaluating the Fake News Problem at the Scale of the Information Ecosystem." *Science Advances* 6 (14). 10.1126/sciadv.aay3539.
- Allen, Jennifer N. L., Antonio A. Arechar, Gordon Pennycook, and David G. Rand.** 2021. "Scaling up Fact-checking Using the Wisdom of Crowds." *Science Advances* 7 (36): 1–10.
- Anderson, Carolyn J., Jay Verkuilen, and Buddy L. Peyton.** 2010. "Modeling Polytomous Item Responses Using Simultaneously Estimated Multinomial Logistic Regression Models." *Journal of Educational and Behavioral Statistics* 35 (4): 422–52.
- Angelucci, Charles, and Andrea Prat.** 2024. "Replication data for: Is Journalistic Truth Dead? Measuring How Informed Voters Are about Political News." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E195601V1>.
- Arana, Gabriel.** 2018. "Decades of Failure." *Columbia Journalism Review* 31 (2): 211–36.
- Arias, Eric, Horacio Larreguy, John Marshall, and Pablo Querubin.** 2018. "Does the Content and Mode of Delivery of Information Matter for Political Accountability? Evidence from a Field Experiment in Mexico." Unpublished.
- Arias, Eric, Horacio Larreguy, John Marshall, and Pablo Querubin.** 2022. "Priors Rule: When Do Malfeasance Revelations Help or Hurt Incumbent Parties." *Journal of the European Economic Association* 20 (4): 1422–77.
- Baek, Young Min, and Magdalena E. Wojcieszak.** 2009. "Don't Expect Too Much! Learning From Late-Night Comedy and Knowledge Item Difficulty." *Communication Research* 36 (6): 783–809.
- Baggini, Julian.** 2017. *A Short History of Truth: Consolations for a Post-Truth World*. London: Quercus Publishing.
- Bago, Bence, David G. Rand, and Gordon Pennycook.** 2020. "Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines." *Journal of Experimental Psychology* 149 (8): 1608–13.
- Ball, James.** 2018. *Post-Truth: How Bullshit Conquered the World*. Hull, UK: Biteback Publishing.
- Banerjee, Abhijit, Selvan Kumar, Rohini Pande, and Felix Su.** 2012. "Do Informed Voters Make Better Choices? Experimental Evidence from Urban India." Unpublished.
- Barrera, Oscar, Sergei M. Guriev, Emeric Henry, and Ekaterina Zhuravskaya.** 2020. "Facts, Alternative Facts, and Fact Checking in Times of Post-Truth Politics." *Journal of Public Economics* 182: 104123.
- Benabou, Roland, and Jean Tirole.** 2002. "Self-Confidence and Personal Motivation." *Quarterly Journal of Economics* 117 (3): 871–915.
- Benabou, Roland, and Jean Tirole.** 2006. "Belief in a Just World and Redistributive Politics." *Quarterly Journal of Economics* 121 (2): 699–746.
- Bock, Darrell R.** 1972. "Estimating Item Parameters and Latent Ability When Responses Are Scored in Two or More Nominal Categories." *Psychometrika* 37: 29–51.

<sup>36</sup>For example, organizations like the National Redistricting Foundation combat gerrymandering ([redistricting-foundation.org](http://redistricting-foundation.org)), the Fair Elections Center promotes voting rights ([fairelectionscenter.org](http://fairelectionscenter.org)), and the Brennan Center for Justice focuses on voting system reform ([brennancenter.org](http://brennancenter.org)).

- Bollinger, Lee C., and Geoffrey R. Stone.** 2018. *The Free Speech Century*. Oxford: Oxford University Press.
- Bullock, John G., Alan S. Gerber, Seth J. Hill, and Gregory A. Huber.** 2015. "Partisan Bias in Factual Beliefs about Politics." *Quarterly Journal of Political Science* 10 (4): 519–78.
- Cagé, Julia.** 2016. *Saving the Media*. Cambridge, MA: Belknap Press.
- Chen, Yuyu, and David Y. Yang.** 2019. "The Impact of Media Censorship: 1984 or Brave New World?" *American Economic Review* 109 (6): 2294–2332.
- d'Ancona, Michael.** 2017. *Post-Truth: The New War on Truth and How to Fight Back*. London: Ebury Press.
- Della Vigna, Stefano, and Ethan Kaplan.** 2007. "The Fox News Effect: Media Bias and Voting." *Quarterly Journal of Economics* 122 (3): 1187–1234.
- Delli Carpini, Michael X., and Scott Keeter.** 1996. *What Americans Know about Politics and Why It Matters*. New Haven, CT: Yale University Press.
- Eisensee, Thomas, and David Strömberg.** 2007. "News Droughts, News Floods, and U.S. Disaster Relief." *Quarterly Journal of Economics* 122 (2): 693–728.
- Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya.** 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review* 101 (7): 3253–85.
- Eurobarometer.** 2017. *Standard Eurobarometer 87*. Brussels: European Commission.
- Ferraz, Claudio, and Frederico Finan.** 2008. "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes." *Quarterly Journal of Economics* 123 (2): 703–45.
- Flaxman, Seth, Sharad Goel, and Justin M. Rao.** 2016. "Filter Bubbles, Echo Chambers, and Online News Consumption." *Public Opinion Quarterly* 80 (S1): 298–320.
- Gaxie, Daniel.** 1978. *Le cens caché: Inégalités culturelles et ségrégation politique*. Editions du Seuil, Paris.
- Gentzkow, Matthew, and Jesse M. Shapiro.** 2011. "Ideological Segregation Online and Offline." *Quarterly Journal of Economics* 126 (4): 1799–1839.
- Gentzkow, Matthew, Jesse M. Shapiro, and Daniel F. Stone.** 2015. "Media Bias in the Marketplace: Theory." In *Handbook of Media Economics*, edited by Simon Anderson, Joel Waldfogel, and David Stromberg, 41–90. Amsterdam: North Holland.
- Gerber, Alan S., Dean Karlan, and Daniel Bergan.** 2009. "Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions." *American Economic Journal: Applied Economics* 1 (2): 35–52.
- Gleick, James.** 2011. *The Information: A History, A Theory, A Flood*. New York, NY: Vintage Books.
- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer.** 2019. "Fake news on Twitter during the 2016 U.S. presidential election." *Science* 363 (6425): 374–78.
- Guess, Andrew M., Brendan Nyhan, and Jason Reifler.** 2020. "Exposure to Untrustworthy Websites in the 2016 US Election." *Nature Human Behaviour* 4: 472–80.
- Guess, Andrew M., Dominique Lockett, Benjamin Lyons, Jacob M. Montgomery, Brendan Nyhan, and Jason Reifler.** 2020a. "'Fake news' may have Limited Effects Beyond Increasing Beliefs in False Claims." *Harvard Kennedy School (HKS) Misinformation Review* 31 (2): 211–36.
- Guess, Andrew M., Michael Lerner, Benjamin Lyons, Jacob M. Montgomery, Brendan Nyhan, Jason Reifler, and Neelanjan Sircar.** 2020b. "A digital media literacy intervention increases discernment between mainstream and false news in the United States and India." *Proceedings of the National Academy of Sciences* 117 (27): 15536–45.
- Guess, Andrew M., Jonathan Nagler, and Joshua Tucker.** 2019. "Less than you think: Prevalence and Predictors of Fake News Dissemination on Facebook." *Science Advances* 5 (1): 10.1126/sciadv.aau4586.
- Guriev, Sergei, and Daniel Treisman.** 2019. "Informational Autocrats." *Journal of Economic Perspectives* 33 (4): 100–127.
- Howard, Philip N.** 2020. *Lie Machines: How to Save Democracy from Troll Armies, Deceitful Robots, Junk News Operations, and Political Operatives*. New Haven, CT: Yale University Press.
- Kakutani, Michiko.** 2018. *The Death of Truth: Notes on Falsehood in the Age of Trump*. New York, NY: Crown.
- Kamarck, Elain, and Alexander R. Podkul.** 2018. *The 2018 Primaries Project: The Demographics of Primary Voters*. Washington, DC: Brookings Institute.
- Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi.** 2015. "How Do Voters Respond to Information? Evidence from a Randomized Campaign." *American Economic Review* 105 (1): 322–53.
- Knight, Brian, and Ana Tribin.** 2022. "Opposition Media, State Censorship, and Political Accountability: Evidence from Chavez's Venezuela." *World Bank Economic Review* 36 (2): 455–87.

- Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva. 2015. "How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments." *American Economic Review* 105 (4): 1478–1508.
- Labonne, Julien, Cesi Cruz, and Philip Keefer. 2021. "Buying Informed Voters: New Effects of Information on Voters and Candidates." *Economic Journal* 131 (635): 1105–34.
- Lazer, David M. J., Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, et al. 2018. "The Science of Fake News." *Science* 359 (6380): 1094–96.
- Le Pennec, Caroline, and Vincent Pons. 2023. "How Do Campaigns Shape Vote Choice? Multicountry Evidence from 62 Elections and 56 TV Debates." *Quarterly Journal of Economics* 138 (2): 703–67.
- Levitin, Daniel J. 2016. *A Field Guide to Lies: Critical Thinking in the Information Age*. New York, NY: Dutton.
- Lizotte, Mary-Kate, and Andrew H. Sidman. 2009. "Explaining the Gender Gap in Political Knowledge." *Politics & Gender* 5 (2): 127–51.
- Martin, Gregory J., and Ali Yurukoglu. 2017. "Bias in Cable News: Persuasion and Polarization." *American Economic Review* 107 (9): 2565–99.
- Meyer, Robinson. 2018. "The Grim Conclusions of the Largest-Ever Study of Fake News." *The Atlantic*, March 8. <https://www.theatlantic.com/technology/archive/2018/03/largest-study-ever-fake-news-mit-twitter/555104/>.
- Mosquera, Roberto, Mofioluwademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie. 2020. "The Economic Effects of Facebook." *Experimental Economics* 23 (2): 575–602.
- Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95 (4): 1031–53.
- Nie, Norman H., Jane Junn, and Kenneth Stehlik-Barry. 1996. *Education and Democratic Citizenship in America*. Chicago, IL: University of Chicago Press.
- Nyhan, Brendan, Ethan Porter, Jason Reifler, and Thomas J. Wood. 2020. "Taking Fact-Checks Literally but Not Seriously? The Effects of Journalistic Fact-Checking on Factual Beliefs and Candidate Favorability." *Political Behavior* 42 (2): 939–60.
- Nyhan, Brendan. 2020. "Facts and Myths about Misperceptions." *Journal of Economic Perspectives* 34 (3): 220–36.
- Palfrey, Thomas R., and Keith T. Poole. 1987. "The Relationship between Information, Ideology, and Voting Behavior." *American Journal of Political Science* 31 (3): 511–30.
- Pennycook, Gordon, and David G. Rand. 2019. "Lazy, Not Biased: Susceptibility to Partisan Fake News is Better Explained by Lack of Reasoning than by Motivated Reasoning." *Cognition* 188: 39–50.
- Pennycook, Gordon, Jonathon McPhetres, Yunhao Zhang, Jackson G. Lu, and David G. Rand. 2020. "Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for a Scalable Accuracy-Nudge Intervention." *Psychological Science* 31 (7): 770–80.
- Pennycook, Gordon, Ziv Epstein, Mohsen Mosleh, Antonio A. Arechar, Dean Eckles, and David G. Rand. 2021a. "Shifting Attention to Accuracy Can Reduce Misinformation Online." *Nature* 592 (2): 590–95.
- Pennycook, Gordon, Ziv Epstein, Mohsen Mosleh, Antonio Arechar, Dean Eckles, and David Rand. 2021b. "Shifting Attention to Accuracy can Reduce Misinformation Online." *Nature* 592: 590–95.
- Pew. 2007. *Public Knowledge of Current Affairs Little Changed by News and Information Revolution*. Washington, DC: Pew Research Center.
- Pew. 2017. *From Brexit to Zika: What Do Americans Know?* Washington, DC: Pew Research Center.
- Pew. 2018. *Wide Gender Gap, Growing Educational Divide in Voters' Party Identification*. Washington, DC: Pew Research Center.
- Pew. 2019. *Many Americans Say Made-Up News Is a Critical Problem That Needs To Be Fixed*. Washington, DC: Pew Research Center.
- Price, Vincent, and John Zaller. 1993. "Who Gets the News? Alternative Measures of News Reception and Their Implications for Research." *Public Opinion Quarterly* 57 (2): 133–64.
- Prior, Markus. 2005. "News vs. Entertainment: How Increasing Media Choice Widens Gaps in Political Knowledge and Turnout." *American Journal of Political Science* 49 (3): 577–92.
- Prior, Markus. 2007. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. Cambridge Studies in Public Opinion and Political Psychology. Cambridge, UK: Cambridge University Press.
- Prior, Markus, and Arthur Lupia. 2008. "Money, Time, and Political Knowledge: Distinguishing Quick Recall and Political Learning Skills." *American Journal of Political Science* 52 (1): 169–83.
- Prior, Markus, Gaurav Sood, and Kabir Khanna. 2015. "You Cannot be Serious: The Impact of Accuracy Incentives on Partisan Bias in Reports of Economic Perceptions." *Quarterly Journal of Political Science* 10 (4): 489–518.

- Rasch, Georg.** 1960. *Probabilistic Models for Some Intelligence and Attainment Tests*. Copenhagen: Nielsen and Lydiche.
- Schaffner, Brian, Stephen Ansolabehere, and Sam Luks.** 2021. "Cooperative Election Study Common Content, 2020." Unpublished.
- Schiffirin, Anya.** 2020. "Mis-And Disinformation Online: A Taxonomy of Solutions." PhD diss. Universidad de Navarra.
- Snyder, James M., and David Strömberg.** 2010. "Press Coverage and Political Accountability." *Journal of Political Economy* 118 (2): 355–408.
- Stengel, Richard.** 2019. *Information Wars. How We Lost the Global Battle Against Disinformation and What We Can Do about It*. New York, NY: Grove Press.
- Strömberg, David.** 2004. "Mass Media Competition, Political Competition, and Public Policy." *Review of Economic Studies* 71 (1): 265–84.
- Strömberg, David.** 2015. "Media and Politics." *Annual Review of Economics* 7 (1): 173–205.
- Sunstein, Cass R.** 2018. *Republic: Divided Democracy in the Age of Social Media*. Princeton, NJ: Princeton University Press.
- Sunstein, Cass R.** 2019. *Too Much Information: Understanding What You Don't Want to Know*. Cambridge, MA: MIT Press.
- Train, Kenneth E.** 2009. *Discrete Choice Methods with Simulation*. Cambridge, MA: Cambridge University Press.
- Tufekci, Zeynep.** 2017. *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. London: Yale University Press.
- US Census Bureau.** 2020. "American Community Survey 2020 5-Year." <https://data.census.gov/mdat/> (accessed August 1, 2022).
- Van der Linden, Wim J., Ronald K. Hambleton, eds.** 1997. *Handbook of Modern Item Response Theory*. New York, NY: Springer.
- Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady.** 1995. *Voice and Equality*. Cambridge, MA: Harvard University Press.
- Vosoughi, Soroush, Deb Roy, and Sinan Aral.** 2018. "The Spread of True and False News Online." *Science* 359 (6380): 1146–51.
- WMC.** 2019. *The Status of Women in the U.S. Media 2019*. Washington, DC: Women's Media Center.
- Wu, Tim.** 2018. "Is the First Amendment Obsolete." *Michigan Law Review* 117 (3): 547–83.