

Persuading the Enemy: Estimating the Persuasive Effects of Partisan Media with the Preference-Incorporating Choice and Assignment Design

JUSTIN DE BENEDICTIS-KESSNER *Boston University*

MATTHEW A. BAUM *Harvard University*

ADAM J. BERINSKY *Massachusetts Institute of Technology*

TEPPEI YAMAMOTO *Massachusetts Institute of Technology*

Does media choice cause polarization, or merely reflect it? We investigate a critical aspect of this puzzle: How partisan media contribute to attitude polarization among different groups of media consumers. We implement a new experimental design, called the Preference-Incorporating Choice and Assignment (PICA) design, that incorporates both free choice and forced exposure. We estimate jointly the degree of polarization caused by selective exposure and the persuasive effect of partisan media. Our design also enables us to conduct sensitivity analyses accounting for discrepancies between stated preferences and actual choice, a potential source of bias ignored in previous studies using similar designs. We find that partisan media can polarize both its regular consumers and inadvertent audiences who would otherwise not consume it, but ideologically opposing media potentially also can ameliorate the existing polarization between consumers. Taken together, these results deepen our understanding of when and how media polarize individuals.

Recently, pundits, politicians, and ordinary citizens have expressed growing concern over political polarization in the United States and the media's purported role in exacerbating the problem. In a January 2018 interview, former President Barack Obama observed that viewers of Fox News are "living on a different planet" from National Public Radio listeners. He added that "We are operating in completely different information universes... At a certain point, you just live in a bubble. And that's part of why our politics is so polarized right now."¹

Perhaps, as Obama suggests, partisan news media cause, or at least exacerbate, polarization between


liberals and conservatives. This perspective places the blame at the feet of the media: by presenting one-sided versions of issues, partisan media outlets like Fox News on the right and MSNBC on the left drive Americans apart. These slanted news sources persuade individuals that a particular one-sided version of issues represents the unvarnished truth.


A second possibility is that in a world of fragmented media, the more polarized our country becomes, the more pre-existing political attitudes drive people's choices of news outlets. If so, this represents a clear example of selective exposure (e.g., Lazarsfeld, Berelson, and Gaudet 1948)—the tendency to seek out information that reinforces preexisting views.


These competing perspectives on the relationship between the media and the public suggest very different roles of partisan media in the modern political system. Do like-minded individuals seek out partisan news sources that support their preexisting beliefs—resulting in a tendency toward a particular perspective among consumers of ideologically narrow partisan media outlets via self-selection? Or do consumers of partisan news alter their views to reflect information encountered in such outlets, resulting in increased polarization between consumers of liberal and conservative news? In the former instance, media choice *reflects* polarization; in the latter, media choice *causes* it. Relatedly, is the rise of partisan media a cause or an effect of public polarization? If the former, how might partisan media exacerbate the polarization of opinion?

We investigate whether and how partisan media contribute to attitude polarization by assessing the extent to which such media can persuade consumers with differing preexisting media preferences. That is, do partisan media cause individuals to change their attitudes in line with the valence of the media message? Adopting a new experimental design, which we call the Preference-Incorporating Choice and Assignment (PICA) design,

Justin de Benedictis-Kessner , Assistant Professor, Department of Political Science, Boston University, jdbk@bu.edu.

Matthew A. Baum , Kalb Professor, John F. Kennedy School of Government, Harvard University, John F. Kennedy School of Government, Matthew_Baum@harvard.edu.

Adam J. Berinsky , Mitsui Professor, Department of Political Science, Massachusetts Institute of Technology, berinsky@mit.edu.

Teppey Yamamoto , Associate Professor, Department of Political Science, Massachusetts Institute of Technology, teppey@mit.edu.

For comments, suggestions, and advice, we thank Kevin Arceneaux, David Broockman, Andres Cambronero, Matthew Gentzkow, Dan Hopkins, Dean Knox, Matthew Levendusky, Robert Pressel, Markus Prior, Jesse Shapiro, and Natalie Stroud. For excellent research assistance, we thank Grace Chao and Henry Feinstein. Financial support was provided by the National Science Foundation (SES-1528487) and the Political Experiments Research Lab (PERL) at MIT. Berinsky's contribution was supported in part by a Joan Shorenstein Fellowship. Replication files are available at the American Political Science Review Dataverse: <https://doi.org/10.7910/DVN/MM2VWC>.

Received: May 18, 2018; revised: February 19, 2019; accepted: June 5, 2019. First published online: July 16, 2019.

¹ Chandran, Nyshka. 2018. "Obama to David Letterman: Media Is Dividing Americans." *CNBC*. Online: <https://www.cnn.com/2018/01/12/former-president-barack-obama-warns-on-polarizing-media-us-electoral-system.html>.

this study combines traditional experimentally randomized exposure to media with a more organic choice process. Our PICA design expands upon the patient preference trial (PPT) framework by not only measuring media consumption preferences but also by accounting, within the same procedure, for the effects of both forced exposure to media and choice among media options. Our design therefore allows us to estimate the degree of selective exposure, while simultaneously disaggregating our estimates of the persuasive power of the media by individuals' preferences for particular media sources.

We estimate the degree of polarization originating from selective exposure by comparing opinions among different groups of news consumers. We distinguish between those individuals who would, if given the choice, read a given partisan media story and those who would elect *not* to read partisan news. This allows us to compute the degree of polarization that exists absent exposure to ideologically dissonant media.

We also compute the treatment effect of partisan news on political opinions and behavior, disaggregated into a series of choice-specific treatment effects. We thus demonstrate the heterogeneity of partisan media's power on different groups of news consumers. We show that partisan media are most likely to persuade individuals who would choose to consume entertainment media rather than partisan media. The treatment effects after one instance of exposure among these inadvertent partisan media consumers are substantial when compared to real-world polarization, reaching close to half the size of existing attitudinal differences between partisan news consumers from opposing sides of the political spectrum.

Those individuals who *do* prefer partisan media are also persuadable, albeit to a limited extent. Consuming opposing partisan media can decrease polarization under some circumstances. We therefore provide evidence that while selective exposure—self-selection by readers into ideologically consonant partisan news sources—limits the exposure of ideological extremists to media that causes them to moderate their opinions, partisan media may ameliorate this polarization. However, we also show that the behavioral effects of partisan media are strongest among consumers of MSNBC, indicating that the downstream effects of exposure to partisan media are likely to have ramifications for how people share information with others in their lives. Additionally, in an advance over other studies that have employed the PPT framework, our design also enables us to estimate the sensitivity of these results to discrepancies between people's stated media preferences and their actual choices, an important source of potential bias unaccounted for in prior research. These sensitivity analyses highlight the consequences of relying on naïve treatment effect estimates in experimental paradigms where there is a danger of large discrepancies between survey-expressed preferences and actual choices. Finally, while we report the results based on a single news topic in the main text, in Online Appendix E we present replications employing three additional news topics, survey modalities, and population samples, thereby demonstrating that our findings are generalizable.

POLARIZATION AS A CAUSE OR AN EFFECT OF PERSUASION?

There is a broad agreement that Americans increasingly have *opportunities* to consume unbalanced news. For instance, the typical US household now receives about 190 television channels, more than a tenfold increase since 1980. The options for different news sources on the Internet are even more numerous. This explosion of media outlets has vastly increased the choices available to consumers and allowed for the development of ideological “niche” news outlets (Hamilton 2005). Partisan media are widely accessible (Baum and Groeling 2010; Groeling 2013) and often present news that is more beneficial to one party than the other (Baum and Groeling 2008). The public widely recognizes this partisan slant, routinely associating Fox News and MSNBC with heavy biases toward the Republican and Democratic parties, respectively (Ladd 2012).

But how do such media actually influence the political preferences of Americans? Whenever social scientists observe a difference between actors exposed to different stimuli in a context where it is not possible to control who gets the treatment, the question arises as to whether any effect results from the treatment itself, or from preexisting differences between the actors exposed to different treatments. This is the problem of self-selection. Many studies designed to determine the effects of partisan media on polarization are thus ill-equipped to disentangle the preexisting differences from the effects of media treatments. The observed differences in attitudes among individuals exposed to partisan information could stem from variation in either the information itself or the kinds of individuals who choose to expose themselves to it.

Most research has addressed only one of these concerns. Some studies treat polarization primarily as a *cause* of media fragmentation, through individuals' decisions to selectively expose themselves to partisan media. Others treat polarization as an *effect*, focusing on the possibility that media fragmentation enhances political polarization because partisan media persuade individuals to adopt more extreme political views. We discuss both approaches, before proposing a corrective for individuals' tendencies to inaccurately recall their media consumption.

INTEGRATING SELECTIVE EXPOSURE WITH PERSUASION

Research dating back to the 1940s (e.g., Campbell et al. 1960; Lazarsfeld, Berelson, and Gaudet 1948) has theorized that selective exposure to information causes divergent political opinions. This research found evidence of such selective exposure to partisan information in media consumption patterns.

Over the last decade, researchers have more precisely identified some of the conditions under which selective exposure occurs. Many researchers (Arceneaux, Johnson, and Murphy 2012; Gaines and Kuklinski 2011; Iyengar and Hahn 2009; Stroud 2011) have shown that Democrats and Republicans—especially the strongest

(Iyengar et al. 2008; Kim 2009) and most politically engaged (Bennett and Iyengar 2008) partisans—prefer news that supports their preexisting beliefs. The implication is that political preferences shape media choice, albeit with some limitations due to increasing online information consumption (Brundidge 2010; Messing and Westwood 2014; Mummolo 2016).

Widespread self-selection into partisan media streams is troubling for democracy because it could lead to increasingly insular partisan information silos among the public (Sunstein 2001; Negroponete 1995; Pariser 2012; but see Gentzkow and Shapiro 2011; Leeper 2014; Prior 2007). If individuals only expose themselves to one side of an argument, they may disproportionately reinforce their attitudes, thereby becoming less inclined to compromise or moderate their views.

This line of research, however, does not account for persuasion by the media. Recent innovations in our understanding of human information processing have upended the longstanding scholarly view that people are largely immune to persuasion (e.g., Berelson, Lazarsfeld, and McPhee 1954; Campbell et al. 1960; Lazarsfeld, Berelson, and Gaudet 1948). Zaller (1992), for instance, argues that researchers have looked for persuasion in the wrong place. He asserts that it is the *moderately* politically aware—not their more highly aware counterparts—who are most amenable to persuasion. These individuals pay enough attention to encounter political messages but not so much that they can effectively counter-argue information that challenges their preexisting beliefs. Those with stronger beliefs are more likely to engage in motivated reasoning; discounting the arguments with which they disagree, while giving undue weight to arguments with which they already agree (Bolsen, Druckman, and Cook 2014; Kunda 1990; Leeper and Slothuus 2014; Taber and Lodge 2006).

Other research has shown that media can be persuasive, at least within controlled experiments. Scholars have begun to identify conditions under which persuasion is more or less likely to occur (Bullock 2011; Dilliplane 2014; Feldman 2011a, 2011b; Levendusky 2013b), as well as under which partisan media may provoke a backlash against the perspective being advocated (Zaller 1992), the outlet presenting the argument (Arceneaux, Johnson, and Murphy 2012; Coe et al. 2008;), or even the media more generally (Ladd 2012).

PROMISE AND LIMITATIONS OF EXISTING RESEARCH

To date, in studying media effects, scholars have struggled to resolve the problem of self-selection underlying the theory of selective exposure. Even controlled experiments designed to identify the effects of partisan media on polarization cannot determine whether real-world-observed differences in attitudes among individuals exposed to different information stem from differences in the information or in the individuals choosing to expose themselves to it.

Typical experiments begin by assessing participants' political partisanship and ideology. Researchers then

follow one of the two different general designs. In the first, *forced-exposure* design, researchers present participants with a single randomly assigned news item and then ask for their opinions on political issues (see e.g., Feldman 2011; Levendusky 2013b). This design allows for causal investigation of persuasion by the single source. Yet, it may produce heterogeneous effects, depending on respondents' underlying preferences about alternative news sources, thereby making it difficult to generalize to real-world polarization.

In the second, far less common, *free-choice* design, researchers present respondents with multiple news items from which to choose and then ask for their opinions on issues over which partisans usually disagree (see, e.g., Ellithorpe, Holbert, and Palmer-Wackerly 2013; Knobloch-Westerwick and Kleinman 2012; Knobloch-Westerwick and Meng 2009). This allows investigators to determine the extent to which participants choose sources compatible with their preexisting beliefs, as well as to measure differences in attitudes between participants who consume different news sources. However, the lack of controlled randomization prevents researchers from identifying a media persuasion effect because the design does not account for selection bias.

Recent research incorporates individual preferences when estimating the persuasive power of the media. Arceneaux and Johnson (2013) and Arceneaux, Johnson, and Murphy (2012) incorporate the elements of both designs in a single study following recommendations from Gaines and Kuklinski (2011). In their “selective exposure experiment,” Arceneaux, Johnson, and Murphy (2012) randomly treat half of their participants with a liberal, conservative, or entertainment news story and then observe the effects of that treatment on subsequent attitudes (the forced-choice component). They allow the other half of their participants to choose any one of the same three program choices (the free-choice component). They then compare the observed effects on attitudes toward the media (Arceneaux, Johnson, and Murphy 2012) and policy opinions (Arceneaux and Johnson 2013). Because this combined design does not capture common information on media preferences from both the free-choice and the forced exposure groups of respondents, they cannot estimate the group-specific effects of media exposure, despite randomization between forced exposure and free choice. Instead, they analyze the free-choice group of respondents as a single-treatment group, limiting the conclusions they can draw.

In a separate experimental design, Arceneaux and his colleagues (Arceneaux, Johnson, and Murphy 2012, 183; Arceneaux and Johnson 2013, 85) incorporate the measurement of media preferences prior to the forced-exposure procedure. These “patient preference designs,” or “participant preference experiments,” measure respondents' preferences for consuming media, but do not incorporate choice. Levendusky (2013a, 2013b) employs a similar design and separately estimates persuasion among individuals who prefer like-minded and oppositional partisan media. Assuming respondents' stated media preferences match their actual media choices, this design can provide a causal estimate of the

effects of media exposure conditional on media choice.² However, people's preferences may differ from their actual choices when given media options beyond just news. Moreover, given the difficulty individuals have in accurately recalling their own past media consumption (Prior 2007, 2013), it is also important to account for the potential discrepancy between self-reported and real-world viewing behavior. Revising their design, as delineated below, allows us to better understand and measure whether and to what extent estimates of the persuasive effects of partisan media depend on untestable assumptions about survey self-reports.

DESIGN

Our study resolves the trade-off between the reliable measurement of preferences offered by selective exposure studies, on the one hand, and the identification of persuasion effects among subgroups of the population that do or do not consume a given media option, on the other. We do so by using our new PICA design (Figure 1). In this design, we randomly assign participants to either a forced exposure or free choice treatment condition. Each of the two conditions individually resembles procedures separately employed by Arceneaux and his colleagues (Arceneaux and Johnson 2013; Arceneaux, Johnson, and Murphy 2012; Gaines and Kuklinski 2011; Levendusky 2013a, 2013b). But the PICA design combines them into a single framework. Participants in both conditions read news reports from either Fox News, MSNBC, or an entertainment network (the Food Network). We derive all reports from real online news stories, which we edited to equalize length and framing.³ For the partisan media treatments, the articles discussed either the economic, social, safety, or public health effects of legalizing marijuana.⁴ On each topic, the Fox News and MSNBC articles were nearly identical except for text that was either supportive of

(MSNBC) or opposed to (Fox) marijuana legalization. The Food Network articles discussed how to save money while grocery shopping, tips for buying meat, and how grocery stores might change in the future.

The PICA design builds upon the general framework known in the medical literature as the patient preference trial (PPT). As noted, some existing political science research employs PPT frameworks, incorporating participants' preferences over treatment options (Torgerson and Sibbald 1998) to examine the effects of experimental treatments that may vary depending on those preferences (Gaines and Kuklinski 2011; Leeper 2017). This step makes it possible to avoid relying on self-reported ideology to identify people's media preferences. Unlike these existing proposals, however, our design also incorporates the stated preferences over treatment options for *all* respondents, enabling us to combine information from the forced exposure and the free choice arms for more nuanced inference. The PICA design thus unifies key elements from Arceneaux and Johnson's two designs as well as the previous two-arm PPT designs.

The overall advantage of the PICA design is therefore twofold. First, the randomization of participants into either the forced exposure or the free choice conditions enables us to draw inferences about the persuasion effects of pro-attitudinal and counter-attitudinal political news among participants who if given a choice would consume those media, as well as among those who would choose nonpolitical media. It also allows us to distinguish between consumers of liberal and conservative partisan media, rather than combining these two groups and estimating the combined net effects of pro- and counter-attitudinal consumption (as Arceneaux and Johnson 2013 do). Second, unlike previous studies, our design measures both stated media preferences and actual selective exposure to the media options. For subjects who are assigned to the free choice condition, we can empirically measure any discrepancy between their stated and actual media preferences. This allows us to undertake sensitivity analyses for the estimated persuasion effect of particular media among subjects who prefer particular partisan news and those who prefer non-political media.⁵

We administered the experiment via an Internet-based survey to a national sample of 7,298 respondents through Survey Sampling International (SSI).⁶ This large sample allows us sufficient power to detect the persuasive effects of media among different groups.

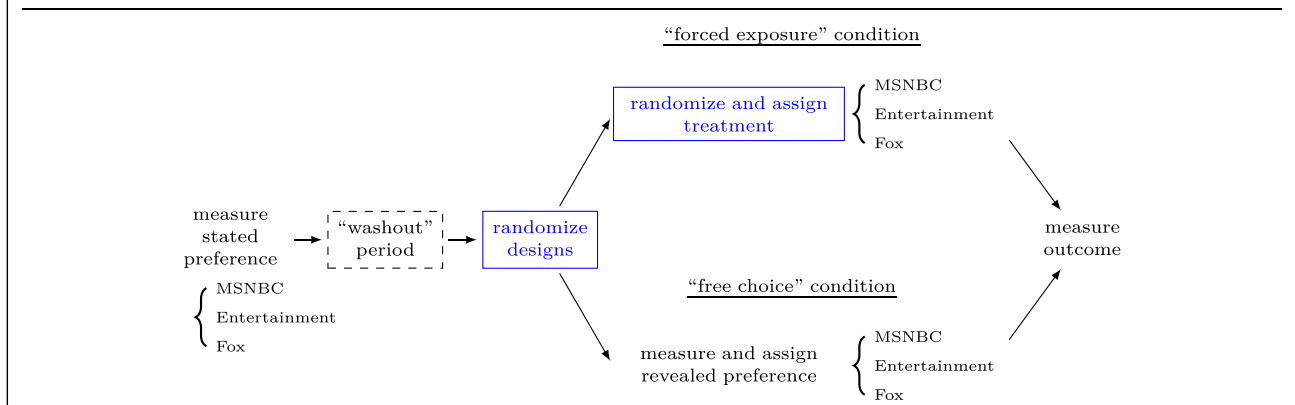
² This is in addition to the assumption that media preferences can be proxied by self-reported ideology. Arceneaux, Johnson, and Murphy (2012) expose research subjects to partisan media and disaggregate results based on the match (or mismatch) between the respondent's ideology and the leanings of the media that they consume. This assumes that experimental subjects accurately report their ideological predispositions in laboratory studies and that ideological predispositions determine real-world media choices. Some research (in the former case, see, e.g., Zell and Bernstein 2014) calls both assumptions, but especially the latter linkage, into question.

³ We pretested all news articles used in the experiment and asked pretest respondents for their perceptions of the ideological direction of the articles, as well as respondents' understanding of the issue after reading the articles. We found that respondents perceived a large degree of divergence in the ideological leanings of the articles but did not have significantly different understanding of the issues after reading articles from opposite sides of the partisan news spectrum. We present the full text of all articles in Online Appendix C.

⁴ We chose marijuana legalization as the policy subject of the news articles due to the relatively small partisan gap in support of the issue, with partisan elites from both parties expressing a mix of support for and opposition to legalization. This may limit respondents' motivated reasoning when consuming information about the topic (Bolsen, Druckman, and Cook 2014). However, we also find patterns largely similar to our main results for three additional policy areas (Online Appendix E).

⁵ Importantly, Gaines and Kuklinski (2011) show that one can identify the treatment effect conditional on actual selective exposure without any additional information—such as stated treatment preferences—if there are only two treatment options (see Knox et al. 2019, for a formal discussion). Our study, however, requires at least three media treatment conditions: pro-attitudinal and counter-attitudinal political media and non-political media. The PICA design addresses the identification problem by use of stated preferences.

⁶ SSI recruits participants through various online communities, social networks, and Web site ads. When deploying a particular survey, SSI randomly selects participants for survey invitations. We asked SSI to recruit a target population that matched the (18 and over) census population on education, gender, age, geography, and income. The result is a diverse national sample, albeit not a probability sample.

FIGURE 1. Preference-Incorporating Choice and Assignment (PICA) Experimental Design

We follow the protocol proposed in Knox et al. (2019) and summarized in Figure 1. First, early in the survey we measure the stated preference of all respondents over the three options. Specifically, we asked: “If you were given the choice of news articles from the following three sources to read, which of the three would you choose?” We presented each choice with an accompanying logo of the network, while randomizing the order in which the options appeared on the screen.

This choice environment is obviously not entirely realistic. The Internet provides access to thousands of different news sources, and, in fact, the average American visits 89 different web domains a month (Nielsen 2010). But to make the experiment tractable and to gain some control over the content of the programs, it was necessary to conduct it in a stylized environment.⁷ Doing so could, admittedly, introduce external validity concerns, as people’s preferences may exhibit more consistency in our experiment than in reality. People’s preference for partisan news rather than an entertainment option may change if there are more entertainment options. But our sensitivity analyses, which identify the extent to which estimates of partisan media persuasion could be affected by discrepancies between the stated and actual preferences, can partially mitigate such concerns by providing conservative estimates of media persuasion.⁸

We then move to a “washout” period, where we asked participants to answer demographic questions not directly related to the media choice, and to complete distraction tasks. The goal was to minimize the possibility that measuring stated preferences might contaminate respondents’ voluntary choices of stories to read in the free choice condition. These distraction tasks enhance external validity by allowing for the instability of preferences over time.

⁷ In pretesting of our experiment and in one replication reported in Online Appendix E, we tried offering participants multiple entertainment options, but found that such a large majority of people selected entertainment media that continuing to do so would have required a dramatic increase in our sample size in order to gain sufficient statistical leverage, thereby rendering the study cost prohibitive.

⁸ Of course, an entirely separate issue may also arise from the potential for consumers’ revealed choices to be different over time (e.g., Sood and Lelkes 2018).

Next, we randomized subjects with equal probability into the forced exposure and free choice conditions. We then randomly assigned those in the forced choice arm to read either the Fox News, MSNBC, or entertainment story, each with probability 1/3. For those in the free choice arm, we instead asked, “Which of these three articles would you like to read now?” We presented these participants with the same three options. Subjects in the free choice arm then read the single story they chose. Finally, we asked a series of questions measuring subjects’ opinions about, as well as their desire to share or learn more about, marijuana policy, as described in more detail in Results section below.

METHODOLOGY

We estimate the degree of persuasion using the average choice-specific treatment effect (ACTE) of the media stories, following Knox et al. (2019).⁹ In the current study, the ACTEs of interest constitute the average effects of exposure to one story relative to another among those who would prefer to read Fox News, MSNBC, or the entertainment article. We focus here on the difference between the two partisan outlets, which most directly estimates the degree of political polarization due to biased partisan media, and in particular the effect of these specific articles on opinions about marijuana legalization.¹⁰ The relationships we report below thus represent the causal effects of treating

⁹ As described in Knox et al. (2019), the ACTE represents the average causal effect of one treatment versus another treatment among those participants who would choose a treatment option—be it the first, second, or another treatment—if given the opportunity to choose. This contrasts with estimation of treatment effects conditional on ideology or partisanship of the viewer, or the match between ideology and the ideological orientation associated with a given treatment option, which are the conditional average treatment effects calculated by Arceneaux, Johnson, and Murphy (2012) and others. We compute alternative treatment effects by partisanship and ideology in Tables B.3 and B.4 in the Online Appendix.

¹⁰ Because our experiment consists of three treatment options, we can estimate three different contrasts for each of these choice subgroups: the comparison between each of the two partisan media options and the entertainment media, and the comparison between the two partisan media options.

individuals with Fox News as opposed to MSNBC, within each choice-specific group.

We take two approaches for the inferences about our ACTEs of interest. First, we use the stated media preferences that we measure at the beginning of the experiment as an approximate measure of actual media choices and estimate the “naïve” ACTEs as the difference between the average opinions among subjects assigned to Fox News and the average among subjects assigned to MSNBC within the forced exposure condition.¹¹ However, our design also allows us to go beyond this naïve estimate. In particular, our two-arm PICA design allows us to use additional information from subjects assigned to the free choice condition to conduct sensitivity analyses via the nonparametric bounds proposed by Knox et al. (2019) and estimate the degree to which the divergence of stated and actual media preferences can bias our naïve estimates of the ACTEs.

Conceptually, these sensitivity analyses involve two steps. First, we can assume the extreme (and unlikely) scenario where respondents whose actual media choices differed from their stated choices could hold *any* opinion about the issue. Under this agnostic assumption, we can calculate the maximum and minimum possible values of the ACTEs for each of the outcome measures. The resulting values represent our “no-assumption” upper and lower bounds for the possible true value of the ACTE, that is, the interval within which we can be confident that the ACTE lies. In the second step, we explore a more plausible “middle ground” by modifying the agnostic assumption to fit a more realistic scenario. We continue to assume that there is a discrepancy between the stated and actual choice groups in their response to a treatment, but we hypothesize the size of the discrepancy to be less than a certain threshold. Following Knox et al. (2019), we call this hypothesized upper bound on the discrepancy the *sensitivity parameter* and denote it by ρ .¹² We then derive the bounds on the ACTE using a procedure

similar to the first step for varying levels of the sensitivity parameter. These bounds represent the largest and the smallest that our treatment effect estimates could be, under the restrictions imposed by a given level of ρ . Finally, we assess the threshold value of ρ at which the bounds contain zero and cease to be informative about the sign of the ACTE. A high value of this threshold ρ value would imply that our conclusion based on the naïve estimate is robust to the possible bias resulting from a relatively large discrepancy between the stated and actual choices.

RESULTS

We first examine evidence for selective exposure by focusing on subjects assigned to the free choice condition (i.e., the bottom portion of Figure 1) and analyzing their revealed media preferences and reported opinions. Our results indicate that polarization in the electorate does indeed correspond with polarization in media consumption. Across all categories of partisanship, 30% of people chose to read the Food Network article, 41% chose Fox News, and 28% chose MSNBC. However, these choices varied significantly when we looked at this same breakdown by party. Among Democrats, 31% preferred the entertainment option, 24% preferred Fox News, and 45% preferred MSNBC. Among Republicans, 28% preferred entertainment, 61% Fox News, and only 12% MSNBC. Among those with no preference for either party, 42% preferred entertainment, 34% Fox News, and 23% MSNBC.¹³ These results demonstrate a strong separation in choices between the ideological left and right: Democrats were 33 percentage points more likely than Republicans to prefer left-leaning MSNBC, while Republicans were 37 percentage points more likely than Democrats to prefer right-leaning Fox News.¹⁴

Polarization by Selective Exposure

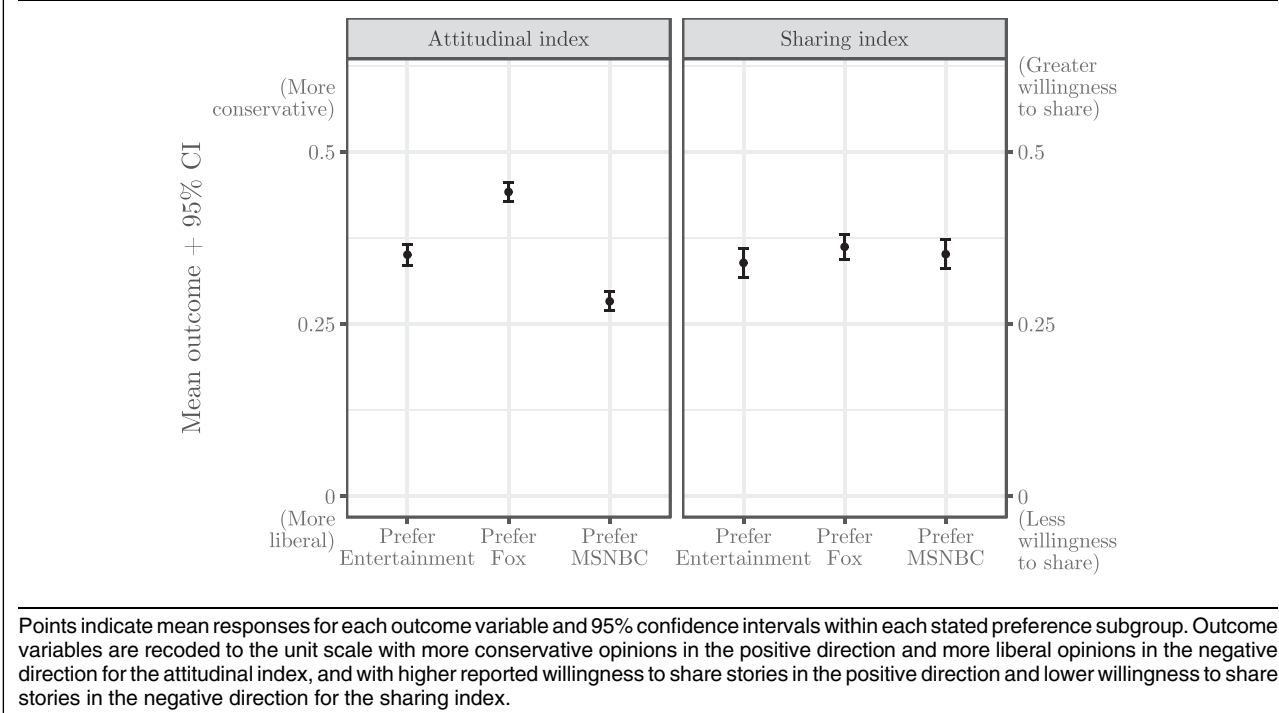
We also find that selective exposure corresponds with differences in political opinions and behaviors. To measure political opinions, we asked respondents ten questions about marijuana and drug policy. We asked them whether they agreed or disagreed with the following statements, “The legalization of marijuana leads to fewer people using more serious drugs, such as heroin and cocaine” and “Marijuana use increases violent crime,” corresponding to the articles about public safety and health consequences of legalization. Similarly, we asked respondents whether legalization would make the economy better or worse, which corresponded directly to

¹¹ As discussed in Knox et al. (2019), these represent naïve estimates of the ACTEs and will be biased unless the deviation between respondents’ stated media preferences and their actual choices are simply stochastic noise. In fact, even stochastic noise has pernicious consequences for inference if we are interested in both the ACTEs per se, and in *differences* among them. Specifically, classical measurement error in the stated preferences will cause the difference between two naïve ACTE estimates to be a downward biased estimate of the difference in the true ACTEs. This immediately follows from the well-known connection between measurement errors and attenuation bias. However, these quantities are a convenient approximation, which we can estimate using the data from the forced exposure arm of our experiment alone, as do many previous experimental studies of media persuasion effects, such as Arceneaux, Johnson, and Murphy (2012).

¹² This assumption is more reasonable than the one underlying the naïve estimator, but, unlike the no-assumption bounds, allows the stated preference to be informative with respect to the actual choice to the extent permitted by the given value of ρ . More precisely, ρ represents the amount of deviation in potential outcomes between stated preferences and actual media choice in units of the outcome measure. For instance, if $\rho = 0.10$ for our dependent variables, then the deviation in potential outcomes caused by the instability in preferences is approximately ten percent of the range of the outcome variable. See Knox et al. (2019) for more detailed discussion.

¹³ We classify partisans as those who strongly or not so strongly identify with a political party, as well as those who lean toward one of the parties. Independents are those who indicate no party preference on either an initial question or a follow-up question asking whether they lean toward one party or the other.

¹⁴ People who preferred different media options also varied across a range of demographic characteristics (see Table A.1 in the Online Appendix).

FIGURE 2. Average Responses in Free Choice Condition

the interventions about the tax and economic implications of legalization.¹⁵ Other questions addressed feelings regarding marijuana more broadly. Respondents placed themselves on 7-point scales between strongly agreeing or strongly disagreeing with the phrases: “Government efforts to enforce marijuana laws cost more than they are worth,” “Using marijuana is morally wrong,” “Marijuana should be legal for medical use,” “Marijuana use is a serious problem today,” and “Marijuana should be legal for recreational use.” Finally, respondents placed themselves along another seven-point scale between, at one end, whether habitual drug use should be thought of as a criminal offense or, at the other, a medical problem.¹⁶ They then placed several substances, including marijuana, on scales from very dangerous to very safe. We recode all opinion variables to 0–1 intervals, with the most liberal or permissive of legalization at 0 and the most conservative or opposing of legalization at 1. To capture variation in the underlying latent attitudes and beliefs about marijuana, we form an additive index of these ten opinion questions. We use this

¹⁵ Exact wording was, “If the sale and possession of marijuana were made legal, do you think it would make the economy better, make the economy worse, or have no effect on the economy?”

¹⁶ Exact wording was, “Some people feel that habitual drug use should generally be considered a criminal offense and dealt with through the courts and criminal justice system. Suppose these people are on one end of the scale, at point 1. Others think that habitual drug use should generally be considered a substance abuse and addiction problem and dealt with through the medical and mental health systems. Suppose these people are at the other end, at point 7. And of course, some other people have opinions somewhere in between. Where would you place YOURSELF on this scale?”

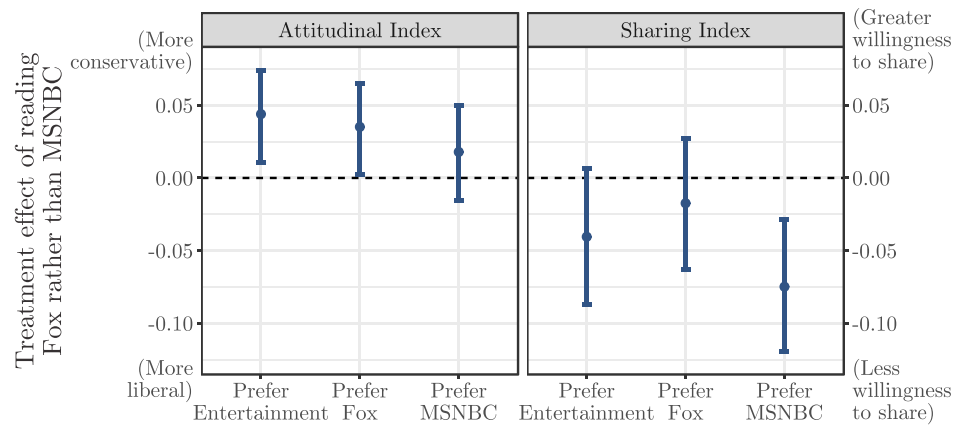
index as the primary dependent variable on which we assess polarization and persuasion.¹⁷

We also asked respondents to indicate their likelihood of forwarding, discussing, posting to social media, or seeking out additional information on the story they had just read. We combined these four measures into an additive index of sharing behavior, which ranges from 0 (least likely to share) to 1 (most likely to share).

Political opinions and behavior differed across the three groups of respondents who self-selected into different media options. Figure 2 shows the average responses among the respondents assigned to the free choice arm (i.e., the bottom portion of Figure 1), separated along the x-axis by media preference group.¹⁸ We plot our attitudinal and sharing indexes in the left- and right-hand panels, respectively. On the left side of each panel, we show the responses among those who preferred to (and did) read the entertainment article. People who both preferred and chose Fox News (in the middle of each panel) reported opinions that were 0.16 greater along the 0–1 scale of our attitudinal index than those who both preferred and chose MSNBC (on the right of each panel). This difference is equivalent to 0.72 standard deviations

¹⁷ The scale is internally consistent, with an alpha of 0.89. We present the full results on each of these individual measures in Online Appendix A (for free-choice respondents) and Online Appendix B (for forced-choice respondents).

¹⁸ We present these estimates of opinion among free choice respondents in tabular format in Online Appendix A. These average opinions also match those of respondents in the forced choice arm of our experiment who were randomly assigned to read the media for which they indicated a preference. We show the comparison between these groups in Online Appendix Table A.4.

FIGURE 3. ACTE Estimates Based on Stated Media Preferences

Points indicate treatment effects of watching Fox rather than MSNBC, which are estimated as the difference in means between these two groups in the forced exposure arm of the experiment.

for this outcome measure. On our sharing index, those who preferred and read Fox reported intended behaviors that were 0.01 greater than those who preferred and read MSNBC—a substantively and statistically insignificant difference. Together these free-choice results show that people who consume different media report different political attitudes but similar behaviors.

Partisan Media's Persuasive Effects

The pattern described above reflects only the descriptive differences between groups of respondents who self-selected into each media treatment. This conflates self-selection and persuasion and cannot be interpreted as polarization necessarily *caused* by partisan media sources.

The forced choice arm of our experiment (the top portion of Figure 1)—where we randomly assign respondents to partisan media or entertainment—accounts for the persuasive effect of media exposure. To measure changes in political opinions, we disaggregate these respondents by the type of media that they would choose if given a choice.¹⁹ The results demonstrate substantial added explanatory value over simpler comparisons more commonly employed in the literature, such as that between partisan groups.²⁰ Here, we present our treatment effects on the aforementioned 10-question additive index, to assess the effect of forced exposure to Fox News rather than MSNBC.²¹

Partisan media have a strong impact on respondents' opinions in the forced-choice arm. In Figure 3, we show the estimated effects of treating respondents with Fox News rather than MSNBC, conditional on media consumption preferences, and based solely on the

respondents assigned to the forced choice condition and their stated media preference.²² The left-hand panel plots the treatment effects with 95% confidence intervals for our attitudinal index along the y-axis, separated by media preferences along the x-axis. Effects above zero represent persuasion in the conservative direction on the 10-attitude-question index, while effects below zero represent persuasion in the liberal direction.

Among respondents who would prefer to read entertainment rather than news—shown on the left sides of both panels in Figure 3—reading Fox instead of MSNBC yields statistically significant movement in the conservative direction on the attitudinal index. Along the 0–1 scale, reading Fox causes these entertainment-preferring respondents to report opinions that are four percentage points more conservative than if they had read MSNBC. This demonstrates the substantively large persuasive effects that partisan media may have on attitudes of those who would not ordinarily consume these news sources.

Among respondents who would choose Fox if given the choice—plotted in the center sections of both panels—the persuasive effect is similar to that for entertainment consumers. Respondents who prefer Fox reported opinions that are three and a half percentage points more conservative (on the 0 to 1 scale) if they read Fox than if they read MSNBC. This effect is statistically significant. So, exposure to more partisan news can sway the opinions of even those respondents who are already exposed to some degree of partisan news.

Among respondents who prefer MSNBC—plotted as the rightmost line within each of the panels in Figure 3—the persuasive effect of reading Fox rather than MSNBC was similar (in a conservative direction) but smaller and statistically insignificant.

Together, these results show that partisan news is persuasive, albeit not for all consumers. In other words, partisan media can move the attitudes of respondents exposed to a given message—both those who choose it

¹⁹ We disaggregate by stated preferences here, which are equivalent to the naive estimates for the ACTE.

²⁰ However, we also conduct the comparison of treatment effects by reported partisanship and ideology in Online Appendix B. These results are largely similar to our primary results.

²¹ Full results for all ten attitudinal measures are in Online Appendix B.

²² We show these results numerically in Table B.2 of Online Appendix B.

and those who would ordinarily not consume political news—in the direction the message advocated. The effects were most consistent among those who preferred entertainment. Among these readers, we found effects in the expected direction for every individual attitudinal question as well as a statistically significant effect on the aggregate index. Among readers who preferred Fox News the effects remained fairly strong and were also statistically significant for our attitudinal index. Finally, for those who preferred MSNBC the effects remained in the expected direction but were somewhat smaller in magnitude and statistically insignificant.

We also found that partisan news has different effects across different subgroups of readers. That is, the media preferences of individuals condition the media's persuasive effects. For individuals who preferred entertainment, partisan media powerfully polarized attitudes. Contingent on exposure, these individuals appear to be most subject to influence by partisan media. However, we also found persuasion among respondents who prefer partisan news. This suggests that partisan media can be persuasive, though perhaps less so because the opinions of their consumers tend to be more solidified and so more resistant to persuasion.

The right-hand panel of [Figure 3](#) turns from attitudes to behavior, by presenting the treatment effects on our sharing index, with positive effects representing a greater intent to share the news story. Across all subgroups, reading Fox News rather than MSNBC had a negative effect on sharing. Respondents were more likely to report a willingness to share the articles after reading MSNBC than after reading Fox News. Reading the Fox News article instead of the MSNBC story reduced respondents' reported willingness to engage in sharing behaviors by between 0.02 and 0.07 along the unit scale for our three subgroups of media consumers. This effect was statistically significant for readers that preferred MSNBC, but not for those that preferred entertainment or Fox News. This indicates that, among those people who preferred MSNBC, exposure to pro-attitudinal media increased their propensity to share its content relative to exposure to counter-attitudinal media.

Overall, the persuasive effect of partisan media is strong across a number of political outcome variables. While other research has suggested that partisan media can induce a backlash among individuals who hold opposing preferences, we find that even counter-attitudinal partisan media can be persuasive. Moreover, such effects are not limited to attitudes. They also appear to affect socially significant political behaviors, such as sharing content. Those who *preferred* either MSNBC or entertainment media were more likely to report an intent to share content after actually consuming MSNBC than after actually consuming Fox News. Our experimental procedure enables us to detect where these persuasive effects occur across different groups of media consumers.²³

²³ Note that these subgroup differences do not represent *causal* differences—i.e., people's media preferences causing differential persuasion—but rather *descriptive* moderation of our main treatment effect by the stated preferences of media consumers.

Polarization from Both Selective Exposure and Persuasion

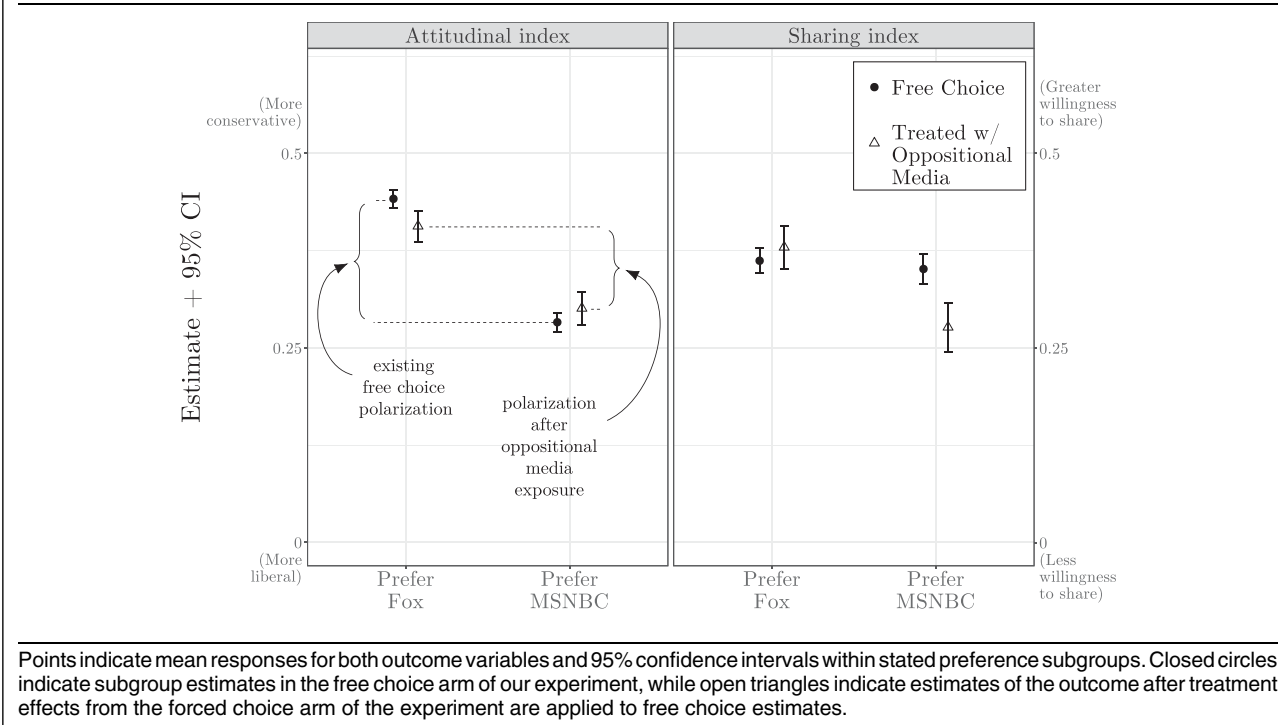
The results thus far do not directly tell us how the persuasive effects of partisan media might bring the attitudes of their consumers closer together in the real world. To put these results in context, we leverage evidence from both the forced exposure and the free choice conditions of our PICA design and simulate how our estimated persuasion effects might change the existing levels of partisan polarization. As a baseline for “real-world” levels of the existing polarization, we first use responses from subjects in the free choice condition to measure the existing differences in opinions between the groups that chose opposing partisan media. Then, using the ACTE estimates obtained from the same media-preference group in the forced choice condition, we see whether a certain amount of “treatment” might at least partially bridge (or widen) these gaps—that is, by hypothetically forcing those individuals to consume opposing media. Specifically, our multi-arm design allows us a window into this hypothetical world by comparing the existing attitude polarization to the situation where we expose individuals to the persuasive power of media from the other side.

[Figure 4](#) presents the gaps in opinion between subgroups in the free-choice arm of our experiment along with the estimates of their opinions had we instead treated them with counter-attitudinal media. We plot the point estimates and corresponding confidence intervals for respondents who read their freely chosen media with filled black circles, and estimates of their opinions if treated with the opposing media with open triangles. We again present the attitudinal and sharing indexes in the left- and right-hand panels, respectively, and separate the subgroups of media consumers along the *x*-axis.

Our experiment indicates that forcing people to read news from the “other side” generally moderates the opinions of partisan media consumers, thereby reducing polarization. On our attitudinal index, treatment from opposing media reduced the opinion gap between respondents who preferred Fox versus MSNBC. We can apply the statistically significant treatment effect of 0.035 among those who prefer Fox News and the smaller and insignificant effect of 0.018 among those who prefer MSNBC to the average outcomes among those two groups in the free-choice arm. Treating respondents that prefer Fox News with the MSNBC article and vice versa would reduce the opinion gap in the free choice arm by about one third—from 0.16 to 0.10 (along the 0–1 scale). This reduction is statistically significant, with a 95% confidence interval of 0.019 to 0.087.²⁴ Applying the treatment effects of oppositional media estimated in the forced-choice arm reduces the existing polarization between the opinions of partisan media consumers in the free-choice arm of our experiment.

Our experiment also yields evidence that consuming counter-attitudinal media may polarize consumers' subsequent behavior. While partisan media consumers may moderate their opinions when exposed to media

²⁴ Confidence intervals calculated based on pooled standard errors as the square root of the sum of squared individual group standard errors.

FIGURE 4. The Effect of Opposing Media Treatment on Polarization

from across the partisan aisle, they appear to respond differently in their sharing behavior. Among people who prefer Fox, exposure to MSNBC made them no more or less likely to share news content than those who chose Fox News. However, among those who preferred MSNBC, exposure to Fox News reduced their willingness to share news content. This means that, were both groups exposed to content from across the aisle, the subsequent sharing of this media would be higher among those who prefer Fox—leading to potential counter-attitudinal consumption of this content by other people in their lives—relative to those who prefer MSNBC. This demonstrates that the immediate short-term behavioral effects of counter-attitudinal partisan media that we find may lead to even larger net effects were they to occur in a situation allowing for content sharing among people’s friends and family. Moreover, the potential for these reinforcing effects may be asymmetrical, primarily affecting those people who prefer MSNBC rather than Fox.

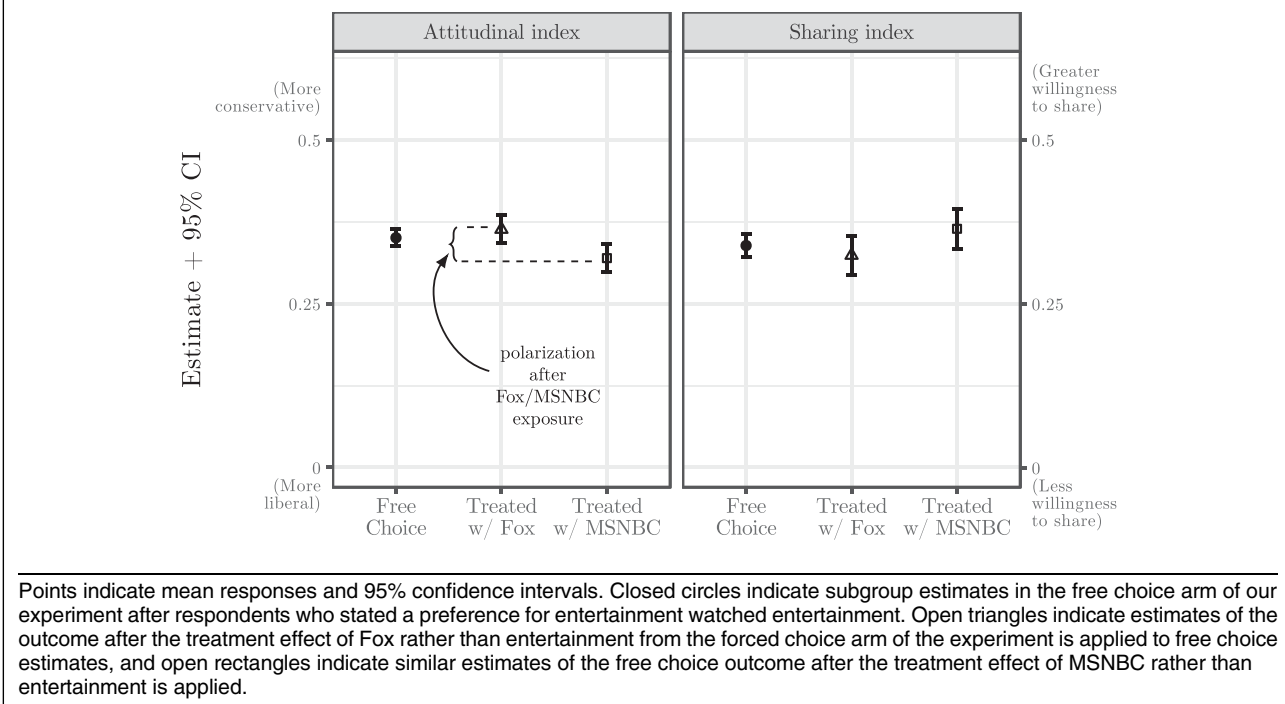
Polarization Among the Inadvertently Exposed

Perhaps the most worrisome effects of partisan media may befall those who would not choose to consume these stories. These inadvertent consumers—the group of citizens who otherwise would prefer entertainment—may change their opinions as a result of chance encounters with partisan media. This could have pernicious consequences for political polarization as partisan media become increasingly prevalent. To show how exposing these individuals to partisan media would change their opinions, we again use both the forced exposure and free choice

conditions to assess how our estimated persuasion effects change the existing opinions. We first measure the baseline of opinions among those who prefer entertainment in the free choice group. We then use our ACTE estimates obtained from the forced choice group to simulate how these opinions would change after “treatment”—that is, by hypothetically exposing entertainment-preferring individuals to partisan media rather than entertainment. As with the groups who preferred partisan media, our experiment allows us to compare their existing opinions with the scenario in which we expose them to partisan media.

Figure 5 presents opinions among those who prefer entertainment in the free-choice arm of our experiment along with the estimates of their opinions had we instead treated them with partisan media. We plot with filled black circles the point estimates and corresponding confidence intervals for respondents who read their freely chosen entertainment story, while plotting with open triangles and open squares the estimates of their opinions if treated with Fox News and MSNBC, respectively. We again present the attitudinal and sharing indexes in the left- and right-hand panels, respectively.

We find that, among inadvertent readers, partisan media can substantially polarize political opinions, driving readers in opposite directions. On our attitudinal index, people who preferred to read entertainment and did so had a mean response of 0.35 along the 0–1 scale in the free choice arm. Applying the treatment effect from the forced choice experiment—equivalent to exposing them to Fox News rather than entertainment—their estimated response would instead be a more conservative 0.36 on average, while if we apply

FIGURE 5. The Effect of Media Treatment on Entertainment-Readers

the treatment effect of MSNBC, their estimated response would be 0.32. This resulting opinion gap of 0.04, as plotted in the left panel of Figure 5 with the open triangle and open square, is statistically significant (with a 95% confidence interval of 0.014 to 0.074). In substantive terms, this level of polarization is approximately one-third of the existing polarization on this index between respondents preferring MSNBC versus Fox in the free choice arm. This demonstrates the polarizing power of just one instance of exposure to partisan media among this segment of the population.

Partisan media may also affect the behavior of these inadvertent partisan media readers. In the right panel of Figure 5, we plot the estimated response on our sharing index among people who prefer entertainment and select it in the free choice arm, as well as their estimated responses if they instead read Fox News or MSNBC. In the free choice arm, these respondents had an average response of 0.34. Had they read MSNBC rather than entertainment, however, they would be even more likely to share the story they read, with an estimated response of 0.36. On the other hand, had they read Fox, they would be less likely to share the story, with an estimated response of 0.32. This estimated gap in sharing behavior of 0.04 (95% confidence interval: $-0.002, 0.083$) demonstrates how the effect of partisan news may change the way that people who prefer entertainment engage with or share the partisan news media that they encounter.

SENSITIVITY ANALYSES

As previously discussed, the ACTE estimates we have presented thus far rely on the rather strong assumption

that the discrepancy between participants' stated preferences over media options and their actual media choice is "ignorable;" that is, unsystematic and unrelated to the way they react to different news media. In this section, we present results of our sensitivity analysis, which relaxes this assumption and assesses how much these estimates of choice-specific persuasion are susceptible to the violation of the assumption that people's reported media preferences match their behavior when choosing media. This analysis requires information from both the forced exposure and the free choice conditions, so our PICA design is essential for assessing such a problem.

The results from the free-choice condition of our experiment show that those subjects who state a preference for a given media option do not always choose that option when given the chance to do so.²⁵ Thanks to random assignment, we know that in expectation the same proportion of those subjects in the forced-exposure condition would have deviated from their

²⁵ We find that 81% of those preferring entertainment actually chose the entertainment article, 84% of those preferring Fox chose that article, and 82% of those preferring MSNBC chose that article. Thus, the actual viewing choice of between 16 and 18% of respondents differed from their stated media preferences. Full tabular display of the match between stated preferences and actual choices from respondents in the free-choice arm of our experiment are in Table A.2 of Online Appendix A. Even with a washout period between our stated preference question and respondents' actual choices, this estimate of the discrepancy between stated preferences and actual choices may be a lower bound due to the limited realism of the survey experiment. This illustrates the need for more analyses estimating the sensitivity of such experimental results to further discrepancies between stated preferences and choices.

stated media preferences if they had been allowed to do so, contrary to our “naïve” assumption. Thus, we begin our sensitivity analysis by asking what values the ACTEs could take if we removed that assumption. That is, if we made *no assumption* about the opinion measures and sharing index for those subjects whose stated and actual media preferences would differ, what could we still conclude about the ACTEs?

We answer this question by calculating the non-parametric bounds on the ACTEs for each dependent variable. In Figure 6, we present the resulting bounds on the far right-hand side of each panel (thick lines) along with their 95% confidence intervals (thin lines). The top and bottom rows report results for our attitudinal and sharing indexes, respectively, while each column disaggregates respondents by their media preferences. These bounds represent a “worst-case” scenario, in that they present the effect of reading a given article when we assume that responses under *stated* preferences are entirely uninformative of responses under *actual* preferences. The estimated bounds suggest that this extreme assumption renders all of our persuasion effects statistically insignificant.

Although these widest bounds represent our most assumption-free estimates, they are also almost certainly too conservative. After all, it is unrealistic to expect that opinions of those individuals who state a preference for (say) Fox News are completely unrelated to the opinions of those who prefer to actually read Fox News. Thus, the next step in our sensitivity analysis is to allow those two groups to have different opinions not arbitrarily but only to a certain degree, which we specify via the sensitivity parameter ρ . That is, if the difference in the average opinions between the stated-preference groups and the actual-choice groups were at most ρ , what can we conclude about the possible values of the ACTEs?

In Figure 6, we also plot the sensitivity bounds for the ACTEs (in dark gray) and corresponding confidence intervals (in light gray) for the summary index of our attitudinal measures (in the top row) and the behavioral index (in the bottom row) as we change the sensitivity parameter ρ between 0 and the value at which it converges with the no-assumption bounds, moving along the x-axis of each panel. As ρ approaches zero, the bounds become tighter (because we allow less divergence of opinions between the stated-preference and actual-choice groups) and eventually collapse to the naïve estimates (on the left). On the right, all the sensitivity bounds coincide with the no-assumption bounds. Our primary interest is in the value on the x-axis at which each effect’s bounds cross zero, which is the value of ρ at which our estimated effect ceases to be informative about the true sign of the ACTE.

We focus first on the statistically significant naïve estimates for the attitudinal index, in the top row of the figure, beginning with those who prefer entertainment (left panel). For this index, the value of ρ at which our bounds around the treatment effect cross zero is estimated to be 0.02. This implies that among people who prefer entertainment, in order to have zero estimated effect of reading Fox News rather than MSNBC, the

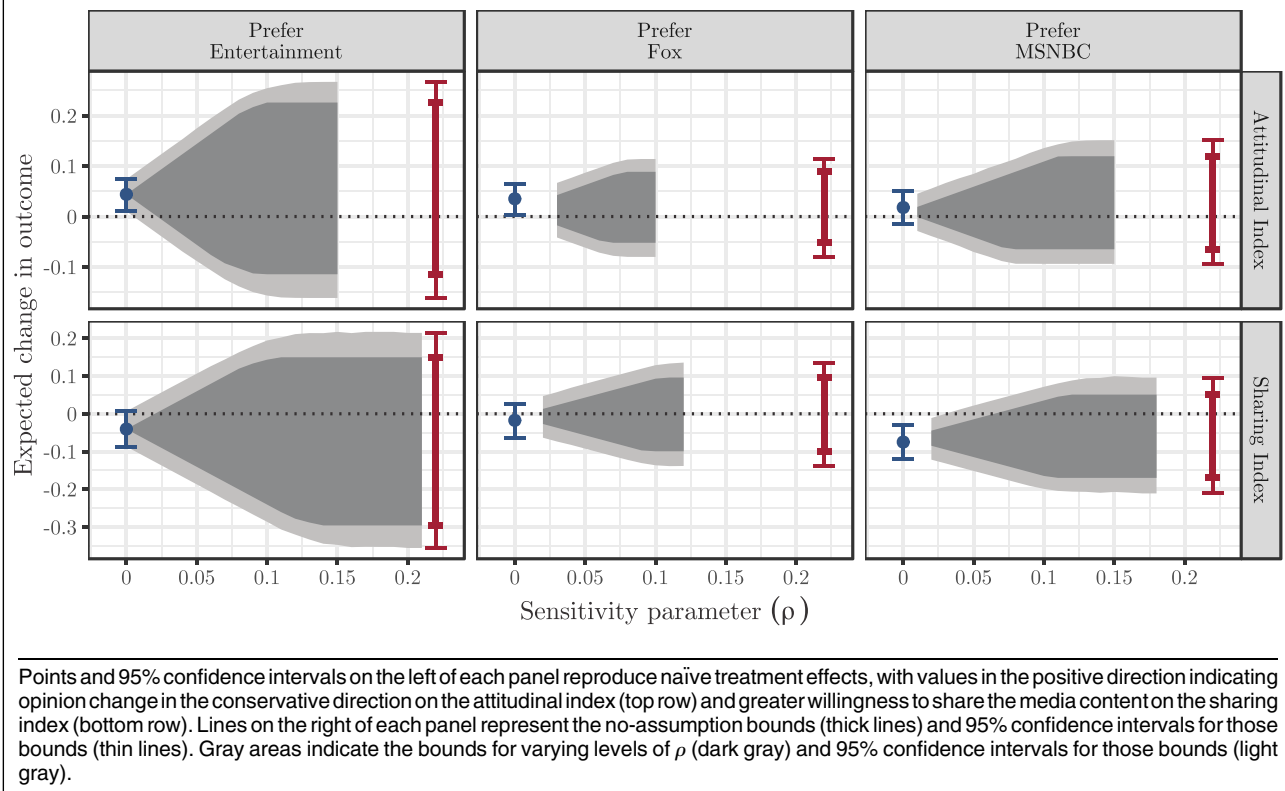
deviation in the measure between stated-preference groups and actual-choice groups would have to be just below half the size of our naïve estimate of 0.04. After incorporating statistical uncertainty into the estimation of these bounds, however, the estimates appear less robust to the discrepancy in the stated and actual preferences: The 95% confidence intervals for the bounds include zero as soon as we allow for a deviation of 0.01 in the average opinions between the stated-preference and actual-choice groups—one quarter the size of our estimated treatment effect. Among those respondents who prefer Fox, for whom we observed a statistically significant treatment effect on the attitudinal index of 0.04, we also observe similar results, with the lower bound of the treatment effect on this measure growing to include zero given a small violation of our assumption about the discrepancy between the stated and actual preferences, making our estimate uninformative about the true direction of the effect.

For our behavioral measure—whether or not respondents would share or discuss with others the news story they read, shown in the bottom row of Figure 6, larger deviations from the naïve assumption are necessary to negate our significant negative estimates of the ACTEs for the respondents who prefer MSNBC. Among these respondents, the upper bound on our treatment effect remains negative until ρ reaches 0.06—nearly equivalent to the size of our treatment effect—and even the 95% confidence interval does not contain zero until ρ is 0.03.

Taken together, these sensitivity analyses indicate that our evidence of media persuasion effects depends on assumptions of relatively modest deviations between stated preferences and actual media choices. When we allow for larger deviations between the opinions of the stated-preference and actual-choice groups, the lower bounds of our estimated persuasive effects of partisan media on attitudes become negative even for the group of people who state a preference for entertainment. This means that we can no longer entirely exclude the possibility that these effects are actually below zero under those scenarios. The estimated levels of sensitivity further increase when we incorporate statistical uncertainty into our analysis. Our estimated effects on intended behavior are robust to relatively larger violations of this assumption.

Thus, we cautiously conclude that our findings are robust to reasonable—but not large—violations of the assumption that the discrepancy between the stated and actual media preferences is ignorable.²⁶ But more broadly, our analysis highlights the importance of

²⁶ Though assessing the reliance on this assumption is important, we have several reasons to believe that the true discrepancy between the stated and actual choice groups is not actually larger than the values of ρ at which our estimates would be uninformative for these outcomes. Scholars have found that stated preferences in online surveys strongly correlate with respondents’ media consumption habits in the real world (e.g., Levendusky 2013b; Sood and Lelkes 2018). Similarly, we find relatively high correspondence between stated and revealed media preferences among the subjects in our survey.

FIGURE 6. Sensitivity Analyses

accounting for the potential deviation between stated preferences and actual behavior in any experimental paradigm aimed at accurately estimating attitudinal change.

CONCLUSION

The explosion of consumer choice over the past several decades and with it the resurgence of an American partisan press, combined with a parallel rise in partisan polarization, has deepened scholarly interest in learning how media influence public opinion. Understanding the role of partisan media is particularly important when building coalitions across partisan and ideological lines seems ever more elusive and citizens increasingly question the capacity of our leaders to overcome partisan polarization.

In contrast to previous research, we demonstrate a strong persuasive impact of partisan media on political attitudes with data collected through our PICA design. While other studies have shown that counter-attitudinal information is unlikely to persuade people, and may even produce backlash effects (e.g., Nyhan and Reifler 2010), we show that both pro- and counter-attitudinal partisan media can, in fact, change people's opinions in the direction of the partisan media's message. This fits with a growing body of research showing limited backlash to opposing information (e.g., Guess and Coppock 2018; Wood and Porter 2019).

Our results indicate that political polarization is not simply a function of selective exposure. To be sure, selective exposure is a real phenomenon—people do prefer media that supports their preexisting attitudes, and those who consume opposing media have polarized political opinions, which we confirm using the free choice arm of our experiment. Yet, this does not preclude partisan media from having a strong polarizing effect as well. In the case of marijuana policy—and on three additional policy areas presented in Online Appendix E²⁷—our results show a strong impact of partisan news on the people who read it—not just among those who would ordinarily choose to read these sources, but also among people who prefer apolitical entertainment news. Indeed, partisan media may have the greatest potential to polarize attitudes among the large segment of people who prefer not to consume partisan news at all. For these consumers, a single exposure to partisan media from one ideological perspective can create polarization in attitudes equivalent to approximately one-third of the polarization that exists between self-selected partisan news consumers on opposite sides of the aisle. Thus, while the consequences are minor if these people maintain an entertainment-only media diet, if they deviate from

²⁷ While there are some differences, as could be expected in different policy areas and on different survey samples, our results largely replicate when the topic is military strikes against ISIS, fracking, or charter school education policy.

such content or are inadvertently exposed to partisan media from one side or the other, their attitudes may change—and by a substantively large amount.

We also find evidence that partisan news can further change the opinions of people who would already self-select into these partisan news silos, exacerbating polarization. Yet, this result also indicates that partisan news media can reduce polarization via exposure to counter-attitudinal media sources among those who would ordinarily choose to consume pro-attitudinal partisan media. Persuasion by oppositional partisan media can help bridge the gap in opinions between groups of consumers. Moreover, the power of partisan media is not limited to opinions, but extends to socially and politically consequential behaviors, as measured by our sharing index. People—particularly those who would prefer to consume MSNBC—are more likely to spread pro- than counter-attitudinal partisan news stories to other people. This result suggests that biases may arise in the interpersonal flow of information. Though endorsements from social connections and other characteristics of specific media content may ameliorate this tendency (e.g., Messing and Westwood 2014; Mummolo 2016), it may still lead to increasingly homogenous information received via other people (cf. Brundidge 2010). Understanding how media consumption affects these types of behavior is especially important given the growing prevalence of online social media as vehicles for information sharing.

We also leveraged our PICA design to bound these persuasion estimates based on the instability between self-reported media consumption preferences and observed behavior. The bounds generated through our sensitivity analyses demonstrate how easily persuasion effects may disappear if assumptions about the stated preferences are not true. These sensitivity analyses illustrate the importance of assessing the treatment effects of partisan media for individuals making real-world choices among media options. Indeed, our experimental findings might look different “in the wild” with violations of our assumption about stated preferences and actual media choices. Further research on persuasion that incorporates real-world behavior along with stated preferences on surveys can better uncover the true effects of media without relying on such assumptions.

Taken together, these results suggest a more nuanced story regarding the effects of ideological media than scholars have previously recognized. Though many people do disproportionately choose to consume news from their own side of the ideological spectrum, many others—both partisans and independents—prefer entertainment over partisan news. Partisan news can be persuasive to these individuals. Separately estimating these effects for people with different media-viewing preferences can help detect important heterogeneity in the effects of partisan media (e.g., Gaines and Kuklinski 2011; Leeper 2017). We demonstrate that reading partisan news on a politically salient topic can change people’s opinions, as well as their actions, to different degrees depending on their media preferences. Indeed, contrary to the “minimal effects” hypothesis, among the

segment of the population who mostly prefers entertainment media, we find strong evidence of attitudinal persuasion.

If, as our findings suggest, partisan news can affect both attitudes and actions, then there is cause for concern. Some people may indeed choose to separate their media consumption patterns according to their ideology, but ironically it is the people who opt for largely apolitical stories who are most vulnerable to the persuasive effects of slanted news. When these individuals consume partisan media content, their opinions can change.

Our findings on the persuasiveness of ideological media are one piece of a puzzle and raise other questions about the broader implications of partisan media. For instance, more research on the persistence of these persuasive effects, or the aggregate effect of repeated treatments in a media-rich environment, could give us some indication of how they will affect politics writ large. Furthermore, estimating how interactions with political media affect subsequent choices among media options could potentially prove an important factor in assessing the total persuasive effect of partisan media. While ideologically slanted media may increase polarization among its natural consumers, our results suggest that it is the large number of relatively apolitical individuals, rather than the much smaller number of regular consumers of these news sources, who are most susceptible to such polarization effects. Indeed, this last finding suggests a powerful incentive for political entrepreneurs to seek out these entertainment-seekers and deliver such slanted political messages directly to them via social network feeds, native ads, or other means not requiring an individual’s choice to consume partisan news. Assessing the extent to which they are able to successfully target these individuals and expose them to partisan news remains a fruitful avenue for future research.

SUPPLEMENTARY MATERIAL

To view supplementary material for this article, please visit <https://doi.org/10.1017/S0003055419000418>.

Replication materials can be found on Dataverse at: <https://doi.org/10.7910/DVN/MM2VWC>.

REFERENCES

- Arceneaux, Kevin, and Martin Johnson. 2013. *Changing Minds or Changing Channels: Partisan News in an Age of Choice*. Chicago: University of Chicago Press.
- Arceneaux, Kevin, Martin Johnson, and Chad Murphy. 2012. “Polarized Political Communication, Oppositional Media Hostility, and Selective Exposure.” *The Journal of Politics* 74 (1): 174–86.
- Baum, Matthew, and Tim J. Groeling. 2008. “New Media and the Polarization of American Political Discourse.” *Political Communication* 25 (4): 345–65.
- Baum, Matthew, and Tim J. Groeling. 2010. *War Stories: The Causes and Consequences of Public Views of War*. Princeton: Princeton University Press.

- Bennett, W. Lance, and Shanto Iyengar. 2008. "A New Era of Minimal Effects?" *Journal of Communication* 58 (4): 707–31.
- Berelson, Bernard R., F. Paul Lazarsfeld, and William N. McPhee. 1954. *Voting: A Study of Opinion Formation in a Presidential Campaign*. Chicago: University of Chicago Press.
- Bolsen, Toby, James N. Druckman, and Fay Lomax Cook. 2014. "The Influence of Partisan Motivated Reasoning on Public Opinion." *Political Behavior* 36 (2): 235–62.
- Brundidge, Jennifer. 2010. "Encountering 'Difference' in the Contemporary Public Sphere: The Contribution of the Internet to the Heterogeneity of Political Discussion Networks." *Journal of Communication* 60 (4): 680–700.
- Bullock, John. 2011. "Elite Influence on Public Opinion in an Informed Electorate." *American Political Science Review* 105 (3): 496–515.
- Campbell, Angus, Philip E. Converse, Warren Miller, and Donald Stokes. 1960. *The American Voter*. Chicago: University of Chicago Press.
- Coe, Kevin, David Tewksbury, Bradley J. Bond, Kristin L. Drogos, Robert W. Porter, Ashley Yahn, et al. 2008. "Hostile News: Partisan Use and Perceptions of Cable News Programming." *Journal of Communication* 58 (2): 201–19.
- Dilliplane, Susanna. 2014. "Activation, Conversion, or Reinforcement? The Impact of Partisan News Exposure on Vote Choice." *American Journal of Political Science* 58 (1): 79–94.
- Ellithorpe, Morgan E., Lance Holbert, and Angela L. Palmer-Wackerly. 2013. "Procrastination and the Shifting Political Media Environment: An Experimental Study of Media Choice Affecting a Democratic Outcome." *Communication Studies* 64 (5): 561–78.
- Feldman, Lauren. 2011a. "The Opinion Factor: The Effects of Opinionated News on Information Processing and Attitude Change." *Political Communication* 28 (2): 161–81.
- Feldman, Lauren. 2011b. "Partisan Differences in Opinionated News Perceptions: A Test of the Hostile Media Effect." *Political Behavior* 33 (3): 407–32.
- Hamilton, James T. 2005. "The Market and the Media." In *The Press*, eds. Geneva Overholser and Kathleen H. Jamieson. Oxford: Oxford University Press, 351–71.
- Gaines, Brian, and James Kuklinski. 2011. "Experimental Estimation of Heterogeneous Treatment Effects Related to Self-Selection." *American Journal of Political Science* 55 (3): 724–36.
- Gentzkow, Matthew, and Jesse Shapiro. 2011. "Ideological Segregation Online and Offline." *Quarterly Journal of Economics* 126 (4): 1799–839.
- Groeling, Tim. 2013. "Media Bias by the Numbers: Challenges and Opportunities in the Empirical Study of Partisan News." *Annual Review of Political Science* 16: 129–51.
- Guess, Andrew, and Alexander Coppock. 2018. "Does Counter-Attitudinal Information Cause Backlash? Results from Three Large Survey Experiments." *British Journal of Political Science* 1–19. Published online 5 November 2018.
- Iyengar, Shanto, and Kyu S. Hahn. 2009. "Red Media, Blue Media: Evidence of Ideological Selectivity in Media Use." *Journal of Communication* 59 (1): 19–39.
- Iyengar, Shanto, Kyu S. Hahn, Jon A. Krosnick, and John Walker. 2008. "Selective Exposure to Campaign Communication: The Role of Anticipated Agreement and Issue Public Membership." *The Journal of Politics* 70 (1): 186–200.
- Kim, Young Mie. 2009. "Issue Publics in the New Information Environment: Selectivity, Domain Specificity, and Extremity." *Communication Research* 36 (2): 254–84.
- Knobloch-Westerwick, Silvia, and Jingbo Meng. 2009. "Looking the Other Way: Selective Exposure to Attitude-Consistent and Counterattitudinal Political Information." *Communication Research* 36 (3): 426–48.
- Knobloch-Westerwick, Silvia, and Steven B. Kleinman. 2012. "Pre-election Selective Exposure: Confirmation Bias versus Informational Utility." *Communication Research* 39 (2): 170–93.
- Knox, Dean, Teppei Yamamoto, Matthew A. Baum, and Adam Berinsky. 2019. "Design, Identification, and Sensitivity Analysis for Patient Preference Trials." *Journal of the American Statistical Association* 1–27. Published online 30 April 2019.
- Kunda, Ziva. 1990. "The Case for Motivated Reasoning." *Psychological Bulletin* 108 (3): 480–98.
- Ladd, Jonathan M. 2012. *Why Americans Hate the Media and How it Matters*. Princeton: Princeton University Press.
- Lazarsfeld, Paul Felix, Bernard Berelson, and Hazel Gaudet. 1948. *The People's Choice: How the Voter Makes up His Mind in a Presidential Campaign*. New York: Columbia University Press.
- Leeper, Thomas J. 2014. "The Informational Basis for Mass Polarization." *Public Opinion Quarterly* 78 (1): 27–46.
- Leeper, Thomas J. 2017. "How Does Treatment Self-Selection Affect Inferences About Political Communication?" *Journal of Experimental Political Science* 4 (1): 21–33.
- Leeper, Thomas J., and Rune Slothuus. 2014. "Political Parties, Motivated Reasoning, and Public Opinion Formation." *Political Psychology* 35 (1): 129–56.
- Levendusky, Matthew S. 2013a. *How Partisan Media Polarize America*. Chicago: University of Chicago Press.
- Levendusky, Matthew S. 2013b. "Why Do Partisan Media Polarize Viewers?" *American Journal of Political Science* 57 (3): 611–23.
- Messing, Solomon, and Sean J. Westwood. 2014. "Selective Exposure in the Age of Social Media: Endorsements Trump Partisan Source Affiliation When Selecting News Online." *Communication Research* 41 (8): 1042–63.
- Mummolo, Jonathan. 2016. "News from the Other Side: How Topic Relevance Limits the Prevalence of Partisan Selective Exposure." *The Journal of Politics* 78 (3): 763–73.
- Nielsen. 2010. "Nielsen Provides Topline U.S. Web Data for March 2010." Accessed April 27, 2010. <http://www.nielsen.com/us/en/insights/news/2010/nielsen-provides-topline-u-s-web-data-for-march-2010.html>.
- Negroponte, Nicholas. 1995. "Being Digital—A Book (P)review." *Wired*. Retrieved April 14, 2017. <https://www.wired.com/1995/02/negroponte-27>.
- Nyhan, Brendan, and Jason Reifler. 2010. "When Corrections Fail: The Persistence of Political Misperceptions." *Political Behavior* 32 (2): 303–30.
- Pariser, Eli. 2012. *Filter Bubble: Wie wir im Internet entmündigt werden*. Munich: Hanser.
- Prior, Markus. 2007. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. New York: Cambridge University Press.
- Prior, Markus. 2013. "Media and Political Polarization." *Annual Review of Political Science* 16: 101–27.
- Sood, Guarav, and Yphtach Lelkes. 2018. "Don't Expose Yourself: Discretionary Exposure to Political Information." In *Oxford Research Encyclopedia of Politics*. <https://oxfordre.com/politics/view/10.1093/acrefore/9780190228637.001.0001/acrefore-9780190228637-e-39>.
- Stroud, Natalie J. 2011. *Niche News: The Politics of News Choice*. New York: Oxford University Press.
- Sunstein, Cass. 2001. *Republic.com*. Princeton: Princeton University Press.
- Taber, Charles S., and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50 (3): 755–69.
- Torgerson, David J., and Bonnie Sibbald. 1998. "Understanding Controlled Trials: What Is a Patient Preference Trial?" *British Medical Journal* 316 (7128): 360.
- Wood, Thomas, and Ethan Porter. 2019. "The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence." *Political Behavior* 41 (1): 135–63.
- Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. Cambridge: Cambridge University Press.
- Zell, Ethan, and Michael J. Bernstein. 2014. "You May Think You're Right Young Adults Are More Liberal Than They Realize." *Social Psychological and Personality Science* 5 (3): 326–33.