

Tilburg University

Biases in information selection and processing: Survey evidence from the pandemic

Faia, Ester; Fuster, Andreas; Pezone, Vincenzo; Zafar, Basit

Published in:
Review of Economics and Statistics

DOI:
[10.1162/rest_a_01187](https://doi.org/10.1162/rest_a_01187)

Publication date:
2024

Document Version
Version created as part of publication process; publisher's layout; not normally made publicly available

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Faia, E., Fuster, A., Pezone, V., & Zafar, B. (2024). Biases in information selection and processing: Survey evidence from the pandemic. *Review of Economics and Statistics*, 106(3), 829-847.
https://doi.org/10.1162/rest_a_01187

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Biases in Information Selection and Processing: Survey Evidence from the Pandemic*

Ester Faia[†] Andreas Fuster[‡] Vincenzo Pezone[§] Basit Zafar[¶]

February 5, 2022

Abstract

We conduct two survey experiments to study which information people choose to consume and how it affects their beliefs. In the first experiment, respondents choose between optimistic and pessimistic article headlines related to the COVID-19 pandemic, and are then randomly shown one of the articles. Respondents with more pessimistic prior beliefs tend to prefer pessimistic headlines, providing evidence of confirmation bias. Additionally, respondents assigned to the less preferred article discount its information. The second experiment studies the role of partisan views, uncovering strong source dependence: news source revelation further distorts information acquisition, eliminating the role of priors in article choice.

Keywords: *Belief updating, confirmatory biases, endogenous information acquisition, media polarization, source dependence, COVID-19*

JEL Codes: D84, D91, E71, I12.

*We thank Stefano Cassella, Chris Roth, Paola Sapienza, Michael Thaler, Laura Veldkamp, Michael Weber, Johannes Wohlfart, two anonymous referees, and seminar participants at various conferences and seminars for useful comments and suggestions. We gratefully acknowledge financial support from the DFG priority grant FA-1022.3-1 and from the Leibniz Institute for Financial Research SAFE.

[†]Goethe University Frankfurt and CEPR. Email: faia@wiwi.uni-frankfurt.de

[‡]EPFL, Swiss Finance Institute and CEPR. Email: andreas.fuster@epfl.ch

[§]Tilburg University. Email: v.pezone@tilburguniversity.edu

[¶]University of Michigan and NBER. Email: basitak@gmail.com

1 Introduction

How people form subjective beliefs in circumstances with high uncertainty is a fundamental question, with far-reaching consequences for economic and social decisions. This is especially the case during a situation such as a pandemic, where people’s subjective beliefs about the health risks and economic consequences may affect their adherence to social distancing guidelines as well as their economic behaviors. While a fully rational agent would form her beliefs by taking into account all available signals, it is well-established that humans are selective in their signal processing, and exhibit various biases.

One broad set of biases is referred to as “motivated beliefs”: people believe what they want to believe. In many contexts, however, this simple insight alone does not make clear predictions. For instance, in the case of the COVID-19 pandemic, people could exhibit “wishful thinking,” and dismiss pessimistic news about health risks or economic consequences—especially if they themselves are more at risk, either health-wise or economically. Alternatively, they could be “ambiguity averse” and put more weight on negative signals, so that they could be better equipped to the realization of worst case scenarios. In addition, people could exhibit “confirmation bias” and choose to pay more attention to signals that are in line with their pre-existing views, regardless of whether they are optimistic or pessimistic. Finally, people may have strong views about the source of a given piece of news, which may affect their information selection and interpretation, even when the tone and content of the news are held fixed.

In this paper, we report the results from two survey experiments (fielded in May 2020, with a total of over 4,000 respondents) designed to study how motivated beliefs affect information selection and processing during the COVID-19 pandemic. The first survey was tailored to shed light on how people select the type of news they consume (in terms of optimistic or pessimistic tone) and examine belief updating when agents *randomly* receive information they did or did not select. The second survey aims at testing the exacerbation of confirmatory biases when news are channelled through sources that are generally aligned with, or opposed to, the respondent’s political orientation. Overall, our main results provide support for theories of confirmatory biases in which agents “overweight” desired signals, which are in turn correlated with priors.¹

¹ When we use the term “overweighting” in our empirical analysis, we simply mean that individuals, on average,

Our second survey shows that, once agents are made aware of the news source, political views become the strongest predictor of signal selection, and even dominate priors.

In both surveys, to measure the extent of motivated information selection and processing, we first elicit respondents' prior beliefs about Covid-related outcomes and their subjective concerns about possible effects on their health and financial situations. They are then informed that they can read a Covid-related article, and are asked to choose among different news headlines related to the pandemic. The headlines differ in terms of tone (pessimistic or optimistic), but are otherwise drafted to be equally informative.² A crucial aspect of our design is that we then randomize whether the respondent receives the preferred news article or not. This allows us to study the causal impact of information on belief updating and different types of (intended as well as actual) actions.

In the first survey, the news articles differ with respect to the domain (economy versus health) and the tone (optimistic versus pessimistic). Our results show that agents' choice of article tends to reflect their priors, independent of whether the latter were on the optimistic or pessimistic side. For instance, a one-standard-deviation increase in a respondent's overall pessimism about the effects of the pandemic (measured as the first principal component of their priors and subjective concerns) is associated with a 5 to 8 percentage point higher likelihood of preferring a health-related article, and a 6 to 9 percentage point higher likelihood of preferring the pessimistic article within each domain, in all cases controlling for a rich set of other respondent characteristics.

We do not find evidence consistent with wishful thinking, namely selection of optimistic prospects by those respondents who are (objectively or subjectively) more exposed to the risks from the pandemic. Wishful thinking has generally been documented in beliefs about oneself (e.g., Weinstein, 1980; Oster et al., 2013), while Thaler (2020a) does not find evidence that individuals wishfully think that the world is a better place for others. Our results thus suggest that Covid-related thinking is not primarily driven by a desire to feel better about oneself (e.g., one's personal health risks).

put more weight on the signal when it is aligned with their priors (that is, they receive the information they selected) versus when they receive information they had not selected. Our empirical setup is not rich enough to investigate whether, relative to a Bayesian benchmark, individuals are putting too little or too much weight on the signals.

2 In an earlier "pilot" survey, we checked that respondents on average found the headlines equally informative and interesting. See Section 2.1 and Appendix A.7.

Moving to information processing, we see that respondents rate a given article as more informative and are able to recall its headline more accurately (a proxy for attention) if they had chosen it. While the headlines were drafted to be equally informative and interesting, one could argue that some of these results may be consistent with rational information processing: if an agent prefers a headline because they are ex-ante uninformed about it,³ or because they find it more informative, such patterns may emerge under Bayesian updating. However, we also find that agents rate the preferred article as more reliable (by about 0.4 standard deviations), even without knowing the source. Moreover, we provide participants with a budget that they can use to advertise the source of the article they read, and show that they spend substantially more money (about 0.25 standard deviations) when receiving their preferred article. Both results are sizable in magnitude, and hard to rationalize under unbiased information processing.

Further support for confirmatory biases emerges from asymmetric updating: respondents incorporate the information from the article more strongly into their beliefs if they had chosen it. For instance, after seeing an article with a pessimistic tone, respondents revise upward their forecasts of the number of Covid-related deaths and of the unemployment rate more strongly (by 0.15-0.30 of the standard deviation in prior beliefs) if the article was the one they had selected.⁴

Other interesting results emerge. First, while a person's relative concern about health vs. economic consequences of the pandemic is important for the choice of the article domain, measures of pessimism and lockdown support are associated with preferring the pessimistic article not only in the health domain but also in the economy domain. While the latter may seem surprising, it suggests that respondents view economic and health effects of the pandemic as complementary.⁵ Second, to verify the consistency of stated beliefs with actual behavior, we also measure policy preferences (support for lockdowns) and two incentivized actions (donation to advertise the news source of the article, and to one of two organizations with opposing views on the policy response to the pandemic, namely the World Health Organization and FreedomWorks).

3 While this is possible in theory, our finding that respondents tend to choose the article that is aligned with their priors suggests that this is not very plausible.

4 In Appendix A.2, we present a theoretical framework featuring optimal belief distortion which can generate predictions in line with these results.

5 Absence of a factual trade-off is documented, for instance, by Fernández-Villaverde and Jones (2020).

Consistently, respondents who choose the pessimistic article and that are selected to read it tend to grant more support to policies and organizations fostering views in line with it.

An immediate consequence of these results is that the ability of individuals to select information in line with their priors is likely to amplify differences in beliefs. These results also have implications for polarization on the political spectrum. We document that priors about the pandemic are significantly correlated with political preferences, with liberals being significantly more pessimistic. Among participants who receive their preferred articles, the differences in the posteriors between conservatives and liberals are further amplified. These results inspire the design of our second survey. The key difference with the first survey is that we now reveal the source of the news articles to some participants.⁶

In the second survey, we provide direct evidence on polarization due to respondents' views about different media sources. After soliciting the same demographics and prior beliefs as in the first survey, we ask respondents to select between two articles differing in their tone. We randomize respondents into one of three groups, which differ in whether we reveal the source of the article or not, and when. Specifically, the first group never sees the sources of the articles (T1), the second group sees the sources after the article choice but before the reporting of posterior beliefs (T2), and the third group sees the sources right from the beginning (T3). Comparison of T3 with the other two groups allows us to study how the article choice is affected by the revelation of the news source (that is, the “selection” effect). Comparison of T2 with T1 allows us to study the impact of source revelation on information processing, keeping selection fixed. We selected a pessimistic article from Fox News and an optimistic article from New York Times. The choice was not coincidental—already in the first survey, we find that respondents view optimistic pandemic-related news as more likely to come from conservative news sources and pessimistic news as more likely to come from liberal news sources (see also Simonov et al., 2021). Thus, we chose the sources to be “counterintuitive” so as to maximize the impact of the source revelation treatment.

Our results are striking. We find that the revelation of the news sources prior to the selection of

⁶ Entrenched views tend to be reinforced when agents interact in groups sharing common values, political inclinations, or cultural traits. Political and media polarization is a manifestation of this tendency. The pandemic has likely exacerbated those societal divisions. For a different view, see Boxell et al. (2020), who find that affective polarization did not increase during the pandemic.

the news article completely offsets the relationship between priors and the respondents' choice of article: respondents appear to ignore their priors and instead choose the article source aligned with their political ideology. For instance, while liberal respondents are 8-9 percentage points more likely to choose the article with the pessimistic title when the source is not provided, they are 16-18 percentage points *less* likely to choose this article in the treatment where they are told upfront that the article comes from Fox News. Regarding information processing, we find that individuals who learn that the article they have read comes from a source aligned with their beliefs rate it more favorably (in terms of reliability, informativeness, attention, and money spent on advertising it). This effect is present, however, only when individuals are especially "surprised" to learn about the news source—that is, if they had *ex ante* thought that the likelihood of the pessimistic article coming from Fox News (or the optimistic article coming from the New York Times) was very low. This is precisely the group of respondents that we expected to be especially sensitive to our source revelation treatment. We dub this effect "source dependence."

While source dependence has large effects on people's selection and assessment of information, we find that impacts on posterior beliefs are muted. Specifically, it is not the case that revelation of the source leads to respondents discounting the information: for example, average revisions for self-reported liberals who read the pessimistic article, and are informed that the source is Fox News, are not different from those of their counterpart liberals who read the same article but are not informed of the source. Taken together, our results on source dependence suggest that policymakers can reach out to the "other side" by sending signals through sources traditionally aligned with that group. For example, lawmakers could have their message broadcast to a broader segment of the population by engaging with news sources that are on the other end of the political ideology spectrum (such as Fox News for Democrats, and MSNBC for Republicans). Our results show that, at least in the short term, such engagement will not change how the subpopulation already ideologically aligned with the lawmaker would interpret the message, and will also provide an opportunity of influencing unaligned groups.

Our paper is related to the literature on the motivated selection and processing of information. In the theoretical literature, several variants of this concept have emerged. They include: cognitive dissonance, namely the tendency to forge contradictory beliefs as consistent (Festinger,

1962); selection of beliefs instrumental in motivating desirable actions or achieving desirable goals (Benabou and Tirole, 2002; Benabou, 2015); self-deception—namely the tendency to deny opposing evidence (Trivers, 2011); wishful thinking and anticipatory feelings (Jevons, 1905; Loewenstein, 1987; Brunnermeier and Parker, 2005; Caplin and Leahy, 2019), and confirmatory biases (Rabin and Schrag, 1999). Our results are in line with confirmatory biases: we do not detect any general tendency for optimism or denial, but we do find strong evidence that respondents select evidence in favor of their prior views of the pandemic and, in particular, of their desired policy responses.

In terms of evidence, there have been numerous lab experiments in the psychology literature.⁷ Field evidence, especially at a large scale, is instead more scant. Most of it focuses on overconfidence and self-esteem; see for instance Eil and Rao (2011), Mobius et al. (2014), and Wiswall and Zafar (2015). Confirmatory biases have received less attention. Another important novelty of our analysis lies in the design, which randomizes information assignment after having solicited respondents' choice of the news they desire to consume.

A distinguishing characteristic of motivated beliefs is asymmetric updating: agents overweight signals that confirm their priors, previous choices or tastes, and underweight or neglect signals that contradict them (Benabou and Tirole, 2016; Benabou, 2015). Incentivized experiments on motivated beliefs include Camerer and Lovallo (1999), Saucet and Villeval (2019), and Zimmermann (2020). In many settings, the signals individuals receive may have different informational content, and so biases may arise due to non-motivated reasons. In our setup, the headlines of the articles are chosen to be equally informative and only differ in their tone (optimistic or pessimistic); see also Thaler (2020b), who finds evidence of (politically-) motivated reasoning in a novel experimental setup where such confounds are arguably absent. While some of our results may be rationalized through other channels, collectively, the patterns clearly show evidence of motivated beliefs and confirmatory bias.

Motivated beliefs can also materialize at the social level, in terms of group-thinking and polarization, which in turn affect economic beliefs (Coibion et al., 2020b). Related to a large literature studying media and political polarization, our second survey experiment studies how

⁷ Weinstein (1980), Lord et al. (1979), Darley and Gross (1983) and, more recently, Ditto et al. (2009) and Talluri et al. (2018) are a few examples. Epley and Gilovich (2016) provide a compact introduction.

media sources may interact with priors in enhancing polarization.⁸ Perhaps closest to our second experiment is Jo (2020), who conducts a field experiment in South Korea (pre-pandemic) where some participants are shown news at random (but with sources revealed) while others get to pick their sources. Similar to us, he finds that the latter group chooses sources aligned with their political preferences; they also update beliefs more and end up with less extreme policy views than the other group. This supports a model in which familiarity with a source allows readers to “de-bias” the news. While some of our results align, we do not find differential updating when the source is revealed after article choice (which Jo’s design does not allow to study), in apparent contrast with the model above. Another recent complementary contribution is Chopra et al. (2019), who also use a survey experiment to study the demand for news. They manipulate respondents’ beliefs about the informativeness of a given news source (the New York Times), finding that people demand less news from this source when it is perceived as more informative/even-handed, independently of whether they are Republican or Democrat.⁹ We instead manipulate whether and when the source of a given article is revealed, and study how this affects the demand for and interpretation of the news.

We believe the pandemic context is useful for a survey on motivated beliefs for several reasons. First, the pandemic was characterized by widespread uncertainty, as even the views of experts and health advisors were evolving as the pandemic progressed. On the one hand, this generates the need for information gathering and amplifies the importance of the media. On the other hand, a wide heterogeneity of views creates the scope for motivated selection of signals. Second, it is a large shock, hence polarized opinions are more likely to emerge from the start. In the US, Covid-related opinions have polarized deeply along the spectrum of political affiliations (see Allcott et al., 2020, and Makridis and Rothwell, 2020), and this political divide was partly driven by the media (Bursztyn et al., 2020). Third, it affects a salient outcome—health—about which people tend to form strong opinions. Fourth, motivated beliefs and wishful thinking have been shown to play a role in health behaviors (Oster et al., 2013; Benabou, 2015). While there

8 Studies on media bias or polarization which are closer to the spirit of our experiment include DellaVigna and Kaplan (2007) and Gentzkow and Shapiro (2006, 2010). For field experiments on the link between media, news consumption and polarization, see also Knobloch-Westerwick et al. (2015), Bail et al. (2018), or Levy (2021).

9 Chopra et al. (2022) find muted demand for fact-checking of ideologically aligned news in a sample of Democrats (with the effect driven by Democrats with strong ideological views).

have been several papers studying different aspects of beliefs about various economic variables during Covid (e.g., Coibion et al., 2020a,c; Hanspal et al., 2020; Dietrich et al., 2020; Bartik et al., 2020), to our knowledge this is the first to focus on motivated beliefs.

Finally, our study is related to a growing literature on information experiments (recently summarized by Haaland et al., 2021). However, few studies have attempted to endogenize the process of information selection (Fuster et al., 2020; Capozza et al., 2021), which is a key feature of our experimental design. Indeed, we show that motivated reasoning affects both the selection into the news articles as well as how they are processed, providing evidence of selective exposure to confirmatory information (see also Cookson et al., 2020, for field evidence on retail investors). This has implications for the design of information interventions, and suggests that experiments that exogenously provide respondents with certain signals may be limited in explaining the process of belief formation in the real world, where individuals sort into media sources and information signals based on their priors and political leanings.

The rest of the paper is structured as follows. Section 2 describes the survey design. Analysis from the first survey is presented in Section 3. Section 4 presents results from the second survey. Finally, Section 5 concludes.

2 Survey Design and Motivation

2.1 Survey Design

We designed two surveys, administered via Qualtrics and directly advertised to their US-based panelists. Details of the two surveys (including question wording) are in Appendices A.8 and A.9. Appendix-Figures A6 and A7 show schematic trees of the two surveys.¹⁰

Survey 1 was conducted online during the third week of May 2020. At the baseline, we first collected information on some demographic characteristics and elicited participants' beliefs with respect to economic outcomes and the health effects of the pandemic (the "priors"). Besides eliciting respondents' level of worry regarding the possible health and its economic consequences of Covid for them and their family/friends, we asked them for point forecasts of

¹⁰ The survey designs and sample sizes were pre-registered (AEA RCT Registry "Motivated beliefs and information selection during a pandemic", ID AEARCTR-0005850; web link: <https://www.socialscienceregistry.org/trials/5850>).

the unemployment rate and the number of Covid-related deaths in the US at the end of the year. In addition, we asked them to provide us with “probabilistic” forecasts, namely the probability that each measure will be higher than a given threshold by the end of the year (20% for the unemployment rate and 200,000 for the number of deaths), and the percent chance that they themselves will have been infected with Covid by the end of the year. Respondents were also asked about their support for social distancing measures.

Next, respondents were told that they had a chance to read a related news article. In Survey 1, they were asked whether they would prefer reading an article about the economic consequences or about the health consequences of the pandemic. Then, they were presented with two article headlines (in random order) and asked to choose their preferred one (separately for both domains). Within each domain, the two headlines differed markedly in their tone, with one being clearly more pessimistic than the other. In the economic domain, the two headlines were “*Hope for swift economic recovery builds as businesses reopen*” and “*Highest unemployment rate since the Depression era, and many jobs may not come back.*” In the health domain we used “*New data suggest the coronavirus is less deadly than we thought*” and “*Two influential forecasting models predict sharp rise in coronavirus deaths.*”¹¹

A key feature of our design is that the article assignment was randomized. While everybody received an article in the desired domain, only half of the participants were assigned the one they had chosen. After respondents had read the article, we elicited their view of the informativeness and reliability of the article, and the probability they assigned to different possible media sources from which the article could have come. Then, we again elicited beliefs with regard to the economic and health effects of the pandemic (we refer to these as “posterior” beliefs in the analysis), using the same point and probabilistic forecast questions discussed above.¹² We also used a qualitative question to measure their support for stay-home orders in case of a resurgence

11 We provided participants with actual news articles, which were slightly edited and condensed such that they had similar length. Some of the headlines were also edited. Respondents were (correctly) told: “The article is from a major U.S. news sites/sources (either top 10 newspaper by circulation, or top 3 news channel organization by viewership). We have slightly edited the article for clarity and brevity.” We had pre-tested the headlines in a pilot we ran on Amazon Turk, in order to validate the optimistic/pessimistic tone and ensure that people’s preference over them was balanced on other dimensions. See Appendix A.7 for further details.

12 Note that we do not incentivize the elicitation of beliefs. Reviewing evidence from various studies, Haaland et al. (2021) conclude that incentives have little effect on beliefs in non-political domains and when responses cannot be readily looked up—conditions that are satisfied in our context. In addition, Armantier and Treich (2013) show that beliefs are less biased (but noisier) in the absence of incentives.

of the pandemic. The survey furthermore featured incentivized choices designed to corroborate the responses. First, respondents were allowed to spend up to \$50 on Facebook ads to promote the news source of the article they just read (at no cost to them). Second, at the end of the survey, participants were asked how they would distribute a \$20-budget between the World Health Organization (WHO), an organization that explicitly supported lockdown measures, and FreedomWorks, a conservative and libertarian advocacy group whose favored policies are opposite to the WHO's.¹³

Survey 1 allows us to investigate the relationship between prior beliefs and news choice, and the nature of updating conditional on being assigned the preferred article or not. However, the survey does not reveal the source of the news article, something that is readily available in the real world and likely plays a role in information selection and processing. In Survey 2, also conducted via Qualtrics during the last week of May 2020 (on a separate set of respondents), we aim to quantify the importance of such source dependence.

The first stage in Survey 2 parallels Survey 1, with demographics, perceptions, and beliefs being elicited at the baseline. Respondents were then given the choice to read a news article. In this case, however, respondents faced a single choice between two articles, again with different tones (a pessimistic and an optimistic one). The two headlines were “*Coronavirus pandemic to likely last two years, new report says,*” and “*Scientists increasingly optimistic that a vaccine can be produced in record time.*” In Survey 2, everybody was shown the article corresponding to the headline they chose. The treatment instead consisted in randomly revealing the source. Participants were randomized into three groups: the first one was never revealed the news source (T1); the second one was revealed the news source when they were shown the article (after having chosen it based solely on the headline) (T2), and the third one was revealed the source right from the beginning (T3). The pessimistic article came from Fox News, whereas the optimistic one came from The New York Times. The choice is not coincidental: the two news outlets have discernible political stances, conservative for the first and liberal for the second.¹⁴

13 Respondents were informed that two choices (for the news source promotion) and ten choices (for the money donation) would be randomly picked and implemented (which we did after data collection was complete). Such ex-post randomized incentives are common in large-scale economic experiments and have been found to generate similar behavior as if every respondent's choices were implemented (March et al., 2016).

14 An alternative design might have randomized the sources of each headline and corresponding article, but to do

As we validate later, respondents would on average find it surprising that a pessimistic article is coming from Fox News.

The final stage of the survey was very similar to that in Survey 1, with respondents reporting their posterior beliefs, attitudes, and subjective assessment of the article they read, and being given the opportunity to advertise the news source.

2.2 Discussion

Before proceeding with the empirical analysis, it is worthwhile to briefly motivate and compare the various features of the two surveys.

In Survey 1, the elicitation of attitudes and beliefs in the baseline stage allows us to investigate the determinants of the choice of the domain and of the tone of the article, and whether the selection of information (domain and tone) is motivated. The random assignment to the article (in the preferred domain) allows us to measure whether and to which extent there is a differential update, conditional on receiving information closer to or further from what respondents wish to hear. Confirmatory bias would predict stronger updating and more attention in the case when the respondent is assigned the preferred article.

Survey 2 brings in further realism relative to Survey 1 (while simplifying other aspects of the design), and allows us to measure respondents' tendency to dismiss or prioritize an article when its source is aligned with their ideology. Specifically, comparison of T3 (the only group of respondents made aware of the article source before making the choice) with the other two groups tells us about how selection of the news article is affected by the revelation of the news source (the "selection" effect). T1 respondents are never informed about the article source, while T2 respondents learn the source after the article choice. Comparison of the two groups informs us about the role of source revelation in information processing, conditional on having chosen an article (that is, keeping the selection fixed). This second survey thus allows us to directly speak to the question of how the (perceived) political stance of media may exacerbate biases in information acquisition and updating.

The empirical analysis for both surveys proceeds in three steps. First, we study the correlation

so without deception would have required an article to appear identically in both sources. We did not find such articles.

between demographics and priors, and assess the determinants of the article domain and tone choice. This provides evidence on how individuals' priors affect information selection, and which theories of motivated beliefs the patterns are consistent with. In Survey 2, we further study how knowing the source of the article affects the choice. Second, we test whether reading an article leads respondents to update their beliefs about the effects of the pandemic, and more importantly, whether the strength of updating depends on whether a respondent saw their preferred article (in Survey 1, where article assignment is randomized) and whether the article source is aligned with their political beliefs (in Survey 2, where source revelation is randomized). Similar analyses are conducted with other judgmental or behavioral variables as outcomes: those include article assessment, attentiveness, and actual behavior.

2.3 Theoretical Framework

Our experimental design is inspired by recent theoretical work on optimal belief distortion. In Appendix A.2, we present a simple theoretical framework where an agent first picks her preferred signal and then, based on its realization, updates her beliefs. Following prior contributions in decision theory and experiments (see, for instance, Baillon et al., 2017) we assume that the agent has a utility function that nests the case of wishful thinking (as in Caplin and Leahy, 2019) and ambiguity aversion (Hansen and Sargent, 2008).

We show that this model predicts that (i) more pessimistic individuals tend to choose signals that assign higher likelihood to states of the world with lower utility, such as articles that emphasize the negative consequences of the pandemic, and (ii) belief updating is stronger in response to the preferred signal. As we discuss below, our empirical evidence is broadly in line with this theoretical framework.

Alternative modeling frameworks could also be consistent with our results. For instance, recent work by Cheng and Hsiaw (2022) presents a model where agents do not have motivated beliefs, but need to learn about the credibility of an information source, and do so based on how far the source's signal is from their priors. The agents then form posteriors by weighing the data using the updated beliefs about credibility; thereby, the agents end up "double-dipping" the data and may overweight signals close to their prior (because they find that source more

credible). This departure from fully Bayesian learning can thus generate patterns that look like confirmation bias.

2.4 Sample Description

As is common in online surveys, the survey included two multiple-choice comprehension questions in the beginning, which a respondent had to answer correctly to proceed.¹⁵ 2,712 participants cleared these attention checks. We further remove the respondents in the bottom 10% of the distribution of the time spent on the questions before reading the article.¹⁶ Our final sample for Survey 1 includes a total of 2,440 US-based participants.¹⁷ The mean (median) survey time was 16.3 minutes (14 minutes). The sample was designed to be representative of the US population in terms of age, gender, and location. We employ similar screenings for our second survey, which consists of 1,571 respondents. Here, the mean (median) survey time was 22.1 minutes (14 minutes). Respondents who had participated in Survey 1 were not invited for the second survey.

Appendix-Table A2 shows the characteristics of our sample. The variables used in this table are described in Appendix A.1. The table shows that the characteristics of respondents between the two surveys are slightly different, which we attribute to sampling error. Looking at the sample for Survey 1, 54% of respondents are females, 23% are Black, Hispanic, or Asian, 51% have a college degree, 69% are homeowners, and 47% rate themselves as liberals (and 40% as conservatives). The average age of our respondents is 49 years. The sample lines up well with national aggregates, except that it has much higher education; this is at least partly due to differential internet access and computer literacy across education groups.¹⁸

15 We asked two questions to measure participants' attention, namely "Are you interested in sports?" and "Are you interested in music?", asking them to answer "Very strongly interested" to the first question and "Not at all interested" to the second. If a participant fails to follow the instructions, the survey ends immediately.

16 We have re-run our key regressions of interest without dropping any respondents, or by dropping respondents in the bottom 1%, 2.5%, 5%, and 7.5% of the time spent on the questions before reading the article. In all these variations, results are very similar to those reported in the main text, and are available upon request.

17 We initially targeted a sample of 2,000 survey participants, who Qualtrics selected based on the total time spent on the survey. Their procedure is as follows: First, the company runs a "pilot" with 50 participants and determines a cutoff time computed as half the median time spent on the survey. Then, for the final launch, responses are recorded until 2,000 individuals complete the survey spending on the survey an amount of time higher than this cutoff. We preferred to use a time cutoff that does not involve the amount of time spent on reading the article, which might be endogenous, and chose instead a restriction based on the time spent *before* reading the article. This results in a sample slightly larger than what was originally planned. Results are qualitatively similar if we instead use the sample imposing the cutoff that was used by Qualtrics.

18 The respective figures from the 2019 American Community Survey are: non-white 21%, college education

The pandemic had far-reaching impacts. 29% of the respondents report that someone in their household had experienced a loss of employment income since February 2020 due to the pandemic. 5% report having been already infected (definitely or very likely). The corresponding proportions are higher in the second survey. 27% (19%) of respondents in Survey 1 (Survey 2) were subject to stay-at-home orders in their area at the time of the survey.

3 Empirical Analysis - Survey 1

3.1 Baseline Concerns about the Pandemic

The first section of our survey included several questions to gauge the respondent's pessimism about health and economic consequences of the pandemic. Specifically, they were asked: "*On the whole, on a scale from 0 (not worried at all) to 10 (extremely worried), how worried are you about the possible effects of a COVID-19 infection on your own health or the health of close family and friends?*" and "*On the whole, on a scale from 0 (not worried at all) to 10 (extremely worried), how worried are you about the effects of COVID-19 and the measures that have been taken to contain it on your personal financial situation or the financial situation of close family and friends?*" The former captures worries related to the health consequences, and the latter those related to the economic consequences of the pandemic. We refer to these variables as "Worried Health" and "Worried Economy", respectively. The average response, as shown in Appendix-Table A2, is 6.61 for Worried Health, and 6.03 for Worried Economy. The large standard deviations, almost 3 points for both measures, indicate that there is substantial variation in respondents' worries about the pandemic. However, at the individual level, the responses to the two questions are strongly positively correlated (correlation of 0.48), suggestive of an underlying unobservable trait of individuals to either always worry or not worry, regardless of the domain. The survey also asked respondents about their support for the stay-home orders and social distancing measures the US has implemented, on a 1-5 scale (with 1 meaning that "the measures went much too far," 3 meaning "the measures were about right," and 5 meaning

32%, homeowners 65%, and average age of 51. In the 2018 United States Congress election, the Democratic and Republican parties obtained 53.4% and 44.8% of the overall popular vote, a 8.6% difference which maps fairly closely to the fraction of survey participants who identified themselves as liberals and conservatives.

“the measures did not go nearly far enough”). This variable, which we refer to as “Lockdown Support”, has an average of 3.31 in Survey 1.

We first explore the relationship between these variables and basic demographics. This allows us to examine the determinants of worries related to the pandemic, and to verify that survey responses follow sensible patterns. Results are shown in Table 1.¹⁹ In column 1, the dependent variable is “Worried H+E”, which is the sum of Worried Health and Worried Economy. In column 2, we focus on the *difference* between these two measures, “Worried H-E.” High values of this variable characterize individuals who are disproportionately more concerned about the health consequences of the pandemic, as opposed to its economic consequences.²⁰ PC1, the dependent variable in column 3, captures participants’ overall pessimism about the effects of the pandemic. It is constructed as the first principal component of seven variables: (i) an estimate of the probability that the number of Covid-related deaths will be more than 200,000 by the end of the year; (ii) an estimate of the probability that the unemployment rate will be more than 20% by the end of the year; (iii) an estimate of the probability of becoming Covid-infected by the end of the year; (iv) a forecast of the number of Covid-related deaths by the end of the year; (v) a forecast of the unemployment rate at the end of the year; (vi) an estimate of the current unemployment rate; and (vii) an estimate of the number of Covid-related deaths so far. The variable is demeaned and divided by its standard deviation for ease of interpretation. Appendix-Table A3 reports correlations between these seven variables and PC1, and shows that this proxy does well in summarizing information about participants’ beliefs.²¹

The last column of Table 1 uses Lockdown Support as the dependent variable. As Appendix-Table A4 shows, Worried H+E, PC1, and Lockdown Support are positively correlated, but their association is far from perfect. Worried H-E is positively correlated with Lockdown Support, as one would expect.²²

Table 1 shows many intuitive correlations. For instance, respondents with underlying health

19 To save space, Table 1 omits standard errors, which are reported for completeness in Appendix-Table A5.

20 Given the strong association between Worried Health and Worried Economy, the analysis uses the sum of the two, rather than running a “horse race” between two variables that are likely to capture the same underlying phenomenon. However, there is still high heterogeneity with regard to which domain individuals are more concerned about, which we capture through the difference Worried H-E.

21 The fraction of variance explained by the first principal component is 34.06% (33.85% in Survey 2).

22 Individually, Worried Health and Worried Economy are both positively correlated with Lockdown Support, though the correlation is stronger for Worried Health.

conditions that make them more vulnerable to a Covid infection are more worried, but especially about health effects. Respondents who are not working or who experienced an employment loss within the household also tend to be worried, but more strongly so about the effects on their financial situation; perhaps surprisingly, those with an employment loss express stronger support for lockdown measures. Respondents with higher incomes are generally less worried, but are relatively more concerned about health effects; those who hold stocks worry more about financial consequences.

Self-reported general preference measures also correlate with Covid worries, pessimism, and lockdown support: higher willingness to take risks is associated with worrying less about health and lower support for lockdowns, while higher patience is associated with higher worry (especially in the health domain) and stronger lockdown support. Stronger trust in others reduces worries and lockdown support. In terms of other personal characteristics, female and married respondents tend to worry more and are more pessimistic and supportive of lockdowns. Perhaps surprisingly, age negatively correlates with being worried and Covid-pessimism.²³ In terms of political beliefs, we find that liberals tend to be much more worried (in general and about health), pessimistic, and supportive of lockdowns. This already provides a first glance at how political stance shapes priors. Finally, local Covid-severity seems to matter relatively little, but more so through deaths per capita than confirmed cases. The local unemployment rate correlates negatively with lockdown support, but otherwise does not enter the regressions significantly.²⁴

While Table 1 shows meaningful relationships, it is worth noting that the rich set of controls can explain less than 20% of the variation in responses. Also, not all the variables that are significant in Table 1 are also significant in the equivalent analysis for Survey 2, shown in Appendix-Table A10, possibly due to the smaller sample size. However, the variables with the strongest effects in Survey 1, such as employment loss during the pandemic, the number of health conditions a respondent has, their political orientation, and self-assessed preference traits, remain strongly significant in Survey 2.

23 This result is consistent with Bordalo et al. (2020) who find that perceived personal health risks associated with Covid fall with age. The bivariate correlations of age and the four dependent variables are also negative, except for Worried H-E, where it is positive (0.16).

24 Note that we have also run all the regressions in the paper with county fixed effects, and obtain qualitatively similar results (although often less precisely estimated, since we lose quite a bit of variation).

3.2 Domain Choice

We next assess whether respondents' article choice is related to their initial beliefs. This could provide a first indication of self-confirmatory selection of information. Our information experiment features binary choices with respect to the article domain (health or economy) and to the article tone (pessimistic or optimistic). The domain choice might be driven by several elements, which include respondents' worries in the specific domain.

Table 2 presents results from linear probability models where the dependent variable is a dummy that equals one if the participant prefers to read an article on the health consequences of the pandemic; this is the case for 65% of respondents. Here and in what follows, whenever the dependent variable is a dummy, we multiply the coefficients by 100 for ease of interpretation.

In column 1, we see that both Worried H+E and Worried H-E positively correlate with the likelihood of preferring an article in the health domain. Both coefficients are economically large and very precisely estimated, with *t*-statistics over 10. In column 2, the regressor is PC1, our proxy for the level of pessimism of the participant's beliefs. The coefficient is again precisely estimated and large: a one-standard-deviation increase in PC1 is associated with an 8 percentage point higher likelihood of preferring the health article. In column 3, the measure for lockdown support also strongly predicts the choice of the health domain, with a coefficient equal to 13.6. Column 4, which includes all the four predictors, shows that all of them remain positive and statistically significant. Columns 5-8 of the table adds controls to the corresponding specifications in the first columns; the set of controls is the same as in Table 1. While the coefficient magnitudes are slightly reduced, the qualitative conclusions are unchanged. In sum, those respondents generally more worried/pessimistic about the effects of the pandemic (higher Worried H+E and PC1), those who worry more about health effects (higher Worried H-E), and those more supportive of lockdown measures are substantially more likely to prefer reading about the health consequences than about the economic consequences of the pandemic.

3.3 Choice of Headline Tone

We now turn to testing whether choosing optimistic or pessimistic articles within either domain is related to pre-existing optimistic or pessimistic attitudes. The dependent variable now is a

dummy that equals one if the participant chooses the pessimistic article. Columns 1 to 4 of Table 3 identify the predictors of choosing the pessimistic article in the health domain, while columns 5 to 8 identify those in the economy domain. The regressors, together with controls, are the same as in Table 2. In both domains, the share of respondents preferring each article is close to balanced, as we had intended. The variable Worried H+E positively and significantly predicts the choice of a pessimistic tone article in both domains, with coefficients fairly similar in magnitude (1.14 in the health domain and 0.96 in the economy domain). The coefficient on Worried H-E is positive in both domains, although much larger in the health domain.

Both PC1 and Lockdown Support enter with a strongly positive and statistically significant coefficient as well (see columns 2 and 3 for the Health domain, and 6 and 7 for the Economy domain). In columns 4 and 8 we run “horse races,” by including our four regressors jointly. PC1 and Lockdown Support retain their predictive power, whereas Worried H+E now becomes insignificant. Worried H-E is positive and significant in the Health domain, and small and insignificant in the Economy domain.²⁵ The larger impact of Worried H-E in choosing the pessimistic article in the Health domain (relative to in the Economy domain) suggests that health concerns matter relatively more in the choice of the tone of the article in the health domain.

The coefficients are economically large. For example, a one-point increase in support for lockdown measures is associated, in the most conservative specifications of columns 4 and 8, with 10.1 and 6.8 percentage points higher likelihoods of choosing the pessimistic articles in the health and economy domains, respectively. Moreover, the estimates are fairly precise: the lower bounds of the 95% confidence intervals are 8.1 and 4.8 percentage points.

To visually appreciate the results, Appendix-Figure A1 plots the fraction of participants who chose the pessimistic headline in the health domain and in the economy domain for each quintile of Worried H+E, Worried H-E, PC1, and for each level of support for lockdown measures. A distinctive monotonically increasing relationship emerges also in these “non-parametric” tests.

These patterns suggest that individuals sort into domains and, in particular, article tone based on their priors. It is certainly not the case that those more worried about health and economic effects of the pandemic make an effort to avoid pessimistic articles in order to “feel better.”

²⁵ In unreported tests, we find that the difference between the coefficients estimated in the two domains is statistically significant at the 5% level.

Instead, the results suggest that people seek information to confirm their pre-existing beliefs and policy stances: for instance, those more worried about health consequences of Covid prefer to read health articles, and ones with pessimistic tones; the same is true for respondents who more strongly support lockdown measures.

It may be surprising that measures of pessimism and lockdown support are also associated with preferring the pessimistic article in the economy domain (and the same is true for Worried H-E in column 5 of Table 3). However, respondents plausibly view economic effects and health effects as complementary (i.e., that there is no perceived “health-economy tradeoff”), which is a view held by many economists as well (e.g., Fernández-Villaverde and Jones, 2020). Alternatively, perhaps the specific optimistic headline may be interpreted as underlining economic costs from lockdowns (“*Hope for swift economic recovery builds as businesses reopen*”), and respondents who support lockdowns and are more worried about health may want to avoid reading about this.

One might be concerned that the patterns in Table 3 may reflect omitted variables or reverse causality. For instance, a respondent might be an innate optimist (or have engaged in wishful thinking prior to entering our survey) and find that pessimistic articles are generally uninformative or misleading, therefore exhibiting a general preference for optimistic articles that they also express in our survey. In Appendix A.5, we report an instrumental variables analysis where Covid-worries/pessimism are instrumented for with predetermined/plausibly exogenous factors (job loss and pre-existing health conditions). While there are some caveats to this analysis, the results are consistent with a causal link from baseline worries and pessimism to a preference for news articles with a pessimistic tone.

3.4 Information Processing

3.4.1 Belief Revision

While theories of optimism bias or wishful thinking predict that agents would tend to prefer positive signals, theories of confirmatory biases feature denial of evidence that contradicts prior attitudes (whether pessimistic or optimistic). Such theories also predict asymmetric updating based on the proximity of the signal to own priors. Hence, independently of whether the respondent is an optimist or a pessimist, she would consistently assign more weight to signals

that confirm her views and less weight to signals that contradict them. We have already shown that there is little evidence of wishful thinking in the choice of the signals. Investigating biases in belief updating is harder because, in most natural settings, individuals select the signals. Our experimental setup features a randomization at this stage which allows us to investigate this: after soliciting the article choice, half the respondents are assigned an article different from the chosen one (though in their preferred domain). Appendix-Table A6 shows that the characteristics of participants who receive the preferred article and those who do not are similar: the differences are either statistically insignificant or economically small.

To measure belief updating, we estimate the following econometric model:

$$\begin{aligned} \Delta y_i = & \alpha Pessimistic_i + \beta Pessimistic_i \times Preferred_i + \\ & \gamma Preferred_i + \theta y_{prior,i} + \delta' X_i + \varepsilon_i, \end{aligned} \quad (1)$$

where Δy is the difference between the posterior belief, measured after the individual has read the article, and the prior. *Pessimistic* is a dummy equal to 1 if the participant reads the article with the pessimistic headline. *Preferred* is a dummy equal to 1 if the participant reads her preferred article. y_{prior} is the prior belief of the survey participant. This is included in all specifications to control for mean reversion. In addition, as all the priors are bounded, we control for dummies corresponding to corner choices (0 and 100). X is the usual vector of controls. This model is estimated for two subgroups of participants, those who chose the Health domain and those who choose the Economy domain (and were consequently shown an article in that domain). Our coefficient of interest is β . A positive value reveals that readers of the pessimistic article become differentially more pessimistic if the article was the preferred one.

Results are reported in Table 4. We focus on the revision in three different measures: the reported probability that the unemployment rate will be higher than 20%; the reported probability that the number of Covid-19-related deaths will be higher than 200,000; and an assessment of the likelihood of being infected by the end of the year.²⁶ In columns 1-3, we include the 1,577

²⁶ We disregard the revision in point forecasts for the number of deaths and the unemployment rate since these exhibit much less variation, with a large mass of observations centered at zero, and, in the case of the number of deaths, some significant outliers. Yet, although the statistical power for these is limited, results are qualitatively similar when analyzing revisions in these belief measures.

survey participants who chose to read an article in the health domain, and columns 4-6 include the 863 participants who selected the economy domain.

We find evidence of self-confirmatory updating. Specifically, the coefficients on the interaction term are all positive. They are also strongly statistically significant, except when the dependent variable is the probability of becoming personally infected.²⁷ One plausible explanation for this discrepancy is that the articles are primarily about aggregate developments and arguably contain little information relevant to one's individual probability of infection. With respect to the unemployment and fatality forecasts, the coefficients are of sizeable magnitude, ranging between 4.2 and 4.6 percentage points in the health domain, and 7.7 and 11.3 percentage points in the economy domain.²⁸ These estimates are roughly about 0.15-0.30 of the underlying standard deviation (in the prior belief). This is a large impact—the effect sizes for the information experimental studies included in the review by Haaland et al. (2021) vary between 0.1-0.2. The fact that reading a pessimistic article in one domain (e.g. health) also makes those respondents who had preferred that article more pessimistic in the other domain (e.g. the economy) again suggests that our respondents tend to view the two dimensions as complementary.

Appendix-Figure A2 displays average revisions for the three dependent variables by group of participants (article choice \times article assignment), for each of the domains.²⁹ Interestingly, most of the participants revise their beliefs downwards. However, the tone of the article clearly matters: average downward revisions are larger for the groups that are assigned the optimistic article. And consistent with confirmatory bias, among respondents who are assigned the optimistic article, average revisions are smaller for those who had preferred the pessimistic article.

27 The uninteracted coefficients on “Preferred” and “Pessimistic” are all negative and in some cases statistically significant. “Preferred” reflects the revisions of participants who wanted the optimistic article and got it; they tend to become more optimistic (relative to the omitted group of people who wanted the pessimistic article but got the optimistic one). “Pessimistic” reflects the revisions of participants who wanted the optimistic article but got the pessimistic one; they tend not to revise differentially compared to the omitted group.

28 The pessimistic article in the health domain presents respondents with a forecast that, by the end of the year, “134,000 people will die of COVID-19 in the United States.” Given that we ask respondents the likelihood that the number of deaths will be higher than 200,000, some participants with especially pessimistic priors may actually perceive this article as optimistic. Notice that this does not necessarily affect the interpretation of our results, to the extent that the optimistic headline is still perceived as less pessimistic. To rule out the possibility of an “anchoring” effect we have also rerun the tests of column 2 by splitting the sample in two subgroups, depending on whether the respondent had a prior belief regarding the number of deaths below or above 134,000, and found very similar estimates.

29 The graphical evidence does not perfectly map to the estimates of Table 4 since regressions include the prior beliefs as well as the control variables.

One potential concern with this result, and with those that follow, is that respondents who receive their non-preferred article may feel antagonized and refuse to read it. To address this concern, we replicate the results of Table 4 by limiting the analysis to respondents who spend a minimum amount of time on the survey page with the article. Panels (a) and (b) of Appendix Table A7 show that our main conclusions are unaffected if we drop people whose time spent on the screen with the article is below either the 10th or the 25th percentiles of the distribution in the full sample, respectively. This exercise should be viewed simply as a robustness check since the time spent on a given screen is endogenous and likely affected by the treatment.³⁰ In addition, in unreported tests we find that the “preferred article” treatment is unrelated to both the time spent reading the article and an assessment of how interesting the participant found the survey.³¹ Again, these results go against the hypothesis that respondents not receiving the preferred article feel antagonized and exert less effort on the survey.

3.4.2 Assessment of Article

The survey included questions on whether respondents found the article that they had read informative and reliable. We also observe an indirect indication of whether a respondent was interested in an article’s content, namely whether they correctly recalled its headline.³² Finally, our survey included an actual behavioral measure: respondents were asked whether they wanted to promote the news source of the article they had read (recall that the news source was never

30 Qualtrics does not record incomplete responses, so we cannot formally test whether being assigned to the less preferred article is associated with a higher rate of exiting the survey midway. However, given that the random assignment to the two groups (receiving the preferred article or not) is determined upon entering the survey and that the data only include respondents who complete it, participants who read their preferred article should be overrepresented in the dataset if that were the case. Instead, we find that the final sample is evenly divided between the groups of participants who did and did not receive their preferred article (49.7% and 50.3%).

31 Specifically, we regress both the minutes spent reading the article and a 1-to-5 rating of how interesting the participant found the survey on a “Preferred Article” dummy and a number of controls, following the specification described in Section 3.4.2. The time spent reading the article is winsorized at the 1% level, given that the variable is extremely skewed. (We reach similar conclusions if we do not winsorize the variable or if we take the logarithm of it.) We estimate coefficients equal to 0.01 (standard error= 0.06) and -0.05 (standard error= 0.05), respectively. In addition, we also find that the time spent on the survey *after* having been presented with the article is not significantly associated with the treatment.

32 Respondents were presented with the list of the four headlines and were asked to identify the headline of the article just read, with the idea that the ability to correctly recall the headline is a reasonable proxy for the interest and attention with which the respondent read the article. 82.6% of respondents answered correctly.

revealed in Survey 1). Specifically, they were provided with a \$50 budget and asked how much of it they wanted to spend on Facebook ads for the news outlet whose article they had read.³³

In Table 5, we test whether these measures of appreciation of the article are significantly different depending on whether the article was the one that the respondent had marked as preferred. For this analysis, all participants were pooled together independently of the domain chosen. The different dependent variables are regressed on a *Preferred* dummy that equals 1 if the participant read her preferred article. We also include four dummies corresponding to the articles shown. In column 1, the dependent variable is a 1-7 assessment of the informativeness of the article. Respondents who receive the preferred article rate it more informative by 0.4 points, about a third of the standard deviation of the measure. Similarly, they rate it more reliable by 0.6 points (column 2), or 40% of a standard deviation. They further appear to read the article with more attention, as they are 5 percentage points more likely to pass the attention check (column 3), and spend, on average, \$4.5 more to advertise the news outlet (column 4). This is a significant amount, considering that the average donation is \$26.5.³⁴

These patterns are consistent with confirmatory biases. On the other hand, an alternative explanation—such as respondents preferring to read an article about which they are ex-ante uninformed—could generate some of these patterns even with perfectly rational individuals. For example, if individuals choose an article because they were ex-ante uninformed about its topic (based on the headline), they could arguably find that article more informative and read it with more attention (and hence be more likely to recall the headline). Such an explanation is unlikely to have an impact on the assessed reliability of the article—however, as discussed above, respondents that receive the preferred article also find it significantly more reliable. Similarly, they allocate more money to promote the media source that published the article. We believe

33 Any amount not donated would be lost. It was also specified that advertising choices would be implemented for two randomly selected participants. We provided a link to a website where we posted copies of the donation receipts once the data collection was completed.

34 One concern could be that this test is capturing how favorably a respondent views the outlet she predicts the article originates from. For example, a participant who considers CNN trustworthy may be inclined to advertise the preferred article not because she appreciates its content, but because she believes that CNN is a likely source for the article read. We can address this concern because we ask respondents to assess the likelihood that the article came from each of eight news outlets, which includes the actual source (the outlets are Breitbart, CNN, Fox News, MSNBC, New York Times, Washington Post, USA Today, and Wall Street Journal. We also include the residual category *Others*). In unreported tests, we find that controlling for this vector of assessed probabilities produces virtually identical results.

these patterns make confirmatory biases the most plausible explanation for how respondents assess and perceive the article.

3.5 Impact on Behavior

Our interest in exploring belief formation stems from the presumption that beliefs matter for behavior. Therefore, any biases in belief updating should also be reflected in behavior and support for different policies. Our survey included two measures at the posterior stage that allow us to test this link. First, respondents were asked for their support for future lockdown policies (in case of a resurgence of the pandemic). Second, participants were endowed with a \$20 budget, and asked to allocate it between WHO and FreedomWorks, two organizations that have promoted opposing views with regard to the need for lockdown policies.³⁵

Equation (1) is estimated again with the dependent variable being either a 1 to 5 measure of support for lockdown policies or the dollar amount allocated to the WHO. The first is a measure of policy support (which would likely affect behavior), whereas the second is an actual incentivized behavior.³⁶ The controls include the support for the current lockdown policies alongside the extensive list of other controls as in our analyses above.

Results are shown in Table 6. In the first two columns results are reported for the subsample of individuals who selected the Health domain, whereas columns 3 and 4 report results for those who chose the Economy domain. Again, in line with confirmatory bias, the interaction between pessimistic prior beliefs and the event of receiving the preferred article positively and significantly affects the choice of supporting lockdown and the dollar amount donated to the WHO. We find that the support for lockdown policies, upon receiving the desired pessimistic article, rises by 0.46-0.56 points (one-third to one-half of the standard deviation of the measure), and the donation to WHO rises by \$1.95-\$3.45.

35 The two organizations' positions were underlined by showing relevant quotes. Specifically, we informed participants that *The WHO has been warning against a rushed end to coronavirus lockdowns. According to its Director-General, "The risk of returning to lockdown remains very real if countries do not manage the transition extremely carefully and in a phased approach"* (Reuters).

Similarly, we reported that *FreedomWorks has claimed, on its website, that "the hidden costs of a closed economy are staggering" and "the sooner we get America back to work and open as much of the economy as we safely can, the fewer of these hidden casualties there will be."*

36 As before, we provided participants with a link where they could see screenshots of the donation receipts.

4 Empirical Analysis - Survey 2

The analysis in the previous section shows that individuals have a tendency to put more weight on the most preferred signals. Those distortions may often be politically driven. In fact, we see suggestive evidence of this in Survey 1. Table 1 shows that liberal respondents tend to be more worried, pessimistic, and supportive of lockdown policies. These factors are significant determinants of domain and article choice (Tables 2 and 3). There is then good reason to expect that assignment to the preferred article widens disagreement between individuals with different political leanings. This is exactly what we find: the difference between liberals and conservatives in average revisions in beliefs and support for lockdown policies are *larger* for those respondents who receive the preferred article, relative to those who do not (see Appendix Table A8).³⁷

These differences are likely to be exacerbated if the source of the news article is revealed. Politically-motivated beliefs are receiving increasing attention as they epitomize people's motivation to justify and support self-serving policy choices, and policy choices have been heavily contested during the pandemic. Moreover, there is an association between political orientation and media sources, as conservatives and liberals choose to consume different media sources and may find news from, for example, The New York Times or Fox News differentially informative.³⁸ In other words, media sources act as an echo-chamber that allows agents to preserve and proselytize their self-serving views and actions. Putting those things together implies that if we were to find source dependence in news choice and interpretation, this would strengthen the argument in favor of motivated reasoning. For this reason, we now go one step further and explore politically-motivated reasoning by randomizing the source revelation.

Specifically, we extend Survey 1 by randomizing whether and when the source of a news article is revealed. Other aspects of the design are simplified relative to Survey 1; in particular, there is no choice of article domain, and all participants are shown their preferred article. The design of the survey, and its rationale, are explained in Section 2. Recall that the survey randomly divided participants into three groups: the first one is never revealed the news source (T1); the

³⁷ The table reports difference-in-differences estimates which are positive for each revision measure, but not always statistically significant. However, the p -value of a test that the estimates are jointly positive is 0.055.

³⁸ See, for example, Taber and Lodge (2006), Nyhan and Reifler (2010), or Thaler (2020b).

second one is revealed the source after having chosen the article but before reporting posterior beliefs (T2), and the third one is revealed the source before making the article choice (T3).³⁹

If we were to select an optimistic article from Fox News, most likely (also according to the results of our previous survey) an optimistic Republican would choose it and update his beliefs even more strongly (than in Survey 1) in the direction of optimism. This would almost be tautological at this stage. To examine the potential “inner tensions” that motivated reasoning brings about, we instead select an optimistic article from the NY Times and a pessimistic one from Fox News. These tone-source combinations were meant to “surprise” the respondents (and we will show later that they did), which in turn will allow us to better assess the strength of the source-dependence effects on information choice and updating.

4.1 Choice of Headline Tone

Before analyzing the effect of source revelation on article choice, we first check whether we can replicate the results of Survey 1, by investigating which variables are correlated with the choice of the article tone. To make results comparable across surveys, participants in group T3 are excluded, so that we include only participants who chose the article without knowing its source. Our findings, reported in Appendix-Table A11, are similar to those obtained when analyzing Survey 1: Worried H+E, PC1, and lockdown support are positively correlated with the choice of the pessimistic article.⁴⁰

To visualize the impact of the predictors for tone choice, Figure 1 displays the fraction of participants who choose the pessimistic headline for each quintile of Worried H+E and PC1, and for each level of lockdown support. The red bars (which only consider participants without knowledge of the source when choosing the headline) show a clear positive relationship between the predictors and the fraction of participants who choose the pessimistic article. For example, 47.8% of respondents in the top quintile of the PC1 distribution choose the pessimistic article,

³⁹ Appendix-Table A9 shows that the randomization across the three groups was successful. The table shows means of selected variables for the three groups, and p -values of Wald tests of the null hypothesis that the three means are equal. Only in one case do we reject the null hypothesis at conventional significance levels, which is what one would expect to happen by chance.

⁴⁰ The only difference is that now Worried H-E is not a significant predictor of the headline choice. A possible explanation for this is that the two headlines, “*Scientists increasingly optimistic that a vaccine can be produced in record time*” and “*Coronavirus pandemic to likely last two years, new report says*” were not explicitly assigned to either the health or economy domain, and arguably had implications in both domains.

compared to 29.6% in the bottom quintile. A similar pattern emerges for Worried H+E and lockdown support.

4.2 Effect of Source Revelation on Article Choice

Next, we use the full sample of participants to test whether revealing the source *ex ante* has an impact on article headline choice, over and above the prior beliefs. For this purpose, we estimate the following model:

$$Pessimistic_i = \alpha Z_i + \beta Z_i \times Revealed\ Before_i + \gamma Revealed\ Before_i + \delta' X_i + \varepsilon_i. \quad (2)$$

The dependent variable is a dummy equal to 1 if the participant chooses to read the article with the pessimistic headline. The variable *Revealed Before_i* is a dummy that equals 1 if participant *i* is revealed the article source before making the choice (which is the case for group T3). *Z* is one of the three predictors that have displayed a significant forecasting ability for the likelihood of choosing the pessimistic article (see Appendix-Table A11): Worried H+E, PC1, and Lockdown Support. Finally, *X* is the usual vector of control variables. The parameter α in this specification shows the impact of *Z* in choosing the pessimistic article when the article source has not been revealed to the participant (groups T1 and T2). The parameter β shows how the role of *Z* in choosing the pessimistic article is impacted when the source is revealed to the participant (prior to choosing the article). The “net effect” of the participant’s belief on article choice for the individuals made aware of the source before making the choice is the sum $\alpha + \beta$.

Results for this econometric specification are shown in columns 1-3 of Table 7. The estimates, based now on the full sample of respondents, show that revealing the source diminishes the predictive power of participants’ priors for the article choice. While the estimates of α are positive, estimates of β are negative. For example, in column 2, the coefficient α on the standalone PC1 variable is equal to 5.36 and highly statistically significant. However, once the news source is revealed, the impact of the standalone PC1 variable is only 1.21 (5.36 – 4.16), and not statistically different from zero. This can also be seen from the “Net Effects” shown at the bottom of the table, equal to –0.53, 1.21, and –1.81 when the regressors are Worried H+E, PC1, and Lockdown Support, respectively. They are all small and statistically insignificant;

moreover, they are relatively precisely estimated (with the exception perhaps of column 2), with upper bounds of the 95% confidence intervals equal to 0.25, 5.1, and 2.1, respectively. Hence, prior beliefs predict the article choice only if participants are unaware of its source. Once the news source is revealed, prior beliefs fail to display predictive power for the article choice—we interpret this result as clear evidence of source dependence.

We see the same patterns in the light bars in Figure 1. While the dark bars in the figure show that there is positive relationship between the fraction of participants who chose the pessimistic headline and prior beliefs, no such relationship is observed for group T3, which consists of participants who are aware of the article source before making the choice.

Given the political inclination of the newspapers chosen for the experiment, it is legitimate to conjecture that it is the heterogeneity in political beliefs that determines the choice of the article, once the news source is revealed. To test this hypothesis, we augment equation (2) by adding, as additional predictors, *Liberal* and *Liberal* \times *Revealed Before*. *Liberal* is a dummy equal to 1 if the respondent defines herself as “very liberal/democrat”, “liberal/democrat”, or “leaning liberal/democrat”.⁴¹ The coefficient on the interaction term *Liberal* \times *Revealed Before* is the one of interest. It measures the extent to which, upon learning that the source of the pessimistic article is Fox News, liberals are either less or more likely to choose it.

Results for this specification are shown in columns 4–6 of Table 7. The coefficient on the *Liberal* dummy is positive, ranging between 7.5 and 9, and statistically significant. Hence, liberal participants are substantially more likely to choose the pessimistic article. This is in line with evidence that in the US, Democrats have been much more concerned about the effects of the pandemic.⁴² Thus, this corroborates the presence of confirmatory beliefs: to the extent that Democrats are more worried about the effects of the pandemic, it is consistent with our evidence from Section 3 that they will be more likely to prefer an article with a pessimistic tone.

This result is conditional, however, on the participants not knowing the article source. Re-

41 In unreported robustness checks, we have replaced *Liberal* alternatively with: (i) a dummy equal to 1 if the participants plans to vote for Biden in the presidential election, and (ii) a dummy equal to 1 if the participants defines herself as strongly Democratic, weakly Democratic, or leaning toward the Democratic Party (in a second, later elicitation of political stance). We find very similar results.

42 For example, a Gallup poll from April 2020 found that 73% of the Democrats were worried about getting Covid-19 and 54% were worried about experiencing financial hardship. The respective figures for conservatives were 36% and 34%, respectively. See “In U.S., More Fear COVID-19 Illness Than Financial Harm” on news.gallup.com, April 16, 2020. See also Painter and Qiu (2021) and Barrios and Hochberg (2020).

member that the pessimistic article is from Fox News and the optimistic one is from the NY Times. While Democrats might prefer the pessimistic article, they might be reluctant to choose it if they know that it comes from a news source not aligned with their political beliefs. Which of the two channels dominates the other is an empirical question.

We find that the coefficient on the interaction term $Liberal \times Revealed\ Before$ is negative and very large, ranging between -25.1 and -26.6 . As a result, the sum of the two coefficients (on $Liberal$ and $Liberal \times Revealed\ Before$) is negative and economically large. It varies between -16.5 and -18.1 , and is significant at the 1% level in all the specifications, with upper bounds of the 95% confidence intervals ranging between -8.3 and -10 .

To summarize, in Survey 2 we also find evidence consistent with confirmatory bias driving news choice. However, survey participants also appear to prefer sources aligned with their political preferences. Once both these effects are present, political beliefs dominate and, in our setting with counterintuitive source/tone combinations, eliminate the effect of priors about the pandemic on the article choice.

4.3 The Effect of Revealing the Article Source on Information Processing

To further examine the role of source dependence in information processing, we now test whether the news source revelation, in combination with political beliefs, affects article assessment. We focus on the subsample of respondents that chose the article without knowing its source (T1 and T2), half of whom were made aware of it after having made the choice but before reporting posterior beliefs and attitudes (T2). We estimate the following model:

$$y_i = \delta' X_i + \alpha Pessimistic_i + \beta Revealed\ After_i + \gamma Liberal_i + \theta Pessimistic_i \times Revealed\ After_i + \phi Revealed\ After_i \times Liberal_i + \rho Pessimistic_i \times Liberal_i + \mu Pessimistic_i \times Revealed\ After_i \times Liberal_i + \varepsilon_i. \quad (3)$$

The variable $Pessimistic$ is a dummy equal to 1 if the individual was shown the article with the pessimistic headline. $Revealed\ After$ is a dummy equal to 1 if the participant is shown the article source after having made the choice, that is, it equals one for participants in T2 and zero for those in T1. $Preferred$ is a dummy equal to 1 if the participant was assigned the preferred

article to read. *Liberal* is defined as in Section 4.2. The usual controls are included in the vector X . Our coefficient of interest is μ , which captures the impact of news source revelation combined with the article tone and the respondent's political beliefs.

For brevity, the coefficient μ for our main dependent variables of interest is shown as solid circle, with its respective 95% confidence intervals, in Figure 2.⁴³ The key coefficient of interest is negative in all the panels: liberals who are informed that the source of the pessimistic article is Fox News find it less reliable and informative, are less likely to recall its headline accurately, and spend less money on advertising it. However, only the last coefficient is statistically significant (shown in Panel d): liberals are willing to spend \$16.9 less to advertise the pessimistic article upon learning that its source is Fox News. These conclusions are mirrored for conservative respondents and the NY Times.

Note that the effect of revealing the news source should be concentrated among participants who have a strong prior that conservative news sources have a more optimistic tone with regard to the effects of the pandemic. We can directly investigate this hypothesis, since the survey included a question that asked participants about the *ex ante* likelihood of the article coming from different news sources (based on the headline only). Using the response to this question, we compute the perceived probability that the article comes from a liberal news source and the perceived probability that the article comes from a conservative news source.⁴⁴ By subtracting the latter from the former, we obtain a proxy for how surprised a participant might be to learn that the pessimistic article she has chosen comes from a conservative news outlet.

Based on this, individuals are then sorted into two groups, depending on whether this measure is above or below the sample median. Appendix-Table A12 shows these statistics for both Survey 2 and, for completeness, Survey 1.⁴⁵ The specification in equation (3) is estimated for each of the two groups separately. Figure 2 plots the coefficient, μ , on the triple interaction term, with confidence intervals, for individuals in the bottom of the distribution ("Low Surprise" group, as squares) and in the top of the distribution ("High Surprise" group, as diamonds). Results are

43 For completeness, Appendix-Table A13 displays the regression coefficients.

44 We define liberal news sources as CNN, The New York Times, MSNBC, and The Washington Post; conservative news sources are Breitbart, Fox News, and The Wall Street Journal.

45 The table also shows these statistics for just NYTimes and Fox News, where we also observe substantial differences in reported probabilities.

as expected: the coefficient μ is never significant for the sub-sample of individuals who are not particularly surprised to learn that the article with the optimistic tone is from a liberal source (or that the article with the pessimistic tone is from a conservative source). Conversely, we find a negative and (at least marginally) statistically significant coefficient for the High Surprise group (except when the dependent variable is the informativeness rating).

Finally, we turn to the impact on beliefs and intended behavior. Appendix-Table A14 reports estimates of equation (3), where the dependent variables are the revisions in the probability that (1) the number of deaths will be higher than 200,000 by the end of the year; (2) unemployment will exceed 20% at the end of the year; (3) the respondent will have been infected by Covid; (4) that they would sign up for a contact tracing app.⁴⁶ Finally, like in Survey 1, we again elicit (5) support of future lockdown policies in case of a new wave of the pandemic.

Source revelation has little impact on belief revision. However, the last column of the table shows that liberals who learn that the pessimistic article comes from Fox News increase their support for lockdown policies by an even larger (economically and statistically significant) amount. A possible explanation for this result is that liberals, who are on average more in favor of lockdown policies, see their support for this policy vindicated once they learn that a news source that they typically avoid has, unexpectedly, views aligned to theirs. Hence, liberal respondents who see that even Fox News has a pessimistic outlook regarding the effects of the pandemic become even more alarmed and increase their support for stricter lockdown measures. Consistent with this interpretation, we find that this effect is concentrated among participants who are particularly surprised to learn about the political stance of the article source (see Panel e in Appendix-Figure A4, which shows the equivalent of Figure 2 for the belief revisions and lockdown support). In Appendix A.6, we build on this insight and show that source revelation leads to greater polarization in beliefs. However, it is important to point out that it is not the case that individuals discount the information entirely upon learning that the source of information is not aligned with their political ideology.

Finally, note that the results are broadly consistent with the predictions of the theoretical

⁴⁶ This behavioral measure was added in Survey 2; it was not elicited in Survey 1. See Appendix A.9 for the question wording.

literature on demand-driven media biases.⁴⁷ In these models, consumers' updates to media news result from a strategic interaction with the media source. The extent of updating in the direction advocated by the source, or how persuasive the source is, depends on many factors (consumer heterogeneity, their priors, the incentives of the sender, and so on).⁴⁸ Source dependence could be more likely when the content involves policy decisions, namely a situation in which the coordinated update of many receivers is more valuable.

5 Conclusions

Casual observation suggests that the COVID-19 pandemic, an enormously consequential event, has generated very wide disagreement across people when it comes to its potential health and economic consequences, and the best policy response to deal with them. While some of this disagreement may reflect differential prior expertise, much of it likely arose from differences in the news about the pandemic that people consume. In some cases, the news that one sees may be exogenously determined, but usually one selects which news sources to consume and how much attention to pay.

In this paper we have shown through two survey experiments how this self-selection of news may enhance existing disagreements, since our respondents seem to be subject to confirmation bias. In our first experiment, we found strong evidence that people with more pessimistic prior beliefs, and those supportive of lockdown policies against the pandemic, tend to prefer reading articles with a pessimistic tone, both in the health domain (which such respondents had a preference for) and the economic domain. Furthermore, respondents updated their beliefs more strongly when presented with their preferred article, rated the article more favorably in terms of informativeness and reliability, and also adjusted their (incentivized) behaviors accordingly.

Our second experiment studied how awareness of the article source (in our case, flagship liberal or conservative newspapers) affected information selection and processing. We found a very strong aversion to consume news "from the other side", even when the headline would otherwise have been appealing. Furthermore, even the content of the article was judged differently upon

47 See for instance Mullainathan and Shleifer (2005) or Bernhardt et al. (2008) for models in which news are shaped based on consumers' confirmatory biases. Similarly, see Gentzkow and Shapiro (2006) for a model in which media news are offered to raise demand from consumers who trust more news close to their priors.

48 See DellaVigna and Kaplan (2007).

learning that it came from an unexpected source. While in our setting, the article sources were picked to be counter-intuitive, in reality it is of course more likely that news media will cater to the preferences of their readers (and indeed, we find evidence that this is what respondents expect). Thus, this source dependence can lead to another type of confirmation bias that further enhances polarization of beliefs.

While we have used the setting of the COVID-19 pandemic as a natural experiment to study these issues, they are certainly of more general interest across different applications. Our results underline that motivated beliefs, and in particular confirmatory bias, are a phenomenon that models of belief formation and updating should accommodate. Furthermore, policy makers can benefit from understanding that reaching out to parts of society with opposing views may be more effective when it occurs through sources these audiences are more familiar with.

References

- Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., and Yang, D. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics*, 191:104254.
- Armantier, O. and Treich, N. (2013). Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging. *European Economic Review*, 62(C):17–40.
- Bail, C., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., Lee, J., Mann, M., Merhout, F., and Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *PNAS*, 37(115):9216–9221.
- Baillon, A., Bleichrodt, H., Huang, Z., and van Loon, R. P. (2017). Measuring ambiguity attitude: (Extended) multiplier preferences for the American and the Dutch population. *Journal of Risk and Uncertainty*, 54(3):269–281.
- Barrios, J. M. and Hochberg, Y. (2020). Risk Perception Through the Lens of Politics in the Time of the COVID-19 Pandemic. Working Paper 27008, NBER.
- Bartik, A. W., Bertrand, M., Cullen, Z. B., Glaeser, E. L., Luca, M., and Stanton, C. T. (2020). How Are Small Businesses Adjusting to COVID-19? Early Evidence from a Survey. Working Paper 26989, NBER.
- Benabou, R. (2015). The Economics of Motivated Beliefs. *Revue d'Economie Politique*, 125(5):665–685.
- Benabou, R. and Tirole, J. (2002). Self-Confidence and Personal Motivation. *Quarterly Journal of Economics*, 1(117):871–915.
- Benabou, R. and Tirole, J. (2016). Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives*, 3(30):141–164.
- Bernhardt, D., Krassa, S., and Polborn, M. (2008). Political polarization and the electoral effects of media bias. *Journal of Public Economics*, 92(5-6):1092–1104.

- Bordalo, P., Coffman, K. B., Gennaioli, N., and Shleifer, A. (2020). Older People are Less Pessimistic about the Health Risks of Covid-19. Working Paper 27494, NBER.
- Boxell, L., Conway, J., Druckman, J. N., and Gentzkow, M. (2020). Affective Polarization Did Not Increase During the Coronavirus Pandemic. Working Paper 28036, NBER.
- Brunnermeier, M. and Parker, J. (2005). Optimal Expectations. *American Economic Review*, 95(4):1092–1118.
- Bursztyjn, L., Rao, A., Roth, C. P., and Yanagizawa-Drott, D. H. (2020). Misinformation during a pandemic. Working Paper 27417, NBER.
- Camerer, C. and Lovallo, D. (1999). Overconfidence and Excess Entry: An Experimental Approach. *American Economic Review*, 89(1):306–318.
- Caplin, A. and Leahy, J. V. (2019). Wishful Thinking. Working Paper 25707, NBER.
- Capozza, F., Haaland, I., Roth, C., and Wohlfart, J. (2021). Studying Information Acquisition in the Field: A Practical Guide and Review. ECONtribute Discussion Papers Series 124.
- Cheng, I.-H. and Hsiaw, A. (2022). Distrust in experts and the origins of disagreement. *Journal of Economic Theory*, 200:105401.
- Chopra, F., Haaland, I., and Roth, C. (2019). Do People Value More Informative News? CESifo Working Paper Series 8026.
- Chopra, F., Haaland, I., and Roth, C. (2022). Do people demand fact-checked news? Evidence from U.S. Democrats. *Journal of Public Economics*, 205:104549.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020a). Does Policy Communication During Covid Work? Working Paper 27384, NBER.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020b). Political Polarization and Expected Economic Outcomes. Working Paper 28044, NBER.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020c). The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending. Working Paper 27141, NBER.
- Cookson, J. A., Engelberg, J. E., and Mullins, W. (2020). Echo Chambers. Socarxiv working paper, Center for Open Science.
- Darley, J. and Gross, P. (1983). A Hypothesis-Confirming Bias in Labeling Effects. *Journal of Personality and Social Psychology*, XLIV:20–33.
- DellaVigna, S. and Kaplan, E. (2007). The Fox News effect: Media Bias and Voting. *Quarterly Journal of Economics*, 122(3):1187–234.
- Dietrich, A. M., Kuester, K., Müller, G. J., and Schoenle, R. S. (2020). News and Uncertainty about COVID-19: Survey Evidence and Short-Run Economic Impact. Working Papers 20-12, Federal Reserve Bank of Cleveland.
- Ditto, P. H., Pizarro, D. A., and Tannenbaum, D. (2009). Motivated Moral Reasoning. *Psychology of Learning and Motivation*, 50:307–338.
- Eil, D. and Rao, J. M. (2011). The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself. *American Economic Journal: Microeconomics*, 3(2):114–38.

- Epley, N. and Gilovich, T. (2016). The Mechanics of Motivated Reasoning. *Journal of Economic Perspectives*, 30(3):133–140.
- Fernández-Villaverde, J. and Jones, C. I. (2020). Macroeconomic Outcomes and COVID-19: A Progress Report. *Brookings Papers on Economic Activity*, 3:111–166.
- Festinger, L. (1962). Cognitive Dissonance. *Scientific American*, 207(4):97–107.
- Fuster, A., Perez-Truglia, R., Wiederholt, M., and Zafar, B. (2020). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *Review of Economics and Statistics*, forthcoming.
- Gentzkow, M. and Shapiro, J. (2010). What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica*, 78(1):35–71.
- Gentzkow, M. and Shapiro, J. M. (2006). Media Bias and Reputation. *Journal of Political Economy*, 114(2):280–316.
- Haaland, I. K., Roth, C., and Wohlfart, J. (2021). Designing Information Provision Experiments. *Journal of Economic Literature*, forthcoming.
- Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press.
- Hanspal, T., Weber, A., and Wohlfart, J. (2020). Income and Wealth Shocks and Expectations during the COVID-19 Pandemic. CESifo Working Paper 8244.
- Jevons, W. (1905). *Essays in Economics*. McMillan.
- Jo, D. (2020). Better the devil you know: Selective exposure alleviates polarization in an online field experiment. Working paper.
- Knobloch-Westerwick, S., Mothes, C., Johnson, B. K., Westerwick, A., and Donsbach, W. (2015). Political online information searching in Germany and the United States: Confirmation bias, source credibility, and attitude impacts. *Journal of Communication*, 65(3):489–511.
- Levy, R. (2021). Social Media, News Consumption, and Polarization: Evidence from a Field Experiment. *American Economic Review*, 111(3):831–70.
- Loewenstein, G. (1987). Anticipation and the Valuation of Delayed Consumption. *Economic Journal*, (97):666–684.
- Lord, C. G., Ross, L., and Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, XXXVII:2098–2109.
- Makridis, C. and Rothwell, J. T. (2020). The Real Cost of Political Polarization: Evidence from the COVID-19 Pandemic. Working paper.
- March, C., Ziegelmeyer, A., Greiner, B., and Cyranek, R. (2016). Pay Few Subjects But Pay Them Well: Cost-Effectiveness of Random Incentive Systems. Working Paper Series No. 5988, CESifo.
- Mobius, M., Niederle, M., Niehaus, P., and Rosenblat, T. (2014). Managing Self-Confidence. Working Paper.
- Mullainathan, S. and Shleifer, A. (2005). The Market for News. *American Economic Review*, 95(4):1031–1053.

- Nyhan, B. and Reifler, J. (2010). When Corrections Fail: The Persistence of Political Misperceptions. *Political Behavior*, 32(2):303–330.
- Oster, E., Shoulson, I., and Dorsey, E. R. (2013). Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease. *American Economic Review*, 103(2):804–830.
- Painter, M. and Qiu, T. (2021). Political beliefs affect compliance with government mandates. *Journal of Economic Behavior & Organization*, 185:688–701.
- Rabin, M. and Schrag, J. (1999). First Impressions Matter: A Model of Confirmatory Bias. *Quarterly Journal of Economics*, 114(1):37–82.
- Saucet, C. and Villeval, M. C. (2019). Motivated memory in dictator games. *Games and Economic Behavior*, 117:250 – 275.
- Simonov, A., Sacher, S. K., Dube, J.-P. H., and Biswas, S. (2021). Frontiers: The Persuasive Effect of Fox News: Noncompliance with Social Distancing During the COVID-19 Pandemic. *Marketing Science*, forthcoming.
- Taber, C. S. and Lodge, M. (2006). Motivated Skepticism in the Evaluation of Political Beliefs. *American Journal of Political Science*, 50(3):755–769.
- Talluri, B. C., Urai, A. E., Tsetsos, K., Usher, M., and Donner, T. H. (2018). Confirmation Bias through Selective Overweighting of Choice-Consistent Evidence. *Current Biology*, 28(19):3128–3135.
- Thaler, M. (2020a). Do People Engage in Motivated Reasoning to Think the World Is a Good Place for Others? Working Paper 2012.01548, arXiv.org.
- Thaler, M. (2020b). The “Fake News” Effect: Experimentally Identifying Motivated Reasoning Using Trust in News. Working Paper 2012.01663, arXiv.org.
- Trivers, R. (2011). *The folly of fools: The logic of deceit and self-deception in human life*. Basic Books.
- Weinstein, N. (1980). Unrealistic Optimism About Future Life Events. *Journal of Personality and Social Psychology*, 5(39):820–839.
- Wiswall, M. and Zafar, B. (2015). Determinants of college major choice: Identification using an information experiment. *The Review of Economic Studies*, 82(2):791–824.
- Zimmermann, F. (2020). The Dynamics of Motivated Beliefs. *American Economic Review*, 110(2):337–61.

Table 1
Determinants of Beliefs (Survey 1)

Table 1 shows regressions where the dependent variable is the sum of Worried Health and Worried Economy (column 1), their difference (col. 2), the principal component of the survey participant's prior beliefs (col. 3) and a 1-5 measure of support for lockdown policies (col. 4). The regressions also include date, region, and article domain order fixed effects. (Coefficients not shown.) All the regressions include day and region fixed effects. ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively. Significance levels are computed using heteroskedasticity-consistent standard errors.

Dep. Var.	Worried H+E	Worried H-E	PC1	Lockdown Support
	(1)	(2)	(3)	(4)
College	0.59***	0.10	-0.02	0.05
Retired	-0.33	0.50**	-0.09	-0.00
Unemployed	0.64	-0.50**	-0.10	0.10
Empl. Loss in HH	1.50***	-0.86***	0.10**	0.11**
Income	-0.08**	0.07***	-0.02***	0.01
Health	-0.46***	-0.05	-0.00	-0.06**
Health Cnds. (0-3)	0.46***	0.39***	0.12***	0.01
Already Infected	1.11**	-0.36	0.39***	-0.08
Liberal	1.87***	0.68***	0.32***	0.70***
Age (\div 100)	-0.97	-0.04	-0.67***	-0.09
Nr Ppl. above 65 in HH	0.11	0.09	0.01	0.10**
Female	0.54***	0.00	0.18***	0.09**
Married	0.59***	-0.12	0.10**	0.01
Black	0.21	-0.07	0.18**	0.14*
Asian	0.36	0.10	0.16*	0.12
Hispanic	1.59***	0.11	0.07	0.18**
Own Primary Residence	0.71***	0.22	0.04	-0.03
Hold Stocks	0.14	-0.41***	-0.05	-0.03
Health Insurance	0.87**	0.60***	0.05	-0.03
Willingness to Take Risks	-0.04	-0.25***	0.01	-0.08***
Willingness to Wait	0.15**	0.11***	0.04***	0.06***
Trust in People	-0.26***	0.06	-0.00	-0.03**
Stay at Home Order	1.05***	-0.28**	0.17***	-0.01
Cases P.C. (\times 1000)	-0.02	-0.02	-0.01**	-0.00
Deaths P.C. (\times 1000)	0.61	0.44	0.27**	0.18
County UR (\div 10)	-0.12	0.02	-0.01	-0.13***
Observations	2,440	2,440	2,440	2,440
Adj. R ²	0.12	0.11	0.12	0.16
Mean of Dep. Var.	12.64	0.58	-0.00	3.31
St. Dev. of Dep. Var.	5.06	2.99	1.00	1.08

Table 2
Domain Choice

Table 2 shows regressions where the dependent variable is a dummy equal to 1 if the survey participant chooses to see a headline related to the health consequences of the Covid-pandemic. Worried H+E is the sum of Worried Health and Worried Economy, Worried H-E is their difference, PC1 is the principal component of the survey participant’s prior beliefs, and lockdown support is a a 1 to 5 measure of support for lockdown policies. Control variables (not shown) are included in columns 5 through 8 and are as in Table 1. All the coefficients are multiplied by 100. Heteroskedasticity-consistent standard errors are reported in parentheses. ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Worried H+E	2.01*** (0.18)			1.11*** (0.20)	1.45*** (0.20)			0.94*** (0.21)
Worried H-E	3.91*** (0.30)			2.90*** (0.32)	3.47*** (0.32)			2.77*** (0.33)
PC1		8.17*** (0.97)		3.88*** (0.98)		5.05*** (1.00)		2.10** (0.99)
Lockdown Support			13.56*** (0.82)	8.59*** (0.93)			10.40*** (0.91)	6.65*** (0.97)
Observations	2,440	2,440	2,440	2,440	2,440	2,440	2,440	2,440
Adj. R ²	0.10	0.03	0.09	0.14	0.15	0.10	0.13	0.17
Mean of Dep. Var.	64.63	64.63	64.63	64.63	64.63	64.63	64.63	64.63
St. Dev. of Dep. Var.	47.82	47.82	47.82	47.82	47.82	47.82	47.82	47.82
Controls	NO	NO	NO	NO	YES	YES	YES	YES

Table 3
Pessimistic Article Choice - Survey 1

Table 3 presents regressions where the dependent variable is a dummy equal to 1 if the participant chooses the article with the pessimistic headline in the health domain (columns 1-4) or in the economy domain (columns 5-8). Worried H+E is the sum of Worried Health and Worried Economy, Worried H-E is their difference, PC1 is the principal component of the survey participant's prior beliefs, and lockdown support is a 1 to 5 measure of support for lockdown policies. Control variables (not shown) are as in Table 1. All the coefficients are multiplied by 100. Heteroskedasticity-consistent standard errors are reported in parentheses. ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively.

<i>Dep. Var.:</i>	<i>Pessimistic Article Choice in</i>							
	<i>Health Domain</i>				<i>Economy Domain</i>			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Worried H+E	1.14*** (0.20)			0.32 (0.21)	0.96*** (0.20)			0.14 (0.22)
Worried H-E	2.15*** (0.32)			1.07*** (0.33)	0.89*** (0.34)			0.07 (0.35)
PC1		6.45*** (1.04)		4.01*** (1.06)		8.76*** (1.06)		7.31*** (1.11)
Lockdown Support			11.98*** (0.89)	10.07*** (1.02)			8.25*** (0.97)	6.82*** (1.05)
Observations	2,440	2,440	2,440	2,440	2,440	2,440	2,440	2,440
Adj. R ²	0.10	0.09	0.13	0.14	0.10	0.11	0.11	0.13
Mean of Dep. Var.	44.63	44.63	44.63	44.63	47.75	47.75	47.75	47.75
St. Dev. of Dep. Var.	49.72	49.72	49.72	49.72	49.96	49.96	49.96	49.96
Controls	YES	YES	YES	YES	YES	YES	YES	YES

Table 4
Belief Revision (Survey 1)

Table 4 shows regressions where the dependent variable is the revision in the probability that the unemployment rate will be higher than 20% (columns 1 and 4), that the number of COVID-19-related deaths will be higher than 200,000 (columns 2 and 5), and that the respondent will be infected (columns 3 and 6). Preferred is a dummy equal to 1 if the individual reads the preferred article; Pessimistic is a dummy equal to 1 if the individual reads the pessimistic article. Control variables (not shown) are included in all the regressions and are as in Table 1. They also include the participant’s prior beliefs, together with dummies for corner choices (0 and 100). Heteroskedasticity-consistent standard errors are reported in parentheses. ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively.

Dep. Var.	Δ Prob. UR>20%	Δ Prob. Deaths>200k	Δ Prob. Infected	Δ Prob. UR>20%	Δ Prob. Deaths>200k	Δ Prob. Infected
	(1)	(2)	(3)	(4)	(5)	(6)
Preferred	-2.39* (1.26)	-2.38* (1.35)	-1.39 (1.16)	-4.41** (2.09)	-4.20** (1.89)	-1.54 (1.22)
Pessimistic	-0.03 (1.24)	-0.71 (1.32)	-1.63 (1.02)	-0.74 (2.17)	-4.78*** (1.80)	-0.89 (1.32)
Pess. \times Pref.	4.55** (1.83)	4.21** (1.95)	1.48 (1.59)	11.30*** (3.12)	7.70*** (2.88)	2.67 (1.86)
Observations	1,577	1,577	1,577	863	863	863
Adj. R ²	0.09	0.13	0.12	0.21	0.12	0.09
Mean of Dep. Var.	-1.40	-2.36	-0.04	-2.90	-2.61	-1.57
St. Dev. of Dep. Var.	17.98	19.52	15.98	20.89	18.36	12.53
Controls	YES	YES	YES	YES	YES	YES
Domain	Health	Health	Health	Economy	Economy	Economy

Table 5
Article Assessment (Survey 1)

Table 5 shows regressions where the dependent variables are: a 1 to 7 assessment of the informativeness of the article (column 1); a 1 to 7 assessment of the reliability of the article (column 2); a dummy equal to 100 if the participant correctly identifies, among four headlines, the one corresponding to the article they have just read (column 3); and the dollar amount spent to advertise the article source (column 4). Preferred is a dummy equal to 1 if the participant was allowed to read the article she chose. Pessimistic is a dummy equal to 1 if the participant read the article with the pessimistic article. Control variables (not shown) are included in all the regressions and are as in Table 1. They also include dummies corresponding to the article chosen. Heteroskedasticity-consistent standard errors are reported in parentheses. ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively.

Dep. Var.	Inform.	Reliability	Headline Check	Money
	(1)	(2)	(3)	(4)
Preferred	0.39*** (0.06)	0.58*** (0.06)	4.74*** (1.48)	4.51*** (0.76)
Observations	2,440	2,440	2,440	2,440
Adj. R ²	0.08	0.10	0.11	0.04
Mean of Dep. Var.	5.06	4.79	82.62	26.50
St. Dev. of Dep. Var.	1.42	1.47	37.90	18.98
Article Dummies	YES	YES	YES	YES
Controls	YES	YES	YES	YES

Table 6
Support and Donation to WHO (Survey 1)

Table 6 shows regressions where the dependent variables are: a measure of support for lockdown policies on a 1 to 5 scale (columns 1 and 3) or the amount donated to the WHO as opposed to FreedomWorks (columns 2 and 4). Columns 1 and 2 include results for survey participants who chose the health article, columns 3 and 4 include participants who chose the economy article. Pessimistic is a dummy equal to 1 if the participant read the article with the pessimistic headline. Preferred is a dummy equal to 1 if the article read was the preferred one. Control variables (not shown) are included in all the regressions and are as in Table 1. They also include prior beliefs with regard to the probability that the unemployment rate will be higher than 20%, that the number of COVID-19-related deaths will be higher than 200,000, and with regard to the probability of being infected. Finally, the regressions control for dummies for support for current lockdown policies in the area where the participant lives. Heteroskedasticity-consistent standard errors are reported in parentheses. ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively.

Dep. Var.	Support	Money to WHO	Support	Money to WHO
	(1)	(2)	(3)	(4)
Preferred	-0.34*** (0.07)	-1.19*** (0.45)	-0.24** (0.11)	-1.74** (0.70)
Pessimistic	-0.18** (0.07)	-0.67 (0.44)	-0.29** (0.11)	-1.53** (0.73)
Pess. × Pref.	0.56*** (0.11)	1.95*** (0.65)	0.46*** (0.17)	3.45*** (1.07)
Observations	1,577	1,577	863	863
Adj. R ²	0.28	0.26	0.45	0.39
Mean of Dep. Var.	3.61	13.45	2.69	8.17
St. Dev. of Dep. Var.	1.18	7.14	1.39	8.19
Controls	YES	YES	YES	YES
Domain	Health	Health	Econ.	Econ.

Table 7
Pessimistic Article Choice - Effect of Source Revelation - Survey 2

Table 7 presents regressions where the dependent variable is a dummy equal to 1 if the participant chooses the article with the pessimistic headline. Worried H+E is the sum of Worried Health and Worried Economy, Worried H-E is their difference, PC1 is the first principal component of the survey participant's prior beliefs, Lockdown Support is a 1 to 5 measure of support for lockdown policies, Revealed Before is a dummy equal to 1 if the participant is shown the article source before making the choice, and Liberal is a dummy equal to 1 if the participants defines herself as being "very liberal/democrat," "liberal/democrat," or "leaning liberal/democrat," and zero otherwise. Control variables (not shown) are as in Table 1. Net effect is the sum of the coefficients on Worried (columns 1 and 4), PC1 (columns 2 and 5), and Lockdown Support (columns 3 and 6) and the coefficients on their interactions with the Revealed Before dummy. All the coefficients are multiplied by 100. Heteroskedasticity-consistent standard errors are reported in parentheses. ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Worried H+E	0.67** (0.30)			0.51* (0.30)		
Worried H+E × Rev. Before	-1.19** (0.48)			-0.81* (0.49)		
PC1		5.36*** (1.58)			4.72*** (1.59)	
PC1 × Source Revealed		-4.16* (2.47)			-1.92 (2.43)	
Lockdown Support			5.71*** (1.43)			4.76*** (1.48)
Lockdown S. × Rev. Before			-7.52*** (2.42)			-3.49 (2.58)
Revealed Before	8.85 (6.45)	-5.84** (2.54)	19.33** (8.42)	15.18** (6.55)	5.71* (3.44)	16.74** (8.42)
Liberal × Rev. Before				-25.42*** (5.10)	-26.57*** (5.07)	-25.10*** (5.36)
Liberal				8.96*** (3.15)	8.48*** (3.15)	7.47** (3.19)
Net Effect	-0.53	1.21	-1.81	-0.30	2.80	1.27
Std. Err.	0.40	1.97	2.01	0.40	1.91	2.14
Observations	1,572	1,571	1,572	1,572	1,571	1,572
R ²	0.05	0.06	0.06	0.07	0.07	0.07
Mean of Dep. Var.	36.96	36.96	36.96	36.96	36.96	36.96
Controls	YES	YES	YES	YES	YES	YES

Figures

Figure 1
Headline Choice - Survey 2

The histograms plot the fraction of participants who chose the pessimistic headline in Survey 2, among individuals who knew (light bars) or did not know (dark bars) the article source before making the choice. The bins correspond to quintiles of the variable Worried H+E (Panel a), PC1 (Panel b), and levels of support for lockdown policies (Panel c).

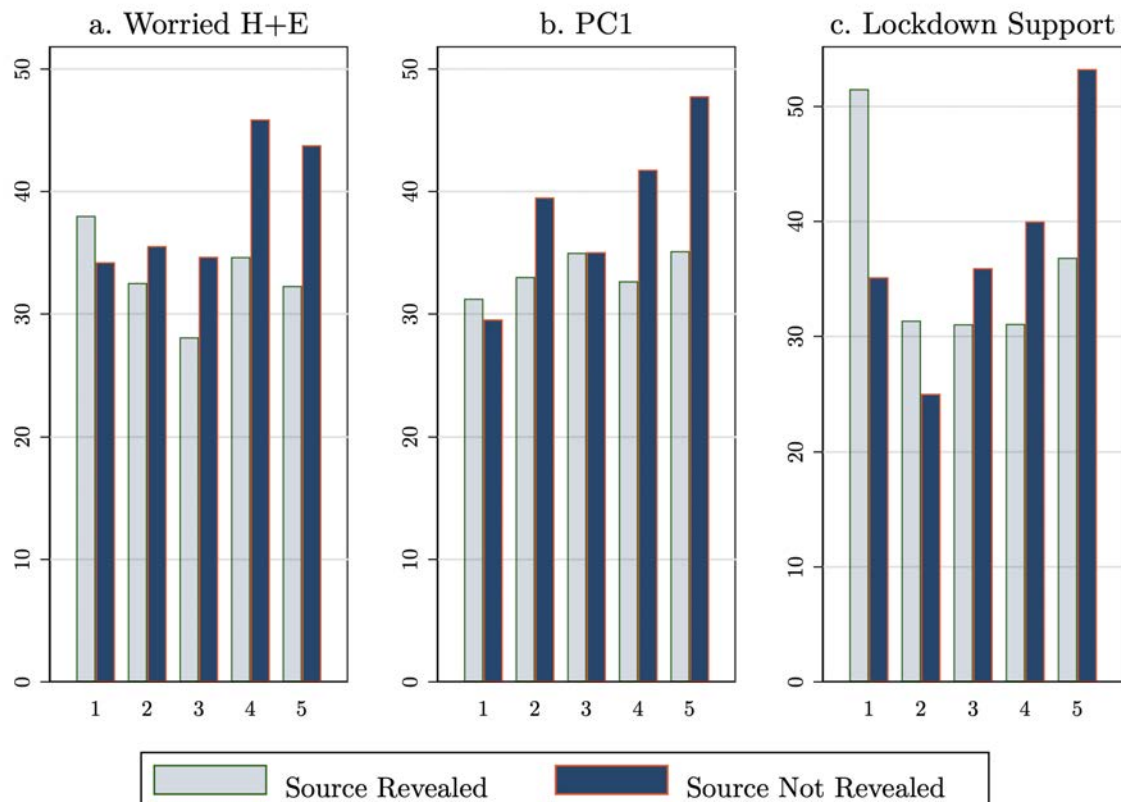


Figure 2
Article Assessment - Survey 2

Figure 2 shows coefficients on the triple interaction term (Liberal \times Revealed After \times Pessimistic) from equation 3. The dependent variable for each regression is indicated at the top of each panel. We compute, for each participant, a proxy capturing how surprised individuals might be to realize that the pessimistic article comes from Fox News, given by the perceived probability that the article chosen comes from a liberal source minus the perceived probability that the article chosen comes from a conservative source. Individuals are then sorted into two halves, depending on this measure. Coefficients for the full sample are shown as circles. Coefficients for individuals in the bottom half and top half are shown as squares and diamonds, respectively. The vertical lines represent 95% confidence intervals.

