



Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Short communication

Do people demand fact-checked news? Evidence from U.S. Democrats[☆]Felix Chopra^a, Ingar Haaland^{b,c}, Christopher Roth^{d,e,f,g,c,h,i,*}^a University of Bonn, Germany^b University of Bergen, Norway^c CESifo, Germany^d University of Cologne, Germany^e ECONtribute, Germany^f C-SEB, Germany^g briq, Germany^h CEPR, United Kingdomⁱ CAGE, United Kingdom

ARTICLE INFO

Article history:

Received 11 July 2021

Revised 13 October 2021

Accepted 27 October 2021

Available online 15 December 2021

JEL:

D83

D91

L82

Keywords:

Fact-checking

News demand

Information

Media bias

Belief polarization

ABSTRACT

In a large-scale online experiment with U.S. Democrats, we examine how the demand for a newsletter about an economic relief plan changes when the newsletter content is fact-checked. We first document an overall muted demand for fact-checking when the newsletter features stories from an ideologically aligned source, even though fact-checking increases the perceived accuracy of the newsletter. The average impact of fact-checking masks substantial heterogeneity by ideology: fact-checking reduces demand among Democrats with strong ideological views and increases demand among ideologically moderate Democrats. Furthermore, fact-checking increases demand among all Democrats when the newsletter features stories from an ideologically non-aligned source.

© 2021 Elsevier B.V. All rights reserved.

[☆] We would like to thank the editor (Ricardo Perez-Truglia) and three anonymous reviewers for excellent and highly constructive suggestions. We also thank Peter Andre, Joshua Dean, Armin Falk, Lukas Hensel, Yucheng Liang, Matt Lowe, Massimo Morelli, Salvatore Nunnari, Kirill Pogorelskiy, Andrea Prat, Carlo Schwarz, Daniel Sgroi, Andreas Stegmann, Joshua Tasoff, Marc Witte, Florian Zimmermann and seminar participants at the Belief-Based Utility Workshop 2021, Bonn, Warwick, Vienna, and the PERICLES Debate on Political Information at Bocconi for very helpful comments. Financial support from BA/Leverhulme Small Research Grants is gratefully acknowledged. Hrishikesh Iyengar, Apoorv Kanoongo, Melisa Kurtis and Kangli Wang provided excellent research assistance. IRB approval was obtained from the German Association for Experimental Economic Research (GfEW) and the University of Warwick HSSREC (86/20-21). The data collections were pre-registered in the AsPredicted registry. Roth: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2126/1-390838866. The usual disclaimer applies.

* Corresponding author at: University of Cologne, Germany.

E-mail addresses: Felix.Chopra@uni-bonn.de (F. Chopra), Ingar.Haaland@uib.no (I. Haaland), roth@wiso.uni-koeln.de (C. Roth).

1. Introduction

Misinformation on mass media is becoming increasingly prevalent (Lazer et al., 2018). Recent examples of misinformation on mass media include false claims about election fraud in the 2020 U.S. Presidential Election that were widely reported in several mainstream news outlets (Pennycook and Rand, 2021). The rise in misinformation coincides with distrust in the media reaching higher levels than ever, with 56% of Americans saying that the mainstream media is purposely trying to mislead the public with inaccurate reporting.¹ Academics and practitioners alike have suggested fact-checking as one of the main tools to combat misinformation and restore trust in the news (Sell et al., 2021). The extent to which fact-checking can be an effective tool to combat misinformation and restore trust in the news crucially depends on the demand

¹ <https://www.axios.com/media-trust-crisis-2bf0ec1c-00c0-4901-9069-e26b21c283a9.html> (accessed July 9, 2021).

for fact-checking services. If consumers—as assumed in many models of media markets—primarily care about the accuracy of the news, news demand should increase when the news content is fact-checked. On the other hand, if consumers also have non-instrumental motives to read news, such as preferences for belief confirmation (Young, 2016; Mullainathan and Shleifer, 2005; Faia et al., 2021; Di Tella, 2015), it is theoretically ambiguous how fact-checking affects the demand for news.

In this paper, we provide evidence on how demand for a newsletter changes when its content is fact-checked. In a large-scale online experiment with more than 4,600 Americans who voted Democratic in the 2020 U.S. Presidential Election, respondents can sign up for a weekly politics newsletter featuring the three top stories about an economic relief plan (the *Biden Rescue Plan*). Whether our respondents sign up for the newsletter is our main outcome of interest. Our key treatment variation is whether respondents are told that all stories featured in the newsletter will be fact-checked. We further cross-randomize whether the newsletter features stories from an ideologically aligned source (*MSNBC*) or a non-aligned news source (*Fox News*). Although focusing exclusively on Democrats limits the generalizability of our results, we made this choice to make sure that the newsletter is equally ideologically aligned for all respondents.

Turning to results, we first establish that our sample of Democrats expects articles featured in the newsletter to contain factual errors and believes that fact-checking increases the accuracy of the newsletter. These results hold irrespective of whether the newsletter features stories from an ideologically aligned or non-aligned source. Our first main result is that demand for a newsletter featuring stories from an ideologically aligned source is largely unaffected by the added fact-checking service: the fact-checking treatment increases newsletter demand by only 1.4 percentage points. The effect is not statistically significant ($p = 0.382$) and corresponds to a modest 2.7% change in demand compared to the control group mean of 49.7%. It is also relatively precisely estimated given our large sample of more than 4,100 respondents, which gives us an ex-post minimum detectable effect size of 4.4 percentage points (at 80% power). We thus have power to detect relatively modest effect sizes.

Our second main result is that the muted average treatment effect masks substantial heterogeneity by ideology: fact-checking decreases newsletter demand by 6.2 percentage points among Democrats with a strong ideology ($p = 0.021$) and increases demand among moderate Democrats by 4.5 percentage points ($p = 0.018$). These effect sizes correspond to a 10.4% reduction in demand among Democrats with a strong ideology and a 9.9% increase in demand among moderate Democrats (compared to control group means of 59.7% and 45%, respectively), underscoring the economic significance of the effects. Our third main result is that fact-checking increases demand among all Democrats when the newsletter features stories from an ideologically non-aligned source. The treatment increases demand by 10 percentage points on average ($p = 0.016$), which corresponds to a 29.1% increase in demand compared to the control group mean of 34.3%. This underscores the economic significance of the effects.

Our results provide a proof of concept that while fact-checking has the potential to increase the demand for news by increasing its perceived accuracy, it could also have the unintended side effect of reducing the demand for ideologically aligned news among consumers with extreme ideological views, who plausibly have a strong preference for belief confirmation. While these findings could potentially inform the optimal regulation of media markets, one should be careful when trying to generalize from a very specific setting with Democrats only. Our results could plausibly have looked differently if we had run the experiment on a different topic where accuracy concerns are likely to be more important, such as

news about COVID-19 vaccine efficacy, or with a sample of Republicans. To draw credible and robust conclusions for policy, future research will need to test the robustness of our findings on the demand for fact-checking across many different settings and samples.

Our paper contributes to several strands of the literature. First, the paper relates to the literature on fact-checking (Barrera et al., 2020), debiasing interventions (Grigorieff et al., 2020; Pennycook et al., 2020; Pennycook and Rand, 2019; Banerjee et al., 2018; Alesina et al., 2018; Galasso et al., 2021; Cruces et al., 2013), and misinformation on mass media (Bursztyn et al., 2021; Pennycook and Rand, 2021). Previous work in this literature has assessed how fact-checking or debiasing interventions affect beliefs and policy views (Nyhan and Reifler, 2010; Nyhan et al., 2019; Haaland et al., forthcoming; Barrera et al., 2020; Haaland and Roth, forthcoming; Fehr et al., forthcoming), trust in fact-checking services (Brandtzaeg and Følstad, 2017; Brandtzaeg et al., 2018), and willingness to share false news on social media (Henry et al., forthcoming).² While these studies have advanced our understanding of how fact-checking affects beliefs and policy views, it is important from a policy perspective to also understand how fact-checking affects people's news demand. We take the first step in this direction by providing evidence on how Democrats' demand for a politics newsletter changes when the newsletter content is fact-checked.

The paper also relates to the literature studying the demand for news (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006; Qin et al., 2018; Gentzkow et al., 2018; DellaVigna and La Ferrara, 2015; Prat and Strömberg, 2013). This literature has debated whether people tend to read ideologically aligned news because they have higher trust in ideologically aligned sources or because they want to confirm their existing beliefs (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006; Druckman and McGrath, 2019). We contribute to this literature by providing a proof of concept that non-instrumental motives, such as preferences for belief confirmation, play a role in driving the demand for ideologically aligned news.

Finally, the paper also relates to the literature on information demand (Zimmermann, 2015; Falk and Zimmermann, 2017; Nielsen, 2020; Golman et al., 2017; Ganguly and Tasoff, 2016; Thaler, 2021; Tappin et al., 2020; Faia et al., 2021; Fuster et al., forthcoming; Chopra et al., 2021).³ We contribute to this literature by providing evidence on whether Democrats have a preference for more accurate news. Compared to much of the previous literature, our design leverages a more natural outcome, namely people's decision to sign up for a real newsletter covering current political and economic news.

2. Sample and experimental design

2.1. Sample

We collected the data for the experiment during January and February 2021 in collaboration with Lucid, a data provider commonly used in economic research (Haaland et al., forthcoming; Bursztyn et al., 2020). The data was collected in four waves, with about 2,000 respondents per wave and 8,399 respondents in total.

² Work in psychology also studies interventions aiming to reduce the spread of misinformation. For example, attaching warnings to news stories disputed by third-party fact-checkers (Pennycook et al., 2020) or using crowdsourcing to generate trust ratings can help consumers identify inaccurate claims (Pennycook and Rand, 2019). While the outcomes considered by this research concern beliefs and trust in news, our focus is on the effects of fact-checking services on the demand for news.

³ See Capozza et al. (2021) for a review of the applied literature on information demand.

Each wave was pre-specified in the AsPredicted registry (see Table B.1 for an overview and additional registry information).⁴ To make sure that the newsletter was equally ideologically aligned for all respondents, we only recruited respondents who had voted for Joe Biden during the 2020 U.S. Presidential Election. Respondents who had voted for another candidate or had not voted at all were immediately screened out of the survey.

One recurring concern about online studies is potentially lower levels of attention among respondents compared to laboratory experiments, which may threaten the internal validity of the study. To address this concern, we included a simple pre-treatment attention check at the beginning of the study (see p. 43 of the Online Appendix for a screenshot). 56% of our respondents passed the attention check, which is very low compared to many other experiments (e.g., 96.4% in [Bottan and Perez-Truglia \(2020\)](#) and 99% in [Nathan et al. \(2020\)](#)). As shown in Section C of the Online Appendix, we also observe much lower data quality among inattentive respondents. We, therefore, focus on attentive respondents in the main specifications, leaving us with a sample of 4,667 respondents.^{5,6}

2.2. Experimental design

All four waves feature two base treatments that are constant across the waves. In the two base treatments, we vary whether we will fact-check a newsletter featuring the three top stories about the *Biden Rescue Plan* featured on *MSNBC*. On top of this, each wave includes a second set of cross-randomized conditions to assess the robustness of our findings to different variations in the newsletter content and to examine potential mechanisms. Specifically, we vary the framing of the plan (wave 1), the perceived instrumental benefits of the plan (wave 2), whether the newsletter features stories from *MSNBC* or *Fox News* (wave 3), and whether the newsletter features news or opinion pieces (wave 4). Each of the cross-randomized conditions includes a version with fact-checking and one without fact-checking, giving us ten treatments in total across the four waves (with 50% of the respondents being assigned to one of the two base treatments).⁷ Section E of the Online Appendix provides screenshots of the full experiment, including all the cross-randomized conditions.

In the experiment, we first measure basic demographics as well as a range of other background characteristics and political views. In the base treatments, respondents are then informed that Congress is debating whether to pass the *Biden Rescue Plan* (the American Rescue Plan Act of 2021) and that the plan has received strong support from liberals but has been criticized by conservatives. We then ask whether they would like to sign up for our weekly newsletter that contains stories about the plan featured on *MSNBC* during the last week.⁸ To fix beliefs about the stories featured in the newsletter, we made it clear to respondents that the newsletter would feature “the three top stories about the Biden Rescue Plan featured on *MSNBC* during the last week.” By always focusing on the

“three top stories” about the plan, our aim was to make sure that treated respondents did not get the impression that fact-checking affected the selection of articles into the newsletter.

We chose to focus on the *Biden Rescue Plan* because it was heavily featured in the news at the time of the experiment and we believed that demand for stories about the plan would be high. Furthermore, since the *Biden Rescue Plan* included a planned \$1,400 stimulus check to all Americans, staying informed about the plan could be instrumentally valuable (e.g. to make optimal saving or investment decisions). We chose to focus on *MSNBC* because it is a well-known liberal outlet that broadly matches the ideological leanings of our respondents. Indeed, in a representative survey of Americans, over 90% who identify *MSNBC* as their primary source of political news are Democrats or lean towards the Democratic party, the highest fraction among any news outlet ([Grieco, 2020](#)).

Respondents are randomized into the fact-checking condition (treatment) or the non-fact-checking condition (control). Respondents in the fact-checking condition are informed that “we will fact-check all stories featured in the newsletter and flag those with inaccuracies.” Respondents in the non-fact-checking condition are offered the same newsletter but without the fact-checking service.⁹ For fact-checking to be valuable, respondents need to have at least some trust in our ability to fact-check the articles. We did not emphasize our affiliation on the decision screen, but the consent form included information about our academic affiliations as “researchers from the University of Bonn, Bergen University, and Warwick University.”

Our main outcome of interest is whether people would like to receive our newsletter featuring the three top stories about the *Biden Rescue Plan*. We chose to focus on newsletter subscriptions because newsletters are a popular way of staying informed about politics, with 21% of Americans receiving news from a newsletter over the course of a week ([Newman et al., 2020](#)). Moreover, by including only the three top articles in the newsletter, we reduce the expected search costs associated with staying up to date about the plan. At the same time, administering the newsletter ourselves allows us to retain sufficient control to vary newsletter characteristics across treatment arms.

We also measure a battery of post-treatment beliefs to assess how fact-checking affected beliefs about different newsletter characteristics, including perceptions of the newsletter’s accuracy, the perceived trustworthiness of the newsletter, as well the newsletter’s entertainment value, political bias, quality, and complexity. We measure these beliefs using five-point Likert scales. Finally, we elicit beliefs about how many articles featured in the newsletter would contain any factual errors, how many articles they expect to be flagged for inaccuracies, and how much they trust our ability to fact-check the news articles. These questions also allow us to check whether fact-checking affected beliefs about the distribution of articles included in the newsletter.

2.2.1. Discussion of the design

Our base treatments exogenously vary the product characteristics of the newsletter similar to conjoint experiments by offering a fact-checking service to a random subset of respondents. This has a few desirable features. First, by providing additional information about the accuracy of the three top *MSNBC* articles on the *Biden Rescue Plan*, our treatment should not affect beliefs about which articles are featured in the newsletter. We are thus holding beliefs about media bias by omission, filtering, or distortion constant between the treatment group and the control group. Since our

⁹ Figure B.1 of the Online Appendix provides screenshots of the treatment and control condition. Section D provides further details about our fact-checking efforts.

⁴ Each pre-registration was submitted to the AsPredicted registry a few hours before the launch of the respective data collection.

⁵ Many experimental studies conducted with similar online samples usually screen out inattentive respondents from the outset (e.g., [Haaland and Roth, 2020](#); [Enke and Graeber, 2019](#); [Haaland et al., forthcoming](#)).

⁶ We had some minor attrition of 1.1% between the main outcome and the subsequent post-treatment belief measures about newsletter characteristics.

⁷ Tables B.4-B.10 in the Online Appendix assess the integrity of randomization for our treatments.

⁸ If respondents indicated that they would like to receive our newsletter, we provided them with a link to a website at the end of the survey. The newsletter was published on this website. To accommodate different versions of the newsletter, we created individual websites for each treatment arm (see Figure D.1 for an example). This procedure allowed us to preserve the anonymity of our respondents by circumventing the need to collect email addresses.

treatment should not affect the expected distribution of articles, our design shuts down mechanisms related to rational delegation of costly information acquisition (Suen, 2004; Chan and Suen, 2008). Second, rational agents without non-instrumental motives should prefer fact-checking because they can freely dispose of the additional information. This allows us to rule out prominent mechanisms based on Bayesian updating about the quality of a source that make it difficult to cleanly identify motives with observational data (Gentzkow and Shapiro, 2006). Third, we deliberately offered the fact-checking service ourselves. We truthfully tell our respondents in the treatment group that we will fact-check the newsletter. Our instructions make it clear that we are independent non-partisan researchers.

3. Results

3.1. Fact-checking of politically aligned news

3.1.1. Descriptives

49.7% of control group respondents signed up for the newsletter featuring stories from *MSNBC*. The high baseline demand for the newsletter likely reflects that our respondents were interested in staying informed about the outcome of the *Biden Rescue Plan* and saw the newsletter as a convenient tool to receive the most important information. Newsletter demand correlates strongly with the perceived accuracy, entertainment value, quality, and trust in the newsletter (as shown in Fig. B.9).

For fact-checking to be valuable in our setting, respondents have to expect at least some factual inaccuracies in the *MSNBC* stories included in the newsletter. Importantly, it is people's subjective expectation of factual inaccuracies—and not the actual prevalence of factual inaccuracies—that determines whether fact-checking should increase the valuation of the newsletter. Fig. 1 uses data from control group respondents to provide descriptive evidence on beliefs about factual inaccuracies in news articles included in the newsletter as well as trust in our ability to fact-check the articles. Fig. 1a shows that 58.8% of the respondents expect at least one article featured in the newsletter with articles from *MSNBC* to contain a factual error. Furthermore, conditional on expecting at least one error, respondents expect 1.6 articles to contain factual errors on average, or slightly more than 50% of all articles.¹⁰

Another necessary condition for fact-checking to be valuable is that respondents trust our ability to identify potential errors in the articles. As shown in Fig. 1b, we find high levels of trust in our fact-checking ability: 94.9% of the respondents report having at least some trust in our ability to fact-check articles from *MSNBC*, suggesting that our fact-checking treatment has scope to change the perceived accuracy of the newsletter.

3.1.2. Empirical specification

In what follows, we assess how demand for the newsletter changes in response to fact-checking. For that purpose, we estimate the following regression specification using OLS:

$$y_i = \alpha_0 + \alpha_1 \text{Treatment}_i + \alpha_2 \mathbf{x}_i + \varepsilon_i \quad (1)$$

where y_i is an indicator taking value 1 if respondent i signs up for

¹⁰ We identified factual errors in the articles that were featured in our newsletter. In our main newsletter featuring articles from *MSNBC*, we identified factual errors in two out of 21 articles. The share of articles with an error was 9.5%, which is lower than people's estimate of 30.2%. In our newsletter with *Fox News* articles, 11% of featured articles included an error, which is far below people's expectation of 71.7%. In comparison, Maier (2005) finds an objective error rate of 48% among 4,800 news sources cited in 14 local newspapers.

the newsletter and value zero otherwise; Treatment_i is an indicator for whether respondent i is in the fact-checking treatment; \mathbf{x}_i is a vector of control variables¹¹; and ε_i is an individual-specific error term. We use robust error terms for inference.

3.1.3. Deviation from the pre-registration

In the main specification, we pool data from all four waves, including the cross-randomized conditions that varied the framing of the plan, the perceived instrumental motives, and whether the newsletter featured news or opinion pieces. These cross-randomized conditions did not differentially affect demand for the newsletter featuring stories from the ideologically aligned source compared to the base treatment (as shown in Table B.14). We deviate from the pre-registration by pooling all results across waves as this allows us to increase the statistical precision of our main estimates and simplify the exposition of our results. A second deviation from the pre-registration is that, motivated by our theoretical model presented in Section A of the Online Appendix, we examine heterogeneity based on the strength of people's ideology. A third deviation from the pre-registration is that, for reasons discussed in Section 2.1, we focus on attentive respondents in our main analysis. All pre-registered regressions are reported exactly as pre-specified in Table B.15.

3.1.4. Main effect

Table 1 presents the main results on how fact-checking affects demand for the newsletter featuring stories from a politically aligned outlet, pooling observations from all waves. Column 1 of Panel A shows the main result of the paper: demand for the newsletter only increases by a non-significant 1.4 percentage points in response to the fact-checking treatment ($p = 0.382$). This effect corresponds to a modest 2.7% change in demand compared to the control group mean of 49.7%. The main effect is relatively precisely estimated given our large sample of more than 4,100 respondents, giving us an ex-post minimum detectable effect size at 80% power of 4.4 percentage points. We thus have power to detect relatively modest effect sizes, suggesting that the average effect of fact-checking on Democrats' demand for news is of relatively low economic importance. Furthermore, as shown in column 2, the muted impact occurs despite a large and statistically significant treatment effect on the perceived accuracy of the newsletter: respondents in the fact-checking condition think that the newsletter has 14.3% of a standard deviation higher accuracy ($p < 0.001$). That treated respondents expect our fact-checking service to increase the overall accuracy of the newsletter is consistent with their high trust in our ability to fact-check the articles (as shown in Fig. 1b). Treated respondents also think that the newsletter has 8.7% of a standard deviation higher trustworthiness ($p = 0.005$). We also see some suggestive evidence that treated respondents associate the newsletter with 4.9% of a standard deviation higher quality ($p = 0.115$) and 5.1% percent of a standard deviation lower left-wing bias ($p = 0.099$), but these effects—while going in the expected direction—are not very large compared to the effect on perceived accuracy. Finally, as shown in columns 6 and 7, it does not seem to be the case that fact-checking affects the perceived complexity ($p = 0.259$) or entertainment value ($p = 0.439$) of the newsletter. Our first main result can be summarized as follows:

¹¹ We include the following control variables: gender, education, employment status, log income, Census region, and race and ethnicity. We include wave fixed effects when pooling observations across waves.

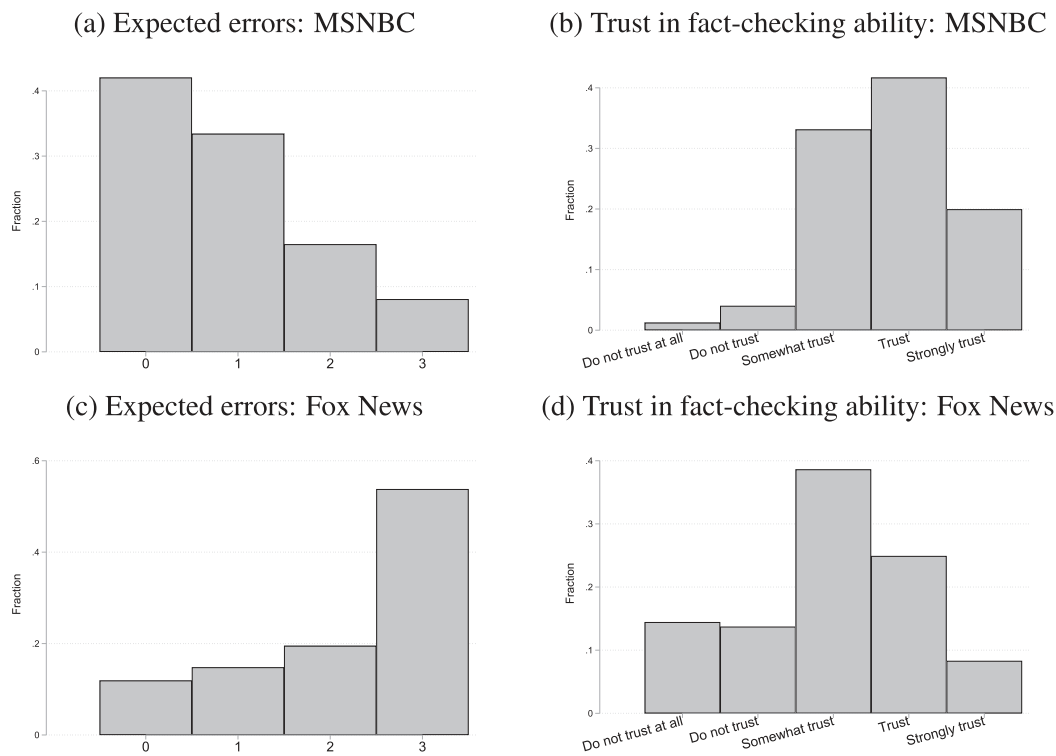


Fig. 1. Expected factual errors and trust in fact-checking. *Note:* This figure uses data from control group respondents who passed the attention check. Panel (a) shows the distribution of responses to the question “How many of the top three articles from MSNBC selected for the newsletter do you expect to contain factual errors?” Panel (b) shows the distribution of responses to the question “How much do you trust our ability to fact check articles from MSNBC?” Panel (c) and Panel (d) show the corresponding figures for Fox News.

Result 1. On average, people have a muted demand for fact-checking of news from politically aligned sources, despite a significant positive effect of fact-checking on the perceived accuracy of the newsletter.

3.1.5. Robustness

We cross-randomized several treatments to assess the robustness of our findings to differences in the content of the newsletter and to examine potential mechanisms. As shown in Table B.14, we find that our main result of a muted demand for fact-checking of ideologically aligned news is robust to varying (i) the framing of the *Biden Rescue Plan* (column 1), (ii) the perceived salience of the financial implications of the plan (column 2), and (iii) the type of articles covered in the newsletter (column 3). Furthermore, as shown in Table B.11, we see very similar point estimates and no significant treatment differences between the base treatments and the pooled cross-randomized treatments. These results suggest that our main finding of an overall muted demand for fact-checking of ideologically aligned news is robust to small variations in the description of the newsletter content.

3.1.6. Heterogeneity by ideology

As discussed in Section A of the Online Appendix, respondents with strong ideological views might assign a larger weight to non-instrumental motives—such as a preference for belief confirmation—than respondents with ideologically moderate views. In this case, we would expect the fact-checking treatment to have an opposite effect on newsletter demand for Democrats with strong and moderate ideological views. To categorize the strength of people’s ideological views, we use a pre-treatment question where people report their ideology on a five-point Likert scale from “very liberal” to “very conservative.” Throughout the paper, we

refer to “very liberal” respondents as those with strong ideological views and to the remaining respondents as moderate respondents.¹² Respondents with strong ideological views hold significantly more extreme policy attitudes than moderate respondents and are, for instance, 54% more likely to “strongly support” the *Biden Rescue Plan*.

Panels B and C of Table 1 show heterogeneity in treatment effects by ideological views (these effects are also displayed graphically in Panel A of Fig. 2). Panel B of Table 1 shows treatment effects for respondents with strong ideological views. These respondents significantly *reduce* their demand for the newsletter by 6.2 percentage points in response to the fact-checking treatment ($p = 0.021$, column 1). This corresponds to a 10.4% decline in demand compared to the control group mean of 59.7%, underscoring the economic significance of the effect. The decline in demand occurs even though the respondents perceive the newsletter as 11.8% of a standard deviation more accurate ($p = 0.028$, column 2). These respondents also perceive the fact-checked newsletter as somewhat less left-wing biased ($p = 0.079$, column 5), providing suggestive evidence for a mechanism where respondents with strong ideological views trade off accuracy against non-instrumental utility. Panel C of Table 1 shows treatment effects for respondents with ideologically moderate views. These respondents significantly *increase* their demand for the newsletter by 4.5 percentage points in response to the fact-checking treatment ($p = 0.018$, column 1), corresponding to a 9.9% increase in demand compared to a control group mean of 45 percent. Ideologically

¹² 31.8% of our sample rated themselves as “very liberal.” Furthermore, consistent with our restriction to focus on respondents who voted for Joe Biden in the 2020 Presidential Election, 93.7% of our respondents rated themselves as either “liberal” or “very liberal.” 5.6% rated themselves as “neither liberal nor conservative” and only 0.6% of respondents rated themselves as “conservative” or “very conservative.”

Table 1
Main results: MSNBC.

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Main effect							
Treatment	0.014 (0.016)	0.143*** (0.031)	0.087*** (0.031)	0.049 (0.031)	-0.051* (0.031)	0.035 (0.031)	0.023 (0.030)
N	4,109	4,069	4,069	4,069	4,069	4,069	4,069
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.497	0	0	0	0	0	0
Panel B: Strong ideology							
Treatment (a)	-0.062** (0.027)	0.118** (0.054)	0.043 (0.053)	0.016 (0.052)	-0.094* (0.054)	0.027 (0.055)	0.023 (0.052)
N	1,307	1,299	1,299	1,299	1,299	1,299	1,299
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.597	0	0	0	0	0	0
Panel C: Moderate ideology							
Treatment (b)	0.045** (0.019)	0.146*** (0.038)	0.097** (0.038)	0.051 (0.038)	-0.006 (0.038)	0.051 (0.038)	0.010 (0.037)
N	2,802	2,770	2,770	2,770	2,770	2,770	2,770
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.450	0	0	0	0	0	0
p-value: a = b	0.001	0.806	0.495	0.638	0.141	0.779	0.808

Note: This table shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use attentive respondents who were offered a newsletter featuring MSNBC articles. Panel A shows results for the full sample of Biden voters. Panel B shows results for respondents with strong ideology (who identify as “very liberal”). Panel C shows results for respondents with moderate ideology (who identify as not “very liberal”). “Treatment” is a binary variable taking value one if the articles in the newsletter are fact-checked. “News demand” is a binary variable taking the value one for respondents who said “Yes” to receive the newsletter and zero for those who said “No.” “Accuracy” of the newsletter is measured on a 5-point scale from “Very inaccurate” to “Very accurate.” “Trust” is the trustworthiness of the newsletter and measured on a 5-point scale from “Not trustworthy at all” to “Very trustworthy.” “Quality” of the newsletter is measured on a 5-point scale from “Very low quality” to “Very high quality.” “Left-wing bias” is measured on a 5-point scale from “Very right-wing biased” to “Very left-wing biased..” “Complexity” of the newsletter articles is measured on a 5-point scale from “Very simple” to “Very complex.” “Entertainment” of the newsletter is measured on a 5-point scale from “Not entertaining at all” to “Very entertaining.” The outcomes in columns 2–7 are z-scored using the control group mean and standard deviation.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

moderate respondents also perceive the fact-checked newsletter as 14.6% of a standard deviation more accurate ($p < 0.001$, column 2).

Comparing treatment effects in Panel B and Panel C of Table 1 reveals that we can reject equality of treatment effects on newsletter demand between respondents with strong and moderate ideological views at any conventional level of statistical significance. By contrast, there are no statistically significant differences in treatment effects between the two groups on beliefs about newsletter characteristics, such as accuracy and trust (columns 2–7). Our second main result follows.

Result 2. Respondents with strong and moderate ideological views respond differently to fact-checking: Despite similar first stage effects on beliefs about newsletter characteristics, respondents with strong ideological views reduce their newsletter demand by 10.4% in response to the fact-checking treatment while ideologically moderate respondents increase their newsletter demand by 9.9%.

3.2. Fact-checking of politically non-aligned news

We next study how fact-checking affects demand for a newsletter featuring stories from a politically non-aligned outlet. According to our theoretical framework (Section A of the Online Appendix), fact-checking only creates a trade-off between accuracy and non-instrumental motives when the articles are selected from a politically aligned news outlet. We would therefore expect fact-checking to increase demand for a newsletter featuring stories from a politically non-aligned outlet (Prediction 2 of Section A). To test this prediction, in wave 3, we cross-randomized whether

the newsletter featured news articles from Fox News instead of MSNBC while at the same time holding all other features of the design constant. We chose to focus on Fox News because it is a well-known outlet with a conservative leaning. Indeed, in a representative survey of Americans, over 90% who identify Fox News as their primary source of political news are Republicans or lean towards the Republican party, the highest fraction among any news outlet (Grieco, 2020).

3.2.1. Descriptives

As expected, we observe a lower demand for news from Fox News: 34.3% of control group respondents sign up for the newsletter featuring stories from Fox News, compared to 49.7% for MSNBC. Given that Biden voters tend to prefer left-wing news, it is reassuring that baseline demand for news from MSNBC is 45% higher than for news from Fox News. Furthermore, newsletter demand correlates strongly with the perceived accuracy of Fox News (as shown in Figure B.10). We next use data from control group respondents to provide descriptive data on beliefs about factual inaccuracies in news articles from Fox News. 88.6% of control group respondents expect at least one article to contain factual errors and 53.8% expect every article to contain some errors (Fig. 1c). Furthermore, 73% of the respondents express having at least some trust in our ability to fact-check articles from Fox News (Fig. 1d). These descriptives demonstrate a large scope for fact-checking to improve the perceived accuracy of the newsletter.

3.2.2. Main results

Panel A of Table 2 shows the treatment effects for the 558 respondents in the Fox News treatments. Column 1 shows that the fact-checking treatment increases newsletter demand by 10

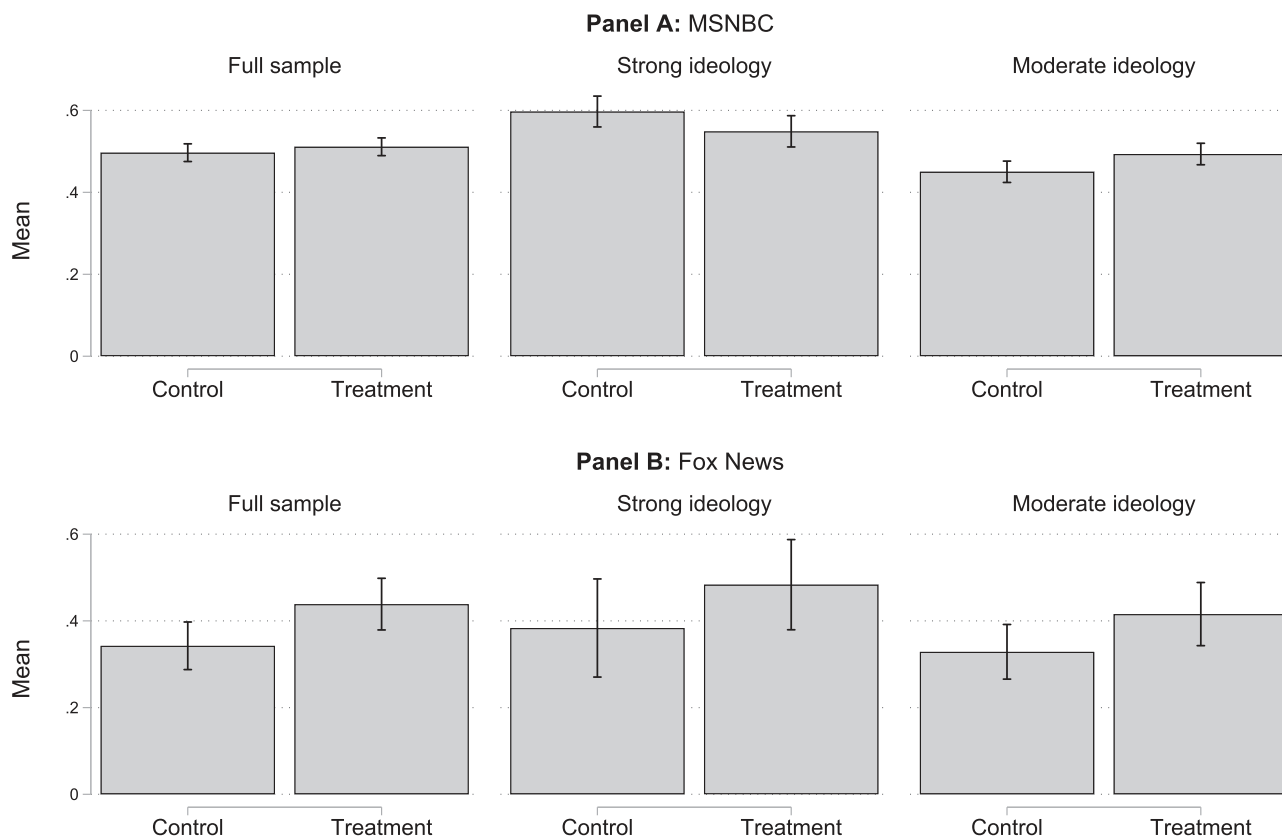


Fig. 2. Treatment effects on demand for the newsletter. *Note:* This figure shows newsletter demand (which is a binary variable taking the value one for respondents who said “Yes” to receive the newsletter and zero for those who said “No”) for MSNBC (Panel A) and Fox News (Panel B) among attentive respondents. Newsletter demand is shown separately by treatment group for the full sample of Biden voters, respondents with a strong ideology (who identify as “very liberal”), and for respondents with a moderate ideology (who identify as not “very liberal”). 95% confidence intervals are indicated.

percentage points ($p = 0.016$). This corresponds to a 29.1% increase in demand relative to the control mean of 34.3%, underscoring the high economic significance of the effect. Respondents in the fact-checking condition also think that the newsletter has 23.1% of a standard deviation higher accuracy ($p = 0.006$, column 2), 15.2% of a standard deviation higher trustworthiness ($p = 0.072$, column 3), and 17.7% of a standard deviation higher quality ($p = 0.038$, column 4).

3.2.3. Heterogeneity by ideology

Table 2 presents treatment effects for Democrats with strong ideology (Panel B) and Democrats with moderate ideology (Panel C). While focusing on these subsamples substantially reduces our power to detect statistically significant effects, especially for respondents with strong ideology, we find broadly similar patterns for both groups. As shown in column 1, treated respondents with strong and moderate ideology increase their demand for the newsletter by 6.4 percentage points ($p = 0.42$) and 9.5 percentage points ($p = 0.056$), respectively (these results are also shown graphically in Panel B of Fig. 2). The increase in demand among both groups is consistent with the theory that the trade-off between instrumental and non-instrumental motives disappears when the newsletter features stories from a politically non-aligned source. Furthermore, as shown in columns 2–7, treatment effects on beliefs about newsletter characteristics, including perceived accuracy, are also similar in magnitude and with no significant differences between the two groups. This leads to our third main result:

Result 3. All respondents, irrespective of their ideological leanings, increase their demand for the newsletter from a politically non-aligned source in response to the fact-checking treatment.

3.3. Alternative mechanisms

In this section, we discuss a series of mechanisms, which might be operating in this setting, but which are unlikely to explain the patterns in our data.

3.3.1. Confidence and ideology

Empirically, we find that both respondents with moderate and strong ideological views expect a more accurate newsletter if it is fact-checked (column 2 of Table 1). However, respondents with strong ideology, who hold strong prior beliefs about the world, might be very confident that they can detect any inaccuracies in reporting themselves. While overconfidence might decrease the perceived added-value of fact-checking services, it cannot strictly decrease the valuation of the newsletter. This would require an additional feature such as a large cost of processing information.

3.3.2. Updating about source quality

People might update about the quality of the underlying source of the newsletter when they learn that the source is fact-checked. For instance, people could think that fact-checking implies that the underlying source is of low quality (hence the need for a fact-check). To address this potential concern, we elicited expected errors from the underlying source of the newsletter. If anything, we actually see that our respondents in the fact-check condition expect fewer errors from the underlying source (Table B.16).

3.3.3. Cognitive constraints

Furthermore, since fact-checking in our context does not affect the selection of articles in the newsletter, we can—to the extent that fact-checking itself is not perceived as cognitively costly—

Table 2
Main results: Fox News.

	(1) News demand	(2) Accuracy	(3) Trust	(4) Quality	(5) Left-wing bias	(6) Complexity	(7) Entertainment
Panel A: Main effect							
Treatment	0.100** (0.041)	0.231*** (0.084)	0.152* (0.084)	0.177** (0.085)	-0.124 (0.081)	-0.076 (0.086)	0.107 (0.087)
N	558	548	548	548	548	548	548
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.343	0	0	0	0	0	0
Panel B: Strong ideology							
Treatment (a)	0.064 (0.079)	0.195 (0.158)	0.117 (0.163)	0.227 (0.172)	-0.208 (0.151)	-0.035 (0.159)	0.265 (0.176)
N	164	163	163	163	163	163	163
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.384	0	0	0	0	0	0
Panel C: Moderate ideology							
Treatment (b)	0.095* (0.049)	0.224** (0.101)	0.141 (0.099)	0.147 (0.097)	-0.062 (0.097)	-0.081 (0.102)	0.022 (0.096)
N	394	385	385	385	385	385	385
Z-scored	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.329	0	0	0	0	0	0
p-value: a = b	0.732	0.953	0.973	0.637	0.381	0.826	0.202

Note: This table shows OLS regression estimates where the dependent variables are demand for the newsletter and different post-treatment beliefs about the newsletter. All regressions use attentive respondents who were offered a newsletter featuring Fox News articles. Panel A shows results for the full sample of Biden voters. Panel B shows results for respondents with strong ideology (who identify as “very liberal”). Panel C shows results for respondents with moderate ideology (who identify as not “very liberal”). “Treatment” is a binary variable taking value one if the articles in the newsletter are fact-checked. “News demand” is a binary variable taking the value one for respondents who said “Yes” to receive the newsletter and zero for those who said “No.” “Accuracy” of the newsletter is measured on a 5-point scale from “Very inaccurate” to “Very accurate.” “Trust” is the trustworthiness of the newsletter and measured on a 5-point scale from “Not trustworthy at all” to “Very trustworthy.” “Quality” of the newsletter is measured on a 5-point scale from “Very low quality” to “Very high quality.” “Left-wing bias” is measured on a 5-point scale from “Very right-wing biased” to “Very left-wing biased.” “Complexity” of the newsletter articles is measured on a 5-point scale from “Very simple” to “Very complex.” “Entertainment” of the newsletter is measured on a 5-point scale from “Not entertaining at all” to “Very entertaining.” The outcomes in columns 2–7 are z-scored using the control group mean and standard deviation.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

change beliefs about accuracy while holding cognitive costs constant. Even if our respondents perceive fact-checking as cognitively costly (which we consider unlikely as column 6 of Table 1 shows that fact-checking does not affect the perceived complexity of the newsletter), the heterogeneity by the strength of people’s ideological views as well as the heterogeneity by the ideological leanings of the outlet suggest that cognitive constraints are not driving the observed patterns in our data.

3.3.4. Demand effects

While the high baseline demand for the newsletter featuring stories from MSNBC to some degree could reflect experimenter demand effects, this is not an issue for estimating treatment effects unless there is differential experimenter demand across treatment and control. While the between-design should not make it salient that we are interested in how fact-checking affects newsletter demand, we cannot rule out that some respondents nonetheless realized that we were studying fact-checking and adjusted their behavior accordingly. However, recent evidence suggests that demand effects are not a major concern in online experiments (de Quidt et al., 2018).

3.4. Expert survey

Lastly, we wanted to examine how experts expect the demand for the newsletter to change in response to fact-checking of the newsletter content. The results from this study can potentially inform a policy maker’s trade-off between following expert advice on fact-checking in a different setting and conducting new exper-

iments (DellaVigna and Pope, 2018). For this purpose, we conducted a survey in March 2021 among leading academic researchers in the areas of media and behavioral economics. We compiled a list of 93 experts who attended major conferences in economics.¹³ Our final sample consists of 65 experts, corresponding to a response rate of 70 percent.¹⁴ After providing the expert participants with information about the sample, design, and experimental instructions (including screenshots of the key treatment screens), we elicit their predictions about the effect of fact-checking on the demand for news for MSNBC and Fox News. For both outlets, we inform experts about baseline demand for the newsletter among respondents in the control group and then elicit their beliefs about newsletter demand among respondents in the treatment group.

Figure B.11 of the Online Appendix shows the results from the expert survey. As shown in Figure B.11, we observe a wide dispersion in expert beliefs about the impact of fact-checking on the demand for news with a mean absolute deviation of seven percentage points between expert opinions and actual treatment effects. The heterogeneity in expert beliefs suggests that there is substantial expert disagreement about the relative importance of different

¹³ These conferences include the briq Workshop on Beliefs, the NBER Summer Institute in Political Economy, and the Stanford Institute for Theoretical Economics (SITE) Summer Workshop (Experimental Economics and Psychology & Economics sessions).

¹⁴ 25% of these experts are Full Professor, 15% are Associate Professor, and 34% are Assistant Professors, 14% are postdoctoral researchers, and only 12% of respondents in our sample are PhD students. Among non-respondents, 65.5% are Full Professors, 14% are Associate Professors, 18% are Assistant Professors, and 4% are PhD students. This suggests lower response rates among full professors compared to assistant professors, PhD students, and postdoctoral researchers.

motives to read the news, such as the importance of accuracy motives versus belief utility motives. As shown in Figure B11b, expert beliefs on average closely resemble the actual treatment effects. As in DellaVigna and Pope (2018), our findings demonstrate a strong wisdom-of-crowds effect: while there is substantial disagreement within the expert sample, experts on average correctly predict the effects of fact-checking on newsletter demand.

4. Concluding remarks

This paper studies how fact-checking affects the demand for news. The main result of the paper is that Democrats have a muted demand for fact-checking of a newsletter featuring ideologically aligned news, even though fact-checking increases the perceived accuracy of the newsletter. This average effect masks substantial heterogeneity: Fact-checking decreases demand for politically aligned news among Democrats with strong ideological views and increases demand among ideologically moderate Democrats. Furthermore, fact-checking increases the demand for a newsletter with politically non-aligned news for all Democrats irrespective of the strength of their ideological leanings.

Our findings provide a proof of concept that non-instrumental motives play a role in driving the demand for ideologically aligned news. These findings have relevance for theories of media markets. In particular, our findings are inconsistent with theories in which all consumers primarily care about the accuracy of the news and point to the relevance of theories incorporating non-instrumental motives, such as a preference for belief confirmation. Furthermore, while one should be careful not to overgeneralize from a very specific setting, our findings suggest that fact-checking services can have very heterogeneous effects on the demand for news. While our study provides the first step to understand how fact-checking affects the demand for news, our results could be specific to our chosen sample and setting. Future research should study how fact-checking affects the demand for news across a range of different settings and samples to generate useful lessons for policy-makers.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.jpubeco.2021.104549>.

References

Alesina, Alberto, Carlana, Michela, La Ferrara, Eliana, Pinotti, Paolo, 2018. Revealing stereotypes: Evidence from immigrants in schools. Technical Report. National Bureau of Economic Research.

Banerjee, Abhijit, Ferrara, Eliana, Orozco, Victor, 2018. The entertaining way to behavioral change: Fighting HIV with MTV. Technical Report. Working paper.

Barrera, Oscar, Guriev, Sergei, Henry, Emeric, Zhuravskaya, Ekaterina, 2020. Facts, alternative facts, and fact checking in times of post-truth politics. *J. Public Econ.* 182, 104123.

Bottan, Nicolas L., Perez-Truglia, Ricardo, 2020. Choosing Your Pond: Location Choices and Relative Income. *Rev. Econ. Stat.* 10, 1–46.

Brandtzaeg, Petter Bae, Følstad, Asbjørn, 2017. Trust and distrust in online fact-checking services. *Commun. ACM* 60 (9), 65–71.

Brandtzaeg, Petter Bae, Følstad, Asbjørn, Chaparro Domínguez, María Ángeles, 2018. How Journalists and Social Media Users Perceive Online Fact-Checking and Verification Services. *Journalism Practice* 12 (9), 1109–1129.

Bursztyn, Leonardo, Rao, Aakaash, Roth, Christopher P., Yanagizawa-Drott, David H., 2021. Opinions as Facts. Working Paper.

Bursztyn, Leonardo, Haaland, Ingar, Rao, Aakaash, Roth, Christopher, 2020. Disguising Prejudice: Popular Rationales as Excuses for Intolerant Expression. University of Chicago, Becker Friedman Institute for Economics Working Paper, 2020-73.

Capozza, Francesco, Haaland, Ingar, Roth, Christopher, Wohlfart, Johannes, 2021. Studying Information Acquisition in the Field: A Practical Guide and Review. Discussion paper 124, ECONtribute.

Chan, Jimmy, Suen, Wing, 2008. A Spatial Theory of News Consumption and Electoral Competition. *Rev. Econ. Stud.* 75 (3), 699–728.

Chopra, Felix, Haaland, Ingar, Roth, Christopher, 2021. Do People Value More Informative News? Working Paper.

Cruces, Guillermo, Perez-Truglia, Ricardo, Tetaz, Martin, 2013. Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment. *J. Public Econ.* 98, 100–112.

de Quidt, Jonathan, Haushofer, Johannes, Roth, Christopher, 2018. Measuring and Bounding Experimenter Demand. *Am. Econ. Rev.* 108 (11), 3266–3302.

DellaVigna, Stefano, Pope, Devin, 2018. Predicting Experimental Results: Who Knows What? *J. Polit. Econ.* 126 (6), 2410–2456.

DellaVigna, Stefano, La Ferrara, Eliana, 2015. Economic and Social Impacts of the Media. In: Anderson, Simon P., Waldfogel, Joel, Strömberg, David (Eds.), *Handbook of Media Economics*, vol. 1B. *Handbook of Media Economics*, North-Holland, pp. 723–768.

Druckman, James N., McGrath, Mary C., 2019. The evidence for motivated reasoning in climate change preference formation. *Nat. Clim. Change* 9 (2), 111–119.

Enke, Benjamin, Graeber, Thomas, 2019. Cognitive Uncertainty. NBER Working Papers 26518. National Bureau of Economic Research.

Faia, Ester, Fuster, Andreas, Pezone, Vincenzo, Zafar, Basit, 2021. Biases in information selection and processing: Survey evidence from the pandemic. Technical Report. National Bureau of Economic Research.

Falk, Armin, Zimmermann, Florian, 2017. Beliefs and Utility: Experimental Evidence on Preferences for Information. Working Paper.

Fehr, Dietmar, Mollerstrom, Johanna, Perez-Truglia, Ricardo, 2021. Your Place in the World: The demand for National and Global Redistribution. *Am. Econ. J. Econ. Policy* (forthcoming).

Fuster, Andreas, Perez-Truglia, Ricardo, Wiederholt, Mirko, Zafar, Basit, 2020. Expectations with Endogenous Information Acquisition: An Experimental Investigation. *Rev. Econ. Stat.* (forthcoming).

Galasso, Vincenzo, Morelli, Massimo, Nannicini, Tommaso, Stanig, Piero, 2021. Fighting Populism on Its Own Turf: Experimental Evidence. Working Paper.

Ganguly, Ananda, Tasoff, Joshua, 2016. Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future. *Manage. Sci.* 63 (12), 4037–4060.

Gentzkow, Matthew A., Shapiro, Jesse M., 2006. Media Bias and Reputation. *J. Polit. Econ.* 114 (2), 280–316.

Gentzkow, Matthew A., Wong, Michael B., Zhang, Allen T., 2018. Ideological Bias and Trust in Information Sources. Working Paper.

Golman, Russell, Hagmann, David, Loewenstein, George, 2017. Information Avoidance. *J. Econ. Lit.* 55 (1), 96–135.

Grieco, Elizabeth, 2020. Americans' main sources for political news vary by party and age". <https://pewrsr.ch/3dHYbmb> (accessed 02-24-2021).

Grigorieff, Alexis, Roth, Christopher, Ufal, Diego, 2020. Does Information Change Attitudes Toward Immigrants? *Demography* 57 (3), 1–27.

Haaland, Ingar, Roth, Christopher, 2020. Labor market concerns and support for immigration. *J. Public Econ.* 191, 104256.

Haaland, Ingar, Roth, Christopher, 2021. Beliefs About Racial Discrimination and Support for Pro-Black Policies. *Rev. Econ. Stat.* (forthcoming).

Haaland, Ingar, Roth, Christopher, Wohlfart, Johannes, 2021. Designing Information Provision Experiments. *J. Econ. Lit.* (forthcoming).

Henry, Emeric, Zhuravskaya, Ekaterina, Guriev, Sergei, 2021. Checking and sharing alt-facts. *Am. Econ. J. Econ. Policy* (forthcoming).

Lazer, David M.J., Baum, Matthew A., Benkler, Yochai, Berinsky, Adam J., Greenhill, Kelly M., Menczer, Filippo, Metzger, Miriam J., Nyhan, Brendan, Pennycook, Gordon, Rothschild, David, Schudson, Michael, Sloman, Steven A., Sunstein, Cass R., Thorson, Emily A., Watts, Duncan J., Zittrain, Jonathan L., 2018. The science of fake news. *Science* 359 (6380), 1094–1096.

Maier, Scott R., 2005. Accuracy Matters: A Cross-Market Assessment of Newspaper Error and Credibility. *Journalism Mass Commun. Quart.* 82 (3), 533–551.

Mullainathan, Sendhil, Shleifer, Andrei, 2005. The Market for News. *Am. Econ. Rev.* 95 (4), 1031–1053.

Nathan, Brad C., Perez-Truglia, Ricardo, Zentner, Alejandro, 2020. My Taxes are Too Darn High: Why Do Households Protest Their Taxes?. Technical Report. National Bureau of Economic Research.

Newman, Nic, Fletcher, Richard, Schulz, Anne, Andi, Simge, Nielsen, Rasmus Kleis, 2020. Reuters Institute Digital News Report 2020. Technical Report, Reuters Institute for the Study of Journalism.

Nielsen, Kirby, 2020. Preferences for the resolution of uncertainty and the timing of information. *J. Econ. Theor.* 189, 105090.

Nyhan, Brendan, Reifler, Jason, 2010. When Corrections Fail: The Persistence of Political Misperceptions. *Polit. Behav.* 42, 303–300.

Nyhan, Brendan, Porter, Ethan, Reifler, Jason, Wood, Thomas J., 2019. Taking fact-checks literally but not seriously? The effects of journalistic fact-checking on factual beliefs and candidate favorability. *Polit. Behav.* 42, 939–960.

Pennycook, Gordon, Bear, Adam, Collins, Evan T., Rand, David G., 2020. The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings. *Manage. Sci.* 66 (11), 4944–4957.

Pennycook, Gordon, Rand, David G., 2019. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Nat. Acad. Sci.* 116 (7), 2521–2526.

Pennycook, Gordon, Rand, David G., 2021. Research note: Examining false beliefs about voter fraud in the wake of the 2020 Presidential Election. Harvard Kennedy School (HKS) Misinformation Review.

- Prat, Andrea, Strömberg, David, 2013. The Political Economy of Mass Media. In: Acemoglu, Daron, Arellano, Manuel, Dekel, Eddie (Eds.), *Advances in Economics and Econometrics: Applied Economics*. *Advances in Economics and Econometrics*, vol. 2. Cambridge University Press, Cambridge, pp. 135–187.
- Qin, Bei, Strömberg, David, Yanhui, Wu., 2018. *Media Bias in China*. *Am. Econ. Rev.* 108 (9), 2442–2476.
- Sell, Tara Kirk, Hosangadi, Divya, Smith, Elizabeth, Trotochaud, Marc, Vasudevan, Prarthana, Gronvall, Gigi Kwik, Rivera, Yonaira, Sutton, Jeannette, Ruiz, Alex, Cicero, Anita, 2021. National Priorities to Combat Misinformation and Disinformation for COVID-19 and Future Public Health Threats: A Call for a National Strategy. Johns Hopkins Center for Health Security, Baltimore, MD.
- Suen, Wing, 2004. The Self-Perpetuation of Biased Beliefs. *Econ. J.* 114 (495), 377–396.
- Tappin, Ben M., Pennycook, Gordon, Rand, David G., 2020. Thinking clearly about causal inferences of politically motivated reasoning: why paradigmatic study designs often undermine causal inference. *Curr. Opin. Behav. Sci.* 34, 81–87.
- Di Tella, Rafael, Perez-Truglia, Ricardo, Babino, Andres, Sigman, Mariano, 2015. Conveniently Upset: Avoiding Altruism by Distorting Beliefs About Others' Altruism. *Am. Econ. Rev.* 105 (11), 3416–3442.
- Thaler, Michael, 2021. The “Fake News” Effect: An Experiment on Motivated Reasoning and Trust in News. Working Paper.
- Young, Eric, 2016. A new understanding: What makes people trust and rely on news. American Press Institute.
- Zimmermann, Florian, 2015. Clumped or Piecewise? Evidence on Preferences for Information. *Manage. Sci.* 61 (4), 740–753.