Consuming cross-cutting media causes learning and moderates attitudes: A field experiment with Fox News viewers*

David E. Broockman† Joshua L. Kalla‡

December 15, 2023

Short Title: A field experiment on cross-cutting media exposure

Abstract

Many Americans consume aligned partisan media, which scholars worry contributes to polarization. Many propose encouraging these Americans to consume cross-cutting media to moderate their attitudes. However, motivated reasoning theory posits that exposure to cross-cutting media could backfire, exacerbating polarization. Building on theories that sustained exposure to novel information can overcome motivated reasoning and that partisan sources on opposite sides cover distinct information, we argue that sustained consumption of cross-cutting media leads voters to learn uncongenial information and moderate their attitudes in covered domains. To test this argument, we used data on actual TV viewership to recruit a sample of regular Fox News viewers and incentivized a randomized treatment group to watch CNN instead for a month. Contrary to predictions from motivated reasoning, watching CNN caused substantial learning and moderated participants' attitudes in covered domains. We close by discussing challenges partisan media may pose for democracy.

Keywords: media effects; motivated reasoning; Fox News; field experiment

*Replication files are available in the JOP Data Archive on Dataverse (https://dataverse.harvard.edu/dataverse/jop).The empirical analysis has been successfully replicated by the JOP replication analyst. Supplementary material for this article is available in the appendix in the online edition. This research was reviewed and approved by the Yale University Human Subjects Committee. This research is based upon work supported in part by the National Science Foundation under Grant #1917993 as well as a grant from the Field Experiments Initiative at the Yale Institution for Social and Policy Studies.

†Associate Professor, Travers Department of Political Science, University of California, Berkeley. dbroockman@berkeley.edu, https://polisci.berkeley.edu/people/person/david-broockman

‡Associate Professor, Department of Political Science and Department of Statistics and Data Science, Yale University. josh.kalla@yale.edu, https://joshuakalla.com.

Many Americans prefer consuming television media that shares their pre-existing political views instead of a more balanced diet containing some cross-cutting sources (Prior 2013; Stroud 2011). Many scholars express concern that congenial partisan media causes political beliefs and attitudes to be more extreme (e.g., Martin and Yurukoglu 2017).

In response, many propose encouraging Americans with one-sided media diets to consume cross-cutting content, hoping this would moderate their attitudes (e.g., Manjoo 2008; Sunstein 2007; Goldman and Mutz 2011). For instance, Barack Obama encouraged Americans to "seek out information that challenges our assumptions."[1] This sentiment dates back centuries: John Stuart Mill (1848) wrote that "It is hardly possible to overstate the value...of placing human beings in contact...with modes of thought...unlike those with which they are familiar."

However, theories of motivated reasoning argue that exposure to cross-cutting content can backfire and actually make beliefs and attitudes more extreme. Cross-cutting content is expected to backfire because individuals generate counterarguments against its uncongenial content which in turn make their beliefs and attitudes more extreme (for review, see Levendusky 2013, p. 614). Therefore, as Coppock (2022, p. 2) reviews, "A key prediction of motivated reasoning theory is backlash: exposure to counter-attitudinal evidence will cause people to hold more strongly to their preexisting positions." Motivated reasoning theory leads many scholars such as Arceneaux and Johnson (2013, p. 74) to warn against cross-cutting exposure, concluding that "exposure to counterattitudinal news can be just as polarizing as exposure to proattitudinal news." Levendusky (2013, p. 614) similarly notes that "cross-cutting media" is often theorized to "increase attitude extremity and polarization."

We argue that sustained consumption of cross-cutting media can lead voters to learn uncongenial information and moderate their attitudes. Motivated reasoning theories argue that motivated reasoning happens in two phases (e.g., Lodge and Taber 2013, p. 152): at the stage of *information search*, when people seek out confirmatory information (selective exposure, i.e., Stroud 2011);

[1]See `https://bit.ly/3ZOBSmP`.

and, second, when people *process information* to which they are exposed.[2] These are analogous to the Receive (selective exposure) and Accept (information processing) steps of Zaller's (1992) RAS model (Nyhan 2014). In this framework, advocates of cross-cutting exposure argue that overcoming selective exposure, the first step in motivated reasoning theories, would moderate beliefs and attitudes. Their critics argue that overcoming selective exposure would actually create backlash—or at best have no effects—due to biases in information processing (the 'Accept' step). We argue against this latter idea in the context of partisan media consumption.

We offer two key reasons why congenial partisan media consumers who engage in sustained exposure to cross-cutting media can overcome motivated reasoning in information processing.

First, partisan sources on opposite sides cover different topics and information, conveying more information favorable to their side and less information unfavorable to their side (e.g., Hayakawa 1940; Mullainathan and Shleifer 2005; Grossman, Margalit and Mitts 2022). For instance, during September 2020 when we conducted the study we present below, CNN provided extensive coverage of COVID-19, which included information about the severity of COVID-19 and poor aspects of then-President Trump's handling of it. Fox News covered COVID-19 much less, and the coverage it did offer provided little of the information CNN did, instead giving viewers information about why the virus was not a serious threat. Research shows that factual information is often able to overcome motivated reasoning because its valence is less immediately apparent (Guess and Coppock 2020; Wood and Porter 2019; Porter and Wood 2022). If individuals balance their media diets to include cross-cutting sources, we therefore expect they would learn some of the

[2]Formal literature offers a related definition of motivated reasoning which predicts that individuals update less in the direction of new information than they should under Bayes' rule (e.g., Little 2022). Lodge and Taber (2013) call this the prior attitude effect and Little (2022) calls this "once-motivated reasoning." However, this definition does not allow for backlash (Little 2022, footnote 22). We discuss this in greater detail in the discussion. For a concise overview of motivated reasoning theory and its relationship to Bayesian updating, see Druckman and McGrath (2019).

cross-cutting information present in these sources (i.e., update their beliefs) and incorporate this information into their attitudes.

Second, research suggests *sustained* exposure to cross-cutting content can lead motivated reasoners to eventually reach a "tipping point" that leads them to revisit their views (Gerber and Green 1998; Redlawsk, Civettini and Emmerson 2010).

In support of this argument, we present results from a field experiment which incentivized individuals who selectively consume congenial partisan television[3] media to instead consume cross-cutting media, leading them to consume a more balanced media diet. This experiment employed a unique design and was conducted among a unique sample: we incentivized a randomized treatment group of regular Fox News viewers to watch CNN instead for four weeks during September 2020, then measured the effects of this consumption on beliefs and attitudes.

Two differences between our research design and previous research on cross-cutting partisan television consumption are particularly theoretically significant. First, we incentivized participants to engage in sustained consumption of cross-cutting partisan media, allowing the topics and information their favored media source (Fox News) and a cross-cutting source (CNN) covered to vary as it does in the real world. By contrast, previous research on televised partisan media has generally exposed participants to brief clips of cross-cutting media while holding constant the topics and even information covered in congenial and cross-cutting sources (for review, see Appendix 1). Previous research is therefore unable to capture any potential effects of the differences in information present in congenial and cross-cutting sources, such as learning or its downstream consequences.

A second difference between our study and previous work is that we conducted it among the population of interest to advocates of balancing media diets and among whom motivated reasoning theories most strongly predict we should expect to find backfire effects: regular congenial partisan media viewers. Motivated reasoning theories argue that partisan media consumers are best

[3]We focus on televised partisan media given that selective exposure is more prevalent in television than online news consumption (Muise et al. 2022). Appendix 1 discusses studies on online media.

equipped and most motivated to counterargue against information from cross-cutting sources, leading backlash to be particularly likely among this population (e.g., Arceneaux and Johnson 2013). Indeed, we find this population harbors extreme attitudes and distrusts cross-cutting sources. However, previous research has only been conducted among members of the general public who state or reveal a preference for various media sources in surveys, not the individuals who actually choose to consume partisan media in the real world (see Appendix 1 for review). This makes our study a hard test of our argument because we conducted it among the population motivated reasoning theories predict should be most likely to exhibit backlash.

To recruit participants, we partnered with a media analytics company (Bully Pulpit Interactive) to identify individuals who regularly watch Fox News and do not watch CNN or MSNBC, using data on their households' actual television viewership. To induce cross-cutting exposure and balance participants' media diets, we offered a randomized treatment group \$15 per hour to watch up to 7 hours of CNN per week during September 2020 at the hours at which they typically watched Fox News. To increase the probability that individuals 'Received' this cross-cutting content and thereby to balance their media diets as much as possible, we enforced compliance with viewership quizzes about non-political features of the coverage (e.g., about which guest had just appeared). Although these quizzes may induce artificially close attention, we interpret any such effects as further increasing reception of cross-cutting content, such that potential biases at the information processing ('Accept') step are all that remain. Indeed, under motivated reasoning theories we would expect these circumstances to be particularly conducive to counterarguing and backfire, since exposure to uncongenial information—the key requisite condition for backfire—is essentially guaranteed. Although it is possible the incentives may have interfered with the information processing step, Khanna and Sood's (2018) results suggest that if anything financial incentives to accurately report information tend to *increase* bias by motivating subjects to reinforce their beliefs.

Consistent with our argument, we found that incentivizing partisan media viewers to balance their media diets towards cross-cutting content led to learning and moderated their attitudes.

First, consistent with our argument that partisan media viewers would learn from cross-cutting media if they 'Received' it despite the predictions of motivated reasoning theories, we found evidence of substantial learning. In particular, we found that watching CNN instead of Fox News affected participants' factual perceptions of current events (i.e., beliefs) and knowledge about the 2020 presidential candidates' positions. It also decreased their knowledge of information covered on Fox News.

Accompanying these shifts, we also found evidence of moderation (i.e., among these conservative participants, leftwards shifts) along a number of dimensions, including attitudes about current events, policy preferences, and evaluations of key political figures and parties. For example, we found leftward shifts in attitudes and preferences about COVID-19, and decreases in evaluations of Donald Trump and Republican candidates and elected officials.

An endline survey two months later found these impacts largely receded as treated participants primarily returned to their prior viewing habits, consistent with participants having a preference for like-minded media (the 'Receive,' information search step of motivated reasoning theories).

We close by elaborating three broader implications of our findings. First, consistent with the benefits of cross-cutting exposure, we find that selective exposure to congenial partisan sources—the information search step of motivated reasoning theories—exacerbates polarization: if individuals were more motivated to consume cross-cutting content, we argue voters would have more moderate, less polarized attitudes. Second, echoing findings on the limits of motivated reasoning in information processing (e.g., Guess and Coppock 2020; Wood and Porter 2019; Redlawsk, Civettini and Emmerson 2010), our results on both learning and attitudes contrast with expectations that Americans—and especially highly engaged partisans—simply reject messages contrary to their partisan loyalties. Third, our findings suggest that partisan media may affect voters' attitudes in part because it selectively reports information. As we elaborate in the discussion, this suggests that partisan media may present a challenge for democratic accountability.

How Sustained Cross-Cutting Exposure May Overcome Motivated Reasoning

As described above, many scholars oppose proposals to encourage cross-cutting media consumption, expecting that it would only make partisan media consumers' attitudes more extreme. For example, Arceneaux and Johnson's (2013) influential research argues that "exposure to counterattitudinal news can be just as polarizing as exposure to proattitudinal news" (p. 74) and that "In spite of the hopeful notion that exposure to alternative views will ameliorate political division," consuming cross-cutting media merely "reinforces preexisting attitudes" (p. 104). Their fears are rooted in theories of motivated reasoning, which argue that when processing uncongenial content, people generate counterarguments in support of their own views, which then lead both their beliefs to change in the *opposite* direction of the content's signal (Nyhan and Reifler 2010) and their attitudes to grow more extreme in the opposite direction, too (Lodge and Taber 2013). As we review in Table OA1, evidence from survey experiments incentivizing brief exposure to cross-cutting partisan media often supports these expectations, particularly among people who prefer consuming congenial partisan media or have strong attitudes.

However, there is an important difference between the media used as experimental stimuli in prior survey experiments and real-world partisan media. In an effort to hold other factors constant, prior research's experimental stimuli essentially always hold constant the issues covered in cross-cutting and congenial media (see Table OA1). However, other research shows that real-world partisan media channels on opposite sides cover dramatically different topics and information (e.g., Baum and Groeling 2008; Grossman, Margalit and Mitts 2022). Researchers have used different terms to refer to this phenomenon, and we use the term *partisan coverage filtering*.[4] We define partisan coverage filtering as when a media outlet conveys more information favorable to its partisan

[4]Hayakawa (1940) and Mullainathan and Shleifer (2005) use "slant"; Besley and Prat (2006) use "bias"; Baum and Groeling (2008) discuss "bias" in story "selection"; Gentzkow, Shapiro and Stone (2016) use "filtering bias"; and Grossman, Margalit and Mitts (2022) use "facts bias." We avoid using "bias" since it is defined with respect to a true parameter (but objective coverage is

or ideological side and less information unfavorable to its side.

Research on the limits of motivated reasoning suggests that the uncongenial information may be able to overcome biases in information processing (Guess and Coppock 2020; Wood and Porter 2019; Porter and Wood 2022). This could lead cross-cutting media to produce moderation instead of backlash. As Wood and Porter (2019) argue, the valence of information is often less obvious than the valence of arguments. In turn, motivated reasoners might be less likely to counterargue against—and thus more likely to 'Accept'—information. Later, this information may be present in mind and inform attitudes. For instance, video CNN played showing that Trump rallies did not require masks during the COVID-19 pandemic may lead to less counterarguing than explicit arguments that Trump performed poorly handling COVID-19. However, if individuals learn that Trump did not require masks at his rallies, when forming attitudes about his performance handling the pandemic, those who support mask-wearing (which the majority of our sample did) might evaluate Trump's performance less favorably. However, existing studies holding topics and issues constant between sources may not capture any such effect.

Our research design, described below, allows the topics and information covered in cross-cutting and congenial partisan media networks to naturally vary. It also evaluates the impact of *sustained* exposure, which evidence also suggests might lead motivated reasoners to reach a "tipping point" that leads them to revisit their views (Gerber and Green 1998; Redlawsk, Civettini and Emmerson 2010). It therefore complements prior survey-based research in conditions that may better resemble the effects of sustained exposure to real-world cross-cutting news sources. We theorize that these conditions may be better able to overcome motivated reasoning.

Our argument that sustained exposure to information in cross-cutting sources can overcome motivated reasoning to change attitudes is distinct from agenda setting, priming, and framing, three common mechanisms for media effects studied in the literature. None of these chiefly consider the role of information in media or its possible effects on learning (for review, see Table OA2).

impractical to define) and avoid using "slant" because it is used to describe news source ideology.

First, the media is thought to influence public opinion through *agenda setting* (McCombs and Shaw 1972). "The idea of agenda setting is that the public's ... beliefs about what is a significant issue or event are determined by the amount of news coverage accorded" to those events (Ansolabehere, Behr and Iyengar 1993, p. 142). Relatedly, agenda setting is also thought to make viewers bring highly-covered topics to mind when evaluating elected officials through *priming* existing attitudes (Krosnick and Kinder 1990). However, even if priming often occurs generally, priming is unlikely to be a primary mechanism by which cross-cutting exposure influences attitudes. For priming to influence attitudes, an individual must have a *pre-existing* mix of liberal and conservative attitudes on different dimensions (e.g., already thinking Trump is poorly handling the COVID-19 pandemic), thereby causing shifts when liberal instead of conservative (or conservative instead of liberal) pre-existing attitudes are primed (e.g., evaluating Trump more poorly overall after existing negative evaluations of his handling of COVID-19 are primed). However, individuals that currently consume one side's partisan media likely have consistently liberal or conservative attitudes across most dimensions (Stroud 2011), meaning which dimension is primed would rarely impact their evaluations.

Another potential mechanism by which cross-cutting media may influence viewers is *framing*. Although definitions vary, we follow definitions of framing as entailing "*emphasizing* which aspect" of a given issue is "relevant for evaluating it *without the frame itself [providing] any new substantive information about the issue*" (Leeper and Slothuus 2020, p. 154, emphasis in original). This paper does not test whether cross-cutting content moderates attitudes through framing, as doing so requires holding information and topics constant; however, as reviewed above, previous research which has done so has largely found that frames in cross-cutting content backfire.

In summary, we argue that balancing partisan media viewers' diets by exposing them to cross-cutting media will cause learning and moderate their attitudes—inconsistent with motivated reasoning and beyond predictions of agenda-setting, priming, and framing. Our argument therefore suggests that the first step in motivated reasoning theories—a preference for selective exposure to congenial information ('reception')—presents a bigger challenge to efforts to depolarize attitudes than the potential for backlash when individuals process cross-cutting information ('acceptance').

Below we present a field experiment where we test two major predictions of our argument: that balancing partisan media consumers' media diets towards cross-cutting sources will (1) lead them to learn uncongenial facts and (2) moderate their attitudes. As we describe in the discussion section, our study's advantage is its relatively greater degree of naturalism. It complements previous lab- and survey-based studies which focus on the effects of brief exposure on attitudes among the general population, and natural experiments which have focused on aggregate behavior such as vote choice (e.g., Martin and Yurukoglu 2017), by studying the effects of sustained exposure to real-world cross-cutting content on beliefs and attitudes. However, a weakness of this approach is that we are unable to tightly control the content of this coverage, meaning we are unable to fully test all the empirical implications of our argument. Nevertheless, we reach starkly different findings than prior survey-based studies, consistent with the differences in the stimuli used in our study and prior studies being theoretically significant.

Experimental Design

Treatment: Incentivizing Frequent Fox News Viewers to Watch CNN

In the fall of 2020, we conducted a pre-registered, randomized experiment that incentivized regular Fox News viewers to consume CNN.[5]

Procedures

We summarize the experimental design briefly in Figure 1 and in more detail in Figure OA1. We drew inspiration for the design from Chen and Yang (2019).

[5] Budget constraints prevented us from studying multiple networks. We do not assume that CNN is more "objective" than Fox News, nor whether it is a liberal or centrist network, only that its ideological slant is more left-wing than Fox News' (as found by Kim, Lelkes and McCrain 2022).

Figure 1: Overview of the Experimental Design

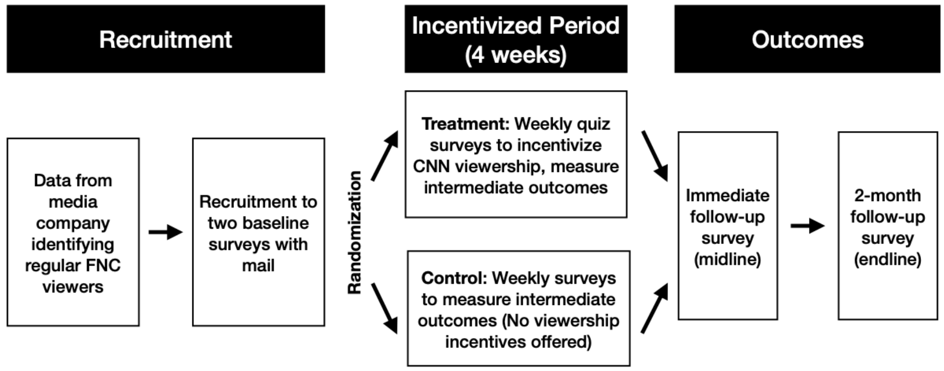

Sample. To understand the impact of consuming cross-cutting media among like-minded media consumers, one must recruit a sample that already consumes like-minded media. Such a sample is difficult to identify because partisans dramatically over-report their consumption of partisan media (Prior 2009). We overcame these challenges with a unique data source. In particular, we began by first identifying current viewers of Fox News using TV viewership data from a media analytics company, Bully Pulpit Interactive. Many modern "smart TVs" are internet-connected devices that, for users who opt-in, use automatic content recognition to measure what a particular TV is watching and then report this information back to the TV manufacturer. This information can then be matched to voter files.

Using this data from one particular smart TV brand, we identified 223,572 registered voters who BPI expected regularly watched Fox News and minimal amounts of CNN or MSNBC. In particular, we selected voters aged 18 to 89 in households which, in the months of January, March, May, and June 2020, averaged watching between 500 and 14,400 minutes of Fox News and less than 30 minutes of both CNN and MSNBC per month.

First Baseline Survey. We then mailed all 223,572 voters in these households a letter inviting them to participate in an online survey. In this initial survey, we obtained informed consent, requested an email address, asked an attention check question, and gathered demographic data. We also asked respondents to self-report their weekly TV viewership, including at which hours of the day they typically watched Fox News, and if they would be willing to participate in a study where they were paid to watch TV. $N = 15,048$ participants responded to this baseline survey.

We then narrowed the sample further to those individuals who self-reported a willingness to participate in a study where they were paid to watch TV, either self-reported at least an hour per week of Fox News or reported regularly watching one of the individual programs that aired at the same time we would later incentivize, and did not report watching more than 15 minutes per week of CNN. This left us with $N = 5,536$ participants who we invited to a second survey.

Offer Survey. This second survey, or "offer survey," asked additional background demographic questions before inviting participants to participate in an experiment. We asked participants: "We are interested in what people think when they watch TV channels different than the channels that they usually watch. Some people may be selected to earn more than $10 per survey in September if they agree to watch a new channel for a few hours and answer questions about what they saw." We then told participants they had been selected to watch CNN and gave them an option to select certain hours to watch CNN during the week. For reasons of practicality, we only gave participants the option of watching CNN during the Monday-Friday prime time hours, when viewership is highest.[6] We first showed participants only the hours during Monday-Friday prime time at which they had told us they watched Fox News during the previous week. If they selected under 7 hours on this screen, we showed them another screen that allowed them to select additional hours, up to 7 per week in total. Participants could select no hours. After participants selected hours, we then confirmed that they would fully participate with the study. We then limited our sample for the experiment to only those participants who agreed to watch at least one hour per week of CNN

[6]The CNN line-up during these hours (the shows we drove participants *to*) was Erin Burnett OutFront, Anderson Cooper 360, Cuomo Prime Time, and the first hour of CNN Tonight with Don Lemon. The Fox News line-up during these hours was The Story with Martha MacCallum, Tucker Carlson Tonight, Hannity, and The Ingraham Angle.

rather than Fox News if assigned. This left 763 individuals living in 695 households in the final sample for the experiment.

Sample Demographics and Representativeness. One potential concern with our design is that it may have selected a group of Fox News viewers whose opinions were unusually open towards CNN, outpartisans, etc. Examining the sample of subjects who participated in the experiment helps assuage these concerns. Appendix Section 4 provides more details on the demographics of the sample at each stage, illustrating both the kinds of voters ultimately included in the experiment and how the process of selecting the sample described above influenced its composition. Overall, our selection process led to a sample that appeared largely representative of the starting sample, but was even more conservative and watched even more Fox News. Relative to the average American, the individuals we selected for the study were, on average, older (average age of 54), whiter (95%), more Republican (92% self-identified as Republican), more supportive of Donald Trump (median Trump feeling thermometer rating of 90), more active voters (90% voted in the 2016 general election), and more frequent Fox News viewers (self-reported watching an average of 840 minutes per week). Moreover, participants showed no signs of being unusually open to influence from CNN: the median feeling thermometer rating of CNN among participants in the experiment was only 1 on a 0-100 scale (with a mean of 11.7). Participants were also *less* likely to say they would enjoy watching CNN (median of 1 on a 1-6 scale) than our starting sample of Fox News viewers.

These statistics underscore that our experiment represents a hard test of our argument, since motivated reasoning theories would expect this sample to be especially likely to counter-argue against cross-cutting content (e.g., Arceneaux and Johnson 2013, see also Table OA1).

Randomization. We then block-randomized at the household level $n = 304$ individuals to a treatment group paid \$15 per hour to watch CNN and $n = 459$ individuals to a control group that received no payment to watch CNN. The treatment group subjects agreed to (and we then incentivized to) watch an average of 5.8 hours of CNN per week (median of 7 hours). The incentivized period to watch CNN began on August 31, 2020 and ended on September 25, 2020.

Treatment Notification and Implementation. Because our experiment sought to test whether participants would 'accept' and not backfire against messages from cross-cutting sources conditional on reception (the information processing step of motivated reasoning theories), we took steps to increase reception of cross-cutting media as much as possible. In particular, we incentivized CNN viewing with quizzes. We told both treatment and control group participants that they would receive a series of short surveys over the course of September 2020 that we would pay them $10 each for completing. We refer to these as "quiz surveys." At the start of each week, we wrote questions probing both beliefs and attitudes about events happening in the news for the prior week. Both treatment and control group subjects received these quiz surveys at the same time, holding constant the number and timing of surveys that treatment and control subjects were invited to take. Individuals received five quiz surveys at randomly assigned times during the incentivized period. Respondents in both conditions received $10 for completing each survey.

To maximize reception of CNN, we also told individuals in the treatment group that these quiz surveys would contain a "pop quiz" about what had happened on CNN when they were supposed to be watching. The pop quiz asked about non-political features of the coverage.[7] This pop quiz came near the beginning of the survey that both treatment and control subjects were asked to complete, and only appeared for treated subjects. Every night, a research assistant watched CNN live during all four incentivized hours, drafted three pop quiz questions per incentivized hour, and sent out these quizzes within 30 minutes of the show ending. Treatment group individuals only received their bonus payment for watching CNN ($15 per hour since the last quiz survey) if they answered at least two out of three quiz questions on that quiz correctly. All perception and attitude items on

[7]For example, we asked "On Monday's program, Anderson Cooper covered the wildfires taking place across the West Coast. Who did Cooper interview about these fires? Kate Brown, Governor of Oregon; Eric Garcetti, Mayor of Los Angeles; Nancy Pelosi, Speaker of the House."

the quiz surveys which we used as outcomes in the experiment appeared *after* treated respondents finished the incentivized items and were told whether they had earned a bonus.

Treatment group subjects also received daily email and text message reminders to watch CNN.

Compliance. Our treatment was expected to both decrease Fox News viewership and increase CNN viewership. We find substantial evidence this occurred. First, compliance with watching CNN was very high in the treatment group. On average, treatment group respondents answered 12.4 out of 15 pop quiz questions correctly (median of 14). Similarly, using the television viewership data, we find that during the incentivized period, CNN viewership was significantly higher in the treatment group than in the control group ($p < 0.001$), although measurement error in the TV viewership data makes it difficult to precisely quantify how much CNN and Fox News consumption changed (see Appendix Section 8.4).

We did not explicitly instruct treatment group subjects to refrain from watching Fox during the incentivized period. However, Appendix Section 8.4 presents evidence that during the incentivized period Fox viewership also decreased in the treatment group, as measured by both the viewership data and self-reported survey data. Furthermore, as we discuss in the results section, our pattern of results suggests that participants consumed less of the prime time Fox shows as, for example, they are less aware of information reported on these shows.

Midline and Endline Surveys. The incentivized period ended on Friday, September 25, 2020. Beginning on Monday, September 28, we invited respondents to participate in a midline survey to measure treatment effects. The midline survey contained a variety of items, described in more detail later, many of which directly corresponded with the topics and information covered on both CNN and Fox News during the incentivized period. Unfortunately, space constraints prohibit us from elaborating on the related literature and motivations for all of the items we asked.

We invited all 763 individuals randomized to treatment or control to respond to this midline survey, with 744 participating (97.5%). We closed this survey on October 14.

Finally, beginning on November 20, we invited individuals to participate in an endline survey. A total of 727 (95.3%) responded. We closed this survey on December 9.

Context: Fox News and CNN Coverage During September 2020

To aid in the interpretation of our experimental results, we next contextualize the coverage on Fox News and CNN during the treatment period (August 31 - September 25, 2020) and the hours when treatment group subjects were incentivized to watch CNN instead of Fox News. To do so, a research assistant read all the transcripts from both networks during this period and totalled the number of words associated with each topic and subtopic (see Appendix 9 for details).

CNN and Fox News covered dramatically different topic areas during this period. For example, Fox News had 2.3 times more coverage of racial protests than CNN while CNN had 2.6 times more coverage of COVID-19 than Fox News (see Figure OA31 for additional topic areas). Furthermore, within topic areas, CNN and Fox News covered different subtopics, corresponding with different information (partisan coverage filtering). The top panel of Figure 2 presents the ten most common subtopics CNN discussed and how often these were discussed on Fox News. The bottom panel does the same with the most common Fox News subtopics.

Consistent with partisan coverage filtering, Fox News was far more likely to report facts favorable to Republicans while CNN was far more likely to do the same for Democrats. For example, CNN extensively reported "Trump's failures to protect US & his supporters from COVID-19," while Fox News spent little time doing so. Likewise, CNN spent 10,251 words discussing the severity of COVID-19, while Fox News devoted only 709 words to this. Instead, Fox News reported information downplaying the severity of COVID-19 and the efforts Trump had undertaken to protect Americans from the virus. On the other hand, Fox News' main focus during this time was on racial issues and related racial protests in American cities during the summer of 2020; Fox News indicated that Joe Biden and Democrats generally supported the protesters' tactics and demands. Both networks covered voting by mail, but provided different information about it.

Figure 2: Transcript Analysis During Treatment Period

(a) Top 10 CNN Subtopics

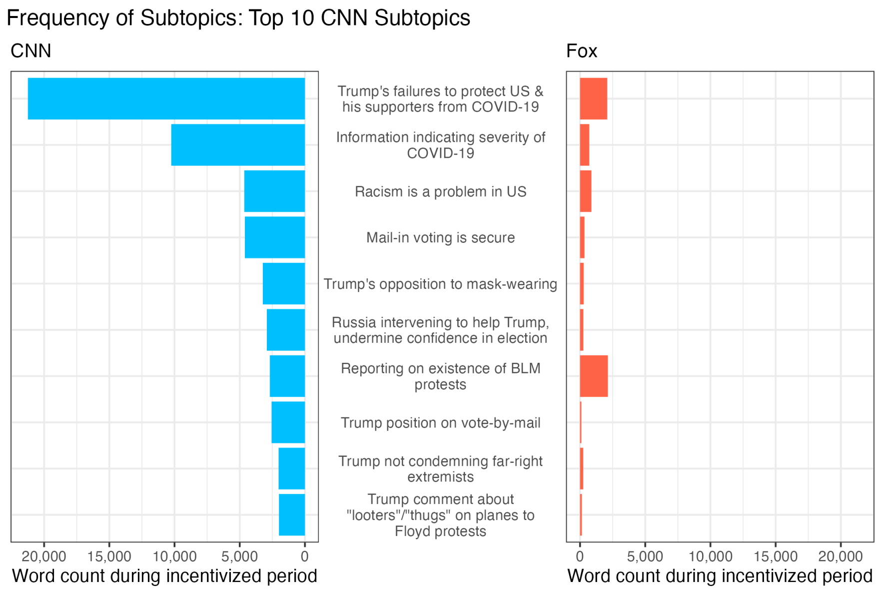

(b) Top 10 Fox News Subtopics

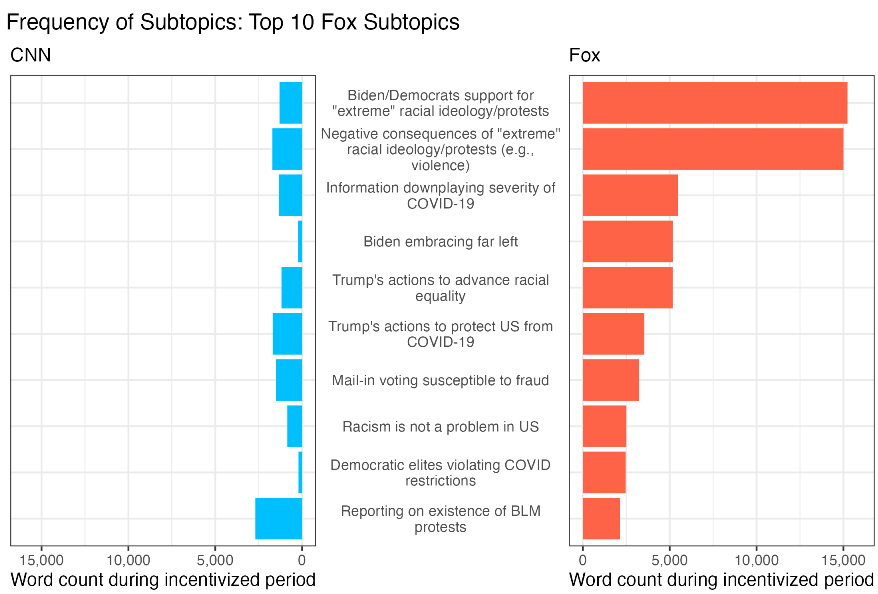

Notes: Table OA15 presents numerical values for the entire set of topics coded. Figure OA31 sums up these results by topic area and shows that CNN and Fox News covered different topics as well.

Analytical Strategy

Following our pre-analysis plan (see Online Appendix Section 2), we estimate the effects of incentivizing CNN viewership by comparing survey responses among those assigned to the treatment group to those assigned to the control group (intent to treat). We use OLS with pre-treatment covariates and standard errors clustered by household. See Appendix Section 7 for details.

As we pre-registered, we report three types of p-values in order to adjust for multiple comparisons. First, we report conventional, unadjusted p-values from covariate-adjusted OLS regression. Second, we report false discovery rate sharpened q-values (Anderson 2008), similar to other recent experiments on media (e.g., Chen and Yang 2019). The q-values are adjusted for false discovery rates across all the items in the entire survey. We separately adjust the results on the indices only. These q-values control the probability of making individual false discoveries; e.g., we should expect only 5% of results with a q-value under 0.05 to be false positives (Type I errors). Finally, in the Online Appendix, we report family-wise error rate adjusted p-values for the individual items. These are much more conservative, and control the probability of making *any* Type I errors at all within each family of outcomes specified in our pre-analysis plans.

We pre-specified that we would form outcome indices by combining multiple survey measures into a single index. We pre-registered which survey items belonged in which index. We formed these indices by first standardizing all individual items to have mean 0 and standard deviation 1 before forming an additive index of these rescaled items, reverse coding items as appropriate. All reverse coding decisions were pre-specified. Full results on all indicies and items are available in the Online Appendix.

Given the number of hypotheses we tested, in the main text, we primarily focus on results on individual items that are statistically significant after applying a pre-registered multiple testing correction, although alongside these results we also note corresponding results on our pre-specified indices and discuss several null results.

Appendix Section 7 presents additional details on our analytical strategy. In addition, Online

Appendices 5 and 6 present tests of design assumptions, in particular tests for covariate balance at each stage and tests for differential attrition.

Experimental Results

Our argument is that sustained exposure to cross-cutting media should cause individuals to learn more information about the events covered on the cross-cutting media and less information about the events covered on like-minded media (since they consume less of it), and that this learning helps moderate attitudes. This stands in contrast to expectations from motivated reasoning theories, which argue that cross-cutting media exposure should lead to counterarguing and backlash, further polarizing attitudes. In the next sections, we present results consistent with our argument. There is not a single case across all the items in the entire paper where we find statistically significant evidence of backlash.

CNN Learning Effects

Motivated reasoning theories argue that individuals' beliefs backlash from exposure to uncongenial information, especially from uncongenial sources. However, we see evidence that the treatment group learned information from consuming cross-cutting media (CNN). Figure 3 provides examples of learning we observed in the quizzes and the midline survey on items that CNN covered substantially but that received minimal coverage on Fox News. The q-values shown on the right side of Figure 3 show the false discovery rate-adjusted q-values for each statistical test. Note that in all figures reverse coded items are coded positively if the treatment group was *less* likely to agree with them; e.g., the third coefficient in Figure 3 shows that participants incentivized to watch CNN were *less* likely to believe that Donald Trump's campaign was taking safety precautions at his rallies.

The evidence in Figure 3 shows multiple examples of learning. For example, as shown in Figure 2a, CNN provided 14 times more coverage on the severity of COVID-19 than Fox News. Consistent with the treatment group learning from CNN exposure, we find that they were 0.18 standard deviations more likely to agree that "The coronavirus causes many people to experience

Figure 3: Effect of Learning on Selected Outcomes Likely Caused by Increasing CNN Viewership

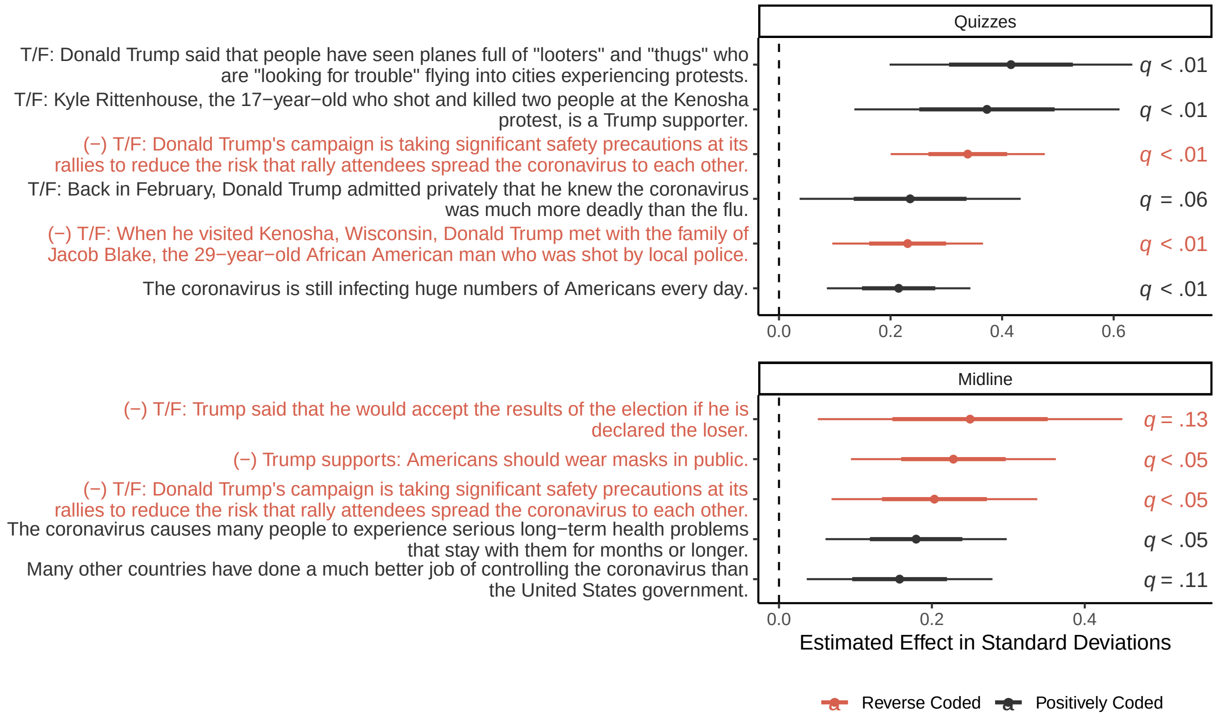

Notes: Standard errors (thick lines) and 95% confidence intervals (thin) surround point estimates. Full results on all pre-registered outcomes and indices are in the Online Appendix.

serious long-term health problems that stay with them for months or longer" ($p_{unadj.} = 0.003$; $q < 0.05$). Participants also learned from CNN about COVID-19, Trump's claims of election fraud, and Trump's role in racial protests. These results on individual items are consistent with the significant effects we find on the Liberal Perceptions of Events CNN Covered (Non-COVID) Index (Figure OA7) and Increased Knowledge of CNN-Covered Trump Positions Index (Figure OA13).

This evidence suggests that sustained cross-cutting media leads individuals to update their beliefs consistent with the cross-cutting source's message, contrary to motivated reasoning theories.

Reduced Fox Learning Effects

Our treatment both increased CNN viewership and decreased Fox News viewership. As a result, we also found decreases in knowledge of information Fox News covered. Results are summarized in Figure 4.

Figure 4: Effect of Learning on Selected Outcomes Likely Caused by Decreasing Fox Viewership

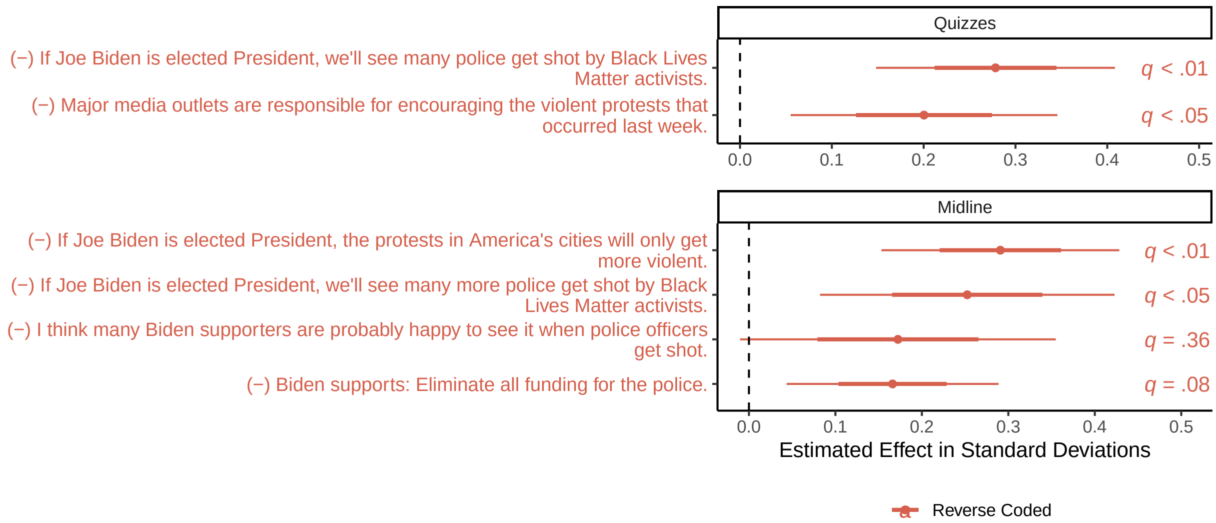

Notes: Standard errors (thick lines) and 95% confidence intervals (thin) surround point estimates. Full results on all pre-registered outcomes and indices are in the Online Appendix.

For example, Fox News was 12 times more likely to cover ties between Biden and racial protests and 9 times more likely to cover the negative consequences (e.g., violence and property damage) of racial protests than CNN. When we reduce the amount of Fox News that the treatment group consumed, we find that the treatment group becomes less likely to believe that negative aspects of racial protests are linked to Biden. The treatment group was 0.29 standard deviations less likely to agree that “If Joe Biden is elected President, the protests in America’s cities will only get more violent” ($p_{unadj.} < 0.001; q = 0.007$). Similarly, the treatment group was 0.25 standard deviations more likely to agree that “If Joe Biden is elected President, we’ll see many more police get shot by Black Lives Matter activists” ($p_{unadj.} < 0.01; q < 0.05$).

These results are consistent with our findings on the Liberal Perceptions of Events Fox Covered (Non-COVID) Index (Figure OA8).

Attitude Change

In the prior sections, we showed that the treatment group learned more about issues and events that CNN covered and learned less about issues and events that Fox News covered. In this section, we examine the consequences of increasing CNN viewership and decreasing Fox News viewership on broader political attitudes. We again fail to find any statistically significant evidence of backlash across any of the items we measured.

Figure 5 provides examples of changes in political attitudes likely caused by learning information from CNN or not learning information from Fox News (or both).

Figure 5: Selected Attitude Changes Likely Caused by Learning

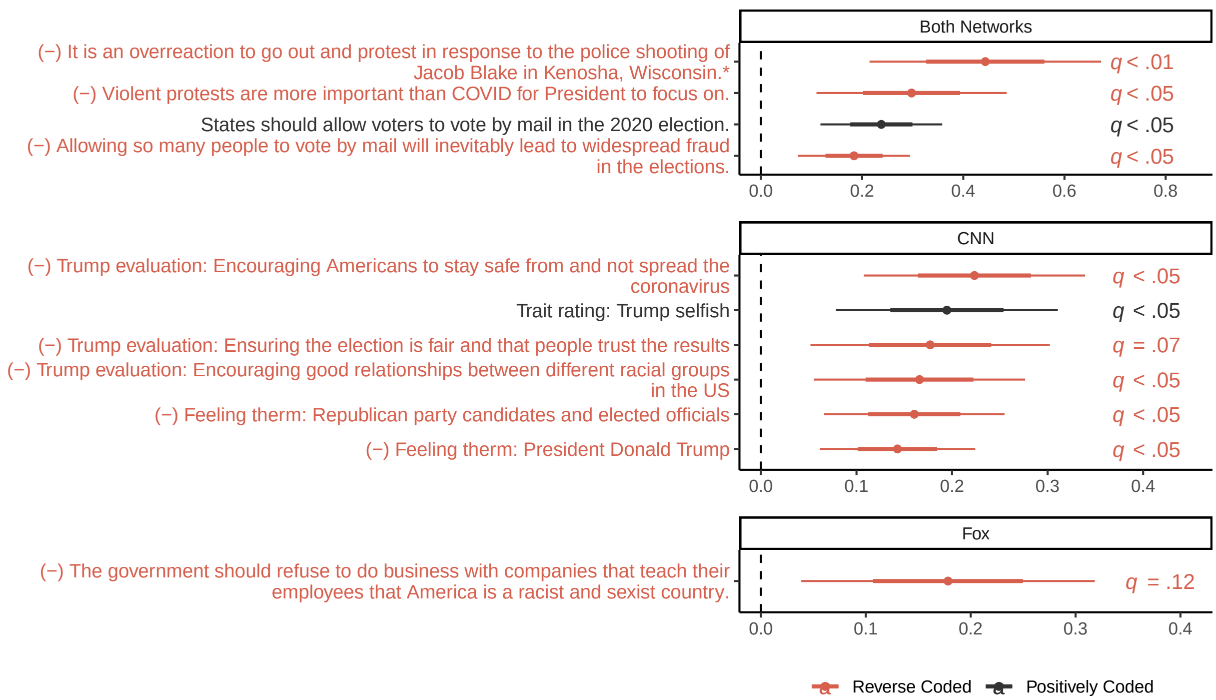

Notes: Standard errors (thick lines) and 95% confidence intervals (thin) surround point estimates. The item ending with ∗ was asked during Quiz 1. All other items come from the midline survey. Full results on all pre-registered outcomes and indices are in the Online Appendix.

For example, COVID-19 received extensive coverage on CNN. As shown in Figure 3, individuals in the treatment group were more likely to learn that COVID-19 was infecting huge numbers of Americans every day and led to serious long-term health problems. They also learned that Trump opposes mask wearing, Trump's campaign was not taking safety precautions, and that many other countries have done a better job at controlling COVID-19 than the U.S. We argued that learning these negative facts about Donald Trump should contribute to to increased negative evaluations of Trump. Theories of motivated reasoning, on the other hand, would predict that exposure to these facts would lead to counterarguing and a backfire effect. Under theories of motivated reasoning, we should expect that exposure to negatively-valenced information about Trump on CNN should lead to more *positive* evaluations of Trump among our sample of Fox News viewers who, at baseline, viewed Trump favorably.

Contrary to theories of motivated reasoning, Figure 5 shows increased negative evaluations of Trump. For example, we find that the treatment group reduced its evaluation of Trump in a feeling thermometer by 0.14 standard deviations ($p_{unadj.} < 0.001; q < 0.05$) and became 0.19 standard deviations more likely to view Trump as selfish ($p_{unadj.} = 0.001; q < 0.05$). Specifically on COVID-19, the treatment group became 0.22 standard deviations less likely to positively evaluate Trump's ability to keep Americans safe from COVID-19 ($p_{unadj.} < 0.001; q < 0.05$). These results are consistent with our findings on the Reduced Trump Evaluation Index (Figure OA24).

Notably, these include effects on a number of items exclusively or nearly exclusively covered on CNN, such as those related to Donald Trump's alleged failure to encourage Americans to stay safe from the coronavirus and refusal to say he would accept the results of the election. This indicates the effects we found are not exclusively due to reducing Fox consumption, but include effects of consuming cross-cutting media (CNN).

Figure 5 also shows that consuming CNN moderated habitual Fox viewers' attitudes on issues, in particular on voting-by-mail and racial protests. For instance, treated participants were more likely to agree that states should allow voters to vote by mail in the 2020 election (0.24 standard deviations; $p_{unadj.} < 0.001; q < 0.05$) and less likely to agree that violent protests rather than COVID were important for the President to focus on (0.30 standard deviations; $p_{unadj.} = 0.002; q < 0.05$).

For readers interested in more interpretable estimates, Table OA16 provides estimates on dichotomized versions of several items.

Null Effects

We did not find effects on all political attitudes—especially on those attitudes that received little coverage during the incentivized period. Consistent with this, Figure OA32 shows the relationship between coverage volume about specific topics and treatment effect estimates on items related to those topics in cases where we could make such a match, finding a positive relationship.

More generally, our results also reveal likely scope conditions on the effects of consuming cross-cutting content. Figure 6 shows additional null findings on several pre-registered indices that were related to the networks' coverage but not the coverage's direct focus. We generally see null effects in these cases. For instance, although racial protests were widely covered, neither network made explicit arguments about the superiority or inferiority of different racial groups. Perhaps as a result, we see no effects on an index of items measuring racial prejudice. We see similar null effects on other items related to issues that did not receive direct coverage, such as the virtues of various democratic norms and on several issues—immigration, free trade, and climate change—that received little to no coverage on CNN and Fox News during the incentivized period.

With that said, despite finding many null effects, there is not a single case across all the items in the entire paper where we find statistically significant evidence of backlash.

Robustness and Mechanisms

First, motivated reasoning theorists might argue that backlash would only be found in the subset of our sample with the strongest attitudes. Our sample already had fairly homogeneously strong attitudes. However, Appendix Section 8.5.1 presents several tests for heterogeneous treatment effects across two pre-registered moderators: an index of baseline items capturing respondents' strength of Republican identification/support and an index of baseline Fox viewership frequency. Reassuringly, we found little evidence of backlash in any subgroup and that the effects generally manifested broadly across the sample.

Second, our argument holds that this cross-cutting exposure was able to overcome motivated reasoning, unlike what has been observed in prior lab- and survey-based studies, due to *sustained*

Figure 6: Null Effects on the Midline Survey

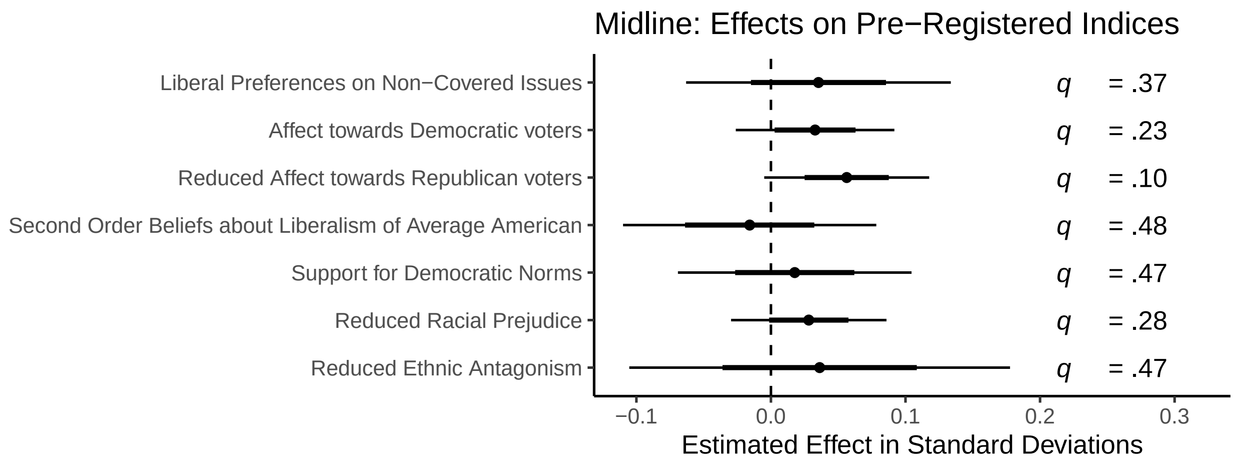

Notes: Standard errors (thick lines) and 95% confidence intervals (thin) surround point estimates. Full results on all pre-registered outcomes and indices are in the Online Appendix. When affective polarization is defined as the difference between the Republican and Democratic voters feeling thermometer ratings, the treatment effect becomes marginally significant ($p_{unadj.} = 0.04$)*, although this would not remain significant after multiple comparison adjustment.*

exposure to uncongenial *information*. Our argument relies on causal mechanisms already proven out in prior work (e.g., Wood and Porter 2019; Redlawsk, Civettini and Emmerson 2010), and, as with any field-based study, cannot definitively determine the role of those mechanisms in generating the effects we find.

However, we did conduct several tests that rule out potential alternative mechanisms. First, the mechanism for the effects of cross-cutting exposure we found on beliefs seems very likely to be information (e.g., information that long-COVID exists on CNN is likely what led treated participants to be aware that it exists). Second, Appendix 8.5.2 presents tests that rule out priming as a mechanism for our findings: the effects are not concentrated among those with baseline liberal attitudes which were primed; as selective exposure predicts, few participants had liberal attitudes available to prime in the first place. Third, agenda setting theory's predictions are limited to effects on issues' perceived importance, but we found results on items beyond this. Fourth, framing the same issues differently cannot be wholly responsible for our findings, as we found effects on issues which were not presented at all on one of the two networks. To be clear, these results do not cast doubt on agenda-setting, framing, or priming theories nor rule out that they may have contributed to some of our results; rather, they indicate that these theories are insufficient to explain all of our findings.

Endline Results

Finally, two months after the end of the incentivized period, we launched the endline survey. The endline survey asked most of the same items as on the midline survey, and no new items.

First, we found little evidence that the treatment affected long-run viewership habits (Figure 7). In neither the endline survey nor the viewership data did we find any long-run effects on Fox viewership. However, in the endline survey, the treatment group reported watching 4.5 additional minutes of CNN yesterday ($p_{unadj.} < 0.01; q < 0.05$) and 16 additional minutes over the course of the past week ($p_{unadj.} = 0.06; q = 0.37$). The television viewership data, though, finds no effects on long-run CNN viewership. This discrepancy could be caused by under-counting in the television viewership data (see Appendix 8.4) or over-reporting in the survey data; these conflicting results are ambiguous. Regardless, we can rule out large effects on long-run CNN viewership, despite ambiguity in whether there were small effects.

Consistent with this at-most-minimal long-run impact on media consumption, Figure OA30 presents largely null results on the attitudes and beliefs measured in the endline survey. However, given that many of the confidence intervals are large and overlap with the effect estimates in the midline survey, these results leave somewhat ambiguous whether the effects partially persisted or entirely decayed. We discuss our interpretation of these results in the discussion.

Discussion

Scholars, civil society leaders, and classic thinkers alike extol the benefits of information and media sources inconsistent with one's beliefs. Consistent with this longstanding view, as fear has grown about the effects of many Americans' near-exclusive consumption of like-minded media sources (e.g., Stroud 2011), there have been increasing calls for Americans to consume cross-

Figure 7: Long-Term Treatment Effect on Television Viewership

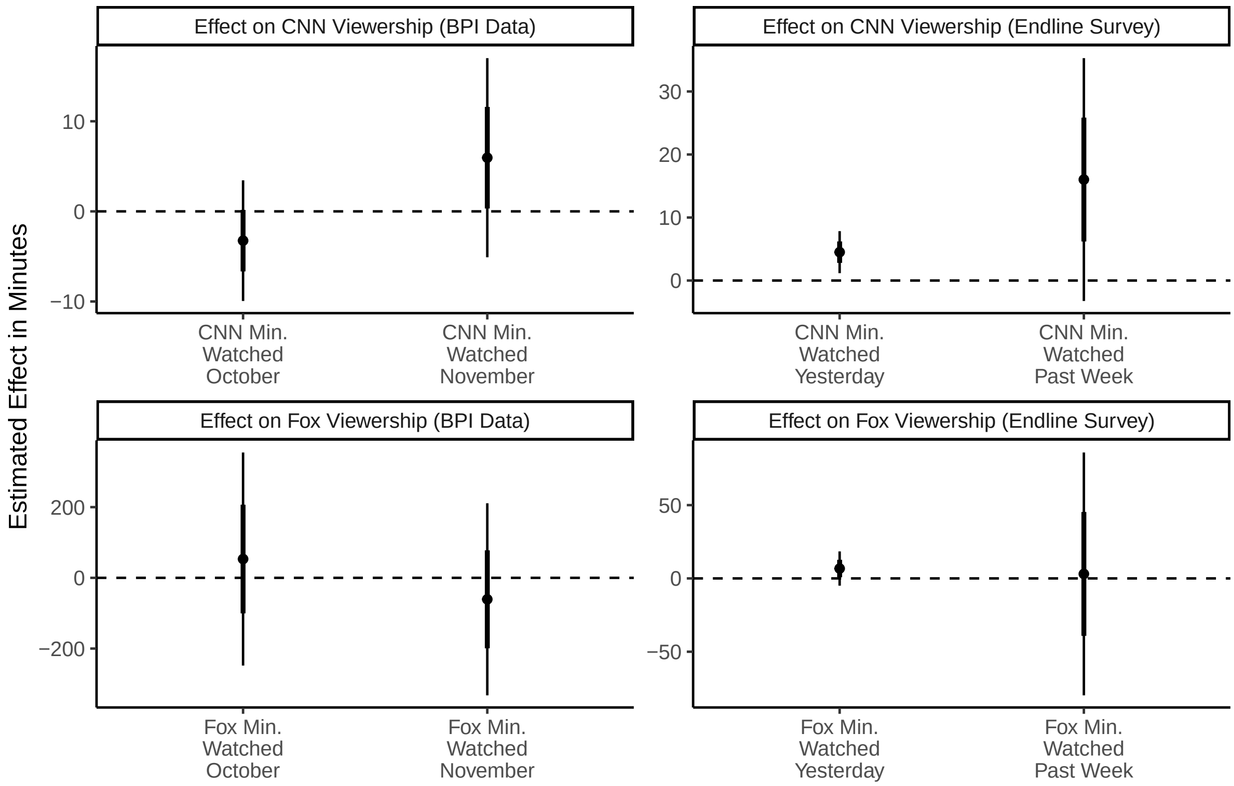

Notes: Standard errors (thick lines) and 95% confidence intervals (thin) surround point estimates.

cutting media that might moderate their beliefs. However, an influential perspective in contemporary scholarship has warned against such cross-cutting exposure. For instance, Arceneaux and Johnson (2013, p. 74) warn that "exposure to counterattitudinal news can be just as polarizing as exposure to proattitudinal news." Or, at best, this perspective argues cross-cutting exposure should have no effects at all. Such warnings are rooted in motivated reasoning theory, which predicts that individuals exposed to cross-cutting sources and information often counter-argue against it, producing backlash in both beliefs and attitudes (e.g., Nyhan and Reifler 2010; Lodge and Taber 2013). Survey- and lab-based experiments on partisan media bear out this warning, finding that exposure to brief clips of outpartisan media can polarize attitudes (see Table OA1 for review).

We argued that sustained exposure to real-world cross-cutting media sources should moderate attitudes, not produce backlash. Our argument is rooted in two key differences between prior lab-based studies and real-world consumption of cross-cutting media: real-world cross-cutting media generally covers different topics and information than congenial media (Baum and Groeling 2008), and real-world exposure can be more sustained than the brief exposure tested in prior studies. We theorized that both of these features of real-world cross-cutting consumption could lead cross-cutting consumption to overcome motivated reasoning and moderate attitudes.

Our experiment supported this argument: Fox News viewers incentivized to watch CNN instead for a month learned information CNN presented, and their attitudes on political issues CNN covered and towards Donald Trump moderated. Despite conducting our study among the population that motivated reasoning theories would most expect to display evidence of backlash or to resist persuasion—highly conservative, regular Fox News viewers—we do not find a single case across all of our measures of statistically significant evidence of backlash, and many cases where we find learning and moderation.

These results were by no means obvious. Classic theories posit that, for a media source to influence Americans' attitudes, they must both 'receive' (i.e., consume) and 'accept' its contents (Zaller 1992). Even if partisan media consumers were to 'receive' messages from cross-cutting sources, many scholars are skeptical that they would 'accept' uncongenial messages from these sources due to motivated reasoning—indeed, as noted, many expect such messages would backfire, "lead[ing] people to become more extreme" (e.g., Arceneaux and Johnson 2013, p. 116). Our findings thus stand in contrast to the expectations of theories of motivated reasoning that exposure to cross-cutting media would have no effects or, more worryingly, produce backlash and further exacerbate polarization.

With this said, our evidence does not conflict with other predictions of motivated reasoning theories, and in fact even supports some of them. Lodge and Taber (2013) identify three mechanisms through which motivated reasoning may operate: first, confirmation bias in the selection of sources (preferring to consume congenial sources, sometimes called selective exposure); second, conditional on reception of content, prior attitude effect (viewing uncongenial content as less credible); and third, again conditional on reception of content, disconfirmation bias (counterarguing that can lead to backlash). Only disconfirmation bias predicts backlash conditional on reception of a source, and thus is the basis of scholars' warnings against encouraging cross-cutting conception (Druckman and McGrath 2019). And it is only disconfirmation bias with which our results are therefore inconsistent: we find that cross-cutting exposure moderates attitudes, and no evidence whatsoever of backlash.

However, we hasten to note that our evidence offers some (albeit quite limited) evidence consistent with the other two predictions of motivated reasoning theory. First, our evidence is consistent with the existence of selective exposure: the regular Fox News viewers who we recruited to our study were nearly all very conservative to begin with. We note, however, that there are interpretations of this pattern other than motivated reasoning (e.g., Fowler and Kim 2022). For example, Fox News viewers may prefer consuming Fox News because they think it is more credible, not because they enjoy having their priors confirmed. We plan to further characterize the extent and nature of selective exposure in future work.

Second, the implications of our evidence for the prior attitude effect are ambiguous. Consistent with it, the sample noted at baseline they viewed CNN as untrustworthy, and, despite potentially small increases in trust during the incentivized period (see Figure OA4), largely still did after the study's conclusion (see Figure OA19). At the same time, they still updated their beliefs (i.e., learned) from CNN despite stating that they did not trust it, suggesting this stated distrust was not fully sincere. This evidence does not allow us to reject the prior attitude effect, though. Recent work in formal theory has studied the prior attitude effect and its implications (often using the shorthand of 'motivated reasoning') (Little 2022). This research conceptualizes the prior attitude effect as leading individuals to update less than they should were they fully Bayesian due to directional motivations. However, Little (2022) shows that it is infeasible to determine in empirical data whether individuals act to some extent as the prior attitude effect predicts or update their beliefs based on new information fully as much as they 'should' under Bayes' rule, in part due to the difficulties of measuring priors and likelihood functions (see also Hill 2017). With respect to motivated reasoning's predictions in this area, then, our results—like potentially all empirical results (Little 2022)—are ambiguous.[8]

Demand effects are unlikely to explain our findings: we found a number of effects decayed, indicating that participants did not simply always give us liberal answers; and we found null effects on many items, especially those which did not receive substantial coverage on Fox or CNN. It also seems unlikely that demand would have led participants to feign ignorance of information reported on Fox, and is unclear how it would have led them to feign knowledge of the information reported on CNN other than by watching CNN and learning this information.

Several limitations of our findings do merit emphasis, however. First, because we sought to maximize reception to cross-cutting media and see whether these messages would be accepted, our experiment may have led individuals to pay unusually close attention to CNN, since they knew we would quiz them on its content (i.e., we manipulated the reception step extremely strongly). This means our results may not speak to how different populations of consumers might consume partisan media under different circumstances; e.g., the impacts of partisan media on low-information voters seeking news just before an election may be different, although other research finds they are substantial (DellaVigna and Kaplan 2007; Hopkins and Ladd 2014; Martin and Yurukoglu 2017).

Second, we measure only the direct effects of the shift in media diets on the individuals in our study and did not measure the potential indirect effects of our study participants' conversations with others not in the study on those non-participants attitudes. At the same time, counter-arguing from

[8]Some may argue that finding persuasion from partisan media is inconsistent with Bayesian learning because individuals aware of a source's slant should discount it. As we note below, the treatment changed individuals' perceptions of Fox News's coverage of Trump, suggesting individuals may not be fully aware of their preferred source's slant without exposure to alternative sources.

others could have attenuated our estimates of the direct effect of the treatment on study participants. Third, although we found that the sample in the experiment appeared fairly representative of the starting sample, was highly conservative, and had extremely negative views of CNN (see Table OA3), the effects we observe may not be generalizable beyond the sample of participants willing to be paid to watch a different news network. Fourth, our experiment was not well-positioned to measure broader impacts of partisan media, such as for outcomes including what other media sources cover, donation behavior, or elite behavior, themes investigated in other research (e.g., Clinton and Enamorado 2014). Finally, due to budget constraints, we only considered the effect of shifting Fox News viewers to CNN. While our argument would expect similar effects among viewers of other partisan media networks, future work should attempt to replicate this, including with attention to local TV news (Martin and McCrain 2019).

Our findings also point to areas for future research. First, many studies of motivated reasoning and media consumption take survey respondents at face value when they say they do not trust various sources. However, despite that most of our sample expressed extremely negative attitudes towards and complete distrust of CNN, we found that they still learned from it. This disjuncture suggests that citizens might portray themselves in surveys as more motivated reasoners than they really are, and merits further inquiry.

Second, our results suggest future research may wish to consider how to encourage consumption of cross-cutting content. Our results on attitudes and long-run consumption suggest that voters have strong preferences for consuming like-minded media, and, relative to consuming a more balanced media diet, consuming like-minded media appears to bolster their partisan loyalties and polarize their attitudes. If individuals were more motivated to consume cross-cutting content, our results suggest that voters would have more moderate, less polarized attitudes—thus raising the question of how to encourage such consumption.

Finally, our findings raise concern about the potential implications of partisan media for democratic accountability. Media outlets plays a central role in helping voters hold elected officials accountable (e.g., Hopkins and Pettingill 2018). By the same token, not covering information—what we call partisan coverage filtering—can undermine voters' ability to hold their elected officials accountable (Besley and Prat 2006). Our evidence in Figures 3 and 4 indicates that partisan media may do exactly this. Participants agreed: we find a 0.20 standard deviation effect on disagreement with the statement "If Donald Trump did something bad, Fox News would discuss it" ($p_{unadj.} < 0.01; q = 0.02$). This may have broader implications for democracy. For instance, even though switching to CNN unsurprisingly did not induce the highly conservative participants in our experiment to prefer Joe Biden in the 2020 presidential election, it did meaningfully reduce evaluations of Trump's performance in key areas and overall. Fox News' coverage therefore likely had important political implications at a nationwide scale: our evidence indicates that Fox News shielded its viewers from information about Trump's mishandling of COVID-19, which would have led them to view Trump's handling of COVID-19 more negatively had they been aware of it. At the same time, our results suggest these effects may also last only as long as individuals are willing to consume cross-cutting content; and our finding that individuals returned to watching Fox News suggests this may prove challenging. Viewed from this vantage point, partisan media is not simply a challenge for the opposing party—it may present a challenge for democracy.

Acknowledgements

We thank Hunt Allcott, Vin Arceneaux, P.M. Aronow, Elliot Ash, Adam Berinsky, Julia Cage, Charles Cameron, Josh Clinton, Jamie Druckman, Ted Enamorado, Marcel Garz, Matt Gentzkow, Don Green, Guy Grossman, Andy Guess, Thomas Fujiwara, Jacob Hacker, Dan Hopkins, Greg Huber, Shanto Iyengar, Jenn Jerit, Samara Klar, Matt Levendusky, Ro'ee Levy, Zhao Li, Andrew Little, Arthur Lupia, Greg Martin, Diana Mutz, Brendan Nyhan, Betsy Levy Paluck, Ethan Porter, Markus Prior, Otis Reid, Todd Rogers, Cyrus Samii, Rune Slothuus, Aaron Strauss, Ian Turner, and David Yang for helpful feedback and Jacob Levy for his work as a research assistant. We are also thankful for useful conversations at the Evidence in Governance and Politics (EGAP) Feedback Sessions and presentations at Columbia, CUNY, DataColada, Harvard, MIT, Princeton, Stanford, UCLA, University of Amsterdam, and Yale. All remaining errors are our own. This paper was previously circulated with the title "The manifold effects of partisan media on viewers' beliefs and attitudes: A field experiment with Fox News viewers."

References

Anderson, Michael L. 2008. "Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American statistical Association* 103(484):1481–1495.

Ansolabehere, Stephen, Roy L Behr and Shanto Iyengar. 1993. *The media game*. Macmillan.

Arceneaux, Kevin and Martin Johnson. 2013. *Changing minds or changing channels?: Partisan news in an age of choice*. University of Chicago Press.

Baum, Matthew A and Tim Groeling. 2008. "New media and the polarization of American political discourse." *Political Communication* 25(4):345–365.

Besley, Timothy and Andrea Prat. 2006. "Handcuffs for the grabbing hand? Media capture and government accountability." *American Economic Review* 96(3):720–736.

Chen, Yuyu and David Yang. 2019. "The Impact of Media Censorship: 1984 or Brave New World?" *American Economic Review* 109(6):2294–2332.

Clinton, Joshua D and Ted Enamorado. 2014. "The national news media's effect on Congress: How Fox News affected elites in Congress." *Journal of Politics* 76(4):928–943.

Coppock, Alexander. 2022. *Persuasion in parallel*. Chicago.

DellaVigna, Stefano and Ethan Kaplan. 2007. "The Fox News effect: Media bias and voting." *Quarterly Journal of Economics* 122(3):1187–1234.

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

Fowler, Anthony and Kisoo Kim. 2022. "An information-based explanation for partisan media sorting." *Journal of Theoretical Politics* 34(4):499–526.

Gentzkow, Matthew, Jesse M Shapiro and Daniel F Stone. 2016. Media bias in the marketplace: Theory. In *Handbook of media economics*. Vol. 1B Elsevier pp. 623–645.

Gerber, Alan and Donald P Green. 1998. "Rational learning and partisan attitudes." *American journal of political science* pp. 794–818.

Goldman, Seth K and Diana C Mutz. 2011. "The friendly media phenomenon: A cross-national analysis of cross-cutting exposure." *Political Communication* 28(1):42–66.

Grossman, Guy, Yotam Margalit and Tamar Mitts. 2022. "How the Ultra-Rich Use Media Ownership as a Political Investment." *Journal of Politics* .

Guess, Andrew and Alexander Coppock. 2020. "Does counter-attitudinal information cause backlash? Results from three large survey experiments." *British Journal of Political Science* 50(4):1497–1515.

Hayakawa, Samuel Ichiyé. 1940. *Language in thought and action*. New York: Hartcourt, Brace and Company.

Hill, Seth J. 2017. "Learning together slowly: Bayesian learning about political facts." *Journal of Politics* 79(4):1403–1418.

Hopkins, Daniel J and Lindsay M Pettingill. 2018. "Retrospective voting in big-city US mayoral elections." *Political Science Research and Methods* 6(4):697–714.

Hopkins, Daniel and Jonathan Ladd. 2014. "The Consequences of Broader Media Choice: Evidence from the Expansion of Fox News." *Quarterly Journal of Political Science* 9(1):115–135.

Khanna, Kabir and Gaurav Sood. 2018. "Motivated responding in studies of factual learning." *Political Behavior* 40:79–101.

Kim, Eunji, Yphtach Lelkes and Joshua McCrain. 2022. "Measuring dynamic media bias." *Proceedings of the National Academy of Sciences* 119(32):e2202197119.

Krosnick, Jon A and Donald R Kinder. 1990. "Altering the foundations of support for the president through priming." *American Political Science Review* 84(2):497–512.

Leeper, Thomas J. and Rune Slothuus. 2020. How the news media persuades: Framing effects and beyond. In *Oxford Handbook of Electoral Persuasion*, ed. Elizabeth Suhay, Bernard Grofman and Alexander H. Trechsel. Oxford: Oxford University Press pp. 151–168.

Levendusky, Matthew S. 2013. "Why do partisan media polarize viewers?" *American Journal of Political Science* 57(3):611–623.

Little, Andrew T. 2022. "Detecting motivated reasoning." Available at `osf.io/b8tvk/`.

Lodge, Milton and Charles S Taber. 2013. *The rationalizing voter*. Cambridge.

Manjoo, Farhad. 2008. *True Enough: Learning to Live in a Post-Fact Society*. Wiley.

Martin, Gregory J and Ali Yurukoglu. 2017. "Bias in cable news: Persuasion and polarization." *American Economic Review* 107(9):2565–99.

Martin, Gregory J and Joshua McCrain. 2019. "Local news and national politics." *American Political Science Review* 113(2):372–384.

McCombs, Maxwell E and Donald L Shaw. 1972. "The agenda-setting function of mass media." *Public opinion quarterly* 36(2):176–187.

Mill, John Stuart. 1848. *Principles of Political Economy with some of their Applications to Social Philosophy*. London: John W. Parker, West Strand.

Muise, Daniel, Homa Hosseinmardi, Baird Howland, Markus Mobius, David Rothschild and Duncan J Watts. 2022. "Quantifying partisan news diets in Web and TV audiences." *Science Advances* 8(28):eabn0083.

Mullainathan, Sendhil and Andrei Shleifer. 2005. "The market for news." *American economic review* 95(4):1031–1053.

Nyhan, Brendan. 2014. "The Rationalizing Voter." *Public Opinion Quarterly* 78(S1):365–367.

Nyhan, Brendan and Jason Reifler. 2010. "When corrections fail: The persistence of political misperceptions." *Political Behavior* 32(2):303–330.

Porter, Ethan and Thomas J Wood. 2022. "Political Misinformation and Factual Corrections on the Facebook News Feed: Experimental Evidence." *Journal of Politics* 84(3):1812–1817.

Prior, Markus. 2009. "The immensely inflated news audience: Assessing bias in self-reported news exposure." *Public Opinion Quarterly* 73(1):130–143.

Prior, Markus. 2013. "Media and political polarization." *Annual Review of Political Science* 16:101–127.

Redlawsk, David P, Andrew JW Civettini and Karen M Emmerson. 2010. "The affective tipping point: Do motivated reasoners ever "get it"?" *Political Psychology* 31(4):563–593.

Stroud, Natalie Jomini. 2011. *Niche news: The politics of news choice*. Oxford University Press.

Sunstein, Cass R. 2007. *Republic 2.0*. Princeton.

Wood, Thomas and Ethan Porter. 2019. "The elusive backfire effect: Mass attitudes' steadfast factual adherence." *Political Behavior* 41(1):135–163.

Zaller, John R. 1992. *The nature and origins of mass opinion*. Cambridge university press.

Biographical Statement

David E. Broockman is Associate Professor of Political Science at the University of California, Berkeley, CA 94720. Joshua L. Kalla is Associate Professor of Political Science and Statistics & Data Science at Yale University, New Haven, CT 06520.

Online Appendix

Contents

1 Review of Survey- and Lab-Based Studies in American Politics

In this section we briefly review the design features and results of previous survey- and lab-based studies on the effects of consuming cross-cutting television partisan media in the American political context. Our review is limited to studies which 1) expose participants to cross-cutting *television* media and 2) measure one of the following outcomes: issue views; issue importance, salience, or priorities; evaluations of in-party politicians or the party as a whole; or evaluations of out-party politicians or the party as a whole. Note that these criteria exclude studies which exclusively use text-based treatments (e.g., de Benedictis-Kessner et al., 2019) and studies which examine other outcomes (for example, some studies focus on evaluations of media itself, e.g., Druckman et al., 2019).

Table OA1 presents our review, comparing features of these studies to features of our study. There are several important differences to note:

- It is rare for existing studies to allow the issues covered in like-minded and cross-cutting sources to vary. All but one existing experiment holds constant which issues are covered across sources, limiting the treatments to clips about issues that are covered on both sources. Because our study exposes current Fox News viewers to actual CNN coverage (and not researcher-selected coverage), the topics covered naturally vary across the networks.

- No existing studies have examined effects among existing partisan media viewers, although many studies have examined effects among those who state a preference for or choose to consume partisan media. Our study is specifically limited to individuals who consume partisan media at baseline.
- Broadly, existing work has argued and found that exposure to cross-cutting media will either be ineffective or actually backfire. Arceneaux and Johnson's (2013) prediction that 'counterattitudinal news can have a polarizing effect' (p. 107) and 'be just as polarizing as exposure to proattitudinal news' (p. 116) finds support in a number of studies. Most studies find that exposure to cross-cutting media either further polarizes viewers or has null effects. Other than this paper, the few studies which report findings consistent with cross-cutting media moderating viewers only reach this finding among subgroups, finding null effects overall.
- No existing studies have measured effects on knowledge of cross-cutting information; authors of existing studies may not have expected information to meaningfully vary across sources when the news outlets were covering the same issues and events.

Overall, these findings underscore our argument that allowing the issue content of cross-cutting media to vary may produce different conclusions than holding it fixed.

There is also a more recent literature on exposure to online partisan media. Most similar to our work, Guess et al. (2021), Levy (2021), and Casas, Menchen-Trevino and Wojcieszak (2022) increase exposure to various online partisan news sources; crucially, participants in these three experiments are exposed to the actual sources as they exist in the real world, not a subset of its content selected by researchers. This means that any differences in what topics or information the networks choose to cover are reflected in these treatments. Their results are broadly consistent with our arguments and findings: Guess et al. (2021) finds effects on knowledge of the information in these sources, Levy (2021) finds that participants' attitudes moderate broadly, and Casas, Menchen-Trevino and Wojcieszak (2022) generally find null effects on attitudinal and affective polarization (i.e., do not find backlash).[1]

[1]See also related work by Bail et al. (2018) and Searles et al. (2022), which do not study partisan *media*. Bail et al. (2018) expose participants to individual tweets researchers chose from elected officials, non-profits, and other sources and find that exposure to these tweets increases polarization. Exposure to individual tweets selected by researchers is distinct from naturalistic exposure to an entire media source. Searles et al. (2022) created a custom online news portal similar to Google News that was randomized to have no Fox News or MSNBC stories (baseline stream), the baseline stream with Fox News stories added, or the baseline stream with MSNBC stories added. Exposure to a news portal is similarly distinct from exposure to the underlying news source. In other related work on online media, Wojcieszak et al. (2021) examines the effect of increasing or decreasing any online news consumption (not necessarily consumption of out-partisan media) and finds null effects on political knowledge, attitude polarization, affective polarization, negative system perceptions, and individual well-being.

Citation	Study Design		Effects of Cross-Cutting Media Exposure on...				
	Treatments Cover Same Issue(s)?	Studies existing partisan media viewers?	Knowledge of cross-cutting information	Moderate issue views	Evaluations of out-party	Evaluations of in-party	Issue importance consistent with content
Prediction: Cross-Cutting Media Polarizes			↓	↓	↓	↑	?
Prediction: Cross-Cutting Media Moderates			↑	↑	↑	↓	?
This Paper	No	Yes	↑	↑	↑	↓	↑
Arceneaux and Johnson (2013, Spring 2009 Selective)	Yes	No*			0		
Arceneaux and Johnson (2013, Fall 2009 Selective)	Yes	No*		↓			↑
Arceneaux and Johnson (2013, Winter 2010 Selective)	No	No*			0		↑
Arceneaux and Johnson (2013, Fall 2011 Selective)	Yes	No*		↓	↓	↑	0
Arceneaux and Johnson (2013, Summer 2010 Preference)	Yes	No*					0
Arceneaux and Johnson (2013, Fall 2011 Preference)	Yes	No*		↓			↑
Druckman, Levendusky and McLain (2018)	Yes	No		↓			
Feldman (2011)	Yes	No		0/↑			
Levendusky (2013*a*, Experiment 2)	Yes	No		0/↑			
Levendusky (2013*a*, Experiment 8)	Yes	No		0			
Levendusky (2013*b*, Experiment 1)	Yes	No			0	0	
Levendusky (2013*b*, Experiment 2)	n/a	No				0/↓	
Levendusky (2013*c*, Experiment 1)	Yes	No		0/↓			
Levendusky (2013*c*, Experiment 2)	Yes	No*		0/↑			
Levendusky (2013*c*, Experiment 3)	Yes	No		0			

*Notes: ↑ = statistically significant positive effect reported. ↓ = statistically significant negative effect reported. 0 = null result reported. 0/↓ = null result reported on average but significant effects reported in a subsample. * = examines effects among those who select or state a preference for consuming like-minded media. The Table does not list Experiments 1, 3, 4, 6 and 7 in Levendusky (2013a) because these experiments are also reported in Levendusky (2013c,b), which are reported in the table.*

Table OA1: Existing Studies of Exposure to Cross-Cutting Televised Media Which Measure Eligible Outcomes

Theoretical Construct	**Media's Action**	**Effect on Viewers**	**Hypothetical Example**
Agenda Setting	A network **covers a topic more**, holding constant the information conveyed about that topic (e.g., McCombs and Shaw, 1972; Iyengar and Kinder, 1987; Krosnick and Kinder, 1990).	This leads viewers to **see this topic as *more important*** and to ***priming* viewers' *pre-existing* attitudes on this topic** when forming political evaluations.	A Republican President launches a new military conflict. Media outlets cover the new conflict every day, leading viewers to see the conflict as important and to base their evaluations of the President on how they think she is handling the conflict.
Framing	A network "provides an interpretation of an issue or policy by *emphasizing* which aspect of the issue is relevant for evaluating it, *without the frame itself [providing] any new substantive information about the issue*" (Leeper and Slothuus, 2020, p.154, emphasis in original)	This leads viewers to think about the issue in a different way, **changing which considerations are salient** to them.	**CNN** refers to local militias fighting back against the US as "freedom fighters," while **Fox News** refers to them as "terrorists." **CNN** refers to civilian casualties as "deaths of unarmed women and children," while **Fox News** refers to them as "collateral damage." When thinking about the conflict, viewers then bring to mind related considerations (e.g., the need to fight terrorists), affecting levels of support for the conflict.
Partisan Coverage Filtering	A network **is more likely to cover information favorable to its partisan or ideological side** and less likely to cover information unfavorable to its partisan or ideological side (e.g., Hayakawa, 1940; Mullainathan and Shleifer, 2005; Besley and Prat, 2006; Gentzkow and Shapiro, 2006).	This leads viewers to ***learn* more information favorable to the network's partisan side**, which could ***change* viewers' attitudes and political evaluations**.	**CNN** gives extensive information about the cost of the conflict, the number of US soldiers who died, and civilian casualties. **Fox News** gives equally extensive information about the severity of the threat that the President's military campaign neutralized and anecdotes of civilians who have greeted US soldiers as liberators. This leaves viewers of each network with different factual understandings of the conflict, and subsequently different levels of support for the conflict and the President.

Table OA2: Overview of Theoretical Constructs

2 Pre-Registration

This experiment was pre-registered at https://osf.io/9fdq2/?view_only=23d11559495b48269e18523e4870ac4f for the midline and https://osf.io/chsf2/?view_only=4ca51b40718944dbaa35825e8b38cb62 for the endline.

3 CONSORT-Style Diagram

Figure OA1 provides an overview of the experimental design.

Figure OA1: CONSORT-Style Diagram

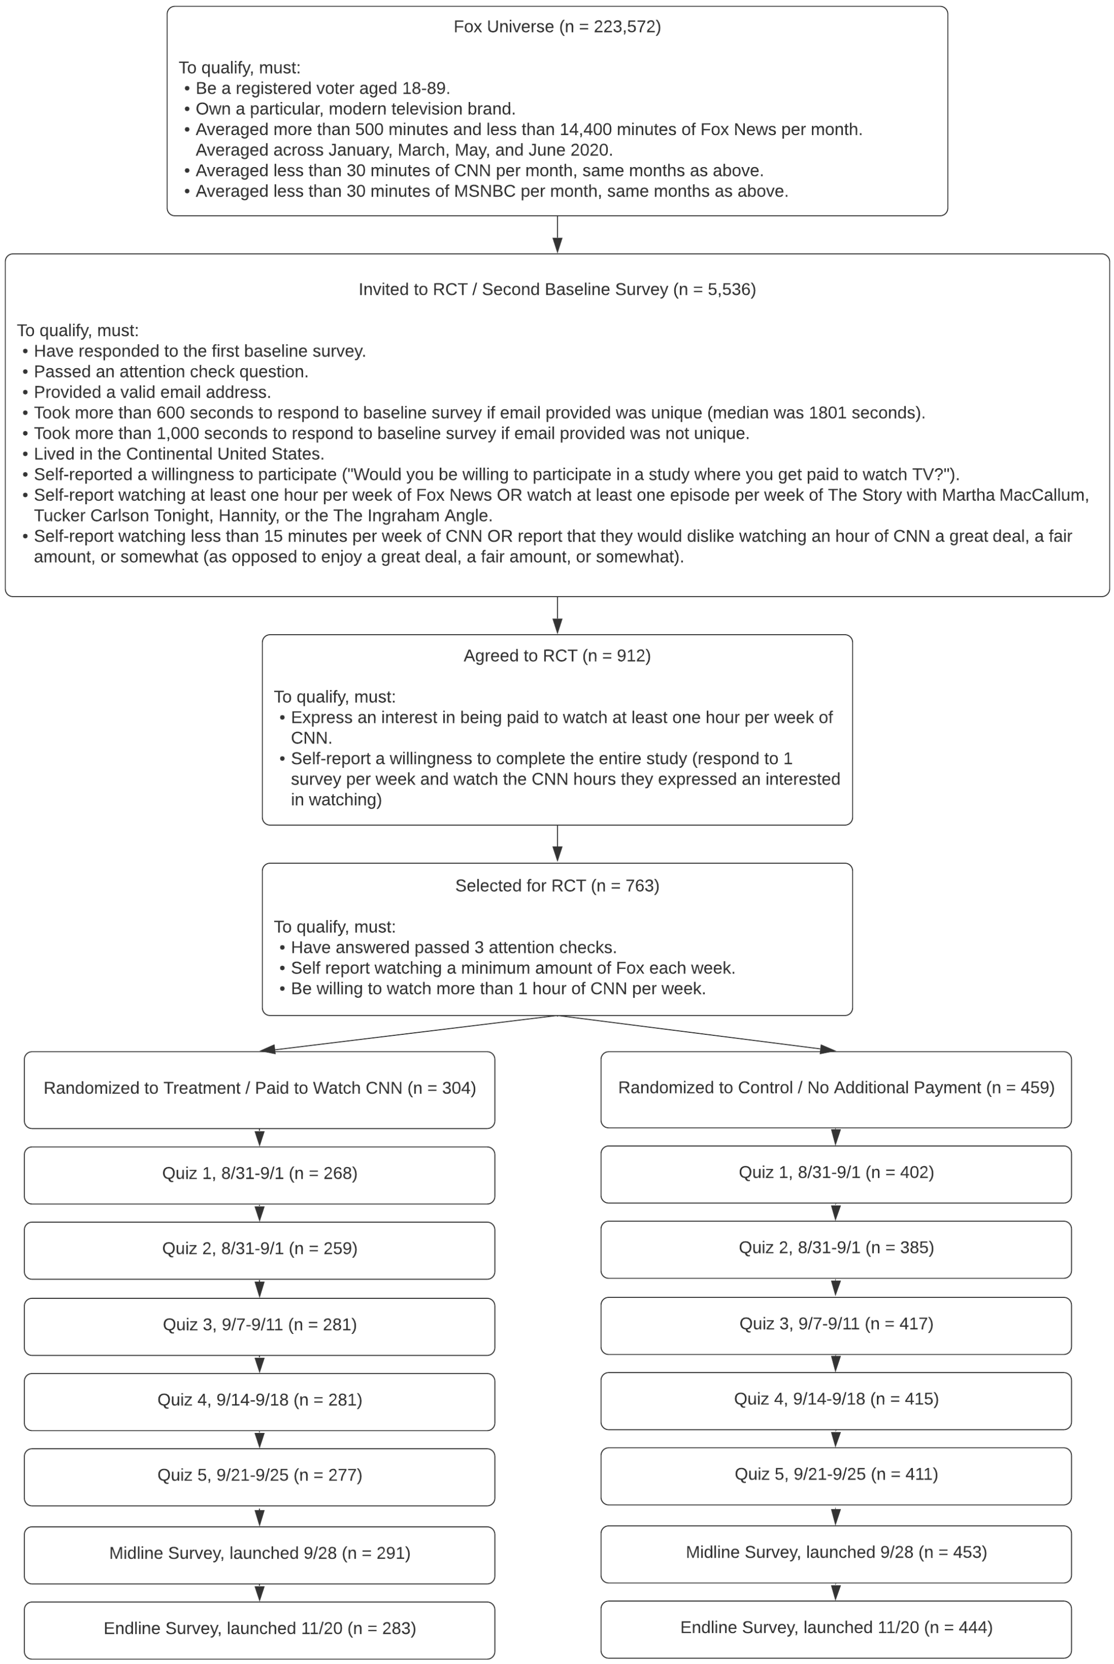

4 Sample Demographics

Table OA3 summarizes the sample demographics at each stage.

Sample	Mailed Invitation	Baseline Survey Respondent	Invited to RCT	Agreed to RCT	Selected for RCT	Midline Survey Respondent	Endline Survey Respondent
Count	223572	14849	5536	912	763	744	727
Age	53.86 (56)	56.47 (59)	58.6 (61)	53.05 (55)	53.51 (56)	53.37 (55)	53.46 (55)
Voted 16G (0/1)	0.8 (1)	0.88 (1)	0.91 (1)	0.89 (1)	0.9 (1)	0.9 (1)	0.89 (1)
Voted 16P (0/1)	0.19 (0)	0.24 (0)	0.27 (0)	0.26 (0)	0.28 (0)	0.28 (0)	0.28 (0)
Reg. GOP (0/1)	0.69 (1)	0.67 (1)	0.82 (1)	0.8 (1)	0.82 (1)	0.82 (1)	0.82 (1)
Reg. Dem. (0/1)	0.18 (0)	0.21 (0)	0.08 (0)	0.1 (0)	0.09 (0)	0.08 (0)	0.09 (0)
White (0/1)	0.91 (1)	0.93 (1)	0.95 (1)	0.95 (1)	0.95 (1)	0.95 (1)	0.95 (1)
Female (0/1)	0.53 (1)	0.52 (1)	0.5 (0)	0.51 (1)	0.51 (1)	0.51 (1)	0.51 (1)
Self-Reported Weekly Fox Min.	NA (NA)	483.58 (150)	856.14 (630)	770.77 (600)	839.61 (630)	841.67 (630)	842.91 (630)
Trump Thermometer (0-100)	NA (NA)	59.32 (70)	83.02 (90)	79.44 (87)	82.54 (90)	82.35 (90)	82.17 (90)
Fox News Thermometer (0-100)	NA (NA)	54.28 (51)	74.09 (79)	72.54 (75)	74.76 (79)	74.86 (79)	74.57 (78)
CNN Thermometer (0-100)	NA (NA)	30.88 (25)	12.15 (1)	13.06 (4)	11.74 (1)	11.65 (1)	11.73 (1)
Ideology (1-9)	NA (NA)	5.98 (6)	7.15 (7)	7.03 (7)	7.16 (7)	7.16 (7)	7.15 (7)
Party Identification (1-7)	NA (NA)	3.02 (3)	1.94 (1)	1.97 (1)	1.85 (1)	1.85 (1)	1.86 (1)
Percent Political Knowledge Qs Correct (0-1)	NA (NA)	0.69 (0.75)	0.74 (0.75)	0.74 (0.75)	0.76 (0.75)	0.76 (1)	0.76 (0.75)
Education (1-7)	NA (NA)	5.17 (6)	5.01 (6)	5.25 (6)	5.22 (6)	5.25 (6)	5.26 (6)
Income over $150k/year	NA (NA)	0.16 (0)	0.15 (0)	0.19 (0)	0.19 (0)	0.19 (0)	0.19 (0)
Would enjoy watching CNN (1-6)	NA (NA)	2.41 (2)	1.6 (1)	1.69 (1)	1.64 (1)	1.64 (1)	1.64 (1)
Would enjoy discussing politics with CNN viewer (1-6)	NA (NA)	2.71 (3)	2.25 (2)	2.41 (2)	2.41 (2)	2.41 (2)	2.42 (2)

Note:
Means are printed first with medians in parantheses. Data labelled as NA were not collected and are unavailable. All other data was measured pre-treatment. Ideology is coded as 1 = Extremely Liberal; 9 = Extremely Conservative. Party identification is coded as 1 = Strong Republican; 7 = Strong Democrat. Education is coded as 1 = no high school; 2 = high school; 3 = associates degree; 4 = vocational degree; 5 = some college; 6 = bachelors degree; 7 = graduate degree. Four political knowledge questions were asked. We report the percent correct across all four. These questions measured House control, John Robert's job title, the issue covered by Dodd-Frank, and the number of years in a Senate term. The 'Would enjoy' items are six point scales ranging from dislike a great deal (1) to enjoy a great deal (6).

Table OA3: Sample Characteristics at Each Stage

5 Covariate Balance at Each Stage

The below tables demonstrate that balance on pre-treatment observable attributes is maintained among the original universe of pre-survey respondents randomized to each group, the sub-sample that responded to the midline survey and the sub-sample that responded to the endline survey. Each table shows the mean value for the covariate (measured in the baseline survey before treatment) under each condition as well as the *p*-value from a one-way ANOVA test. The first table considers all voters who were randomly assigned after having taken the pre-surveys; the second table considers all voters who responded to the midline survey; the third table considers all voters who responded to endline survey.

	Control (No Incentive)	Treatment (CNN Incentive)	p-value
Baseline Partisanship Factor	0.01	-0.02	0.6
Baseline CNN Factor	0.01	-0.02	0.61
Baseline Fox Factor	-0.01	0.01	0.76
Number of Survey Respondents in Household	1.18	1.17	0.76
Number of Hours Available for Incentivizing	5.85	5.85	0.99
Age	53.21	53.95	0.51
Voted 2016 General (0/1)	0.91	0.88	0.34
Voted 2016 Primary (0/1)	0.27	0.30	0.36
Registered Democrat (0/1)	0.08	0.09	0.85
White (0/1)	0.95	0.95	0.7
Female (0/1)	0.51	0.50	0.81
Trump Thermometer (0-100) at Baseline	81.89	83.52	0.3
Ideological Self-Placement (1-9) at Baseline	7.15	7.17	0.85
Education (1-7) at Baseline	5.29	5.12	0.17
Party Identification (1-7) at Baseline	1.84	1.88	0.6
Reinterview Rate from Baseline Survey	100.00	100.00	-
N	459.00	304.00	-

Table OA4: Covariate Balance among Pre-Survey Respondents.

	Control (No Incentive)	Treatment (CNN Incentive)	p-value
Baseline Partisanship Factor	0.02	-0.01	0.68
Baseline CNN Factor	0.01	-0.02	0.6
Baseline Fox Factor	0.00	0.02	0.77
Number of Survey Respondents in Household	1.18	1.18	0.9
Number of Hours Available for Incentivizing	5.86	5.87	0.92
Age	53.20	53.65	0.69
Voted 2016 General (0/1)	0.91	0.88	0.34
Voted 2016 Primary (0/1)	0.26	0.30	0.36
Registered Democrat (0/1)	0.08	0.09	0.84
White (0/1)	0.95	0.96	0.61
Female (0/1)	0.51	0.51	0.86
Trump Thermometer (0-100) at Baseline	81.78	83.23	0.36
Ideological Self-Placement (1-9) at Baseline	7.16	7.17	0.93
Education (1-7) at Baseline	5.29	5.18	0.35
Party Identification (1-7) at Baseline	1.84	1.88	0.6
Reinterview Rate from Baseline Survey	98.69	95.72	-
N	453.00	291.00	-

Table OA5: Covariate Balance among Midline Survey Respondents

	Control (No Incentive)	Treatment (CNN Incentive)	p-value
Baseline Partisanship Factor	0.02	0.00	0.77
Baseline CNN Factor	0.01	-0.01	0.7
Baseline Fox Factor	0.00	0.01	0.79
Number of Survey Respondents in Household	1.18	1.18	0.94
Number of Hours Available for Incentivizing	5.88	5.88	0.95
Age	53.23	53.83	0.61
Voted 2016 General (0/1)	0.90	0.88	0.32
Voted 2016 Primary (0/1)	0.27	0.30	0.44
Registered Democrat (0/1)	0.08	0.09	0.69
White (0/1)	0.95	0.95	0.89
Female (0/1)	0.50	0.52	0.76
Trump Thermometer (0-100) at Baseline	81.53	83.19	0.3
Ideological Self-Placement (1-9) at Baseline	7.15	7.14	0.95
Education (1-7) at Baseline	5.30	5.21	0.5
Party Identification (1-7) at Baseline	1.84	1.89	0.55
Reinterview Rate from Baseline Survey	96.73	93.09	-
N	444.00	283.00	-

Table OA6: Covariate Balance among Endline Survey Respondents

6 Test of Differential Attrition and Differential Attrtition by Pre-Treatment Covariates

The above subsection demonstrated that at each stage, there was covariate balance. We next examine whether there is evidence of differential attrition.

First, we do find evidence of a small amount of average differential attrition: the control group response rate to the midline survey was 98.7% while the treatment group response rate was 95.7%, a difference of 3.0 percentage points ($p = 0.02$). This difference is substantively small, however. To test whether attrition patterns are similar by covariates in treatment and control, we use a linear regression of whether or not an individual responded to each follow-up survey on treatment, baseline covariates used in blocking, and treatment-covariate interactions. We then perform a heteroskedasticity-robust F-test of the hypothesis that all the interaction coefficients are zero. This procedure was pre-specified in our pre-analysis plan and is standard practice Gerber and Green (2012). Below we report the p-value of this F-test. Based on the results presented in the Table below, there does not appear to be evidence of asymmetrical attrition.

Midline Survey	0.77
Endline Survey	0.81

Table OA7: p-value by Survey Wave Test of Differential Attrition by Covariates Used in Blocking.

7 Description of Estimation and Inference Procedure

In order to improve the precision of our treatment effect estimates, we use regression with covariate adjustment. To select the covariates we include, we employ a machine learning approach to model building (Belloni, Chernozhukov and Hansen, 2014; Bloniarz et al., 2016). Specifically, following our pre-analysis plan, to estimate all treatment effects, for each outcome variable, we first used a 20-fold elastic net regression with this outcome and our pre-registered, pre-treatment covariates. These covariates were:

- The log of the pre-treatment minutes watched of CNN.

- The log of the pre-treatment minutes watched of Fox News
- The log of the pre-treatment minutes watched of total TV.
- A baseline partisanship factor.
- A baseline CNN factor.
- A baseline factor on Fox watching behavior.
- Household size.
- The number of hours each person was eligible to be incentivized.
- All variables on the pre-treatment t0 and t1 surveys.

This first regression does not include the treatment indicator. We then extract the variables with non-zero coefficients from this eleastic net regression and use them in a second regression. This second regression regresses the outcome variable on the treatment indicator and the pre-treatment covariates extracted from the elastic net regression. The coefficient on the treatment indicator is the covariate-adjusted treatment effect. This second regression is also clustered at the household level for our standard errors.

This procedure is described in Belloni, Chernozhukov and Hansen (2014) and Bloniarz et al. (2016). For a similar application in political science, see Fang, Guess and Humphreys (2019).

We pre-registered that we would report three types of p-values. These different p-values use different approaches to adjust for multiple comparisons:

- Family-wise error rate (FWER) adjusted p-values: This uses the `wyoung` command in Stata Jones, Molitor and Reif (2019). This command controls the family-wise error rate (FWER) when performing multiple hypothesis tests using the free step-down resampling method of Westfall and Young (1993). We pre-registered which outcome variable belongs to which family. We run the `wyoung` command once per family. We do not compute FWER-adjusted p-values for indices themselves. The FWER is the probability of making any Type 1 errors at all. This analysis is very computationally intensive and required the use of a cluster over the course of multiple days in order to run. Each iteration of `wyoung` had 10,000 bootstraps.
- False discovery rate (FDR) sharpened q-values: This uses Anderson (2008)'s code, following the approaches used in Chen and Yang (2019) and Allcott et al. (2020). The FDR is the expected proportion of false rejections out of all rejections.
- Conventional, unadjusted p-values: These are the p-values from the second regression described above. Reporting these p-values follows the convention used by Jones, Molitor and Reif (2019) to report both adjusted and unadjusted p-values.

In order to simplify the main text, we do not report FWER adjusted p-values there. Given the large number of hypotheses tested, these p-values are often too conservative Michler and Josephson (2022).

Effect estimates, standard errors, and all three types of p-values for all the individual items in the quiz, midline, and endline surveys, as well as the text of the items, can be found here: https://osf.io/dqv96/?view_only=796462de006b4e0d86b389089aaf63b7.

8 Full Results

This section reports the effects on pre-registered indices and their constituent items. For all Figures in this section, standard errors (thick lines) and 95% confidence intervals (thin) surround point estimates.

8.1 Quiz Figures

Figure OA2: Treatment Effect on Current Event Perceptions During Incentivized Period (Individual Items)

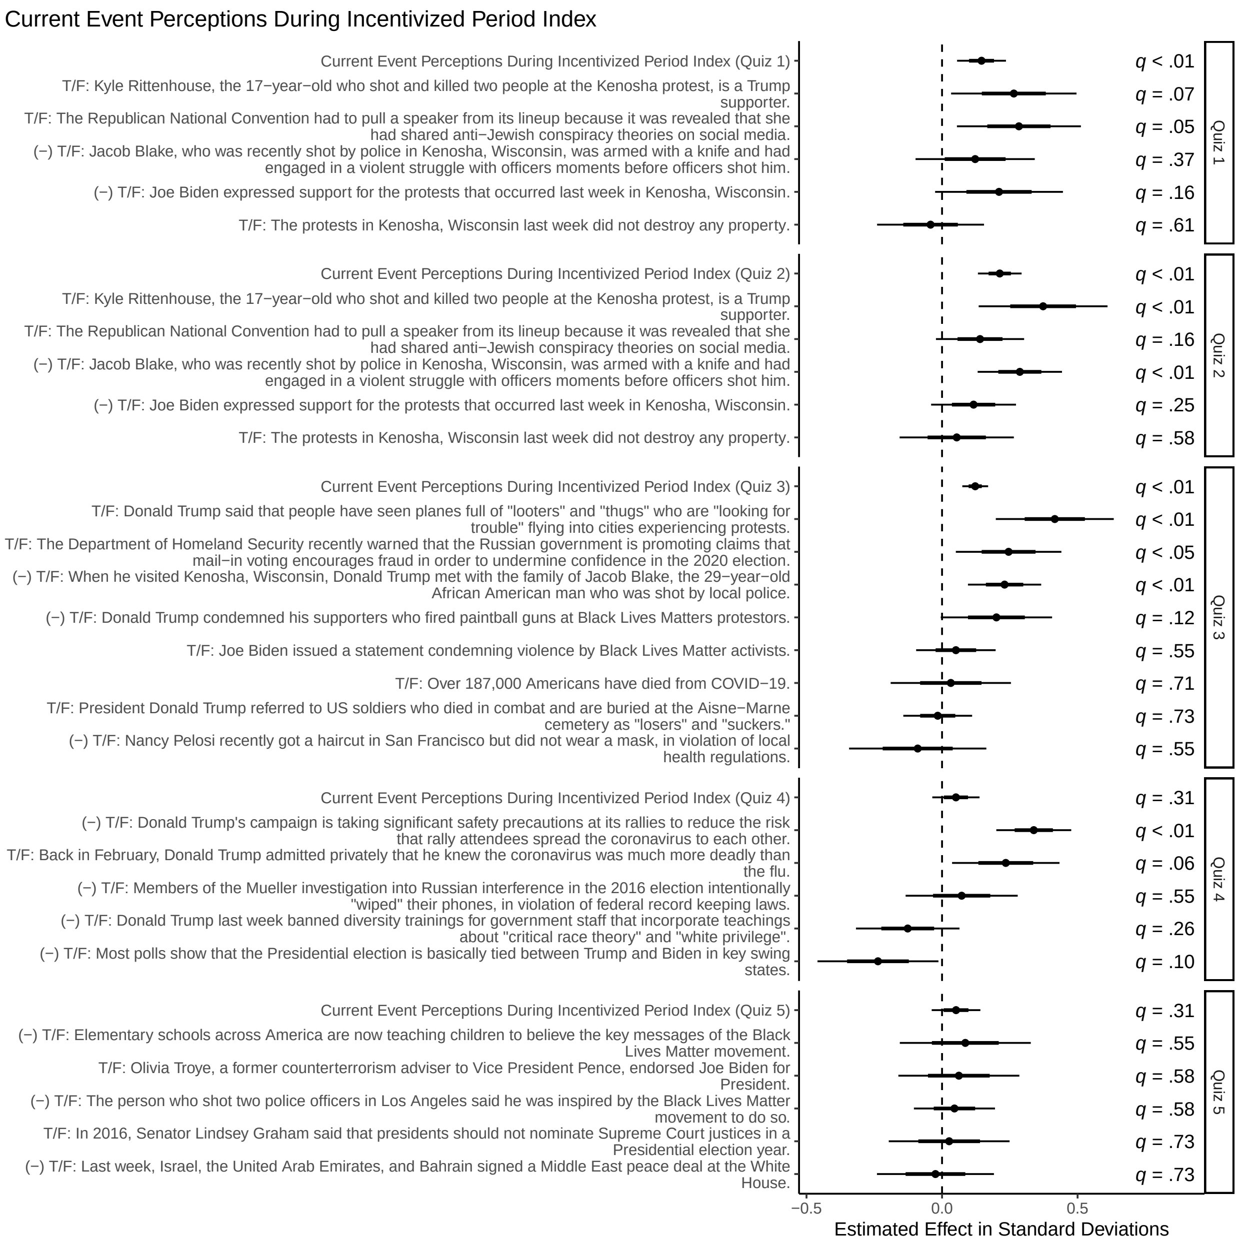

Figure OA3: Treatment Effect on Attitudes Towards Events During Incentivized Period (Individual Items)

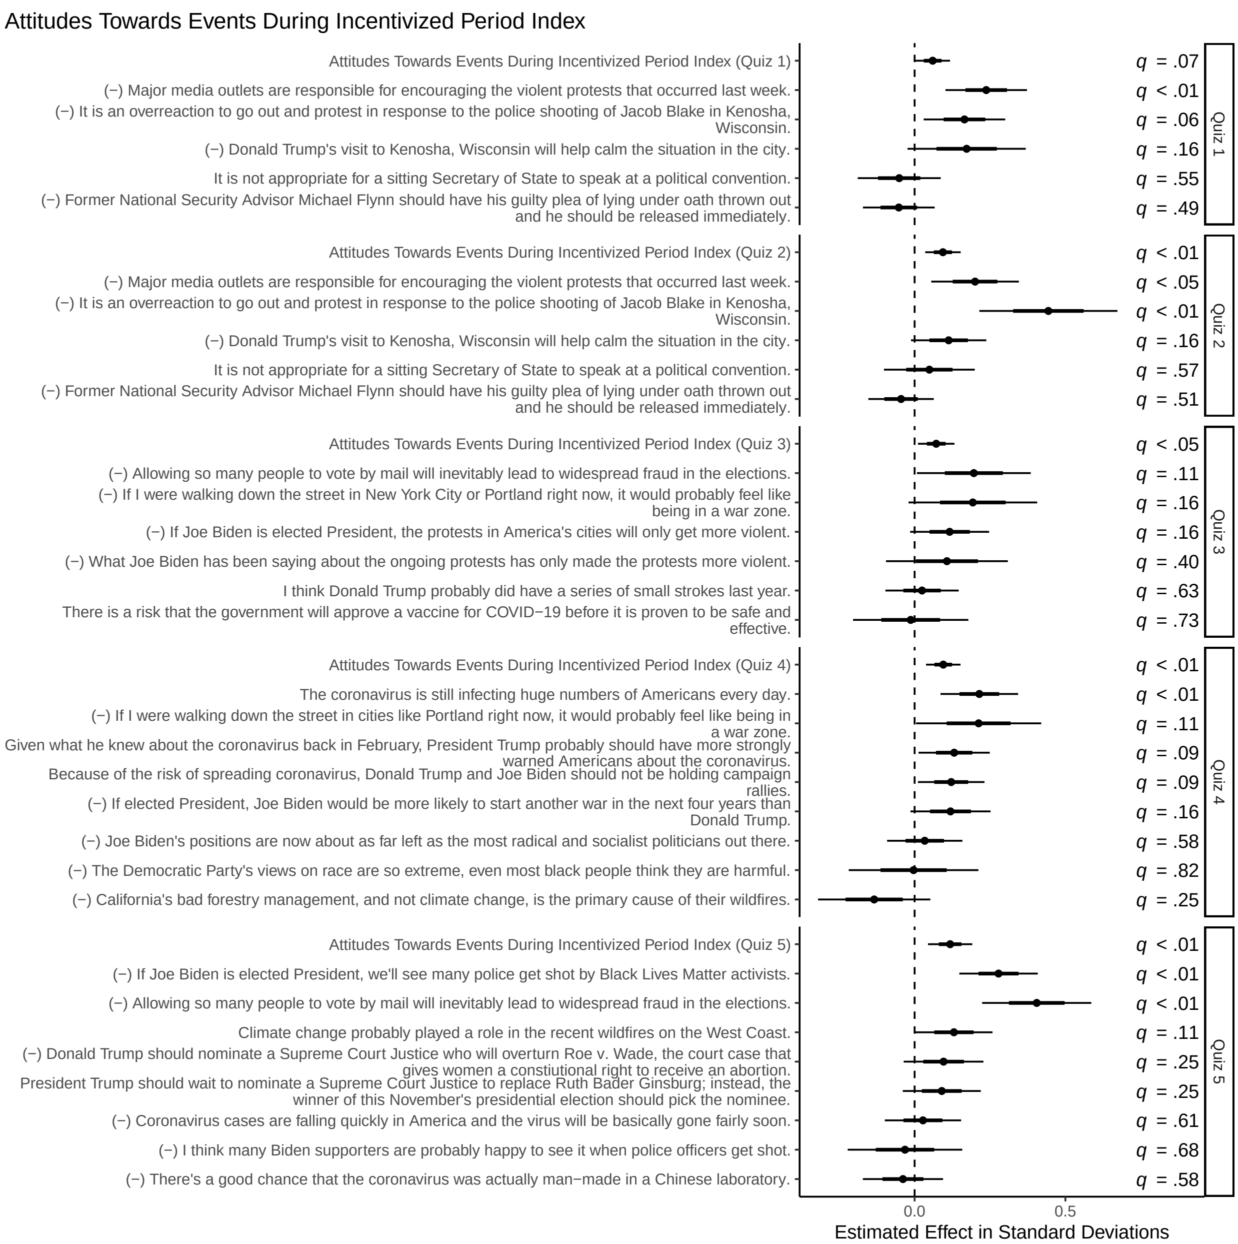

Figure OA4: Treatment Effect on Indices During Incentivized Period (Pooled Across Quizzes)

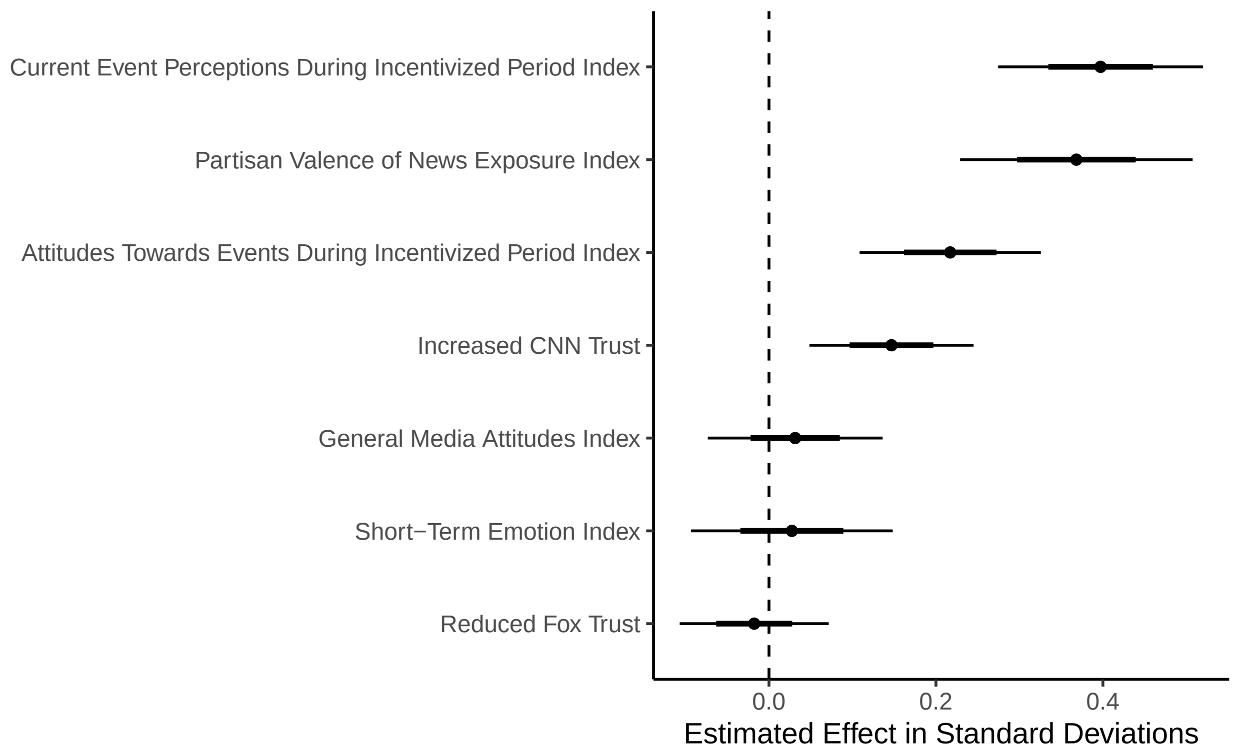

8.2 Midline Figures

Figure OA5: Effects on Pre-Registered Indices (Midline)

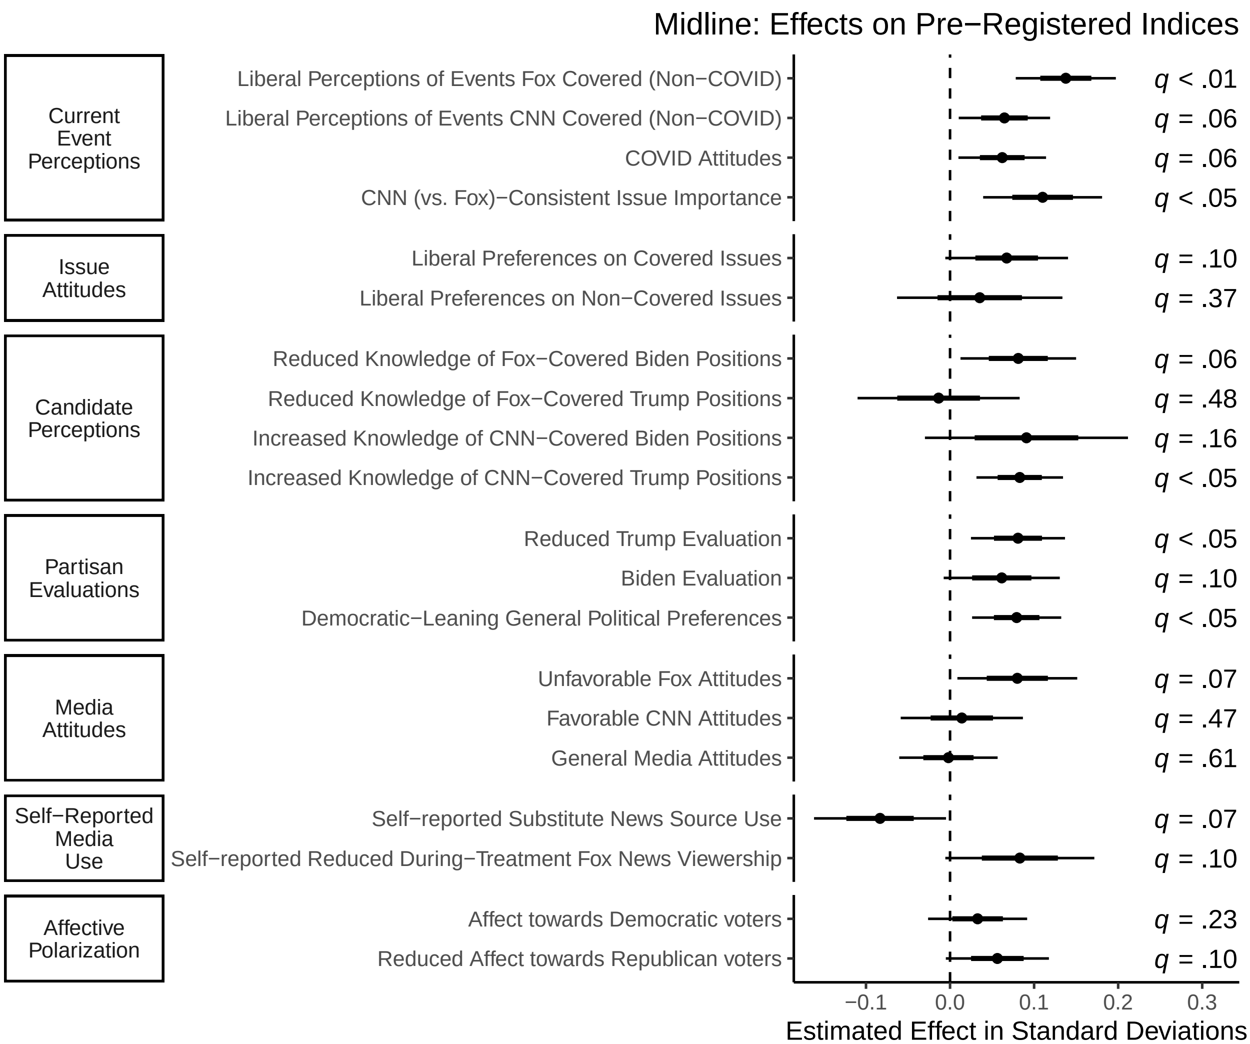

Figure OA6: Effects on Issue Importance and Agenda Setting

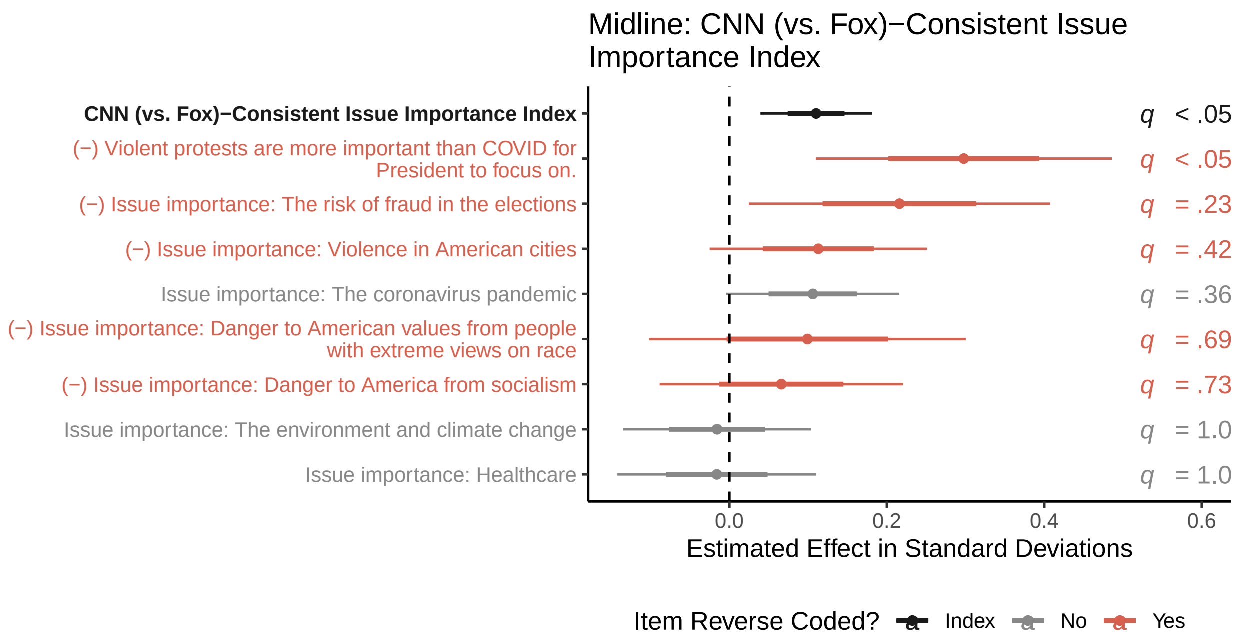

Figure OA7: Liberal Perceptions of Events CNN Covered (Non-COVID) Index

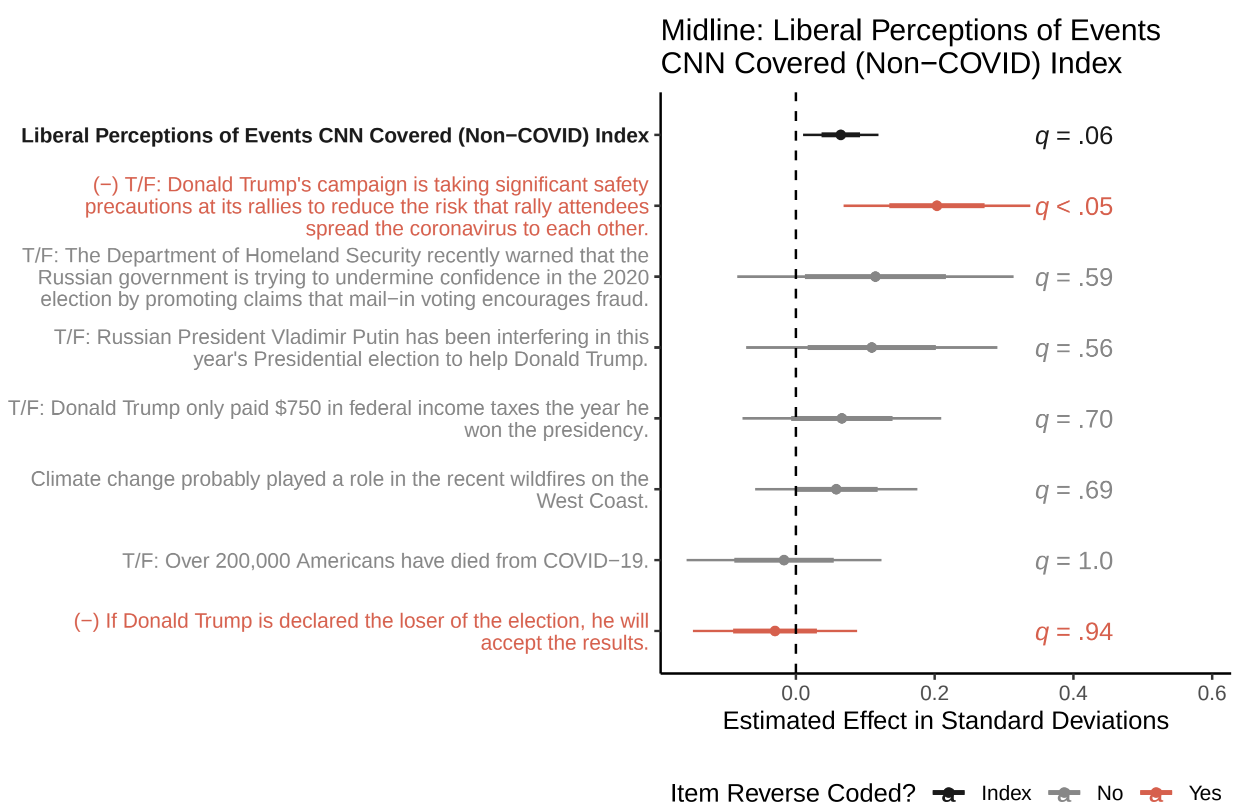

Figure OA8: Liberal Perceptions of Events Fox Covered (Non-COVID) Index

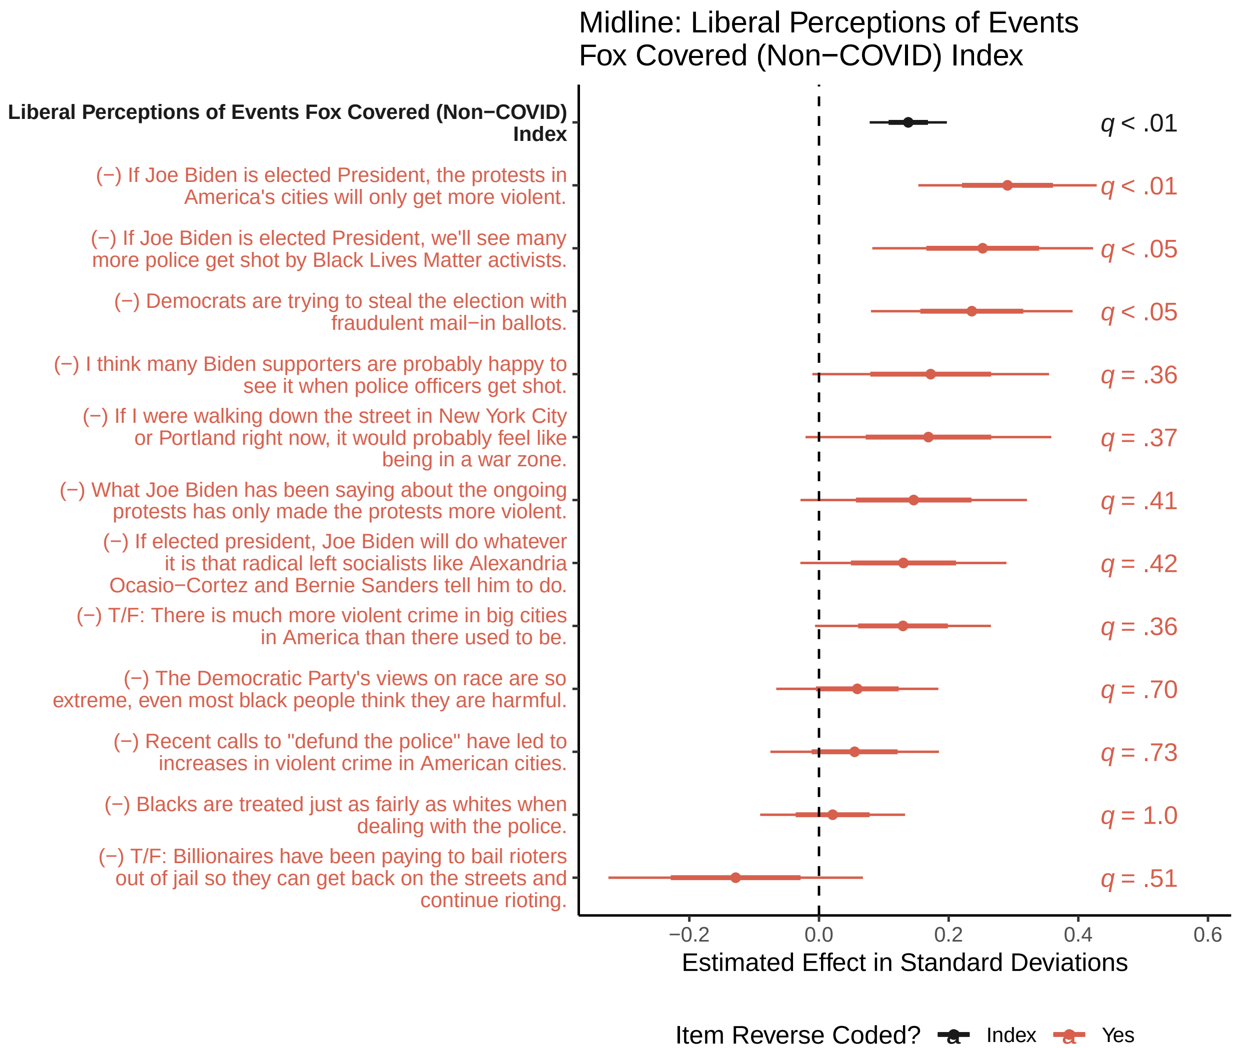

Figure OA9: COVID Attitudes Index

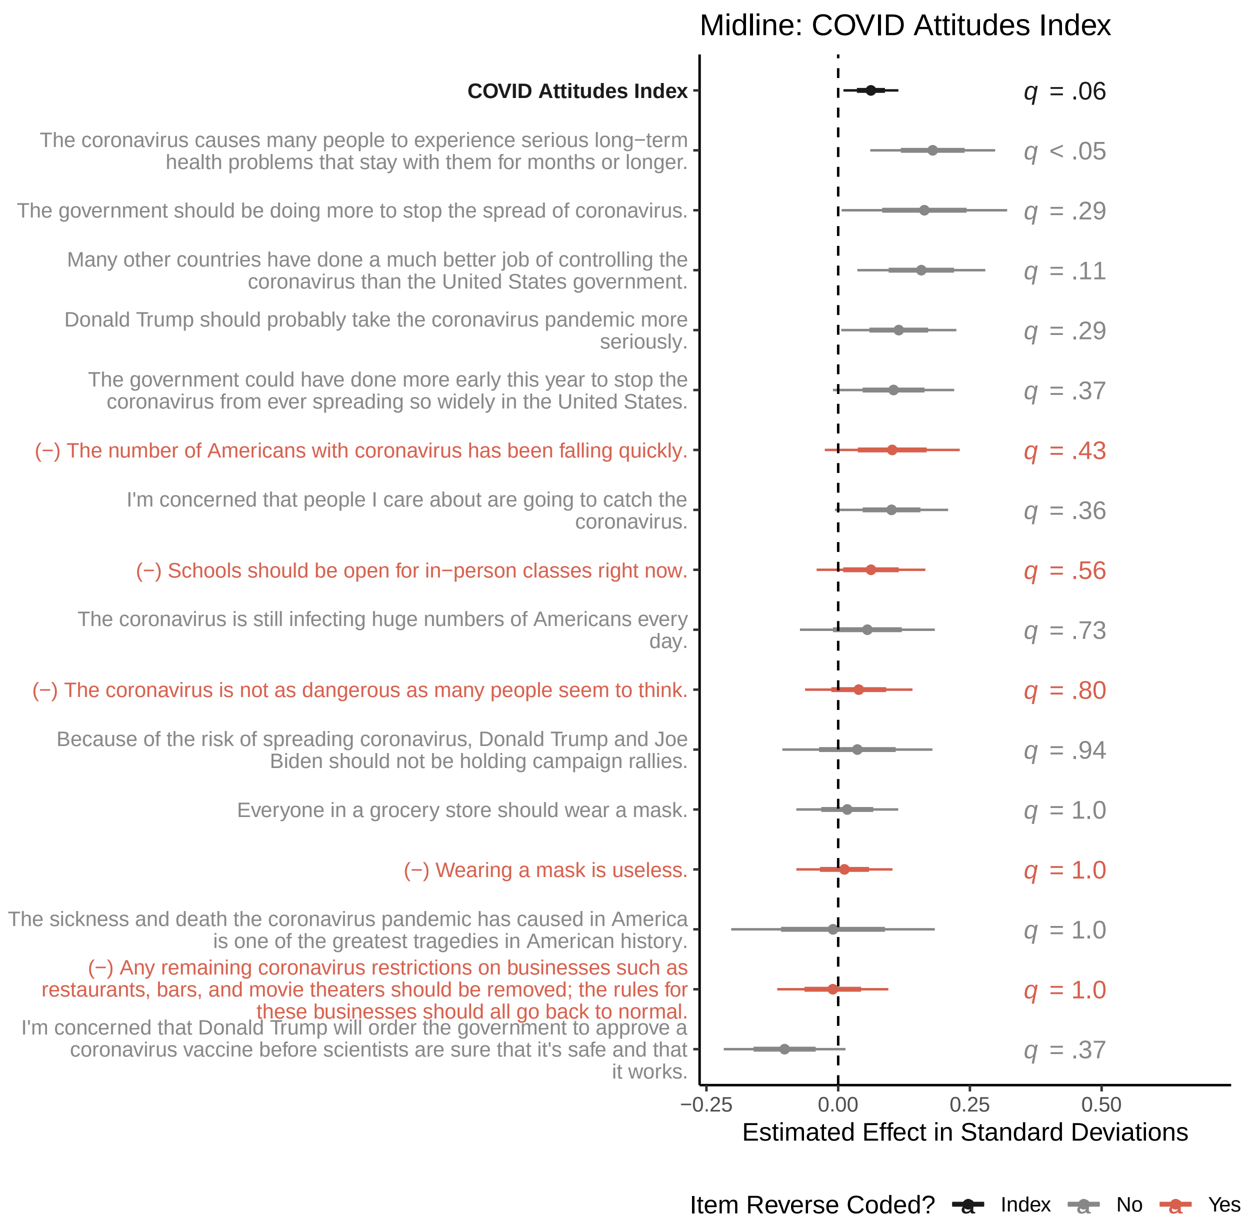

Figure OA10: Reduced Knowledge of Fox-Covered Biden Positions Index

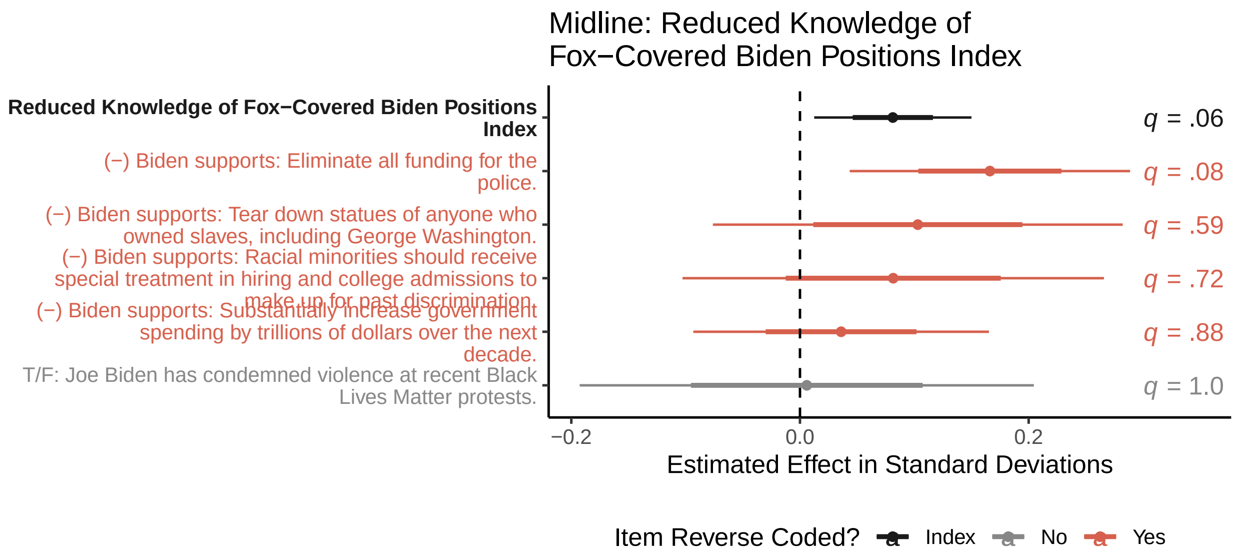

Figure OA11: Reduced Knowledge of Fox-Covered Trump Positions Index

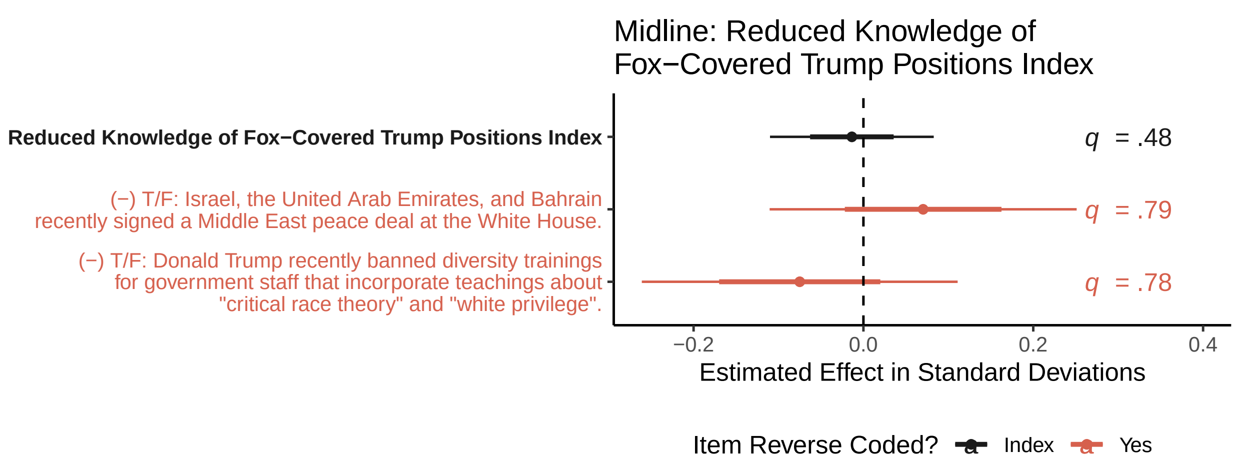

Figure OA12: Increased Knowledge of CNN-Covered Biden Positions Index

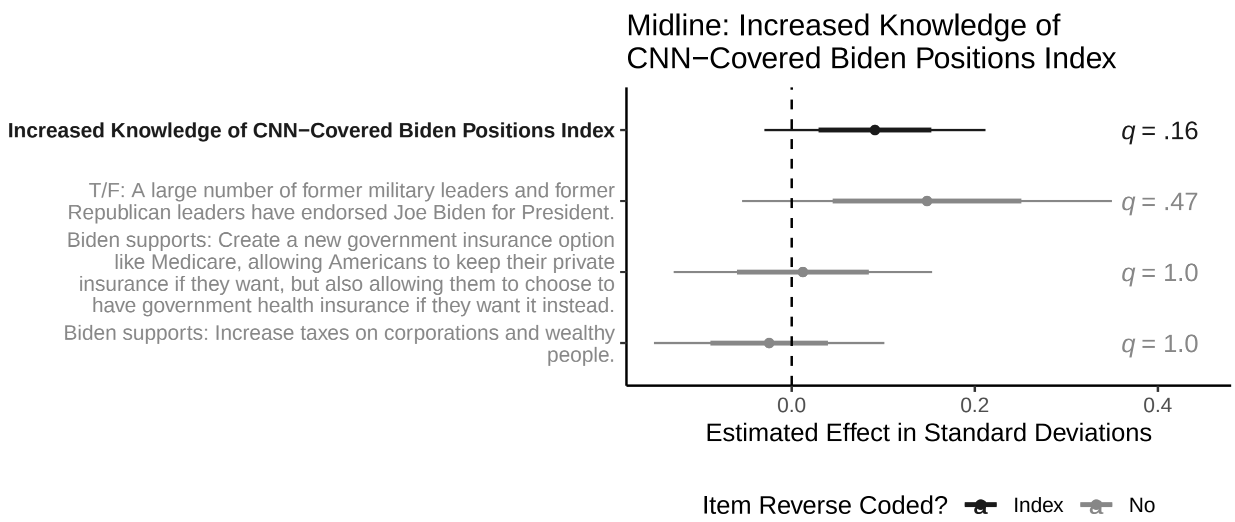

Figure OA13: Increased Knowledge of CNN-Covered Trump Positions Index

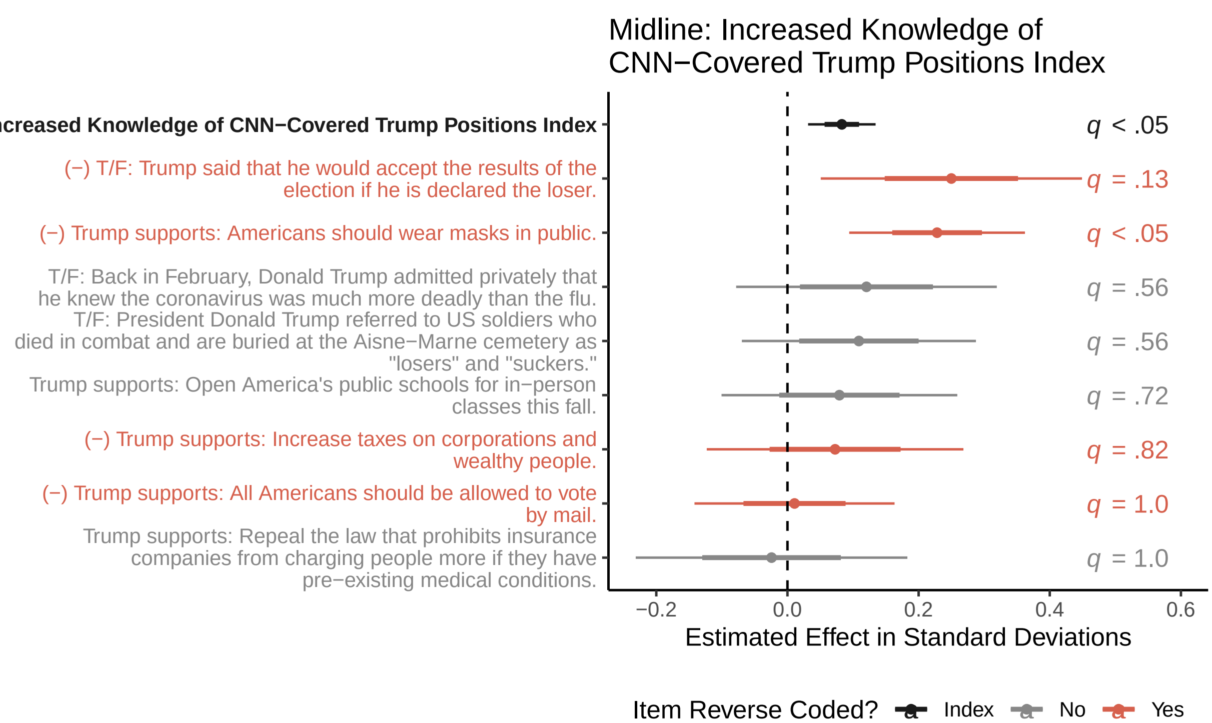

Figure OA14: Liberal Preferences on Covered Issues Index

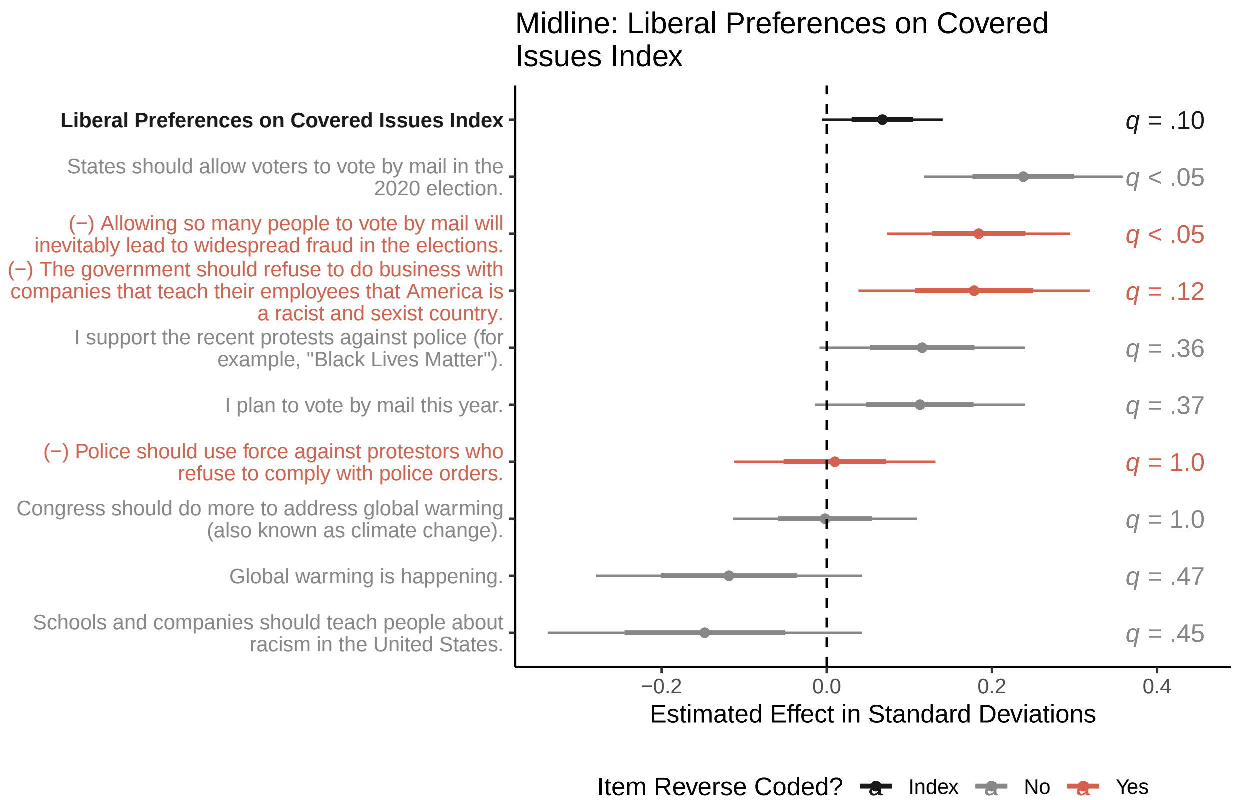

Figure OA15: Reduced Ethnic Antagonism Index

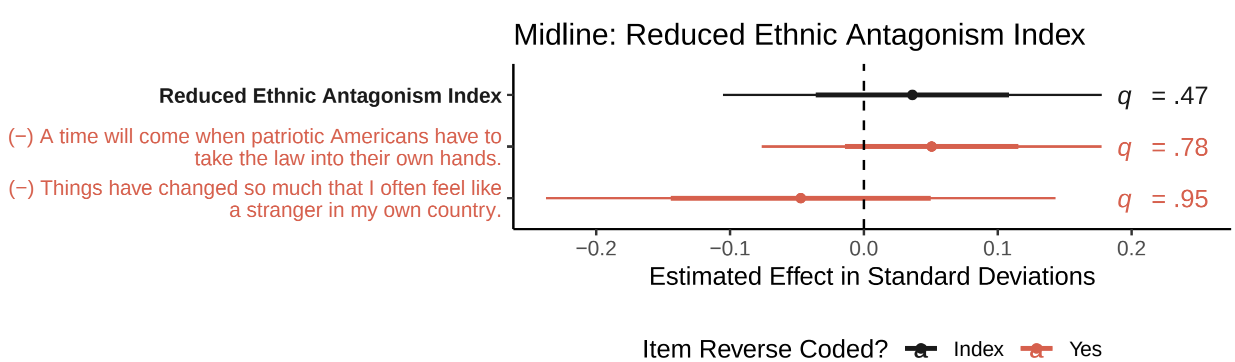

Figure OA16: Reduced Racial Prejudice Index

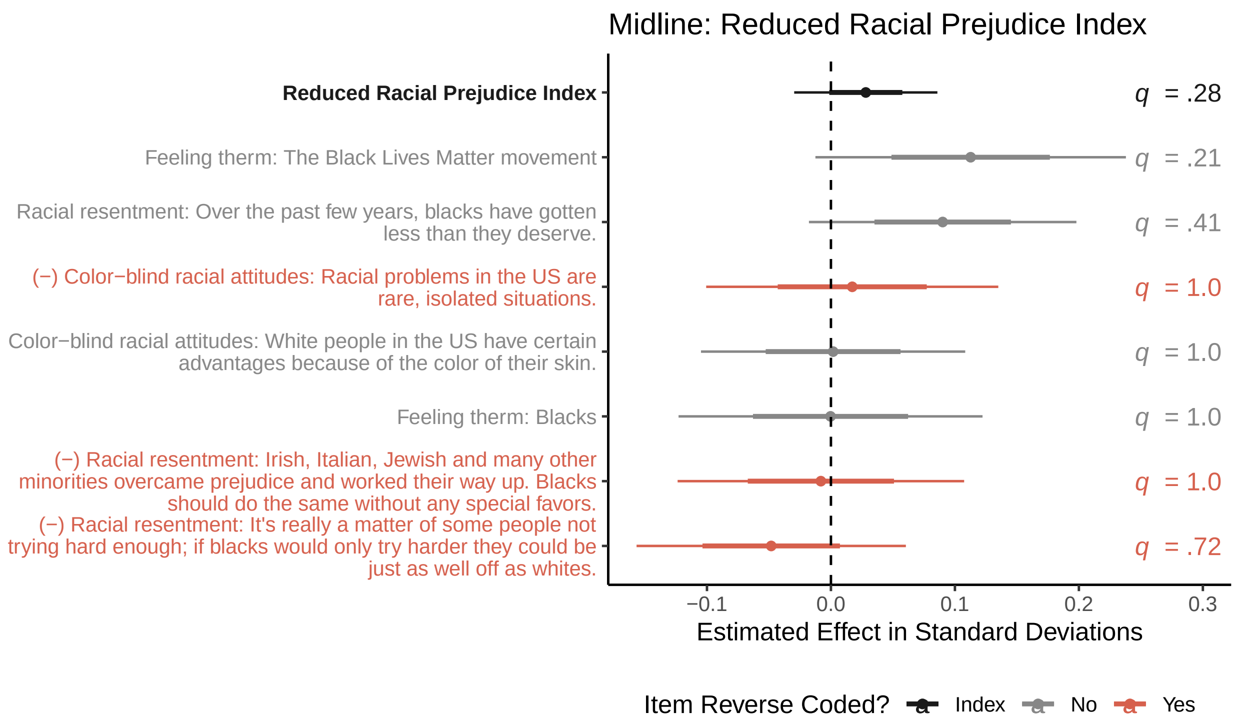

Figure OA17: Democratic-Leaning General Political Preferences Index

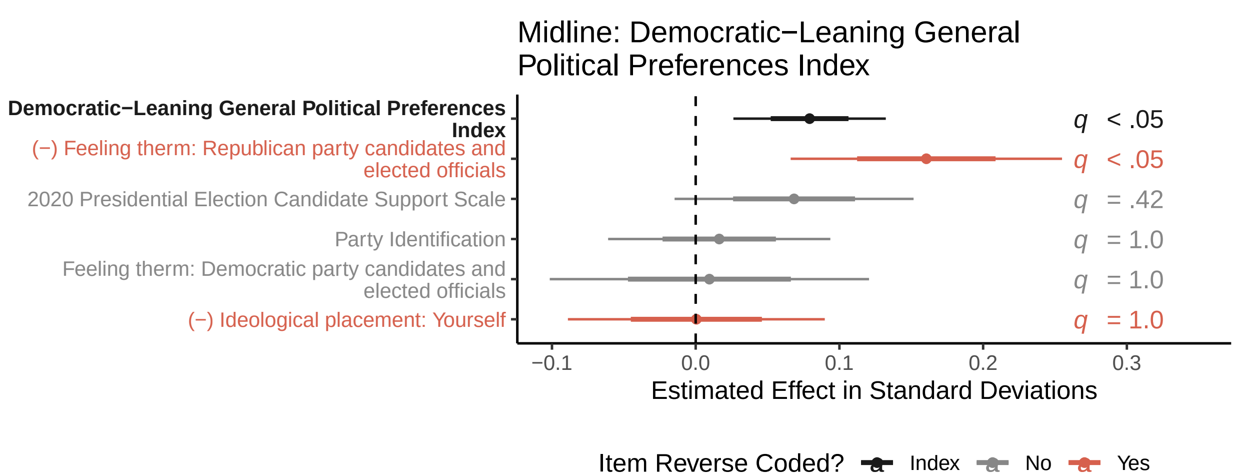

Figure OA18: Biden Evaluation Index

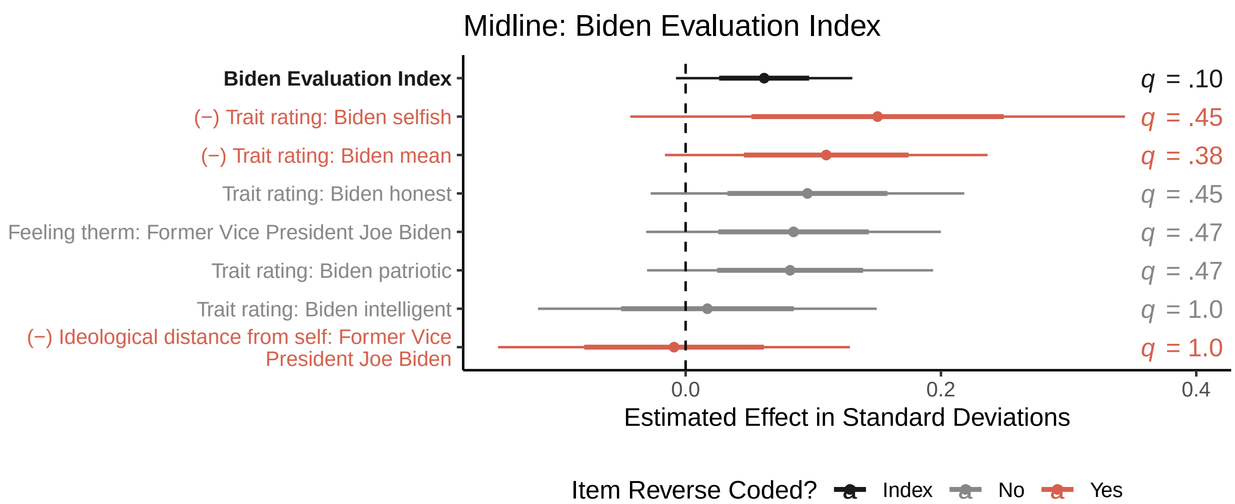

Figure OA19: Favorable CNN Attitudes Index

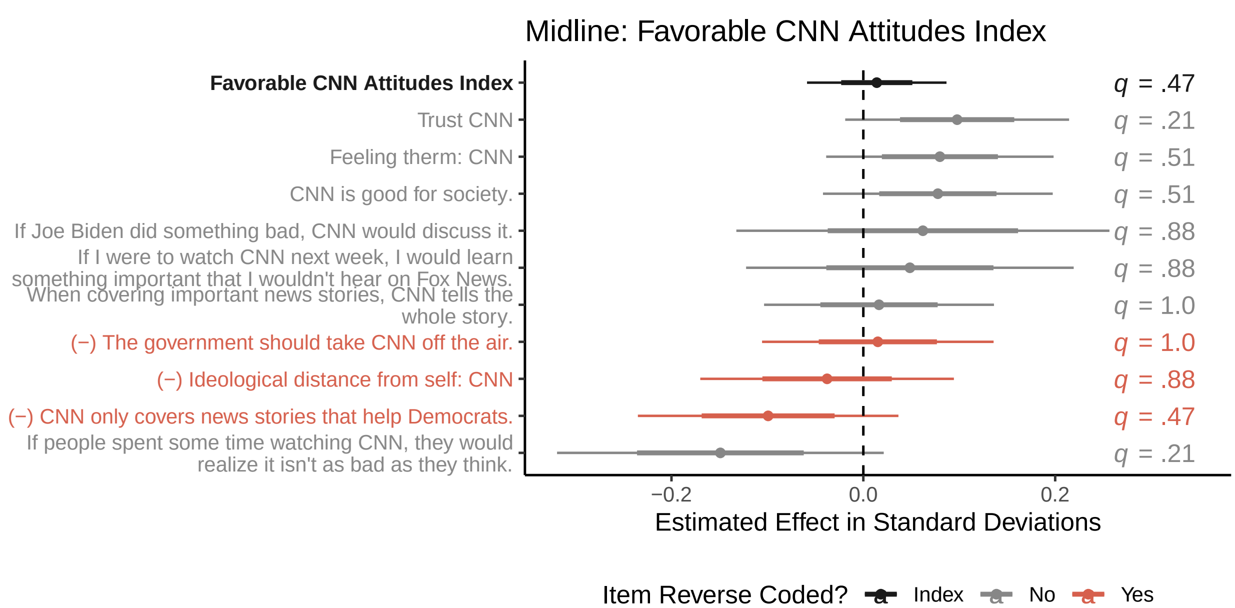

Figure OA20: Affect Towards Democratic Voters Index

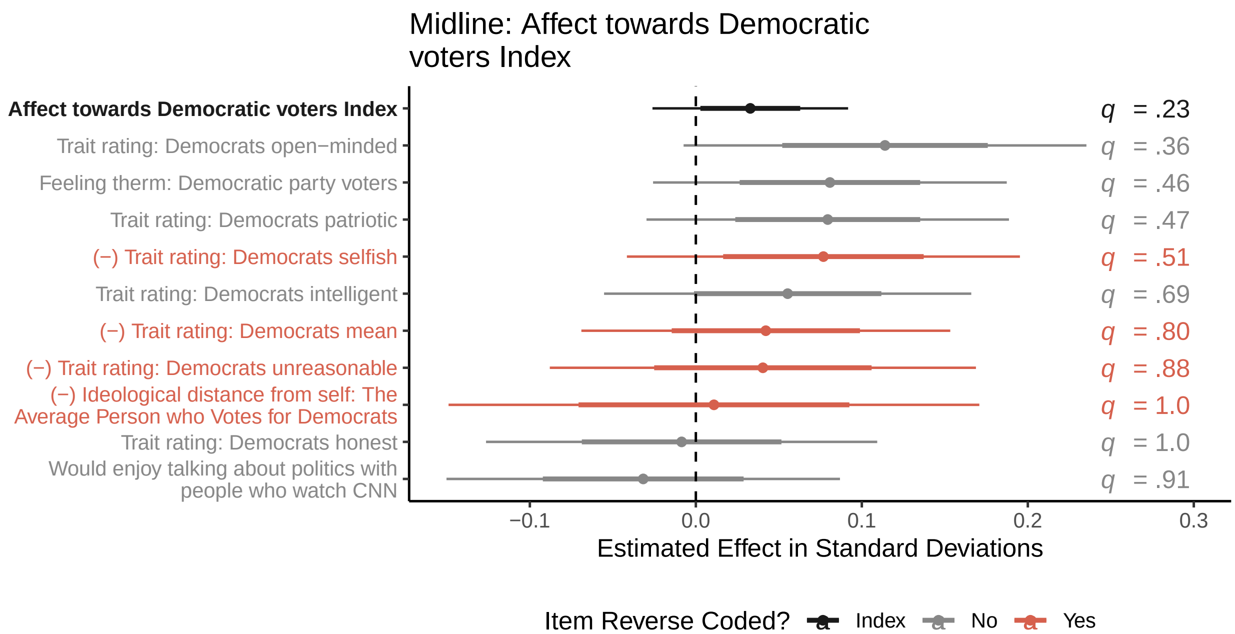

Figure OA21: General Media Attitudes Index

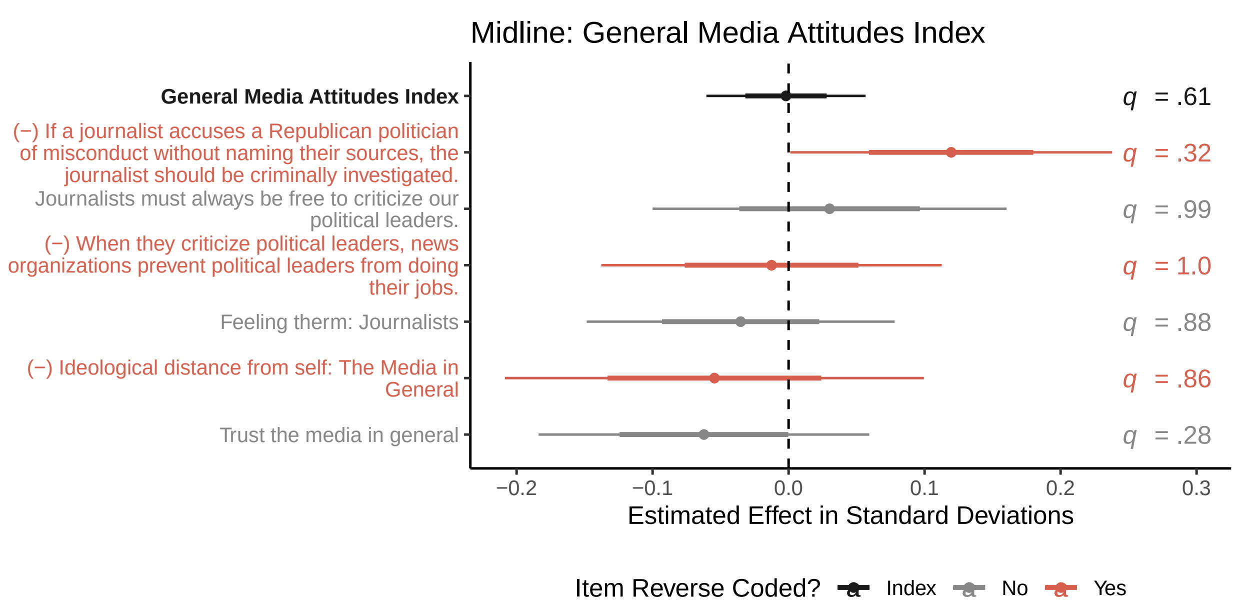

Figure OA22: Liberal Preferences on Non-Covered Issues Index

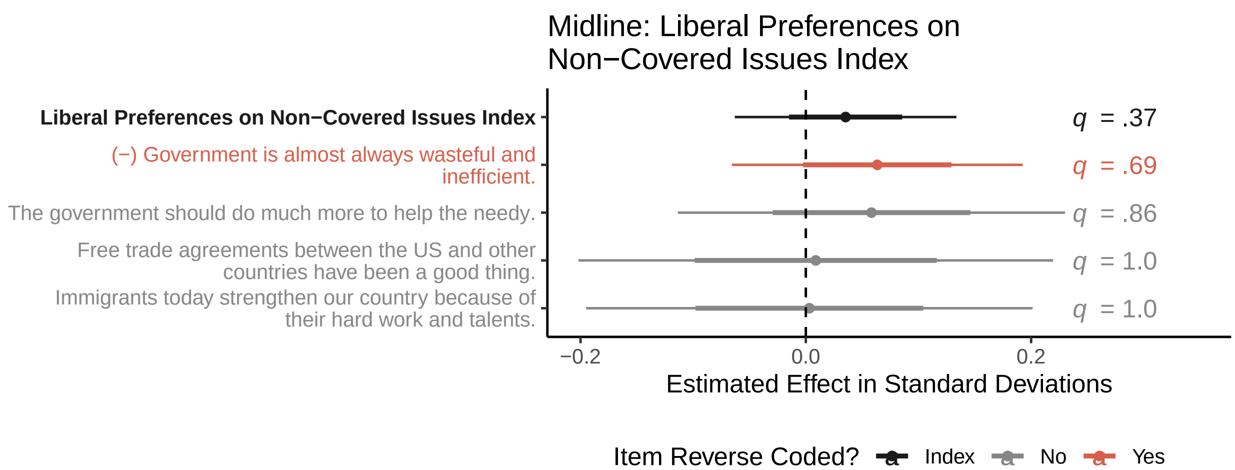

Figure OA23: Reduced Affect Towards Republican Voters Index

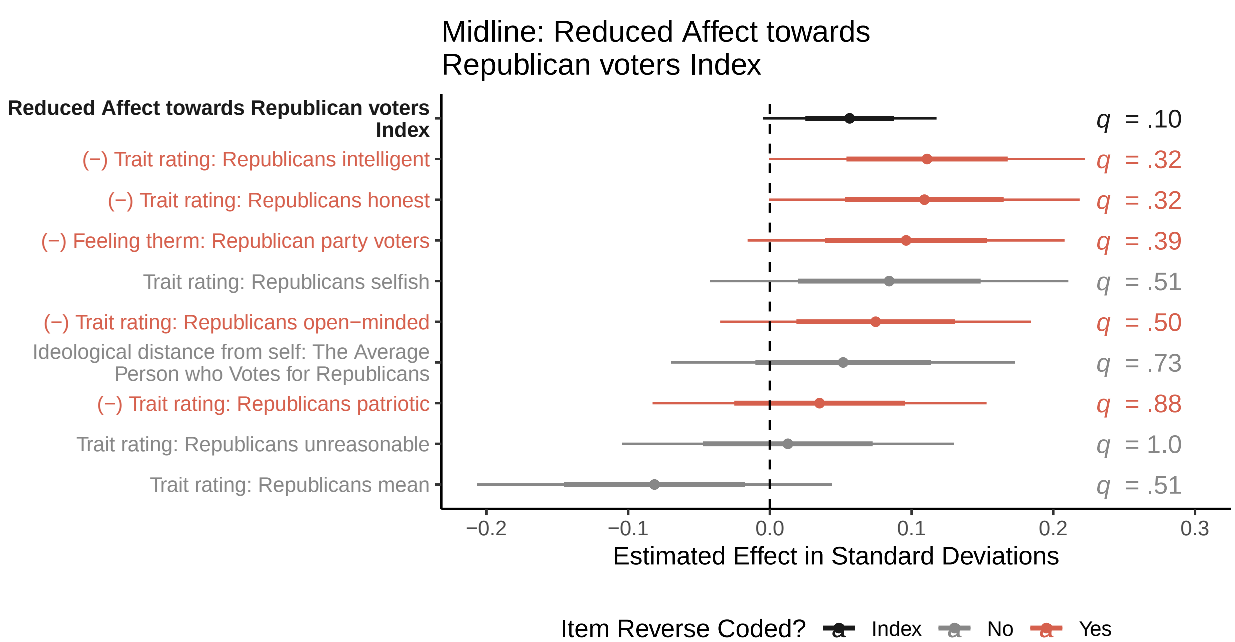

Figure OA24: Reduced Trump Evaluation Index

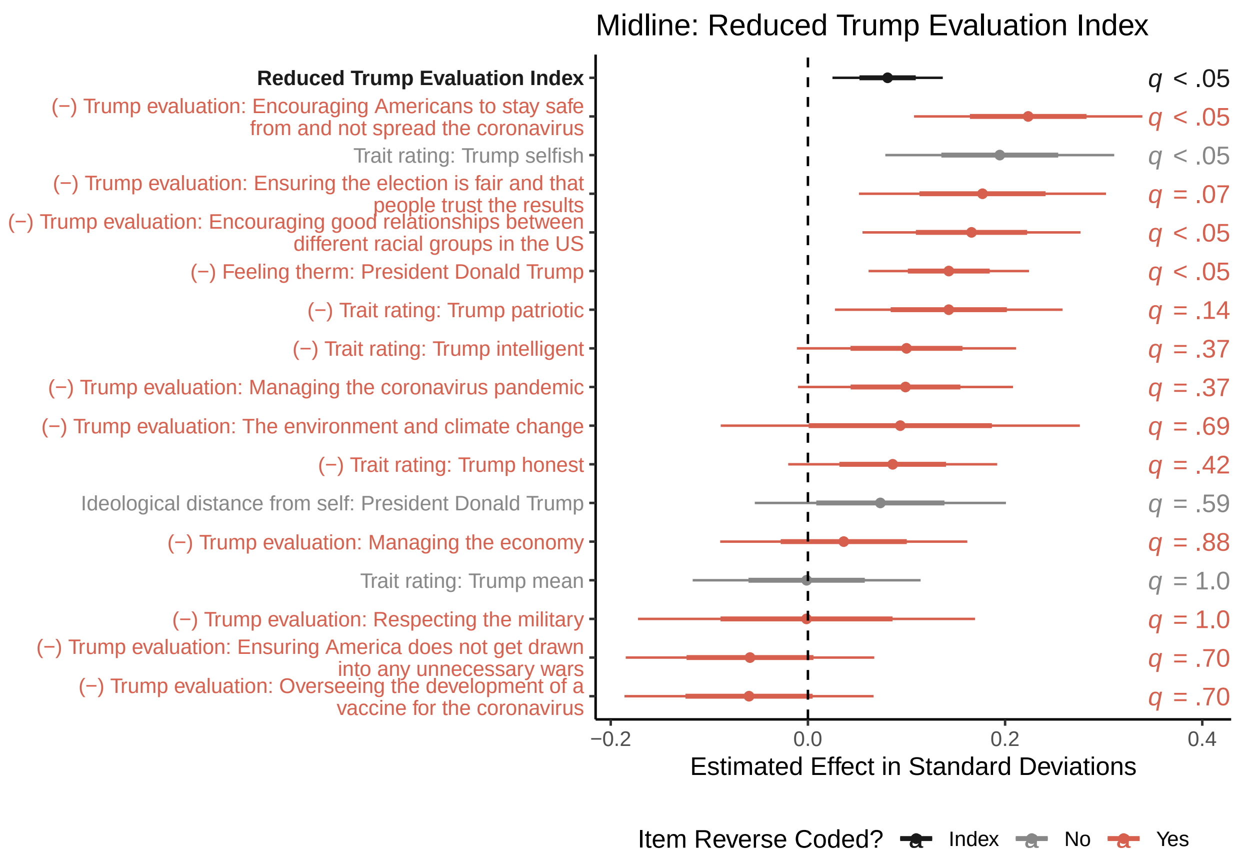

Figure OA25: Second Order Beliefs Index

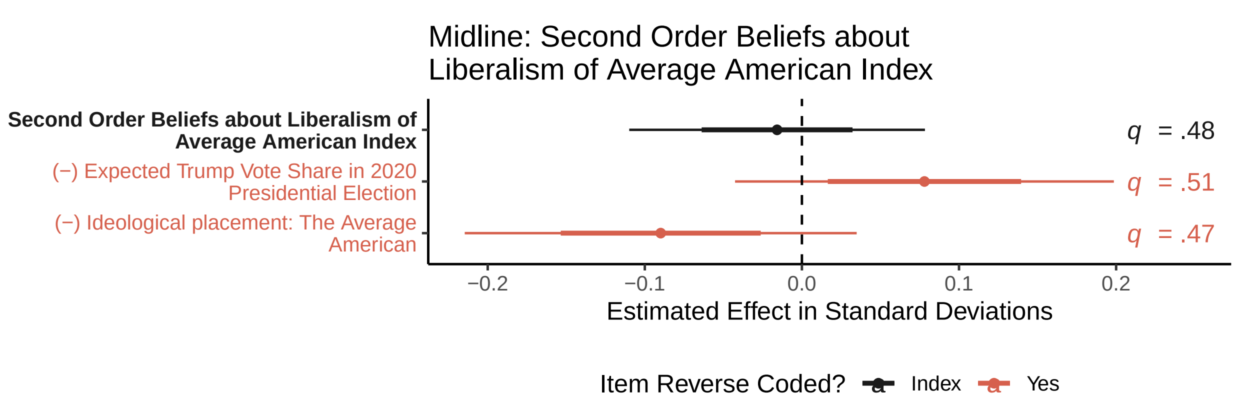

Figure OA26: Self-Reported Fox News Viewership Index

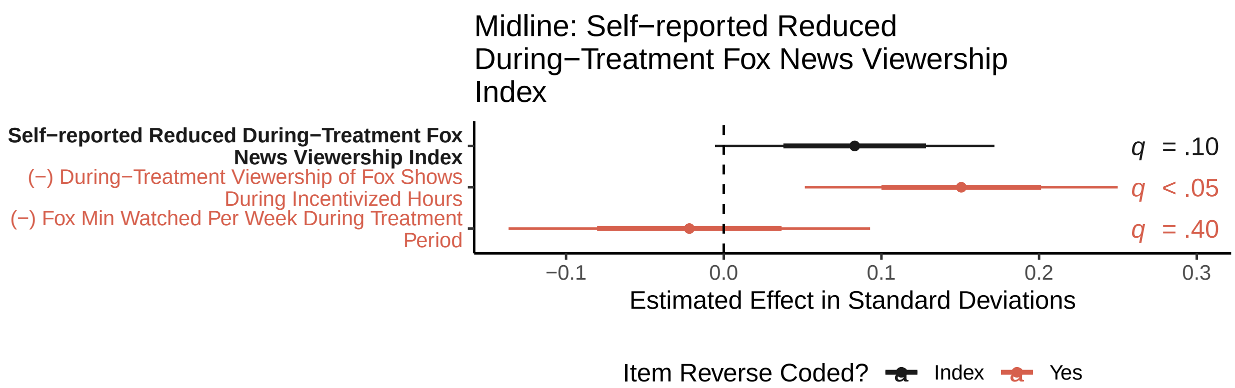

Figure OA27: Self-Reported Substitute News Source Index

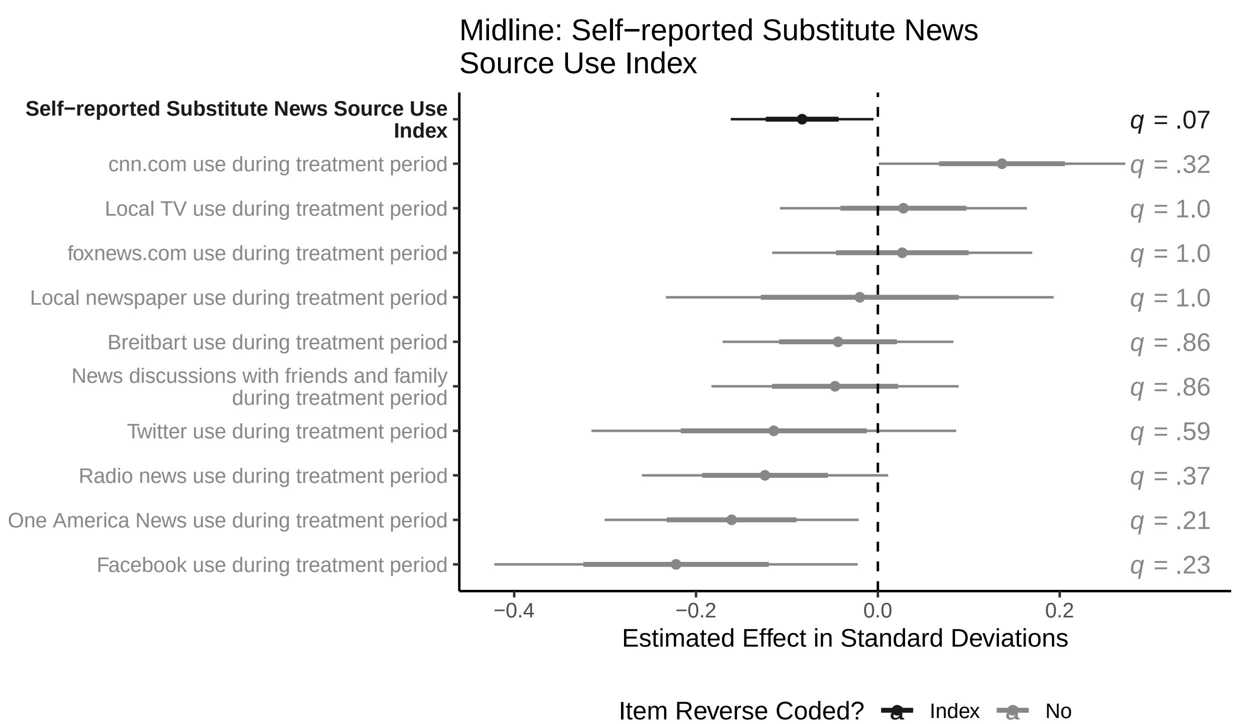

Figure OA28: Support for Democratic Norms Index

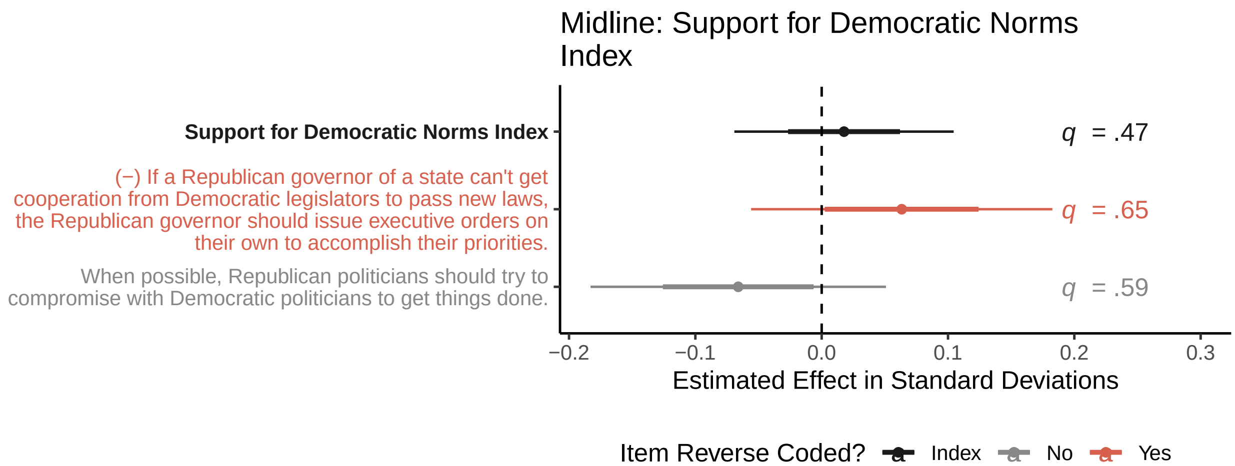

Figure OA29: Unfavorable Fox Attitudes Index

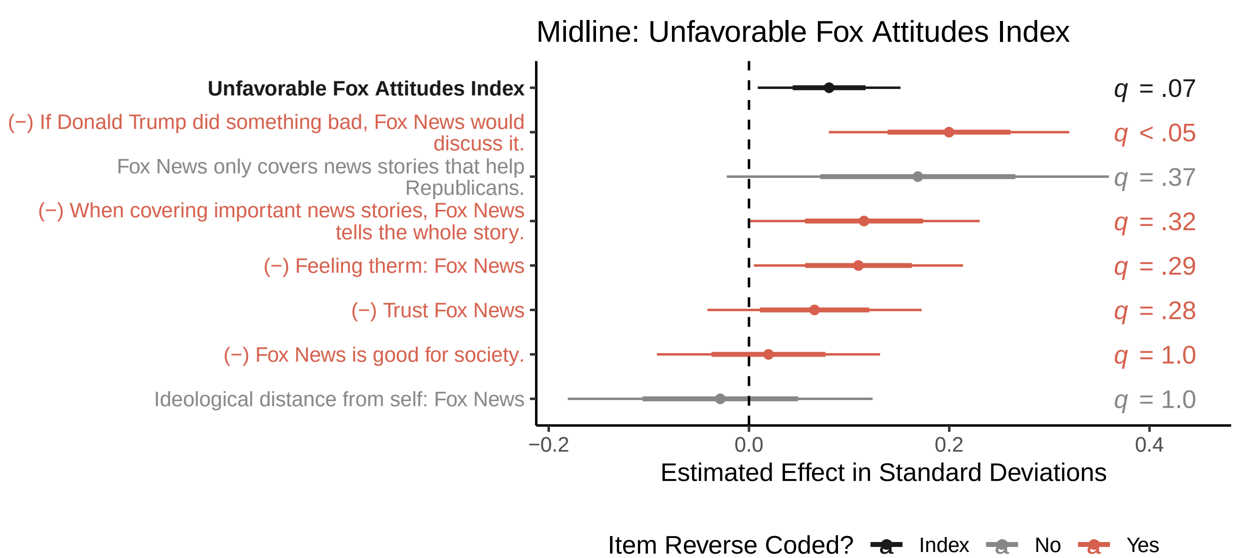

8.3 Endline Figures

Figure OA30: Effects on Pre-Registered Endline Indices

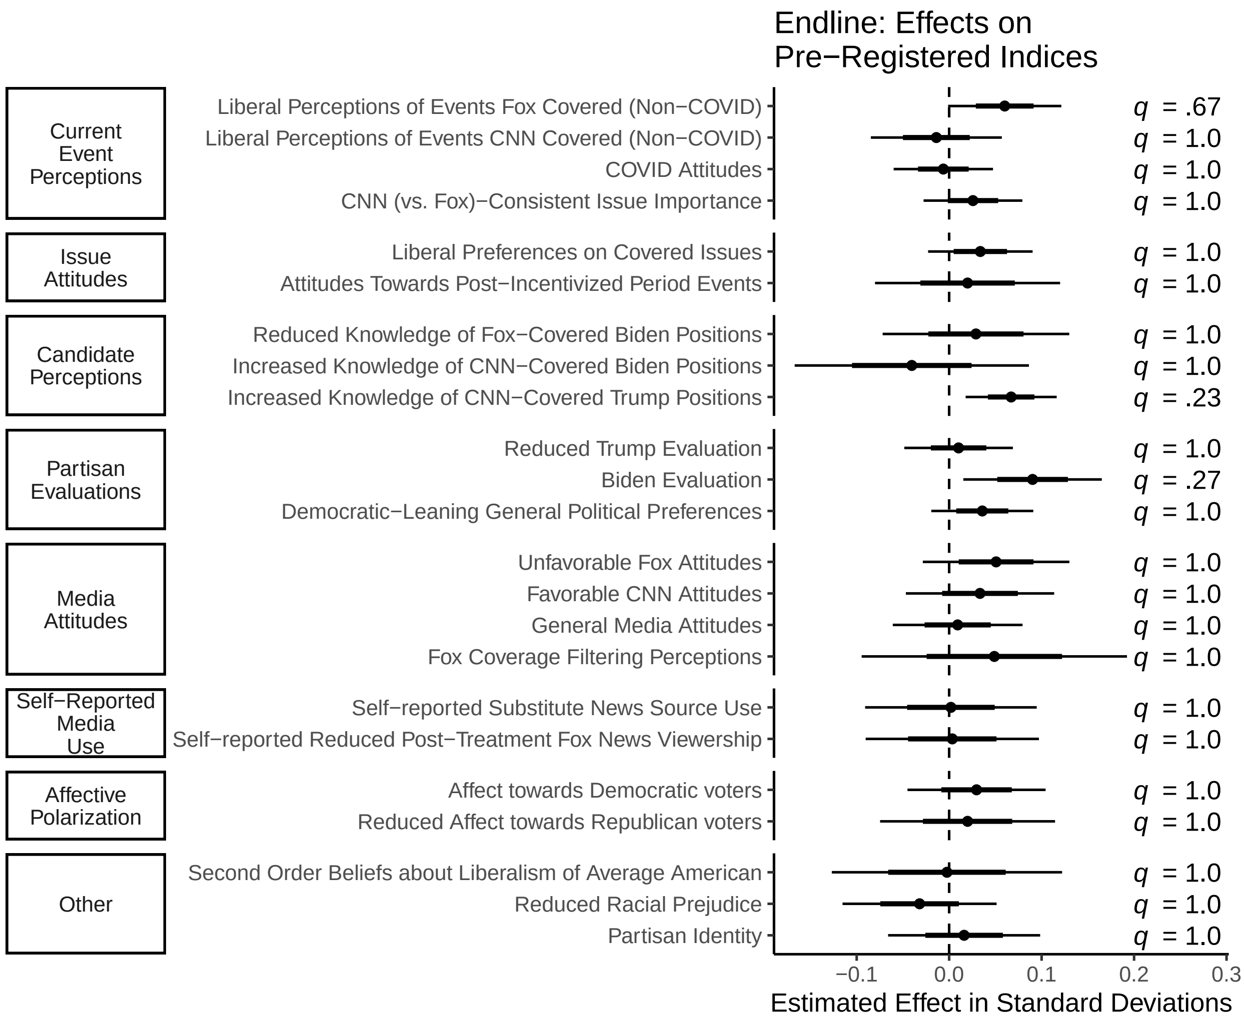

8.4 TV Viewership Data

The TV viewership data came from a media analytics company that linked the IP addresses of internet-connected televisions to household addresses and the voter file.

This data was necessary to conduct the experiment because, without it, it would have been impractical to locate a large enough sample of current Fox News viewers. However, although the viewership data appears to have succeeded in allowing us to recruit a sample that was much more likely to watch Fox News than the general population, the viewership data still contains substantial measurement error. In discussions with the media analytics company, this arises from two sources. First, the process which matches internet-connected televisions to voter file records contains some error. Second, households have multiple televisions, only some of which are captured in our data, meaning that the viewership data is likely to both undercount the true number of minutes watched and include viewership behavior from other household members.

This is why, to qualify for the experiment, individuals had to *both* be identified as likely Fox viewers in the viewership data *and* self-report watching a substantial amount of Fox News at baseline. Each of these sources has measurement error, but combining them allowed us to identify a final group for the experiment that would be very likely to actually watch Fox News. Consistent with our experimental subjects being by and large regularly Fox News viewers, as reported in the main paper, we found substantial treatment effects on a number of items related to content that was only present on Fox News.

At the same time, we would expect, and indeed find, that these sources of measurement error would lead us to substantially underestimate effects on the viewership outcomes themselves during the experiment. As an example of this, we estimate with the viewership data that the treatment group watched an additional 180 minutes of CNN during the incentivized period than the control group ($p < 0.001$). However, based on the quiz results, we believe this number to be higher. On average, individuals in the treatment group were incentivized to watch 5.9 hours per week of CNN for four weeks. The average treatment group respondent passed 4.1 out of 5 quizzes. If we assume the average treatment group respondent watched 80% of the CNN they were randomized to watch, then over the four week period, they would have watched 1,133 additional minutes of CNN compared to the 180 minutes we estimate from the viewership data (implying the viewership data only captures as little as approximately 16% of true viewership). This means that the viewership results are heavily attenuated towards zero.

Nevertheless, with these caveats in mind, below are the treatment effects on the television viewership data. We do not adjust these for multiple testing. Note that the incentivized period refers to the dates 2020-08-31 to 2020-09-25.

Table OA8: Effect on TV Viewership Data (log minutes)

Time Period	Network	Effect	SE	p-val
Incentivized Period	CNN	1.694	0.188	0.000
Oct.	CNN	0.055	0.094	0.555
Nov.	CNN	0.099	0.121	0.415
Incentivized Period	Fox News	-0.244	0.160	0.127
Oct.	Fox News	-0.174	0.165	0.291
Nov.	Fox News	-0.287	0.170	0.093
Incentivized Period	Total TV Viewership	-0.016	0.077	0.833
Oct.	Total TV Viewership	-0.072	0.083	0.386
Nov.	Total TV Viewership	-0.071	0.087	0.414

Table OA9: Effect on TV Viewership Data (raw minutes)

Time Period	Network	Effect	SE	p-val
Incentivized Period	CNN	184.143	30.033	0.000
Oct.	CNN	-3.250	3.417	0.342

Table OA9: Effect on TV Viewership Data (raw minutes) *(continued)*

Time Period	Network	Effect	SE	p-val
Nov.	CNN	5.956	5.638	0.291
Incentivized Period	Fox News	7.492	102.855	0.942
Oct.	Fox News	53.234	153.742	0.729
Nov.	Fox News	-60.715	138.566	0.661
Incentivized Period	Total TV Viewership	-23.467	196.238	0.905
Oct.	Total TV Viewership	-101.951	313.981	0.746
Nov.	Total TV Viewership	-364.508	313.723	0.246
t2_hours_yesterdayfoxnews	u	-8.256	5.400	0.127
t2_hours_yesterdaycnn	u	6.930	2.943	0.019
t2_minutes_perweek_fox	u	8.006	43.523	0.854
t2_minutes_perweek_cnn	u	394.488	35.684	0.000
t3_hours_yesterdayfoxnews	u	6.735	5.960	0.259
t3_hours_yesterdaycnn	u	4.500	1.701	0.008
t3_minutes_perweek_fox	u	3.074	42.261	0.942
t3_minutes_perweek_cnn	u	16.021	9.832	0.104

8.5 Heterogenous Treatment Effects (HTEs)

8.5.1 Pre-Registered HTEs

In the pre-analysis plan, we stated:

> Our main heterogenous treatment effect of interest is with respect to a binary variable coded as 1 if a person said in the baseline survey that the reason they watch Fox News is because "They share my point of view" and 0 otherwise.

This variable is `hte_pov` below, and is set to 1 if the person said they like watching Fox News because it shares their point of view and 0 otherwise.

We also wrote:

> We are particularly interested in heterogeneous treatment effects on the CNN Attitudes index, but this is of some interest across all the indices.

Given the large number of indices and individual items of interest, we subsequently decided to limit tihs analysis to only CNN Attitudes index. We did not look at any other results.

We also wrote:

> Because our sample is relatively homogenous on many baseline political dimensions, we do not have a priori strong expectations regarding heterogenous treatment effects by the other covariates that we measured. However, for strictly exploratory purposes only, we may compute heterogenous treatment effects by the following variables that we blocked on: num_hrs_incentivizing (Number of CNN hours assigned to watch / would have been assigned to watch (which was determined pre-treatment for all subjects)); baseline_partisan_factor; baseline_cnn_factor; baseline_fox_watch_factor
>
> We will look for HTEs by splitting the variables into terciles and calculating the conditional average treatment effects within each tercile of the variables named above. We will only plan to look for HTEs on indices and items listed as "Individual Item of Interest".

Given the large number of indices and individual items of interest, we subsequently decided not to investigate all of these potential heterogenous treatment effects (HTEs). Instead, we limit our HTE analysis to only `baseline_partisan_factor` and `baseline_fox_watch_factor` on the indices `t2_i_trump_evaluation` and

Tercile of Baseline Partisan Factor	Effect	SE	p-val.
Outcome = Trump Evaluation			
Tercile 1	0.111	0.054	0.042
Tercile 2	0.126	0.053	0.018
Tercile 3	-0.002	0.045	0.972
Outcome = Attitudes Fox Covered			
Tercile 1	0.147	0.060	0.015
Tercile 2	0.107	0.051	0.037
Tercile 3	0.131	0.048	0.007

Table OA10: Heterogenous Treatment Effects by Baseline Partisan Factor

Tercile of Baseline Fox Viewership Factor	Effect	SE	p-val.
Outcome = Trump Evaluation			
Tercile 1	0.138	0.052	0.008
Tercile 2	0.016	0.050	0.751
Tercile 3	0.177	0.051	0.001
Outcome = Attitudes Fox Covered			
Tercile 1	0.178	0.054	0.001
Tercile 2	0.128	0.058	0.029
Tercile 3	0.008	0.045	0.852

Table OA11: Heterogenous Treatment Effects by Baseline Fox Viewership Factor

`t2_i_attitudes_fox_covered`. We did not investigate any other HTEs. This is limited to only the midline survey. Given the exploratory nature of this analysis, we do not adjust for multiple comparisons.

We orient both variables such that higher values capture more conservative attitudes and more Fox News viewership. To provide more context, for the `baseline_partisan_factor`:

- Tercile 1 has an average baseline Trump thermometer rating of 64.1, an average Biden thermometer rating of 29.2, an average 7-point party identification of 2.6, and self-reported watching 1.5 hours yesterday and 622 minutes per week of Fox News.
- Tercile 2 has an average baseline Trump thermometer rating of 86.5, an average Biden thermometer rating of 11.7, an average 7-point party identification of 1.7, and self-reported watching 2.0 hours yesterday and 873 minutes per week of Fox News.
- Tercile 3 has an average baseline Trump thermometer rating of 96.5, an average Biden thermometer rating of 3.5, an average 7-point party identification of 1.3, and self-reported watching 2.4 hours yesterday and 1030 minutes per week of Fox News.

For the `baseline_fox_watch_factor`:

- Tercile 1 has an average baseline Trump thermometer rating of 73.5, an average Biden thermometer rating of 18.8, an average 7-point party identification of 2.1, and self-reported watching 0.7 hours yesterday and 289 minutes per week of Fox News.
- Tercile 2 has an average baseline Trump thermometer rating of 83.8, an average Biden thermometer rating of 14.2, an average 7-point party identification of 1.8, and self-reported watching 1.6 hours yesterday and 701 minutes per week of Fox News.
- Tercile 3 has an average baseline Trump thermometer rating of 89.7, an average Biden thermometer rating of 11.5, an average 7-point party identification of 1.7, and self-reported watching 3.5 hours yesterday and 1534 minutes per week of Fox News.

The below table summarizes the results.

Watches Fox for Point of View	Effect	SE	p-val.
Outcome = CNN Attitudes			
Watches Fox for POV0	0.014	0.067	0.834
Watches Fox for POV1	0.027	0.046	0.553

Table OA12: Heterogenous Treatment Effects on CNN Attitudes, by Watches Fox for POV

8.5.2 Can Priming Alone Explain The Results? (Not Pre-Registered)

As we note in the main text, the mechanism for the effects we uncover is impossible to definitively determine. We suspect that the changes in information participants learned on each of the two networks (i.e., the consequences of partisan coverage filtering) is chiefly responsible for the effects we found. However, a potential alternative explanation for our findings is *priming* due to *agenda setting*—i.e., that participants simply hearing more about topics on which they have liberal opinions at baseline caused their overall attitudes to become more liberal, as these pre-existing liberal attitudes were primed by CNN's agenda-setting efforts. For example, did respondents evaluate Trump more negatively because their pre-existing dissatisfaction with his handling of COVID-19 became more top-of-mind (priming) due to CNN's greater coverage of COVID (agenda-setting)?

Although we cannot rule out that this priming alternative contributes to the results, we can conduct a test of whether priming alone can entirely explain the results. In particular, the classic test for media priming is to examine how effects of consuming media vary by the baseline attitude being primed. In our data, we therefore test how some of our mind findings differ by baseline attitudes capturing people's views about topics that were extensively covered on Fox and CNN during the intervention.

From this vantage point, it should be noted at the outset that the composition of the sample we examine makes priming ex ante unlikely. As detailed above, the sample for the experiment began with strongly conservative attitudes, making it unlikely that priming any particular dimension would liberalize their attitudes; they have conservative attitudes on every dimension. However, we can examine this in more detail on particular attitudes dimensions.

First, since COVID-19 represented by far the most commonly discussed topic on CNN, we examine whether the effects are due to respondent's pre-existing views towards Trump's handling of COVID-19 being primed. The baseline survey asked "Do you agree or disagree with the following statements?", with one of the statements being "Donald Trump has done a good job handling the coronavirus pandemic." The response options were "Strongly agree", "Somewhat agree", "Neither agree nor disagree", "Somewhat disagree", and "Strongly disagree". We group the two "agree" and "disagree" responses together so there are three categories: those who approved of Trump's handling of COVID-19 at baseline ($n = 645$), those who neither approve nor disapprove ($n = 59$), and those who disapproved ($n = 40$).

Table OA13 shows the effects by whether respondents approved of Trump's handling of COVID-19 at baseline. As in the above analyses, we examine effects on two dependent variables, post-treatment (midline) evaluations of Trump and post-treatment (midline) attitudes towards the events that Fox News covered. Were priming to solely account for the findings, we would expect the effects to be limited to the subset of the sample that did not approve of Trump's handling of COVID-19 at baseline, whose negative attitudes on this dimension were then primed. We might even expect those who approved of Trump's handling of COVID-19 at baseline to approve of him more often after the treatment if this dimension were primed. This is not what we find. Table OA13 finds that the effects on Trump evaluation are actually smallest for those who disapproved of Trump's handling of COVID-19 at baseline, and are clearly present for those who approved of Trump's handling of COVID-19 at baseline. The second half of the table shows that the effects on attitudes towards events Fox covered are similar for all three categories, although only statistically significant among those who supporting Trump's handling at baseline, the largest group.

Since the intervention involved switching respondents away from consuming Fox News, it is also possible that the intervention could have had its effects by *reducing* the salience of issues related to protests against police violence, which were, as shown above, by far the most common topic on Fox News. Were priming to drive the results, we would therefore expect the effects to the limited to those who opposed the protests at baseline. This was most of the sample: at baseline, 654 people did not agree that they supported the recent protests against the police (i.e., opposed the protests), 45 neither agreed nor disagreed, and 44 agreed that they supported them. However, Table OA14 shows

Trichotomized Baseline Support for Protests	Effect	SE	p-val.
Outcome = Trump Evaluation			
Disapprove at Baseline	0.028	0.152	0.854
Neither at Baseline	0.105	0.121	0.387
Approve at Baseline	0.069	0.031	0.025
Outcome = Attitudes Fox Covered			
Disapprove at Baseline	0.194	0.231	0.406
Neither at Baseline	0.313	0.151	0.043
Approve at Baseline	0.142	0.031	0.000

Table OA13: Heterogenous Treatment Effects by Baseline Approval of Trump's Handling of COVID

Trichotomized Baseline Support for Protests	Effect	SE	p-val.
Outcome = Trump Evaluation			
Supportive	0.400	0.174	0.026
Neither	0.462	0.197	0.024
Opposed	0.080	0.030	0.008
Outcome = Attitudes Fox Covered			
Supportive	0.337	0.174	0.059
Neither	0.338	0.215	0.123
Opposed	0.129	0.030	0.000

Table OA14: Heterogenous Treatment Effects by Baseline Support for Protests

that the effects remain statistically significant, albeit imprecisely estimated, among those who actually supported the protests at baseline.

Tables OA13 and OA14 therefore suggest that priming is not driving the results. Priming entails holding constant one's evaluation of a politican or an issue on a certain dimension and merely changing the weight assigned to that dimension. Yet, if the treatment simply changed the weight respondents placed on protests or COVID-19 and worked through no other mechanism, we should not have seen the pattern of results we did above, with those who supported Trump's actions on COVID-19 still being affected by the treatment despite CNN's greater coverage of this issue, and those who supported the protests against police still being affected by the treatment despite them consuming less Fox News during this time. Taken together, then, although these results do not rule out some role for priming, especially among respondents who were more moderate or left-leaning at baseline, these results suggest priming is unlikely to be responsible for driving the preponderance of the effects we find.

One potential alternative explanation is that events outside the experiment changed respondents' attitudes between the baseline and the follow-up survey, and that these changed attitudes were then primed by the treatment. For example, a rising COVID caseload during September 2020 could have changed attitudes towards Trump's handling of COVID, which in turn the treatment could have primed. Were this the case, we would expect to see attitudes in the control group change between the baseline survey in August and the endline survey in September. Given that this is such a short time period, we would not expect meaningful changes in respondents' attitudes. And indeed, examining items that appeared on both the baseline and endline survey, we find no meaningful changes over time in the control group. (Note that the treatment changing those attitudes would not be a case of priming, that would be an instance of the treatment changing attitudes not priming them.)

9 Transcript Analysis During Incentivized Period

To help contextualize the coverage on each network during the incentivized period, we first provided a research assistant blind to our hypotheses with a list of broad topics (e.g., COVID-19) and, within these, subtopics, capturing specific information (see below table; e.g., long COVID exists). We based the subtopics on notes the research assistant had taken on the most common information covered on the two networks during the incentivized period. The research assistant then read every Fox News and CNN transcript during the incentivized period. If a subtopic was discussed, the research assistant then copied any quotes directly relevant to that subtopic into a spreadsheet. We finally manually audited a subset of this coding. We then counted and totaled the number of words in those quotes as a measure of media attention to that subtopic.

9.1 Word Counts by Topic and Subtopic

Table OA15 shows the total counts of words for each statement by network. The final column is the ratio of Fox coverage to CNN coverage. If CNN did not cover this statement (a 0 value), then this ratio is infinity.

Topic	Subtopic (I.e., Information)	Total CNN Words	Total Fox Words	Fox:CNN Ratio
Election Integrity	Mueller investigation: negative coverage (e.g., bias, witch hunt)	0	1619	Inf
Election Integrity	Trump will accept election results	0	101	Inf
Race/Protests	Biden position on tearing down statues	0	104	Inf
Economy/Taxes	Biden position on government spending	0	313	Inf
2020 Election	Biden embracing far left	236	5187	21.98
Other	Israel, the United Arab Emirates, and Bahrain recently signed a Middle East peace deal at the White House.	84	1278	15.21
Other	Global warming / climate change is NOT a problem	116	1739	14.99
COVID	Democratic elites violating COVID restrictions	209	2464	11.79
Race/Protests	Biden/Democrats support for extreme racial ideology/protests	1300	15236	11.72
2020 Election	Biden is NOT the favorite to win the Presidential election (e.g., polls are tied)	107	1085	10.14
Race/Protests	Negative consequences of extreme racial ideology/protests (e.g., violence)	1712	15003	8.76
Economy/Taxes	Biden position: Increase taxes on corporations and wealthy people.	44	234	5.32
Race/Protests	Trump's actions to advance racial equality	1186	5172	4.36
Election Integrity	Russa has NOT interfered to help Trump	46	190	4.13
COVID	Information downplaying severity of COVID-19	1332	5477	4.11
Healthcare (non-COVID)	Biden position on Medicare for All/expansion	96	311	3.24
Race/Protests	Racism is not a problem in US	854	2504	2.93
Election Integrity	Mail-in voting susceptible to fraud	1507	3242	2.15
COVID	Trump's actions to protect US from COVID-19	1695	3538	2.09
Race/Protests	Trump opposed to extreme racial ideology	942	1892	2.01
COVID	Description of Trump actions on COVID	178	317	1.78
Race/Protests	Biden condemnation of extreme racial ideology/protests	886	1405	1.59
Race/Protests	Trump condemning own supporters' extreme tactics	0	0	1.00
2020 Election	RNC anti-Semitic speakers	0	0	1.00
Economy/Taxes	Trump paying only $750 in taxes in 2016	0	0	1.00
Healthcare (non-COVID)	Trump position on pre-existing condition protections	0	0	1.00
Economy/Taxes	Trump position on taxes	49	49	1.00
Election Integrity	Mueller investigation: positive coverage	0	0	1.00
Race/Protests	Reporting on existence of BLM protests	2693	2136	0.79
Race/Protests	Jacob Blake was armed with knife	869	581	0.67
Other	Trump offensive comments about fallen soldiers	844	465	0.55
2020 Election	Biden is the favorite to win the Presidential election	455	201	0.44
COVID	Trump's position on in-person schooling	270	91	0.34
Election Integrity	Descriptions of voting by mail	1415	424	0.30
Race/Protests	Rittenhouse, 17-year-old who killed at Kenosha protest, is Trump supporter	183	47	0.26
Other	Global warming / climate change is a problem	1793	427	0.24
Race/Protests	Racism is a problem in US	4652	872	0.19
Race/Protests	Trump not condemning far-right extremists	2020	243	0.12
Election Integrity	Trump will not accept election results	1947	212	0.11
COVID	Trump's failures to protect US & his supporters from COVID-19	21244	2086	0.10
COVID	Trump's opposition to mask-wearing	3232	283	0.09
Election Integrity	Russia intervening to help Trump, undermine confidence in election	2923	254	0.09
Race/Protests	Supportive comments about BLM protests	715	56	0.08
Election Integrity	Mail-in voting is secure	4606	342	0.07
Race/Protests	Trump comment about looters/thugs on planes to Floyd protests	1993	138	0.07
COVID	Information indicating severity of COVID-19	10251	709	0.07
2020 Election	Former military/national security leaders, former Republican leaders praising/endorsing Joe Biden	1487	97	0.07
Election Integrity	Trump position on vote-by-mail	2559	102	0.04

Table OA15: Word Count of Statements Covered by Network (Aug 31 - 25 Sept)

9.2 Topic Areas

We also group subtopics into their overall topic areas. The below figure shows how topics vary by network.

Figure OA31: Topic Areas by Network

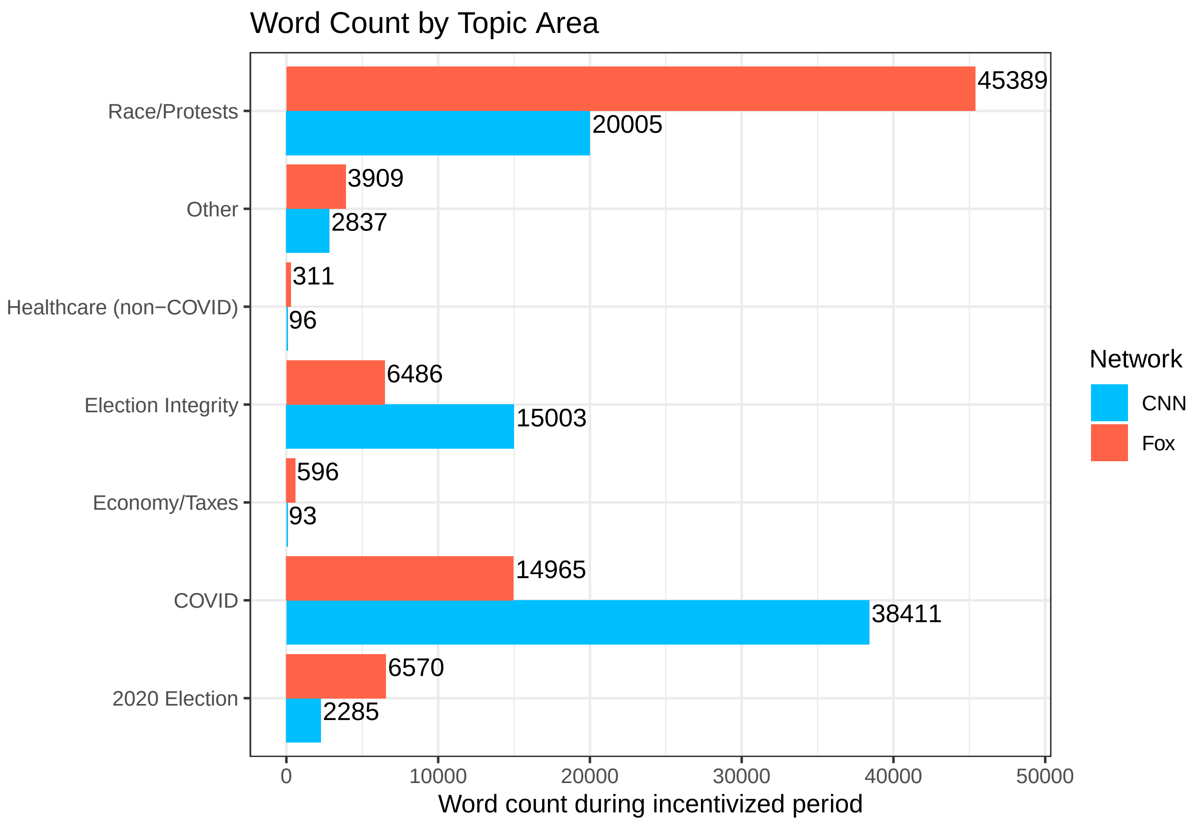

9.3 Item-Level Relationship Between Effects at Midline and CNN/Fox Coverage Content

We conducted an exploratory analysis to test whether we observe larger treatment effects on issues that FNC and CNN covered more (as measured using the word count from our transcript analysis) than on issues they covered less.

This analysis was not pre-registered and our transcript analysis was not conducted with this analysis in mind.

To conduct this analysis, we first mapped survey outcome measures to the topics included in our transcript analysis (e.g., the survey measure "The coronavirus causes many people to experience serious long-term health problems that stay with them for months or longer" was mapped to the transcript code "Information on or emphasis of the severity of the COVID-19 pandemic (e.g., infection rate, long COVID facts, number of deaths)"). We mapped some survey measures to multiple topics (e.g., the survey measure "Democrats are trying to steal the election with fraudulent mail-in ballots" was mapped to "Voting by Mail Negative Coverage (fraud, etc.)" and "Voting by Mail Positive Coverage (safe, etc.)"). In these cases, we added up the total words from each topic. In cases where a single topic in the transcript analysis mapped to multiple survey measures, we evenly distributed those word counts across the different survey measures.

Figure OA32 shows the precision-weighted correlation between treatment effect estimates and the amount of coverage that survey measure received on FNC and CNN, combined, during the incentivized period. For the non-COVID items, we find a positive relationship, suggesting that we observe larger treatment effects on those measures that received more coverage. However, for the COVID items, we observe no relationship. This could be related to measurement error in the COVID items, as mapping the survey measures to the transcript analysis was more difficult than the non-COVID items.

Figure OA32: Relationship Between Coverage Volume and Treatment Effect Estimates

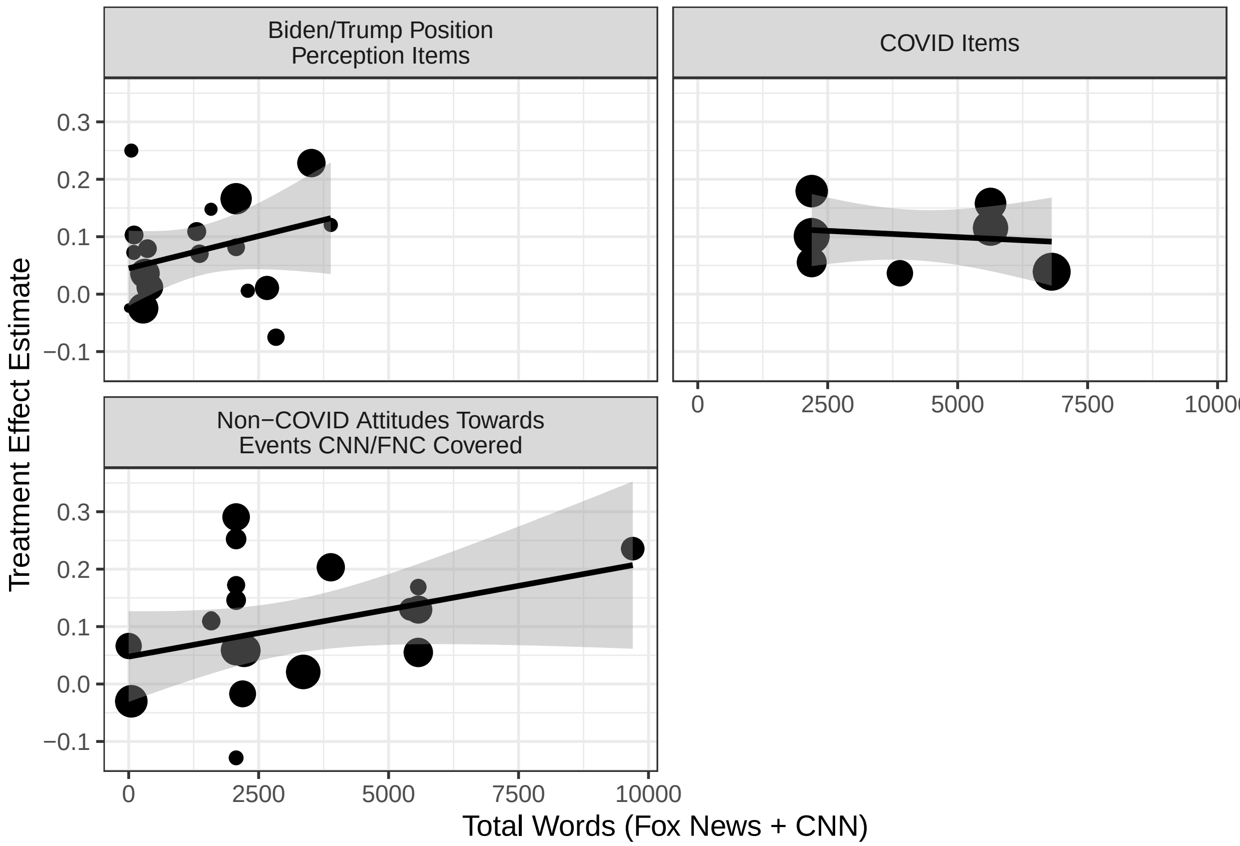

10 Additional Appendix Figures and Tables

Table OA16: Effect on Example Dichotomized Items

Variable Name	Control (Fox) Mean	Treatment (CNN) Mean	Treatment Effect	Standard Error	Item Wording	Coding
t2_selfview_biden_blm_shot_u	0.634	0.504	-0.130	0.033	Do you agree or disagree with the following statements? If Joe Biden is elected President, we'll see many more police get shot by Black Lives Matter activists.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
q2_truefalse_blake_armed_u	0.804	0.678	-0.126	0.034	Below are several statements. Some statements below are completely true. Others are at least partly false. Some are completely false. Which of the below statements do you think are completely true? If you think part of the statement is false, select false. If you don't know, just make your best guess. Jacob Blake, who was recently shot by police in Kenosha, Wisconsin, was armed with a knife and had engaged in a violent struggle with officers moments before officers shot him.	1 = I think this is COMPLETELY TRUE; 0 = I think this is FALSE
q2_selfview_protests_blake_u	0.754	0.636	-0.118	0.033	Do you agree or disagree with the statements below? It is an overreaction to go out and protest in response to the police shooting of Jacob Blake in Kenosha, Wisconsin.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
t2_selfview_biden_elect_protes	0.600	0.486	-0.114	0.036	Do you agree or disagree with the following statements? If Joe Biden is elected President, the protests in America's cities will only get more violent.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
t2_trump_eval_safe_covid_u	0.389	0.263	-0.126	0.032	How would you evaluate President Trump's performance in these areas? Encouraging Americans to stay safe from and not spread the coronavirus	1 = performed somewhat or way above expectations; 0 = at or below expectations
t2_agenda_covid_vs_blm_u	0.514	0.395	-0.119	0.034	On a scale of 1 to 10, what problem do you think it is more important for the President to focus on, the coronavirus pandemic (1), violent protests (10), or somewhere in between?	6-10 = 1 (violent protests more important than COVID); 1-5 = 0 (COVID at least as important as violent protests)

Table OA16: Effect on Example Dichotomized Items *(continued)*

Variable Name	Control (Fox) Mean	Treatment (CNN) Mean	Treatment Effect	Standard Error	Item Wording	Coding
t2_truefalse_trump_rally_safe_	0.581	0.473	-0.108	0.035	Which of the below statements do you think are completely true? If you think part of the statement is false, select false. If you don't know, just make your best guess. Donald Trump's campaign is taking significant safety precautions at its rallies to reduce the risk that rally attendees spread the coronavirus to each other.	1 = I think this is COMPLETELY TRUE; 0 = I think this is FALSE
t2_selfview_biden_sup_pol_shot	0.567	0.464	-0.103	0.034	Do you agree or disagree with the following statements? I think many Biden supporters are probably happy to see it when police officers get shot.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
t2_issue_trump_masks_u	0.556	0.437	-0.120	0.035	If you had to guess, where do you think Republican Donald Trump stands on each of these proposals? Americans should wear masks in public.	1 = Trump supports; 0 = Trump opposes or not sure
t2_selfview_dem_steal_election	0.788	0.695	-0.093	0.027	Do you agree or disagree with the following statements? Democrats are trying to steal the election with fraudulent mail-in ballots.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
q4_selfview_covid_infecting_ma	0.412	0.498	0.086	0.034	Do you agree or disagree with the statements below? The coronavirus is still infecting huge numbers of Americans every day.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
t2_issueview_votebymail_u	0.175	0.244	0.069	0.026	State and national leaders have debated many important issues recently. For each of the following, tell us whether you agree or disagree with the statement in principle States should allow voters to vote by mail in the 2020 election.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
t2_issue_biden_policefund_u	0.391	0.333	-0.058	0.030	If you had to guess, where do you think Democrat Joe Biden stands on each of these proposals? Eliminate all funding for the police.	1 = Joe Biden supports; 0 = Joe Biden opposes or not sure

Table OA16: Effect on Example Dichotomized Items *(continued)*

Variable Name	Control (Fox) Mean	Treatment (CNN) Mean	Treatment Effect	Standard Error	Item Wording	Coding
t2_trump_eval_race_relation_u	0.479	0.426	-0.054	0.033	How would you evaluate President Trump's performance in these areas? Encouraging good relationships between different racial groups in the US	1 = performed somewhat or way above expectations; 0 = at or below expectations
t2_covid_other_countries_u	0.157	0.206	0.049	0.026	Do you agree or disagree with the following statements about the coronavirus (COVID-19) pandemic? Many other countries have done a much better job of controlling the coronavirus than the United States government.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
t2_covid_longterm_u	0.456	0.504	0.048	0.032	Do you agree or disagree with the following statements about the coronavirus (COVID-19) pandemic? The coronavirus causes many people to experience serious long-term health problems that stay with them for months or longer.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)
t2_news_eval_fox_trump_bad_u	0.874	0.822	-0.052	0.023	Do you agree or disagree with the below statements? If Donald Trump did something bad, Fox News would discuss it.	1 = Somewhat or strongly agree; 0 = otherwise (neither agree/disagree or disagree)

Note:

The Table shows the effects on some of the items with the largest point estimates when those items are dichotomized to make the results more interpretable. This was not pre-registered. We offer these estimates to help aid interpretability.

References

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer and Matthew Gentzkow. 2020. "The welfare effects of social media." *American Economic Review* 110(3):629–76.

Anderson, Michael L. 2008. "Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American statistical Association* 103(484):1481–1495.

Arceneaux, Kevin and Martin Johnson. 2013. *Changing minds or changing channels?: Partisan news in an age of choice*. University of Chicago Press.

Bail, Christopher A, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout and Alexander Volfovsky. 2018. "Exposure to opposing views on social media can increase political polarization." *Proceedings of the National Academy of Sciences* 115(37):9216–9221.

Belloni, Alexandre, Victor Chernozhukov and Christian Hansen. 2014. "High-dimensional methods and inference on structural and treatment effects." *Journal of Economic Perspectives* 28(2):29–50.

Besley, Timothy and Andrea Prat. 2006. "Handcuffs for the grabbing hand? Media capture and government accountability." *American Economic Review* 96(3):720–736.

Bloniarz, Adam, Hanzhong Liu, Cun-Hui Zhang, Jasjeet S Sekhon and Bin Yu. 2016. "Lasso adjustments of treatment effect estimates in randomized experiments." *Proceedings of the National Academy of Sciences* 113(27):7383–7390.

Casas, Andreu, Ericka Menchen-Trevino and Magdalena Wojcieszak. 2022. "Exposure to extremely partisan news from the other political side shows scarce boomerang effects." *Political Behavior* pp. 1–40.

Chen, Yuyu and David Y Yang. 2019. "The impact of media censorship: 1984 or brave new world?" *American Economic Review* 109(6):2294–2332.

de Benedictis-Kessner, Justin, Matthew A Baum, Adam J Berinsky and Teppei Yamamoto. 2019. "Persuading the enemy: Estimating the persuasive effects of partisan media with the preference-incorporating choice and assignment design." *American Political Science Review* 113(4):902–916.

Druckman, James N, Matthew S Levendusky and Audrey McLain. 2018. "No Need to Watch: How the Effects of Partisan Media Can Spread via Interpersonal Discussions." *American Journal of Political Science* 62(1):99–112.

Druckman, James N, SR Gubitz, Matthew Levendusky and Ashley Lloyd. 2019. "How incivility on partisan media (de-) polarizes the electorate." *Journal of Politics* 81(1).

Fang, Albert H, Andrew M Guess and Macartan Humphreys. 2019. "Can the government deter discrimination? Evidence from a randomized intervention in New York City." *The Journal of Politics* 81(1):127–141.

Feldman, Lauren. 2011. "The opinion factor: The effects of opinionated news on information processing and attitude change." *Political Communication* 28(2):163–181.

Gentzkow, Matthew and Jesse M Shapiro. 2006. "Media bias and reputation." *Journal of Political Economy* 114(2):280–316.

Gerber, Alan S and Donald P Green. 2012. *Field experiments: Design, analysis, and interpretation*. WW Norton.

Guess, Andrew M, Pablo Barberá, Simon Munzert and JungHwan Yang. 2021. "The consequences of online partisan media." *Proceedings of the National Academy of Sciences* 118(14).

Hayakawa, Samuel Ichiyé. 1940. *Language in thought and action*. New York: Hartcourt, Brace and Company.

Iyengar, Shanto and Donald R Kinder. 1987. *News That Matters: Television and American Opinion*. University of Chicago Press.

Jones, Damon, David Molitor and Julian Reif. 2019. "What do workplace wellness programs do? Evidence from the Illinois workplace wellness study." *The Quarterly Journal of Economics* 134(4):1747–1791.

Krosnick, Jon A and Donald R Kinder. 1990. "Altering the foundations of support for the president through priming." *American Political Science Review* 84(2):497–512.

Leeper, Thomas J. and Rune Slothuus. 2020. How the news media persuades: Framing effects and beyond. In *Oxford Handbook of Electoral Persuasion*, ed. Elizabeth Suhay, Bernard Grofman and Alexander H. Trechsel. Oxford: Oxford University Press pp. 151–168.

Levendusky, Matthew. 2013*a*. *How partisan media polarize America*. University of Chicago Press.

Levendusky, Matthew. 2013*b*. "Partisan media exposure and attitudes toward the opposition." *Political communication* 30(4):565–581.

Levendusky, Matthew S. 2013*c*. "Why do partisan media polarize viewers?" *American Journal of Political Science* 57(3):611–623.

Levy, Ro'ee. 2021. "Social media, news consumption, and polarization: Evidence from a field experiment." *American Economic Review* 111(3):831–70.

McCombs, Maxwell E and Donald L Shaw. 1972. "The agenda-setting function of mass media." *Public opinion quarterly* 36(2):176–187.

Michler, Jeffrey D and Anna Josephson. 2022. Recent developments in inference: practicalities for applied economics. In *A Modern Guide to Food Economics*, ed. Jutta Roosen and Jill E. Hobbs. Northampton, MA: Edward Elgar Publishing pp. 235–268.

Mullainathan, Sendhil and Andrei Shleifer. 2005. "The market for news." *American economic review* 95(4):1031–1053.

Searles, Kathleen, Joshua P Darr, Mingxiao Sui, Nathan Kalmoe, Raymond Pingree and Brian Watson. 2022. "Partisan media effects beyond one-shot experimental designs." *Political Science Research and Methods* 10(1):206–214.

Westfall, Peter H and S Stanley Young. 1993. *Resampling-based multiple testing: Examples and methods for p-value adjustment*. Vol. 279 John Wiley & Sons.

Wojcieszak, Magdalena, Sjifra de Leeuw, Andreu Casas, Xudong Yu, Ericka Menchen-Trevino, Bernhard Clemm von Hohenberg and Miriam Boon. 2021. "Null effects of news exposure: A causal test of the (un) desirable effects of a 'news vacation'and 'news binging'." *ResearchSquare Working Paper* .