Manuscript version: Working paper (or pre-print)

The version presented here is a Working Paper (or 'pre-print') that may be later published elsewhere.

Persistent WRAP URL:

http://wrap.warwick.ac.uk/139675

How to cite:

Please refer to the repository item page, detailed above, for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

warwick.ac.uk/lib-publications

Do People Value More Informative News?

Felix Chopra, Ingar Haaland & Christopher Roth

July 2020　　　　　　　　No: 1285

Warwick Economics Research Papers

Do People Value More Informative News?[*]

Felix Chopra[†]   Ingar Haaland[‡]   Christopher Roth[§]

July 18, 2020

Abstract

Drawing on representative samples of the U.S. population with more than 15,000 respondents in total, we measure and experimentally vary people's beliefs about the informativeness of news. Inconsistent with the desire for more information being the dominant motive for people's news consumption, treated respondents who think that a newspaper is less likely to suppress information reduce their demand for news from this newspaper. Furthermore, treated respondents who think that a news outlet is more likely to make false claims do not reduce their demand for this outlet. These findings strongly suggest that people have other motives to read news that sometimes conflict with their desire for more information. We discuss the implications of our findings for the regulation of media markets. (*JEL* D83, D91, L82)

Keywords: News Consumption, Information, Media Bias, Belief Polarization, Informativeness.

---

1 Introduction

What motivates people to consume economic and political news? Consistent with a core principle in standard economics—that more information is always better—Americans cite getting the facts right as their single most valued factor when choosing a specific news source (Young, 2016). While an overwhelming majority of Americans also say that the news media should be unbiased in its coverage of political issues (Mitchell, 2018), a large literature has documented that newspapers report news in a biased way by slanting their news stories toward the beliefs of their readers (Gentzkow and Shapiro, 2010).

There are two main competing explanations for why people consume biased news. The first explanation is that people want more informative news but perceive news that are closer to their prior beliefs as more informative (Gentzkow and Shapiro, 2006). Consistent with this explanation, both liberals and conservatives consider politically-aligned news outlets as more trustworthy (Mitchell and Weisel, 2014). The second explanation is that people have other motives to read news that sometimes conflict with expanding their knowledge, such as a preference for belief confirmation (Mullainathan and Shleifer, 2005).

The relative importance of these two explanations has important implications for the optimal regulation of media markets, such as the welfare effects of regulations to increase competition. Despite its relevance for optimal policy, there is no direct evidence on whether the desire for more information dominates other motives for reading news. Previous studies on news consumption cannot identify whether the desire for more information dominates other motives because perceptions of informativeness are unobserved in observational data. To circumvent this challenge, we propose an experimental approach that allows us to measure and exogenously change perceptions about the informativeness of news.

In a series of experiments with over 15,000 Americans, we use two complementary designs to achieve exogenous variation in the perceived informativeness of news articles. We then measure news demand using real articles on economic and political news. The experiments were designed to change beliefs about informativeness without changing beliefs about the complexity or technicality of reporting. The experiments thus allow us to

test whether people value more informative news in a setting where cognitive constraints are not binding. Our criterion for comparing the informativeness of news outlets is Blackwell's (1951) ranking of information structures—the benchmark for evaluating the information content of signals. While it is usually not possible to compare the Blackwell informativeness of different news articles, we designed the experiments such that the news articles should be perceived as strictly more or less Blackwell informative by treated respondents. We use a simple model to outline the theoretical foundation for the predictions of the more-information-is-better principle in both experimental designs.

In the main experiment, we generate exogenous variation in the perceived informativeness of news by informing treated respondents that the *New York Times* did not strategically suppress information from a news article. Specifically, we first tell our respondents that the Congressional Budget Office (CBO), Congress's official nonpartisan provider of cost and benefit estimates for legislation, published a report about the "Trump Healthcare Plan" (the American Health Care Act of 2017). Respondents are told that the CBO estimated that the plan would decrease the federal deficit by $119 billion (contradicting claims made by Democrats) and leave 23 million more people uninsured (contradicting claims made by Republicans). To elicit beliefs about the informativeness of news, we then tell respondents that the *New York Times* wrote an article about the CBO report and ask about the subjective percent chance that it only reported the statistic on the federal deficit, only the statistic on the number of uninsured, or both statistics. This allows us to quantify beliefs about the informativeness of news articles in the *New York Times*: Reporting both statistics is strictly more informative (in the Blackwell sense) than selectively reporting only information that favors one of the parties. To generate exogenous variation in perceptions of informativeness, we then inform treated respondents that the *New York Times* reported both statistics from the CBO report. To measure how the information affects the demand for news, we offer all respondents free access to an article in the *New York Times* covering a CBO report on a different topic, namely the "Trump Tax Plan" (the Tax Cuts and Jobs Act of 2017). We also ask a series of belief questions to shed light on mechanisms.

The treatment generates a strong and significant effect on perceptions of the informa-

tiveness of news in the *New York Times*: Treated respondents are 6.8 percentage points more likely to think that the *New York Times* did not suppress a key statistic contradicting claims by Democrats in a news article about the CBO report on the Trump Tax Plan. The treatment also affects more general perceptions of the *New York Times*: Treated respondents are less likely to think that it is politically biased and more likely to think that it provides high-quality news. Given our treatment effects on perceptions, the more-information-is-better principle predicts a strict increase in the demand for news from the *New York Times*. In stark contrast to this prediction, the main result of the paper is that respondents who learn that the *New York Times* does not suppress information significantly *reduce* their demand for news from this newspaper by 3.5 percentage points. This corresponds to a 13 percent reduction in the demand for news.

Examining heterogeneity in treatment responses, we find that the negative treatment effect is mainly driven by respondents who thought the *New York Times* was more likely to suppress facts contradicting claims made by their own political party, a result that is broadly consistent with people having a preference for belief confirmation. We further run a series of mechanism experiments to rule out alternative explanations of our main finding. First, we run a placebo experiment suggesting that cognitive constraints are not binding in our setting. In this experiment, we inform respondents that the CBO highlighted two key statistics in a report on healthcare without providing any further information about the content of the report. Treated respondents are then informed that the *New York Times* reported both key statistics in its coverage of the CBO report. In this experiment, where other potentially conflicting psychological motives, such as a preference for belief confirmation, arguably play no role, we do not observe a decrease in the demand for news. Second, we run a replication experiment that uses a different set of news articles and also addresses concerns about differential curiosity across treatment arms about how the *New York Times* covers CBO reports. In this experiment, we find a 4.7 percentage point decrease in demand for news—an even larger effect size than in the main experiment. Third, we run additional experiments to show that rational mechanisms cannot explain the main result.

In the second main experiment, we provide another test of the more-information-is-better principle by studying how perceptions of false claims affect the demand for news. The design complements the first main experiment by changing perceptions of a different source of variation in informativeness and by studying a different news outlet. Specifically, we first tell our respondents that the U.S. recently extracted a high-level spy from Russia. We then inform our respondents that some media outlets correctly reported that the CIA extracted the spy because of widespread media speculation about its sources, whereas other media outlets falsely reported that CIA extracted the spy because it feared that President Trump would mishandle classified information.[1] To elicit pre-treatment beliefs about the likelihood of false claims made by *CNN*, we ask respondents about their subjective percent chance that *CNN* reported that the CIA extracted the spy due to media speculation about its sources or due to concerns about President Trump. To generate exogenous variation in perceptions of false claims, we then inform treated respondents that *CNN* reported that the reason for the extraction was concerns about President Trump. To measure how the information affects the demand for news, we offer all respondents free access to a news article in *CNN* about the impeachment process against President Trump. Finally, we ask a series of questions that shed light on mechanisms.

The treatment generates a highly significant first stage on perceptions of false claims made by *CNN*. Treated respondents are more likely to think that *CNN* articles about President Trump contain false claims, more likely to think that *CNN* generally makes false claims in its political reporting, and more likely to think that *CNN* publishes stories about President Trump based on unverified and potentially misleading sources. The treatment also affects more general perceptions of *CNN*: Treated respondents display lower trust in *CNN*, think its news articles are of lower quality, and think it is more likely to be politically biased. Given these results, the more-information-is-better principle predicts a strict decrease in the demand for news from *CNN*. The main result from this experiment is that despite a strong first stage on perceptions of false claims by *CNN*, treated respondents

---

[1] All information presented to the respondents was truthful. For more background, see the following article in Associated Press News: `https://www.apnews.com/b432a20d85ca48a5bca6c3394920c7fe` (accessed October 2, 2019) and our discussion in Section 4.

do not reduce their demand for news from *CNN*.

Taken together, the results from the two main experiments suggest that some of the motivations people have to read news directly conflict with the desire for more information. One such motive could be a preference for belief confirmation (Mullainathan and Shleifer, 2005), which leads people to avoid news that is less likely to confirm their prior beliefs. A second such motive could be a desire for entertainment, which could lead people to seek news that are more likely to contain extreme and surprising facts (Ely et al., 2015)—even if the underlying facts are less likely to be true.

Our results contribute to the literature on media bias (DellaVigna and Ferrara, 2015; DellaVigna and Hermle, 2017; Gentzkow et al., 2018; Jo, 2019; Mullainathan and Shleifer, 2005; Perego and Yuksel, 2018; Pogorelskiy and Shum, 2019; Qin et al., 2018; Szeidl and Szucs, 2017) and fake news (Allcott and Gentzkow, 2017; Lazer et al., 2018; Vosoughi et al., 2018). The focus of this literature has largely been on the supply side of media bias (Gentzkow and Shapiro, 2006, 2010; Gentzkow et al., 2014) or the persuasive effect of biased news on voting decisions (DellaVigna and Kaplan, 2007; Enikolopov et al., 2011), and there is only indirect evidence on the demand side. A common approach for studying the demand side of media bias is to use aggregate data on news demand and examine how different political groups adjust their demand for news in response to changes in media bias (Durante and Knight, 2012; Garz et al., 2018; Gentzkow and Shapiro, 2010). These studies are consistent with people having other motives to read news that conflict with their desire for more information, such as a preference for like-minded news, but are also consistent with standard preferences and unobserved differences in the perceived informativeness of news. To reject that people have standard preferences as in Gentzkow and Shapiro (2006), these studies therefore require model-based identification assuming that people do not change their beliefs about the informativeness or quality of news in response to changes in media bias. The main contribution of this paper is to provide the first direct evidence on whether the desire for more information dominates other motives to read news. Our main results—that respondents who think that a newspaper is less likely to suppress information reduce their demand for news and that respondents who think that

a news outlet is more likely to make false claims do not reduce their demand—cannot be rationalized with standard preferences that would imply a demand for more informative news.

Furthermore, we also contribute to a small but growing literature on people's demand for information and information avoidance (Charness et al., 2019; Falk and Zimmermann, 2017; Ganguly and Tasoff, 2016; Golman et al., 2017; Nielsen, 2019; Zimmermann, 2015).[2] Studying the demand for more informative news is of particular importance because of the political externalities of news consumption, including its effects on political accountability, electoral efficiency, and political polarization (Strömberg, 2015; Sunstein, 2018). Understanding whether people demand more informative news is also of critical importance for the debate on whether policy makers should introduce regulations to increase competition in media markets (Gentzkow and Shapiro, 2006; Mullainathan and Shleifer, 2005). Our key contribution to this literature is thus to provide clean evidence on information avoidance in the context of news consumption. To identify information avoidance, we employ a new identification strategy by varying perceptions about the accuracy of the signal provided by a news outlet.[3] In contrast to much of the previous experimental literature on information avoidance, we vary perceptions about a real-world source rather than features of an artificial lab environment.

The remainder of the paper proceeds as follows. Section 2 provides a simple theoretical framework that formalizes the more-information-is-better principle and relates it to our main experimental design. Section 3 covers the main experiment on strategic suppression of information in the *New York Times* in addition to several robustness and mechanisms experiments. Section 4 covers the second main experiment on false claims in *CNN*. Section 5 discusses the implications of our findings for models of news consumption. Section 6 concludes and discusses implications for the regulation of media markets. The Online Appendix provides additional theoretical and empirical results and the full set of

---

[2]More broadly our evidence relates to a relatively large literature on motivated belief updating (Exley, 2015; Exley and Kessler, 2018; Schwardmann and van der Weele, 2019; Di Tella et al., 2015; Thaler, 2019).

[3]We thus also contribute to a literature on information provision experiments. For a review of this literature, see Haaland et al. (2020).

experimental instructions.

2 Theoretical framework: Filtering

We now present a simple framework that formalizes the implications of strategic information suppression of news outlets for the Blackwell informativeness of news in our empirical design. This provides us with a theoretical benchmark for how learning that a newspaper is less likely to strategically suppress information should affect the demand for news according to the more-information-is-better principle.

There is a binary state space $\Theta = \{L, R\}$ with a typical element denoted by θ and an agent with prior belief $q \in \Delta(\Theta)$ about the hidden state. The agent has the option to acquire information from a newspaper. The newspaper provides information about θ by publishing an article $n \in N$ whose content is revealed only upon acquiring it. To introduce scope for information suppression, we assume that the newspaper receives a set of private signals $s = \{s_1, \ldots, s_K\} \in S$ from its information source. The set consists of K binary bits of information $s_i \in \Theta$ about the state of the world θ, where K is randomly drawn and independent of θ. The individual bits, s_i, are drawn from a state-dependent distribution, F_θ. We assume that F_L places higher weight on L compared to F_R, implying that s_i is informative about θ. The source signal can thus be represented as an information structure (S, π) with state-dependent likelihood $\pi : \Theta \rightarrow \Delta(S)$. In our main empirical design, the CBO is the newspaper's source, providing two conflicting bits $s = \{L, R\}$ about the desirability, θ, of the Trump Healthcare Plan.

The newspaper can disclose any subset of s in its article n, i.e. $n \subseteq s$, implying that it cannot distort individual bits. Information suppression occurs whenever $n \neq s$. We are agnostic about the newspaper's incentives to suppress information, subsuming them in the reader's belief $\rho : S \rightarrow \Delta(N)$ about how the newspaper reports conditional on s.

From the agent's perspective, the informativeness of an article n should be an invariant of the state-dependent distribution over news articles, $\sigma : \Theta \rightarrow \Delta(N)$, induced by the

agent's belief about the quality of the newspaper's source, π, and the belief about how the newspaper reports, ρ. Specifically, consider two articles n and n' with distributions $\sigma, \sigma' : \Theta \to \Delta(N)$. We use Blackwell's (1951) notion of informativeness and say that n is *(Blackwell) more informative* than n' if (n, σ) is *sufficient* for (n', σ'), that is: there is a stochastic transformation τ such that n' and $\tau(n)$ are identically distributed. Intuitively, we obtain n' by adding noise to n. This is the benchmark for evaluating the informativeness of an information structure: any agent with access to an article n that is more informative than n' can attain an expected payoff at least as large as the maximal expected payoff attainable with n', regardless of the prior q and the decision problem $a \in A$ with payoffs $u(a, \theta)$ (Blackwell, 1953). This provides the prediction that the demand for news should be strictly increasing in the perceived informativeness of the news.

How does strategic suppression affect the informativeness of news? Suppose the newspaper received the signals $s = \{s_1, \ldots, s_K\}$ and let $\sigma(s' \mid s)$ denote the agents' belief that the newspaper would report $s' \subseteq s$ after receiving s. Intuitively, the informativeness of the article n should be increasing in the probability of fully conveying the signal. Indeed, the Blackwell informativeness of an article strictly increases if we decrease the probability $\sigma(s' \mid s)$ of reporting a filtered signal $s' \subsetneq s$ and instead increase the probability of full information transmission, $\sigma(s \mid s)$.[4]

This provides us with an empirical test of the more-information-is-better principle. Specifically, we leverage an information treatment that decreases respondents' expectation that the *New York Times* strategically filters information from CBO reports and increases their expectation that the *New York Times* reveals all information from the report. By the previous discussion, treated respondents should perceive articles from the *New York Times* as more informative, especially for articles covering reports from the CBO. If the desire for more information is the dominant motive for economic and political news consumption, the theoretical benchmark prediction is that treated respondents should strictly increase their demand for news. However, if motives that conflict with becoming better informed dominate this desire, we would expect the demand to decrease (Mullainathan and Shleifer,

---

[4]See Proposition 1 in the Online Appendix for a proof.

2005).

3 Filtering experiment

3.1 Sample

We collected the data for the main filtering experiment in three waves in collaboration with *Dynata*, formerly called *Research Now SSI*, a leading market research company commonly used in social science research (de Quidt et al., 2018; Enke, 2020). We have a sample of 4,631 respondents that is broadly representative of the U.S. population in terms of education, age, income, region, gender and political affiliation (column 1 of Table C.1). The treatment and control group are balanced in terms of observable characteristics (Table C.2).

3.2 Experimental design

This section outlines our experimental design. Figure 1 provides a summary of the structure and Section E of the Online Appendix provides the full experimental instructions.

Pre-treatment characteristics and beliefs We first measure basic demographics, namely income, age, gender, and region of residence. Furthermore, we ask for people's political preferences and beliefs, how often they read the *New York Times*, and the three newspapers they are most likely to read. Thereafter, we elicit people's beliefs about how the *New York Times* reports about the Trump Healthcare Plan. Specifically, we tell our respondents that the Congressional Budget Office (CBO), Congress' official nonpartisan provider of cost and benefit estimates for legislation, published a report about the Trump Healthcare Plan(the American Health Care Act of 2017). Respondents are told that the CBO estimated that the plan would decrease the federal deficit by $119 billion (contradicting claims made by Democrats) and leave 23 million more people uninsured (contradicting claims made by Republicans). Subsequently, we measure respondents' beliefs about how the *New York*

Times covered the CBO report by asking our respondents to estimate the percent chance that the *New York Times* reported only the figure on the number of uninsured people, only the figure on the deficit decrease, or both figures.

We chose to focus on the *New York Times* for two main reasons. First, the *New York Times* is a well-known newspaper with a national coverage. Second, it tends to lean toward the Democratic Party (for instance, it has consistently supported Democratic candidates for president since 1960). Choosing a newspaper with a clear partisan stance ensures that beliefs about political reporting of the newspapers are shifted in the same direction for the majority of respondents. Furthermore, we focused on the *New York Times*'s reporting strategy about news from the CBO for the following reasons: First, the CBO is Congress's official provider of cost and benefit estimates for legislation and is known to be nonpartisan (to stay politically neutral, it only assesses the consequences of proposed policies and does not make its own policy recommendations). Second, all major newspapers in the U.S. generally feature CBO reports in their news reporting.

Information treatment   We provide a random subset of respondents with information about the *New York Times* (treatment group).[5] Specifically, we provide treated respondents with the following information treatment, which is framed in a neutral way to minimize experimenter demand effects:

> In its article about the CBO estimates, The New York Times reported **both** that the federal budget deficit would decrease by $119 billion **and** that the number of people without health insurance would increase by 23 million.

Respondents in the control group proceed without receiving any information.

Post-treatment outcomes   To mitigate concerns about consistency bias in survey response (Falk and Zimmermann, 2012), a subset of respondents of Experiment 1.1 are

---

[5]We stratify the assignment into treatment and control group by whether respondents identify as Republicans, Democrats, or Independents.

cross-randomized to receive either (i) a question on the demand for news (2,250 respondents) or (ii) the post-treatment beliefs block (755 respondents).[6] In all other experiments, all respondents proceed to the question on demand for news.

Main outcome: Article demand   We collect a behavioral outcome measure on people's demand for news by providing them with an opportunity to read an article from the *New York Times*. This article is unrelated to the Trump Healthcare Plan and instead covers a CBO report about a different policy: the Trump Tax Plan. However, we do not provide any additional information about the content of the article or the corresponding CBO evaluation of the Trump Tax Plan. We make this distinction salient to respondents, thus ensuring that respondents expect to receive an article containing new information not previously mentioned in the survey. Specifically, we tell respondents that the CBO analyzed the consequences of the Trump Tax Plan over the next decade and ask them whether they want to read an article about its findings in *New York Times*. We tell respondents that if they decide not to receive access to the article, they will proceed with the survey without receiving access to the article. If they decide to receive access to the article, they will receive access at the end of the survey. We thus decrease the cost of accessing the *New York Times* article both in terms of search costs and in terms of avoiding the *New York Times* paywall.[7]

There are several reasons why we choose this as our main outcome. First, the decision on whether or not to read a real news article in the *New York Times* has high external validity as most online news consumption decisions are low stakes in nature. Second, our setting allows us to hold some beliefs about article characteristics across the treatment and control group constant. For instance, to fix perceptions of cognitive effort required to read the article across treatment arms, we directly tell respondents that the article contains about 1,100 words. Third, by embedding the article in our online survey, we can measure

---

[6]The main experiment was conducted in three waves. We only cross-randomized respondents into the beliefs block in the second wave. See Table 1 for more details.

[7]While the *New York Times* is technically behind a paywall, it offers non-subscribers free access to 10 articles per month. If our respondents were perfectly aware of this, they could thus save one of their 10 free articles if they decided to receive access through our survey rather than visiting the *New York Times* directly.

not only the extensive margin, that is, whether people want to read the news article, but also the intensive margin, that is, how much time they spend reading the news article.

Post-treatment belief I: Filtering   To study whether our treatment intervention affected people's beliefs about informativeness, we collect a post-treatment measure of strategic information suppression, but only for the subset of respondents who were randomized not to receive the demand for news block. To do so, we provide respondents with information about cost and benefit estimates from a CBO report regarding the Trump Tax Plan. We tell respondents about the opposing predictions made by Republicans and Democrats about the plan's impact on the federal debt and job creation. To avoid consistency bias in survey responses, we employ a different way of measuring beliefs about information suppression compared to the elicitation of prior beliefs. We inform respondents that the *New York Times* reported that the Trump Tax Plan would increase the federal debt as claimed by Democrats. We then ask our respondents to estimate the percent chance that the *New York Times* also reported that the Trump Tax Plan would create 1.1 million jobs.

Post-treatment belief II: Omission   Furthermore, we measure beliefs about the extensive margin of political news coverage by eliciting people's beliefs about whether the *New York Times* strategically decides not to publish articles about certain CBO reports. This allows us to test for treatment effects on beliefs about news coverage that are conceptually distinct from within-article filtering of information. Specifically, we ask respondents to estimate the percent chance that the *New York Times* wrote any article at all about a CBO report estimating that a signature policy proposed by Democrats would add $27 billion to the federal debt. This signature policy would grant citizenship status to 1.8 million young undocumented immigrants (known as the Dreamers), and we inform respondents that Democrats claimed that it would not increase the federal debt.

Post-treatment beliefs III: Article characteristics   We also measure additional beliefs about (i) the quality of news articles in the *New York Times*, (ii) whether the *New York*

Times article about the Trump Tax Plan will be dry and technical, and (iii) whether the article about the Trump Tax Plan will be complex.

Additional beliefs and demographics   We also separately elicit people's perception of whether the *New York Times* and the CBO are politically biased, their trust in the *New York Times* and the CBO, and general trust in the media. Furthermore, we measure beliefs about the accuracy of CBO forecasts. Finally, we ask some additional demographic questions.

3.3   Main results

In this section, we study the causal effect of learning that a newspaper is less likely to suppress information on beliefs and demand for news.

3.3.1   Empirical specification

Our main empirical specification for different outcomes, y_i, is given as follows:

$$y_i = \alpha_0 + \alpha_1 T_i + \alpha_2 \mathbf{x}_i + \varepsilon_i \tag{1}$$

where T_i is an indicator for whether subject *i* received the information treatment; $\mathbf{x}_i$ is a vector of controls[8]; and ε_i is an individual-specific error term. We use robust error terms for inference. y_i is the outcome variable of interest.

3.3.2   Post-treatment beliefs about reporting

We provide evidence that treated respondents expect more informative news reports from the *New York Times* and hold more favorable views of the newspaper. First, treated

---

[8]We use the following pre-specified controls: gender (male indicator), age (continuous), log income (continuous), region (three indicators), race (white indicator), education (college indicator), employment status (indicator for full-time work), frequency of reading the *New York Times* (continuous, elicited pre-treatment), beliefs about the consequences of the Trump Tax Plan and the Trump Healthcare Plan (both continuous and elicited pre-treatment), pre-treatment beliefs about the probability that the *New York Times* would report unbiasedly, and experiment fixed effects (two indicators). We also have a few respondents in the sample who did not complete all demographic questions; we include indicators for missing values for these respondents.

respondents positively update about the informativeness of the newspaper as they think it is 6.8 percentage points more likely that the *New York Times* does not suppress any information about the CBO report on the Trump Tax Plan ($p < 0.01$, column 1 of Table 2). Second, column 2 provides suggestive evidence that the treatment also affects beliefs about the extensive margin of news coverage: Treated respondents are 3.3 percentage points more likely to think that the *New York Times* covers a CBO report which contradicts claims by Democrats that granting citizenship status to undocumented immigrants would not have negative fiscal consequences. However, this effect is not statistically significant at conventional levels. Third, treated respondents are 4.1 percentage points more likely to think that the *New York Times* is not politically biased ($p < 0.01$, column 6). This provides indirect evidence that people perceive the article containing both statistics as more balanced. Finally, perceived quality of *New York Times* articles is 10.3 percent of a standard deviation higher among treated respondents ($p < 0.10$, column 3). Taken together, treated respondents think that the *New York Times* provides more information, is less politically biased and writes higher quality articles. Any of these beliefs should imply a subsequent increase in demand for news in rational models of information demand.

[Insert Table 2 here]

3.3.3 Treatment effects on demand for news

Our main object of interest is people's demand for news, which takes value one if our respondents want free access to the news story about the Trump Tax Plan and zero otherwise. The main finding of this paper is that respondents who learn that the *New York Times* is more informative than they thought reduce their demand for news. Column 1 of Table 3 highlights that treated respondents on average significantly reduce their demand for news by 3.5 percentage points. This corresponds to a reduction in the demand for news of approximately 13 percent, i.e. one third of the control group difference in demand for reading the article between Republicans and Democrats.[9] A comparison of the treatment

---

[9]The baseline demand for news of 27.4 percent in the control group reflects people's opportunity cost of reading the article and is in itself not indicative of a violation of the more information is better principle.

effect with the magnitude of the first-stage—i.e., the 6.8 percentage point increase in the perceived likelihood that the *New York Times* does not suppress information—suggests that people's demand for news is relatively elastic to changes in perceptions of informativeness.

The median time spent reading the article about the Trump Tax Plan is 66 seconds, suggesting that a substantial fraction of our respondents read at least some parts of the article.[10] The time spent reading the article does not vary significantly across treatment arms, indicating that the treatment did not affect how carefully people read the article.

[Insert Table 3 here]

3.3.4 Preference for belief confirmation and news demand

The negative treatment effect on demand for news is inconsistent with people being rational unbiased readers who only seek news to get information. One mechanism that potentially explains the negative treatment effect is that people get disutility from reading news inconsistent with their prior beliefs (Mullainathan and Shleifer, 2005). If our respondents generally believe that the *New York Times* slants to the left, a preference for belief confirmation predicts a clear polarization in news consumption between Democrats and Republicans. However, as shown in Figure B.1, there is substantial heterogeneity in beliefs about whether the *New York Times* slants to the right or left. This heterogeneity makes it necessary to take prior beliefs into account when testing predictions based on a preference for belief confirmation.

Furthermore, while party affiliation strongly predicts people's beliefs about the Trump Tax Plan, a substantial fraction of both Republicans and Democrats hold beliefs that differ from their own political party (as shown in Figure B.2). Since people might receive disutility from reading news inconsistent with their own beliefs as well as news inconsistent with claims made by their political party, it is not clear what we should predict when people hold divergent beliefs from claims made by their own political party. To examine whether the main result from Experiment 1.1 is consistent with a preference for belief

---

[10]Reading the full article takes between four and five minutes.

confirmation, we therefore focus on respondents with party-consistent beliefs and estimate treatment effects separately for Democrats and Republicans.[11]

Among Democrats, we observe a general decrease in the demand for news by 3.2 percentage points (column 1 of Table C.4). Consistent with predictions based on a preference for belief confirmation, we observe a negative but non-significant interaction effect between the treatment and beliefs about left-wing slant (column 2 of Table C.4). To allow for a more flexible heterogeneity analysis, we also examine heterogeneous treatment effects based on a nonparametric kernel smoothing estimator (Hainmueller et al., 2019). As shown in Panel A of Figure B.4, there is a U-shaped pattern in the heterogeneous responses: Democrats who think the *New York Times* is quite likely to slant to the left significantly reduce their demand for news in response to the information, but there is no response among respondents with more extreme beliefs in either direction. This U-shaped pattern is consistent with a preference for belief confirmation. As we show in Section A.2 of the Online Appendix, a simple model predicts that people with more extreme beliefs are less likely to be marginal, implying a non-monotonic relationship between pre-treatment beliefs and news demand. The intuition behind this observation is that respondents who are predicted to receive the largest negative changes in consumption utility had higher expected utility from reading the article to begin with.

However, as shown in column 3 of Table C.4, we also see a negative non-significant interaction effect between the treatment and pre-treatment beliefs about right-wing slant. The negative sign of the interaction effect is clearly not consistent with a preference for belief confirmation. As shown in Panel B of B.4, the negative point estimate is generally driven by Democrats who believe that the *New York Times* is very likely to slant to the right. However, the effect for this subgroup is imprecisely estimated as very few Democrats believe that the *New York Times* slants to the right.

For Republicans, we observe a general decrease in the demand for news by 4 percentage points (column 5 of Table C.4). There is essentially no treatment heterogeneity by

---

[11]I.e., we exclude Democrats who think that the plan has overall positive consequences and Republicans who think it has overall negative consequences from the regressions. The regressions include Independents who lean toward the Democratic Party or the Republican Party.

beliefs about left-wing slant (column 6 of Table C.4 and Panel C of Figure B.4). Assuming that news demand is elastic to changes in beliefs for all respondents, a preference for belief confirmation would predict a strict increase in demand among Republicans who thought the *New York Times* was very likely to slant to the left. On the other hand, consistent with a preference for belief confirmation, we observe a strong and statistically significant negative interaction effect between the treatment and pre-treatment and beliefs right-wing slant (column 7). The nonparametric analysis shows that the reduction in demand among Republicans is driven by respondents who initially believed that the *New York Times* was very likely to be right-wing biased (Panel D of Figure B.4)

To conclude, the heterogeneity analysis does not provide very conclusive evidence either in favor or against the relevance of a preference for belief confirmation in driving our negative treatment effect. Consistent with a preference for belief confirmation, the negative treatment effect is broadly driven by respondents who learn that the *New York Times* is less likely to confirm their prior beliefs. On the other hand, we do not observe a similar increase in demand for news among respondents who learn that the *New York Times* is less likely to only report statistics that contradict their beliefs. However, if respondents have a very strong preference for belief confirmation and are thus motivated to avoid articles that contain any information that conflicts with their prior beliefs, it is not clear that these respondents should increase their demand in response to the information.

3.3.5 Article spin

We also provide evidence on how the demand for news responds to more subtle forms of media bias, such as "spin," namely newspapers' tendency to systematically emphasize facts that favor a particular interpretation of an event. We conducted an additional experiment where we manipulate respondents' beliefs about the facts emphasized in the headline of an article in the *New York Times* to study whether perceptions of spin also generate belief-based utility motives for reading news. In this experiment, we provide some respondents with information highlighting that the *New York Times* slanted a news report to the left. In line with people having a preference for belief confirmation, the treatment significantly

polarizes news consumption between those who, in light of their prior beliefs and political affiliation, anticipate their beliefs to be confirmed or to be challenged by the *New York Times* article. We provide details on the design and results from this experiment in Section D of the Online Appendix.

3.4 Discussion of alternative explanations

One explanation for our main finding that respondents reduce their demand for news that they perceive as more informative is that people have a preference for belief confirmation that sometimes dominates people's desire for more information. It could also be case that people perceive more informative news as less entertaining, explaining why we generally see a decrease in the demand for news across different subgroups.

We now provide evidence against a series of potential alternative mechanisms. We first discuss behavioral explanations based on cognitive constraints, curiosity, and experimenter demand effects. We then discuss rational explanations based on beliefs about quality, diversification of news sources, and delegation incentives for filtering. Finally, we provide evidence that our results are robust to various changes in the experimental design.

3.4.1 Behavioral explanations

Cognitive constraints The more-information-is-better principle might not hold in our setting if the marginal cognitive cost of processing one additional statistic from a CBO report exceeds the expected informational value from the additional statistic. Our main finding could then be driven by treated respondents who expect the article about the Trump Tax Plan to contain too much statistical information to justify the cognitive cost of reading the article.

To assess whether cognitive constraints are likely to drive our main result, we conduct a placebo experiment in collaboration with Dynata (Experiment 1.3, $n = 930$; see Table 1). As in the main experiment (Experiment 1.1), we inform respondents that the CBO analyzed

18

the impact of the Trump Healthcare Plan.[12] We then inform respondents that the CBO highlighted two key statistics in its report and that the *New York Times* subsequently wrote an article about the report. However, in contrast to the main experiment, we do not tell respondents what the statistics are about. We then ask our respondents to state the percent chance they assign to the *New York Times* citing zero, one, or two of these two key statistics from the report. To exogenously vary beliefs about how many statistics the *New York Times* is likely to mention from CBO reports, we inform a random subset of respondents that the *New York Times* reported both key statistics in its coverage of the Trump Healthcare Plan. We then measure demand for news exactly as in the main experiment by asking respondents whether they want free access to an article in the *New York Times* about the Trump Tax Plan. In this neutral setting, people do not reduce their demand for news when learning that the *New York Times* provides more informative news. If anything, the treatment increases people's demand for news by 1.5 percentage points (column 1 of Table C.7)—the opposite of the prediction of the cognitive constraints account.

Explanations based on cognitive constraints are also inconsistent with several patterns in the data from the main experiment. First, if we use educational attainment as a proxy for cognitive costs with respondents, we do not find any statistically significant differences in treatment effects for people with low or high cognitive costs (column 2 of Table C.5). Second, for respondents who were cross-randomized into not being offered free access to the article, we collected a series of post-treatment measures related to perceived cognitive costs of reading the article about the Trump Tax Plan. Respondents in the treatment group think the article about the Trump Tax Plan would be equally complex, dry and technical as respondents in the control group (columns 4 and 5 of Table 2).

Curiosity   In the main experiment (Experiment 1.1), we elicit pre-treatment beliefs about suppression of information in the *New York Times* and only inform treated respondents about whether it actually suppressed information. This creates two potential curiosity mo-

---

[12]In this experiment, we used the term "GOP Health Bill" instead of "Trump Healthcare Plan."

tives that could differ between the treatment and control group. First, treated respondents might be curious about whether the information we provided was accurate and perceive the article about the Trump Tax Plan as a chance to validate the information, thereby increasing demand relative to the control group. Second, control group respondents might be curious to learn whether or not the *New York Times* tends to suppress information from CBO reports and view the article about the Trump Tax Plan as an opportunity to learn about this, thereby increasing demand relative to the treatment group. However, it is worth emphasizing that learning about whether the *New York Times* suppressed information from the Trump Tax Plan is not straightforward as it would require both reading the published article in the *New York Times* as well as retrieving and reading the original CBO report about the Trump Tax Plan.[13] The net directional effect of the two curiosity motives is difficult to predict since they work in opposite directions, but curiosity could potentially explain our negative main effect if the motive is present and stronger in the control group.

To address concerns about curiosity, we conducted an additional experiment on a broadly representative sample recruited in collaboration with Lucid ($n = 3,189$, Experiment 1.2; see Table 1).[14] In this experiment, we tell all respondents that we will ask them a question about how the *New York Times* covered the findings from a CBO report and highlight the following text in bold: "We will tell you how The New York Times covered these findings at some later point in the survey" (Section E.2.2 in the Online Appendix provides a screenshot). If control respondents were curious to find out whether the *New York Times* tends to suppress information from CBO reports, they would no longer have an extra incentive to read the article to find out.[15] As column 2 of Table 3 illustrates, we find a quantitatively similar effect size in this experiment (if anything, we uncover a larger treatment effect than in the main experiment)—suggesting a limited role for curiosity in driving the treatment effects.[16]

---

[13]Our respondents do not receive any information about the CBO report about the Trump Tax Plan.

[14]Lucid is a survey provider commonly used in social science research (Wood and Porter, 2019).

[15]Expected learning about biases in reporting is thus constant across the treatment and control group, and learning about any possible bias in the second article would again require respondents to read both the article and the underlying CBO report.

[16]As we discuss in Section 3.4.3 on robustness of design choices, this experiment also used different articles from the main experiment to elicit beliefs and article demand.

Several patterns in the data from the other experiments are also inconsistent with a strong curiosity motive driving the negative treatment effect on demand for news. First, in the main experiment (Experiment 1.1), we collected a post-treatment measure on how interested people were in learning whether the *New York Times* "reports unbiasedly about political issues." Column 8 of Table 2 shows that treated respondents are not more curious to learn about this. The effect is close to zero and relatively precisely estimated. Second, curiosity motives might also be present in the placebo experiment (Experiment 1.3; see Table 1) in which we measure people's beliefs about whether the *New York Times* suppressed any key statistics from a CBO report without giving respondents any information about what the statistics were. Instead of being curious about bias in reporting, respondents in the control group might be curious to learn how much statistical information the *New York Times* tends to report. However, we find, if anything, that the treatment increases people's demand for news—inconsistent with the predictions of the curiosity motive (column 1 of Table C.7).

Experimenter demand effects It is possible that treated respondents form different beliefs about the experimenter's expectations compared to control group respondents. However, we do not believe that experimenter demand is a major concern in our setting: First, it seems more likely that learning that a newspaper provides more informative news should create the expectation that demand for news should be increased—the opposite of what we find. Additionally, the treatment was framed in a neutral way specifically to avoid concerns about experimenter demand effects. Second, we do not observe a decline in the demand for news in the placebo experiment (see p. 18 for a description) where we also inform people that the *New York Times* provides more information. Third, recent evidence suggests that experimental subjects respond only moderately to explicit signals about the experimenter's expectations, indicating a limited quantitative importance of experimenter demand effects (de Quidt et al., 2018; Mummolo and Peterson, 2018).

3.4.2 Rational explanations

Trust and quality   In the context of news consumption, a theoretically important explanation is based on Bayesian updating about the quality of the *New York Times* and its information sources. Gentzkow and Shapiro (2006) show that if there is uncertainty about both the state of the world and the quality of a newspaper, a Bayesian consumer will update negatively about the quality of the newspaper after reading an article that conflicts with his prior belief about the state of the world. While this exact mechanism cannot explain our main result since we keep the source of the information constant across articles, the treatment may still affect people's evaluation of the *New York Times* along the quality dimension. We therefore collect a battery of post-treatment measures of quality as well as on trust and political bias. Empirically, treated respondents are more likely to say the *New York Times* provides high-quality articles (column 3 of Table 2) and display identical levels of trust in the *New York Times* (columns 7). There are no obvious reasons why our treatment should affect beliefs about the CBO. In line with this, we find no treatment effect on respondents' perceptions of the political bias of the CBO (column 9) or their level of trust in the CBO (column 10). Rather, respondents do update positively about the accuracy of the CBO (column 11). This effect in isolation should increase the demand for news about the CBO if people want more informative news, the opposite of what we find.

Diversification   Another potential explanation is based on the idea that people might consume a diverse set of news articles to extract a more informative signal by combining the different pieces of information (Mullainathan and Shleifer, 2005).[17] Accordingly, a newspaper is particularly valuable if it provides information that is complementary to the information contained in the consumer's news portfolio. Our treatment might then reduce the value of the *New York Times* in balancing out right-leaning news sources because it is perceived as more even-handed.

---

[17]This portfolio motive hinges on people's perceived ability to debias themselves. However, empirical evidence suggests that it may be difficult for people to fully debias themselves (Enke, 2019).

We asked people pre-treatment to list up to three newspapers they are likely to read from a list of 20 popular newspapers across the political spectrum. For 46 percent of our respondents, the diversification motive is not present as they selected only left-leaning or right-leaning newspapers. Moreover, treatment effects are similar for respondents that only consume newspapers on one side of the political spectrum compared to those who read both at least one left-wing newspaper and one right-wing newspaper (column 1 of Table C.5).

Delegation   Consumers delegate costly information acquisition to newspapers. If demand-side or supply-side constraints limit newspapers' ability to communicate all the information available to them, Suen (2004) and Chan and Suen (2008) show that it can be rational for consumers to have a demand for articles that primarily contain information that confirm their prior beliefs. Delegation incentives are psychologically different from a behavioral preference for belief confirmation, but make similar predictions.

We think that delegation incentives do not drive our treatment effects. First, supply-side constraints are unlikely given that all major newspapers reported both findings.[18] Moreover, all respondents were informed that the *New York Times*'s article contains 1,100 words, which is sufficient to discuss all key results from CBO reports. Second, our placebo experiment provides evidence against demand-side constraints based on cognitive constraints (see p. 18 for the discussion). Third, one implication of delegation is that people from different political groups may have differential demand for different pieces of information. In an additional experiment, we test empirically whether Democrats and Republicans exhibit such patterns of differential demand with data from a representative online panel in which we measure people's demand for learning about the CBO estimates about the Trump Healthcare Plan and the Trump Tax Plan (Experiment 1.4; see Table 1). We find no differential demand for different pieces of information within each political group (as shown in Figure B.6).

---

[18]We verified that all top 15 newspapers by circulation (as of June 2019) reported both findings.

3.4.3 Robustness

We also conducted additional experiments using samples from Lucid (Experiment 1.2) and Amazon Mechanical Turk (Experiment 1.5, 1.6, and 1.7) to assess the robustness of our results (Table 1 provides an overview of all experiments).

Article choice and monetary incentives   We first assess the robustness of our main result to a different choice of articles. In Experiment 1.2 with Lucid ($n = 3,189$; see Table 1) we use two new CBO reports to elicit beliefs about information suppression and measure article demand.[19] To elicit beliefs about information suppression, we rely on a CBO report about the consequences of "Democrats' \$15 Minimum Wage Bill" (the Raise the Wage Act, a Democratic bill to to raise the federal minimum wage to \$15). Respondents are told that the CBO estimated that the bill would lift 1.3 million people out of poverty (contradicting claims made by Republicans) and decrease the number of jobs by 1.3 million (contradicting claims made by Democrats). To measure article demand after an information treatment in which respondents are informed that the *New York Times* reported both statistics, we offer all respondents free access to an article in the *New York Times* covering a CBO report about the consequences of establishing a single-payer health care system. As shown in column 2 of Table 3, demand for news declines by 4.7 percentage points in this experiment ($p < 0.01$), confirming that our main result is robust to using different articles.

In an experiment on Amazon Mechanical Turk (Experiment 1.6, $n = 723$), conducted with Democrats and Democrat-leaning Independents, we reverse the order of articles used in Experiment 1.1. Furthermore, in this experiment, we elicit pre-treatment beliefs about how the *New York Times* covered the findings from the CBO report about the Trump Tax Plan using monetary incentives and a quadratic scoring rule.[20] We subsequently measure people's demand for the article about the CBO evaluation of the Trump Healthcare Plan. The patterns of beliefs and treatment effects are very similar to those in our main filtering

---

[19]In this experiment, we also address concerns about curiosity as discussed on page 19.

[20]We randomly selected one in ten respondents to be paid up to \$1 according to their guess.

24

experiment (column 1 of Table C.6), suggesting that monetary incentives and reversed article order do not substantially affect our results.[21]

Platform   Across the experiments, we recruit respondents from three different platforms: Dynata, Lucid, and Amazon Mechanical Turk. These platforms are extensively used in social science research.[22] Table 3 shows that the main treatment effect on demand for news is very stable across platforms, which includes an experiment on Amazon Mechanical Turk with an identical design as the main experiment (Experiment 1.7; $n = 1{,}332$). If anything, the estimated treatment effects are larger in our experiment using a representative sample from Lucid (column 2) and in our experiments on Amazon Mechanical Turk (column 3).

External validity   We conducted an additional experiment on Amazon Mechanical Turk (Experiment 1.5; $n = 199$) in which we assess the external validity of our behavioral measure of article demand. Specifically, we measure in randomized order both people's demand for news and (incentivized) willingness to pay for a 3-month subscription to the *New York Times* using a multiple price list.[23] We find that our measure of article demand is significantly correlated with people's willingness to pay ($\rho = 0.298$, $p < 0.01$). Despite being a binary variable, article demand has greater explanatory power for people's willingness to pay compared to a saturated regression controlling for political affiliation, gender, income, and people's beliefs about how the *New York Times* covered a CBO report.

---

[21]The point estimate of -0.04 is very close to the change in the demand for news in the main filtering experiment of -0.035, but would require larger samples to be statistically significant.

[22]Coppock and McClellan (2019) find that samples from Lucid score similarly to the American National Election Study's (ANES) on the Big-5 personality inventory, show similar levels of political knowledge, and recover framing effects similar to the ones observed in the General Social Survey. Horton et al. (2011) find that experiments on MTurk closely replicate results from traditional lab experiments.

[23]Respondents decide between varying amounts of U.S. dollars and a subscription to the *New York Times*. We informed respondents that one out of ten randomly chosen participants would get one of their choices implemented. We used the following monetary amounts: 50 cents, $1, $2, $3, $4, $5, $10. Screenshots of the willingness to pay elicitation are provided in Section E.5.

4 False claims experiment

This section provides another test of the more-information-is-better principle by studying how perceptions of false claims affect the demand for news. The design complements the main experiment by changing perceptions of a different source of informativeness. Conceptually, changing perceptions of false claims amounts to changing beliefs about the noise in the signal reported by a news outlet. The more-information-is-better principle predicts that the demand for news should strictly decrease in the perceived level of noise.

4.1 Theoretical framework: False claims

We briefly formalize how false claims by news outlets affect the informativeness of news, which provides us with a theoretical benchmark for our treatment effect. Motivated by the empirical challenge to disentangle genuine mistakes from intentional distortion, we adopt a reduced-form perspective and directly focus on the agent's belief about the state-dependent distribution $\sigma : \Theta \to \Delta(N)$ over news articles.

The newspaper is constrained to binary messages $n \in \{L, R\}$ about the binary state $\theta \in \Theta$. The probability of reporting $n = R$ in state θ is σ_θ for each $\theta \in \Theta$. Specifically, σ_R is the likelihood of a correct report in state R, while σ_L represents the likelihood of a false claim in state L. Intuitively, an agent should thus prefer larger σ_R and smaller σ_L if his objective is to learn about the truth. Assume that $\sigma_L < \sigma_R$, i.e. the newspaper is more likely to report $n = R$ if the state is R instead of L.[24]

Consider two distributions over articles given by $\sigma = (\sigma_L, \sigma_R)$ and $\sigma' = (\sigma'_L, \sigma_R)$ where $\sigma_L < \sigma'_L < \sigma_R$. An article n from the distribution σ is Blackwell more informative than an article n' from the distribution σ'.[25] Intuitively, we obtain σ' by adding noise to σ in state L, thereby reducing its informativeness. This suggests a second empirical test of whether people value more informative news that is independent of people's prior

---

[24]This is without loss of generality because the article is completely uninformative if $\sigma_L = \sigma_R$, and we can relabel the articles if $\sigma_L > \sigma_R$.

[25]See Proposition 2 in the Online Appendix for a proof.

beliefs by inducing exogenous variation in the level of noise. In our experimental design, we exogenously increase treated respondents' perceived likelihood that *CNN* makes false claims in its political reporting. Theoretically, this should decrease the demand for news irrespective of whether these changes in perceived informativeness arise from people updating about the degree of distortion by *CNN* or about the likelihood of making a mistake in its reporting.

4.2 Experimental design and results

Design and sample   We collected data for the experiment in two waves in collaboration with Lucid. We have a sample of 2,081 respondents that is broadly representative of the U.S. population in terms of age, race, gender, income, race, and region (column 5 of Table C.1). We first measure basic demographics as well as a range of other background characteristics and political views (Table E.1 of the Online Appendix provides a full overview). We then introduce exogenous variation in perceptions of false claims made by *CNN* in its political reporting. Specifically, we first tell our respondents that the U.S. recently extracted a high-level CIA spy from Russia. We then tell our respondents that some media outlets correctly reported that CIA extracted the spy because of widespread media speculation about its sources, whereas other media outlets falsely reported that CIA extracted the spy because it feared that President Trump would mishandle classified information.[26]

We subsequently measure beliefs about the percent chance that *CNN* reported that the spy was extracted because (i) the CIA feared that President Trump would mishandle

---

[26]The experiment involves no deception. For background information, see the following article in Associated Press News: `https://www.apnews.com/b432a20d85ca48a5bca6c3394920c7fe` (accessed October 2, 2019). In the article, CIA's Director of Public Affairs, Brittany Bramell, is quoted as saying: "CNN's narrative that the Central Intelligence Agency makes life-or-death decisions based on anything other than objective analysis and sound collection is simply false" and that "Misguided speculation that the president's handling of our nation's most sensitive intelligence, which he has access to each and every day, drove an alleged exfiltration operation is inaccurate." The *New York Times* also wrote an article challenging *CNN*'s reporting, writing that "former intelligence officials said there was no public evidence that Mr. Trump directly endangered the source, and other current American officials insisted that media scrutiny of the agency's sources alone was the impetus for the extraction" (https://www.nytimes.com/2019/09/09/us/politics/cia-informant-russia.html, accessed October 2, 2019).

classified information, or because (ii) of widespread media speculation about the CIA's sources. To introduce exogenous variation in perceptions of false claims in *CNN* news reports, we provide a random subset of respondents with the following information treatment:[27]

> In its article, CNN reported that the spy was extracted because the CIA feared that President Trump would mishandle classified information.

The treatment was purposefully framed in a neutral way to minimize experimenter demand effects. We thereafter measure people's demand for reading a different news article from *CNN* on a related but different topic, namely the Trump impeachment inquiry.[28] As in the filtering experiment, respondents are told that they will receive access to the article at the end of the survey if they say "Yes" to the article and that they will proceed with the survey without receiving access to the article if they say "No."

Finally, we ask a series of post-treatment beliefs about *CNN* and its reporting. Specifically, we measure trust, perceptions of quality and political bias in reporting, beliefs about false claims made by *CNN* in its political reporting, and the percent chance that a *CNN* article about President Trump would contain any false claims. We further measure beliefs about strategic risk taking by separately asking respondents how likely they think it is that *CNN* would publish articles based on "unverified and potentially misleading sources" about (i) President Trump and (ii) Joe Biden (the then Democratic front-runner for the 2020 presidential race).

---

[27] To address concerns about differential curiosity between the treatment and control group, we inform all respondents that we will tell them how the CNN covered the story "at some point later in the survey" just before eliciting their prior beliefs.

[28] Wave 1 ($n = 1{,}427$) and wave 2 ($n = 645$) use identical instructions except for the introductory sentence to the *CNN* article we use to measure the demand for news. We started data collection for wave 1 on September 24, 2019, the day Speaker of the House Nancy Pelosi announced the formal impeachment inquiry. Subsequent rapid developments made us concerned that our article would be perceived as "old news," which is why we made an adjustment. Specifically, in wave 1, we tell respondents that "A recent whistleblower complaint following a phone conversation between President Trump and a foreign leader led Democrats in Congress to start an impeachment inquiry against the president." We then inform them that *CNN* wrote an article about the impeachment inquiry and measure their demand. In wave 2, we tell respondents that "CNN today published a new story about the Trump impeachment inquiry" and measure their demand for this new article. We include wave fixed effects in all specifications.

Results As shown in Table 4, the treatment generates a highly significant first stage on perceptions of false claims in *CNN*. Treated respondents are 2.8 percentage points more likely to think that *CNN* articles about President Trump contains false claims (column 1, $p < 0.01$), 6.6 percent of a standard deviation more likely to think that that *CNN* generally makes false claims in its political reporting (column 2, $p < 0.01$), and 9.5 percent of a standard deviation more likely to think that *CNN* publishes stories about President Trump based on unverified and potentially misleading sources (column 3, $p < 0.01$). We find suggestive evidence that treated respondents are also more likely to think that *CNN* publishes unverified stories about Joe Biden (column 4), but the effect is smaller than for President Trump and not statistically significant at conventional levels. Treated respondents are also 3.3 percent of a standard deviation more likely to think that *CNN* intentionally tries to hurt President Trump by publishing false claims (column 5), but this effect is not statistically significant and also small compared to the treatment effect on perceptions of whether *CNN* publishes stories based on unverified and misleading sources. These results thus suggest that our respondents are more likely to interpret the false claim as a mistake than as strategic distortion.

The treatment also affects more general perceptions of *CNN*: Treated respondents think that news articles from *CNN* are 8.7 percent of a standard deviation lower in quality (column 6, $p < 0.01$), 5.6 percentage points more likely to be politically biased (column 7, $p < 0.01$), and 4.7 percent of a standard deviation less trustworthy (column 8, $p < 0.10$). The effects on perceptions of quality and political bias are quantitatively similar to the treatment effect observed in the main filtering experiment. The treatment effect on perceptions of trust is, if anything, stronger than in the main filtering experiment.

Taken together, these results suggest that treated respondents perceive news articles in *CNN* as less informative and more noisy—in particular when it comes to political reporting about President Trump. Given these results, the more-information-is-better principle predicts a strict decrease in the demand for news from *CNN*.[29] Contrary to this

---

[29] As discussed in Section 4.1, informativeness is negatively affected irrespective of the reason for the false claim.

theoretical prediction, the main result from this experiment is that treated respondents do not reduce their demand for news from CNN. The point estimate on demand for news is an increase of 0.8 percentage points with a standard error of 1.8 percentage points (column 9 of Table 4); i.e., close to zero and relatively precisely estimated. This finding suggests that people's demand for news is inelastic to perceptions of false claims. Furthermore, we know from the filtering experiment that this result is not due to demand for news being generally inelastic to beliefs about reporting.[30]

Are the patterns of heterogeneity we observe in the data consistent with respondents having a preference for belief confirmation?[31] To study the role of preference for belief confirmation, we proxy people's attitudes toward President Trump by their prior belief that President Trump betrayed his oath of office and should thus be impeached. Treated respondents who think that Trump is guilty should increase their demand for news because they expect the article to contain more negative information about President Trump, while treated respondents who believe that Trump is not guilty should decrease their demand for news from *CNN*. Consistent with these predictions, treated respondents who think that Trump is not guilty decrease their demand for news from *CNN*, while for respondents who think that Trump is guilty, we observe a positive point estimate (column 9 of Table C.9). While the patterns we uncover are qualitatively consistent with people having a preference for belief confirmation, the effects are far from being statistically significant at conventional levels.

5   Discussion

Across two complementary experimental designs, we find that people's demand for news does not respond to exogenous variation in informativeness in ways predicted by the

---

[30]In the main filtering experiment, we find a treatment effect of -3.5 percentage points using a very similar setup for the experimental design which results in comparable magnitudes of changes in perceptions of article quality.

[31]Tables C.8 and C.9 show that there is little evidence of differential updating of beliefs about the news source by political affiliation and beliefs about whether Trump is guilty. The relatively homogeneous first-stage in beliefs justifies a reduced form approach for analyzing heterogeneous treatment effects.

more-information-is-better principle. We show that this is not due to cognitive constraints, suggesting an important role for psychological utility in news consumption.

Our findings imply that people's preferences over news articles reflect more than a pure taste for more information. While models of news consumption differ in their assumptions, preferences over news articles probably result from trade-offs between three key dimensions: (i) psychological utility from belief confirmation, (ii) the entertainment value of news, and (iii) the value of information. We discuss how our results relate to each dimension in turn.

Psychological utility arising from people's intrinsic desire to confirm their prior beliefs seems particularly relevant for models of political news consumption. Political polarization and strong identification with one's own political party create an incentive to avoid cognitive dissonance by seeking like-minded news. Indeed, this mechanism has been recognized in theoretical work on media markets. For example, Mullainathan and Shleifer (2005) analyze the market for news assuming that consumers prefer news that are closer to their partisan beliefs—a preference for cognitive consistency also present in theories of mass communication (Severin and Tankard, 2000). Relatedly, Bernhardt et al. (2008) assume that consumers receive consumption utility from political news containing negative information about the other party's candidates. We find some evidence consistent with a preference for belief confirmation in driving political news consumption across the different experiments. First, treated respondents in the main filtering experiment decrease their demand for news. Moreover, the patterns of heterogeneity in treatment responses from the filtering experiment are somewhat consistent with the predictions made by this mechanism.[32] Second, the result that treated respondents who think that a news outlet is more likely to make false claims do not reduce their demand for this outlet is also consistent with explanations based on psychological utility from belief confirmation.

A second motive that might shape people's demand for political news is the pure consumption or *entertainment* value of it. People may be motivated to read news for

---

[32] We do not observe a decrease in the demand for news in the placebo experiment where other potentially conflicting psychological motives, such as a preference for belief confirmation, are arguably not affected by the treatment.

the entertainment value from witty commentary or surprising revelations about public figures (Ely et al., 2015). The entertainment value of political news could compensate consumers for the cost of gathering information from news articles (Dyck et al., 2013). Our findings from the filtering experiment suggest that the entertainment value of balanced news might be perceived as lower. Moreover, the results from the false claims experiment are consistent with people valuing more extreme news stories even if they are based on unverifiable or even false information.

Finally, we consider the desire for more informative news as a driver of news consumption. Even though we provide evidence that the desire for more informative news is not the dominant motive for news acquisition in the context we study, it might still be an important motive at play. Conceptually, information can have both instrumental value if it informs decision-making and non-instrumental value arising from intrinsic information preferences, i.e. people wanting to learn about the truth. One important caveat for the interpretation of our results is that the value of information might vary across decision-making contexts. We believe that future research should explore whether people also violate the more-information-is-better principle in news domains in which people have potentially stronger instrumental motives to acquire informative news. However, whether people value more informative political news is of particular interest nonetheless because it is a key input for the functioning of democracies (Strömberg, 2004).

6   Concluding remarks

Our paper provides the first causal evidence on whether the desire for more information is people's dominant motive for reading economic and political news in a context where cognitive constraints are not binding. Our main finding is that respondents who learn that a news outlet does not strategically suppress information, and thus expect to receive more information, reduce their demand for articles from this outlet. Moreover, learning that a news outlet makes false claims in its reporting does not decrease people's demand for news from this outlet. These findings are inconsistent with the theoretical benchmark

prediction of the more-information-is-better principle and suggest that people hold other motives for consuming news that sometimes dominate their desire for more information.

Understanding how both demand and supply side factors shape media content is of high relevance due to the media's influence on public discourse (King et al., 2017) and political outcomes (DellaVigna and Kaplan, 2007). Our findings suggest an important role for demand-side explanations of media bias. The distinction between demand and supply side factors has important policy implications as competition generally reinforces the incentives to deliver the product consumers want. If people's dominant motive for reading news is to obtain more information, market regulations designed to increase competition should reduce bias in reporting and improve information aggregation (Chan and Suen, 2009). However, our finding that people do not respond to variation in informativeness in the way predicted by the more-information-is-better principle suggests that the effects of regulation are more nuanced. The distinction between demand and supply side factors also has important implications for demand-side policy interventions that aim to correct consumers' misperceptions of the informativeness of news, such as transparency initiatives to inform consumers about the extent of media bias in markets and efforts of fact-checking organizations to debunk false claims. Under supply-driven media bias, increasing consumer knowledge about media bias leads to welfare improvements by steering consumers toward more informative news. Under demand-driven media bias, by contrast, such interventions might backfire and actually increase political belief polarization by shifting people toward more biased sources. Our findings thus demonstrate the complexity of optimal regulation and policy.

References

Allcott, Hunt and Matthew A. Gentzkow, "Social Media and Fake News in the 2016 Election," *Journal of Economic Perspectives*, 2017, *31* (2), 211–36.

Bernhardt, Dan, Stefan Krasa, and Mattias Polborn, "Political polarization and the electoral effects of media bias," *Journal of Public Economics*, 2008, *92* (5-6), 1092–1104.

Blackwell, David, "Comparison of Experiments," in Jerzy Neyman, ed., *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press Berkeley, CA 1951, pp. 93–102.

_ , "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 1953, pp. 265–272.

Chan, Jimmy and Wing Suen, "A Spatial Theory of News Consumption and Electoral Competition," *Review of Economic Studies*, 07 2008, *75* (3), 699–728.

_ **and** _ , "Media as watchdogs: The role of news media in electoral competition," *European Economic Review*, 2009, *53* (7), 799 – 814.

Charness, Gary, Ryan Oprea, and Sevgi Yuksel, "How Do People Choose Between Biased Information Sources? Evidence from a Laboratory Experiment," *Working Paper*, 2019.

Coppock, Alexander and Oliver A. McClellan, "Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents," *Research and Politics*, 2019, *6* (1), 1–14.

de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth, "Measuring and Bounding Experimenter Demand," *American Economic Review*, 2018, *108* (11), 3266–3302.

DellaVigna, Stefano and Eliana La Ferrara, "Economic and Social Impacts of the Media," in Simon P. Anderson, Joel Waldfogel, and David Strömberg, eds., *Handbook of Media Economics*, Vol. 1B of *Handbook of Media Economics*, North-Holland, 2015, pp. 723 – 768.

_ **and Ethan Kaplan**, "The Fox News Effect: Media Bias and Voting," *Quarterly Journal of Economics*, 2007, *122* (3), 1187–1234.

_ **and Johannes Hermle**, "Does Conflict of Interest Lead to Biased Coverage? Evidence from Movie Reviews," *Review of Economic Studies*, 2017, *84* (4), 1510–1550.

Durante, Ruben and Brian Knight, "Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy," *Journal of the European Economic Association*, 2012, *10* (3), 451–481.

Dyck, Alexander, David Moss, and Luigi Zingales, "Media versus Special Interests," *Journal of Law and Economics*, 2013, *56* (3), 521–553.

Ely, Jeffrey, Alexander Frankel, and Emir Kamenica, "Suspense and Surprise," *Journal of Political Economy*, 2015, *123* (1), 215–260.

Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya, "Media and Political Persuasion: Evidence from Russia," *American Economic Review*, 2011, *101* (7), 3253–85.

Enke, Benjamin, "What you see is all there is," *Available at SSRN 2691907*, 2019.

_ , "Moral Values and Voting," *Journal of Political Economy*, 2020.

Exley, Christine L, "Excusing Selfishness in Charitable Giving: The Role of Risk," *Review of Economic Studies*, 2015.

Exley, Christine L. and Judd B Kessler, "Motivated Errors," *Working Paper*, 2018.

Falk, Armin and Florian Zimmermann, "A Taste for Consistency and Survey Response Behavior," *CESifo Economic Studies*, 2012, *59* (1), 181–193.

_ **and** _ , "Beliefs and Utility: Experimental Evidence on Preferences for Information," *Working Paper*, 2017.

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

Garz, Marcel, Gaurav Sood, Daniel F. Stone, and Justin Wallace, "What Drives Demand for Media Slant?," *Working Paper*, 2018.

Gentzkow, Matthew A. and Jesse M. Shapiro, "Media Bias and Reputation," *Journal of Political Economy*, 2006, *114* (2), 280–316.

_ **and** _ , "What Drives Media Slant? Evidence From U.S. Daily Newspapers," *Econometrica*, 2010, *78* (1), 35–71.

_ **, Michael B. Wong, and Allen T. Zhang**, "Ideological Bias and Trust in Information Sources," *Working Paper*, 2018.

Gentzkow, Matthew, Jesse M. Shapiro, and Michael Sinkinson, "Competition and Ideological Diversity: Historical Evidence from US Newspapers," *American Economic Review*, 2014, *104* (10), 3073–3114.

Golman, Russell, David Hagmann, and George Loewenstein, "Information Avoidance," *Journal of Economic Literature*, 2017, *55* (1), 96–135.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart, "Designing Information Provision Experiments," *CESifo Working Paper No. 8406*, 2020.

Hainmueller, Jens, Jonathan Mummolo, and Yiqing Xu, "How much should we trust estimates from multiplicative interaction models? Simple tools to improve empirical practice," *Political Analysis*, 2019, *27* (2), 163–192.

Horton, John J, David G Rand, and Richard J Zeckhauser, "The Online Laboratory: Conducting Experiments in a Real Labor Market," *Experimental Economics*, 2011, *14* (3), 399–425.

Jo, Donghee, "Better the Devil You Know: Selective Exposure Alleviates Polarization in an Online Field Experiment," Technical Report, Working Paper 2019.

King, Gary, Benjamin Schneer, and Ariel White, "How the news media activate public expression and influence national agendas," *Science*, 2017, *358* (6364), 776–780.

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

Mitchell, Amy, "A new understanding: What makes people trust and rely on news," Pew Research Center January 2018.

_ **and Rachel Weisel**, "Political Polarization & Media Habits," Pew Research Center October 2014.

Mullainathan, Sendhil and Andrei Shleifer, "The Market for News," *American Economic Review*, 2005, *95* (4), 1031–1053.

Mummolo, Jonathan and Erik Peterson, "Demand Effects in Survey Experiments: An Empirical Assessment," *American Political Science Review*, 2018, *113* (2), 517–529.

Nielsen, Kirby, "Preferences for the Resolution of Uncertainty and the Timing of Information," *Working Paper*, 2019.

Perego, Jacopo and Sevgi Yuksel, "Media Competition and Social Disagreement," *Working Paper*, 2018.

Pogorelskiy, Kirill and Matthew Shum, "News We Like to Share: How News Sharing on Social Networks Influences Voting Outcomes," *Working Paper*, 2019.

Qin, Bei, David Strömberg, and Yanhui Wu, "Media Bias in China," *American Economic Review*, 2018, *108* (9), 2442–2476.

Schwardmann, Peter and Joël van der Weele, "Deception and self-deception," *Nature Human Behavior*, 2019, *3* (10), 1055–1061.

Severin, Werner J. and James W. Tankard, *Communication Theories: Origins, Methods, and Uses in the Mass Media*, 5 ed., Addison Wesley Longman, 2000.

Strömberg, David, "Mass Media Competition, Political Competition, and Public Policy," *Review of Economic Studies*, 2004, *71* (1), 265–284.

_ , "Media Coverage and Political Accountability: Theory and Evidence," in Simon P. Anderson, Joel Waldfogel, and David Strömberg, eds., *Handbook of Media Economics*, Vol. 1, North-Holland, 2015, chapter 14, pp. 595–622.

Suen, Wing, "The Self-Perpetuation of Biased Beliefs," *The Economic Journal*, 2004, *114* (495), 377–396.

Sunstein, Cass R, *# Republic: Divided Democracy in the Age of Social Media*, Princeton University Press, 2018.

Szeidl, Adam and Ferenc Szucs, "Media Capture through Favor Exchange," *Working Paper*, 2017.

Tella, Rafael Di, Ricardo Perez-Truglia, Andres Babino, and Mariano Sigman, "Conveniently Upset: Avoiding Altruism by Distorting Beliefs About Others' Altruism," *American Economic Review*, 2015, *105* (11), 3416–3442.

Thaler, Michael, "The "Fake News" Effect: An Experiment on Motivated Reasoning and Trust in News," *Working Paper*, 2019.

Vosoughi, Soroush, Deb Roy, and Sinan Aral, "The spread of true and false news online," *Science*, 2018, *359* (6380), 1146–1151.

Wood, Thomas and Ethan Porter, "The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence," *Political Behavior*, 2019, *41* (1), 135–163.

Young, Eric, "A new understanding: What makes people trust and rely on news," American Press Institute 2016.

Zimmermann, Florian, "Clumped or Piecewise? Evidence on Preferences for Information," *Management Science*, 2015, *61* (4), 740–753.

Figures and tables

Figure 1: Design features: Filtering (Experiment 1.1)

Enter the experiment ($n = 4{,}631$)

Pre-treatment questions
(i) Socioeconomic questions
(ii) Newspapers most likely to read
(iii) Beliefs about a NYT article

Randomization
Stratified assignment
based on political affiliation

Control group
$n = 2{,}312$

Treatment group
$n = 2{,}319$

Information treatment
NYT reported both statistics.

**Cross-randomized
post-treatment outcomes**

Demand for news ($n = 3{,}864$)
Willingness to read NYT article

Beliefs about reporting ($n = 767$)
Belief about filtering and omission

Post-treatment beliefs and characteristics
(i) Beliefs about the NYT and the CBO
(ii) Beliefs about article characteristics
(iii) Additional demographic questions

Exit the experiment

Notes: This figure shows the main features of our filtering experiment conducted with Dynata (Experiment 1.1). Our other filtering experiments have a similar structure but do not cross-randomize post-treatment outcomes (Experiments 1.2, 1.3, 1.6, and 1.7; see Table 1).

Table 1: Overview of experiments

Experiment	Sample	Treatments Arms	Main Outcomes
Filtering			
Experiment 1.1 *Main filtering design* Wave 1: Jan 2019 Wave 2: Jan/Feb 2019 Wave 3: Feb 2019	Dynata: representative sample (region, income, gender, education, and age); $n = 4,631$	**Treatment**: Information about how the NYT covered the CBO report on the Health Bill **Control**: No information	Demand for reading a NYT article about the Tax bill; Post-treatment beliefs about reporting
Experiment 1.2 *Robustness curiosity* September 2019	Lucid: representative sample (region, income, gender, education, and age); $n = 3,387$	**Treatment**: Information about how the NYT covered the CBO report on the Minimum Wage Bill **Control**: No information	Demand for reading a NYT article about a single-payer health care system
Experiment 1.3 *Cognitive constraints placebo* April 2019	Dynata: representative sample (region, income, gender, and age); $n = 930$	**Treatment**: Information about how many statistics from the CBO report on the Health Bill the NYT reported **Control**: No information	Demand for reading a NYT article about the Tax Bill
Experiment 1.4 *Information demand* May 2019	Lucid: representative sample (region, income, gender, education, and age); $n = 703$	None	Demand for information about CBO estimates for the Tax Bill and the Health Bill
Experiment 1.5 *External validity* April 2019	MTurk: $n = 199$	None	Demand for reading a NYT article about the Tax Bill; Incentivized WTP for a digital NYT subscription
Experiment 1.6 *Incentives and reversed article order* September 2018	MTurk: Democrats and Democrat-leaning respondents; $n = 723$	**Treatment**: Information about how the NYT covered the CBO report on the Tax Bill **Control**: No information	Demand for reading a NYT article about the Health Bill
Experiment 1.7 *Platform robustness* January 2019	MTurk: $n = 1,332$	**Treatment**: Information about how the NYT covered the CBO report on the Health Bill **Control**: No information	Demand for reading a NYT article about the Tax Bill
False claims			
Experiment 2 *False claims* Wave 1: Sep 2019 Wave 2: Oct 2019	Lucid: representative sample (region, income, gender, education, and age); $n = 2,081$	**Treatment**: Information about a false statement in an *CNN* article **Control**: No information	Demand for reading a *CNN* news article about the Trump Impeachment
Deceptive spin			
Experiment 3 *Deceptive spin* April 2019	MTurk: $n = 1,503$	**Treatment**: Information about how the NYT covered the CBO report on the cost of the government shutdown (2019) **Control**: No information	Demand for reading a NYT article about the Tax Bill

Notes: This table provides an overview of all experiments. Wave 2 and 3 of Experiment 1.1 ($n = 4,025$) was registered in the AEA RCT Registry as trial 3855.

Table 2: Post-treatment beliefs: Filtering experiment

	Beliefs: Less suppression		Article characteristics			The New York Times			Congressional Budget Office		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Filtering	Omission	Quality	Dryness	Complex	No bias	Trust	Curious	No bias	Trust	Accuracy
Treatment	0.068***	0.033	0.103*	-0.004	0.048	0.041***	-0.018	-0.008	-0.013	0.017	0.063**
	(0.020)	(0.020)	(0.060)	(0.074)	(0.074)	(0.014)	(0.027)	(0.028)	(0.015)	(0.029)	(0.029)
N	749	742	737	737	737	4563	4547	4547	4523	4523	4523
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Z-scored	No	No	Yes	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Control group mean	0.479	0.528	0	0	0	0.376	0	0	0.533	0	0

Note: This table displays main treatment effects on a series of post-treatment beliefs using data from Experiment 1.1 (see Table 1). Columns 1 to 5 use respondents who were cross-randomized into not receiving the option to read an article in the *New York Times*, while columns 6 to 11 use all respondents. "Filtering" refers to the percent chance that the *New York Times* reported that the Trump Tax Plan would create 1.1 million jobs. "Omission" refers to the percent chance that the *New York Times* wrote an article about the CBO's analysis of granting citizenship to the Dreamers. "Quality" refers to people's perception of the quality of articles in the *New York Times*. "Dryness" captures people's perception of whether reporting of the *New York Times* is dry and technical. "Complex" measures people's perception of whether reporting of the *New York Times* is complex. "No bias" is a dummy variable taking value one if our respondents think that the *New York Times* is not politically biased (column 6), and is defined similarly for the CBO (column 9). "Trust" measures people's trust in the *New York Times* (column 7) and the CBO (column 10). "Curious" measures people's interest in learning whether the *New York Times* is biased. "Accuracy" measures people's perception of the accuracy of the forecasts of the CBO. The outcomes in columns 3, 4, 5, 7, 8, 10, and 11 are measured on five-point Likert scales and then z-scored. Regressions include the following controls: gender, age, income, region, race, education, employment status, frequency of reading the *New York Times*, pre-treatment beliefs about the probability that the *New York Times* would report both statistics, beliefs about the consequences of the policy bills, and experiment fixed effects.

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

Table 3: Treatment effects on demand for news: Filtering experiments

	(1) Main experiment	(2) Curiosity experiment	(3) Robustness experiments	(4) Pooled across all experiments
Treatment	-0.035**	-0.047***	-0.053***	-0.043***
	(0.014)	(0.015)	(0.019)	(0.009)
N	3858	3189	2169	9216
Controls	Yes	Yes	Yes	Yes
Control group mean	0.274	0.280	0.325	0.286

Note: This table shows OLS regressions where the dependent variable is an indicator that takes the value one for respondents who wanted to read an article in the *New York Times* about a CBO report. Column 1 includes respondents from our main experiment (conducted with Dynata; Experiment 1.1; see Table 1). Column 2 includes respondents from a robustness experiment to alleviate concerns about curiosity as a mechanism (conducted with Lucid; Experiment 1.2; see Table 1). Column 3 includes respondents from additional robustness experiments conducted on Amazon Mechanical Turk (Experiments 1.6 and 1.7; see Table 1). Column 4 pools all respondents from Columns 1 to 3. "Treatment" is an indicator that takes the value one for respondents who received information that the *New York Times* did not suppress any key facts from the CBO report. All regressions include the set of controls from Table 2.

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

Table 4: Treatment effects on beliefs and article demand: False claims experiment

	Beliefs: Less informative reporting in CNN					Perceptions of CNN			CNN
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	False claims: Trump	False claims: News	Unverified: Trump	Unverified: Biden	Intention: Hurt Trump	Quality	No bias	Trust	Article demand
Treatment	0.028***	0.066***	0.095***	0.057	0.033	-0.087***	-0.056***	-0.047*	0.008
	(0.009)	(0.025)	(0.029)	(0.042)	(0.027)	(0.027)	(0.019)	(0.026)	(0.018)
N	2069	2069	2069	2069	2069	2076	2069	2076	2081
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Z-scored	No	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Control group mean	0.489	0	0	0	0	0	0.382	0	0.237

Note: This table shows main treatment effects on a series of post-treatment beliefs and demand for a *CNN* article from OLS regressions. "Treatment" is an indicator that takes the value one for respondents who received information that the *CNN* falsely reported that the CIA decided to extract a spy from Russia because it feared that President Trump would mishandle classified information. "False claims: Trump" is the subjective percent chance that an article from *CNN* about President Trump would contain any false claims. "False claims: News" measures how often people think *CNN* makes false claims in its political reporting. "Unverified: Trump" is the belief about how likely *CNN* is to publish stories based on unverified and potentially misleading sources about Trump . "Unverified: Biden" is the analogous question about Joe Biden. "Intention: Hurt Trump" is the belief about whether *CNN* intentionally makes false claims to hurt President Trump. "Quality" refers to people's perception of the quality of articles in *CNN*. "No bias" is a dummy variable taking value one if our respondents think that the *CNN* is not politically biased. "Trust" measures people's trust in *CNN*. "Article demand" is an indicator variable that takes the value one for respondents who wanted to read an article from *CNN* about the impeachment process against President Trump. The outcomes in all columns except for columns 1, 7 and 9 are measured on five-point Likert scales and then z-scored. Regressions include the following controls: gender, age, income, region, race, education, employment status, political views, frequency of reading/watching *CNN*, pre-treatment beliefs about the probability that the *CNN* would cite Trump as the reason for the extraction, views on whether Trump deserves impeachment, pre-treatment beliefs about how often *CNN* makes false claims in its political reporting, beliefs about the acceptability of publishing false claims stories, and wave fixed effects.

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

For online publication only:
Do People Value More Informative News?

Felix Chopra, Ingar Haaland, and Christopher Roth

Summary of the Online Appendix

Section A contains proofs related to claims made in Section 2 and a discussion of prior beliefs about reporting and patterns of treatment effect heterogeneity.

Section B contains additional figures. Figure B.1 shows pre-treatment beliefs about information suppression in the *New York Times*, separately for Democrats and Republicans, for Experiment 1.1. Figure B.2 shows pre-treatment beliefs about the consequences of the Trump Tax Plan, separately for Democrats and Republicans, for Experiment 1.1. Figure B.3 shows treatment effect on demand for news for respondents in Experiment 1.1. Figure B.4 shows heterogeneity in treatment effects by prior beliefs separately for Democrats and Republicans for Experiment 1.1. Figure B.5 shows heterogeneity in treatment effects by priors separately for Democrats and Republicans for Experiment 2. Figure B.6 shows the demand for information from the CBO about the Trump Tax Plan and the Trump Healthcare Plan from Experiment 1.4.

Section C contains additional tables. Table C.1 provides the summary statistics. Table C.2 examines the integrity of randomization for the main filtering experiment 1.1. Table C.3 provides evidence on heterogeneity by political affiliation in experiment 1.1. Table C.4 shows heterogeneous effects by prior beliefs in experiment 1.1. Table C.5 provides further heterogeneity analysis to rule out potential confounds in Experiment 1.1. Table C.6 shows the treatment effects separately for the filtering robustness experiment (Experiments 1.6 and 1.7). Table C.7 shows the main treatment effects in the placebo experiment (Experiment 1.3). Table C.8 shows heterogeneous treatment effects by political affiliation on beliefs and article demand. Table C.9 provides evidence on heterogeneous effects by people's beliefs about whether Trump is guilty in the false claims experiment

1

(Experiment 2).

Section D describes the experimental design and the results for the experiment on deceptive spin. Table D.1 provides an overview of the main results. Table D.2 shows heterogeneous effects by prior beliefs in the deceptive spin experiment 3. Figure D.1 shows heterogeneity in treatment effects by priors separately for Democrats and Republicans for Experiment 3.

Section E contains screenshots of the instructions for all experiments. Table E.1 provides an overview of variables collected by experiment. Section E.1 shows the full set of experimental instructions for Experiments 1.1 and 1.7. In Section E.2, we provide the instructions for Experiment 1.2. In Section E.3, we show the instructions for Experiment 1.3. Section E.4 provides instructions for experiment 1.4. In Section E.5, we show the instructions for Experiment 1.5. In Section E.6, we show the instructions for Experiment 1.6. Section E.7 shows the instructions for Experiment 2. In Section E.8, we show the instructions for Experiment 3.

A Model appendix

A.1 Proofs

The following proposition provides the theoretical justification for the claims made in Section 2 on how strategic information suppression affects the Blackwell informativeness of news.

Proposition 1 (Filtering). Using the notation from Section 3, fix $s = \{s_1, \ldots, s_K\} \in S$ and two reporting strategies $\rho, \rho' : S \to \Delta(N)$. Let $\sigma, \sigma' : \Theta \to \Delta(N)$ be the information structures induced by combining the source signal $\pi : \Theta \to \Delta(S)$ with the reporting strategies, respectively. Suppose that (i) $\rho(t \mid s) \leq \rho'(t \mid s)$ for all $t \subsetneq s$, (ii) $\rho(s \mid s) > \rho'(s \mid s)$, and that (iii) $\rho(\cdot \mid s') = \rho'(\cdot \mid s')$ for all $s' \neq s$. Then the information structure σ is Blackwell more informative than σ'.

Proof. It suffices to show that the conclusion obtains if we strengthen the assumption by additionally assuming that $\rho(t \mid s) < \rho'(t \mid s)$ for some $t \subsetneq s$ and that for all other $t' \subsetneq s$ with $t' \neq t$, we have $\rho(t' \mid s) = \rho'(t' \mid s)$. The general case then follows by applying the result to the sequence $\rho = \rho_1, \ldots, \rho_L = \rho'$ where ρ_k and ρ_{k+1} differ at most on the set $\{s, s'\}$ for some $s' \subseteq s$ and $L = |\mathscr{P}(s)|$. Suppose that $n \in N$ is a random variable with state-dependent distribution σ. To show that σ is Blackwell more informative than σ', it suffices to construct an n-measurable random variable $n' \in N$ with state-dependent distribution σ', thereby establishing statistical sufficiency. We construct n' as follows: let $n' = n$ whenever $n \neq s$ and set $\beta = \rho'(s \mid s)/\rho(s \mid s)$. If $n = s$, then n' takes value s with probability β and value t with probability $1 - \beta$. One can then verify that conditional on the state $\theta \in \Theta$, the distribution of n' is $\sigma'(\cdot \mid \theta)$. This concludes the proof. $\qquad\square$

The following proposition provides the theoretical justification for the claims made in Section 4.1 that increasing the likelihood of false claims reduces the Blackwell informativeness of a newspaper's article.

Proposition 2 (False claims). Suppose there is a hidden state $\theta \in \Theta = \{L, R\}$. Let σ_θ

denote the probability of a news article reporting R if the state is θ and consider any triple $\sigma_L < \sigma_L' < \sigma_R$. Then a newspaper article n with distribution (σ_L, σ_R) is Blackwell more informative than an article n' with distribution (σ_L', σ_R).

Proof. Consider any triple $\sigma_L < \sigma_L' < \sigma_R$. We have to show that there is a stochastic transformation τ such that the distribution of $\tau(n)$ is (σ_L', σ_R). This amounts to finding two probabilities $\alpha, \beta \in [0, 1]$ such that we obtain the desired distributional equivalence if we set $\alpha = P(n' = R \mid n = R)$ and $\beta = P(n' = R \mid n = L)$. The probabilities α and β are the solution to the following two equations:

$$\sigma_L' = \alpha \sigma_L + \beta(1 - \sigma_L) \tag{2}$$

$$\sigma_R = \alpha \sigma_R + \beta(1 - \sigma_R) \tag{3}$$

We can algebraically verify that a solution is given by

$$\alpha = 1 - \left(\frac{\sigma_L' - \sigma_L}{\sigma_R - \sigma_L} \right)(1 - \sigma_R) \quad \text{and} \quad \beta = \left(\frac{\sigma_L' - \sigma_L}{\sigma_R - \sigma_L} \right) \sigma_R. \tag{4}$$

Moreover, $\sigma_L < \sigma_L' < \sigma_R$ implies that $0 \leq \beta \leq \alpha \leq 1$. This concludes the proof. $\qquad\square$

A.2 Heterogeneity by beliefs about reporting

This section discusses the difficulties of making predictions about the patterns of heterogeneity with respect to prior beliefs about reporting in our experimental design. For simplicity, suppose there is a binary state space $\Theta = \{L, R\}$ and focus on a consumer who identifies as Democrat. This consumer holds a belief $q > \frac{1}{2}$ that $\theta = L$, and expects a left-leaning outlet to provide a confirmatory signal with probability $p \in [0, 1]$. In the absence of a standard model of confirmatory preferences, we make the reduced form assumption that the utility from news consumption $U(p, q)$ depends of the consumer's prior belief about the state θ and the belief that the newspaper reports L. The consumer prefers to read the article N if $U(p, q) > v$ where v is the value of his outside option.

Models of confirmatory preferences would plausibly predict that $\frac{\partial U}{\partial p} \geq 0$ for consumers with prior belief $q > 0.5$. Thus, an information treatment that weakly decreases p will reduce the expected utility from news consumption, but affects the demand for news only for marginal consumers with U sufficiently close to v. However, without additional assumptions on (i) the marginal utility from p, (ii) the value of the outside option, and (iii) the extent of Bayesian updating about p, the theoretical prediction for the interaction effect between the information treatment and the prior belief p are ambiguous.

Figure A.1: Utility from news and beliefs about reporting

To illustrate this point, Figure A.1 provides a stylized example of the utility from news consumption. Suppose the consumer starts out with a prior p_0 and strongly revises her belief to p_1 after receiving the treatment. While her utility U strongly declines, she also started from a higher level, and thus still prefers reading the article. However, the treatment could move respondents with a prior belief p_1 below the cutoff p^*. This example thus predicts a U-shaped pattern of treatment effect heterogeneity by people's prior beliefs about reporting: People with very low or very high beliefs that a newspaper reports L will not adjust their demand (no treatment effect), whereas people with moderate priors will reduce their demand (negative treatment effect).

B   Additional figures

Figure B.1: Distribution of pre-treatment beliefs about information suppression in the *New York Times*, by political views

Notes: This figure uses data from Experiment 1.1 (see Table 1). It shows pre-treatment beliefs about information suppression in the *New York Times* separately for Democrats (including Independents leaning toward the Democratic Party) and Republicans (including Independents leaning toward the Republican Party). **Panel A** shows data on beliefs about whether the *New York Times* slants to left (by suppressing the statistic on the positive fiscal consequences of the Trump Healthcare Plan). **Panel B** shows data on beliefs about whether the *New York Times* slants to right (by suppressing the statistic on the negative social consequences of the Trump Healthcare Plan).

Figure B.2: Distribution of pre-treatment beliefs about the impact of the Trump Tax Plan, by political views
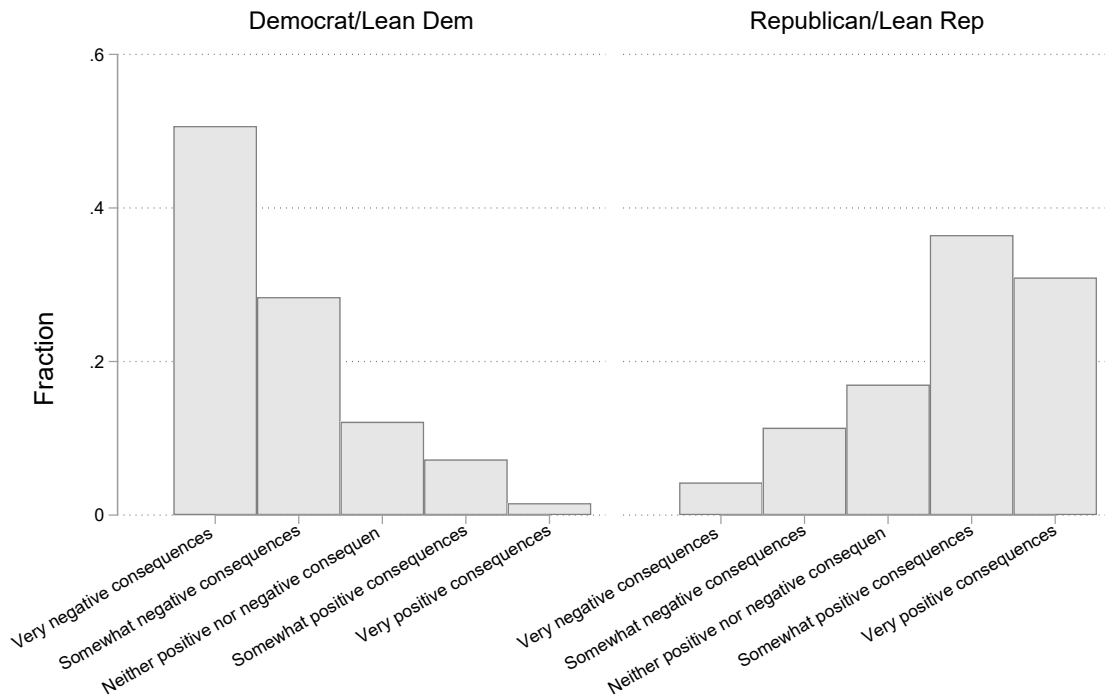
Notes: This figure uses data from Experiment 1.1 (see Table 1). It shows pre-treatment beliefs about the consequences of the Trump Tax Plan separately for Democrats (including Independents leaning toward the Democratic Party) and Republicans (including Independents leaning toward the Republican Party).

Figure B.3: Demand for news: Treatment versus control

N=3858. P-value (equal means) = 0.046

Notes: This figure uses data from Experiment 1.1 (see Table 1). It shows, separately for respondents in
the treatment and control group, the fraction of respondents who wanted to read an article in the *New
York Times* about the Trump Tax Plan. Error bars indicate the standard error of the mean.

Figure B.4: Treatment effects by pre-treatment beliefs: Filtering
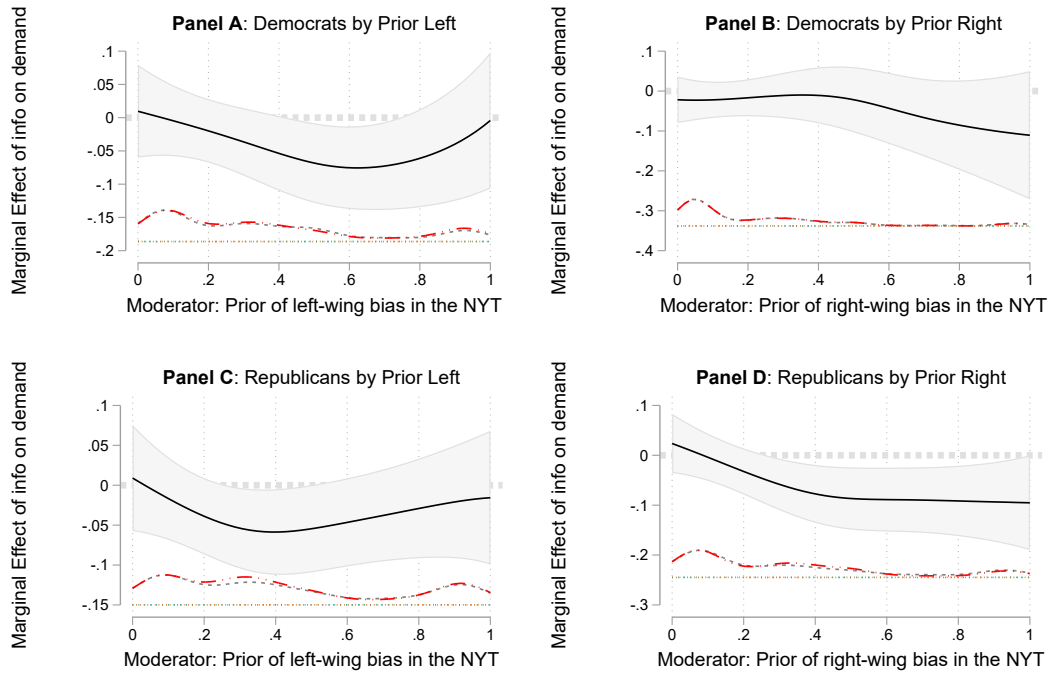
Notes: This figure displays heterogeneous treatment effects by respondents' prior beliefs about how the *New York Times* covered a CBO report about the consequences of the Trump Healthcare Plan (Hainmueller et al., 2019). "Prior left" is the percent chance (from 0 to 1) that the *New York Times* wrote an article about the CBO estimates on the Trump Healthcare Plan that mentioned only the number of people who would lose health coverage. "Prior right" is the percent chance (from 0 to 1) that the *New York Times* wrote an article about the CBO estimates on the Trump Healthcare Plan that mentioned only the impact on the federal debt. This figures uses respondents from experiment 1.1 (see Table 1). Democrats included in this figure are respondents who either identify with the Democratic Party or identify as Independents leaning toward the Democratic Party, and excludes respondents who think that the Trump Tax Plan will have somewhat positive or very positive consequences. Republicans included in this figure are respondents who either identify with the Republican Party or identify as Independents leaning toward the Republican Party, and excludes respondents who think that the Trump Tax Plan will have somewhat negative or very negative consequences. The dashed red curve shows smoothed kernel density estimates for the moderating prior variable.

Figure B.5: Treatment effects by pre-treatment beliefs: False claims

Panel A: Democrats   **Panel B**: Republicans

Notes: This figure displays heterogeneous treatment effects by respondents' belief about false claims in *CNN* articles (Hainmueller et al., 2019). The figures uses respondents from Experiment 2 (the false claims experiment; see Table 1). "Prior of *CNN* making false claims" is respondents' pre-treatment beliefs (from 0 to 1) that *CNN* falsely reported that the CIA made the decision to extract the spy because it feared that President Trump would mishandle classified information and their political affiliation. Democrats included in this figure are respondents who either identify with the Democratic Party or identify as Independents leaning toward the Democratic Party. Republicans included in this figure are respondents who either identify with the Republican Party or identify as Independents leaning toward the Republican Party. The dashed red curve shows smoothed kernel density estimates for the moderating prior variable.

10

Figure B.6: Demand for information about CBO statistics

Panel A: Demand for info: Tax Plan   **Panel B:** Demand for info: Healthcare Plan

Democrats/lean Democratic   Republicans/Lean Republican

Notes: This figure uses data from an experiment with Lucid (Experiment 1.4, see Table 1). The figure shows, separately for Democrats/Democrat-leaners and Republicans/Republican-leaners, the fraction of respondents who wanted information about different statistics from the CBO reports. Specifically, respondents were either asked about their demand for information about the Trump Tax Plan (see Panel A) or about their demand for information about the Trump Healthcare Plan (see Panel B). Respondents were then asked separately and in randomized order for each of the two headline statistics from the respective CBO report, whether they want to receive the CBO's point estimate or not. However, respondents do not know anything about the value of the statistics prior to making the decision. If they selected "Yes", we provided them with the information at the end of the survey. "Debt" is the share of respondents who want to learn how the Trump Tax Plan will affect the federal debt. "Jobs" is the share of respondents who want to learn how the Trump Tax Plan will affect the number of jobs. "Deficit" is the share of respondents who want to learn how the Trump Healthcare Plan will affect the federal deficit. "Uninsured" is the share of respondents who want to learn how the Trump Healthcare Plan will affect the number of people with health coverage.

11

C Additional tables

Table C.1: Summary statistics

	(1) Main exp.	(2) Curiosity	(3) Robustness	(4) Spin	(5) False claims	(6) Cognitive
Male	0.442	0.459	0.508	0.446	0.421	0.463
Age (midpoint)	48.546	42.712	36.063	39.249	46.170	44.774
White	0.871	0.779	0.797	0.783	0.772	0.846
Log income	3.355	3.099	3.349	3.333	3.353	3.499
College education	0.450	0.301	0.675	0.662	0.554	0.556
Full-time work	0.392	0.477	0.782	0.685	0.483	0.490
Northeast	0.195	0.161	0.217	0.178	0.216	0.218
Midwest	0.245	0.232	0.225	0.209	0.255	0.218
West	0.192	0.169	0.147	0.254	0.179	0.208
South	0.369	0.438	0.411	0.359	0.350	0.356
Republican	0.325	0.300	0.163	0.250	0.308	0.320
Democrat	0.340	0.355	0.553	0.437	0.399	0.343
Observations	4625	3189	2169	1503	2081	930

Note: This table displays the mean value of basic covariates for each experiment (see Table 1 for an overview). Specifically, "Main exp." refers to experiment 1.1, "Curiosity" refers to experiment 1.2, "Robustness" refers to experiment 1.6 and 1.7, "Spin" refers to experiment 3, "False Claims" refers to experiment 2, and "Cognitive" refers to experiment 1.3. "Male" is a binary variable with value one for male respondents. "Age" is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). "White" is a binary variable with value one if the respondent selected "Caucasian/White". "Log income" is coded continuously as the logarithm of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "College education" is a binary dummy variable taking value one if the respondent selected "Some college, no degree", "Associates degree", "Bachelor's degree", or "Post-graduate degree". "Full-time work" is a binary dummy variable taking value one if the respondent is working full-time. "Northeast", "Midwest", "West" and "South" are binary dummy variables with value one if the respondent lives in the respective region. "Republican" and "Democrat" are binary dummy variables with value one if the respondent identifies as Republican or Democrat.

Table C.2: Test of balance

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Gender	0.45	0.44	0.703	4631
Age	48.30	48.78	0.314	4631
Log income	10.91	10.88	0.248	4025
South	0.36	0.37	0.473	4631
West	0.20	0.18	0.196	4631
Northeast	0.19	0.20	0.216	4631
Republicans	0.33	0.32	0.907	4631
Democrats	0.34	0.34	0.916	4631
White	0.87	0.87	0.963	4458
College education	0.45	0.45	0.563	4488

Notes: This table provides a balance test for the main filtering experiment (Experiment 1.1; see Table 1). The *p*-value of a joint F-test regressing the treatment indicator on a series of observables is given by $p = 0.62$. "Gender" is a binary variable with value one for male respondents. "Age" is coded as the continuous midpoint of the age bracket (18 to 24, 25 to 34, 35 to 44, 45 to 54, 55 to 64, 65 or older). "Log income" is coded continuously as the logarithm of the income bracket's midpoint (Less than $15,000, $15,000 to $24,999, $25,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $149,999, $150,000 to $200,000, $200,000 or more). "South", "West", and "Northeast" are binary dummy variables with value one if the respondent lives in the respective region. "Republican" and "Democrat" are binary dummy variables with value one if the respondent identifies as Republican or Democrat. "White" is a binary variable with value one if the respondent selected "Caucasian/White". "College education" is a binary dummy variable taking value one if the respondent selected "Some college, no degree", "Associates degree", "Bachelor's degree", or "Post-graduate degree".

Table C.3: Political heterogeneity in treatment responses: Filtering experiment

	Beliefs: Less suppression		Article characteristics			The New York Times			Congressional Budget Office			NYT
	(1) Filtering	(2) Omission	(3) Quality	(4) Dryness	(5) Complex	(6) No bias	(7) Trust	(8) Curious	(9) No bias	(10) Trust	(11) Accuracy	(12) Demand
Treatment	0.027	0.048*	0.122*	0.033	0.074	0.055***	-0.011	-0.012	0.007	0.042	0.109***	-0.035*
	(0.028)	(0.026)	(0.072)	(0.091)	(0.096)	(0.020)	(0.036)	(0.037)	(0.020)	(0.043)	(0.042)	(0.020)
Treatment × Rep.	0.083**	-0.033	-0.067	-0.080	-0.050	-0.028	-0.017	0.008	-0.041	-0.049	-0.095	-0.001
	(0.040)	(0.040)	(0.115)	(0.147)	(0.147)	(0.027)	(0.054)	(0.055)	(0.029)	(0.057)	(0.058)	(0.027)
N	749	742	737	737	737	4563	4547	4547	4523	4523	4523	3858
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Z-scored	No	No	Yes	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No
Control group mean	0.479	0.528	0	0	0	0.376	0	0	0.533	0	0	0.274

Note: This table displays heterogeneous treatment effects by political affiliation on a series of post-treatment beliefs in addition to article demand using data from Experiment 1.1 (see Table 1). "Rep." is an indicator that takes the value one for self-identified Republicans and Independents who lean toward the Republican Party. Columns 1 to 5 use respondents who were cross-randomized into not receiving the option to read an article in the *New York Times*, while column 6 to 11 use all respondents. "Filtering" refers to the percent chance that the *New York Times* reported that the Trump Tax Plan would create 1.1 million jobs. "Omission" refers to the percent chance that the *New York Times* wrote an article about the CBO's analysis of granting citizenship to the dreamers. "Quality" refers to people's perception of the quality of articles in the *New York Times*. "Dryness" captures people's perception of whether reporting of the *New York Times* is dry and technical. "Complex" measures people's perception of whether reporting of the *New York Times* is complex. "No bias" (column 6) is a dummy variable taking value one if our respondents think that the *New York Times* is not politically biased. "Trust" (column 7) measures people's trust in the *New York Times*. "Curious" measures people's interest in learning whether the *New York Times* is biased. "Accuracy" measures people's perception of the accuracy of the forecasts of the CBO. "No bias" (column 9) measures people's perception of whether the CBO is biased. "Trust" (column 10) measures people's trust in the CBO. "Demand" is a dummy variable taking value one if our respondents wanted to read an article in the *New York Times* about the Trump Tax Plan. The outcomes in columns 3, 4, 5, 7, 8, 10, and 11 are measured on five-point Likert scales and then z-scored. All regressions include the set of controls from Table 2.

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

Table C.4: Heterogeneity by beliefs about reporting: Filtering experiments

	Democrats				Republicans			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-0.032	-0.016	-0.015	0.012	-0.040**	-0.043	-0.006	0.036
	(0.021)	(0.030)	(0.026)	(0.037)	(0.020)	(0.028)	(0.026)	(0.045)
Treatment × Prior Left		-0.049		-0.068		0.007		-0.073
		(0.061)		(0.063)		(0.053)		(0.064)
Treatment × Prior Right			-0.098	-0.118			-0.120**	-0.165**
			(0.080)	(0.083)			(0.054)	(0.065)
N	1796	1796	1796	1796	1589	1589	1589	1589
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean	0.317	0.317	0.317	0.317	0.227	0.227	0.227	0.227

Note: This table shows OLS regressions where the dependent variable is an indicator that takes the value one for respondents who wanted to read an article in the *New York Times* about a CBO report. The sample includes respondents from our main filtering experiment (Experiment 1; see Table 1). Columns 1 to 4 use respondents who either identify with the Democratic Party or identify as Independents leaning toward the Democratic Party, and excludes respondents who think that the Trump Tax Plan will have somewhat positive or very positive consequences. Columns 5 to 8 use respondents who either identify with the Republican Party or identify as Independents leaning toward the Republican Party, and excludes respondents who think that the Trump Tax Plan will have somewhat negative or very negative consequences. "Prior Left" is the percent chance (from 0 to 1) that the *New York Times* wrote an article about the CBO estimates on the Trump Healthcare Plan that mentioned only the number of people who would lose health coverage. "Prior Right" is the percent chance (from 0 to 1) that the *New York Times* wrote an article about the CBO estimates on the Trump Healthcare Plan that mentioned only the impact on the federal debt. All regressions include the set of controls from Table 2.

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

Table C.5: Treatment effects on demand for news in the filtering experiment: Interaction effects

	(1)	(2)
Treatment	-0.027	-0.041**
	(0.018)	(0.017)
Interactant	0.015	0.042*
	(0.020)	(0.022)
Treatment × Interactant	-0.018	0.015
	(0.028)	(0.028)
Interactant	Portfolio	College
N	3858	3858
Controls	Yes	Yes

Note: This table uses data from our main filtering experiment (Experiment 1.1; see Table 1). This table displays heterogeneous treatment effects on people's demand for reading an article in the *New York Times*. "Portfolio" takes value one for respondents who read both at least one left-wing newspaper and one right-wing newspaper. "College" takes value 1 for respondents who received at least some college education. All regressions include the set of controls from Table 2.

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

Table C.6: Treatment effects on demand for news: Robustness experiments

	(1) **Experiment 1.6** Incentives and reverse order	(2) **Experiment 1.7** Platform robustness
Treatment	-0.040	-0.063***
	(0.036)	(0.023)
N	752	1417
Controls	Yes	Yes
Control group mean	0.392	0.289

Note: This table shows OLS regressions where the dependent variable is an indicator that takes the value one for respondents who wanted to read an article in the *New York Times* (see Table 1). Regressions include the following controls: gender, age, income, region, race, education, employment status, frequency of reading the *New York Times*, pre-treatment beliefs about the probability that the *New York Times* would report both statistics, and beliefs about the consequences of the policy bills.

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

Table C.7: Treatment effects: Placebo experiment

	(1) Article demand	(2) Quality	(3) No bias
Treatment	0.015	0.099*	0.019
	(0.029)	(0.054)	(0.030)
N	930	928	928
Controls	Yes	Yes	Yes
Z-scored	No	Yes	No
Control group mean	0.265	0	0.384

Note: This table shows OLS regressions using data from the "Placebo experiment" conducted with Dynata (Experiment 1.3; see Table 1). "Treatment" is an indicator taking the value one for respondents who were informed that the *New York Times* reported two out of two statistics from the CBO report. "Article demand" is a binary variable with value one if the respondent wanted to read the *New York Times* article about the Trump Tax Plan. "Quality" refers to perceptions of quality in the *New York Times* and is measured on a 5-point Likert scale and then z-scored by the mean and standard deviation of control group respondents. "No bias" is a binary variable with value one if the respondent thinks that the *New York Times* is not politically biased. Regressions include the following controls: gender, age, income, region, race, education, employment status, and prior beliefs that the *New York Times* would report both statistics.

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

Table C.8: Political heterogeneity in treatment effects on beliefs and article demand: False claims experiment

	Beliefs: Less informative reporting in CNN					Perceptions of CNN			CNN
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	False claims: Trump	False claims: News	Unverified: Trump	Unverified: Biden	Intention: Hurt Trump	Quality	No bias	Trust	Article demand
Treatment	0.032***	0.062*	0.120***	0.128**	0.046	-0.105***	-0.078***	-0.086**	-0.003
	(0.012)	(0.033)	(0.044)	(0.050)	(0.040)	(0.035)	(0.028)	(0.035)	(0.025)
Treatment × Republican/lean Rep	-0.013	0.006	-0.063	-0.159*	-0.035	0.043	0.050	0.092*	0.023
	(0.018)	(0.050)	(0.058)	(0.087)	(0.054)	(0.054)	(0.036)	(0.053)	(0.036)
N	2069	2069	2069	2069	2069	2076	2069	2076	2081
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Z-scored	No	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Control group mean	0.489	0	0	0	0	0	0.382	0	0.237

Note: This table shows main treatment effects on a series of post-treatment beliefs in addition to article demand (OLS regressions), using data from experiment 2 (see Table 1). "Treatment" is an indicator that takes the value one for respondents who received information that the *CNN* falsely reported that the CIA decided to extract a spy from Russia because it feared that President Trump would mishandle classified information. "Republican/Lean Rep" is an indicator that takes the value one for respondents who identify with the Republican Party or identify as Independents leaning toward the Republican Party. "False claims: Trump" is the subjective percent chance that an article in *CNN* about President Trump would contain any false claims. "False claims: News" measures how often people think *CNN* makes false claims in its political reporting. "Unverified: Trump" is beliefs about how likely *CNN* is to publish stories based on unverified and potentially misleading sources about President Trump. "Unverified: Biden" is the analogous question about Joe Biden. "Intention: Hurt Trump" is beliefs about whether *CNN* intentionally makes false claims to hurt Trump. "Quality" refers to people's perception of the quality of articles in *CNN*. "No bias" is a dummy variable taking value one if our respondents think that *CNN* is not politically biased. "Trust" measures people's trust in *CNN*. "Article demand" is an indicator variable that takes the value one for respondents who wanted to read an article in the *CNN* about the impeachment process against President Trump. The outcomes in all columns except for columns 1, 7 and 9 are measured on five-point Likert scales and then z-scored. Regressions include the same set of controls as Table 4.

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

Table C.9: Heterogeneity in treatment effects on beliefs and article demand by views on impeachment: False claims experiment

	Beliefs: Less informative reporting in CNN					Perceptions of CNN			CNN
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	False claims: Trump	False claims: News	Unverified: Trump	Unverified: Biden	Intention: Hurt Trump	Quality	No bias	Trust	Article demand
Treatment	0.030**	0.102***	0.097***	-0.021	0.060*	-0.098**	-0.043*	-0.038	-0.013
	(0.012)	(0.037)	(0.038)	(0.071)	(0.036)	(0.042)	(0.024)	(0.040)	(0.026)
Treatment × Support Impeachment	-0.003	-0.062	-0.004	0.138	-0.047	0.019	-0.023	-0.016	0.038
	(0.017)	(0.050)	(0.057)	(0.087)	(0.053)	(0.054)	(0.036)	(0.053)	(0.036)
N	2069	2069	2069	2069	2069	2076	2069	2076	2081
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Z-scored	No	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Control group mean	0.489	0	0	0	0	0	0.382	0	0.237

Note: This table shows main treatment effects on a series of post-treatment beliefs in addition to article demand (OLS regressions), using data from experiment 2 (see Table 1). "Treatment" is an indicator that takes the value one for respondents who received information that the *CNN* falsely reported that the CIA decided to extract a spy from Russia because it feared that President Trump would mishandle classified information. "Support Impeachment" is an indicator that takes the value one for respondents who said "Yes" to the question of whether President Trump should be impeached (this question was elicited pre-treatment). "False claims: Trump" is the subjective percent chance that an article in *CNN* about President Trump would contain any false claims. "False claims: News" measures how often people think *CNN* makes false claims in its political reporting. "Unverified: Trump" is beliefs about how likely *CNN* is to publish stories based on unverified and potentially misleading sources about President Trump. "Unverified: Biden" is the analogous question about Joe Biden. "Intention: Hurt Trump" is beliefs about whether *CNN* intentionally makes false claims to hurt Trump. "Quality" refers to people's perception of the quality of articles in *CNN*. "No bias" is a dummy variable taking value one if our respondents think that *CNN* is not politically biased. "Trust" measures people's trust in *CNN*. "Article demand" is an indicator variable that takes the value one for respondents who wanted to read an article in the *CNN* about the impeachment process against President Trump. The outcomes in all columns except for columns 1, 7 and 9 are measured on five-point Likert scales and then z-scored. Regressions include the same set of controls as Table 4.

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

D   Experiment on deceptive spin

The filtering and false claims experiments leveraged variation in the quantity of information and the truthfulness of statements. We now consider a more subtle but prevalent form of media bias, which weakly decreases the informativeness of news, "deceptive spin."[1] By deceptive spin we mean newspapers' tendency to emphasize facts that favor a particular interpretation of an event which may be misleading. Specifically, we exogenously vary people's perceptions of the extent to which the *New York Times* spins news reports in politically biased way by using potentially misleading statistics. We argue that the "more information-is-better" principle would predict that people should reduce their demand for news when realizing that a newspaper provided a misleading statistic.

Experimental design and sample   We recruited approximately 1,500 respondents from Amazon Mechanical Turk in April 2019 (Experiment 3 in Table 1). We first elicit people's beliefs about the economic consequences of the five-week government shutdown that started after disagreement about funding for the proposed U.S.–Mexico wall and ended in January 2019. We then inform our respondents that the CBO analyzed the economic consequences of the shutdown and concluded that it had cost the U.S. economy $11 billion, of which $3 billion would be permanently lost. Since President Donald Trump was largely seen as responsible for the shutdown,[2] newspapers could spin the story in a left-wing way by emphasizing the "high" short-term cost estimate ($11 billion) or in a right-wing way by emphasizing the "low" long-term cost estimate ($3 billion).[3]

---

[1]While the theoretical link to informativeness is not as strong as in the previous two cases, the prevalence of deceptive spin provides a strong empirical justification for studying how the demand for news responds to perceptions of spin.

[2]See, for instance, the following article: https://www.bloomberg.com/news/articles/2019-01-14/trump-took-responsibility-for-shutdown-and-voters-give-it-to-him (accessed April 08, 2019).

[3]Newspapers writing about the findings emphasized different statistics in their headlines. For instance, the *Wall Street Journal*, which leans center-right, reported only the $3 billion statistic in its headline ("CBO: Shutdown Will Cost Government $3 Billion of Projected 2019 GDP," https://www.wsj.com/articles/cbo-shutdown-will-cost-government-3-billion-of-projected-2019-gdp-11548688574), whereas the *New York Times*, which leans center-left, only reported the $11 billion statistic ("Government Shutdown Cost U.S. Economy $11 Billion, C.B.O. Says," https://nyti.ms/2S75xrK). Both articles were accessed on April 08, 2019.

We thereafter inform our respondents that the *New York Times* wrote an article about the CBO findings using one of the following two headlines: (i) "Government Shutdown Cost U.S. Economy $11 Billion, C.B.O. Says"; (ii) "Government Shutdown Cost U.S. Economy $3 Billion, C.B.O. Says". We then ask respondents to state the percent chance they assign to the *New York Times* using each headline for its story about the CBO estimates (which needed to add up to 100 percent). To introduce exogenous variation in people's belief about left-wing spin in the *New York Times*, we then inform a random subset of respondents that the *New York Times* used the headline featuring the $11 billion statistic. We use the same behavioral outcome measure as in the filtering experiment, namely people's demand for a *New York Times* article about the consequences of the Trump Tax Plan. We also ask a series of post-treatment questions such as perceptions of quality and bias, or trust in the *New York Times*.

Results   Table D.1 shows that the information provision about the headline changed people's perceptions of bias and quality of news reporting. Treated respondents perceive the quality of reporting as lower (column 1) and think that the *New York Times* is more politically biased (column 3). If anything, both Democrats and Republicans perceive the *New York Times* as less trustworthy (columns 6 and 10). This suggests that both Democrats and Republicans alike perceive the long-term permanent cost of the shutdown as the more relevant and more informative measure of economic costs.

On average, we observe an insignificant 2 percentage point decrease in demand for news (column 4), which is driven by Republican respondents (column 12). For Democrats, however, we see no average decrease in demand for news (column 8), even though they think that the quality of reporting in the *New York Times* is lower. The lack of a decrease among Democrats is inconsistent with the predictions of "more-information-is-better principle."

Heterogeneity by belief confirmation   Our treatment should (weakly) increase respondents' belief that the *New York Times* slants to the left. This in turn means that treated

22

Democratic respondents with a more biased pre-treatment belief about the slant of the *New York Times* will learn that the newspaper is more likely to confirm their existing beliefs; on the other hand, treated Republican respondents with more biased pre-treatment beliefs about the slant of the *New York Times* will learn that the *New York Times* is less likely to confirm their existing beliefs. We non-parametrically examine treatment effects by prior beliefs separately for Republicans and Democrats in Figure D.1. We find patterns consistent with respondents having a preference for belief confirmation. The treatment polarizes news consumption between Democrats and Republicans. Democrats who initially placed a higher percent chance that the *New York Times* slants to the right do not decrease their demand for news, while Republicans who initially placed a higher percent chance that the *New York Times* slants to the right strongly decrease their demand once they learn that the *New York Times* is more likely to use a left-wing biased headline.

Table D.1: Treatment effects on demand for news: Spin experiment

	Full sample				Democrats/lean Democrat				Republicans/lean Republican			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Quality	Trust	No bias	Demand	Quality	Trust	No bias	Demand	Quality	Trust	No bias	Demand
Treatment	-0.098**	-0.065	-0.040*	-0.020	-0.103**	-0.070	-0.058*	-0.003	-0.093	-0.080	-0.009	-0.048
	(0.044)	(0.042)	(0.023)	(0.023)	(0.050)	(0.049)	(0.032)	(0.030)	(0.078)	(0.073)	(0.030)	(0.034)
N	1498	1500	1500	1503	929	930	930	931	569	570	570	572
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Z-scored	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes	No	No
Control group mean	0	0	0.340	0.290	0.330	0.350	0.450	0.320	-0.580	-0.620	0.170	0.250

Note: This table displays main treatment effects on article demand and a series of post-treatment beliefs for the spin experiment (Experiment 3; see Table 1). Columns 1 to 4 use the full sample. Columns 5 to 8 use respondents who identify with the Democratic Party and respondents who identify as Independents but lean toward the Democratic Party. Column 9 to 12 use respondents who identify with the Republican Party and respondents who identify as Independents but lean toward the Republican Party. "Treatment" is an indicator that takes the value one for respondents who received information that the *New York Times* used a headline featuring the $11 billion short-term cost (rather than the $3 billion long-term cost) of the government shutdown. "Quality" refers to people's perception of the quality of articles in the *New York Times*. "Trust" measures people's trust in the *New York Times*. The quality and trust outcomes are measured on five-point Likert scales and then z-scored using the mean and standard deviation for control group respondents. "No bias" is a dummy variable taking value one if our respondents think that the *New York Times* is not politically biased. "Demand" is an indicator variable that takes the value one for respondents who wanted to read an article in the *New York Times* about the G.O.P. Tax Bill. Regressions include the following controls: gender, age, income, region, race, education, employment status, frequency of reading the *New York Times*, pre-treatment beliefs about the probability that the *New York Times* would use a headline featuring the $3 billion number, and pre-treatment beliefs about the consequences of the shutdown.

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

Table D.2: Political heterogeneity in demand for news: Spin experiment

	Democrats/lean Dem		Republicans/lean Rep	
	(1)	(2)	(3)	(4)
Treatment	-0.003	-0.098**	-0.048	0.003
	(0.030)	(0.049)	(0.034)	(0.047)
Treatment × Prior: 3 billion		0.303**		-0.198
		(0.126)		(0.141)
N	931	931	572	572
Controls	Yes	Yes	Yes	Yes
Control group mean	0.318	0.318	0.245	0.245

Note: This table displays OLS regressions using respondents from Experiment 3 (see Table 1). The dependent variable in all columns is an indicator that takes the value one for respondents who wanted to read an article in the *New York Times* about the Trump Tax Plan. "Treatment" is an indicator that takes the value one for respondents who received information that the *New York Times* used a headline featuring the $11 billion short-term cost (rather than the $3 billion long-term cost) of the government shutdown. "Prior: 3 billion" is the percent chance (from 0 to 1) that the *New York Times* wrote an article about the CBO's evaluation of the government shutdown with a headline mentioning the long-run cost of $3 billion rather than the short-run costs of $11 billion. Column 1 and 2 use respondents who identify with the Democratic Party and respondents who identify as Independents but lean toward the Democratic Party. Column 3 and 4 use respondents who identify with the Republican Party and respondents who identify as Independents but lean toward the Republican Party. Regressions include the following controls: gender, age, income, region, race, education, employment status, frequency of reading the *New York Times*, pre-treatment beliefs about the probability that the *New York Times* would use a headline featuring the $3 billion number, and pre-treatment beliefs about the consequences of the shutdown.

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

Figure D.1: Treatment effects by pre-treatment beliefs: Spin

Panel A: Democrats   **Panel B**: Republicans

Notes: This figure displays heterogeneous treatment effects (Hainmueller et al., 2019) for the deceptive spin experiment (Experiment 3; see Table 1) by respondents' prior beliefs about how the *New York Times* covered a CBO report about the consequences of the government shutdown in early January, 2019. "Prior of right-wing bias in the NYT" is the percent chance (from 0 to 1) that the *New York Times* wrote an article about the CBO's evaluation of the government shutdown with a headline mentioning the long-run cost of $3 billion rather than the short-run costs of $11 billion. "Democrats/Lean Dem" are respondents who identify with the Democratic Party and respondents who identify as Independents but lean toward the Democratic Party. "Republican/Lean Rep" are respondents who identify with the Republican Party and respondents who identify as Independents but lean toward the Republican Party. The dashed red curve shows smoothed kernel density estimates for the moderating prior variable.

E Instructions

This section contains screenshots from all experiments. Table E.1 provides an overview of measures collected for each experiment. We provide screenshots of the full experimental instructions for the main filtering design (Experiment 1.1; see Table 1). For all other experiments, we always provide screenshots of the instructions used to measure pre-treatment beliefs about reporting, the information treatment, and the measure of article demand used in the experiment. We also provide screenshots of elements that differ from the main design. For example, the robustness curiosity experiment (Experiment 1.2; see Table 1) informs all respondents pre-treatment that they will learn about how the *New York Times* reported about a CBO report at the end of the survey. To avoid repetition, we do not include screenshots of demographic variables and post-treatment measures for other experiments if they are measured similarly to the main filtering design.

Table E.1: Overview of measures collected by experiment

Experiment:	1.1	1.2	1.3	1.4	1.5	1.6	1.7	2	3
Attention check	Yes	Yes	Yes	Yes	Yes			Yes	
Gender	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes	Yes			Yes	Yes
State						Yes	Yes		No
Household size						Yes	Yes		
Household income	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ethnicity	Yes			Yes		Yes	Yes	Yes	Yes
Race	Yes	Yes	Yes	Yes		Yes	Yes	Yes	Yes
Employment status	Yes	Yes	Yes	Yes		Yes	Yes	Yes	Yes
Education	Yes	Yes	Yes	Yes		Yes	Yes	Yes	Yes
Subscription to the NYT						Yes	Yes		
Frequency of reading the NYT	Yes	Yes		Yes	Yes		Yes		Yes
Frequency of reading CNN								Yes	
3 newspapers most likely to read	Yes					Yes	Yes		Yes
Acceptable to use unverified sources								Yes	
Political affiliation (3-point scale)	Yes		Yes	Yes	Yes	Yes	Yes		Yes
Political affiliation (5-point scale)		Yes						Yes	
Political leaning (for Independents)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Voting: 2016		Yes	Yes					Yes	
Voting: 2012		Yes	Yes						
Approval: Trump's policy agenda	Yes	Yes					Yes	Yes	
Impeach Trump								Yes	
Pre-treatment beliefs: Tax	Yes		Yes		Yes	Yes	Yes		
Pre-treatment beliefs: Health	Yes		Yes		Yes	Yes	Yes		
Pre-treatment beliefs: Dreamers							Yes		
Pre-treatment beliefs: Shutdown									Yes
Pre-treatment beliefs: False claims								Yes	
Beliefs about reporting	H	MW	H		H	T	H	CIA	GS
Article demand	T	SP	T		T	H	T	I	T
Posterior: Filtering (cont.)	Yes						Yes		
Posterior: Omission (cont.)	Yes						Yes		
Posterior: Quality (5-point)	Yes	Yes	Yes		Yes		Yes	Yes	Yes
Posterior: Dry and technical (5-point)	Yes	Yes					Yes		Yes
Posterior: Complex (5-point)	Yes						Yes		Yes
Political bias (5-point)	Yes	Yes	Yes		Yes		Yes	Yes	Yes
Trust (5-point)	Yes				Yes			Yes	Yes
Curiosity about NYT's bias (4-point)	Yes								Yes
CBO: Trust (5-point)	Yes						Yes		
CBO: Accuracy (5-point)	Yes						Yes		
CBO: Political bias (5-point)	Yes	Yes					Yes		
Top three reasons for reading news	Yes								
Section: Most interesting	Yes								
Platform to read news	Yes								
Willingness to pay					Yes				
Info demand: Debt (Tax)				Yes					
Info demand: Jobs (Tax)				Yes					
Info demand: Deficit (Health)				Yes					
Info demand: Insured (Health)				Yes					
Posterior: False claims (5-point)								Yes	
Posterior: False claims (cont.)								Yes	
Intention: Hurt Trump (5-point)								Yes	
Intention: Hurt Biden (5-point)								Yes	

Notes: This table provides an overview of collected variables by experiment (see Table 1). Rows above "Beliefs about reporting" contain measures collected pre-treatment. "Beliefs about reporting" refers to the pre-treatment belief elicitation about how a news outlet reported underlying facts. "Article demand" is the main outcome of interest. Rows below "Article demand" contain measures collected post-treatment. "H" refers to the *New York Times* article about the CBO evaluation of the Trump Healthcare Plan. "T" refers to the *New York Times* article about the CBO evaluation of the Trump Tax Plan. "GS" refers to the *New York Times* article about the CBO evaluation of the government shutdown in 2019. "MW" refers to the *New York Times* article about the CBO evaluation of raising the minimum wage to $15. "SP" refers to the *New York Times* article about the CBO evaluation a single-payer health care system. "I" refers to the *CNN* article about the impeachment inquiry. "CIA" refers to the *CNN* article about the extraction of a CIA spy from inside Russia.

E.1 Experiment 1.1 and 1.7

E.1.1 Attention Check (Experiment 1.1)

NHH

The next question is about the following problem. In questionnaires like ours, sometimes there are participants who do not carefully read the questions and just quickly click through the survey. This means that there are a lot of random answers which compromise the results of research studies. To show that you read our questions carefully, please choose **both** "Extremely interested" and "Not at all interested" on the question below.

Given the text above, how interested are you in sports?

Extremely interested

Very interested

A little bit interested

Very little interested

Not at all interested

>>

E.1.2 Pre-treatment beliefs and characteristics (Experiment 1.1 and 1.7)

NHH

Please indicate your gender.

Male

Female

What is your age?

18–24

25–34

35–44

45–54

55–64

65 or older

3

What is your region of residence?

Northeast   (CT, ME, MA, NH, RI, VT, NJ, NY,PA),

Midwest   (IL, IN, MI, OH, WI, IA, KS, MN, MO, NE, ND, SD)

South   (DE, DC, FL, GA,MD, NC, SC, VA, WV, AL, KY, MS, TN, AR, LA, OK, TX)

West   (AZ, CO, ID, NM, MT, UT,NV, WY, AK, CA, HI, OR, WA)

What was your family's gross household income in 2017 in US dollars?

Less than $15,000

$15,000 to $24,999

$25,000 to $49,999

$50,000 to $74,999

$75,000 to $99,999

$100,000 to $149,999

$150,000 to $200,000

More than $200,000

In politics, as of today, do you consider yourself a Republican, a Democrat, or an Independent?

Republican

Democrat

Independent

>>

NHH

In politics, as of today, do you lean towards the Republican Party or lean towards the Democratic Party?

The Republican Party

The Democratic Party

>>

4

How often do you read The New York Times (in print or digital)?

Daily

4-6 times a week

2-3 times a week

Once a week

Monthly

Once a year

Never

>>

Please rank the three newspapers that you are most likely to read from the list below (where 1 is the one you are most likely to read), and drag them to the appropriate position.

Items

Breitbart

BuzzFeed News

Chicago Sun-Times

Daily Mail

Drudge Report

InfoWars

Los Angeles Times

New Republic

New York Daily News

New York Post

Palmer Report

The Denver Post

The Huffington Post

The Mercury News

The New York Times

The Wall Street Journal

The Washington Post

The Washington Times

USA Today

3 newspapers most likely to read

>>

5

NHH

Do you approve or disapprove of Donald Trump's policy agenda?

Strongly approve

Approve

Disapprove

Strongly disapprove

>>

NHH

President Trump and Republicans in Congress have suggested two major legislative reforms:

- The **Trump Tax Plan** (to cut corporate taxes by $1.5 trillion)
- The **Trump Healthcare Plan** (to repeal and replace Obamacare)

On balance, do you think that the **Trump Tax Plan** will have positive or negative consequences?

Very positive consequences

Somewhat positive consequences

Neither positive nor negative consequences

Somewhat negative consequences

Very negative consequences

On balance, do you think that the **Trump Healthcare Plan** would have positive or negative consequences?

Very positive consequences

Somewhat positive consequences

Neither positive nor negative consequences

Somewhat negative consequences

Very negative consequences

>>

E.1.3 Prior: Filtering (Experiment 1.1 and 1.7)

NHH

The Congressional Budget Office (CBO) is Congress's nonpartisan provider of cost and benefit estimates for legislation. In 2017, the CBO analyzed the consequences of the Trump Healthcare Plan.

When debating the Trump Healthcare Plan, Republicans claimed that the plan would decrease the federal deficit, but would not increase the number of people without health coverage. The Democrats, by contrast, claimed that the plan would fail to decrease the deficit and massively increase the number of people without health coverage.

In its published report, the CBO estimated that the Trump Healthcare Plan would **decrease the deficit by $119 billion** and leave **23 million more people uninsured**.

What do you think?

After the CBO published its report, **The New York Times** wrote an article about its findings.

What would you say is the percent chance that **The New York Times** reported...
(Please note: The numbers need to add up to 100%)

that the deficit would decrease by $119 billion **but not** that the number of uninsured people would increase by 23 million.	0	%.
that the number of uninsured people would increase by 23 million **but not** that the deficit would decrease by $119 billion.	0	%.
that the deficit would decrease by $119 billion **and** that the number of uninsured people would increase by 23 million.	0	%.
Total	0	%.

>>

E.1.4 Treatment (Experiment 1.1 and 1.7)

NHH

In its article about the CBO estimates, The New York Times reported **both** that the federal budget deficit would decrease by $119 billion **and** that the number of people without health insurance would increase by 23 million.

>>

7

E.1.5 Main outcome (Experiment 1.1 and 1.7)

NHH

Last year, the Congressional Budget Office analyzed the consequences of the **Trump Tax Plan** over the next decade.

Do you want to read an article about its findings in **The New York Times**?

Yes
No

If you click "Yes" we will provide you with free access to the article (1100 words) at the end of the survey. If you click "No" you will proceed with the survey without receiving access to the article.

>>

E.1.6 Posterior beliefs (Wave 2 of Experiment 1.1 only)

NHH

When debating the Trump Tax Plan, Republicans claimed that the plan would create new jobs without increasing the federal debt. By contrast, Democrats claimed that the plan would fail to create more jobs and massively increase the federal debt.

In its report, the Congressional Budget Office estimated that the Trump Tax Plan would add **$1.6 trillion to the federal debt** and **create 1.1 million jobs**.

What do you think?

After the CBO published its report, **The New York Times** wrote an article about its findings.

In its article, The New York Times reported that the Trump Tax Plan would add $1.6 trillion to the federal debt.

What would you say is the percent chance that The New York Times **also reported** that the Trump Tax Plan would create 1.1 million jobs?

|	Extremely unlikely	Somewhat unlikely	Neither likely nor unlikely	Somewhat likely	Extremely likely	
	0	10	20	30	40	50	60	70	80	90	100
Percent						

>>

Democrats have urged President Trump to grant citizenship status for up to 1.8 million young undocumented immigrants, known as the Dreamers. Democrats have claimed that this would decrease the federal debt, whereas Republicans have claimed that it would increase the federal debt.

The CBO recently also analyzed the impact of granting citizenship status for the Dreamers. In its report, the CBO estimated that this would **add $27 billion to the federal debt**.

What do you think?

What would you say is the percent chance that **The New York Times** did **not** write an article about the findings from this CBO report?

	Extremely unlikely	Somewhat unlikely	Neither likely nor unlikely	Somewhat likely	Extremely likely
	0 10 20	30	40 50 60	70 80	90 100

Percent ●

`>>`

In general, how do you rate the **quality of news articles** in The New York Times?

Very low
Low
Medium
High
Very high

When thinking about its coverage of the CBO report about the Trump Tax Plan, how **dry and technical** do you expect the article to be?

Not at all dry and technical
Not dry and technical
Somewhat dry and technical
Very dry and technical
Extremely dry and technical

9

When thinking about The New York Times' coverage of the CBO report about the Trump Tax Plan, do you expect a **very simple message** or a **very complex message**?

Very simple

Simple

Neither simple nor complex

Complex

Very complex

>>

E.1.7 Perceptions: NYT and CBO (Experiment 1.1)

NHH

In general, do you think **The New York Times** is politically biased?

Very right-wing biased

Somewhat right-wing biased

Not biased

Somewhat left-wing biased

Very left-wing biased

>>

How much do you trust The New York Times?

| Strongly trust |
| Trust |
| Somewhat trust |
| Do not trust |
| Do not trust at all |

How much do you trust the media in general?

| Strongly trust |
| Trust |
| Somewhat trust |
| Do not trust |
| Do not trust at all |

How interested would you be in learning about statistics on whether The New York Times reports unbiasedly about political issues?

| Very interested |
| Interested |
| Not interested |
| Not interested at all |

>>

How much do you trust the forecasts of the Congressional Budget Office?

Strongly trust

Trust

Somewhat trust

Do not trust

Do not trust at all

In your opinion, how accurate are the forecasts of the Congressional Budget Office?

Very accurate

Accurate

Somewhat accurate

Inaccurate

Very inaccurate

Do you think the Congressional Budget Office is politically biased?

Very right-wing biased

Somewhat right-wing biased

Not biased

Somewhat left-wing biased

Very left-wing biased

>>

E.1.8   Post-treatment beliefs and characteristics (Experiment 1.1)

NHH

Why do you usually read political news? Please rank the three most important reasons (where 1 is the the most important one for you)

Items	Main 3 reasons
To improve my knowledge about political issues	
To be able to follow the national conversation	
To expose myself to different points of view	
For the entertainment value	
To make more informed voting choices	
Because it is important for my job	

>>

NHH

In newspapers, which section are you most interested in?

Entertainment

Advice columns

Editioral & opinion pages

Lifestyle

Political news

Which of these platforms are you most likely to use as news sources?

Radio

Social media

Print newspapers

News websites

Television

>>

Which of the following best describes your race or ethnicity?

African American/Black

Asian/Asian American

Caucasian/White

Native American, Inuit or Aleut

Native Hawaiian/Pacific Islander

Other

Are you of Hispanic, Latino, or Spanish origin?

Yes

No

Which category best describes the highest level of education you have completed?

Eighth grade or less

Some high school

High school degree/GED

Some college

2-year college degree

4-year college degree

Master's degree

Doctoral degree

Professional degree (JD, MD, MBA)

What is your current employment status?

Full-time employee

Part-time employee

Self-employed or small business owner

Unemployed and looking for work

Student

Not in labor force (for example: retired or full-time parent)

>>

E.1.9 End of Survey

NHH

Please see below to find your free access to The New York Times article covering the CBO estimates of the Trump Tax Plan.

The New York Times

Federal Budget Deficit Projected to Soar to Over $1 Trillion in 2020

By Thomas Kaplan

April 9, 2018

WASHINGTON — The federal government's annual budget deficit is set to widen significantly in the next few years, and is expected to top $1 trillion in 2020 despite healthy economic growth, according to new projections from the nonpartisan Congressional Budget Office released Monday.

The national debt, which has exceeded $21 trillion, will soar to more than $33 trillion in 2028, according to the budget office. By then, debt held by the public will almost match the size of the nation's economy, reaching 96 percent of gross domestic product, a higher level than any point since just after World War II and well past the level that economists say could court a crisis.

The fear among some economists is that rising deficits will drive up interest rates, raise borrowing costs for the private sector, tank stock prices and slow the economy, which would only drive the deficit higher.

"Such high and rising debt would have serious negative consequences for the budget and the nation," said Keith Hall, the director of the budget office. "In particular, the likelihood of a fiscal crisis in the United States would increase."

The budget office forecast is the first since President Trump signed a sweeping tax overhaul, then signed legislation to significantly increase military and domestic spending over the next two years. The figures are sobering, even in a political climate where deficit concerns appear to be receding.

The tax overhaul, which includes permanent tax cuts for corporations and temporary ones for individuals, will increase the size of the economy by an average of 0.7 percent from 2018 to 2028, according to the budget office.

But that added economic growth does not come close to paying for the tax overhaul, which the budget office said would add more than $1.8 trillion to deficits over that period, from lost tax revenue and higher interest payments.

E.2 Experiment 1.2 – Robustness curiosity

E.2.1 Pre-treatment beliefs and characteristics

Which of these describes you more accurately?

Male

Female

What is your age?

18–24

25–34

35–44

45–54

55–64

65 or older

What is your region of residence?

Northeast (CT, ME, MA, NH, RI, VT, NJ, NY,PA),

Midwest (IL, IN, MI, OH, WI, IA, KS, MN, MO, NE, ND, SD)

South (DE, DC, FL, GA,MD, NC, SC, VA, WV, AL, KY, MS, TN, AR, LA, OK, TX)

West (AZ, CO, ID, NM, MT, UT,NV, WY, AK, CA, HI, OR, WA)

What was your family's gross household income in 2018 in US dollars?

Less than $15,000

$15,000 to $24,999

$25,000 to $49,999

$50,000 to $74,999

$75,000 to $99,999

$100,000 to $149,999

$150,000 to $200,000

More than $200,000

In politics, as of today, do you consider yourself a Republican, a Democrat, or an Independent?

Republican

Democrat

Independent

>>

Do you approve or disapprove of Donald Trump's policy agenda?

Strongly approve

Approve

Disapprove

Strongly disapprove

Who did you vote for in the 2016 Presidential election?

Donald Trump

Hillary Clinton

Other

I did not vote

Who did you vote for in the 2012 Presidential election?

Barack Obama

Mitt Romney

Other

I did not vote

>>

Which of the following best describes your race or ethnicity?

White	Asian
Black or African American	Native Hawaiian or Pacific Islander
American Indian or Alaska Native	Other

Which category best describes the highest level of education you have completed?

12th grade or less

Graduated high school or equivalent

Some college, no degree

Associate degree

Bachelor's degree

Post-graduate degree

Which of these describes your current situation most accurately?

Employed full-time

Employed part-time

Self-employed

Unemployed and looking for a job

Unemployed but not looking for a job

Retired

Student

Other

>>

How often do you read **The New York Times** (in print or digital)?

Daily

4-6 times a week

2-3 times a week

Once a week

Monthly

Once a year

Never

>>

E.2.2 Belief elicitation

Information

The Congressional Budget Office (CBO) is Congress's nonpartisan provider of cost and benefit estimates for legislation.

We will now ask you a question about how The New York Times covered the findings from a recent major policy report from the CBO.

We will tell you how The New York Times covered these findings at some later point in the survey.

`>>`

Democrats' $15 Minimum Wage Bill

In July, the CBO analyzed the consequences of a bill to increase the federal minimum wage to $15 an hour.

When debating the bill, **Democrats** claimed that the bill would lift more people out of poverty without reducing the number of jobs.
Republicans, by contrast, claimed that the bill would fail to lift people out of poverty and massively reduce the number of jobs.

In its report, the CBO estimated that the bill would **lift 1.3 million people out of poverty** and that the bill would **decrease the number of jobs by 1.3 million**.

What do you think?

After the CBO published its report, The New York Times wrote an article about its findings.

What would you say is the percent chance that **The New York Times** reported that...
(Please note: The numbers need to add up to 100%)

1.3 million people would be lifted out of poverty **but not** that the number of jobs would decrease by 1.3 million.	0	%.
the number of jobs would decrease by 1.3 million **but not** that 1.3 million people would be lifted out of poverty.	0	%.
1.3 million people would be lifted out of poverty **and** that the number of jobs would decrease by 1.3 million.	0	%.
Total	0	%.

22

E.2.3 Treatment

Information

In its article about the CBO estimates, The New York Times reported **both** that 1.3 million people would be lifted out of poverty **and** that the number of jobs would decrease by 1.3 million.

E.2.4 Main outcome

The Congressional Budget Office recently also analyzed the consequences of establishing a **single-payer health care system** (to achieve universal health insurance coverage).

Do you want to read an article about its findings in **The New York Times**?

Yes

No

If you click "Yes" we will provide you with free access to the article (1100 words) at the end of the survey. If you click "No" you will proceed with the survey without receiving access to the article.

E.3 Experiment 1.3 – Cognitive constraints placebo

Figure E.1: Belief elicitation

The Congressional Budget Office (C.B.O.) is Congress's nonpartisan provider of cost and benefit estimates for legislation.

In 2017, the C.B.O. analyzed the consequences of the G.O.P. Health Bill.

When the C.B.O. published its report about the G.O.P. Health Bill, the **C.B.O. highlighted two key statistics** from the report.

After the C.B.O. published its report, **The New York Times** wrote an article about the report.

What do you think?

What would you say is the percent chance that the **The New York Times** cited zero, one, or two of the two key statistics from the C.B.O. report?

(Please note: The numbers must total 100%)

The New York Times cited 0 of the 2 key statistics.	0	%
The New York Times cited 1 of the 2 key statistics.	0	%
The New York Times cited 2 of the 2 key statistics.	0	%
Total	0	%

Next >>

Figure E.2: Treatment screen

In its article, **The New York Times** cited **2 of the 2 key statistics** from the C.B.O. report.

Next >>

25

Figure E.3: Willingsness to read about GOP Tax Bill in the NYT

Last year, the C.B.O. analyzed the consequences of the **G.O.P. Tax Bill** over the next decade.

Do you want to read an article about its findings in **The New York Times**?

○ Yes

○ No

If you click "Yes" we will provide you with free access to the article (1100 words) at the end of the survey. If you click "No" you will proceed with the survey without receiving access to the article.

Next >>

E.4 Experiment 1.4 – Information demand

For this experiment, we provide the instructions for the measurement of people's demand for information about the CBO estimates about the consequences of the (i) Trump Tax Plan and (ii) the Trump Healthcare Plan. The order of these two blocks was randomized.

Figure E.4: Demand for information: Tax Bill

The Congressional Budget Office (CBO) is Congress's nonpartisan provider of cost and benefit estimates for legislation. In 2018, the CBO analyzed the economic consequences of the Trump **Tax** Plan (to cut corporate taxes by $1.5 trillion).

Would you like to receive the estimate from the CBO about how the Trump Tax Plan would affect the **federal debt** over the next decade?

If you click "Yes" you will receive the estimate at the end of the survey.

Yes
No

Would you like to receive the estimate from the CBO about how the Trump Tax Plan would affect the **number of jobs** over the next decade?

If you click "Yes" you will receive the estimate at the end of the survey.

Yes
No

>>

Figure E.5: Demand for information: Health Bill

The Congressional Budget Office (CBO) is Congress's nonpartisan provider of cost and benefit estimates for legislation. In 2018, the CBO analyzed the economic consequences of the Trump **Healthcare** Plan (to repeal and replace Obamacare).

Would you like to receive the estimate from the CBO about how the Trump Healthcare Plan would affect the **federal budget deficit** over the next decade?

If you click "Yes" you will receive the estimate at the end of the survey.

Yes
No

Would you like to receive the estimate from the CBO about how the Trump Tax Plan would affect the **number of people without health insurance** over the next decade?

If you click "Yes" you will receive the estimate at the end of the survey.

Yes
No

>>

E.5 Experiment 1.5 – External validity

We provide the instructions for the measurement of people's willingness to pay for a digital subscription to the *New York Times*. The demand for news was measured as in the main filtering design (Experiment 1.1; see Table 1).

NHH

You will now make multiple decisions that can have **real financial consequences for you**. Please consider each decision carefully.

In each decision, we will ask you to choose one of two options:

- Option A: 3-month digital subscription to **The New York Times**.
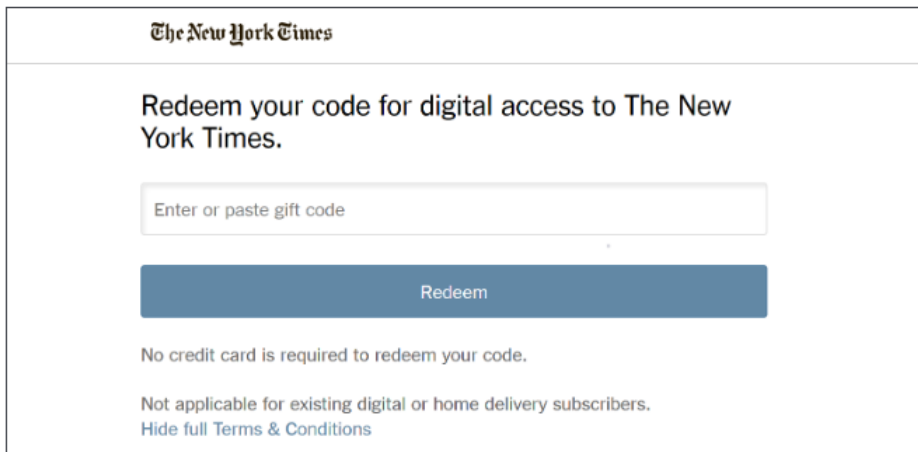- Option B: Varying amounts of money.

We will randomly select **1 out of 10** participants of this study. If we select you, we will randomly choose one of your decisions and implement the option you chose. Each decision has the same chance of being implemented.

If we implement Option A, you will receive a unique gift code for a 3-month subscription. If we implement Option B, you will receive an amount of money (paid out as a bonus to your MTurk account).

>>

How would you receive your digital subscription to The New York Times?

1. We will send you a **unique gift code**. It looks like this: *4432a7af8c83b5da72gb*

2. We provide you with a **link to a website** where you can redeem it. It looks like this:

The New York Times

Redeem your code for digital access to The New York Times.

Enter or paste gift code

Redeem

No credit card is required to redeem your code.

Not applicable for existing digital or home delivery subscribers.
Hide full Terms & Conditions

---

3. Enter the code and **create an account**. You only need an email address for this.

4. **That's all!**

No credit card information is required to create an account. The subscription will automatically be canceled after the 3 month period. You can also cancel the subscription at any time if you want.

The code is completely anonymous and cannot be used to identify your email or any other of your personal characteristics.

We will now give you the opportunity to decide between two options:

- **Option A:** 3-month digital subscription to **The New York Times (NYT)**.
- **Option B:** Varying amounts of money.

Which option do you prefer?

Option A	Option B
NYT subscription	50 cents
NYT subscription	$1
NYT subscription	$2
NYT subscription	$3
NYT subscription	$4
NYT subscription	$5
NYT subscription	$10

E.6 Experiment 1.6

We provide the instructions for the (i) financially incentivized pre-treatment belief elicitation, (ii) the information treatment, and (iii) the demand for news.

Figure E.6: Explanations of probability

In some of the following questions, we will ask you to think about the percent chance of something happening in the future. Your answers can range from 0 to 100, where 0 means there is absolutely no chance, and 100 means that it is absolutely certain.

For example, numbers like:

- 2 or 5 percent may indicate "almost no chance."
- 18 percent or so may mean "not much chance."
- 47 or 52 percent chance may be a "pretty even chance."
- 83 percent or so may mean a "very good chance."
- 95 or 98 percent chance may be "almost certain."

>>

Figure E.7: Explanations of incentive payment

In what follows, we will ask you to make some estimates on factual statements which are either true or false. One out of ten participants can earn additional money based on their estimates. For those participants, we will randomly pick one of the questions in which they can earn money, and pay them according to their estimate. They can earn up to an additional $1.

We will ask you to think about the percent chance that of different statements being true. The below formula explains in detail how the payout is determined. While this formula may appear to be complicated, the important take-away message from the formula is that participants will earn more money the closer they are to the truth. If the statement is true then participants will receive a higher payoff the higher their estimate. If the statement is false then they will receive a higher payoff the lower their estimate. Moreover, they can never make a loss by giving an estimate.

Your payment depends on your estimate in the following way:

Payment (in US dollars) $= 1 - 1 \times (\text{estimate}/100 - \text{truth})^2$

where **truth** takes the value 1 if the statement is true, and zero otherwise.

>>

Figure E.8: Beliefs about reporting I

The Congressional Budget Office (CBO), Congress's nonpartisan provider of cost and benefit estimates for legislation, recently analyzed the impact of the GOP Tax Bill on the economy.

When debating the GOP Tax Bill, Republicans claimed that the Tax Bill would create new jobs without increasing the federal debt. By contrast, Democrats claimed that the GOP Tax Bill would fail to create more jobs and massively increase the federal debt.

In April 2018, CBO published its report about the impact of the GOP Tax Bill on the economy. The CBO estimated that the GOP Tax Bill would add **$1.6 trillion to federal debt** and **create 1.1 million jobs** over the next decade.

What do you think?

After the Congressional Budget Office published its report in April 2018, **The New York Times** wrote an article about its findings.

What would you say is the percent chance that **The New York Times** reported...
(Please note: The numbers need to add up to 100 percent)

that jobs would increase by 1.1 millions **but not** that the federal debt would increase by $1.6 trillion	0	%.
that the federal debt would increase by $1.6 trillion **but not** that jobs would increase by 1.1 millions	0	%.
that the federal debt would increase by $1.6 trillion **and** that jobs would increase by 1.1 millions	0	%.
Total	0	%.

Figure E.9: Treatment screen

In its coverage of the CBO estimates, The New York Times reported both that the federal debt would increase by $1.6 trillion **and** that jobs would increase by 1.1 millions.

34

Figure E.10: Article demand

The Congressional Budget Office also analyzed the economic impact of the **GOP Health Care Bill.** Do you want to read a story about its findings in **The New York Times**?

If you click "Yes" we will provide you with free access to the article. If you click "No" you will proceed with the survey without receiving access to the article.

Yes

No

>>

E.7 Experiment 2 – false claims

Figure E.11: Policy views

Do you approve or disapprove of Donald Trump's policy agenda?

Strongly approve

Approve

Disapprove

Strongly disapprove

Who did you vote for in the 2016 Presidential election?

Donald Trump

Hillary Clinton

Other

I did not vote

Democrats in Congress just started an impeachment inquiry after accusing Donald Trump of betraying his oath of office.

Do you personally think Donald Trump has betrayed his oath of office and should be impeached?

Yes

No

>>

Figure E.12: Acceptability of using unverified sources

To what extent do you agree or disagree with the following statement:

"It is acceptable for news organizations to publish stories based on unverified sources if the stories are potentially important but could lead to fake news."

Strongly agree

Somewhat agree

Neither agree nor disagree

Somewhat disagree

Strongly disagree

>>

Figure E.13: Prior CNN consumption

How often do you watch or read **CNN** (on TV or online)?

Daily

4-6 times a week

2-3 times a week

Once a week

Monthly

Once a year

Never

>>

Figure E.14: Prior beliefs about false claims made by CNN

How often do you think CNN makes false claims in its political reporting?

Very often

Often

Sometimes

Rarely

Never

>>

Figure E.15: Addressing curiosity

Information

We will now ask you a question about how CNN covered the extraction of a CIA spy from Russia.

We will tell you how CNN covered this story at some later point in the survey.

>>

Figure E.16: Beliefs about CNN's reporting

CIA Spy Extracted From Russia

Two years ago, the US extracted a high-level CIA spy from inside Russia. The spy had provided the US with valuable intelligence on the inner workings of Putin's regime and how Moscow could threaten America. The news media shared details about the story in early September.

Some media outlets **correctly reported** that the CIA made the decision to extract the spy in late 2016 following widespread media speculation about CIA's sources.

Other media outlets **falsely reported** that the CIA made the decision to extract the spy because it feared that President Trump would mishandle classified information and potentially reveal the identity of the spy.

What do you think?

In its news article about the CIA spy extracted from Russia, what would you say is the percent that CNN reported that the spy was extracted ...

(Please note: The numbers need to add up to 100%)

because of widespread media speculation about the CIA's sources.	0	%
because the CIA feared that President Trump would mishandle classified information.	0	%
Total	0	%

Notes: The order of the two options was randomized.

Figure E.17: Information treatment

Information

In its article, **CNN** reported that the spy was extracted because the CIA feared that President Trump would mishandle classified information.

Figure E.18: Article demand

A recent whistleblower complaint following a phone conversation between President Trump and a foreign leader led Democrats in Congress to start an impeachment inquiry against the president.

After the impeachment inquiry was made public, CNN wrote a news article about the impeachment process with the full background story.

Do you want to receive free access to the CNN article?

Yes

No

If you click "Yes" we will provide you with free access to the article (1,600 words) at the end of the survey. If you click "No" you will proceed with the survey without receiving access to the article.

Figure E.19: Perceptions of CNN I

In general, how do you rate the quality of news articles by CNN?

Very low

Low

Medium

High

Very high

How much do you trust CNN?

Strongly trust

Trust

Somewhat trust

Do not trust

Do not trust at all

How much do you trust the media in general?

Strongly trust

Trust

Somewhat trust

Do not trust

Do not trust at all

>>

42

Figure E.20: Perceptions of CNN II

Do you think CNN is politically biased?

Very right-wing biased

Somewhat right-wing biased

Not biased

Somewhat left-wing biased

Very left-wing biased

How often do you think CNN has made false claims in its political reporting?

Never

Rarely

Sometimes

Often

Very often

To what extent do you agree or disagree with the following statement:

"CNN has intentionally made false claims to hurt President Trump."

Strongly agree

Somewhat agree

Neither agree nor disagree

Somewhat disagree

Strongly disagree

Figure E.21: Perceptions of CNN III

To what extent do you agree or disagree with the following statement:

"CNN has intentionally made false claims to hurt President Trump."

Strongly agree

Somewhat agree

Neither agree nor disagree

Somewhat disagree

Strongly disagree

Imagine that you read an article about Donald Trump from CNN. What would you say is the percent chance that this article would contain any false claims?

Extremely unlikely	Somewhat unlikely	Neither likely nor unlikely	Somewhat likely	Extremely likely
0 10 20	30 40	50 60	70 80	90 100

Percent ⬤

How likely do you think CNN is to publish stories about President Trump based on unverified and potentially misleading sources?

Very likely

Somewhat likely

Neither likely nor unlikely

Somewhat unlikely

Very unlikely

How likely do you think CNN is to publish stories about Joe Biden based on unverified and potentially misleading sources?

Very likely

Somewhat likely

Neither likely nor unlikely

Somewhat unlikely

Very unlikely

E.8 Experiment 3 – Deceptive spin

We provide the instructions for (i) beliefs about the shutdown, (ii) beliefs about reporting, (iii) the information treatment, and (iv) the measure of demand for news.

Figure E.22: Beliefs about shutdown

President Donald Trump's standoff with Democrats over funding for his proposed U.S.-Mexico border wall resulted in a five-week partial government shutdown that that ended on January 25, 2019.

On balance, do you think that the government shutdown will have positive or negative consequences for the U.S. economy in 2019?

Very positive consequences

Somewhat positive consequences

Neither positive nor negative consequences

Somewhat negative consequences

Very negative consequences

>>

Figure E.23: Belief elicitation of reporting

The Congressional Budget Office (C.B.O.) is Congress's nonpartisan provider of cost and benefit estimates for legislation.

After the federal government shutdown ended on January 25, 2019, the C.B.O. analyzed the consequences of the shutdown on the U.S. economy.

In its published report, the C.B.O. estimated that the shutdown had cost the U.S. economy **$11 billion**, of which **$3 billion** would be permanently lost.

After the C.B.O. published its report, **The New York Times** wrote an article about its findings.

In its article, **The New York Times** used one of the two headlines below. What would you say is the percent chance that it used each of the following headlines...

(Please note: The numbers need to add up to 100%)

Government Shutdown Cost U.S. Economy $3 Billion, C.B.O. Says	0	%
Government Shutdown Cost U.S. Economy $11 Billion, C.B.O. Says	0	%
Total	0	%

Figure E.24: Treatment screen

The New York Times used the following headline in its article:

Government Shutdown Cost U.S. Economy $11 Billion, C.B.O. Says

46

Figure E.25: Willingness to read about GOP Tax Bill in the NYT

Last year, the C.B.O. analyzed the consequences of the **G.O.P. Tax Bill** over the next decade.

Do you want to read an article about its findings in **The New York Times**?

Yes

No

If you click "Yes" we will provide you with free access to the article (1100 words) at the end of the survey. If you click "No" you will proceed with the survey without receiving access to the article.