

Ester Faia | Andreas Fuster | Vincenzo Pezone | Basit Zafar

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

SAFE Working Paper No. 307

Leibniz Institute for Financial Research SAFE
Sustainable Architecture for Finance in Europe

info@safe-frankfurt.de | www.safe-frankfurt.de

Electronic copy available at: <https://ssrn.com/abstract=3783215>

Non-Technical Summary

How people form subjective beliefs in situations of high uncertainty is a fundamental question, with far-reaching consequences for economic and social decisions. This is especially the case during 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 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 channeled through sources that are generally aligned with, or opposed to, the respondent's political orientation.

In the first survey, we present different article headlines to participants and ask them which article they would like to read. The news articles differ with respect to the domain (economic versus health) and the tone (optimistic versus pessimistic). Our results show, consistently with motivated beliefs theories in general, that agents' article choices tends to reflect their prior. Specifically, individuals with more pessimistic priors are substantially more likely to select an article with a more negative headline.

We then randomize whether participants are allowed to read their preferred article or not, and find that respondents rate a given article as more informative and are able to recall its headline more accurately (our proxy for attention) if they had chosen it. Moreover, 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 if the article was the one they originally selected. 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, an observation that motivate the design of our second survey.

In the second survey we provide direct evidence on polarization due to views about different media sources. After soliciting the same demographics and prior beliefs as in the first survey, we ask respondents to select among 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 choice of article 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.

Our results are striking. We find that the revelation of the news sources prior to the selection of 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. 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. While source dependence has large effects on people's selection and assessment of information, we find that impacts on posterior beliefs are muted.

This evidence suggests that self-selection of news may enhance existing disagreements, since our respondents seem to be subject to confirmation bias. 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 the society with opposing views may be more effective when it occurs through sources these audiences are more familiar with.

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

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

February 9, 2021

Abstract

How people form beliefs is crucial for understanding decision-making under uncertainty. This is particularly true in a situation such as a pandemic, where beliefs will affect behaviors that impact public health as well as the aggregate economy. We conduct two survey experiments to shed light on potential biases in belief formation, focusing in particular on the tone of information people choose to consume and how they incorporate this information into their beliefs. In the first experiment, people express their preferences over pandemic-related articles with optimistic and pessimistic headlines, and are then randomly shown one of the articles. We find that respondents with more pessimistic prior beliefs about the pandemic are substantially more likely to prefer pessimistic articles, which we interpret as evidence of confirmation bias. In line with this, respondents assigned to the less preferred article rate it as less reliable and informative (relative to those who prefer it); they also discount information from the article when it is less preferred. We further find that these motivated beliefs end up impacting incentivized behavior. In a second experiment, we study how partisan views interact with information selection and processing. We find strong evidence of source dependence: revealing the news source further distorts information acquisition and processing, eliminating the role of prior beliefs 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, Michael Thaler, Michael Weber, Johannes Wohlfart 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. The views expressed are those of the authors and do not necessarily reflect those of the Swiss National Bank.

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

[‡]Swiss National Bank and CEPR. Email: andreas.fuster@gmail.com

[§]Goethe University Frankfurt. Email: pezone@safe.uni-frankfurt.de

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

1. Introduction

How people form subjective beliefs in situations of high uncertainty is a fundamental question, with far-reaching consequences for economic and social decisions. This is especially the case during 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 give higher weight to more 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, 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 is 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.

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 their financial situations. They are then informed that they can read a Covid-related article, and 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.¹ 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.

1 In an earlier “pilot” survey, we checked that respondents on average found the headlines equally informative and interesting. See Appendix A.6.

In the first survey, the news articles differ with respect to the domain (economic versus health) and the tone (optimistic versus pessimistic). Our results show that agents' choice of article tends to reflect their prior, independent of whether the latter was on the optimistic or pessimistic side. 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 result thus suggests that Covid-related thinking is not primarily driven by a desire to feel better about oneself (e.g. one's personal health risks).

Moving to information processing, we see that respondents rate a given article as more informative and are able to recall its headline more accurately (our 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 patterns at the individual level may be consistent with rational information processing: if an agent prefers a headline because they are ex-ante uninformed about it,² 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. Moreover, we provide participants with a budget that they can use to advertise the source of the article they are assigned to read, and show that they spend substantially more money when receiving their preferred article. Both results are 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 if the article was the one they originally selected.³

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 effects and health effects as complementary.⁴ 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

2 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.

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

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

WHO and Freedom Works). 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. 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.⁵

In the second survey we provide direct evidence on polarization due to views about different media sources. After soliciting the same demographics and prior beliefs as in the first survey, we ask respondents to select among 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 choice of article 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., 2020). 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 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.⁶ Regarding information processing,

5 For a different view, see Boxell et al. (2020), who find that affective polarization did not increase during the pandemic.

6 This may appear to contrast with Chopra et al. (2019) who find that individuals demand news

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 policy-makers can reach out to the “other side” by sending signals through sources traditionally aligned with that group.

Our paper is related to the literature on the motivated selection and processing of information. This departure from rational expectations, which has received much attention in psychological research, encompasses several variants.⁷ Motivated beliefs 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.

A distinguishing characteristic of motivated beliefs is asymmetric updating: agents overweight signals that confirm their priors or their previous choices and 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), Hoelzl and Rustichini (2005), Mobius et al. (2014), 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

consistent with their pre-existing beliefs. However, note that in our context, by design, the news article that is aligned with one’s political ideology is less likely to be aligned with one’s priors.

7 Weinstein (1980), Lord et al. (1979), Darley and Gross (1983) and, more recently, Ditto et al. (2009) and Talluri et al. (2018) are among the numerous experiments in the psychology literature. Epley and Gilovich (2016) provide a compact introduction.

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, as manifested by phenomena such as group-thinking, political partisanship, or media polarization. On the role of partisanship, Alesina et al. (2020) and Coibion et al. (2020b) find that polarization affects individuals' perceptions not only about policy options, but also about reality. Past literature has included several experiments on group-thinking more generally. Related to a large area of research associated with media and political polarization, our experiment also includes such an extension through the second survey, which aims at testing the effects of media polarization.⁸ The role of media during an event as dramatic as a pandemic is crucial, with anxiety leading many people to the continuous search of information. In the US, opinions on the severity and policy response of the pandemic were also largely polarized from the start. We find clear evidence of source dependence as our respondents rate the same information differently, depending on which source it comes from.

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 role of the media. On the other hand, such a context with a wide variety 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, biased beliefs and wishful thinking have been shown to play a role in health behaviors (Oster et al., 2013), and it has been pointed out that the neglect of contradictory evidence is particularly visible in situations of health stress (Benabou,

8 Studies on media bias which are closer to the spirit of our experiment include DellaVigna and Kaplan (2007), Gentzkow and Shapiro (2006, 2010), and Mullainathan et al. (2008). A study implementing an information treatment experiment close in spirit to our second survey is Bail et al. (2018).

2015). Among the many survey studies on the Covid outbreak,⁹ 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., 2020). However, few studies have attempted to endogenize the process of information selection (see Fuster et al., 2020, for a related discussion), 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

Our data come from two surveys that we designed. The surveys were administered on Qualtrics, and were directly advertised to their US-based panelists. Details of the two surveys (including question wording) are in Appendices A.7 and A.8. Figures 1 and 2 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 economic consequences of Covid for them and their family/friends, we asked them for point forecasts of the unemployment rate and the number of Covid-related deaths in the

9 See e.g. Coibion et al. (2020a) on employment expectations, Hanspal et al. (2020) on households' debt, working hours and retirement age, Dietrich et al. (2020) on uncertainty, Bartik et al. (2020) on firms' expectations, Adams-Prassl et al. (2020) on inequality, Bu et al. (2020) on risk-taking, Fetzer et al. (2020) on contagion estimates and misinformation, or Caria et al. (2020) on global preferences. Surveys on Covid are being conducted by scholars in other fields as well: Raude et al. (2020) measure optimism or pessimism about infection risk; Brooks et al. (2020) summarize psychology survey studies conducted during Covid. Finally, a recent survey by Alsan et al. (2020) examines how the policy response to the Covid crisis was related to the preference for civil liberties across various countries.

10 Note that 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.socialsciscience.org/trials/5850>).

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 (on a 5-point scale) for stay-home orders and social distancing measures implemented in the US.

Next, respondents were told that before they reported their views on the consequences of the pandemic, 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 in this domain. 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. We also used a qualitative question to measure their support for stay-home orders in a case of a resurgence of the pandemic. The survey furthermore featured incentivized choices designed to confirm the reliability of 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 allocate a donation (by us) of \$20 between the World Health Organization (WHO), an organization that explicitly supported

11 These articles were actual news articles that were slightly edited and condensed such that they had similar length. Some of the headlines were also edited by the authors. 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. See Appendix A.6 for further details.

lockdown measures, and FreedomWorks, a conservative and libertarian advocacy group whose favored policies are opposite to the WHO's.¹²

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 by 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 (see Figure 2). 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 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. 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. Respondents were also given the opportunity to decide how much, out of \$50, they wanted to spend on advertising the news source on Facebook.

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,

¹² These choices were incentivized in that respondents were informed that for the news source promotion, two of them, and for the money donation, ten of them, would be randomly picked and their choices be implemented (which we did after data collection was complete).

and whether the selection of information (domain and tone) is motivated. For example, wishful thinking would imply that respondents worried about the health impacts of the pandemic may avoid the negative health-related article. 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 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). Comparison of T1 with T2 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 between demographics and priors, and assess the determinants of the article domain and tone choice. This provides evidence on whether and 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, time spent on reading the article, attentiveness (which is measured by whether the respondent failed an attention check), and actual behavior, as proxied by the choice to donate to one of the proposed organizations and to sponsor the news organization that has published the article.

2.2.1. 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. We follow Baillon et al. (2017) and 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 (optimistic) individuals tend to choose signals that assign higher (lower) 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.

2.3. Sample Description

As is common in online surveys that are run on platforms like Qualtrics, the survey included two multiple-choice comprehension questions in the beginning, which a respondent had to answer correctly to proceed.¹³ 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. Our final sample for Survey 1 includes a total of 2,440 US-based participants.¹⁴ 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).

Table 1 shows the characteristics of our sample. The variables used in this table are described in Table A.1. Respondents who had participated in Survey 1 were not invited for the second survey. The table shows that the characteristics of respondents between the two surveys are slightly different. We attribute this to sampling error. Looking at the sample for Survey 1, 54% 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 is much higher educated; this is at least partly due to differential internet access and computer literacy across education groups.¹⁵

13 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.

14 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 in a time more 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.

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

The pandemic has had far-reaching impacts. 29% of the respondents report that someone in their household has 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% of the respondents, at the time of the first survey, were subject to stay-at-home orders in their area; the proportion falls to 19% for participants in the second survey.

We now proceed with the analysis. Results are described progressively first for Survey 1 and then for Survey 2.

3. Empirical Analysis - Survey 1

3.1. Baseline Concerns about the Pandemic

At the baseline, 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 economic consequences of the pandemic. We refer to these variables as “Worried Health” and “Worried Economy”, respectively. The average response, as shown in Table 1, is 6.61 for Worried Health, and 6.03 for Worried Economy. The large standard deviations of almost 3 points in both measures indicates 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 “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

education 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.

allows us to explore the determinants of worries related to the pandemic, and to verify that survey responses follow sensible patterns. Results are shown in Table 2. 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.¹⁶ 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 A2 reports correlations between these seven variables and PC1, and shows that this proxy does well in summarizing information about participants’ beliefs.

The last column of Table 2 uses Lockdown Support as the dependent variable. As Appendix-Table A3 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.¹⁷

Table 2 shows many intuitive correlations. For instance, respondents with underlying 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 worry relatively more 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 less worrying about health and less support for lockdowns, while higher patience is associated with more worries (esp. about health) and stronger lockdown support. Stronger trust in

16 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.

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

others reduces worries and lockdown support. In terms of other personal characteristics, female and married respondents tend to worry more, are more pessimistic and supportive of lockdowns. Perhaps surprisingly, age negatively correlates with being worried and Covid-pessimism; this result is consistent with Bordalo et al. (2020) who find that perceived personal health risks associated with Covid fall with age.¹⁸ 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.¹⁹

While Table 2 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 variables that are significant in Table 2 are also significant in the equivalent analysis for Survey 2, shown in Appendix-Table A7. 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.

3.2. Domain Choice

We next assess whether respondents choose the article based on their proximity 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 dictated by several elements, which include respondents' worries in the specific domain, or their area of expertise.

Table 3 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,

¹⁸ The bivariate correlations of age and the four dependent variables are also negative, except for Worried H-E, where it is positive (0.16).

¹⁹ 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).

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. This reflects the fact that respondents who thought measures were too strict (Lockdown Support of 1 or 2) only have a probability of preferring the health article of 36%, while for those who think measures should have been stricter (Lockdown Support of 4 or 5), the corresponding share is 79%. 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 2. 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 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 through 4 of Table 4 serve to identify the predictors of choosing the pessimistic article in the health domain, while columns 5 through 8 identify those in the economy domain. The regressors, together with controls, are the same as in Table 3. 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.²⁰ The larger impact

²⁰ In unreported tests, we find that the difference between the coefficients estimated in the two domains is statistically significant at the 5% level.

of the Worried H-E variable 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.

To visually appreciate the results, Figure 3 plots the fraction of participants who chose the pessimistic headline in the health domain (green bars) and in the economy domain (red bars) for each quintile of Worried H+E (Panel a), Worried H-E (Panel b), PC1 (Panel c), and for each level of support for lockdown measures (Panel d). A distinctive monotonic 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”. 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 4). However, plausibly respondents 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 4 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. To address this point and show a plausible *causal* link from Covid-worries/-pessimism to article tone preference, we instrument for the covariates of interest using “pre-determined/exogenous” factors that affect worries and pessimism, namely whether the respondent’s household had experienced a job loss since the onset of Covid, and the number of pre-existing health conditions. Both of these variables were found in Table 2 to strongly affect worries and prior beliefs (though with less of an effect on lockdown support). Also, intuitively, having experienced an employment loss has stronger effects on worries about economic consequences, while pre-existing health conditions have stronger effects on health worries. Furthermore, both

variables are “objective” measures that should not be affected by a respondent’s general optimism/pessimism.

Results are shown in Table 5. The proposed instruments do provide meaningful variation as shown by the first-stage Kleibergen-Paap Wald F -statistics. However, they are not very high and around the rule-of-thumb value of 10 (except for Lockdown Support, where the first stage is weaker). Our results are very similar if we instead use a limited information maximum likelihood (LIML) estimator, which is more robust in the presence of weak instruments.²¹ We see that the instrumented covariates are generally statistically and economically significant. Exceptions are the imprecise estimate in column 3, and the effects of Worried H-E.²² Coefficient magnitudes are larger than in the corresponding OLS regressions in Table 4, which could reflect measurement error in the worry measures and prior beliefs, or the fact that the IV provides a LATE.

In sum, we can conclude that there is a plausibly causal effect from stronger baseline worries and pessimism on the 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 A4 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.

²¹ We further note that for the 2SLS regressions where PC1 is instrumented, the effective F -statistic of Montiel Olea and Pflueger (2013) of 8.43 exceeds the critical values at which the null of weak instruments can be rejected at the 1% confidence level, for the most conservative threshold τ (5%) reported according to the Stata routine of Pflueger and Wang (2015).

²² Interestingly, Worried H-E enters positively in column 1, meaning those more worried about health more strongly prefer the pessimistic health headline, but negatively in column 4, meaning those more worried about the economy prefer the pessimistic economy headline. However, both effects are imprecisely estimated.

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

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

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, including a constant. 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 6. 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.²³ In columns 1-3, we include the 1,577 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. 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 economically meaningful, 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. 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.

Figure 4 displays average revisions for the three dependent variables for each of the four group of participants (article choice \times article assignment), for each of the domains.²⁴

²³ We disregard the 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.

²⁴ The graphical evidence does not perfectly map to the estimates of Table 6 since regressions include the prior beliefs as well as the control variables.

Looking at the revisions for beliefs regarding the number of deaths in the Health domain, we see that all four groups revise their beliefs downwards. The tone of the article clearly matters: average downward revisions are larger for the groups that are assigned the optimistic article (dark green and red bars). However, consistent with confirmatory bias, among respondents who are assigned the optimistic article, average revisions are smaller for those who had preferred the pessimistic article (red bar).

3.4.2. Assessment of Article

The survey included questions on whether respondents found the article that they had read informative, and whether they thought the information it contained was 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.²⁵ 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 revealed in Survey 1. Specifically, they were provided with a \$50 budget (at no cost to themselves) and asked how much of it they wanted to spend on Facebook ads for the news outlet whose article they had read.²⁶

In Table 7, 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 are 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.

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

25 Respondents were presented with the list of the four article 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 the participants passed this “attention check”.

26 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.

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). However, such an explanation is unlikely to have an impact on the assessed reliability of the article—but 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 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 relevant 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) on a 1 to 5 scale. 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.²⁷

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 (as participants were informed that 10 respondents will be picked at random and their choices will be implemented).²⁸ 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 8. In the first two columns results are reported for the subsample of individuals who selected the Health domain, while 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 donating

²⁷ 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."*

²⁸ As before, we provided participants with a link where they could see screenshots of the donation receipts.

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 in the measure), and the donation to WHO rises by \$1.95-\$3.45.

4. Empirical Analysis - Survey 2

The analysis in the previous section shows that individuals have a tendency to put more weight on signals that they find more attractive. Those distortions may often be politically driven. In fact, we see suggestive evidence of this in Survey 1. Table 2 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 3 and 4). 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 A5).²⁹

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, New York Times or Fox News differentially informative.³⁰ 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, even more so if the source-dependence were to affect updating differently or against prior beliefs. 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; the visual depiction of the design is presented in Figure 2. Recall that the survey randomly divided participants into three groups: the first one is never revealed the news source (T1);

²⁹ 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.

³⁰ See, for example, Thaler (2020b), Taber and Lodge (2006), Kahan et al. (2012), Nyhan and Reifler (2010), or Nyhan et al. (2014).

the second one is revealed the source after having chosen the article but before reporting the posterior beliefs (T2), and the third one is revealed the source before making the article choice (T3).³¹

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.³² 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

We first investigate which variables are correlated with the choice of the article tone. The dependent variable in Table 9 is a dummy that equals 1 if the participant chose the pessimistic headline. The regressors are as in Table 4, where we conducted a similar analysis for Survey 1. To make results comparable across surveys, participants in group T3 are excluded here, so that we include only participants who chose the article without knowing its source. We see that Worried H+E, PC1, and lockdown support are positively correlated with the choice of the pessimistic article. Overall, despite the smaller sample size, results are qualitatively similar to those obtained when analyzing Survey 1. The only difference is that now Worried H-E does not appear to be a significant predictor of the headline choice.³³

To visualize the impact of the predictors for tone choice, Figure 5 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

31 Appendix-Table A6 shows that 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.

32 One could also consider randomizing the source while holding article headline and content fixed. As we did not want to use deception in our experiment (and therefore use actual articles from these sources) we did not pursue this path.

33 A possible explanation for this is that now the headlines were not explicitly assigned to either the health or economy domain; arguably 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*” both had implications in both domains.

distribution choose the pessimistic article, 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; this variable equals 1 for participants assigned to group T3. *Z* is one of the three predictors that, in Table 9, have displayed a significant forecasting ability for the likelihood of choosing the pessimistic article: Worried H+E, PC1, and Lockdown Support. Finally, *X* is the usual vector of control variables (including a constant). The parameter α in this specification shows the impact of the covariate *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 10. 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 5.40 and highly statistically significant, in line with the results from Table 9. However, once the news source is revealed, the impact of the standalone PC1 variable is only 1.27 ($5.40 - 4.13$), and not statistically different from zero. This can also be seen from the “Net Effects” shown at the bottom of the table; they are all small and statistically insignificant in the first three columns. 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 green bars in Figure 5. While the red bars in the figure show that there is positive relationship between the fraction of participants who chose the pessimistic headline and prior beliefs (quintiles of the three variable distributions), 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) as follows:

$$\begin{aligned}
 Pessimistic_i = & \alpha Z_i + \beta Z_i \times Revealed\ Before_i + \gamma Revealed\ Before_i + \\
 & \theta Liberal_i + \phi Liberal_i \times Revealed\ Before_i + \delta' X_i + \varepsilon_i.
 \end{aligned}
 \tag{3}$$

Here, *Liberal* is a dummy equal to 1 if the respondent defines herself as “very liberal/democrat”, “liberal/democrat”, or “leaning liberal/democrat”.³⁴ This variable is included as a standalone to account for the fact that, ceteris paribus, political beliefs might affect the article choice. The coefficient ϕ on the interaction term *Liberal* \times *Revealed Before* is the one that is of interest. It measures the extent to which, upon learning that the source of the pessimistic article is Fox News, liberals are either less ($\phi < 0$) or more likely to choose it.

Results for this specification are shown in columns 4-6 of Table 10. The coefficient θ is positive, ranging between 7.6 and 9.1, 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.³⁵ Thus, this is further confirmation of 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. Remember 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 ϕ 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 $\theta + \phi$ is

³⁴ 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.

³⁵ 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 (2020) and Barrios and Hochberg (2020).

negative and economically large. It varies between -16.3 and -18, and is significant at the 1% level in all the specifications.

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, has any effect on article assessment. We focus on the subsample that includes all respondents who choose the article without knowing its source (T1 and T2), half of whom are made aware of it after having made the choice but before reporting posterior beliefs and attitudes (T2). We estimate the following model:

$$\begin{aligned}
 y_i = & \alpha Pessimistic_i + \beta Revealed After_i + \gamma Liberal_i + \\
 & \theta Pessimistic_i \times Revealed After_i + \phi Revealed After_i \times Liberal_i + \\
 & \rho Pessimistic_i \times Liberal_i + \mu Pessimistic_i \times Revealed After_i \times Liberal_i + \\
 & \delta' X_i + \varepsilon_i.
 \end{aligned} \tag{4}$$

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 (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 reported in green, with its respective 95% confidence intervals, in Figure 6.³⁶ 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.

³⁶ For completeness, Appendix Table A9 displays the regression coefficients.

Note that the source revelation should have a muted effect on participants who were not particularly “surprised” to learn that the article with the positive tone came from the New York Times (or that the pessimistic article came from Fox News). Instead, the effect of revealing the news source should be concentrated among those 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 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 first compute the perceived probability that the article comes from a liberal news source as the sum of the percent chance that the individual assigns to the article being from CNN, The New York Times, MSNBC, and The Washington Post. Similarly, the perceived probability that the article comes from a conservative news source is the sum of the percent chance that the individual assigns to the article being from Breitbart, Fox News, and The Wall Street Journal. Once the latter is subtracted from the former, a proxy is obtained 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. Our hypothesis is that individuals in the above-median subsample should be more surprised by the revelation of the information source; hence, the source revelation treatment should have a stronger effect. Note that the mean probability assigned to the pessimistic headline coming from liberal news sources in the sample was 57.5%, and the mean probability assigned to it coming from conservative news sources was 22.7%. The respective figures for the optimistic article are 44.7% and 34.2%. Hence, on average participants assign a higher probability to the article coming from a liberal news source (partly mechanically, given that we have four liberal and three conservative sources), but much more so for the pessimistic article. Thus, the average respondent should be disproportionately more surprised by the news that a pessimistic article came from Fox News. Appendix-Table A8 shows these statistics for both Survey 2 and, for completeness, Survey 1, where similar patterns are observed.³⁷

The specification in equation (4) is estimated for each of the two groups separately. Figure 6 plots the coefficient, μ , on the triple interaction term, with confidence intervals, for individuals in the bottom of the distribution (“Low Surprise” group) and in the top of the distribution (“High Surprise” group). The former are indicated with an orange bar and the latter with a blue bar. Results are exactly as one would have 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, a negative

³⁷ The table also shows these statistics for just NYTimes and Fox News, where substantial differences in probabilities are also observed.

and (at least marginally) statistically significant coefficient is found 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. Table 11 reports estimates of equation (4), 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.³⁸ Finally, like in Survey 1, we again elicit (5) support of future lockdown policies in case of a new wave of the pandemic. As before, the parameter of interest is μ , the parameter on the triple interaction term *Pessimistic* \times *Revealed After* \times *Liberal*.

Source revelation has little impact on belief revision. For example, liberals who read the pessimistic article and learn that the source is Fox News do not discount the information, in the sense that their mean revisions are very similar to those of their counterpart liberals who read the same article but are not informed of the article source (in fact, in column 1, they revise their beliefs upward by an additional 1.8 percentage points, although this estimate is quite noisy). 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. This is indeed surprising: Figure 6 showed that source revelation led liberals to rate the pessimistic article as less reliable and informative.

One possibility 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. The opposite holds for conservatives learning that their chosen optimistic article was actually from a liberal source, the New York Times. Learning that even a supposedly alarmist news source displays a more optimistic tone induces conservatives to further reduce their support for containment 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 Figure A2, which shows the equivalent of Figure 6 for the belief revisions and lockdown support). In Appendix A.5, 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.

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

³⁸ This behavioral measure was added in Survey 2; it was not elicited in Survey 1. See Appendix A.8 for the question wording.

literature on demand-driven media biases.³⁹ In those 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).⁴⁰ 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 possible 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 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.

³⁹ 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.

⁴⁰ See DellaVigna and Kaplan (2007).

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

- Adams-Prassl, A., Boneva, T., Golin, M., and Rauh, C. (2020). Inequality in the impact of the coronavirus shock: Evidence from real time surveys. *Journal of Public Economics*, 189(C).
- Alesina, A., Miano, A., and Stantcheva, S. (2020). The polarization of reality. *AEA Papers and Proceedings*, 110:324–28.
- 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.
- Alsan, M., Braghieri, L., Eichmeyer, S., Kim, M. J., Stantcheva, S., and Yang, D. Y. (2020). Civil liberties in times of crisis. Working Paper 27972, National Bureau of Economic Research.
- 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.
- Baillon, A., Huang, Z., Selim, A., and Wakker, P. P. (2018). Measuring ambiguity attitudes for all (natural) events. *Econometrica*, 86(5):1839–1858.
- 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, National Bureau of Economic Research.
- 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, National Bureau of Economic Research.
- 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., Krasa, 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, National Bureau of Economic Research.
- Boxell, L., Conway, J., Druckman, J. N., and Gentzkow, M. (2020). Affective polarization did not increase during the coronavirus pandemic. Working Paper 28036, National Bureau of Economic Research.
- Brooks, S. K., Webster, R. K., Smith, L. E., Woodland, L., Wessely, S., Greenberg, N., and Rubin, G. J. (2020). The psychological impact of quarantine and how to reduce it: rapid review of the evidence. *Lancet*.
- Brunnermeier, M. and Parker, J. (2005). Optimal Expectations. *American Economic Review*, 95(4):1092–1118.
- Bu, D., Hanspal, T., Liao, Y., and Liu, Y. (2020). Risk taking during a global crisis: Evidence from Wuhan. *Covid Economics*, 5:106–146.
- Bursztyn, L., Rao, A., Roth, C. P., and Yanagizawa-Drott, D. H. (2020). Misinformation during a pandemic. Working Paper 27417, National Bureau of Economic Research.
- 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, National Bureau of Economic Research.
- Caria, S., Fetzer, T., Fiorin, S., Goetz, F., Gomez, M., Haushofer, J., Hensel, L., Ivchenko, A., Jachimowicz, J., Kraft-Todd, G., Reutskaja, E., Roth, C., Witte, M., and Yoeli, E. (2020). Global Behaviors and Perceptions in the COVID-19 Pandemic. Discussion Paper 14631, CEPR.
- Chopra, F., Haaland, I. K., and Roth, C. (2019). Do People Value More Informative News? CESifo Working Paper Series 8026.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020a). Labor Markets During the COVID-19 Crisis: A Preliminary View. Working Paper 27417, National Bureau of Economic Research.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020b). Political polarization and expected economic outcomes. Working Paper 28044, National Bureau of Economic Research.
- Cookson, J. A., Engelberg, J. E., and Mullins, W. (2020). Echo Chambers. SocArXiv n2q9h, 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.
- 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. Working Paper 28004, National Bureau of Economic Research.
- Festinger, L. (1962). Cognitive dissonance. *Scientific American*, 207(4):97–107.
- Fetzer, T., Hensel, L., Hermle, J., and Roth, C. (2020). Coronavirus Perceptions and Economic Anxiety. *Review of Economics and Statistics*, forthcoming.
- 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.
- Ghirardato, P., Maccheroni, F., Marinacci, M., et al. (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(2):133–173.
- Haaland, I. K., Roth, C., and Wohlfart, J. (2020). Designing Information Provision Experiments. CESifo Working Paper Series 8406.
- 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.
- Hoelzl, E. and Rustichini, A. (2005). Overconfident: Do You Put Your Money on It? *Economic Journal*, 115(503):305–318.
- Jevons, W. (1905). *Essays in Economics*. *McMillian*.

- Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D., and Mandel, G. (2012). The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change*, 2(10):732–735.
- 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.
- Mobius, M., Niederle, M., Niehaus, P., and Rosenblat, T. (2014). Managing Self-Confidence. Working paper.
- Montiel Olea, J. L. and Pflueger, C. (2013). A Robust Test for Weak Instruments. *Journal of Business and Economic Statistics*, 31(358-369).
- Mullainathan, S., Schwartzstein, J., and Shleifer, A. (2008). Coarse Thinking and Persuasion. *Quarterly Journal of Economics*, 123(2):577–619.
- 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.
- Nyhan, B., Reifler, J., Richey, S., and Freed, G. L. (2014). Effective messages in vaccine promotion: a randomized trial. *Pediatrics*, 133(4):e835–e842.
- 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. (2020). Political beliefs affect compliance with COVID-19 social distancing orders. Working paper.
- Pflueger, C. E. and Wang, S. (2015). A robust test for weak instruments in Stata. *The Stata Journal*, 15(1):216–225.
- Rabin, M. and Schrag, J. (1999). First Impressions Matter: A Model of Confirmatory Bias. *Quarterly Journal of Economics*, 114(1):37–82.
- Raude, J., Debin, M., Souty, C., Guerris, C., Turbelin, i., Falchi, A., Bonmarin, I., Paolotti, D., Moreno, Y., Obi, C., Duggan, J., Wisnia, A., Flahault, A., Blanchon, T., and Colizza, V. (2020). Are people excessively pessimistic about the risk of coronavirus infection? PsyArxiv Preprint.

- 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. (2020). The Persuasive Effect of Fox News: Non-Compliance with Social Distancing During the Covid-19 Pandemic. Working Paper 27237, National Bureau of Economic Research.
- Strzalecki, T. (2011). Axiomatic Foundations of Multiplier Preferences. *Econometrica*, 79(1):47–73.
- 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 (AZ).
- Weinstein, N. (1980). Unrealistic Optimism About Future Life Events. *Journal of Personality and Social Psychology*, 5(39):820–839.
- Zimmermann, F. (2020). The Dynamics of Motivated Beliefs. *American Economic Review*, 110(2):337–61.

6. Tables

Table 1
Descriptive Statistics

Table 1 has descriptive statistics for the main variables used in the paper.

Variable	Survey 1 ($N=2,440$)		Survey 2 ($N=1,571$)		Difference
	Mean	St. Dev.	Mean	St. Dev.	p -value
College (0/1)	0.51	0.50	0.47	0.50	0.01
Retired (0/1)	0.28	0.45	0.19	0.39	0.00
Unemployed (0/1)	0.06	0.23	0.08	0.26	0.02
Empl. Loss in HH (0/1)	0.29	0.46	0.32	0.47	0.10
HH Income (1–13)	6.36	3.66	6.36	3.58	1.00
Health (1–5)	3.61	0.95	3.47	0.99	0.00
Health Cnds. (0–3)	0.37	0.60	0.45	0.63	0.00
Already Infected (0/1)	0.05	0.22	0.07	0.25	0.06
Liberal (0/1)	0.47	0.50	0.44	0.50	0.09
Age (18–81)	49.13	17.92	45.72	17.29	0.00
Nr. Ppl. above 65 in HH (0/2)	0.40	0.66	0.29	0.57	0.00
Female (0/1)	0.54	0.50	0.51	0.50	0.05
Married (0/1)	0.58	0.49	0.51	0.50	0.00
Black (0/1)	0.09	0.28	0.09	0.28	0.78
Asian (0/1)	0.07	0.26	0.08	0.28	0.21
Hispanic (0/1)	0.07	0.26	0.09	0.29	0.04
Own Primary Residence (0/1)	0.69	0.46	0.62	0.48	0.00
Hold Stocks (0/1)	0.57	0.50	0.49	0.50	0.00
Health Insurance (0/1)	0.91	0.29	0.87	0.34	0.00
Willingness to Take Risks (1–7)	3.89	1.50	3.78	1.62	0.04
Willingness to Wait (1–7)	4.59	1.50	4.61	1.50	0.62
Trust in People (1–7)	3.63	1.70	3.58	1.76	0.42
Stay at Home Order	0.27	0.44	0.19	0.40	0.00
Cases P.C. ($\times 1000$) (0–39)	3.84	5.19	4.54	5.54	0.00
Deaths P.C. ($\times 1000$) (0–2)	0.21	0.33	0.24	0.36	0.00
County UR (1.7–18.3)	14.77	4.47	14.70	4.57	0.63
Worried - Health (0-10)	6.61	2.90	6.40	2.96	0.02
Worried - Economy (0-10)	6.03	2.98	5.97	2.99	0.55
Lockdown Support (1–5)	3.31	1.08	3.35	1.07	0.25
Pr(Deaths $>$ 200k) (0–100) (Prior)	54.68	26.66	56.88	26.26	0.01
Pr(UR $>$ 10%) (0–100) (Prior)	49.38	27.41	49.42	27.17	0.96
Pr(COVID inf.) (0–100) (Prior)	38.05	24.70	37.84	25.94	0.80

Table 2
Determinants of Beliefs (Survey 1)

Table 2 shows regressions where the dependent variable is the sum of Worried Health and Worried Economy (column 1), their difference (column 2), the principal component of the survey participant's prior beliefs (column 3) and a 1 to 5 measure of support for lockdown policies (column 4). The regressions also include date, region, and article domain order fixed effects. (Coefficients not shown.) Heteroskedasticity-consistent standard errors are reported in parentheses. 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.

Dep. Var.	Worried H+E	Worried H-E	PC1	Lockdown Support
	(1)	(2)	(3)	(4)
College	0.59*** (0.23)	0.10 (0.13)	-0.02 (0.04)	0.05 (0.05)
Retired	-0.33 (0.31)	0.50** (0.21)	-0.09 (0.06)	-0.00 (0.06)
Unemployed	0.64 (0.43)	-0.50** (0.25)	-0.10 (0.09)	0.10 (0.09)
Empl. Loss in HH	1.50*** (0.22)	-0.86*** (0.13)	0.10** (0.05)	0.11** (0.05)
Health Cnds. (0-3)	0.46*** (0.17)	0.39*** (0.11)	0.12*** (0.03)	0.01 (0.04)
Income	-0.08** (0.04)	0.07*** (0.02)	-0.02*** (0.01)	0.01 (0.01)
Health	-0.46*** (0.12)	-0.05 (0.07)	-0.00 (0.02)	-0.06** (0.02)
Already Infected	1.11** (0.45)	-0.36 (0.29)	0.39*** (0.11)	-0.08 (0.11)
Liberal	1.87*** (0.20)	0.68*** (0.12)	0.32*** (0.04)	0.70*** (0.04)
Age (\div 100)	-0.97 (0.91)	-0.04 (0.53)	-0.67*** (0.20)	-0.09 (0.20)
Nr Ppl. above 65 in HH	0.11 (0.18)	0.09 (0.12)	0.01 (0.03)	0.10** (0.04)
Female	0.54*** (0.21)	0.00 (0.12)	0.18*** (0.04)	0.09** (0.04)
Married	0.59*** (0.23)	-0.12 (0.13)	0.10** (0.05)	0.01 (0.05)
Black	0.21	-0.07	0.18**	0.14*

Table 2 – *Continued from previous page*

Dep. Var.	Worried H+E	Worried H-E	PC1	Lockdown Support
	(0.40)	(0.19)	(0.09)	(0.08)
Asian	0.36	0.10	0.16*	0.12
	(0.38)	(0.21)	(0.09)	(0.08)
Hispanic	1.59***	0.11	0.07	0.18**
	(0.38)	(0.20)	(0.08)	(0.09)
Own Primary Residence	0.71***	0.22	0.04	-0.03
	(0.25)	(0.14)	(0.05)	(0.05)
Hold Stocks	0.14	-0.41***	-0.05	-0.03
	(0.25)	(0.14)	(0.05)	(0.05)
Health Insurance	0.87**	0.60***	0.05	-0.03
	(0.38)	(0.20)	(0.07)	(0.08)
Willingness to Take Risks	-0.04	-0.25***	0.01	-0.08***
	(0.08)	(0.04)	(0.02)	(0.02)
Willingness to Wait	0.15**	0.11***	0.04***	0.06***
	(0.07)	(0.04)	(0.01)	(0.02)
Trust in People	-0.26***	0.06	-0.00	-0.03**
	(0.06)	(0.04)	(0.01)	(0.01)
Stay at Home Order	1.05***	-0.28**	0.17***	-0.01
	(0.23)	(0.13)	(0.05)	(0.05)
Cases P.C. ($\times 1000$)	-0.02	-0.02	-0.01**	-0.00
	(0.04)	(0.02)	(0.01)	(0.01)
Deaths P.C. ($\times 1000$)	0.61	0.44	0.27**	0.18
	(0.62)	(0.36)	(0.13)	(0.12)
County UR	-0.01	0.00	-0.00	-0.01***
	(0.02)	(0.01)	(0.00)	(0.00)
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 3
Domain Choice

Table 3 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 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 2. 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 4
Pessimistic Article Choice - Survey 1

Table 4 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 2. 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>							
	Health Domain				Economy Domain			
	(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 5
Pessimistic Article Choice - Survey 1 - IV

Table 5 presents instrumental variables 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 1-3) or in the economy domain (columns 4-6). 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 2. The instruments are a dummy equal to 1 if the individual experienced an employment loss in the household and a 0 to 3 measure of health conditions. 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)
Worried H+E	4.02**			4.49**		
	(1.74)			(1.83)		
Worried H-E	3.47			-1.70		
	(2.90)			(3.04)		
PC1		27.94**			32.17**	
		(12.80)			(13.00)	
Lockdown Support			31.56			75.94**
			(21.59)			(35.57)
Observations	2,440	2,440	2,440	2,440	2,440	2,440
Adj. R ²	0.02	-0.07	-0.02	-0.04	-0.08	-1.68
First Stage F-Stat	11.27	8.46	2.63	11.27	8.46	2.63
Mean of Dep. Var.	44.63	44.63	44.63	47.75	47.75	47.75
St. Dev. of Dep. Var.	49.72	49.72	49.72	49.96	49.96	49.96
Controls	YES	YES	YES	YES	YES	YES

Table 6
Belief Revision (Survey 1)

Table 6 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 2. 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 7
Article Assessment (Survey 1)

Table 7 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 donated 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 2. 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 8
Support and Donation to WHO (Survey 1)

Table 8 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 the article read was the preferred one. Control variables (not shown) are included in all the regressions and are as in Table 2. 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 9**Pessimistic Article Choice (Unknown Source) - Survey 2**

Table 9 presents regressions where the dependent variable is a dummy equal to 1 if the participant chooses the article with the pessimistic headline. The sample only includes participants that were not shown the source before choosing the article. 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, and lockdown support is 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 2 (except for the domain order dummies). 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	0.71** (0.29)			0.05 (0.32)	0.56* (0.30)			0.12 (0.33)
Worried H-E	-0.08 (0.55)			-0.92 (0.57)	-0.11 (0.56)			-0.73 (0.59)
PC1		5.92*** (1.55)		4.29** (1.72)		4.56*** (1.61)		3.36* (1.76)
Lockdown Support			5.83*** (1.41)	5.60*** (1.56)			4.55*** (1.49)	4.29*** (1.63)
Observations	1,039	1,038	1,039	1,038	1,039	1,038	1,039	1,038
Adj. R ²	0.00	0.01	0.02	0.02	0.04	0.05	0.05	0.05
Mean of Dep. Var.	38.79	38.73	38.79	38.73	38.79	38.73	38.79	38.73
St. Dev. of Dep. Var.	48.75	48.74	48.75	48.74	48.75	48.74	48.75	48.74
Controls	NO	NO	NO	NO	YES	YES	YES	YES

Table 10
Pessimistic Article Choice - Effect of Source Revelation - Survey 2

Table 10 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 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 9. 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)			0.52* (0.30)		
Worried H+E × Rev. Before	-1.20** (0.48)			-0.82* (0.49)		
PC1		5.40*** (1.60)			4.72*** (1.59)	
PC1 × Rev. Before		-4.13* (2.48)			-1.90 (2.44)	
Lockdown Support			5.76*** (1.46)			4.76*** (1.48)
Lockdown S. × Rev. Before			-7.58*** (2.43)			-3.56 (2.58)
Revealed Before	8.93 (6.47)	-5.89** (2.54)	19.47** (8.44)	15.27** (6.56)	5.70* (3.44)	16.94** (8.42)
Liberal × Rev. Before				-25.44*** (5.10)	-26.64*** (5.07)	-25.10*** (5.37)
Liberal				9.08*** (3.15)	8.63*** (3.15)	7.61** (3.19)
Net Effect	-0.53 (0.40)	1.27 (1.99)	-1.82 (2.06)	-0.30 (0.40)	2.83 (1.92)	1.20 (2.14)
Std. Err.						
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

Table 11

Belief Revision, Lockdown Support and Covid Tracing App (Survey 2)

Table 11 shows regressions where the dependent variables are: the revision in the probability that the unemployment rate will be higher than 20% (column 1), that the number of COVID-19-related deaths will be higher than 200,000 (column 2), that the respondent will be infected (column 3), that the respondent will sign up for a Covid tracing app (column 4), and a 1 to 5 measure of lockdown support (column 5). The sample only includes participants who were not shown the source before choosing the article. Pessimistic is a dummy equal to 1 if the individual is 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. 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 9. Columns 1 through 4 also includes the priors of each measure and dummies for the corner choices (0% and 100%). Column 5 also includes a 1 to 5 measure of support for lockdown measures in the area where the respondent 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.	Δ Prob. UR>20%	Δ Prob. Deaths>200k	Δ Prob. Infected	Δ Prob. CT app	Lockdown Support
	(1)	(2)	(3)	(4)	(5)
Pessimistic	5.98** (2.33)	6.40*** (2.27)	4.78*** (1.73)	-0.01 (1.86)	0.07 (0.14)
Rev. After	-0.50 (1.51)	3.46** (1.76)	0.49 (1.40)	0.38 (1.35)	0.05 (0.11)
Liberal	0.48 (2.47)	-0.04 (2.63)	1.90 (2.32)	1.88 (2.15)	0.08 (0.15)
Pess. \times Liberal	-1.17 (3.31)	0.19 (3.26)	-2.15 (2.46)	2.29 (2.75)	-0.25 (0.20)
Pess. \times Rev.	-3.84 (3.12)	-3.35 (3.15)	-3.31 (2.57)	-0.85 (2.69)	-0.29 (0.20)
Liberal \times Rev.	-0.80 (2.67)	-2.35 (2.58)	0.85 (2.47)	0.19 (2.38)	-0.16 (0.16)
Liberal \times Rev. \times Pess.	1.82 (4.54)	0.03 (4.32)	-1.28 (3.82)	-1.00 (3.83)	0.63** (0.27)
Observations	1,039	1,039	1,039	1,039	1,039
Adj. R ²	0.10	0.10	0.10	0.03	0.34
Mean of Dep. Var.	-0.53	-1.50	-0.32	-0.96	3.47
St. Dev. of Dep. Var.	17.71	17.60	15.24	14.64	1.26
Controls	YES	YES	YES	YES	YES

7. Figures

Figure 1
Structure of Survey 1

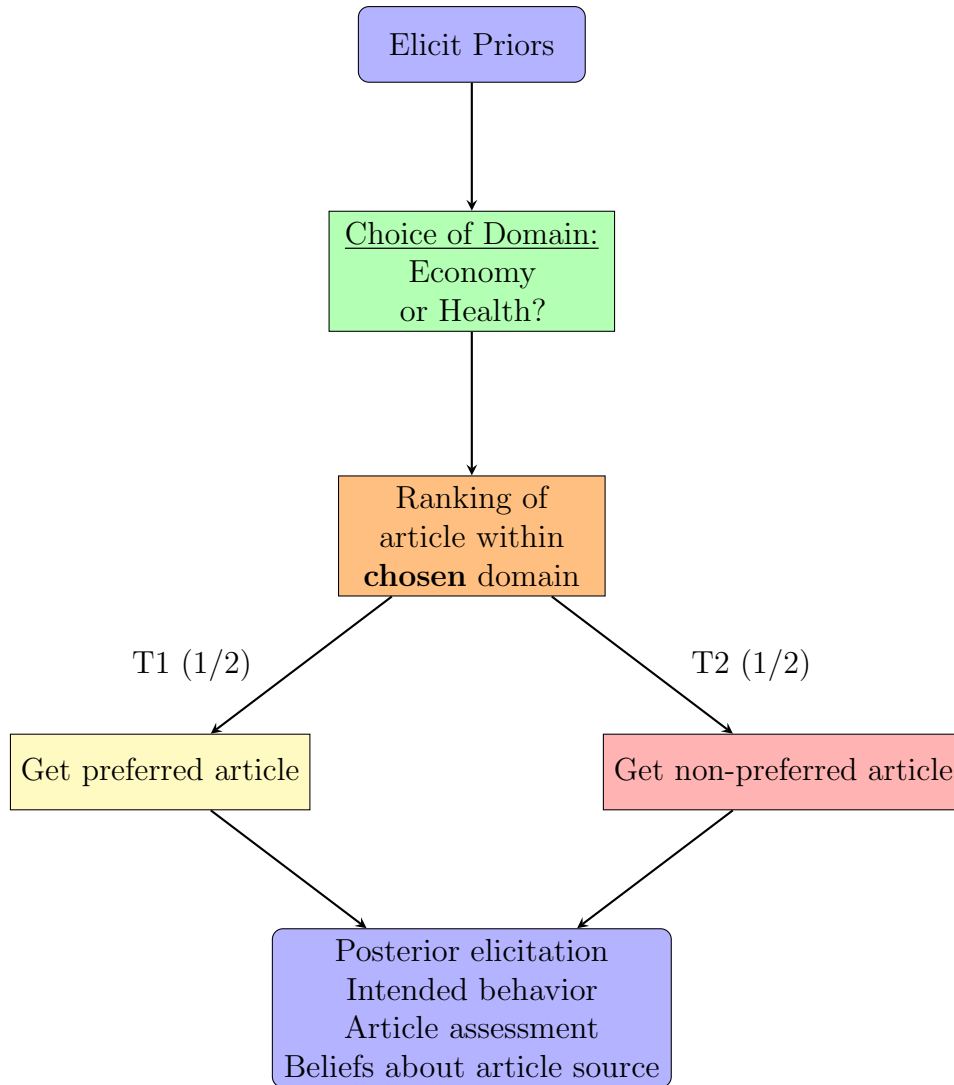


Figure 2
Structure of Survey 2

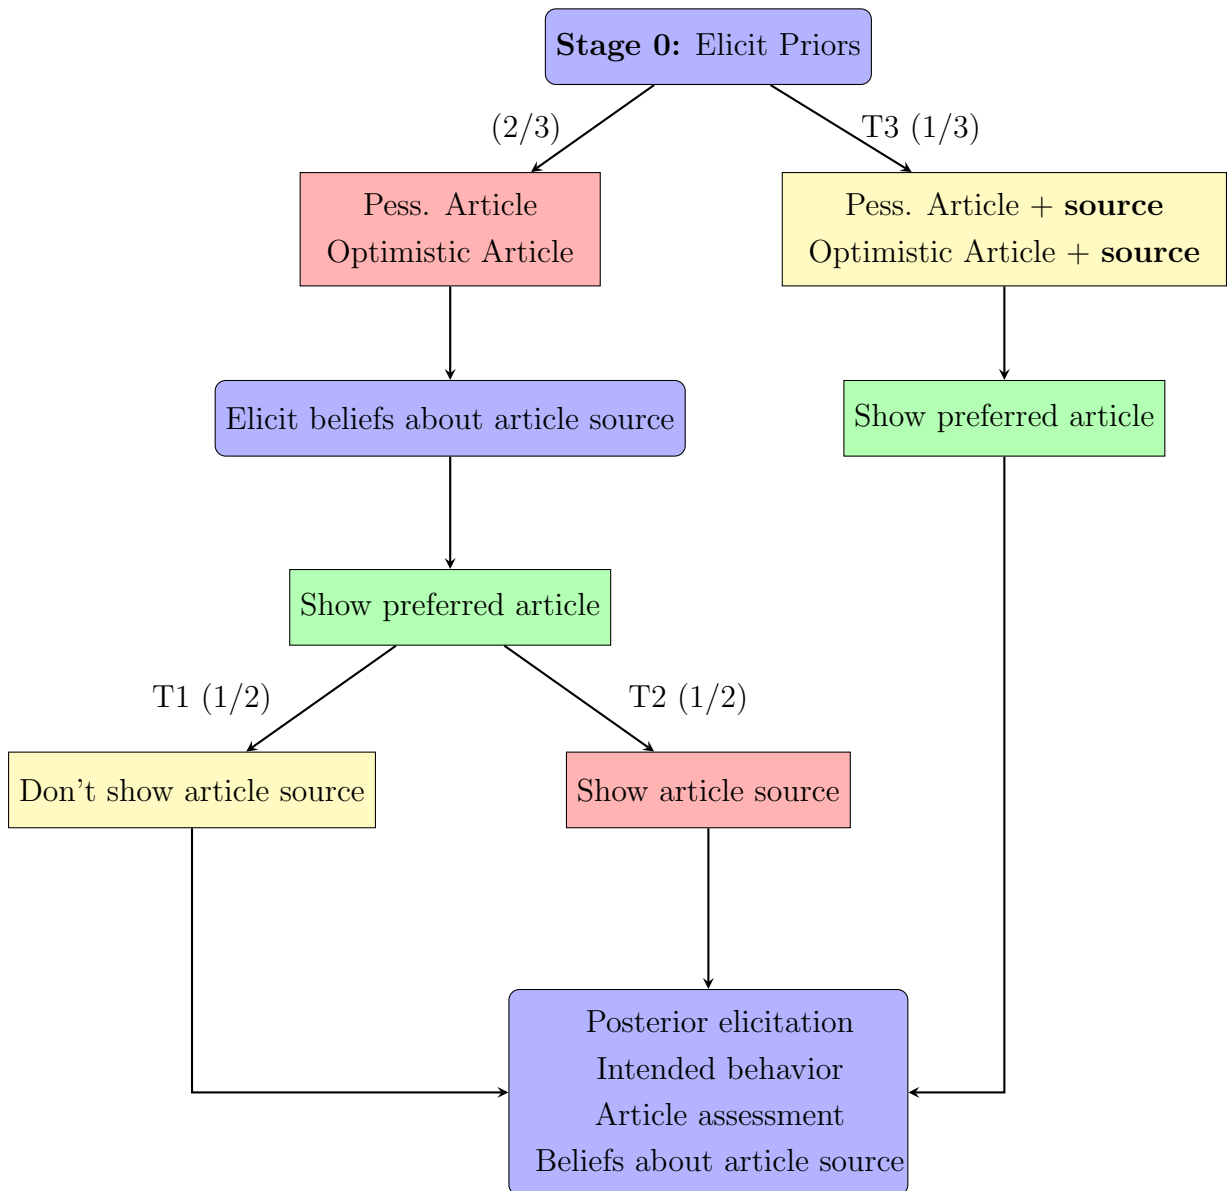


Figure 3
Determinants of Headline Choice (Survey 1)

The histograms in Figure 3 plot the fraction of participants who, in Survey 1, chose the pessimistic headline in the health domain (green bars) and economy domain (red bars). The bins are quintiles of the variables Worried H+E, (Panel a), Worried H-E (Panel b), PC1 (Panel c), and the level of support for lockdown policies (Panel d).

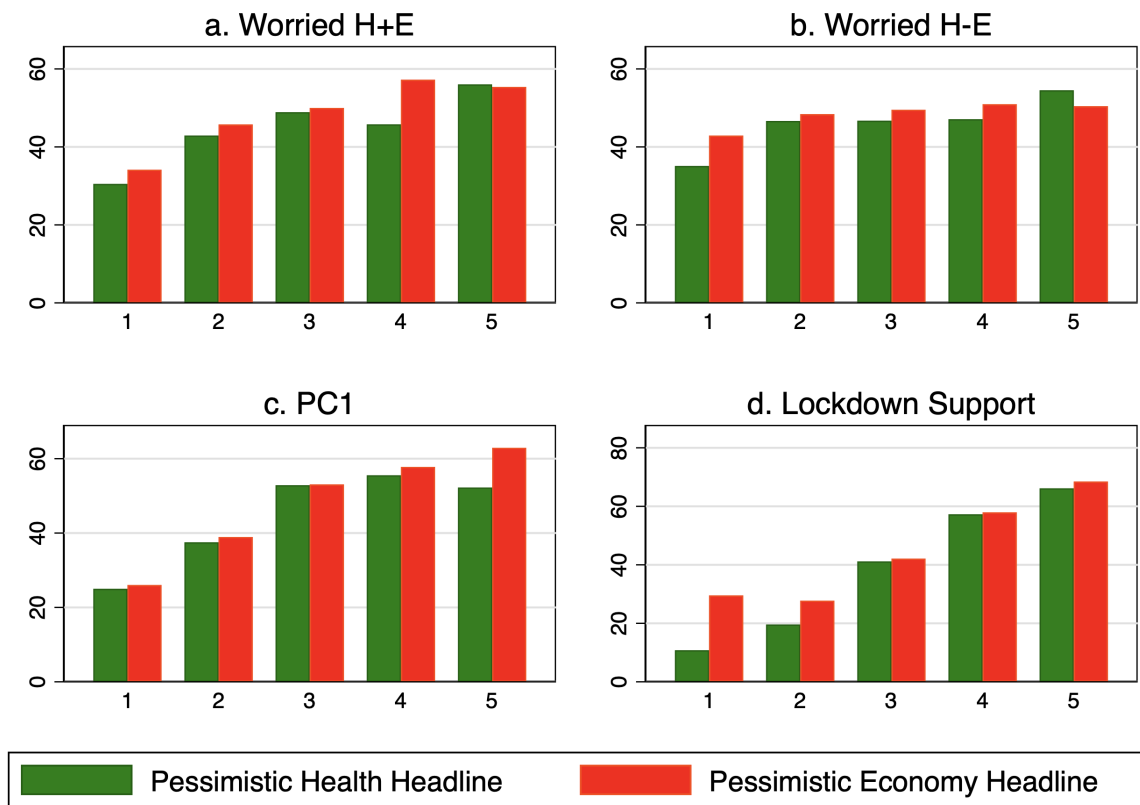



Figure 4
Belief Revision (Survey 1)

The histograms in Figure 4 plot the average of each of the variables indicated in the panels' titles for four participants' subgroups, depending on whether they chose the pessimistic or the optimistic article, and whether they were assigned the preferred or less preferred article. The variables are: the revision in the probability that the number of deaths will be higher than 200,000 by the end of the year, the revision in the probability that the unemployment rate will be higher than 10% by the end of the year, and the change in the probability of being infected by the coronavirus. Panel (a) includes participants who chose the health domain, and Panel (b) includes participants who chose the economy domain.

(a) Health Domain



(b) Economy Domain




Figure 5
Headline Choice (Survey 2)

The histograms plot the fraction of participants who chose the pessimistic headline in Survey 2, among individuals who knew (green bars) or did not know (red 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).


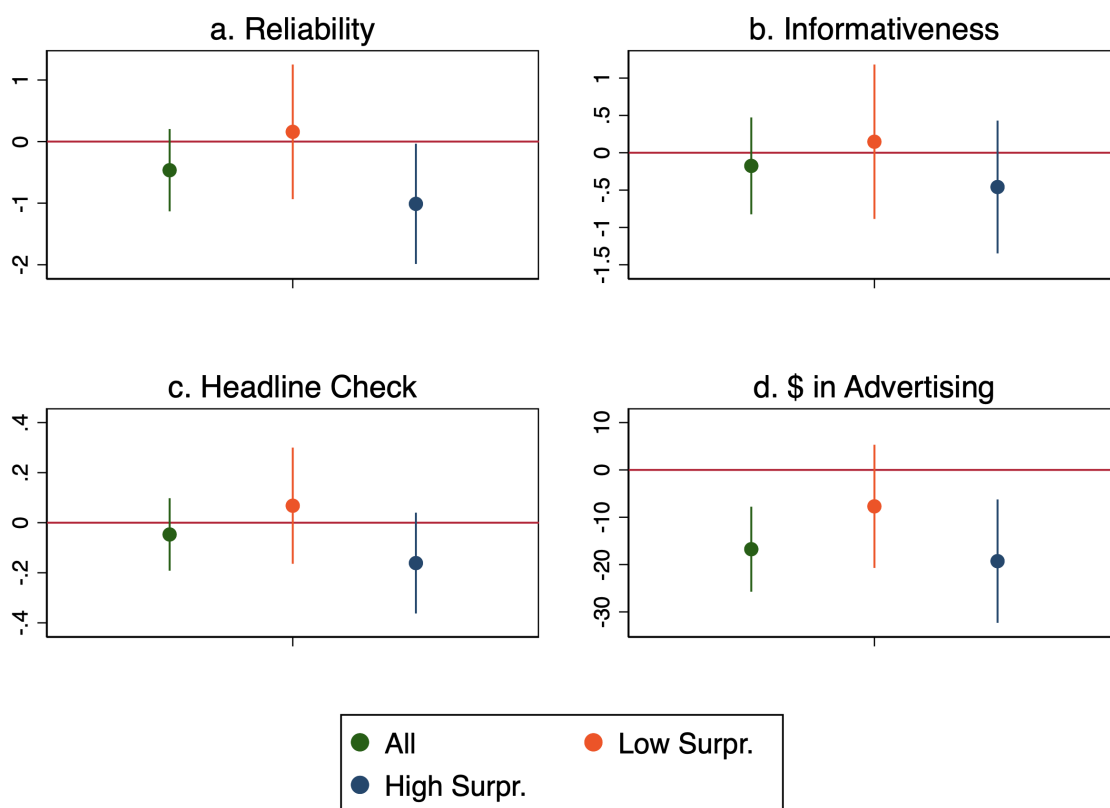


Figure 6
Article Assessment (Survey 2)

Figure 6 shows coefficients on the triple interaction term (Liberal \times Revealed After \times Pessimistic) from equation 4. 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 to the article chosen comes from a liberal source minus the perceived probability that the article assigned comes from a conservative source. Individuals are then sorted into two halves, depending on this measure. Coefficients for individuals in the bottom half are in blue, together with confidence intervals; coefficients for individuals in the top half are in red.



A. Appendix

This appendix includes extensions, details, and results omitted from the main text for brevity. Appendix A.1 has definitions of the main variables used in the paper. Appendix A.2 presents a simple model that rationalizes the empirical evidence. Appendix A.3 has additional results. Appendix A.4 presents additional figures. Appendix A.5 presents descriptive evidence on belief dispersion across the political spectrum. Appendix A.6 describes the pilot studies conducted on MTurk. Appendices A.7 and A.8 provide additional detail on the design and question wording of Survey 1 and 2, respectively.

A.1. Variables Definitions

Table A1
Variables Definitions and Sources

This table has definitions of the main variables used in the regressions.

Variable	Definition
<i>Survey/explanatory variables:</i>	
Preferred	Dummy equal to 1 if the participant reads the article they had ranked as preferred (Survey 1).
Pessimistic	Dummy equal to 1 if the participant reads the article with the pessimistic headline. In Survey 1, when individuals who chose the health or the economy domain are pooled, we distinguish between a "Pessimistic - Health" and a "Pessimistic - Economy" dummy.
Source Revealed	A dummy equal to 1 if the participant is revealed the source of each article before making the choice (Survey 2).
Revealed After	A dummy equal to 1 if the participant is revealed the article source after having made the choice (Survey 2).
<i>Measured in prior and posterior stage:</i>	
Forecast Deaths	The participant's forecast of the number of Covid-related deaths at the end of the year, winsorized at the 2.5% level.
Forecast UR	The participant's forecast of the unemployment rate at the end of the year, winsorized at the 2.5% level.
Pr. Deaths > 200k	The participant's estimate of the probability that the number of Covid-related deaths will be higher than 200,000 by the end of the year.
Pr. UR > 20%	The participant's estimate of the probability that the unemployment rate will be higher than 20% by the end of the year.

Continued on next page

Table A1 – *Continued from previous page*

Variable	Definition
Pr. Infection	An estimate of the probability that the participant will be infected by the coronavirus by the end of the year.
Pr. Contact Tracing (CT) App	The likelihood that the participant will voluntarily sign up for a contact tracing app (Survey 2).
<i>Measured in prior stage:</i>	
Worried - Health	A measure, on a 1 to 10 scale, of how worried the participant is about the health consequences of the pandemic.
Worried - Economy	A measure, on a 1 to 10 scale, of how worried the participant is about the economic consequences of the pandemic.
Worried H+E	The sum of Worried - Economy and Worried - Health.
Worried H-E	The difference between Worried - Health and Worried - Economy.
Nowcast Deaths	The participant's estimate of the number of deaths related to the Covid pandemic at the moment of the survey, winsorized at the 2.5% level.
Nowcast UR	The participant's estimate of the current unemployment rate, winsorized at the 2.5% level.
PC1	The first principal component of Worried - Health, Worried - Economy, Pr. Deaths > 200k, Pr. UR > 20%, Forecast Deaths, Forecast UR, and Pr. Infection (all measured in prior stage).
Lockdown Support	An evaluation of lockdown measures taken so far, from 1 ("The measures went much too far") to 5 ("The measures did not go nearly far enough").
<i>Measured in posterior stage (after article is shown):</i>	
Informativeness	An assessment of the informativeness of the article read on a scale from 1 to 7.
Reliability	An assessment of the reliability of the article read on a scale from 1 to 7.
Headline Check	A dummy equal to 1 if the participant correctly identifies, among four headlines, the one corresponding to the article just read.
Amount Spent for Advertising	The dollar amount that the participant is willing to spend to advertise the news source on Facebook, from a \$50 budget.

Continued on next page

Table A1 – *Continued from previous page*

Variable	Definition
Support	A measure of support for future lockdown policies from 1 (“No intervention, just let the virus run its course”) to 5 (“Stricter stay-home orders than this time, even if the economic cost is very large”).
Money to WHO	The dollar amount the participant is willing to donate to the World Health Organization, as opposed to FreedomWorks, out of a budget of 20 dollars (Survey 1).
<i>Respondent characteristics:</i>	
College	A dummy equal to 1 if the participant has a bachelor, master, or doctoral degree.
Retired	A dummy equal to 1 if the participant is a retiree or an early retiree.
Unemployed	A dummy equal to 1 if the participant is unemployed.
Employment Loss in HH	A dummy equal to 1 if somebody in the participant’s household has suffered an employment loss related to the pandemic.
Income	A measure of household income that takes values 0 (“Less than \$10,000”), 1 (“\$10,000 - \$19,999”), 2 (“\$20,000 - \$29,999”), 3 (“\$30,000 - \$39,999”), 4 (“\$40,000 - \$49,999”), 5 (“\$50,000 - \$59,999”), 6 (“\$60,000 - \$69,999”), 7 (“\$70,000 - \$79,999”), 8 (“\$80,000 - \$89,999”), 9 (“\$90,000 - \$99,999”), 10 (“\$100,000 - \$149,999”), 11 (“\$150,000 - \$199,999”), 12 (“\$200,000 or more”).
Health	Self-assessed overall health, from 1 (“poor”) to 5 (“excellent”).
Nr Health Conditions	The number of conditions the participant suffers or has suffered from, between (i) Lung disease such as chronic bronchitis or emphysema; (ii) Diabetes or high blood sugar; and (iii) A heart attack, coronary heart disease, angina, congestive heart failure, or other heart problems. List of conditions was taken from Allcott et al. (2020).
Already Infected	A dummy equal to 1 if the participant indicates that they have been tested positive for COVID-19 or that they have very likely been infected.
Liberal	A dummy equal to 1 if the participant defines herself as “very liberal/democrat”, “liberal/democrat”, or “leaning liberal/democrat”, and 0 otherwise.
Age	The participant’s age.
People in HH	Number of individuals in the household, capped at 7.

Continued on next page

Table A1 – *Continued from previous page*

Variable	Definition
Nr People above 65 in HH	Number of individuals older than 65 in the household.
Female	A dummy equal to 1 if the participant is a woman.
Married	A dummy equal to 1 if the participant is married.
Black, Asian, Hispanic	Dummies corresponding to the participant's race/ethnicity.
Own Primary Residence	A dummy equal to 1 if the participant owns her primary residence.
Hold Stocks	A dummy equal to 1 if the participant owns stocks.
Health Insurance	A dummy equal to 1 if the participant owns a health insurance.
Willingness to Take Risks	Self-reported willingness to take risks, from 1 (“Absolutely unwilling to take risks”) to 7 (“Fully prepared to take risks”).
Willingness to Wait	Self-reported willingness to give up something that is beneficial today in order to benefit more from that in the future, from 1 (“Completely unwilling to give up”) to 7 (“Very willing to give up”).
Trust in People	Self-reported trust in other people, from 1 (“Most people can be trusted”) to 7 (“Most people can be trusted”).
<i>Local characteristics:</i>	
Stay at Home Order	A dummy equal to 1 if there is a stay at home order in the area where the participant lives.
Cases Per Capita	Number of Covid cases per capita multiplied by 1 million, measured the day before the participant takes the survey.
Deaths Per Capita	Number of Covid-related deaths per capita multiplied by 1 million, measured the day before the participant takes the survey.
County UR	Unemployment rate at the county level as of April, 2020.

A.2. A Model of Confirmatory Biases with Endogenous Information Acquisition

In this appendix, we outline a model that rationalizes our design and the results. The model features two main stages. In stage one, agents choose their desired belief, after observing a signal. In stage zero, they choose which signal they want to see, based on partial information about its content. Technically, we formalize motivations through multiplier preferences.⁴¹ Agents entertain different subjective beliefs, choosing the ones that maximize their (subjectively weighted) expected utility, but minimizing the cost of being wrong. The subjective probability bias, relatively to the objective probability, stems from utility gains. The bounds on the extent of bias is obtained by imposing a classical cost on entropy, or a Kullback-Leibler measure on the distance between the desired belief and the correct one. In other words, agents are willing to choose what to believe to the extent that this is not too wrong. For this stage, we build on Caplin and Leahy (2019), who offer a model where motivation twists posterior probabilities toward optimism. We depart from them in that we consider the more general category of confirmatory biases, according to which agents, whether optimist or pessimist, update by reinforcing their prior beliefs. Knowing that in stage one they will choose optimally motivated beliefs, in stage zero agents select a signal that they expect to provide them the highest utility (this stage is solved last by backward induction). Below we provide the main derivations of the model and show how its implications match our empirical evidence.

Stage 1 - Choice of Motivated Beliefs. In stage 1, the choice of the motivated beliefs is derived. Priors are defined as π , the state as ω , the signal as s (which in our design is represented by the content of the articles), agent's utility as $u(\omega)$,⁴² the objective probability as $\bar{p}(s | \omega)$, and the subjective probability as $p(s | \omega)$.

The agents choose their desired subjective belief so as to maximize their utility, embedding the motivation, while minimizing the distance between their motivated belief and the objective one. As a result, it is never optimal to be “too wrong.” Following Baillon et al. (2017), the optimal motivated belief, $p(s | \omega)$, is chosen to maximize the following value function:

$$V(\pi, s) = \begin{cases} \max_p \sum_{\omega \in \Omega} \frac{p(s|\omega)\pi(\omega)}{\sum_{\omega' \in \Omega} p(s|\omega')\pi(\omega')} u(\omega) + \sigma R(p|\omega) & \text{if } \sigma < 0 \\ \sum_{\omega \in \Omega} \frac{p(s|\omega)\pi(\omega)}{\sum_{\omega' \in \Omega} p(s|\omega')\pi(\omega')} u(\omega) & \text{if } \sigma = 0 \\ \min_p \sum_{\omega \in \Omega} \frac{p(s|\omega)\pi(\omega)}{\sum_{\omega' \in \Omega} p(s|\omega')\pi(\omega')} u(\omega) - \sigma R(p|\omega) & \text{if } \sigma > 0, \end{cases} \quad (\text{A.1})$$

41 Under certain conditions, those can be mapped into maxmin or variational expectation preferences; see Strzalecki (2011)

42 In general, set-ups with endogenous beliefs also feature an action choice a , after the choice of the optimal belief, so utility is written as $u(a, \omega)$; here we simplify the set-up and assume that utility only depends on the state of the world directly. The underlying motivations might take several forms (individual, political, etc.); we are not explicit about that in the model.

where $R(p|\omega) \equiv \sum_{\omega \in \Omega} \sum_{s \in \mathcal{S}} p(s' | \omega) \ln \frac{p(s'|\omega)}{\bar{p}(s'|\omega)}$ is the Kullback-Leibler measure, and the multiplier σ determines the relative importance of the two terms (and can be interpreted as the penalty term). As σ approaches positive or negative infinity the model delivers rational Bayesian updating, as we show further below. Intuitively, if the cost of being wrong becomes very large, the agent exhibits near-rational beliefs. Importantly, we do not restrict the multiplier σ to be either positive (as it is typically assumed in robust control; see Hansen and Sargent, 2008) or negative (see, for instance, the wishful thinking model of Caplin and Leahy, 2019), but assume that it can vary in the cross-section of respondents. A positive multiplier is typically interpreted as capturing ambiguity aversion, which leads to pessimistic optimal beliefs, meaning that the assign higher probability to low-utility states. A negative value of the multiplier, which one can interpret as ambiguity or opportunity seeking, leads to more optimistic optimal subjective beliefs. As we will show below, respondents in our survey vary widely in how relatively optimistic/pessimistic they are, but the relative optimism in different dimensions (health, economy) is quite strongly positively correlated within a respondent. Within the model, variation in the multiplier σ captures the variation of optimism and pessimism in our population of respondents and, as shown below, also determines asymmetric updating (overweighting of desired signals).⁴³

Taking the first order condition of (A.1) with respect to $p(s | \omega)$ leads to:

$$\frac{dE_\gamma u(\omega)}{dp(\omega)} + \sigma \ln \frac{p(\omega)}{\bar{p}(\omega)} - \sigma - \lambda_\omega = 0, \quad (\text{A.2})$$

where $E_\gamma u(\omega) = \sum_{\omega \in \Omega} \frac{p(s|\omega)\pi(\omega)}{\sum_{\omega' \in \Omega} p(s|\omega')\pi(\omega')} u(\omega)$, with γ denoting the Bayesian posterior, and where, for a given signal realization, we can replace $p(s | \omega)$ with $p(\omega)$ and $\bar{p}(s | \omega)$ with $\bar{p}(\omega)$. Upon solving for the optimal beliefs from (A.2), one obtains:

$$p(\omega) = \bar{p}(\omega) \exp \left[-\frac{1}{\sigma} \left(\frac{\pi(\omega)}{E_p \pi(\omega)} [u(\omega) - E_\gamma u(\omega)] \right) + 1 + \frac{1}{\sigma} \lambda_\omega \right], \quad (\text{A.3})$$

where we have substituted $\frac{dE_\gamma u(\omega)}{dp(s|\omega)} = \frac{\pi(\omega)}{E_p \pi(\omega)} [u(\omega) - E_\gamma u(\omega)] \equiv m$, the ‘‘motivational’’ term.⁴⁴ First, note that (A.3) implies that the bias relates to signals for which the derivative of the expected utility is not zero. At this stage it is useful to define a pessimistic signal as one for which m is negative, i.e. $u(\omega) < E_\gamma u(\omega)$ and, conversely, an optimistic signal as one that delivers $m > 0$. Note that if $\sigma < 0$ the optimal belief underweights pessimistic signals compared to the objective distribution, $\bar{p}(\omega)$, and overweights optimistic ones. The opposite is true when $\sigma > 0$. Next, we derive the updated belief. Given all possible signals, the Bayesian posterior is given by $\gamma(\omega) = \frac{p(\omega)\pi(\omega)}{\sum_{\omega} p(\omega')\pi(\omega')}$. To fix ideas we

⁴³ There is increasing evidence that σ varies within the population and also over time (e.g., Baillon et al., 2017, 2018). Varying ambiguity attitudes leading to optimism and pessimism are also axiomatized in Ghirardato et al. (2004).

⁴⁴ Note that by imposing the constraint $\sum_{s'} p(s'|\omega) = 1$, the constant term in the first order condition is set to zero, so that for $\sigma = 0$ the optimal belief equals the rational expectation benchmark.

consider only two states of the world, ω and ω' , which in our design would correspond to unemployment or deaths being lower or higher than a certain threshold. Computing the optimal posterior through the Bayes rule $\gamma(\omega)$, delivers:

$$\gamma(\omega) = \frac{\bar{p}(\omega)\exp(-\frac{m}{\sigma}) \cdot \pi(\omega)}{\bar{p}(\omega)\exp(-\frac{m}{\sigma}) \cdot \pi(\omega) + (1 - \bar{p}(\omega)) \cdot (1 - \pi(\omega))}. \quad (\text{A.4})$$

First note that for $\frac{1}{\sigma}$ approaching 0, equation (A.4) results in $\gamma(\omega) = \frac{\bar{p}(\omega)\pi(\omega)}{\bar{p}(\omega)\pi(\omega) + (1-\bar{p}(\omega))(1-\pi(\omega))}$, hence the classical rational Bayesian updating, where signals are equally weighted. Consider now a negative signal, $m < 0$. If $\sigma < 0$, the posterior *underweights* pessimistic signals compared to a rational Bayesian updater and *overweights* optimistic signals. The opposite is true with $\sigma > 0$, that is, when the respondent is a pessimist.⁴⁵ Hence, optimists overweight signals that assign higher likelihood to states with utility above the mean (optimistic signals) and underweight pessimistic signals. Conversely, pessimists overweight pessimistic signals and underweight optimistic signals. From now on we define γ^+ as the posterior of an agent with $\sigma < 0$ (an optimist) and γ^- as the posterior of an agent with $\sigma > 0$ (a pessimist).

Stage zero - Choice of Signal. Our survey design also features a stage zero in which the respondent chooses among two headlines with differing tones, an optimistic one and a pessimistic one. The headlines effectively represent “signals of signals,” in that they provide hints on the possible content of the article, which is the signal considered in stage 1.

Anticipating the motivated beliefs that she will choose in stage 1, the agent is now going to choose the signal that she expects to deliver the highest expected utility under the optimal motivated beliefs. The latter is defined as follows: $\tilde{V}^+(\pi, s) = \sum_{\omega \in \Omega} \gamma^+(\omega)u(\omega)$ for an optimist and as $\tilde{V}^-(\pi, s) = \sum_{\omega \in \Omega} \gamma^-(\omega)u(\omega)$ for a pessimist. Let us consider two signals, $\xi_1 \sim \mathcal{N}(\mu^+, \frac{1}{\tau})$ and $\xi_2 \sim \mathcal{N}(\mu^-, \frac{1}{\tau})$, with $\mu^+ > \mu^-$. Assume that a higher value signifies more positive news (for example, a higher employment rate or speed in developing an effective care). Since the first signal is more likely to deliver positive news, $\int \tilde{V}^+(\pi, s)d\xi_1 > \int \tilde{V}^+(\pi, s)d\xi_2$, an optimist derives higher expected utility from it, as in expectation they can obtain an optimistic posterior with a smaller cost from belief distortion. The opposite is true for pessimists: for them, $\int \tilde{V}^-(\pi, s)d\xi_2 > \int \tilde{V}^-(\pi, s)d\xi_1$.

Implications. This theoretical framework has several implications.

(1) The signal choice reveals one’s type: optimists prefer the positive headline, pessimists the negative headline.

⁴⁵ That is to say if there is a pessimistic signal which makes m negative and σ is positive, the result is a positive number and the exponential of a positive number is larger than one; hence, the agent overweights negative signals.

(2) Since, in reality, belief updating occurs continuously, we expect the optimists to already enter our survey with more optimistic priors, so that we predict a positive relationship between optimistic priors and the choice of the positive headline.

(3) Turning to the updating, the model implies that the posterior should put more weight on a signal that an agent wanted to see, and that is more in line with their priors. Thus, both the headline choice and updating display “confirmatory bias”.

Our empirical results are in line with these implications.

A.3. Additional Results

Table A2
Correlation Matrix (Survey 1): Components of PC1

Table A2 shows a correlation matrix of the following variables: the probability that unemployment rate will be higher than 20% by the end of the year; the probability that the number of deaths will be higher than 200,000; the forecasts for the unemployment rate and the number of Covid-19-related deaths; the estimates of the *current* unemployment rate and Covid-19-related deaths; the probability of being infected by the end of the year. In addition, the matrix includes the first principal component of the seven variables (PC1). ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively.

	Pr(Deaths > 200k)	Pr(UR > 10%)	Pr(COVID inf.)	Forecast Deaths	Forecast UR	Nowcast Deaths	Nowcast UR	PC1
Pr(Deaths > 200k)	1.00							
Pr(UR > 10%)	0.49***	1.00						
Pr(COVID inf.)	0.36***	0.31***	1.00					
Forecast Deaths	0.28***	0.11***	0.07***	1.00				
Forecast UR	0.15***	0.45***	0.11***	0.13***	1.00			
Nowcast Deaths	0.18***	0.07***	0.03	0.81***	0.08***	1.00		
Nowcast UR	0.04	0.23***	0.06**	0.08***	0.72***	0.05**	1.00	
PC1	0.62***	0.68***	0.42***	0.59***	0.68***	0.52***	0.55***	1.00

Table A3
Correlation Matrix (Survey 1)

Table A2 shows a correlation matrix of the following variables: Worried H+E, Worried H-E, PC1, and Lockdown Support. ***, **, and * indicate statistically different from zero at the 1%, 5%, and 10% level of significance, respectively.

	Worried H+E	Worried H-E	PC1	Lockdown Support
Worried H+E	1.00			
Worried H-E	-0.03	1.00		
PC1	0.35***	0.02	1.00	
Lockdown Support	0.33***	0.30***	0.23***	1.00

Table A4
Randomization Tests (Survey 1)

The first column of Table A4 displays means of the main variables used in the text for survey participants who were selected to read their preferred article and less preferred article, respectively. The third column displays their difference. The fourth column displays the p -value computed under the null hypothesis that the difference is zero.

	Preferred	Non-Pref.	Difference	p -value
College	0.53	0.50	0.02	0.24
Retired	0.27	0.28	-0.00	0.89
Unemployed	0.06	0.06	0.00	0.74
Empl. Loss in HH	0.28	0.31	-0.03	0.12
HH Income	6.28	6.44	-0.16	0.28
Health	3.65	3.57	0.08	0.04
Health Cnds. (0-3)	0.35	0.38	-0.03	0.28
Already Infected	0.05	0.06	-0.01	0.26
Liberal	0.45	0.48	-0.04	0.07
Age (\div 100)	0.49	0.49	-0.00	0.93
Nr Ppl. above 65 in HH	0.37	0.42	-0.05	0.05
Female	0.53	0.55	-0.01	0.53
Married	0.56	0.59	-0.03	0.18
White	0.80	0.80	0.00	0.83
Own Primary Residence	0.70	0.68	0.02	0.26
Hold Stocks	0.57	0.56	0.01	0.75
Health Insurance	0.91	0.91	0.00	0.91
Willingness to Take Risks	3.91	3.87	0.05	0.42
Willingness to Wait	4.62	4.56	0.05	0.41
Trust in People	3.73	3.53	0.20	0.00
Stay at Home Order	0.27	0.27	0.00	0.86
Cases P.C. (\times 1000)	3.81	3.87	-0.06	0.78
Deaths P.C. (\times 1000)	0.21	0.21	0.00	0.83
County UR	14.77	14.78	-0.01	0.97
Worried H+E	12.20	13.07	-0.87	0.00
Worried H-E	0.64	0.52	0.12	0.32
PC1	-0.02	0.02	-0.05	0.23
Lockdown Support	3.29	3.34	-0.05	0.28
Pr(Deaths > 200k)	54.02	55.33	-1.31	0.23
Pr(UR > 10%)	48.14	50.60	-2.46	0.03
Pr. Infected	37.29	38.80	-1.50	0.13

Table A5
Political Preferences and Belief Revisions (Survey 1)

The first four rows of Table A5 shows averages of the revision in 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 that the respondent will be infected for four groups of respondents, depending on their political preferences (liberals or conservative) and on whether they got their preferred article. The fifth row is the difference between the averages of each variables, including only participants who got their less preferred article; the sixth row includes only participants who got their preferred article. The seventh row is the difference between the sixth and the fifth row. The last column reports *p*-values of the differences.

Political Preferences	Preferred Article?	ΔProb. UR>20%	ΔProb. Deaths>200k	ΔProb. Infected
Liberal	NO	-1.77	-1.24	-0.41
Liberal	YES	0.00	-1.21	0.45
Conservative	NO	-2.39	-3.39	-0.43
Conservative	YES	-3.20	-3.64	-1.72
Liberal – Conservative	NO	0.62	2.15	0.02
Liberal – Conservative	YES	3.19	2.43	2.17
Difference in Differences	–	2.58	0.28	2.15
<i>p</i> -value	–	0.09	0.85	0.07

Table A6
Randomization Tests (Survey 2)

The first, second, and third column of Table A4 display means of the main variables used in the tests for survey participants who were selected in treatment 1 (article source never revealed), 2 (article source revealed after the choice), and 3 (article source revealed immediately). Heteroskedasticity-consistent standard errors are reported in parentheses. The last column displays p -values for χ^2 -statistics computed under the null hypothesis that the means are all equal.

	T1	T2	T3	p-value
College	0.50	0.46	0.46	0.46
Retired	0.20	0.20	0.18	0.67
Unemployed	0.08	0.08	0.08	1.00
Empl. Loss in HH	0.33	0.32	0.31	0.79
HH Income	6.34	6.25	6.48	0.57
Health	3.47	3.51	3.44	0.49
Health Cnds. (0–3)	0.43	0.43	0.48	0.40
Already Infected	0.06	0.07	0.07	0.95
Liberal	0.43	0.45	0.43	0.80
Age ($\div 100$)	0.46	0.45	0.46	0.72
Nr Ppl. above 65 in HH	0.30	0.30	0.26	0.39
Female	0.51	0.48	0.53	0.27
Married	0.54	0.52	0.49	0.30
White	0.79	0.76	0.79	0.46
Own Primary Residence	0.63	0.62	0.63	0.95
Hold Stocks	0.49	0.49	0.48	0.96
Health Insurance	0.88	0.86	0.86	0.69
Willingness to Take Risks	3.79	3.86	3.70	0.25
Willingness to Wait	4.57	4.69	4.59	0.35
Trust in People	3.71	3.52	3.52	0.13
Lockdown Severity	2.98	3.01	2.92	0.16
Cases P.C. ($\times 1000$)	4.66	4.53	4.43	0.82
Deaths P.C. ($\times 1000$)	0.25	0.24	0.23	0.74
County UR	14.72	14.63	14.75	0.92
Worried H+E	12.22	12.66	12.25	0.32
Worried H-E	0.56	0.34	0.37	0.38
PC1	-0.03	0.02	0.01	0.72
Lockdown Support	3.34	3.36	3.36	0.95
Pr(Deaths > 200k)	57.22	58.87	54.63	0.03
Pr(UR > 10%)	49.70	50.22	48.38	0.54

Table A7
Determinants of Beliefs (Survey 2)

Table A7 shows regressions where the dependent variable is the sum of Worried Health and Worried Economy (column 1), their difference (column 2), the principal component of the survey participant's prior beliefs (column 3) and a 1 to 5 measure of support for lockdown policies (column 4). The regressions also include date and region fixed effects. (Coefficients not shown.) Heteroskedasticity-consistent standard errors are reported in parentheses. 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.

Dep. Var.	Worried H+E	Worried H-E	PC1	Lockdown Support
	(1)	(2)	(3)	(4)
College	0.21 (0.29)	0.15 (0.15)	0.08 (0.06)	0.06 (0.06)
Retired	-0.66 (0.47)	0.53** (0.26)	-0.09 (0.08)	0.02 (0.09)
Unemployed	0.04 (0.49)	-0.20 (0.22)	-0.04 (0.10)	-0.05 (0.09)
Empl. Loss in HH	2.07*** (0.28)	-1.20*** (0.15)	0.19*** (0.05)	0.03 (0.06)
Health Cnds. (0-3)	0.69*** (0.20)	0.33*** (0.11)	0.19*** (0.04)	0.08* (0.04)
Income	0.05 (0.04)	0.08*** (0.02)	-0.02** (0.01)	0.00 (0.01)
Health	-0.42*** (0.15)	-0.19** (0.08)	0.00 (0.03)	-0.09*** (0.03)
Already Infected	0.71 (0.48)	-0.25 (0.28)	0.38*** (0.10)	-0.01 (0.12)
Liberal	1.43*** (0.26)	0.60*** (0.14)	0.30*** (0.05)	0.56*** (0.05)
Age (\div 100)	-0.35 (1.10)	-0.19 (0.58)	-0.30 (0.21)	-0.12 (0.23)
Nr Ppl. above 65 in HH	-0.22 (0.26)	0.02 (0.14)	-0.04 (0.04)	-0.05 (0.05)
Female	0.81*** (0.29)	-0.09 (0.15)	0.14** (0.06)	0.07 (0.06)
Married	0.02 (0.29)	0.05 (0.15)	0.01 (0.05)	-0.04 (0.06)
Black	-0.12	-0.61***	0.01	-0.07

Continued on next page

Table A7 – *Continued from previous page*

Dep. Var.	Worried H+E	Worried H-E	PC1	Lockdown Support
	(0.49)	(0.22)	(0.10)	(0.09)
Asian	1.21***	-0.02	0.06	0.18**
	(0.44)	(0.20)	(0.09)	(0.09)
Hispanic	0.39	0.03	0.09	0.17*
	(0.44)	(0.21)	(0.09)	(0.09)
Own Primary Residence	0.45	-0.01	0.08	-0.03
	(0.30)	(0.16)	(0.06)	(0.06)
Hold Stocks	0.23	-0.09	-0.04	0.03
	(0.29)	(0.16)	(0.05)	(0.06)
Health Insurance	0.03	0.13	0.05	-0.04
	(0.38)	(0.21)	(0.08)	(0.08)
Willingness to Take Risks	-0.24**	-0.27***	-0.01	-0.10***
	(0.10)	(0.05)	(0.02)	(0.02)
Willingness to Wait	0.47***	0.15***	0.05***	0.11***
	(0.10)	(0.05)	(0.02)	(0.02)
Trust in People	-0.08	0.07*	-0.02	-0.01
	(0.08)	(0.04)	(0.02)	(0.02)
Stay at Home Order	1.51***	0.09	0.25***	0.04
	(0.34)	(0.16)	(0.07)	(0.07)
Cases P.C. ($\times 1000$)	0.12**	0.01	0.01	0.01
	(0.05)	(0.02)	(0.01)	(0.01)
Deaths P.C. ($\times 1000$)	-1.74**	0.02	-0.02	-0.09
	(0.72)	(0.39)	(0.13)	(0.12)
County UR	0.03	0.14**	-0.02	0.01
	(0.11)	(0.07)	(0.03)	(0.02)
Observations	1,572	1,572	1,571	1,572
Adj. R ²	0.12	0.13	0.09	0.14
Mean of Dep. Var.	12.37	0.42	-0.00	3.35
St. Dev. of Dep. Var.	5.28	2.75	1.00	1.07

Table A8**Beliefs about Article Sources**

Table A8 reports, for each article tone (pessimistic or optimistic) the average estimate of the probability that the article comes from a conservative source (Breitbart, Fox News, and The Wall Street Journal), from a liberal source (CNN, The New York Times, MSNBC, and The Washington Post), from The New York Times, and from Fox News. These values are labeled Pr(Liberal), Pr(Conservative), Pr(NYT), and Pr(Fox), respectively. The table also reports the difference between Pr(Liberal) and Pr(Conservative) and the difference between Pr(NYT) and PR(Fox). The statistics are reported for participants in Survey 1 who chose the health domain, participants in Survey 1 who chose the economy domain, and participants in Survey 2.

Article Tone	Prob(Lib.)	Prob(Cons.)	Prob(Lib.) - Prob(Cons.)	Prob(NYT)	Prob(Fox)	Prob(NYT) - Prob(Fox)
<i>Survey 1 (Health Domain)</i>						
Pessimistic	55.49	27.37	28.12	14.82	10.39	4.43
Optimistic	37.83	38.41	-0.58	10.53	22.76	-12.23
Pessimistic - Optimistic	17.66	-11.03	28.69	4.29	-12.36	16.65
<i>Survey 1 (Economy Domain)</i>						
Pessimistic	51.44	23.59	27.84	14.07	10.75	3.32
Optimistic	29.85	38.60	-8.75	7.59	24.32	-16.73
Pessimistic - Optimistic	21.58	-15.00	36.59	6.49	-13.57	20.06
<i>Survey 2</i>						
Pessimistic	57.51	22.73	34.78	14.32	11.74	2.57
Optimistic	44.74	34.19	10.55	11.93	20.18	-8.25
Pessimistic - Optimistic	12.77	-11.46	24.24	2.39	-8.43	10.82

Table A9
Article Assessment (Survey 2)

Table A9 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 1 if the participant correctly identifies, among four headlines, the one of the article they have just read (column 3) the dollar amount spent donated to advertise the article source (column 4). The sample only includes participants who were not shown the source before choosing the article. Pessimistic is a dummy equal to 1 if the individual is 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. Preferred is a dummy equal to 1 if the participant is assigned the preferred article to read. 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 A7. 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)
Pessimistic	0.40** (0.16)	0.32** (0.16)	3.43 (4.02)	-0.99 (2.21)
Rev. After	0.13 (0.13)	0.09 (0.14)	-0.06 (3.21)	1.38 (1.83)
Liberal	-0.18 (0.19)	-0.28 (0.19)	1.30 (4.86)	-1.29 (2.55)
Pess. × Liberal	0.01 (0.23)	0.18 (0.23)	2.10 (5.28)	2.91 (3.23)
Pess. × Rev.	-0.46** (0.23)	-0.27 (0.24)	-0.38 (5.68)	0.24 (3.18)
Liberal × Rev.	-0.07 (0.20)	0.20 (0.21)	4.78 (4.73)	2.56 (2.84)
Liberal × Rev. × Pess.	-0.19 (0.33)	-0.48 (0.34)	-4.65 (7.38)	-16.81*** (4.58)
Observations	1,039	1,039	1,039	1,039
Adj. R ²	0.08	0.09	0.12	0.06
Mean of Dep. Var.	5.42	5.05	89.32	27.82
St. Dev. of Dep. Var.	1.30	1.37	30.91	18.14
Controls	YES	YES	YES	YES


A.4. Additional Figures

Figure A1

Article Assessment (Survey 1)

The histograms in Figure A1 plot the average of each of the variables indicated in the panels' titles for four participants' subgroups, depending on whether they chose the pessimistic or the optimistic article, and whether they were assigned the preferred or less preferred article. The variables are: a 1 to 7 measure of the article reliability, a 1 to 7 measure of informativeness, the time spent reading the article in minutes, the amount spent to advertise the news source, out of a \$50 budget, and a dummy equal to 1 if the individual identifies the headline of the article just read, out of four options. Panel (a) includes participants who chose the health domain and Panel (b) includes participants who chose the economy domain.

(a) Health Domain



(b) Economy Domain


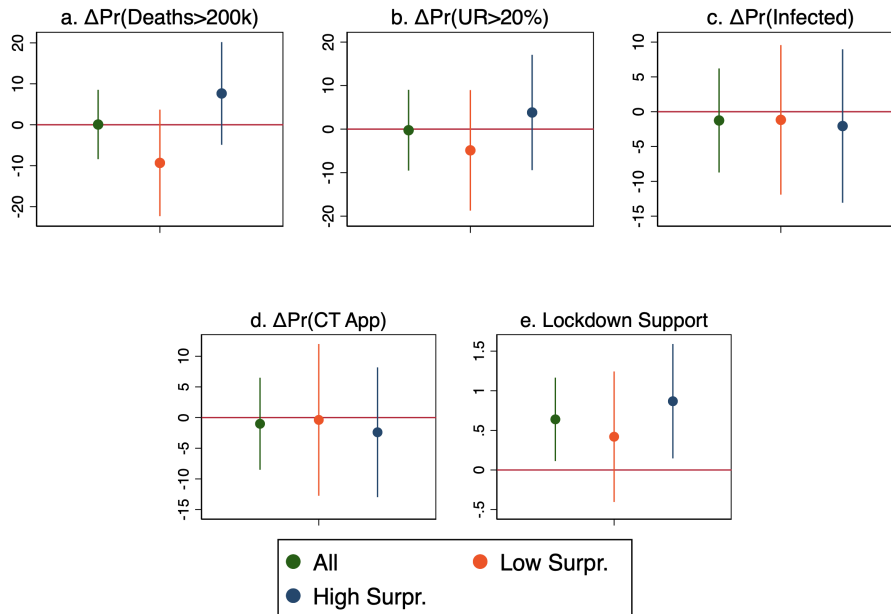


Figure A2
Additional Outcomes (Survey 2)

Figure A2 shows coefficients on the triple interaction term ($\text{Liberal} \times \text{Revealed After} \times \text{Pessimistic}$) from equation 4. 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 assigned comes from a conservative source. Individuals are then sorted into two halves, depending on this measure. Coefficients for individuals in the bottom half are in blue, together with confidence intervals; coefficients for individuals in the top half are in red.



A.5. Suggestive Evidence on Source Dependence and Polarization


In this Appendix we discuss how source revelation may lead to greater polarization in beliefs. This can be seen in the left-most panels of Figure A3 which show the mean revisions in the probability of deaths exceeding 200,000 for self-identified liberals and conservatives separately, for the four subgroups (whether they chose the pessimistic or the optimistic article, and whether the source was revealed). For liberals, when assigned the pessimistic article, the red bar shows that their beliefs turn more pessimistic (that is, are revised upwards) when they do not learn about the source, while they are revised down (become less pessimistic) when they learn that this information has been disseminated by Fox News (brown bar). On the other hand, the reverse patterns are found for conservatives: upward revisions are larger for conservatives assigned to the pessimistic article when they know the source of the article (brown bar) than when they do not (red bar). In fact, mean revisions are very similar for conservatives and liberals, who chose the pessimistic article, when the source is not revealed. However, dispersion in beliefs between liberals and conservatives (and, hence, polarization) increases when they know the source. It should be pointed out that the magnitudes are economically small (and none of the differences are statistically significant).

The increased polarization can also be seen in the mid-panels of Figure A3, which show the difference in disagreement (more specifically, standard deviation) in posterior beliefs for conservatives and liberals, conditional on choice of article, for those who were revealed the source and those who were not (that is, the difference in standard deviations of posterior beliefs is presented for those who are shown the source versus those who are not, conditional on article choice and political ideology). We see that disagreement goes up for those liberals who had chosen the pessimistic article when they are informed of the source (relative to those who were not). The opposite is the case amongst the conservatives who had chosen the pessimistic article: being informed that the article came from Fox News leads to less dispersion in beliefs in this subgroup (relative to their conservative counterparts who were not informed about the source). These patterns are reversed when we look at the set of respondents who had chosen the optimistic article: disagreement goes down (up) for liberals (conservatives) upon learning that the originator of the optimistic article was the New York Times, relative to when the source is not revealed. That is, being exposed to an article misaligned with one's political ideology leads to greater disagreement among respondents.



Figure A3
Belief Revision (Survey 2)

The histograms in Figure A3 plot the average of each of the variables indicated in the panels' titles for four participants' subgroups, depending on whether they chose the pessimistic or the optimistic article, and whether they were assigned the preferred or less preferred article. The variables are: a. the revision in the probability that the number of deaths will be higher than 200,000 by the end of the year; b. the difference in the standard deviations of the posterior death forecasts, for those shown the source and those not shown the source, conditional on article choice; c. the revision in lockdown support. Panel (a) includes participants who identify themselves as liberal, and Panel (b) includes participants who identify themselves as conservative.

(a) Liberals



(b) Conservatives



A.6. Headline Choice – Preliminary Tests

Prior to launching the surveys, we conducted pilot studies on Amazon Mechanical Turk to test people’s perceptions of the proposed headlines. In order to maximize statistical power of our analysis, we wanted the different headlines to be roughly equally attractive to respondents. We furthermore wanted to ensure that perceptions of how relatively optimistic/pessimistic the tone of different headlines was perceived by respondents aligned with our intuition.

Therefore, we fielded two short surveys, each to 100 participants, on Amazon Mechanical Turk. Each respondent was shown either the two health headlines or the two economy headlines for Survey 1 (in each case, the order in which the headlines were shown was randomized). We separately fielded a similar survey to 100 different participants for the two headlines for Survey 2. We then asked four questions comparing the two headlines. Subjects were paid \$0.10 upon completion. The four questions were:

Q1 Which of these two headlines do you think a typical news reader would find more interesting (meaning it will make them want to read the article)?

Q2 Which of these two headlines do you think a typical news reader would find more informative (telling them about the content of the article)?

Q3 Which of these two headlines do you think has a more optimistic/positive tone about the future?

Q4 Which of these two articles would you personally be more likely to read?

Subjects were required to answer on a 5-point scale, where 1 signifies “Definitely article A,” 5 signifies “Definitely Article B,” and 3 means “Both equally.” Below, for each survey and domain the average responses are shown, together with p -values for t -tests computed under the null hypothesis that the mean response is equal to 3 (that is, indifference between the two headlines). In what follows, A corresponds to the optimistic headline and B corresponds to the pessimistic headline (as noted above, in the actual survey subjects were shown the headlines in random order).

Question	Mean Response (<i>p</i> -value of null hypothesis that mean = 3)		
	Survey 1, Health	Survey 1, Economy	Survey 2
Q1	3.24 (0.085)	3.23 (0.075)	2.58 (0.003)
Q2	3.11 (0.419)	3.24 (0.062)	2.76 (0.048)
Q3	1.69 (0.000)	1.70 (0.000)	1.60 (0.000)
Q4	2.75 (0.868)	2.94 (0.677)	2.75 (0.073)

The average responses to questions Q1, Q2, and Q4 cluster quite closely around 3, ranging between 2.58 and 3.24. The null hypothesis of indifference between the two choices cannot be rejected at the 5% significance level, except for Q1 and Q2 of Survey 2. For Q3, instead, participants clearly rate headline A as more positive and optimistic, precisely as expected and intended.

To sum up, the tone differences is clearly perceived by respondents, while the headlines are perceived as comparable with respect to how interesting and informative they are.

A.7. Survey 1 - Motivated Beliefs

The survey starts by collecting information on demographics that were targeted by sample quotas, namely the region where the respondent lives, age, and gender. The demographics questions are followed by an attention test, in order to screen out bots or respondents who simply try to complete the survey as quickly as possible; such respondents would add considerable noise to our data.

Next, the survey solicits respondents' priors about their risk of infection, past and future, how worried they are, their "nowcast" and forecast of unemployment and Covid-related deaths, and their support for existing containment measures. Specifically, these survey questions read:

- *Do you think you may have already been infected with COVID-19 (also known as Coronavirus)?* (Options: Yes, definitely - confirmed by test; Very likely - did not get tested, but had symptoms; Possible - had some symptoms; Unlikely; Definitely not)
- *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?*
- *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?*
- *What do you think is the percent chance (or chances out of 100) that you will become infected by COVID-19 by the end of the year?* (Scale from 0 - Absolutely no chance - to 100 - Absolutely certain. Respondents were first provided with instructions regarding "percent chance" questions, closely following what is done in the Federal Reserve Bank of New York's Survey of Consumer Expectations.)
- *What is your best guess regarding the official total number of recorded deaths attributed to COVID-19 in the U.S. as of today?*
- *And what is your best guess regarding the official total number of deaths that will be attributed to COVID-19 in the U.S. by the end of the year?*
- *What do you think is the percent chance that the official total number of deaths that will be attributed to COVID-19 in the U.S. by the end of the year will be more than 200,000?* (Scale from 0 - Absolutely no chance - to 100 - Absolutely certain.)
- *What is your best guess of the current US unemployment rate?*
(The unemployment rate is defined as the number of unemployed people as a percentage of the labor force (which includes people who are either working or actively seeking work). For reference, as of February 2020, the US unemployment rate was 3.5%.)
- *What do you think the US unemployment rate will be at the end of this year?*
- *What do you think is the percent chance the US unemployment rate at the end of*

this year will be higher than 20%? (Scale from 0 - Absolutely no chance - to 100 - Absolutely certain.)

- *To reduce the spread of the virus, the US has implemented stay-home orders and social distancing measures. How do you evaluate the measures that have been taken? The measures...* (5-point scale, with labels 1 - ...went much too far, 3 - ...were about right, 5 - ...did not go nearly far enough.)

After soliciting these priors we propose to the respondents to choose among newspapers articles in the health or economic domain. Specifically, the survey text reads as follows:

- *Below, we will ask you further questions about your views regarding the health and economic consequences of the virus.*

However, before you do so, you will be given the chance to read a recent news article related to COVID-19.

You first have to decide whether you would prefer the article to be related to the health impacts OR the economic impacts of COVID-19. If given the choice, which kind of article would you prefer to read?

Note that both types of article have approximately the same length.

(Options: I would prefer an article related to the health impacts; I would prefer an article related to the economic impacts)

Independently of the domain choice, each respondent then is asked to choose their preferred headline for articles from both domains:

- *Here are two news articles on the economic impacts of COVID-19 that you could read. Both articles are from major U.S. news sites/sources (either top 10 newspaper by circulation, or top 3 news channel organization by viewership). Based on the headlines, which article would you prefer to read?* (Options (in random order): Hope for swift economic recovery builds as businesses reopen; Highest unemployment rate since the Depression era, and many jobs may not come back.)

- *Here are two news articles on the health impacts of COVID-19 that you could possibly read. Both articles are from major U.S. news sites/sources (either top 10 newspaper by circulation, or top 3 news channel organization by viewership). Based on these headlines, which article would you prefer to read?* (Options (in random order): New data suggest the coronavirus is less deadly than we thought; Two influential forecasting models predict sharp rise in coronavirus deaths.)

Following their choice, the assignment of articles is randomized, so that some respondents get their desired article and some do not (though they all get the article from the domain they had ranked as preferred). The survey text reads as follows:

- *The computer has randomly chosen which article you will get to see. You will get to see the article with the headline "..."*

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.

After respondents have read the article, we test whether they remember the headline of the article they read (asking a multiple choice question with the four article headlines as possible choices). We also ask them whether they found the article informative and reliable:

- *How informative did you find the article you just read?* (7-point scale from 1 – Not informative at all – to 7 – Very informative.)
- *Does this information feel reliable to you?* (7-point scale from 1 – Not reliable at all – to 7 – Completely reliable.)

They are also asked what they think the percent chance is that the article they just read is from each of the following news sources: Breitbart, CNN, Fox News, MSNBC, New York Times, USA Today, Wall Street Journal, Washington Post, Other news sources. Respondents were asked to enter a number greater than or equal to 0 for each news source, with the answers needing to add up to 100.

We further implement a “revealed preference” measure to assess their view of the article. The survey text reads as follows:

- *Consider the situation where you have the option to advertise the news outlet whose article you just read. By paying 50 dollars, the news outlet will be advertised roughly 5,000 times on Facebook. You can decide how much to spend on advertising the news outlet. This can be anywhere between zero and 50 dollars. We will pay for this, so there is no cost to you.*
Note that we will pick 2 survey participants at random and implement their choice. If you are picked, we will actually spend the amount you chose to advertise the news outlet. At the end of the survey, we will give you further details on how to check whether you were picked and your choice was implemented.
How much (out of 50 dollars) would you like to go towards advertising the news outlet? Any amount that is not donated will be lost.

In the final part we re-elicite their (posterior) beliefs about both health and economic effects of the pandemic: using the same wording as above, we again ask for point forecasts of the total number of deaths that will be attributed to COVID-19 in the U.S. by the end of the year and the US unemployment rate at the end of the year; and probabilities (percent chance) that the official total number of deaths attributed to COVID-19 in the U.S. will be more than 200,000 by the end of the year; the US unemployment rate at the end of the year will be higher than 20 percent; and that the respondent will become infected by COVID-19 by the end of the year.

These questions are followed by a question on the preferred policy response to a possible new outbreak of the pandemic:

- *Assume that in the coming weeks, the spread of the virus begins to slow, and the stay-home (lockdown) orders are lifted across the US. However, suppose that in the fall of this year, infections start increasing again, in a way similar to what happened earlier this year. If that happens, what should politicians do?* (Options: 1 - Stricter stay-home orders

than this time, even if the economic cost is very large; 2 - Slightly stricter stay-home orders than this time; 3 - Same stay-home orders as this time; 4 - Weaker stay-home orders than this time; 5 - No intervention, just let the virus run its course.)

Other demographic questions follow, on education, current employment status, whether anyone in the household experienced a loss of employment income since February 2020, ethnicity/race, health status (in general and whether they have three specific conditions), whether they have health insurance, political affiliation/leaning (incl. candidate they will likely vote for in presidential election), the news sources they usually consult about COVID-19, marital status, ZIP code, ownership of primary residence, number of households members in different age groups, current status of stay-home orders in respondent's location, stock ownership, qualitative measures of risk aversion, patience, and general trust, and their households' income bracket.

The survey closes with a second revealed preference check, which reads as follows:

- *Among all the respondents to this survey, 10 will be randomly picked. Each of these respondents gets to allocate a \$20 donation to two organizations that have been active in the COVID-19 policy discussion. The two organizations are:*

The World Health Organization - It 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." (Source: Reuters)

FreedomWorks - It 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."

Please select how you would like to allocate \$20 between the two organizations. If you get randomly selected, then the amounts you choose will be (anonymously) donated according to your selection below.

(Answers: \$ to World Health Organization; \$ to FreedomWorks; both numbers need to be in [0,20] and the total of the two numbers needs to sum to 20.)

A debriefing screen then provides the respondent with an ID number and a link to a website where they can check whether their ID number was drawn for either of the incentivized choices, and where we posted receipts of the resulting donations once the data collection was complete.

The debriefing screen also contains links to all four articles that a respondent could have seen, and links to official guidelines regarding COVID-19, from the World Health Organization and the Centers for Disease Control and Prevention.

Links to the original articles and to their edited versions are the following:

– **Optimistic Economy:**

Original – *Why the stock market is up even with historic job losses* (CNBC, May 8, 2020): <https://web.archive.org/web/20200508203712/https://www.cnbc.com/2020/05/08/why-the-market-is-up-even-with-historic-job-losses.html>

Edited – *Hope for swift economic recovery builds as businesses reopen*: https://safefrankfurt.fra1.qualtrics.com/jfe/form/SV_a8En9PrELt96vsx

– **Pessimistic Economy:**

Original – *U.S. unemployment rate soars to 14.7 percent, the worst since the Depression era* (The Washington Post, May 8, 2020): https://web.archive.org/web/20200518135858if_/https://www.washingtonpost.com/

Edited – *Highest unemployment rate since the Depression era, and many jobs may not come back*: https://safefrankfurt.fra1.qualtrics.com/jfe/form/SV_7a0VEkgx7WWfiLj

– **Optimistic Health:**

Original – *New Data Suggest the Coronavirus Isn't as Deadly as We Thought* (The Wall Street Journal, April 17, 2020): <https://web.archive.org/web/20201125031043/https://www.wsj.com/articles/new-data-suggest-the-coronavirus-isnt-as-deadly-as-we-thought-11587155298>

Edited – *New data suggest the coronavirus is less deadly than we thought*: https://safefrankfurt.fra1.qualtrics.com/jfe/form/SV_egQspiwTDZMTgI1

– **Pessimistic Health:**

Original – *Coronavirus model projects 134,000 deaths in US, nearly double its last estimate* (CNN, May 5, 2020): <https://web.archive.org/web/20201125023601/https://edition.cnn.com/2020/05/04/health/us-coronavirus-monday/index.html>

Edited – *Two influential forecasting models predict sharp rise in coronavirus deaths*: https://safefrankfurt.fra1.qualtrics.com/jfe/form/SV_bsCw2Cut7jys27b

A.8. Survey 2 - Source Dependence

The first part of this second survey mirrors Survey 1. Specifically there are demographic questions, followed by an attention check and by questions to elicit respondents' worries and priors on health and economics risks. Relative to the first survey, we added this question to the first part:

-At the moment, apps and systems are being developed for "contact tracing": identifying and notifying all those who come in contact with a COVID-19 carrier. What is the percent chance that you would voluntarily sign up for such a service and provide information about yourself?

Then, respondents were informed that they would get to read a recent news article related to COVID-19, and were asked to choose which article they would like to read based on the headlines (shown in random order):

Scientists increasingly optimistic that a vaccine can be produced in record time

Coronavirus pandemic to likely last two years, new report says

Contrary to the previous survey we now randomly reveal to some of the respondents the source of the article, which is the New York Times for the first article and Fox News for the second article. Respondents are randomly allocated to one of three equal-sized groups. In group T1, the source is never revealed; respondents are simply told "Both articles are from major U.S. news sites/sources (either top 10 newspaper by circulation, or top 3 news channel organization by viewership)." They are also asked for the percent chance that the article they selected is from one of nine news sources (like in Survey 1). In group T2, the source of the article is revealed when the article is shown to the respondent; that is, after the article was chosen (based on the headline) and the respondent provided their assessment of where the article likely came from. In group T3, the source of the article is revealed before the article choice; it is presented alongside the headline (e.g., "Coronavirus pandemic to likely last two years, new report says," published by Fox News").

After the article is shown, the survey again proceeds like Survey 1: we check whether respondents remember the headline they just read; ask for their ratings of informativeness and reliability; and give them the possibility to allocate between \$0 and \$50 to Facebook ads (at no cost to them) for the news source that had published the article they just read (with this source unknown to T1 but known to T2 and T3). Then, we elicit the same posterior forecasts and support for future stay-home orders as in Survey 1. In addition, we re-elicited the probability they would sign up for a contact-tracing app.

Finally, the same additional demographics were elicited as in Survey 1, along with a few additional questions about (current and intended) pandemic-related behaviors:

-How often do you wear a mask when outside? (Answer on a 7-point scale with labels 1 - Never - to 7 - All of the time.)

-Which of the following precautionary measures have you taken in the last one month? Please select all that apply. (Options: I have attended fewer social gatherings; I have reduced the number of trips I make outside of my house; I keep a distance of at least two meters (6 feet) with people outside my home; I wear a mask to cover my mouth and nose when I leave my home; I have reduced the number of times I attend religious gatherings

(church etc.).)

Unlike in Survey 1, there was no donation allocation question (WHO vs. Freedom-Works) in Survey 2. The survey ended with a similar debriefing screen as Survey 1.

Links to the original articles and to their edited versions are the following:

– **Optimistic:**

Original – *A New Entry in the Race for a Coronavirus Vaccine: Hope* (The New York Times, May 20, 2020): <https://web.archive.org/web/20201212032018/https://www.nytimes.com/2020/05/20/health/coronavirus-vaccines.html>

Edited – *Scientists increasingly optimistic that a vaccine can be produced in record time*: https://safefrankfurt.fra1.qualtrics.com/jfe/form/SV_6nAGcJr80T8iCKF

– **Pessimistic:**

Original – *Coronavirus pandemic could last 2 years, 70 percent of world population needs immunity: researchers* (Fox News, May 1, 2020): <https://web.archive.org/web/20200607161457/https://www.foxnews.com/science/coronavirus-pandemic-could-last-2-years-70-percent-of-the-worlds-population-need-immunity-researchers>

Edited – *Coronavirus pandemic to likely last two years, new report says*: https://safefrankfurt.fra1.qualtrics.com/jfe/form/SV_bKP30e81hPbw3Vb

Recent Issues

No. 306	Aljoscha Janssen, Johannes Kasinger	Obfuscation and Rational Inattention in Digitalized Markets
No. 305	Sabine Bernard, Benjamin Loos, Martin Weber	The Disposition Effect in Boom and Bust Markets
No. 304	Monica Billio, Andrew W. Lo, Loriana Pelizzon, Mila Getmansky Sherman, Abalfazl Zareei	Global Realignment in Financial Market Dynamics: Evidence from ETF Networks
No. 303	Ankit Kalda, Benjamin Loos, Alessandro Previtero, Andreas Hackethal	Smart (Phone) Investing? A Within Investor-Time Analysis of New Technologies and Trading Behavior
No. 302	Tim A. Kroencke, Maik Schmeling, Andreas Schrimpf	The FOMC Risk Shift
No. 301	Di Bu, Tobin Hanspal, Yin Liao, Yong Liu	Risk Taking, Preferences, and Beliefs: Evidence from Wuhan
No. 300	Dennis Gram, Pantelis Karapanagiotis, Jan Krzyzanowski, Marius Liebald, Uwe Walz	An Extensible Model for Historical Financial Data with an Application to German Company and Stock Market Data
No. 299	Ferdinand A. von Siemen	Motivated Beliefs and the Elderly's Compliance with COVID-19 Measures
No. 298	Calebe de Roure, Emanuel Moench, Loriana Pelizzon, Michael Schneider	OTC Discount
No. 297	Dimitrios Kostopoulos, Steffen Meyer, Charline Uhr	Ambiguity and Investor Behavior
No. 296	Reint Gropp, Thomas Mosk, Steven Ongena, Ines Simac, Carlo Wix	Supranational Rules, National Discretion: Increasing Versus Inflating Regulatory Bank Capital?
No. 295	Besart Avdiu, Alfons J. Weichenrieder	Financing Costs and the Efficiency of Public-Private Partnerships
No. 294	Christian Alemán, Christopher Busch, Alexander Ludwig, Raül Santaeulàlia-Llopis	Evaluating the Effectiveness of Policies Against a Pandemic