

The Long-lasting Effects of Newspaper Op-Eds on Public Opinion

Alexander Coppock¹, Emily Ekins² and David Kirby^{3*}

¹ *Yale University, New Haven, CT 06520, USA*

^{2,3} *Cato Institute, 1000 Massachusetts Ave, N.W. Washington, DC, USA*

ABSTRACT

Do newspaper opinion pieces change the minds of those who read them? We conduct two randomized panel survey experiments on elite and mass convenience samples to estimate the effects of five op-eds on policy attitudes. We find very large average treatment effects on target issues, equivalent to shifts of approximately 0.5 scale points on a 7-point scale, that persist for at least one month. We find very small and insignificant average treatment effects on non-target issues, suggesting that our subjects read, understood, and were persuaded by the arguments presented in these op-eds. We find limited evidence of treatment effect heterogeneity by party identification: Democrats, Republicans, and independents all appear to move in the predicted direction by similar magnitudes. We conduct this study on both a sample of Amazon Mechanical Turk workers and a sample of elites. Despite large differences in demographics and initial political beliefs, we find that op-eds were persuasive to both the mass public and elites, but marginally more persuasive among the mass public. Our findings add to the growing body of evidence of the everyday nature of persuasion.

Keywords: Persuasion; persistence; survey experiment

*This research was reviewed and approved by the Columbia University Institutional Review Board (IRB-AAAP9305). Except where indicated, all analyses were pre-registered at egap.org. The authors are grateful to the Cato Institute and its staff for their assistance.

Online Appendix available from:

http://dx.doi.org/10.1561/100.00016112_app

Supplementary Material available from:

http://dx.doi.org/10.1561/100.00016112_supp

MS submitted on 1 October 2016; final version received 27 October 2017

ISSN 1554-0626; DOI 10.1561/100.00016112

© 2018 A. Coppock, E. Ekins and D. Kirby

“As the world has grown smaller, the nation more powerful, the problems besetting man infinitely more complex, the pressures more intense, the health of this democracy has increasingly depended on deeper public understanding of difficult issues. Through the new page opposite the Editorial Page that we inaugurate today, we hope that a contribution may be made toward stimulating new thought and provoking new discussion on public concerns.”

New York Times, September 21, 1970

The modern “opposite the Editorial Page,” or op-ed, debuted in the *New York Times* on September 21, 1970. At its launch, opinion pieces were designed to provide an intellectual arena to provoke new ideas and discussion on public policies.¹ John B. Oakes, editor of the *Times*, had long argued that the central function of newspaper should be to “interpret [the] age to the general public” while avoiding the the inaccessibilities of elitism (quoted in Socolow (2010)). The *Times* intended to create a forum to host a variety of outside experts to articulate their arguments and engage in the “exchange and clash of ideas” for the benefit of the general public, their intended audience. In fact, when deciding whether to launch the op-ed pages, the *Times* publisher believed an op-ed page would help maintain readership as the paper raised prices. Thus, from its inception, the op-ed was intended to be written by a diverse array of elite experts on the salient issues of the day with the purpose of prompting civic discourse and learning among the general public. Indeed, many view the op-ed pages today as a mechanism that continues to shape public opinion (Porpora and Nikolaev, 2008). In the more than 40 years since, the op-ed format has grown widely, with nearly every major print and online newspaper publishing two to three op-eds per day (Sommer and Maycroft, 2008).

Present-day newspaper op-eds are very similar in form to those envisioned in the 1970s: expert elites publish policy opinions intended to provoke debate among the general public (Sommer and Maycroft, 2008). However, today it is acknowledged that op-ed pages of major papers set the agenda not only for the general public but also for the fellow elites² (Alexander, 2004; Golan, 2013; Nico Calavita, 2003; Rosenfeld, 2000; Sommer and Maycroft, 2008). For

The Cato Institute funded the remuneration of survey respondents but explicitly allowed the authors to publish findings regardless of the results.

¹Newspaper pieces that bear some resemblance to the modern op-ed can be traced back as early as 1912 when the *Chicago Tribune* featured a separate page dedicated to outside opinion. *New York World* editor Herbert Bayard Swope was the first to structure a section of commentary opposite the editorial page. Several newspapers followed suit such as the *Washington Post* who featured an “op-ed page” in the 1930s, and the *Los Angeles Times* in the 1950s (Socolow, 2010).

²We take a broad view of “elites,” among whom we include lawmakers, journalists, Hill staffers, economists, lawyers, and political and policy professionals of all stripes.

instance, Sommer and Maycroft (2008) claim that during legislative debate, lawmakers routinely circulate op-eds in efforts to persuade colleagues to their point of view.

A large and sophisticated op-ed industry has developed in Washington, DC. For instance, in 2015, scholars at the Cato Institute, a Washington think tank where two of the present study's co-authors work, published 944 op-eds, including 73 in top 10 newspapers (Cato Institute, 2016). Cato is by no means unique among think tanks. The Brookings Institute website catalogs 116 op-eds in top publications in the 2015³ and the American Enterprise Institute reports 3,385 (American Enterprise Institute, 2015). These think tanks employ dedicated staff whose job it is to edit and place op-eds. Of course, the White House and U.S. House and Senate Leadership Offices have members of their communication staff who serve similar functions, not to mention trade associations and advocacy groups. Organizations that are too small to employ own in-house op-ed writing teams often engage one of the many public relationships firms located in Washington DC. Indeed, former opinion editors at national newspapers can make lucrative second careers ghost-writing and placing op-eds for CEOs of major companies, politicians, and celebrities. Our informal inquiries reveal that the going rate for an op-ed can range from \$5,000 to \$25,000, depending on the sophistication of the piece and prestige of the placement.

Given all the time, energy, and money expended on the production of op-eds, a natural question arises: Do op-eds accomplish either the goal of persuading the mass public or persuading elites? A large research literature has shown that members of the mass public know little about the nuanced details of politics and public policy (Berelson *et al.*, 1954; Campbell *et al.*, 1960; Converse, 1964; Delli Carpini and Keeter, 1996; Neuman, 1986; Somin, 1998). Given low levels of political knowledge, one might expect that complex arguments that depend on understanding nuanced policy details would be ineffective among a general audience that at best has a hazy grasp of the terms of the debate. Some scholars have suggested that op-ed pages are "off limits" to average people because of elite authorship and focus on complex policy topics (Ciofalo and Traveso, 1994). If op-eds are ineffective at changing mass opinion, should they be expected to change the minds of elites? Since elites are more politically aware and exposed to more policy arguments, they may be more selective of which policy considerations they decide to accept and process (Zaller, 1992). They may also be more resistant to information that conflicts with their political identities.

If it were the case that neither of the intended audiences, mass or elite, actually updates their views as a result of exposure to op-eds, why would

³This figure calculated from the search results at Brookings.edu (Content Type = Op-Ed, Topic = Any, Start Date = 2015-01-01, End Date = 2015-12-31).

anyone write them? Perhaps it flatters the egos of op-ed writers to see their names in print. Perhaps op-eds bolster the authors' influence within elite policy circles by virtue of increased name recognition or "buzz." Or perhaps op-eds are ineffective at changing minds, but newspapers publish them anyway because they attract readers who already agree with the message: Op-eds may preach to the choir, but at least the choristers generate ad revenue.

We take up the question of whether newspaper op-eds actually achieve their original objective of changing minds. If so, do they have greater impact among the mass public or elites? Do op-eds change the minds of readers ideologically opposed to the author's argument? Are opinion changes persistent or ephemeral?

In this paper, we present results from two large-scale randomized panel survey experiments that estimate the effects of newspaper policy op-eds on public opinion, one among a mass sample and the other among elite opinion leaders, with three main results. For both mass public and elite samples, we randomly assigned respondents to receive one of several newspaper policy op-eds or not. All respondents then took the same public opinion survey that included policy questions regarding policies related to each of the several op-ed treatments used. We then administered two follow-up surveys to the mass public sample and one follow-up survey to the elite sample to determine if any measured persuasive effects endured.

We find, first, large average treatment effects among both the mass public and elites, despite large differences in demographics and initial political beliefs. Consistent with expectations, these effects are somewhat smaller among elite readers. Second, we find limited evidence of heterogeneous treatment effects by partisan identification. This indicates that op-eds do not merely "preach to the choir" by only changing the opinions of aligned ideological adherents. Instead, we find that Democrats, Republicans, and independents update their opinions in the predicted direction of the op-ed and by roughly similar magnitudes. Third, we find that policy op-eds have a very large treatment effects on target issues that persist for at least one month, suggesting that effects are due to underlying attitude change and not simply experimenter demand or survey artifacts.

Our study differs from previous work in several ways. First, we use real, unmodified opinion pieces as treatments. This choice means that we cannot describe what particular feature of the op-ed was the causal agent that changed minds, but what we lose in our ability to separate out mechanisms, we gain back in realism. Second, we unbundle the opinion piece from the rest of the newspaper, allowing us to understand the specific impact of the op-ed separate from the news of the day. Third, as recent changes in survey technology have made panel studies cheaper and more feasible, we measure outcomes over the course of an entire month, allowing us to estimate the persistence of the persuasive effects of these treatments.

Previous Literature

As Ciofalo and Traveso (1994) note, there is little scholarly research of the origins, role, and effect of newspaper op-eds. Furthermore, to our knowledge, no previous study has specifically investigated the persuasive effects of newspaper op-eds on policy opinions of readers. The extant research has tended to focus on op-eds' use of issue-framing, not persuasion specifically (e.g. Golan, 2010, 2013; Porpora and Nikolaev, 2008).

More has been written about the effects of newspapers generally, not just the op-ed section. These investigations into the persuasive effect of newspaper media content on public opinion generally fall into two categories. The first is observational studies that use aggregate data to try and demonstrate a causal link between newspaper exposure and political outcomes. Exposure is often operationalized as newspaper "slant." Overall, these studies tend to conclude that newspaper content impacts voters' candidate evaluations and political knowledge. For instance, Dalton *et al.* (1998) examine the correlation between media content and voter preferences in the 1992 election and find that newspaper editorials were a statistically significant predictor of candidate favorability. Druckman and Parkin (2005) focus on a single Senate race find that editorial slant of two local newspapers is predictive of both candidate evaluations and vote choice. Jerit *et al.* (2006) find that variations in quantities of newspaper and broadcast coverage of political issues is correlated with political knowledge, with this relationship appearing strongest among the more educated. Nicholson (2003) find that voters were more aware of ballot initiatives if featured in a major newspaper, controlling for other aspects of the political environment leading up to an election.

Drawing causal inferences from these studies can require strong assumptions, in particular that people who are exposed one level of "slant" are otherwise similar to those who are exposed to a different level. This problem is compounded by measurement error in the independent variable. People have poor recall and inaccurately report what they have seen and read (Guess, Forthcoming), and furthermore, errors associated with respondent recall can be correlated with political attitudes (Vavreck *et al.*, 2007). For instance, individuals who are more interested in politics are more likely to receive and recall media messages and thus differ systematically from those who don't recall receiving media content. Ladd and Lenz (2009) endeavor to address these selection issues by comparing those who do and do not read British newspapers who switched their endorsement of political parties to Labour in the 1997 elections, finding large effects of newspaper endorsements. Page *et al.* (1987) compared identical pairs of survey questions fielded at two points in time and TV news content that occurred in between fielding. They find that TV news content significantly predicts aggregate opinion change, particularly

content from experts and political commentators, precisely the types of people who typically author newspaper op-eds.

The second category of studies examining the effects of newspapers generally are randomized experiments in which the researcher directly controls the content to which subjects are exposed. DeFleur *et al.* (1992) measure the effect of exposing respondents to news stories from various types of media on recall and knowledge of the news stories, finding that people learned the most from newspapers. Norris and Sanders (2003) found campaign issue information transmitted through various platforms increased political knowledge of the parties, but did not find newspapers to outperform other types of media. In their book, Neumann *et al.* (1992) expose respondents to parallel news stories from different types of media and measure the impact on political knowledge, finding that newspapers did increase knowledge although broadcast media was more effective among those with average cognitive ability. Nyhan and Reifler (2010) and Wood and Porter (2018) use altered newspaper op-eds as part of their treatments. They measure the effect of adding corrective information to altered newspaper op-eds on readers' perceptions of facts. They find that by and large readers update their misperceptions when corrected with reputable facts and that some corrections are more effective than others.

Two previous studies have examined the impact of newspapers in a field experimental setting. Gerber *et al.* (2009) randomly sent Virginia residents either the left-leaning *Washington Post*, the right-leaning *Washington Times*, or neither, and found little evidence of differences in political knowledge, opinions about political events, or voter turnout. However, they do find that receiving *either* paper led to increased support for the Democratic gubernatorial candidate. The treatments in that study were the entire newspapers, including but not limited to the opinion pages. Jerit *et al.* (2013) conduct and contrast both laboratory and field experiments of newspaper article stories on knowledge and policy attitudes. In the laboratory experiment they find large treatment effects on some attitudes but find weak evidence of newspaper impact on attitudes or knowledge in the field experiment. Of the 17 outcome measures reported, only one attitudinal measure registered a significant effect.

Taken together, the existing scholarly record on the effects of newspapers on political attitudes and behavior is mixed. Observational studies estimate positive effects, but those estimates may be prone to bias due to measurement error and unobserved heterogeneity. The survey and laboratory experimental literature finds that newspapers can indeed increase political knowledge. The existing field experimental evidence on this point shows no evidence of a consistent relationship, though it should be noted that compliance with the experimental treatments may have been low, indicating that these studies were underpowered to detect the modest effects measured by other research designs.

Our reading of the previous literature leads us to five hypotheses:

- **(H1) We expect that newspaper op-eds can influence the opinions of readers on target issues in the direction intended by the author.** Expectancy value and accessibility models of how mass media information impacts public opinion contend that survey response is a product of respondents' considerations and the weights associated with the respective considerations (Ajzen and Fishbein, 1980; Chong and Druckman, 2012; Iyengar, 1990; Nelson *et al.*, 1997; Zaller, 1992). Accordingly, survey response results from the balance of political messages, or considerations, received and resisted and how recently those considerations have been brought to mind (Zaller, 1992). Applying the expectancy value model to previous studies of newspaper effects suggests that such studies' newspaper treatments influenced the combination of considerations available to impact opinion response. Consequently, we would expect newspaper op-eds to influence readers' opinion on target issues in the intended direction.
- **(H2) We do not expect op-eds to impact attitudes on non-target issues.** Since the scope of treatment is a single policy opinion piece, we avoid introducing competing considerations and thus we do not expect op-eds to impact attitudes on non-target issues.
- **(H3) We expect the effects of op-eds to be smaller among elites compared to the mass public.** Existing research would lead us to expect smaller treatment effects among elites, since elites are more politically sophisticated, are exposed to more policy arguments than the mass public, and have more stable attitudes (Feldman, 1989; Jennings, 1992; Putnam *et al.*, 1979). Expectancy models of public opinion anticipate that survey responses are a product of respondents' top-of-the-head considerations and the weights associated with the respective considerations (Ajzen and Fishbein, 1980; Chong and Druckman, 2012; Nelson *et al.*, 1997; Zaller, 1992). Accordingly, influential considerations result from the balance of political messages received and resisted and how recently those considerations have been brought to mind (Zaller, 1992). Consequently, we would expect that if anything, effects of op-eds would be smaller among elites compared to the mass public.
- **(H4) We expect heterogeneous treatment effects by respondent party identification.** A similar line of reasoning to H3 suggests that we may also find smaller effects among subjects whose political identities are at odds with the op-ed's policy argument. Such readers may resist information inconsistent with their ideology, suggesting heterogeneous treatment effects across partisan groups.

- **(H5) We expect op-eds to have lasting persuasive effects.** Baden and Lecheler (2012) argue that informational treatments should have more enduring effects. Since our treatments are full-length real op-eds, and thus laden with policy-relevant information, we anticipate that treatment effects should persist.

Study 1: Mechanical Turk

We enrolled 3,567 subjects on Amazon’s Mechanical Turk (MTurk) in a three-wave panel survey.⁴ Participants were offered \$1.00 for each wave of the survey they completed. In Wave 1, we collected pre-treatment background variables, exposed subjects to one of five treatment op-eds (or nothing), and collected immediate outcomes. Ten days later, we recontacted subjects for Wave 2, in which we collected outcomes a second time. Wave 2 included a cross-cutting distraction experiment (the results of which are detailed in the Appendix) to guard against the possibility that respondents simply remember how they answered the questions last time or that respondents imagined that the researchers were looking for particular answers. Wave 3, conducted 30 days after treatment, collected outcomes a third and final time.

This design is summarized in Table 1. We obtained a recontact rate of 84% in Wave 2 and 69% in Wave 3. The treatment does not appear to have influenced whether subjects respond in subsequent waves. Using chi-square tests, we fail to reject the null hypothesis that treatment status and response in follow-up waves are independent (Wave 2: $p = 0.10$; Wave 3: $p = 0.61$).

We analyze our experiment in two main ways. The first simply compares mean outcomes in each treatment group to the control group. This approach has the advantage of making no assumptions about how the op-eds might influence attitudes on non-target issues. The second approach compares mean outcomes in each treatment group to all other treatment conditions, essentially rolling the other treatment groups in with the control subjects. The main advantage of this approach is increased statistical precision, but the interpretation of the resulting estimates can be tricky. If we assume that treatments do not affect attitudes towards non-target issues, then pooling these groups together presents no problem. If treatment can affect non-target issues, then this effect can be interpreted as the effect of treatment in a “noisy” information environment. As shown in Section 2.3, this distinction turns out

⁴Some research conducted on MTurk has been criticized on external validity grounds. While we are sensitive to such concerns, we also note that recent replication efforts (Coppock, 2018; Mullinix *et al.*, 2015) have shown that experimental estimates obtained on Mechanical Turk correspond very closely with those obtained on nationally representative samples. For this reason, we believe that our MTurk sample will provide estimates of the effects of newspapers among the “mass” public to a first approximation.

Table 1: MTurk experimental design.

Wave 1	Treatment condition						Totals
	Control	Amtrak	Climate	Flat Tax	Veterans	Wall Street	
<i>N</i>	622	597	570	587	592	603	3571
Wave 2							
Distraction	263	240	252	243	256	243	1497
No distraction	253	247	240	243	257	261	1501
Responded	516	487	492	486	513	504	2998
Did not respond	106	110	78	101	79	99	573
Wave 3							
Responded	433	412	386	386	412	422	2451
Did not respond	189	185	184	201	180	181	1120

not to matter much, either for precision or interpretation: we obtain similar answers using both approaches.

Study 1: Treatments

Our treatments consisted of op-eds advocating for libertarian policy positions. Table 2 presents the op-ed topics, titles, authors, and publication outlets. Cato Institute scholars authored four of the op-eds and a presidential candidate, Sen. Rand Paul (R-KY), authored the fifth op-ed. The op-ed on Amtrak argues that government spends transportation funds inefficiently, using new revenue to build new projects rather than to maintain and repair existing infrastructure, such as old rail lines. The author argues for Amtrak and other infrastructure to be funded through user fees rather than general taxes to ensure money goes toward its intended purpose. The op-ed on climate change suggests that natural causes also contribute toward climate change and that politicians unjustly wield their political power to bully climate scientists who challenge “alarmist claims about the climate.” The op-ed on the Department of Veterans Affairs criticizes the mismanagement of the VA and argues to replace the government run veterans health care system with a voucher system in which the government would give money to veterans to purchase private health insurance. The op-ed on Wall Street argues that Wall Street bankers are not all bad and help efficiently allocate investment funds to the companies that need it to produce valued consumer goods. The op-ed on the flat tax suggests that we “blow up the tax code” and adopt a flat tax of roughly 14%

Table 2: Op-ed treatments.

Treatment	Title	Author	Publication
Amtrak	The Amtrak Crash: Is More Spending the Answer?	Randal O’Toole	Newsweek
Climate	The Political Assault on Climate Skeptics	Richard S. Lindzen	The Wall Street Journal
Flat Tax	Blow Up the Tax Code and Start Over	Rand Paul	The Wall Street Journal
Veterans	The Other Veterans Scandal	Michael F. Cannon and Christopher Preble	The New York Times
Wall Street	Wall Street Offers Very Real Benefits	Thaya Knight	USA Today

on incomes over \$50,000, arguing that this would help all Americans not just wealthy Americans by eliminating tax loopholes.

To ensure the experience of reading the op-eds was as similar as possible to how respondents would encounter them in print or online, we preserved key visual elements such as the publication masthead logo, title and subtitle typesetting, font, font size and color, and byline with position and affiliation information. We include the full text of the treatments, as they were seen by our subjects, in the Appendix. Our Mechanical Turk subjects spent an average of 2 to 4 minutes reading, depending on the length of the op-ed.

Study 1: Outcome Measures

Each of the issue areas addressed by our treatment opinion pieces is complex and multifaceted. In order to measure our subjects’ policy attitudes, we endeavored to select questions that were as closely related to the specific arguments made by the op-eds. When possible, we selected standard survey questions to measure relevant attitudes. We chose four or five questions for each issue area. The wording for all dependent variables are given in the Appendix.

We will present the effects of treatment on our dependent variables in three ways. First, in our pre-analysis plan, we selected one question per topic to be the “main” dependent variable. These variables are all 7-point scales, with higher values indicating more libertarian attitudes. Second, we

constructed composite scales for each attitude, using all of the questions in the corresponding issue area. The scale was constructed by estimating a factor analysis model with two factors each, subjecting the resulting scores to the varimax rotation, and extracting the first dimension. This model was estimated on the control group only. The scores have unit variance in the control group, so any average shift can be interpreted directly in standard deviations. We use the same model to generate scores for the second and third waves, that is, the measurement model does not change over time (or across experiments). Third, we calculate dichotomous “agreement” dependent variables by splitting the composite scale variable at the median in the control group. We will use this variable when discussing the cost per “mind changed” in the section so titled. By construction, exactly 50% of the control group “agrees” with the Op-ed author; we assess the effects of treatment on increasing this agreement score.⁵ We will assess the robustness of this measure by dichotomizing at the 25th and 75th percentiles as well.

The main dependent variables are listed below. In each case, outcomes are coded so that higher values correspond to the predicted direction of the treatment effect due to the corresponding opinion piece.

- Amtrak Main DV: Do you think the government should spend more, less, or about what it does now on transportation and infrastructure? [7-point scale, 1: A lot more to 7: A lot less]
- Climate Main DV: Would you say that climate change is best described as a... [7-point scale, 1: Crisis to 7: Not a problem at all]
- Flat Tax Main DV: Would you favor or oppose changing the federal tax system to a flat tax, where everyone making more than \$50,000 a year pays the same percentage of his or her income in taxes? [7-point scale, 1: Strongly oppose to 7: Strongly favor]
- Veterans Main DV: How much confidence do you have in the Department of Veterans Affairs’ ability to care for veterans? [7-point scale, 1: A great deal, 7: None at all]
- Wall Street Main DV: How much confidence do you have in Wall Street bankers and brokers to do the right thing? [7-point scale, 1: None at all, 7: A great deal]

Study 1: Results

In this section, we present the immediate effects of treatment on our main dependent variables (Table 3) and composite scale variables (Table 4). For a

⁵In this spirit of full disclosure, we note that the “agreement” measurement strategy was not included in the pre-analysis plan.

Table 3: MTurk experiment: Treatment effects on main dependent variables.

	Amtrak	Climate	Flat Tax	Veterans	Wall Street
Op-ed: Amtrak	0.440*** (0.085)	-0.020 (0.095)	0.056 (0.108)	-0.056 (0.082)	-0.035 (0.079)
Op-ed: Climate	0.054 (0.079)	0.427*** (0.098)	0.132 (0.110)	-0.009 (0.082)	0.026 (0.080)
Op-ed: Flat Tax	0.195** (0.078)	0.130 (0.097)	0.850*** (0.109)	-0.033 (0.082)	-0.018 (0.079)
Op-ed: Veterans	0.047 (0.079)	0.032 (0.095)	-0.106 (0.110)	0.770*** (0.078)	-0.124 (0.078)
Op-ed: Wall Street	0.038 (0.078)	0.079 (0.095)	0.151 (0.107)	-0.210** (0.084)	0.915*** (0.082)
Constant (Constant)	2.902 (0.054)	2.773 (0.065)	3.738 (0.076)	4.502 (0.058)	2.616 (0.054)
N	3,571	3,571	3,571	3,571	3,571
R^2	0.012	0.008	0.026	0.048	0.061

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Models estimated via OLS. Robust standard errors in parentheses.

visualization of the effects of our treatments on all 21 dependent variables, see the Online Appendix.

All five treatments appear to have had large, robust effects on attitudes, either as measured by the main dependent variables or by the composite scale. Turning first to the main dependent variables, the size of the treatment effects of the op-eds on their target issues varied from 0.427 scale points (on a 1–7 scale) for the climate piece to 0.915 scale points for the Wall Street piece. All five of the effects on their target dependent variables are statistically significant at $p < 0.001$. The effects of the op-eds on the non-target issues are all close to zero. Of the 20 treatment effects on non target issues, only two are statistically significant (Flat Tax op-ed on Amtrak outcome and Wall Street op-ed on Veterans outcome).

The treatment effects reported in Table 4 are in terms of the composite scales, and can be interpreted in standard units. The effect sizes on the target issues are quite large. The smallest is again the effect of the climate op-ed on the climate scale, at 0.276 standard deviations, while the remainder range from about 0.5 to 0.7 standard deviations. By this measure as well, the cross-issue effects are small to non-existent.

In summary, the effects of the newspaper opinion pieces were large, positive, and statistically significant. Our hypothesis H1, that op-eds would affect atti-

Table 4: MTurk experiment: Treatment effects on composite scale dependent variables.

	Amtrak	Climate	Flat Tax	Veterans	Wall Street
Op-ed: Amtrak	0.501*** (0.059)	0.018 (0.056)	0.005 (0.056)	0.051 (0.056)	-0.034 (0.054)
Op-ed: Climate	0.078 (0.056)	0.276*** (0.057)	0.045 (0.057)	-0.007 (0.056)	0.008 (0.055)
Op-ed: Flat Tax	0.080 (0.055)	0.110** (0.056)	0.488*** (0.060)	0.048 (0.056)	-0.023 (0.056)
Op-ed: Veterans	0.044 (0.058)	0.003 (0.057)	-0.026 (0.057)	0.646*** (0.054)	-0.090* (0.054)
Op-ed: Wall Street	0.055 (0.055)	0.067 (0.057)	0.097* (0.057)	-0.113** (0.057)	0.698*** (0.056)
Constant (Constant)	-0.125 (0.039)	-0.079 (0.039)	-0.101 (0.039)	-0.103 (0.040)	-0.096 (0.038)
<i>N</i>	3,571	3,571	3,571	3,571	3,571
<i>R</i> ²	0.029	0.009	0.030	0.063	0.076

p* < 0.1; *p* < 0.05; ****p* < 0.01.
 Models estimated via OLS. Robust standard errors in parentheses.

tudes on target issues, is strongly supported in the MTurk sample. Hypothesis H2, that the op-eds would not move attitudes on non-target issues, was also strongly supported.

Study 2: Elites

In order to explore whether the results from the Mechanical Turk study would generalize to the other presumed target of op-eds in national newspapers, we conducted a nearly identical study among individuals one might characterize as “elite.” We began with a database of 32,498 email addresses of individuals that we considered, broadly, to be political or policy professionals. Our list included journalists, op-ed editors, television news producers, and other media professionals; think tank scholars, law professors, and other policy-focused academics; trade association, advocacy, and other government affairs professionals; Congressional legislative staffers in the U.S. House and Senate; state legislative policymakers and their staff; and, Wall Street traders, bankers, analysts, and other financial professionals. We did not offer these subjects any incentives for participation because we were cautioned that many government officials are specifically prohibited from accepting “gifts” of any kind.

Table 5: Elite experimental design.

Wave 1	Treatment condition					Totals
	Control	Amtrak	Flat Tax	Veterans	Wall Street	
<i>N</i>	448	407	463	438	425	2181
Distraction	139	135	126	132	138	670
No distraction	141	126	141	148	132	688
Responded	280	261	267	280	270	1358
Did not respond	168	146	196	158	155	823

We assured all subjects that their answers would be kept fully anonymous. The experimental treatments and outcome measures were identical to the Mechanical Turk study, but we dropped the “climate” treatment arm because we anticipated that our final sample would be too small to support the full six-arm design.

We invited these elite subjects to participate in our study by first sending a pre-invitation email indicating that we would be sending a survey link in a few days. We sent this link and two reminders to those who had not yet responded. This procedure yielded 2,169 subjects who completed our survey. As in the MTurk study, we asked subjects to participate in a follow-up survey after 10 days; after two reminders, we obtained 1,349 complete responses. Table 5 displays the experimental design as well as the number of subjects in each arm that we were able to recontact.⁶

The resulting sample is by no means a probability sample of all “elites” in the United States. First, there exists no commonly accepted definition of elite, nor does there exist a list of such elites from which to draw a representative sample. Second, many subjects declined to participate because they “don’t take political surveys” or were “too busy.” We heard from many journalists in particular who reported that they were not permitted by their employers to take surveys. Our elite subjects, like our MTurk subjects, constitute a convenience sample.

The elites differ in substantively important ways from the Mechanical Turk sample. They are older, whiter, more male, better educated, and hold stronger partisan attachments.⁷ On average, our elite sample spent 25–45 seconds less

⁶See the Appendix for a demonstration that attrition was unrelated to treatment assignment and for a description of our elite sample by sector.

⁷These differences are all highly statistically significant, as determined by chi-square tests. See the Appendix for a complete demographic breakdown of both samples.

Table 6: Elite experiment: Treatment effects on main dependent variables.

	Amtrak	Flat Tax	Veterans	Wall Street
Op-ed: Amtrak	0.438*** (0.110)	-0.094 (0.150)	0.069 (0.093)	0.071 (0.103)
Op-ed: Flat Tax	-0.023 (0.101)	0.411*** (0.144)	-0.060 (0.091)	-0.004 (0.101)
Op-ed: Veterans	-0.004 (0.103)	-0.156 (0.147)	0.045 (0.091)	0.119 (0.101)
Op-ed: Wall Street	0.042 (0.109)	0.067 (0.148)	-0.055 (0.093)	0.791*** (0.105)
Constant (Constant)	2.304 (0.073)	3.578 (0.103)	4.585 (0.064)	2.926 (0.071)
N	2,181	2,181	2,181	2,181
R^2	0.012	0.008	0.001	0.037

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Models estimated via OLS. Robust standard errors in parentheses.

time reading our treatment articles than the MTurk sample.

Study 2: Results

In this section, we present the immediate effect of treatment on our main dependent variables (Table 6) and composite scale variables (Table 7). As expected, we find moderately large, statistically significant effects of our treatments on policy attitudes for each of the four policy op-eds among our elite sample. Turning first to main dependent variables, we found statistically significant treatment effects for three of the four treatments (Amtrak, Flat Tax, Wall Street) at $p < 0.001$, but did not for the Veterans op-ed treatment. The size of the treatment effects of the op-eds on their target issues ranged from 0.411 scale points (on a 1–7 scale) for the Flat Tax treatment to 0.791 scale points for the Wall Street op-ed. We find no evidence of cross-issue effects of the treatments on non-target issues for the main dependent variables.

Table 7 presents the treatment effects in terms of the standardized composite issue scales. We find statistically significant treatment effects for three of the four treatments (Amtrak and Wall Street at $p < 0.001$; Veterans at $p < 0.05$), but did not for the Flat Tax treatment. Thus, we find significant treatment effects for the Flat Tax op-ed using the main dependent variable but not the composite scale variable, and the reverse is true for the Veterans Affairs op-ed. For a visualization of the effects of our treatments on all 16 dependent

Table 7: Elite experiment: Treatment effects on composite scale dependent variables.

	Amtrak	Flat Tax	Veterans	Wall Street
Op-ed: Amtrak	0.303*** (0.073)	-0.014 (0.070)	0.131** (0.065)	0.010 (0.063)
Op-ed: Flat Tax	-0.084 (0.067)	0.104 (0.070)	-0.0005 (0.062)	-0.039 (0.062)
Op-ed: Veterans	-0.041 (0.069)	-0.044 (0.070)	0.153** (0.063)	0.066 (0.063)
Op-ed: Wall Street	-0.021 (0.071)	0.077 (0.069)	-0.032 (0.064)	0.571*** (0.060)
Constant (Constant)	0.012 (0.048)	0.002 (0.049)	-0.043 (0.044)	-0.123 (0.044)
N	2,181	2,181	2,181	2,181
R^2	0.016	0.003	0.006	0.057

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Models estimated via OLS. Robust standard errors in parentheses.

variables, see the Online Appendix. The treatment effect sizes on their target issues range from 0.104 for the Flat Tax treatment to 0.571 for the Wall Street treatment. The effects of the op-eds on non-target issues are all close to zero. Of the 16 treatment effects on non-target issues, only one is statistically significant (Amtrak op-ed on Veterans outcome).

The elite experiment shows that even individuals with presumably well-formed and consistent opinion about politics can nevertheless change their minds in response to fact-based, reasoned arguments. Hypotheses H1 and H2 are supported in the elite sample as well.

Studies 1 & 2: Heterogeneous Effects by Experimental Sample

Our hypothesis H3 posited that while we predicted that op-eds would sway the opinions of both elite and mass subjects, the treatment effects would be smaller for the elite sample. Table 8 provides some evidence in support of that claim. The difference between the effects obtained on MTurk and the elite sample are represented by the interaction terms. On each of the target dependent variables, the treatment effect for elites was smaller. For example, the effect of the Amtrak op-ed on the Amtrak composite scale dependent variable was 0.198 scale points smaller for elites. This difference is statistically significant, as it is for the effects of the Flat Tax and Veterans treatments on their target outcomes. While the interaction is negative for the effect

Table 8: Comparison of treatment effects on composite scale dependent variables.

	Amtrak	Flat Tax	Veterans	Wall Street
Op-ed: Amtrak	0.501*** (0.059)	0.005 (0.056)	0.051 (0.056)	-0.034 (0.054)
Op-ed: Flat Tax	0.080 (0.055)	0.488*** (0.060)	0.048 (0.056)	-0.023 (0.056)
Op-ed: Veterans	0.044 (0.058)	-0.026 (0.057)	0.646*** (0.054)	-0.090* (0.054)
Op-ed: Wall Street	0.055 (0.055)	0.097* (0.057)	-0.113** (0.057)	0.698*** (0.056)
Elite Experiment	0.137** (0.062)	0.103 (0.063)	0.060 (0.059)	-0.027 (0.058)
Elite X Amtrak	-0.198** (0.094)	-0.018 (0.090)	0.080 (0.086)	0.044 (0.083)
Elite X Flat Tax	-0.164* (0.087)	-0.384*** (0.092)	-0.048 (0.084)	-0.016 (0.083)
Elite X Veterans	-0.085 (0.090)	-0.019 (0.090)	-0.493*** (0.083)	0.156* (0.083)
Elite X Wall Street	-0.076 (0.090)	-0.020 (0.089)	0.081 (0.086)	-0.127 (0.082)
Constant (Constant)	-0.125 (0.039)	-0.101 (0.039)	-0.103 (0.040)	-0.096 (0.038)
N	5,182	5,182	5,182	5,182
R^2	0.026	0.021	0.045	0.076

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Models estimated via OLS. Robust standard errors in parentheses.

of the Wall Street op-ed on its target dependent variable, the difference is not statistically significant. This provides an indication that on some issues, elites are more resistant to accepting considerations that run contrary to their existing views, comporting with hypothesis H3. Nevertheless, we still find moderately large treatment effects on policy attitudes, indicating that even elites are persuaded by newspaper op-eds, albeit to a lesser degree.

Studies 1 & 2: Heterogeneous Effects by Partisanship

In the previous section, we explore treatment effect heterogeneity across our experimental samples to see if elites and members of the mass public process our treatments differently. In this section, we consider another di-




Figure 1: Effects of treatment, by party and experimental sample.

mention along which treatment effects might vary: partisanship. Because our treatments were all from libertarian sources, one might expect Republicans, Democrats, and Independents to respond differently to the treatments. Hypothesis H4 held that if anything, Democrats should experience lower treatment effects.⁸

⁸We pre-registered a heterogeneous effects analysis using a machine-learning method (Bayesian additive regression trees). This method, while excellent for exploratory analysis of heterogeneous effects (Green and Kern, 2012), can sometimes obscure relatively straightforward questions such as whether effects differ by a particular subject characteristic. We opt instead to present a heterogeneous effects analysis by partisanship because of the clear relevance of respondents' partisanship to their political views.

Table 9: Joint tests of treatment effect heterogeneity by partisanship.

	Mechanical Turk		Elites	
	Main DV	Scale DV	Main DV	Scale DV
Amtrak	0.045	0.029	0.436	0.342
Climate	0.845	0.243		
Flat Tax	0.009	0.001	0.146	0.293
Veterans	0.878	0.562	0.434	0.291
Wall Street	0.310	0.371	0.316	0.265

Entries are p -values from F tests in which the null hypothesis is that the average treatment effects for Democrats, Independents, and Republicans are equal.

Figure 1 provides some evidence in support of this expectation. Some cases of treatment effect moderation are clear cut: in the Mechanical Turk sample, the effect of the Flat Tax treatment is larger among Republicans than it is among Democrats. Other cases are murkier. We can overcome some of the noise inherent in our estimates by averaging across treatments. The precision weighted average of the treatment effects on the Main DVs for is 0.58 (SE = 0.05) among Democrats on MTurk and 0.88 (SE = 0.08) among Republicans. The standard error of this 0.3 point difference is 0.09, indicating that the difference-in-difference is statistically significant. On average, Republicans experience higher treatment effects than Democrats.

That being said, we wish to be careful not to overstate the extent to which partisans respond differently to treatment. On the whole, subjects in all three subgroups update their attitudes by moderate amounts in response to the op-eds. Table 9 presents the results of a formal statistical test of the null hypothesis that the treatment effects among Republicans, Democrats, and Independents are equal. With the exceptions of the Amtrak and Flat Tax treatments in the Mechanical Turk sample, we fail to reject this null: treatment effects do not appear to vary dramatically by partisanship.

Studies 1 & 2: Long Term Effects

Next we measure the persistence of newspaper op-eds’ persuasive effects. As Gerber *et al.* (2011) point out, large treatment effects followed by steep decay undermine the idea that op-eds are actually changing minds, but rather making particular considerations more accessible when taking the survey. However, if newspaper op-eds have a persuasive effect lasting over several weeks time this would be suggestive of some degree of underlying attitudinal change (Coppock, 2016).

To measure the lasting effects of newspaper op-eds, we measured outcomes at two additional waves, after 10 days and again at 30 days, allowing us to estimate the persistence of the treatment effects caused by our op-eds, over time.⁹ Respondents were not shown the op-ed again, but instead were only presented with the survey of policy questions. An added benefit of the subsequent waves is to separate treatment from outcome measurement. Thus, respondents taking the survey in Wave 2 and Wave 3 were responding to opinion questions only, and had not been immediately primed to consider arguments prior to taking the surveys.

Figure 2 shows our results. Each facet groups together the effects of the treatments by outcome. Time in days since treatment is plotted on the horizontal axis and the average level of the composite scale by treatment groups is plotted on the vertical axis. At time zero, the large, robust effects of treatment can be easily discerned by comparing the outcomes of the target issue treatment group to all other conditions. The separation between the target issue treatment group and the other groups persists overtime. While the treatment effects are indeed smaller (approximately 50% the original magnitudes in each case), they remain statistically significant in most cases. These data support our hypothesis H5, that treatment effects would persist over time.

On Mechanical Turk, we measured outcomes at three points in time. We were surprised by the “hockey stick” pattern of results. We had expected that the 30-day results would be diminished relative to the 10-day results just as the 10-day results were diminished relative to the immediate results.¹⁰ However, we do not observe much decay at all after the initial decline. This phenomenon requires much further study, but an initial explanation of this pattern might be that immediate effects are a combination of two factors: information and priming. The priming effects fade, but the information effects persist.

Cost Per Mind Changed

In the previous sections, we demonstrated that op-eds have large, long lasting effects on attitudes among both elites and members of the mass public. We now turn to a brief discussion of the cost-effectiveness of op-eds. Analogous to calculations of the cost per vote of get-out-the-vote experiments, we consider the cost per mind changed of newspaper op-eds.

We calculate the cost per mind changed using Equation (1). It has three arguments, the cost of producing an op-ed, the number of people who read it,

⁹We restrict our attention to subjects who responded in all waves. We rely on an assumption that these subjects are “always-reporters,” or that whether or not they respond is unrelated to their treatment status.

¹⁰In the Appendix, we provide evidence that Democrats, Republicans, and Independents appear to experience a similar patterns of decay.




Figure 2: Long term effects of treatment.

and the probability of the op-ed changing a reader’s mind. We estimate the first two arguments using industry figures and the final argument using the results of our experiments.

$$\text{Cost per mind changed} = \frac{\text{Cost}}{\text{Op-Ed readership} * \text{Percentage point change in agreement}} \quad (1)$$

We take two approaches to estimating the cost of producing an op-ed. The first is the going market rate for ghost-writing and placing an op-ed, which can range between \$5,000 and \$25,000, depending on the complexity of the topic and the prestige of the placement. The second is the cost of an op-ed-sized advertisement, which costs approximately \$50,000 for the weekday opinion section of the *New York Times*.

Table 10 displays our attempt to back out the number of readers the average op-ed can expect to reach. The first column shows the total circulation numbers (print and online) for the *New York Times*, the *Wall Street Journal*, *USA Today*, and *Newsweek*, the outlets where our treatment op-eds were originally published (Alliance for Audited Media, 2014). The second column is an estimate of the proportion of readers who read the opinion section. Figures for *New York Times* and the *Wall Street Journal* are derived from internal surveys conducted at those institutions, the topline results of which were shared with us via email. Similar surveys were not available from *USA Today* and *Newsweek*, so we conservatively estimate that only 50% of readers of those publications read the opinion pages. We imagine that at most, one in four readers reads one of the three to six op-eds published per weekday by these publications. The final column shows our back-of-the-envelope guess for the number of unique views received by the average op-ed: about half a million for the *New York Times*, *Wall Street Journal*, and *USA Today* and approximately

Table 10: Estimated number of Op-Ed readers.

	Circulation (print + online)	% Reading opinion section	Average number of op-eds read	Unique op-ed readers
New York Times	2,134,150	95%	0.25	506,861
Wall Street Journal	2,276,207	97%	0.25	551,980
USA Today	4,139,380	50%	0.25	517,423
Newsweek	433,333	50%	0.25	54,167
			Average	407,608

50,000 for *Newsweek*. For use in our cost-per-mind changed formula, we will approximate the audience of an op-ed as comprising 400,000 readers.

The final number we need is the percentage point change in agreement. Unfortunately, we do not have a binary dependent variable that represents whether a subject “agrees” with the author from which we can derive a direct estimate of the number of minds that changed. Instead, we have a summary index of subjects’ attitudes in the five issue areas touched on by our treatments. We can dichotomize that index at any arbitrary value of the index in order to approximate whether subjects, in a summary fashion, agree or disagree with the author. The choice of split is consequential for the results. Because we did not pre-register a split in advance, we present three reasonable values at which to dichotomize the scale: the 25th, 50th, and 75th percentiles of the composite scale DVs in the Mechanical Turk control group. This choice normalizes the baseline level of “agreement” in the control group and we assess the impact of the treatments in increasing this proportion.

Table 11 shows the estimated treatment effects of the op-eds on their target issues in terms of agreement. On Mechanical Turk, the precision weighted average of the estimated effects (using the 50th percentile split) indicates that on average, the treatments increased agreement with the authors by 20 percentage points. That is, if 50% of the control group agrees with the author, then we can say that approximately 70% of the target issue treatment group agrees with the author. In the elite sample, the estimate is smaller, but is still impressive at 12 percentage points. Our results are somewhat sensitive to the choice of split. Minimally, the range of estimates presented in Table 11 help us to calibrate the plausible range of the proportion of the sample whose minds “changed” as a result of our treatments.

Under the rosier scenario, an op-ed costs \$5,000 to produce, reaches 400,000 people, and changes the mind of 20% of them. Plugging these values

Table 11: Effects of Op-eds on “agreement” dependent variables.

	25th Percentile		50th Percentile		75th Percentile	
	Elite	MTurk	Elite	MTurk	Elite	MTurk
Amtrak	0.04 (0.03)	0.12 (0.02)	0.10 (0.03)	0.15 (0.03)	0.09 (0.03)	0.23 (0.03)
Climate		0.10 (0.02)		0.11 (0.03)		0.10 (0.03)
Flat Tax	0.02 (0.03)	0.08 (0.02)	0.02 (0.03)	0.20 (0.03)	0.06 (0.03)	0.18 (0.03)
Veterans	0.02 (0.03)	0.17 (0.02)	0.08 (0.03)	0.28 (0.03)	0.09 (0.03)	0.29 (0.03)
Wall Street	0.14 (0.02)	0.16 (0.02)	0.26 (0.03)	0.28 (0.03)	0.21 (0.03)	0.27 (0.03)
Average	0.06 (0.01)	0.13 (0.01)	0.12 (0.02)	0.20 (0.01)	0.11 (0.02)	0.21 (0.01)

Entries are estimated average treatment effects.
 Robust standard errors are in parentheses.
 Agreement DVs are dichotomized at the 25th, 50th, or 75th percentile of the MTurk control group.

into Equation (1), we obtain that the cost per mind changed is a mere 6 cents. Under a more conservative set of assumptions, an op-ed costs 10 times as much (\$50,000), reaches half as many people (200,000), and changes half as many minds (10%). The resulting cost per mind changed would work out to \$2.50. The cost-effectiveness of an op-ed clearly varies with the persuasiveness of the author, but we think that cost-per-mind change figures that range between 50 cents and 3 dollars are reasonable.

Discussion

Our paper has sought to answer the question previously left unanswered by existing scholarship: do op-eds change the minds of the people who read them? Many remain skeptical that op-eds achieve the initial vision set out by the *New York Times* 40 years ago to meaningfully enhance debate and inform the general public, or even elites, on substantive policy matters. The pessimistic view would hold that both voters are ill-informed and thus perhaps unable to grasp detailed policy subjects, and that since well-informed elites are better able to resist information that conflicts with their partisan identities that they too would remain impervious to arguments running counter to existing views. Our study indicates that both pessimistic views of op-eds' impact are unwarranted.

In both studies of the mass public and elites, we find large, statistically significant average treatment effects of op-eds between 0.30 and 0.50 standard deviations on policy attitudes. These findings comport with previous scholarship that has found significant effects of newspapers on political knowledge, candidate evaluations, and turnout. However, unlike previous scholarship, we focus on the specific effects that newspaper op-ed pages have specifically on policy attitudes. A strength of our study is the self-replication: we repeated the same experimental design with both mass public and elite samples, finding large treatment effects on target issues and relatively small or no effects on non-target issues in both samples. As predicted, the effects among the elite sample, while robust, were smaller in magnitude than those estimated on the mass public sample.

Since these op-eds were written from the libertarian economic perspective, we would expect Democrats to be less persuaded than Republicans. We find some evidence that this is the case, but hasten to note that Democrats, Independents, and Republicans all move in the direction intended by the author and the differences across party are not enormous.

Furthermore, op-ed treatment effects are not ephemeral. They appear to persist for at least one month after the initial exposure to the op-ed, suggesting underlying attitudinal change. Over time results have the potential to shed light on the primary mechanism by which op-eds change attitudes. We expect

that treatments that operate by priming particular considerations or by framing arguments in one light or another to have fleeting effects, whereas treatments that operate by providing subjects new information are expected to last longer. The “hockey stick” pattern of decay between Waves 1 and 2 but the leveling off between Waves 2 and 3 leads us to consider that perhaps immediate treatment effects may be a combination of both priming as well as information impacting attitudes. Over the course of a month, however, priming effects may fade. Because the average effects of our treatments stayed at the approximately the same level from Wave 2 to Wave 3, we infer that these effects operate primarily by information rather than priming.

Lastly, these results indicate that op-eds are remarkably cost effective. We estimate, based on the cost of producing an op-ed, the number of people who read it, combined, and the probability of changing a reader’s mind, that the cost-per-mind changed ranges from approximately 50 cents to 3 dollars.

We acknowledge that survey experimental research is well-equipped to answer some questions while leaving others unanswered. As with any survey experiment, the extent to which these results generalize to other settings is of concern. Because we found similar findings across two very different populations (MTurk respondents and political elites), we would expect that if *this specific* experiment were conducted on new samples, the results would be quite similar. We also believe our results are robust to the choice of specific survey questions used to measure outcomes because we found similar results regardless of whether we used only one main dependent variable as the primary outcome measure or a composite scale of the 4–5 questions in each corresponding issue area. We also expect our findings to generalize to other op-ed treatments, including op-eds that were written by liberal or conservative (rather than libertarian) authors.

However, the inferential target might not be the effect that would obtain if this same experiment (or minor variants) were conducted on a new population. Instead, we might imagine that the main question of interest is what happens when people encounter op-eds in a real-world, naturalistic setting. In the “wild” of real-world media communications, op-eds are typically encountered in a bundle: a physical newspaper or email digest of the day’s news. People may selectively choose the content they read, skipping over op-eds with titles or authors that they find boring or that they anticipate will conflict with their existing points of view. Subjects in survey experiments may pay attention differently than those who would encounter an op-ed in their physical newspaper or online. Additionally, our treatments did not offer two-sided messages as are often found in newspapers offering “point-counterpoint” content. These differences between the survey experimental setting and the field suggest that the estimates of the effects of op-eds on opinion may represent an upper bound. Finally, we do not wish to suggest that when a newspaper publishes an op-ed on a particular issue that opinion shifts by 10–20 percentage points. Since

op-eds typically reach fewer than half a million readers, we estimate that approximately 0.2% of the US adult population reads even the most widely shared and successful opinion piece. If 10% of those readers change their minds, aggregate opinion would shift by a mere 0.02 percentage points. This perspective on the political impact of op-eds helps to reconcile our moderate-to-large treatment effect estimate and relatively stable over time public opinion.

We conclude from these results that well-argued, long-form opinion pieces have the ability to change minds of even ideological opponents, and contributes to the growing body of evidence of the everyday nature of persuasion. Although our increasingly polarized political environment poses challenges to substantive and respectful political discourse, our study shows that at least in the setting of newspaper op-eds, individuals are capable of considering diverse views and may perhaps even change their minds.

References

- Ajzen, I. and M. Fishbein. 1980. *Understanding Attitudes and Predicting Social Behavior*. Vol. 278. Englewood Cliffs, NJ: Prentice Hall.
- Alexander, E. R. 2004. "Capturing the Public Interest: Promoting Planning in Conservative Times". *Journal of Planning Education and Research* 24(1): 102–6.
- Alliance for Audited Media. 2014. "September 2014 U.S. Newspaper Data Now Available. Technical report". Available at: <http://auditedmedia.com/news/blog/2014/october/september-2014-us-newspaper-data-now-available/>.
- American Enterprise Institute. 2015. "Annual Report. Technical report".
- Baden, C. and S. Lecheler. 2012. "Fleeting, Fading, or Far-Reaching? A Knowledge-Based Model of the Persistence of Framing Effects". *Communication Theory* 22(4): 359–82.
- Berelson, B. R., P. F. Lazarsfeld, and W. N. McPhee. 1954. *Voting: A Study of Opinion Formation in a Presidential Campaign*. Chicago: University of Chicago Press.
- Campbell, A., P. E. Converse, W. E. Miller, and D. E. Stokes. 1960. *The American Voter*. New York: Wiley.
- Cato Institute. 2016. "Annual Report. Technical report". Available at: <http://object.cato.org/sites/cato.org/files/pubs/pdf/annual-report-2015-update.pdf>.
- Chong, D. and J. N. Druckman. 2012. "Dynamics in Mass Communication Effects Research". In: *The Sage Handbook of Political Communication*. Ed. H. Semetko and M. Scammell. Los Angeles, CA: Sage Publications, 307–23.
- Ciofalo, A. and K. Traveso. 1994. "Does the Op-ed Page Have A Chance to Become a Public Forum?" *Newspaper Research Journal* 15(4): 51–63.

- Converse, P. E. 1964. "The Nature of Belief Systems in Mass Publics". In: *Ideology and Discontent*. Ed. D. E. Apter. Ann Arbor: University of Michigan Press.
- Coppock, A. 2016. "The Persistence of Survey Experimental Treatment Effects". Unpublished Manuscript.
- Coppock, A. 2018. "Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach". *Political Science Research and Methods*, in Press.
- Dalton, R. J., P. A. Beck, and R. Huckfeldt. 1998. "Partisan Cues and the Media: Information Flows in the 1992 Presidential Election". *American Political Science Review* 92(1): 111–26.
- DeFleur, M. L., L. Davenport, M. Cronin, and M. DeFleur. 1992. "Audience Recall of News Stories Presented by Newspaper, Computer, Television and Radio". *Journalism and Mass Communication Quarterly* 69(4): 1010–22.
- Delli Carpini, M. X. and S. Keeter. 1996. *What Americans Know About Politics and Why it Matters*. New Haven: Yale University Press.
- Druckman, J. N. and M. Parkin. 2005. "The Impact of Media Bias: How Editorial Slant Affects Voters". *Journal of Politics* 67(4): 1030–49.
- Feldman, S. 1989. "Measuring Issue Preferences: The Problem of Response Instability". *Political Analysis* 1: 25–60.
- Gerber, A. S., J. G. Gimpel, D. P. Green, and D. R. Shaw. 2011. "How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment". *American Political Science Review* 105(1): 135–50.
- Gerber, A. S., D. Karlan, and D. Bergan. 2009. "Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions". *American Economic Journal* 1(2): 35–52.
- Golan, G. J. 2010. "Editorials, Op-ed Columns Frame Medical Marijuana Debate". *Newspaper Research Journal* 31(3): 50–61.
- Golan, G. J. 2013. "The Gates of Op-Ed Diplomacy: Newspaper Framing the 2011 Egyptian Revolution". *International Communication Gazette* 75(4): 359–73.
- Green, D. P. and H. L. Kern. 2012. "Modeling Heterogeneous Treatment Effects in Survey Experiments with Bayesian Additive Regression Trees". *Public Opinion Quarterly* 76(3): 491–511.
- Guess, A. M. Forthcoming. "Measure for Measure: An Experimental Test of Online Political Media Exposure". *Political Analysis*.
- Iyengar, S. 1990. "The Accessibility Bias In Politics: Television News And Public Opinion". *International Journal of Public Opinion Research* 2(1): 1–15.
- Jennings, M. K. 1992. "Ideological Thinking Among Mass Publics and Political Elites". *Public Opinion Quarterly* 56(4): 419–41.

- Jerit, J., J. Barabas, and T. Bolsen. 2006. "Citizens, Knowledge, and the Information Environment". *American Journal of Political Science* 50(2): 266–82.
- Jerit, J., J. Barabas, and S. Clifford. 2013. "Comparing Contemporaneous Laboratory and Field Experiments on Media Effects". *Public Opinion Quarterly* 77(1): 256–82.
- Ladd, J. M. and G. S. Lenz. 2009. "Exploiting a Rare Communication Shift to Document the Persuasive Power of the News Media". *American Journal of Political Science* 53(2): 394–410.
- Mullinix, K. J., T. J. Leeper, J. N. Druckman, and J. Freese. 2015. "The Generalizability of Survey Experiments". *Journal of Experimental Political Science* 2: 109–38.
- Nelson, T. E., R. A. Clawson, and Z. M. Oxley. 1997. "Media Framing of a Civil Liberties Conflict and Its Effect on Tolerance". *American Political Science Review* 91(3): 567–83.
- Neuman, W. R. 1986. *The Paradox of Mass Publics: Knowledge and Opinion in the American Electorate*. Cambridge: Harvard University Press.
- Neumann, W. R., M. R. Just, and A. N. Crigler. 1992. *Common Knowledge: News and the Construction of Political Meaning*. Chicago: University of Chicago Press.
- Nicholson, S. P. 2003. "The Political Environment and Ballot Proposition Awareness". *American Journal of Political Science* 47(3): 403–10.
- Nico Calavita, N. K. 2003. "Capturing the Public Interest: Using Newspaper Op-Eds to Promote Planning in Conservative Times". *Journal of Planning Education and Research* 22(4): 400–6.
- Norris, P. and D. Sanders. 2003. "Message or Medium? Campaign Learning During the 2001 British General Election." *Political Communication* 20(3): 233.
- Nyhan, B. and J. Reifler. 2010. "When Corrections Fail: The Persistence of Political Misperceptions". *Political Behavior* 32(2): 303–30.
- Page, B. I., Y. S. Robert, and D. Glenn. 1987. "What Moves Public Opinion?" *American Political Science Review* 81.
- Porpora, D. V. and A. Nikolaev. 2008. "Moral Muting in US Newspaper Op-eds Debating the Attack on Iraq". *Discourse and Communication* 2(2): 165–84.
- Putnam, R. D., R. Leonardi, and R. Y. Nanetti. 1979. "Attitude Stability Among Italian Elites". *American Journal of Political Science* 23(3): 463.
- Rosenfeld, S. S. 2000. "The Op-Ed Page: A Step to a Better Democracy". *The Harvard International Journal of Press/Politics* 5(3): 7–11.
- Socolow, M. J. 2010. "A Profitable Public Sphere: The Creation of the New York Times Op-Ed Page". *Journalism and Mass Communication Quarterly* 87(2): 281–96.
- Somin, I. 1998. "Voter Ignorance and the Democratic Ideal". *Critical Review* 12(4): 413–58.

- Sommer, B. and J. R. Maycroft. 2008. "Influencing Public Policy: An Analysis of Published Op-Eds by Academics". *Politics and Policy* 36(4): 586–613.
- Vavreck *et al.*, L. 2007. "The Exaggerated Effects Of Advertising On Turnout: The Dangers Of Self-reports". *Quarterly Journal of Political Science* 2(4): 325–43.
- Wood, T. and E. Porter. 2018. "The Elusory Backfire Effect: Mass Attitudes' Steadfast Factual Adherence". *Political Behavior*, in press.
- Zaller, J. R. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.